

Controllable Guarantees for Fair Outcomes via Contrastive Information Estimation

Umang Gupta¹, Aaron M Ferber², Bistra Dilkina², Greg Ver Steeg¹

¹ Information Sciences Institute, University of Southern California

² University of Southern California

{umanggup, gregv}@isi.edu, {aferber, dilkina}@usc.edu

Abstract

Controlling bias in training datasets is vital for ensuring equal treatment, or parity, between different groups in downstream applications. A naive solution is to transform the data so that it is statistically independent of group membership, but this may throw away too much information when a reasonable compromise between fairness and accuracy is desired. Another common approach is to limit the ability of a particular adversary who seeks to maximize parity. Unfortunately, representations produced by adversarial approaches may still retain biases as their efficacy is tied to the complexity of the adversary used during training. To this end, we theoretically establish that by limiting the mutual information between representations and protected attributes, we can assuredly control the parity of any downstream classifier. We demonstrate an effective method for controlling parity through mutual information based on contrastive information estimators and show that they outperform approaches that rely on variational bounds based on complex generative models. We test our approach on *UCI Adult* and *Heritage Health* datasets and demonstrate that our approach provides more informative representations across a range of desired parity thresholds while providing strong theoretical guarantees on the parity of any downstream algorithm.

1 Introduction

Decisions based on biased data can promote biased outcomes. Learning algorithms often exploit and exaggerate biases present in the training dataset. One way to prevent algorithms from reproducing bias in the data is to pre-process the data so that information about protected attributes is removed (Song et al. 2019; Madras et al. 2018; McNamara, Ong, and Williamson 2017). Ideally, users of this transformed data can focus on maximizing performance for their tasks using any available method without the risk of producing unfair outcomes (Cisse and Koyejo 2019). The strongest requirement for fair representations is to be statistically independent of sensitive attributes, but this may lead to large drops in predictive performance as sensitive attributes are often correlated with the target. Therefore, it is desirable to produce representations that can trade-off some measure of fairness (e.g., statistical parity in this work) with utility (Menon and Williamson 2018; Dutta et al. 2019).

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Representative Methods	Adversarial Guarantee	Controllable Parity
Song et al. (2019)	Weak ¹	Heuristic
Moyer et al. (2018)	Strong	No ²
Madras et al. (2018)	None	Heuristic
Roy and Boddeti (2019)	None	No
Ours	Strong	Provable

Table 1: Fair representation learning methods

Many recent approaches for learning fair representations leverage adversarial learning to remove unwanted biases from the data by limiting the ability of an adversary to reconstruct sensitive attributes from the representation during training (Jaiswal et al. 2020; Roy and Boddeti 2019; Song et al. 2019; Madras et al. 2018; Edwards and Storkey 2016). While adversarial methods have been shown to be useful in learning fair representations, they are often limited by the adversary’s model capacity. A more powerful adversary than the one used during training may reveal hidden biases that are present in the representations (Xu et al. 2020). As a result, a model trained to control fairness against one adversary has no guarantee to control fairness against an arbitrary adversary. Other methods for learning fair representations focus on inducing statistical independence (Moyer et al. 2018; Louizos et al. 2016). While fairness guarantees of these methods are agnostic to the downstream algorithms, including adversarial attempts to exploit bias, they are inefficient at trading-off between fairness and informativeness in the representation, which is often desirable for fairness applications. Some methods rely on complex generative modeling of data to discover representations that are invariant to the protected attributes while preserving as much information as possible about the data. However, the quality of the generative model is a performance bottleneck. We summarize some of these approaches and their properties in Table 1.

¹Song et al. (2019) minimize two different bounds on $I(\mathbf{z};c)$ — one is a very loose upper bound, and another uses adversarial learning. So the adversarial guarantee is unclear or, at best weaker.

²Moyer et al. (2018) designed their method for enforcing independence; however, we consider a modification of their method for controlling parity, based on our Thm. 2.

In this paper, we focus on a widely used fairness measure known as statistical parity, though other measures may be more appropriate for specific problems. The difference in outcomes for two groups, or parity, is often used to quantify the fairness of a decision procedure. It has also been shown to best capture people’s perception of fairness (Srivastava, Heidari, and Krause 2019). One of the main objectives of fair representation learning is to limit the parity of any possible downstream decision algorithm. To this end, we relate the parity of any possible decision algorithm to an algorithm-agnostic measure of dependence. We show that we can provably control the parity of any possible decision algorithm that acts only on these representations by limiting mutual information between the representations and the sensitive attributes.

Estimating and bounding mutual information from data is a challenging and active field of research (Poole et al. 2019). We propose practical ways to limit the mutual information via contrastive estimates of conditional mutual information, thereby bypassing the need to use complex generative models that require explicit assumptions about the input distribution. Contrastive information measures have demonstrated state-of-the-art performance on many representation learning tasks (Oord, Li, and Vinyals 2018; Veličković et al. 2019; Mikolov et al. 2013) but not for fair representation learning. We evaluate our approach on two fairness benchmark datasets — *UCI Adult* and *Heritage Health* dataset and show that representations provided by our method preserve more information at the desired fairness threshold compared to other adversarial as well as non-adversarial baselines.

Our main contributions are — a) We theoretically show that mutual information between the representations and the sensitive attributes bounds the statistical parity of any decision algorithm; b) We propose practical ways to limit mutual information leveraging contrastive information estimators to efficiently trade-off predictability and accuracy.

2 Mutual Information Bounds Parity

We consider a dataset of triplets $\mathcal{D} = \{x_i, y_i, c_i\}_{i=1}^N$, where x_i, y_i, c_i are iid samples from data distribution $p(\mathbf{x}, \mathbf{y}, \mathbf{c})$. \mathbf{c} are the sensitive or protected attributes, \mathbf{y} is the label, \mathbf{x} are features of the sample which may include sensitive attributes and $\hat{\mathbf{y}}$ denotes predicted label, according to some algorithm. We may also interpret $\hat{\mathbf{y}}$ as the outcome of some decision procedure. We use bold letters to denote the random variable, and the regular font represents corresponding samples. In this work, we consider stochastic representations of data i.e., $z(x) \sim q(\mathbf{z}|\mathbf{x} = x)$. We learn d -dimensional representations z of input x , such that any classifier trained on only z is guaranteed to be fair, i.e., it has parity within some δ' . In this work, we focus on statistical parity, a popular measure of group fairness, and it is defined as:

Definition 1. Statistical Parity: (Dwork et al. 2012) *It is the absolute difference between the selection rates of two groups. Mathematically,*

$$\Delta_{DP}(\mathcal{A}, \mathbf{c}) = |P(\hat{\mathbf{y}} = 1|\mathbf{c} = 1) - P(\hat{\mathbf{y}} = 1|\mathbf{c} = 0)|$$

where $\hat{\mathbf{y}}$ denotes decisions produced by some decision algorithm \mathcal{A} . When there are more than two groups, we define

statistical parity to be the maximum parity between any two groups (as implemented in Bird et al. (2020)).

Statistical parity is an algorithm dependent measure of fairness, whereas we require our representations to be fair for any downstream decision algorithm. Bearing this in mind, we show that mutual information between the representations and the protected attributes, denoted as $I(\mathbf{z}:\mathbf{c})$, can be used to limit the statistical parity via the following result.

Theorem 2. *For some $z, c \sim p(\mathbf{z}, \mathbf{c})$, $z \in \mathbb{R}^d$, $c \in \{0, 1\}$, and any decision algorithm \mathcal{A} that acts on z , we have*

$$I(\mathbf{z}:\mathbf{c}) \geq g(\pi, \Delta_{DP}(\mathcal{A}, \mathbf{c}))$$

where $\pi = P(c = 1)$ and g is a monotonically increasing function.

We defer the proof of the above statement, the expression of g , and generalization of the theorem for multinomial \mathbf{c} to Appendix A.1³. We also visualize the bound on parity from Thm. 2 in Appendix A.1.

We only require the following information for our arguments — g is monotonically increasing in Δ_{DP} , and π is constant for a specific dataset. We can see from Thm. 2 that $I(\mathbf{z}:\mathbf{c})$ bounds the parity of any downstream decision algorithm. Therefore, if \mathbf{z} is a representation with bounded mutual information with \mathbf{c} , then any algorithm relying only on \mathbf{z} to make a decision will also have bounded parity. Moreover, since g is a monotonically increasing function, any reduction in $I(\mathbf{z}:\mathbf{c})$ will decrease Δ_{DP} too.

We remark that $I(\mathbf{z}:\mathbf{c})$ has been used as a proxy objective to control statistical parity previously (Edwards and Storkey 2016; Song et al. 2019; Moyer et al. 2018). It is often justified via the data processing inequality and the intuition that both statistical parity and mutual information are measures of dependence. However, using the data processing inequality, we can only guarantee that if we limit information about \mathbf{c} in \mathbf{z} , then no subsequent operations on \mathbf{z} can increase information about \mathbf{c} , i.e., $I(\hat{\mathbf{y}}:\mathbf{c}) \leq I(\mathbf{z}:\mathbf{c})$, but this fact alone implies nothing about statistical parity. Our result (Thm. 2) demonstrates that limiting $I(\mathbf{z}:\mathbf{c})$ will monotonically limit statistical parity, which had not been theoretically demonstrated until now.

3 Practical Objectives for Controlling Parity

Equipped with an algorithm agnostic upper bound to parity, we now discuss practical objectives for learning fair representations. Along with limiting parity, we also want the latent representation to be highly predictive (informative) about the label, which is often realized by maximizing mutual information between \mathbf{y} and \mathbf{z} , i.e., $I(\mathbf{y}:\mathbf{z})$ implicitly (Edwards and Storkey 2016; Madras et al. 2018; Jaiswal et al. 2020) or explicitly (Moyer et al. 2018).

$$\begin{aligned} \mathcal{O}_1: \max_q I(\mathbf{y}:\mathbf{z}) \text{ s.t. } I(\mathbf{z}:\mathbf{c}) \leq \delta \\ \text{or, } \max_q I(\mathbf{y}:\mathbf{z}) - \beta I(\mathbf{z}:\mathbf{c}) \end{aligned} \quad (1)$$

where, $\mathbf{z}, \mathbf{x} \sim q(\mathbf{z}|\mathbf{x})p(\mathbf{x})$ and $\beta > 0$ is a hyperparameter.




Figure 1: Venn diagram to show interference between $I(\mathbf{y}:\mathbf{z})$ and $I(\mathbf{z}:\mathbf{c})$. See Sec. 3.1 for details.

3.1 Interference Between $I(\mathbf{y}:\mathbf{z})$ and $I(\mathbf{z}:\mathbf{c})$

While $I(\mathbf{y}:\mathbf{z})$ has been commonly proposed as a criterion to enforce the desiderata of representations being informative about labels, we argue that when the data is biased, i.e., $I(\mathbf{y}:\mathbf{c}) > 0$, maximizing $I(\mathbf{y}:\mathbf{z})$ is in direct contradiction with minimizing $I(\mathbf{z}:\mathbf{c})$. To illustrate this point, we refer to the information Venn diagram in Fig. 1. The goal of fair representation learning is to move the circle representing information about the representation, \mathbf{z} , to have high overlap with \mathbf{y} and low overlap with \mathbf{c} . However, there is a conflict in the gray region where we cannot increase overlap with \mathbf{y} without also increasing overlap with \mathbf{c} . In our experiments, we observe that this conflict hurts the model performance and makes it hard to achieve lower parity values at a reasonable accuracy (Fig. 5b). However, this conflict is not necessary. Since fair learning aims to capture information about \mathbf{y} that is *not* related to the protected attribute \mathbf{c} , we want to maximize the overlap between \mathbf{z} and the region of \mathbf{y} that excludes \mathbf{c} . This quantity is precisely the conditional mutual information, $I(\mathbf{y}:\mathbf{z}|\mathbf{c})$, which we propose to maximize instead of $I(\mathbf{y}:\mathbf{z})$. This leads us to the following objective:

$$\begin{aligned} \mathcal{O}_2: \max_q I(\mathbf{y}:\mathbf{z}|\mathbf{c}) \text{ s.t. } I(\mathbf{z}:\mathbf{c}) \leq \delta \\ \text{or, } \max_q I(\mathbf{y}:\mathbf{z}|\mathbf{c}) - \beta I(\mathbf{z}:\mathbf{c}) \end{aligned} \quad (2)$$

where, $\mathbf{z}, \mathbf{x} \sim q(\mathbf{z}|\mathbf{x})p(\mathbf{x})$ and $\beta > 0$ is a hyperparameter. Concurrent to this work, Rodríguez-Gálvez, Thobaben, and Skoglund (2020) have also proposed similar conditional mutual information term for retaining information about labels in the representation. Eq. 2 defines our approach, but the information-theoretic terms that appear are challenging to estimate directly. In the next sections, we derive practical variational bounds for these terms.

Before proceeding, we briefly discuss an ambiguity in the information-theoretic argument above. The triple overlap region of the information Venn diagram, sometimes called the “interaction information”, can be negative (Williams and Beer 2010). This corresponds to the case that \mathbf{z} and \mathbf{c} have *synergistic* information about \mathbf{y} . That is, their combination allows us to predict \mathbf{y} , even though individually they may not be predictive at all. The classic example of a synergistic interaction is the binary XOR relationship among three variables (example given in Appendix B). In that case, no variable can predict any other, but knowing any two variables

perfectly predicts the third. If synergies are present, we may be able to attain a large value of $I(\mathbf{y}:\mathbf{z}|\mathbf{c})$ while $I(\mathbf{y}:\mathbf{z}) = 0$. In other words, maximizing $I(\mathbf{y}:\mathbf{z}|\mathbf{c})$ may include synergistic information, even if it hurts the ability to predict \mathbf{y} from \mathbf{z} alone. While maximizing $I(\mathbf{y}:\mathbf{z})$ will not include synergistic information, it will prefer large positive interaction information. Positive interaction information can be interpreted as increasing the redundant information shared among $\mathbf{z}, \mathbf{y}, \mathbf{c}$, even though this conflicts with our goal of excluding information about protected attributes. Different choices of objective will alter the prioritization of various multivariate information terms, but ultimately our choice is justified by the empirical performance shown in Sec. 4.

3.2 Lower Bound for $I(\mathbf{y}:\mathbf{z})$ and $I(\mathbf{y}:\mathbf{z}|\mathbf{c})$

Fact 3. For any distribution $r(\mathbf{a}|\mathbf{b})$ and $a, b \sim p(\mathbf{a}, \mathbf{b})$,

$$H(\mathbf{a}|\mathbf{b}) = -\mathbb{E}_{\mathbf{a}, \mathbf{b}} \log r(\mathbf{a}|\mathbf{b}) - KL(p(\mathbf{a}|\mathbf{b}) \| r(\mathbf{a}|\mathbf{b}))$$

and therefore,

$$H(\mathbf{a}|\mathbf{b}) \leq -\mathbb{E}_{\mathbf{a}, \mathbf{b}} \log r(\mathbf{a}|\mathbf{b})$$

and equality holds when $r(\mathbf{a}|\mathbf{b}) = p(\mathbf{a}|\mathbf{b})$

By a direct application of above identity, we have:

$$\begin{aligned} I(\mathbf{y}:\mathbf{z}) &= H(\mathbf{y}) - H(\mathbf{y}|\mathbf{z}) \\ &\geq H(\mathbf{y}) + \max_r \mathbb{E}_{\mathbf{y}, \mathbf{z}} \log r(\mathbf{y}|\mathbf{z}) \end{aligned} \quad (3)$$

and similarly,

$$I(\mathbf{y}:\mathbf{z}|\mathbf{c}) \geq H(\mathbf{y}|\mathbf{c}) + \max_r \mathbb{E}_{\mathbf{y}, \mathbf{z}, \mathbf{c}} \log r(\mathbf{y}|\mathbf{z}, \mathbf{c}) \quad (4)$$

$H(\mathbf{y})$ and $H(\mathbf{y}|\mathbf{c})$ are properties of data and, therefore, constant from the optimization perspective. When \mathbf{y} is a one-dimensional variable denoting the target class, this is equivalent to minimizing cross-entropy. To this end, we will parametrize the variational distribution r using a neural network with parameters ψ , but other models can also be used.

3.3 Upper Bound for $I(\mathbf{z}:\mathbf{c})$

Our technique for upper-bounding $I(\mathbf{z}:\mathbf{c})$ is similar to Moyer et al. (2018) and makes use of the following observations:

$$I(\mathbf{z}:\mathbf{c}) = I(\mathbf{z}:\mathbf{c}|\mathbf{x}) + I(\mathbf{z}:\mathbf{x}) - I(\mathbf{z}:\mathbf{x}|\mathbf{c}) \quad (5)$$

$I(\mathbf{z}:\mathbf{c}|\mathbf{x}) = 0$, as \mathbf{z} is a function of \mathbf{x} and some independent noise. As a result, we have $I(\mathbf{z}:\mathbf{c}) = I(\mathbf{z}:\mathbf{x}) - I(\mathbf{x}:\mathbf{z}|\mathbf{c})$. The first term is the information bottleneck term (Alemi et al. 2017) and limits the information about \mathbf{x} in \mathbf{z} , and we will bound it by specifying a prior over \mathbf{z} . Second term tries to preserve information about \mathbf{x} but not in \mathbf{c} and we will lower bound it via contrastive estimation.

Upper Bound for $I(\mathbf{z}:\mathbf{x})$ by Specifying a Prior: In order to upper-bound $I(\mathbf{z}:\mathbf{x})$, we use the following observation:

Fact 4. For any $a, b \sim p(\mathbf{a}, \mathbf{b})$ and distribution $q(\mathbf{a})$

$$I(\mathbf{a}; \mathbf{b}) = \mathbb{E}_{\mathbf{a}, \mathbf{b}} \log \frac{p(\mathbf{a}|\mathbf{b})}{q(\mathbf{a})} - KL(p(\mathbf{a}) \| q(\mathbf{a}))$$

³Appendix is available at <https://arxiv.org/abs/2101.04108>.

and therefore,

$$I(\mathbf{a}:\mathbf{b}) \leq \mathbb{E}_{\mathbf{a},\mathbf{b}} \log \frac{p(\mathbf{a}|\mathbf{b})}{q(\mathbf{a})} = \mathbb{E}_{\mathbf{b}} \text{KL}(p(\mathbf{a}|\mathbf{b}) \parallel q(\mathbf{a})) \quad (6)$$

and equality holds when $p(\mathbf{a}) = q(\mathbf{a})$.

Therefore, we have

$$I(\mathbf{z}:\mathbf{x}) \leq \mathbb{E}_{\mathbf{x}} \text{KL}(q(\mathbf{z}|\mathbf{x}; \phi) \parallel p(\mathbf{z})) \quad (7)$$

where $p(\mathbf{z})$ is any distribution. This is similar to the rate term in the VAE or information bottleneck approach (Alemi et al. 2017; Higgins et al. 2017). Motivated by this similitude, we let $p(\mathbf{z})$ be standard normal distribution and $q(\mathbf{z}|\mathbf{x}; \phi)$ be a diagonal gaussian distribution whose mean and variance are parametrized as $\mu(x) = f_{\mu}(x; \phi)$, $\Sigma(x) = f_{\Sigma}(x; \phi)$. Other parameterizations like normalizing flows (Rezende and Mohamed 2015) or echo noise (Brekelmans et al. 2019) can be used too.

Lower Bound for $I(\mathbf{x}:\mathbf{z}|\mathbf{c})$ via Contrastive Estimation: We propose to lower bound $I(\mathbf{x}:\mathbf{z}|\mathbf{c})$ via contrastive estimation of mutual information, and we use the following proposition to derive our estimator.

Proposition 5. (Poole et al. 2019) For any $u, v, w \sim p(\mathbf{u}, \mathbf{v}, \mathbf{w})$, $\tilde{u} \sim p(\mathbf{u}|\mathbf{w})$, and function f , we have

$$I(\mathbf{u}:\mathbf{v}|\mathbf{w}) \geq \mathbb{E}_{\mathbf{u},\mathbf{v},\mathbf{w}} \log \frac{e^{f(u,v,w)}}{\frac{1}{M} \sum_{j=1}^M e^{f(\tilde{u}_j,v,w)}} \quad (8)$$

where, $u, \tilde{u} \in \mathcal{U}$, $v \in \mathcal{V}$, $w \in \mathcal{W}$, $f : \mathcal{U} \times \mathcal{V} \times \mathcal{W} \rightarrow \mathbb{R}$ and M is the number of samples from $p(\mathbf{u}|\mathbf{w})$.

The proof is similar to the non-conditional version (Poole et al. 2019), but for completeness, we present the proof in the Appendix A.2. As a direct application of this result, we can lower-bound and maximize $I(\mathbf{z}:\mathbf{x}|\mathbf{c})$. However, there is a caveat that we need to sample from $P(\mathbf{z}|\mathbf{c})$. Sampling from this conditional distribution in the general case can be hard, but it can be easily accomplished for our problem. Often for fairness applications, \mathbf{c} is a discrete random variable with low cardinality; in fact, it is often a binary random variable. Therefore, we can consider $\{z_j : (z_j, c_j = i)\}$ to be the samples from $p(\mathbf{z}|\mathbf{c} = i)$. In our experiments, we parametrize $f(z, x, c)$ as a bilinear function (similar to Oord, Li, and Vinyals (2018)).

3.4 Overall Variational Objective

Using Eq. 5, we can write our objective, i.e., Eq. 2 in terms of three information theoretic quantities as:

$$\max_q \underbrace{I(\mathbf{y}:\mathbf{z}|\mathbf{c})}_{\text{Eq. 4}} - \beta \left(\underbrace{I(\mathbf{z}:\mathbf{x})}_{\text{Eq. 7}} - \underbrace{I(\mathbf{z}:\mathbf{x}|\mathbf{c})}_{\text{Eq. 8}} \right)$$

The first term is lower-bounded using Eq. 4 and it is the same as the cross-entropy loss. Second term is upper bounded using Eq. 7. For this, we set $p(\mathbf{z})$ to be standard normal distribution and $q(\mathbf{z}|\mathbf{x})$ is a normal distribution whose parameters are generated by a neural network. The KL divergence term in Eq. 7 can then be expressed in the closed form. Finally, the last term is lower-bounded using the contrastive conditional

mutual information estimator from Proposition 5 (Eq. 8). We call our objective **Fair Contrastive Representation Learner (FCRL)**.

Objective 1 can be similarly computed and optimized using Eqs. 3, 7, and 8. We provide comprehensive details to compute the proposed variational objective in Appendix C.1.

Even though objective 1 & 2 require a single parameter β to be varied, we found in our experiments that it is necessary to use different multipliers for $I(\mathbf{x}:\mathbf{z}|\mathbf{c})$ and $I(\mathbf{z}:\mathbf{x})$ to effectively tradeoff fairness and accuracy. Therefore, we introduce another parameter λ . The $I(\mathbf{z}:\mathbf{c})$ term in both the objectives then becomes $I(\mathbf{z}:\mathbf{x}) - \lambda I(\mathbf{x}:\mathbf{z}|\mathbf{c})$. We found that $\lambda = 2$ performs better empirically and so we set $\lambda = 2$, unless specified otherwise.

3.5 Other Bounds for $I(\mathbf{z}:\mathbf{c})$ and Their Caveats

One may observe that $I(\mathbf{z}:\mathbf{c})$ could also be upper-bounded as a direct consequence of Eq. 7. This is possible, however upper bounding $I(\mathbf{z}:\mathbf{c})$ simply by introducing a prior marginal $p(\mathbf{z})$ as in Eq. 6 is not tractable, as that would require computing $p(\mathbf{z}|\mathbf{c})$ (Song et al. 2019). Therefore, it is necessary to decompose $I(\mathbf{z}:\mathbf{c})$ as the difference of two information-theoretic quantities (Eq. 5).

Another way to bound $I(\mathbf{x}:\mathbf{z}|\mathbf{c})$ is to use the lower bound from Eq. 4 which leads to a conditional reconstruction term.

$$I(\mathbf{x}:\mathbf{z}|\mathbf{c}) \geq H(\mathbf{x}|\mathbf{c}) + \max_{\theta} \mathbb{E}_{\mathbf{x},\mathbf{z},\mathbf{c}} \log q(\mathbf{x}|\mathbf{z}, \mathbf{c}; \theta) \quad (9)$$

where $q(\mathbf{x}|\mathbf{z}, \mathbf{c}; \theta)$ is some distribution parametrized by θ . $H(\mathbf{x}|\mathbf{c})$ is a constant. The last term can be seen as reconstructing x from c and its latent representation z and is similar to the distortion term in VAE, albeit with conditioning on \mathbf{c} (Moyer et al. 2018). While this is a tractable bound, it involves training a conditional decoder, which may be hard for certain domains and almost always requires either complex models or the restrictive assumption that feature dimensions are independent given \mathbf{z} . It also requires explicitly stating the distribution of \mathbf{x} , which is bypassed using contrastive estimation. In our experiments, we found that this approximation is limiting (Fig. 5a); therefore, we propose using a decoder-free variational approximation via contrastive estimation of mutual information.

Remarks on Adversarial Approach: One may also use Eq. 3 to approximate $I(\mathbf{z}:\mathbf{c})$, which leads to a common technique to bound $I(\mathbf{z}:\mathbf{c})$ via adversarial learning. This is a lower bound to $I(\mathbf{z}:\mathbf{c})$, but since we want to minimize $I(\mathbf{z}:\mathbf{c})$, it should ideally be upper bounded. However, Eq. 3 is often used and leads to adversarial min-max approaches (Madras et al. 2018; Edwards and Storkey 2016; Song et al. 2019), where maximization is over parameters of the classifier $r(\mathbf{c}|\mathbf{z})$ and minimization is over $q(\mathbf{z}|\mathbf{x})$. Other than the difficulty of optimization, we can see that adversarial approaches have a fundamental problem: they minimize a *lower bound* at each iteration. The gap between mutual information and its lower bound can be arbitrarily bad if the set of adversarial classifiers is restricted, as it must be in practice. Moreover, rather than using exact maximization, most of the methods approximate it using one or a few steps of

SGD. This also puts into question one of the common approaches for evaluating invariant representations by predicting c from z . Indeed, in our experiments, we find that this evaluation can be misleading and may pass unfair representations as fair.

4 Experiments

We validate our approach on two datasets — *UCI Adult* (Dua and Graff 2017) and *Heritage Health*⁴ Dataset. *UCI Adult* is 1994 census data with 30K samples in the train set and 15K samples in the test set. The target task is to predict whether the income exceeds \$50K, and the protected attribute is considered gender (which is binary in this case). We use the same preprocessing as Moyer et al. (2018). *Heritage Health* dataset is data of around 51K patients (40K in the train set and 11K in the test set), and the task is to predict the Charleson Index, which is an indicator of 10-year survival of a patient. We consider age as the protected attribute, which has 9 possible values. We use the same preprocessing as Song et al. (2019).

A fair representation learning algorithm aims to produce representations such that any downstream decision algorithm that uses these representations will produce fairer results. Therefore, similar to Madras et al. (2018), we train the representation learning algorithm on training data and evaluate the representations by training classifiers for downstream prediction tasks. Since our purpose is to assess the representations, we report average accuracy (as an indicator of most likely performance) and maximum parity (as an indicator of worst-case bias) computed over 5 runs of the decision algorithm with random seeds. Unlike Madras et al. (2018), we also allow for preprocessing to be done on representations. Preprocessing steps like min-max or standard scaling are common and often precede training of classifiers in a regular machine learning pipeline. The held-out test set is used to evaluate representations on downstream tasks only; we use the training set for all the other steps.

We compare with a number of recent approaches, including information-theoretic and adversarial methods from the recent literature. MIFR (Song et al. 2019) is a recent approach demonstrating competitive performance. MIFR combines information-theoretic and adversarial methods and generalizes several previous fair representation learning approaches (Louizos et al. 2016; Edwards and Storkey 2016; Madras et al. 2018; Zemel et al. 2013). A related approach, CVIB (Moyer et al. 2018), is based solely on information-theoretic optimization without adversarial training. We also compare with recent state-of-the-art adversarial methods. In particular, we compare with Adversarial Forgetting (Jaiswal et al. 2020), which is the state-of-the-art adversarial method for learning invariant representations. Roy and Boddeti (2019) (MaxEnt-ARL) proposes a theoretically superior min-max objective by proposing to train the encoder to maximize the entropy of sensitive attributes. In contrast, traditional adversarial approaches try to minimize the likelihood of the discriminator. Finally, we compare with

LAFTR (Madras et al. 2018), which minimizes a more directed adversarial approximation to parity and has been designed bearing in mind parity-accuracy trade-offs. However, LAFTR is only applicable when c is a binary variable. As a baseline, we also train a one hidden layer MLP predictor directly on the data without regards to fairness (*Unfair MLP*).

We visualize the trade-offs between fairness and task performance by plotting parity vs. accuracy curves for each representation learning algorithm. We vary each method’s inherent hyperparameters over the ranges specified in the original works to get different points on this curve. For a fair comparison, we set $d = 8$ for all the methods and use model components like encoder, decoder, etc., of the same complexity. We provide a detailed list of the hyperparameter grids, network architectures, and other training setups for all the methods in Appendix D. The code to reproduce all the experiments is available at <https://github.com/umgupta/fairness-via-contrastive-estimation>. We give a quantitative summary of performance across the entire spectrum of trade-offs by reporting the area over the Parity-Accuracy curve. This section uses a 1-hidden-layer MLP with ReLU non-linearity and 50 neurons in the hidden layer as the decision algorithm and representations are preprocessed by standard scaling. Results with other decision algorithms, i.e., Random Forest, SVM, 2-hidden-layer MLP, and logistic regression, are shown in the Appendix D.

Improved Accuracy Versus Parity Trade-offs: For different fair representation learners, we compare accuracy versus parity achieved for a specific downstream classifier. The goal is to push the frontier of achievable trade-offs as far to the bottom-right as possible, i.e., to achieve the best possible accuracy while maintaining a low parity. From a visual inspection of Fig. 2, we can see that our approach preserves more information about label y , across a range of fairness thresholds for both *UCI Adult* and *Heritage Health* datasets. We observed improved trade-offs for our method even when different downstream classifiers are used (results are shown in Appendix D), matching expectations since our method gives theoretical bounds on parity that hold regardless of the classifier used. In contrast, methods that use adversaries in training see increases in parity when a more powerful classifier is used for downstream classification.

Metric Quantifying Performance Across Trade-offs: Fair representation learning approaches often provide a single number to quantify fairness by predicting c from z at a single accuracy level or only qualitatively demonstrate re-

Method	UCI Adult	Heritage Health
FCRL (Ours)	0.307	0.319
CVIB	0.182	0.191
MIFR	0.251	0.202
MaxEnt-ARL	0.144	0
LAFTR	0.155	N/A
Adversarial Forgetting	0.087	0.156

Table 2: Area Over Parity Accuracy Curve (Higher is better)

⁴<https://www.kaggle.com/c/hhp>

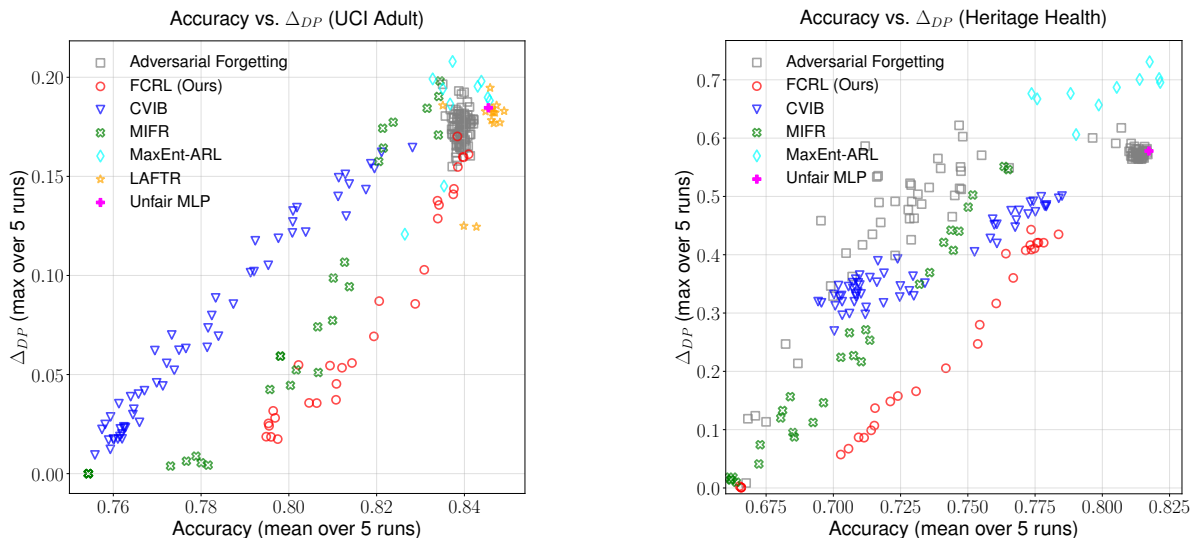


Figure 2: Parity vs. Accuracy trade-off for *UCI Adult* and *Heritage Health* dataset using a 1-hidden-layer MLP. Lower Δ_{DP} is better, and higher accuracy is better. We use representations generated by varying each method’s inherent loss hyperparameters to get different points. See Table 2 for quantitative results. Results with other decision algorithms are shown in Appendix D.

sults by showing the parity-accuracy curve. We report the normalized area over the parity-accuracy curve in Table 2 (1 being the maximum value) to enable quantitative comparison. It is the normalized hypervolume of the feasible parity-accuracy region, and an efficient fair representation learning method should maximize it. We provide more intuition and details to compute this metric in Appendix D.

Controlling Parity: Our approach has a single intuitive hyperparameter β , which can be used to control $I(\mathbf{z};\mathbf{c})$ directly and, therefore, via Thm. 2, to monotonically control parity (see Fig. 3). For the *UCI Adult* dataset, MIFR (Song et al. 2019) is competitive with our approach near low parity; however, it fails to achieve higher accuracy at higher demographic parity. This is because MIFR proposed to use $I(\mathbf{z};\mathbf{x})$ as an upper bound to $I(\mathbf{z};\mathbf{c})$, which is very loose, and this penalizes information about x as well, which is not desirable. CVIB (Moyer et al. 2018) is able to consistently trade-off accuracy while using a reconstruction based bound. But it maximizes $I(\mathbf{y};\mathbf{z})$ which conflicts with desired minimization of $I(\mathbf{z};\mathbf{c})$ (see Sec. 3.1).

Fine-tuning the Parity: To achieve different points on parity-accuracy trade-off, one must train models with different loss hyperparameters (β in our case). To a practitioner, parity-accuracy curves are essential, as it helps in deciding which representations should be used to satisfy the desired fairness constraints. The additional computational cost of training models, as well as the introduction of additional hyperparameters for fair learning pose significant hindrances towards the adoption of fair models (Sylvester and Raff 2018).

Recently, Gao and Chaudhari (2020) showed that it is possible to explore different equilibrium points on the rate-distortion curve with iso-accuracy constraints by adjusting the loss coefficients according to some dynamical process




Figure 3: Parity and Accuracy variation with β for *UCI Adult* dataset. Our method can explore feasible regions of parity and accuracy by varying only a single parameter β .

and finetuning the model parameters. Similarly, we show that we can finetune a trained model to produce different points on the parity-accuracy curve. Since our loss function has a single coefficient (β), we do not need to derive the dynamical process. We can simply vary β by sufficiently small steps and explore the parity-accuracy trade-off by finetuning the trained model. We can see that our method performs equally well when finetuning compared to training from scratch (See Fig. 4), reducing the computational cost drastically. We expect that the reduced computational cost of achieving the parity-accuracy curve and introduction of only a single intuitive hyperparameter along with improved parity-accuracy trade-off should reduce resistance towards




Figure 4: Exploring parity-accuracy trade-off for *UCI Adult* dataset by varying β and fine-tuning the model for 20 epochs. When training from scratch, we need to train for 6000 epochs to generate 30 points (200 epochs each). However, with fine-tuning we can get the same result with only 780 epochs ($200 + 29 \times 20$).

the adoption of fair models.

Uncovering Hidden Bias: In our experiments, we find that adversarial approaches that use deterministic representations (e.g., Roy and Boddeti (2019); Jaiswal et al. (2020)) are particularly susceptible to hidden bias, which can be revealed by slightly modified classifiers that consistently result in large parity values. Roy and Boddeti (2019) demonstrated that the representations produced were fair w.r.t to c by showing that an MLP trained to predict c from z performs poorly. This evaluation approach of predicting c from z is often used to demonstrate fairness or invariance of the representation w.r.t to c (Moyer et al. 2018; Jaiswal et al. 2020; Xie et al. 2017). However, we emphasize that due to Eq. 3, this is only a lower bound on $I(z:c)$ and, therefore, any claims relying solely on this metric are weak. Further, we found that these methods do not remove information about sensitive attributes and, at best, obfuscate it such that an adversary cannot easily recover the protected attribute during training. We found that even if a classifier trained directly on z cannot predict c , after preprocessing z by standard scaling, it becomes very easy to predict the sensitive attribute even with a model of similar or lower complexity as the original adversary used during the training. As a result, some of these approaches could not achieve low parities and even exhibited higher parity than the baseline *Unfair MLP* trained without fairness constraints (see Fig. 2). We investigate this further in the Appendix E.

Ablation Study of Information-Theoretic Methods

Next, we will show that our proposed objective of maximizing $I(y:z|c)$ is more efficient at exploring parity-accuracy trade-offs compared to maximizing $I(y:z)$ (i.e., objective 1 vs. 2). We also compare the effect of optimizing the reconstruction based information bound (Eq. 9) with our proposed

contrastive information estimator (Eq. 8).

Maximizing $I(y:z|c)$ vs. $I(y:z)$: We see in Fig. 5b that when maximizing $I(y:z)$, we can still reduce the parity, but this comes at the cost of predictive performance. When using $I(y:z|c)$, we see that the parity sharply drops with only a slight drop in the accuracy, and we see much better trade-offs between predictability and fairness. By conditioning mutual information on c , we see that z can easily retain information about y without conflicting with $I(z:c)$ objective as explained in Sec. 3.1.

Reconstruction vs. Contrastive Estimation: Fig. 5a compares optimizing reconstruction and contrastive estimation based bounds for $I(x:z|c)$. Reconstruction based bounds on information rely on explicitly modeling the data distribution through a decoder, $q(x|z, c)$, that reconstructs the data. For these methods to be effective, they must give reasonable reconstructions and therefore preserve most of the information in the data. We see in the results that preserving information leads to high accuracy but also high parity. Contrastive estimation is better at achieving lower parity than reconstruction based bounds because it can directly model high-level features that are predictive and fair without requiring a model to reconstruct data.

5 Related Work

As machine learning becomes more enmeshed in everyday life, a vibrant field of fair machine learning research has emerged to question potential risks (Mehrabi et al. 2019). Different applications may prescribe different notions of fairness, and these notions may even conflict with each other (Kleinberg, Mullainathan, and Raghavan 2017). In this work, we focus on fairness between groups concerning decisions that may be cast as classification problems. In particular, our work is focused on controlling statistical parity or disparate impact, which was proposed in Dwork et al. (2012) and is widely adopted in the literature (Edwards and Storkey 2016; McNamara, Ong, and Williamson 2017). Statistical parity is preferred as a fairness measure when the goal is to ensure equal selections from different groups. One such example is the US government’s *Uniform Guidelines on Employee Selection Procedure* (United States Equal Opportunity Employment Commission 1978) adoption of 80% rule to ensure fairness in selection procedures. Statistical parity as a fairness measure also captures a lay person’s perception of fairness (Srivastava, Heidari, and Krause 2019). Incorporating other popular notions of group fairness like Equalized-Opportunity and Equalized-Odds (Hardt et al. 2016), and notions of individual fairness (e.g., Sharifi-Malvajerdi, Kearns, and Roth (2019)) in our approach is left as future work.

Fair classification methods are often categorized based on which stage of the machine learning pipeline they target. Post-processing methods work by modifying the trained classifier (e.g. Hardt et al. (2016)), in-processing methods work by regularizing the classifier during training to ensure fairness (e.g. Zafar et al. (2017)), and pre-processing methods transform the original dataset to reduce bias present in




Figure 5: Ablation studies on *UCI Adult* dataset: Reconstruction (Eq. 9) and Contrastive (Eq. 8) mutual information estimation bounds are compared in Fig. a. Fig. b compares Objective 2 and Objective 1. λ was set to 1 for these experiments.

the data. Pre-processing methods are beneficial when the onus of fairness is on a third party or the data controller, and the end-user may be oblivious to fairness constraints (McNamara, Ong, and Williamson 2017; Cisse and Koyejo 2019). Our work also belongs to this category, and we specifically focus on controlling parity.

While post-processing and in-processing methods have to ensure fairness for a specific classifier, pre-processing methods must ensure fairness with respect to any downstream classification algorithm. Many pre-processing methods have discussed the desiderata of ensuring strong guarantees on fairness so that any downstream classifier may be used freely (McNamara, Ong, and Williamson 2017; Song et al. 2019; Edwards and Storkey 2016; Madras et al. 2018). However, their operationalization often leads to an approach that may not ensure guarantees (due to limits of adversarial methods, for instance). Our approach to these problems bounds parity in terms of mutual information, and then we optimize tight upper bounds on this information. Other works have explored information-theoretic objectives for learning invariant representations (Moyer et al. 2018) and fair representations (Song et al. 2019). Dutta et al. (2019) use tools from information theory to analyze the trade-off between fairness and accuracy. He, Burghardt, and Lerman (2020) minimize correlation between the sensitive attributes and the representations to achieve fairness. Correlation is a linear measure of dependence. Therefore, their method provides no guarantees when the downstream classifier is non-linear. In contrast, our approach controls mutual information between the sensitive attribute and the representation, a more robust measure of dependence.

Contrastive learning and its variants have shown promising results for learning representations for many applications, e.g., images and speech (Oord, Li, and Vinyals 2018), text (Mikolov et al. 2013) and graphs (Veličković et al. 2019). We are the first to explore its application for learn-

ing fair representations. Contrastive learning has been most actively explored in self-supervised learning, where the information to optimize is chosen by hand to be similar to some target task (Chen et al. 2020; Oord, Li, and Vinyals 2018), while in our work, we demonstrated a natural connection between parity and mutual information. Other variational bounds on information (Poole et al. 2019) and estimators like MINE (Belghazi et al. 2018) and NWJ (Nguyen, Wainwright, and Jordan 2010) could also be leveraged for parity control using our results.

6 Conclusion

As the negative repercussions of biased data become increasingly apparent, governments and other organizations will require tools capable of controlling this bias. In many cases, compromises between fairness and task objectives will be desired and sometimes even legally enshrined in terms of required bounds on fairness measures like statistical parity. Many methods do not provide a way to control parity, and even if they do, often, it is only in a heuristic way. Adversarial classifiers are popular for checking bounds on parity, but these methods are not guaranteed against the possibility of a more powerful classifier either intentionally or accidentally exploiting bias still hidden in the data. By proving a one-to-one relationship between information-theoretic quantities and statistical parity of arbitrary classifiers, we can finally see how varying a single hyper-parameter controlling information can explore the entire fairness versus accuracy spectrum. This information-theoretic characterization is algorithm-independent so that our control of parity can be guaranteed regardless of downstream applications.

Acknowledgements

We thank anonymous reviewers, especially reviewer 3 and 6 for their constructive feedback towards improving this pa-

per. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00112090104 and the DARPA NGS2 research award W911NF-16-1-0575. AF & BD were supported by NSF grant #1935451.

7 Ethical & Broader Social Impact

Machine learning is increasingly ubiquitous in daily life and can affect small things like what websites and products an individual sees to big things like who is approved for a loan or released on bail. Because machine learning relies on training data that may reflect historical bias, the outcomes of small and large decisions based on these systems may propagate that bias. The sudden emergence of a large amount of research on fair learning reflects the urgency and complexity of the issues. While clear-cut solutions to age-old ethical conundrums are unlikely, this should not prevent us from identifying spheres of human activity where machine-assisted decision-making can be made more fair. Our work follows in this vein by looking for ways to improve parity when decisions affect two groups, while maintaining as much ability to predict good outcomes as possible.

Fair machine learning research, including this paper, is not without pitfalls. There is a danger that mathematically attractive principles for fairness can lead to a sense of false security. Depending on the context, these measures can lead to counter-intuitive results. It is difficult for humans to agree on a fair outcome in many situations, and we cannot expect a single formula to settle the debate. For instance, Hardt et al. (2016) make two arguments against parity, the measure investigated in our paper. First, they argue that parity does not ensure intuitive notions of fairness because one way it can be satisfied is by taking qualified candidates from group A and an equal number of random candidates from group B. Even though parity is achieved, this may seem unfair as the most qualified candidate from group B could be omitted. Second, they argue that perfect parity hobbles the utility of machine learning. Because of subtle correlations among protected attributes and other predictive attributes, we end up completely destroying our ability to predict useful outcomes.

While the criticisms of parity given by Hardt et al. (2016) are valid, we believe that our work makes substantial contributions toward addressing both criticisms. First, we showed that trade-offs between parity and predictive performance could be significantly improved. This is important for widespread adoption of fair learning techniques, as voluntary adoption may be facilitated if the cost to performance can be demonstrated to be very low. Second, we introduced a tight connection between an information-theoretic perspective on fairness by limiting mutual information and parity. The generality of mutual information precludes bad faith solutions to achieve perfect parity like choosing qualified people from one group and random people from another. The reason for this is that, after learning a representation with bounded mutual information, information theory guarantees that it is impossible for any method to distinguish which group an individual is from in a way that would allow unequal application of selection criteria.

Although we believe that the information-theoretic approach we have provided in this paper is a more effective approach for guaranteeing parity, it should be recognized that parity is not necessarily the “fair” solution in every context. For instance, to address the negative repercussions that historical bias has had on a group may require treatment that purposely violates parity (Mehrabi, Huang, and Morstatter 2020).

References

- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2017. Deep Variational Information Bottleneck. In *International Conference on Learning Representations*.
- Belghazi, M. I.; Baratin, A.; Rajeshwar, S.; Ozair, S.; Bengio, Y.; Courville, A.; and Hjelm, D. 2018. Mutual Information Neural Estimation. In *Proceedings of the 35th International Conference on Machine Learning*, 531–540.
- Bird, S.; Dudík, M.; Edgar, R.; Horn, B.; Lutz, R.; Milan, V.; Sameki, M.; Wallach, H.; and Walker, K. 2020. Fairlearn: A toolkit for assessing and improving fairness in AI. Technical Report MSR-TR-2020-32, Microsoft. URL <https://www.microsoft.com/en-us/research/publication/fairlearn-a-toolkit-for-assessing-and-improving-fairness-in-ai/>.
- Brekelmans, R.; Moyer, D.; Galstyan, A.; and Ver Steeg, G. 2019. Exact Rate-Distortion in Autoencoders via Echo Noise. In *Advances in Neural Information Processing Systems*, volume 32, 3889–3900.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv preprint arXiv:2002.05709*.
- Cisse, M.; and Koyejo, S. 2019. Fairness and representation learning. *NeurIPS Invited Talk* 2019.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository. URL <http://archive.ics.uci.edu/ml>.
- Dutta, S.; Wei, D.; Yueksel, H.; Chen, P.-Y.; Liu, S.; and Varshney, K. R. 2019. An Information-Theoretic Perspective on the Relationship Between Fairness and Accuracy. *arXiv preprint arXiv:1910.07870*.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12. ISBN 9781450311151.
- Edwards, H.; and Storkey, A. 2016. Censoring Representations with an Adversary. In *International Conference on Learning Representations*.
- Gao, Y.; and Chaudhari, P. 2020. A Free-Energy Principle for Representation Learning. *arXiv preprint arXiv:2002.12406*.
- Hardt, M.; Price, E.; Price, E.; and Srebro, N. 2016. Equality of Opportunity in Supervised Learning. In *Advances in Neural Information Processing Systems*, volume 29, 3315–3323.
- He, Y.; Burghardt, K.; and Lerman, K. 2020. A Geometric Solution to Fair Representations. In *Proceedings of the*

- AAAI/ACM Conference on AI, Ethics, and Society, AIES '20, 279–285. New York, NY, USA.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2017. beta-vae: Learning Basic Visual Concepts with a Constrained Variational Framework. In *International Conference on Learning Representations*.
- Jaiswal, A.; Moyer, D.; Ver Steeg, G.; AbdAlmageed, W.; and Natarajan, P. 2020. Invariant Representations through Adversarial Forgetting. In *AAAI*, 4272–4279.
- Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2017. Inherent Trade-Offs in the Fair Determination of Risk Scores. In *Leibniz International Proceedings in Informatics, LIPIcs*, volume 67.
- Louizos, C.; Swersky, K.; Li, Y.; Welling, M.; and Zemel, R. 2016. The Variational Fair Autoencoder. In *International Conference on Learning Representations*.
- Madras, D.; Creager, E.; Pitassi, T.; and Zemel, R. 2018. Learning Adversarially Fair and Transferable Representations. In *Proceedings of the 35th International Conference on Machine Learning*, 3384–3393.
- McNamara, D.; Ong, C. S.; and Williamson, R. C. 2017. Provably Fair Representations. *arXiv preprint arXiv:1710.04394*.
- Mehrabi, N.; Huang, Y.; and Morstatter, F. 2020. Statistical Equity: A Fairness Classification Objective. *arXiv preprint arXiv:2005.07293*.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2019. A Survey on Bias and Fairness in Machine Learning. *arXiv preprint arXiv:1908.09635*.
- Menon, A. K.; and Williamson, R. C. 2018. The cost of fairness in binary classification. In *Conference on Fairness, Accountability and Transparency*, 107–118.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems*, volume 26, 3111–3119.
- Moyer, D.; Gao, S.; Brekelmans, R.; Galstyan, A.; and Ver Steeg, G. 2018. Invariant Representations without Adversarial Training. In *Advances in Neural Information Processing Systems*, volume 31, 9084–9093.
- Nguyen, X.; Wainwright, M. J.; and Jordan, M. I. 2010. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory* 56(11): 5847–5861.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *arXiv preprint arXiv:1807.03748*.
- Poole, B.; Ozair, S.; Van Den Oord, A.; Alemi, A.; and Tucker, G. 2019. On Variational Bounds of Mutual Information. In *Proceedings of the 36th International Conference on Machine Learning*.
- Rezende, D.; and Mohamed, S. 2015. Variational Inference with Normalizing Flows. In *Proceedings of the 32nd International Conference on Machine Learning*, 1530–1538.
- Rodríguez-Gálvez, B.; Thobaben, R.; and Skoglund, M. 2020. A Variational Approach to Privacy and Fairness. *arXiv preprint arXiv:2006.06332*.
- Roy, P. C.; and Boddeti, V. N. 2019. Mitigating Information Leakage in Image Representations: A Maximum Entropy Approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2586–2594.
- Sharifi-Malvajerdi, S.; Kearns, M.; and Roth, A. 2019. Average Individual Fairness: Algorithms, Generalization and Experiments. In *Advances in Neural Information Processing Systems*, volume 32, 8242–8251.
- Song, J.; Kalluri, P.; Grover, A.; Zhao, S.; and Ermon, S. 2019. Learning Controllable Fair Representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, 2164–2173.
- Srivastava, M.; Heidari, H.; and Krause, A. 2019. Mathematical Notions vs. Human Perception of Fairness: A Descriptive Approach to Fairness for Machine Learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2459–2468.
- Sylvester, J.; and Raff, E. 2018. What About Applied Fairness? In *Machine Learning: The Debates (ML-D) organized as part of the Federated AI Meeting (FAIM 2018)*.
- Tange, Ole. 2020. GNU Parallel 20200522 ('Kraftwerk'). doi:10.5281/zenodo.3841377.
- United States Equal Opportunity Employment Commission. 1978. Uniform Guidelines on Employee Selection Procedures.
- Veličković, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. In *International Conference on Learning Representations*.
- Williams, P. L.; and Beer, R. D. 2010. Nonnegative Decomposition of Multivariate Information. *arXiv preprint arXiv:1004.2515*.
- Xie, Q.; Dai, Z.; Du, Y.; Hovy, E.; and Neubig, G. 2017. Controllable Invariance through Adversarial Feature Learning. In *Advances in Neural Information Processing Systems*, volume 30, 585–596.
- Xu, Y.; Zhao, S.; Song, J.; Stewart, R.; and Ermon, S. 2020. A Theory of Usable Information under Computational Constraints. In *International Conference on Learning Representations*.
- Zafar, M. B.; Valera, I.; Ródriguez, M. G.; and Gummadi, K. P. 2017. Fairness Constraints: Mechanisms for Fair Classification. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 962–970.
- Zemel, R.; Wu, Y.; Swersky, K.; Pitassi, T.; and Dwork, C. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on Machine Learning*, 325–333.