

Classification Utility, Fairness, and Compactness via Tunable Information Bottleneck and Rényi Measures

Adam Gronowski¹, William Paul², Fady Alajaji³, *Senior Member, IEEE*,
Bahman Ghahesifard⁴, *Senior Member, IEEE*, and Philippe Burlina, *Senior Member, IEEE*

Abstract—Designing machine learning algorithms that are accurate yet fair, not discriminating based on any sensitive attribute, is of paramount importance for society to accept AI for critical applications. In this article, we propose a novel fair representation learning method termed the Rényi Fair Information Bottleneck Method (RFIB) which incorporates constraints for utility, fairness, and compactness (compression) of representation, and apply it to image and tabular data classification. A key attribute of our approach is that we consider - in contrast to most prior work - both demographic parity and equalized odds as fairness constraints, allowing for a more nuanced satisfaction of both criteria. Leveraging a variational approach, we show that our objectives yield a loss function involving classical Information Bottleneck (IB) measures and establish an upper bound in terms of two Rényi measures of order α on the mutual information IB term measuring compactness between the input and its encoded embedding. We study the influence of the α parameter as well as two other tunable IB parameters on achieving utility/fairness trade-off goals, and show that the α parameter gives an additional degree of freedom that can be used to control the compactness of the representation. Experimenting on three different image datasets (EyePACS, CelebA, and FairFace) and two tabular datasets (Adult and COMPAS), using both binary and categorical sensitive attributes, we show that on various utility, fairness, and compound utility/fairness metrics RFIB outperforms current state-of-the-art approaches.

Index Terms—Deep learning, fair representation learning, equalized odds, demographic parity, classification, information bottleneck (IB), Rényi divergence, Rényi cross-entropy.

I. INTRODUCTION

MACHINE learning algorithms are used for a variety of high stake applications such as loan approvals,

Manuscript received 6 October 2022; revised 21 May 2023 and 21 August 2023; accepted 20 November 2023. Date of publication 5 December 2023; date of current version 19 December 2023. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Joerg Kliewer. An earlier version of this paper was presented in part at the 2022 Canadian Workshop on Information Theory [DOI: 10.1109/CWIT55308.2022.9817669]. (*Corresponding author: Adam Gronowski.*)

Adam Gronowski and Fady Alajaji are with the Department of Mathematics and Statistics, Queen's University, Kingston, ON K7L 3N6, Canada (e-mail: adam.gronowski@queensu.ca; fa@queensu.ca).

William Paul and Philippe Burlina are with the Johns Hopkins University Applied Physics Laboratory, Laurel, MD 20723, USA (e-mail: william.paul@jhuapl.edu; philippe.burlina@jhuapl.edu).

Bahman Ghahesifard is with the Electrical and Computer Engineering Department, University of California at Los Angeles, Los Angeles, CA 90095 USA (e-mail: ghahesifard@ucla.edu).

Digital Object Identifier 10.1109/TIFS.2023.3340094

police allocation, admission of students, and disease diagnosis. In spite of its vast benefits, the use of automated algorithms that are not designed to also address potential bias and fairly serve members of diverse groups can lead to harm and exacerbate social inequities. The problem of developing algorithms that are both accurate and *fair*, i.e., algorithms that do not discriminate against individuals because of their gender, race, age, or other sensitive (protected) attributes, has now become paramount to the effective deployment of production-grade AI systems that could be accepted and adopted by society as a whole.

Fair machine learning methods have been developed for multiple domains such as automated healthcare diagnostics and treatments delivery [2], natural language processing [3], finance [4], and others [2], [5], [6], [7], [8]. One way to create such fair machine learning methods is through learning *fair representations* that can be used with existing architectures. These representations would allow for making accurate predictions while ensuring fairness. However, this leads to the difficulty that developing fair representations may involve trade-offs and is further complicated by the existence of various metrics for measuring fairness outcomes, often tailored towards different applications and settings.

Our work entails the development of a fair representation learning method that addresses a number of trade-offs. We consider the case where the sensitive attributes are directly accessible, and for the majority of experiments focus on the case where bias is caused by a lack of data for a protected subgroup. Unlike most prior studies that tend to focus on satisfying a single type of fairness constraint, we consider here how to jointly address and balance two of the arguably most important and widely used definitions for fairness, *demographic parity* and *equalized odds* [5], [9], [10]. We also examine different classical trade-offs between fairness and utility (commonly measured via accuracy). Trade-offs between utility and fairness arise as a result of interventions on models or data to make models more fair, which may yield decreased bias, but may also result in affecting utility. We study how these trade-offs are impacted by “compactness.” More specifically, we develop a variational approach taking into account different information-theoretic metrics that balance the above two constraints on fairness with utility and compactness. We also show how to analytically simplify the resulting loss function and relate it to the *Information Bottleneck (IB)* principle [11], and then exploit bounds to compute these metrics.

This work thus makes the following novel contributions:

- 1) We develop a novel variational method that balances a triplet of objectives, consisting of utility (accuracy), fairness (itself balancing two types of fairness constraints), and compression/compactness. Specifically, in contrast to prior work on fairness which narrowly uses either the demographic parity or equalized odds constraints, our loss includes both types of constraints. We relate analytically the resulting loss function to the classical IB method.
- 2) Operationally, we derive a simple upper bound on the mutual information between the data and its representation in terms of the Rényi divergence [12] and the Rényi cross-entropy [13], [14], [15] of order α . We study the effect of this added flexibility on achieving a balance between fairness and accuracy.
- 3) Using various datasets such as CelebA, EyePACS, FairFace, Adult, and COMPAS, and working with both binary and categorical sensitive variables, we compare with methods of record that intervene on the model or methods that intervene on training data. We show that our method overall performs best. We establish these comparisons via a number of metrics that measure utility and fairness individually or in a combined metric, including two different types of fairness constraints.

The rest of the article is organized as follows. We present related work in Section II and derive a cost function for our method in Section III. In Section IV, we describe the details of our implementation, describe the metrics and datasets used, and present extensive experimental results. Finally, we conclude the article with Section V.

II. RELATED WORK

A. Fairness Approaches

We summarize existing work on fair machine learning in three categories. For a broad-strokes categorization of fairness approaches, one can think along the lines of *interventions* made either on: 1) the model output, 2) the training data, or 3) the model itself. Each of these are motivated by different inductive biases.

1) *Interventions on Model Outputs Including Recalibration and Thresholding*: The inductive bias here is that irrespective of the cause of bias, fairness can be addressed at the output of the model. Some methods adjust the decision threshold to achieve equal odds across groups, as in [9], or re-calibrate the model for different subpopulations, as in [16]. These methods work well for categorical data, but they do not address the underlying causes of bias in the data and the model, which may be more important for image/video applications, as argued in [9]. Therefore, we explore a different approach.

2) *Data Interventions: Generative Models and Style Transfer*: The idea here is to modify the data to reduce bias since data imbalance can lead to biased models. Methods that alter the training data can censor sensitive information, blind the model to protected factors, augment the data, or re-balance the data by re-sampling or re-weighting. Data augmentation methods use generative models or style transfer/image translation. Some examples of generative models are based

on generative adversarial networks (GANs) [17], [18], [19], variational autoencoders [20], [21], or latent space optimization [19]. An example of image translation is [22]. However, these methods may have limitations in preserving the diversity of styles within subpopulations, which may be important for image/video applications. These methods are also related to domain adaptation and distributional shift scenarios.

Generating data with these methods has drawbacks as it may be hard to isolate and manipulate the protected factors without affecting other features (a problem called *entanglement*). Generative methods may also introduce artifacts in synthetic images, which could lower the performance and utility of the debiased models without significantly improving fairness. Based on these observations, [19] suggests that model interventions may be better than data interventions using generative models. These limitations motivate our approach, which belongs to the category of model interventions, discussed next.

3) *Interventions on Models via Adversarial and Variational Approaches*: Such methods are grounded on the inductive bias that dependence of the prediction on protected factors is a cause of lack of fair predictions; as a result, these methods generally aim to remedy this dependence at the encoding of the data, rendering them blind to protected factors.

Studies such as [23] and [24] use an adversarial network to penalize the prediction network if it could predict a protected factor. Other studies employ adversarial representation learning to remove protected factor information from latent representations, including [10], [25], [26], [27], [28]. Applications using this principle include [29] which uses adversarial learning to develop fair models of cardiovascular disease risk, while [30] explores the statistical properties of fair representation learning and [31] applies an adversarial approach for continuous features. Similar to our work, [32] and the more general recent work in [33] employ an information-theoretic approach to learn fair representations. More specifically in [33], the authors investigate an adversarial censoring method to generate universal fair representations that are least informative about sensitive attributes given a utility constraint. In contrast to the above techniques, our method does not use adversarial training.

In addition to the methods mentioned above, there exist non-adversarial methods that modify the model via incorporation of multiple constraints in a variational setting. Many of these are closely related to the IB method that we discuss next.

B. Information Bottleneck Methods

The IB method, originally proposed by Tishby et al. [11], is a method that seeks to develop representations that are both compact and expressive by minimizing and maximizing two mutual information terms. Alemi et al. first developed a variational approximation of the IB method by parameterizing it using neural networks and this was followed by multiple variations such as the nonlinear information bottleneck [34] and conditional entropy bottleneck [35]. Many recent generalizations of the IB method have been developed including [36],

[37], [38], [39], and [40] while [41] and [42] investigated its connections to deep learning theory and privacy applications.

Techniques related to the information bottleneck have been used for fair representation learning. First proposed by [5], fair representation learning consists of mapping input data to an intermediate representation that remains informative but discards unwanted information that could reveal the protected sensitive factors. This is related to the IB problem and several works have explored the connection between the two, such as [43] and [44], where fair representations are acquired through the minimization and maximization of various mutual information terms.

Our work is closest to [44] but we depart from it in several important ways, including through the use of a more general loss formulation, satisfying broader constraints of fairness, classifying images rather than focusing solely on tabular data, and entailing the use of Rényi divergence and Rényi cross-entropy [13], [14], [15]. To our knowledge we are the first to use Rényi divergence and cross-entropy for fair representation learning but there have been several recent works based on Rényi information measures and their variants. These include an IB problem under a Rényi entropy complexity constraint [40], bounding the generalization error in learning algorithms [45], Rényi divergence variational inference [46], Rényi differential privacy [47] and the analysis and development of GANs [14], [48], [49], [50], [51], [52]. In addition, Baharlouei et al. [53] developed a fair representation method but one based on Rényi correlation rather than divergence.

C. Connections to Other Work

Fairness is related to the problem of domain adaptation (DA) [54], [55], [56], [57] that consists of training a neural network on a source dataset to obtain good accuracy on a target dataset that is different from the source. This is especially true for the case of a severe data imbalance that we consider in this work where training data is completely missing for a protected subgroup. In this case, achieving fairness is similar to the DA problem of improving performance on a complete target dataset that includes all groups after training on an incomplete source dataset.

Our work is also related to group distributionally robust optimization (GDRO) methods [58], [59] that address the problem of performance disparity among different subgroups by minimizing the worst-case loss among different subgroups. Reducing these accuracy differences among subgroups is something our method also addresses, but while this is the sole objective for GDRO methods, we consider this problem in relation to multiple other fairness criteria.

There are other studies about finding a balance between accuracy and fairness such as Zhang et al. [60]. However, they achieved this objective by finding early stopping criteria rather than through a fair representation preprocessing method like we use here; also, unlike our method, finding a balance between multiple fairness constraints is not investigated.

Finally, there are hybrid methods, such as the one proposed by Paul et al. in [19], that use a combination of the previously discussed techniques of interventions on data, model,

or output. However, such techniques are quite limited in scope compared to the more varied objectives considered here which involve jointly achieving utility, fairness, and compactness.

III. METHODS

We take an information-theoretic approach to fairness (refer to [61] and [62] for background knowledge on standard information measures such as Shannon entropy and mutual information). We represent input data as a random variable $X \in \mathcal{X}$ and sensitive information as a random variable $S \in \mathcal{S}$. Our goal is to use the data X to predict a target $Y \in \mathcal{Y}$ but in a way that is uninfluenced by the sensitive information S . To reach this objective, we adopt a variational approach, which we call *Rényi Fair Information Bottleneck* (RFIB), to encode the data into a new representation $Z \in \mathcal{Z}$. The representation can then be used with existing deep learning model architectures to draw inferences about Y . As we compute Z from X without accessing data from Y or S , we can make a standard assumption similar to [63] and assume that the Markov chain

$$(Y, S) \rightarrow X \rightarrow Z \rightarrow \hat{Y}$$

holds, where \hat{Y} is the prediction of Y based on Z (\hat{Y} is a function of Z). For simplicity, we assume in this section that all random variables are discrete, though a similar derivation holds for a mix of continuous and discrete random variables.

A. Fairness Defined

Among the three principal definitions of fairness – *demographic parity*, *equalized odds*, and *equality of opportunity* – we focus on addressing both demographic parity and equalized odds since a) equalized odds is a stronger constraint than equality of opportunity (equalized odds requires both an equal true positive rate and equal false positive rate across the sensitive variable, whereas equality of opportunity only requires an equal false positive rate), and b) demographic parity, also called statistical parity, is an altogether different type of constraint compared to the former two constraints in that the requirement of independence does not involve the actual target label value.

For demographic parity, the goal is for the model's prediction \hat{Y} to be independent of the sensitive variable S [5], i.e.,

$$P(\hat{Y} = \hat{y}) = P(\hat{Y} = \hat{y} | S = s), \quad (1)$$

for all s, \hat{y} , while for equalized odds the goal is to achieve this independence by conditioning on the actual target Y [9], i.e.,

$$P(\hat{Y} = \hat{y} | Y = y) = P(\hat{Y} = \hat{y} | S = s, Y = y) \quad (2)$$

for all s, \hat{y}, y .

B. Lagrangian Formulation

We formulate a Lagrangian by minimizing and maximizing various mutual information terms, similar to [63] and [64]. To encourage equalized odds, we want as per (2) that \hat{Y} and S are conditionally independent given Y (written as $(\hat{Y} \perp\!\!\!\perp S) | Y$); this is equivalent to requiring that $I(\hat{Y}; S | Y) = 0$.

To achieve this objective, we minimize the proxy measure $I(Z; S|Y)$ over $P_{Z|X}$. This is justified by the following: by the chain rule for mutual information, we can write

$$\begin{aligned} I(Z, \hat{Y}; S|Y) &= I(Z; S|Y) + \underbrace{I(\hat{Y}; S|Z, Y)}_{=0} \\ &= I(\hat{Y}; S|Y) + \underbrace{I(Z; S|\hat{Y}, Y)}_{\geq 0}, \end{aligned} \quad (3)$$

where $I(\hat{Y}; S|Z, Y) = 0$, since by the Markov chains $(Y, S) \rightarrow Z \rightarrow \hat{Y}$ and $Y \rightarrow Z \rightarrow \hat{Y}$,¹ we have for all \hat{y}, z, s, y that

$$\begin{aligned} P(\hat{Y} = \hat{y}|Z = z, S = s, Y = y) &= P(\hat{Y} = \hat{y}|Z = z) \\ &= P(\hat{Y} = \hat{y}|Z = z, Y = y) \end{aligned}$$

and hence $(\hat{Y} \perp\!\!\!\perp S) | (Z, Y) \Leftrightarrow I(\hat{Y}; S|Z, Y) = 0$. We therefore obtain from (3) that

$$0 \leq I(\hat{Y}; S|Y) \leq I(Z; S|Y).$$

So by minimizing $I(Z; S|Y)$ over $P_{Z|X}$, we are squeezing $I(\hat{Y}; S|Y)$ closer to zero. In the ideal case where we can drive $I(Z; S|Y)$ to zero, we readily obtain that

$$I(\hat{Y}; S|Y) = 0 \Leftrightarrow (\hat{Y} \perp\!\!\!\perp S) | Y$$

as ultimately desired in equalized odds.

To both obtain good classification accuracy (utility) and help promote demographic parity, we maximize $I(Z; Y|S)$. For demographic parity, we want as per (1) that $\hat{Y} \perp\!\!\!\perp S \Leftrightarrow I(\hat{Y}; S) = 0$, while for utility we want high $I(Y; \hat{Y})$ which we address by maximizing $I(Z; Y)$. To achieve $I(\hat{Y}; S) = 0$ we require $I(Z; S) = 0$ (i.e., $Z \perp\!\!\!\perp S$) since

$$0 \leq I(S; \hat{Y}) \leq I(S; Z)$$

by the Data Processing Inequality [61], [62] as $S \rightarrow Z \rightarrow \hat{Y}$ holds. Now by the mutual information chain rule, we can write

$$\begin{aligned} I(Z; Y, S) &= I(Z; Y) + I(Z; S|Y) \\ &= I(Z; S) + I(Z; Y|S). \end{aligned} \quad (4)$$

In the ideal condition where $I(Z; S | Y) = 0$ (i.e., equalized odds is achieved), a necessary condition for $I(Z; S) = 0$ is that $I(Z; Y|S) = I(Z; Y)$ which we want to maximize. Thus maximizing $I(Z; Y|S)$ promotes both demographic parity and high utility.

Finally, we minimize $I(Z; X|S, Y)$, a compression term similar to one from the IB problem [63]. This minimization further encourages Z to discard information irrelevant for drawing predictions about Y , hence improving generalization capability and reducing the risk of keeping nuisances (irrelevant data that can degrade performance).

Combining all above mutual information terms leads to a Lagrangian, \mathcal{L} , that we seek to minimize over the encoding conditional distribution $P_{Z|X}$. The Lagrangian is given by

$$\mathcal{L} = I(Z; S|Y) + I(Z; X|S, Y) - \lambda_1 I(Z; Y) - \lambda_2 I(Z; Y|S),$$

¹Note that these two Markov chains are implied by the original Markov chain formulation $(Y, S) \rightarrow X \rightarrow Z \rightarrow \hat{Y}$.

$$- \lambda_2 I(Z; Y|S), \quad (5)$$

where λ_1 and λ_2 are hyper-parameters. Developing this Lagrangian, we have that

$$\begin{aligned} \mathcal{L} &= H(Z|Y) - H(Z|S, Y) + H(Z|S, Y) \\ &\quad - H(Z|X, S, Y) - \lambda_1 I(Z; Y) - \lambda_2 I(Z; Y|S) \\ &= H(Z|Y) - H(Z|X) - \lambda_1 I(Z; Y) - \lambda_2 I(Z; Y|S) \\ &= H(X) - H(Z, X) - [H(Y) - H(Z, Y)] \\ &\quad - \lambda_1 I(Z; Y) - \lambda_2 I(Z; Y|S) \\ &= I(Z; X) - I(Z; Y) - \lambda_1 I(Z; Y) - \lambda_2 I(Z; Y|S) \\ &= I(Z; X) - (\lambda_1 + 1)I(Z; Y) - \lambda_2 I(Z; Y|S), \end{aligned} \quad (6)$$

where $H(\cdot)$ denotes entropy, and the second equality follows from the Markov chain assumption $(Y, S) \rightarrow X \rightarrow Z$. Hence, we have shown that the Lagrangian \mathcal{L} admits a simpler equivalent expression given by

$$\mathcal{L} = I(Z; X) - \beta_1 I(Z; Y) - \beta_2 I(Z; Y|S), \quad (7)$$

where $\beta_1 = \lambda_1 + 1 \geq 1$ and $\beta_2 = \lambda_2 \geq 0$. This simpler Lagrangian is easier to compute while exactly maintaining the properties of the original one. It also reveals a direct relation of the original Lagrangian with the first two terms being exactly equivalent to the ‘‘classical IB’’ formulation. Note that both the $I(Z; X)$ term and $I(Z; Y|S)$ term contribute to fairness. Minimizing $I(Z; X)$ results in increased compression which can lead to increased fairness at the possible expense of accuracy. The $I(Z; Y|S)$ term contributes to both fairness through the conditioning on S and to accuracy through the maximization of the information between Z and Y . The hyper-parameters β_1 and β_2 control trade-offs between accuracy (or ‘‘utility’’) and fairness, as a higher weight on β_1 and β_2 reduces the influence of the $I(Z; X)$ compression term. As $I(Z; Y)$ is partially derived from the $I(Z; S|Y)$ term designed to improve equalized odds, using a higher β_1 over β_2 should give more priority to improving equalized odds, whereas a higher β_2 should result in improved demographic parity. This allows for more nuanced outcomes compared to other methods that focus rigidly on a single fairness metric. It is also possibly an interesting tool for policy makers to translate those more balanced and nuanced versions of fairness into an ‘‘engineered system.’’

C. Variational Bounds

We use a variational approach to develop bounds on the three terms in the Lagrangian \mathcal{L} in (7), finding lower bounds for the terms to be maximized and an upper bound for the term to be minimized. The Markov chain property $(Y, S) \rightarrow X \rightarrow Z$ results in the joint distribution P_{SYXZ} factoring as $P_{SYX}P_{Z|X}$.

The distribution $P_{Z|X}$ is a parametric stochastic encoder to be designed while all other distributions are fully determined by the joint data distribution $P_{S, X, Y}$, the encoder, and the Markov chain constraint. To simplify notation, we simply write $P_{Z|X}$ rather than including the parameter $P_{Z|X, \theta}$, with θ denoting network weights. Computing the mutual information

terms requires the usually intractable or difficult to compute distributions $P_{Y|S,Z}$, $P_{Y|Z}$, and P_Z [63], [65]; we thus replace them with variational approximations $Q_{Y|S,Z}$, $Q_{Y|Z}$ and Q_Z , respectively. We next derive a simple upper bound for the compression term $I(Z; X)$ with the novel use of Rényi's divergence and cross-entropy of order α , which provide a tunable extra degree of freedom for the variational version of \mathcal{L} :

$$\begin{aligned} I(Z; X) &= \min_{Q_Z} D_{KL}(P_{XZ} \| P_X Q_Z) \\ &\leq E_{P_X} D_{KL}(P_{Z|X} \| Q_Z) \\ &\leq \begin{cases} E_{P_X} D_\alpha(P_{Z|X} \| Q_Z) & \text{if } \alpha \geq 1 \\ E_{P_X} H_\alpha(P_{Z|X}; Q_Z) & \text{if } \alpha < 1. \end{cases} \end{aligned} \quad (8)$$

The first identity follows from the non-negativity of the Kullback-Leibler (KL) divergence (see for example [62, Problem 2.18], [32], [44], [63]).

For the case of $\alpha \geq 1$ in the final inequality given in (8), we take the Rényi divergence $D_\alpha(\cdot \| \cdot)$ of order α (e.g., see [12], [66]), rather than the KL divergence as typically done in the literature, where

$$D_\alpha(P \| Q) = \frac{1}{\alpha - 1} \log \left(\sum_{x \in \mathcal{X}} P(x)^\alpha Q(x)^{1-\alpha} \right) \quad (9)$$

for $\alpha > 0$, $\alpha \neq 1$ and distributions P and Q with common support \mathcal{X} .² By the continuity property of D_α in α , see [66], we define its extended orders at $\alpha = 1$ and $\alpha = 0$ as

$$D_1(P \| Q) := \lim_{\alpha \rightarrow 1} D_\alpha(P \| Q) = D_{KL}(P \| Q) \quad (10)$$

and

$$D_0(P \| Q) := \lim_{\alpha \rightarrow 0} D_\alpha(P \| Q) = -\log Q(x : P(x) > 0), \quad (11)$$

which is equal to 0 when P and Q share a common support. This upper bound in (8) readily holds for $\alpha \geq 1$ since D_α is non-decreasing in α and $D_1(P \| Q) = D_{KL}(P \| Q)$.

For the case of $\alpha < 1$ in (8), we take the Rényi cross-entropy $H_\alpha(\cdot; \cdot)$ of order α , used recently in [14] to generalize the original loss function of GANs [67]:³

$$H_\alpha(P_{Z|X=x}; Q_Z) = \frac{1}{1-\alpha} \log \left(\sum_{z \in \mathcal{Z}} P_{Z|X}(z|x) Q_Z(z)^{\alpha-1} \right). \quad (12)$$

We hence justify the inequality in (8) when $\alpha < 1$ as follows:

$$\begin{aligned} &E_{P_X} D_{KL}(P_{Z|X} \| Q_Z) \\ &= E_{P_X} \sum_{z \in \mathcal{Z}} P_{Z|X}(z|x) \log \frac{P_{Z|X}(z|x)}{Q_Z(z)} \\ &= E_{P_X} \underbrace{\left[\sum_{z \in \mathcal{Z}} P_{Z|X}(z|x) \log \frac{1}{Q_Z(z)} \right]}_{=H(P_{Z|X=x}; Q_Z)} \end{aligned}$$

²When P and Q are both probability density functions, then $D_\alpha(P \| Q) = \frac{1}{\alpha-1} \log \left(\int_{\mathcal{X}} P(x)^\alpha Q(x)^{1-\alpha} dx \right)$.

³When $P_{Z|X=x}$ and Q_Z are both probability density functions, then the (differential) Rényi cross-entropy is given by $h_\alpha(P_{Z|X=x}; Q_Z) = \frac{1}{1-\alpha} \log \left(\int_{\mathcal{Z}} P_{Z|X}(z|x) Q_Z(z)^{\alpha-1} dz \right)$.

$$\begin{aligned} &+ E_{P_X} \underbrace{\left[\sum_{z \in \mathcal{Z}} P_{Z|X}(z|x) \log P_{Z|X}(z|x) \right]}_{=-H(Z|X=x) \leq 0} \\ &\leq E_{P_X} H(P_{Z|X=x}; Q_Z) \\ &\leq E_{P_X} H_\alpha(P_{Z|X=x}; Q_Z), \end{aligned} \quad (13)$$

where $H(\cdot; \cdot)$ is the Shannon cross-entropy and the final inequality follows since Rényi's cross-entropy is non-increasing in α and $\lim_{\alpha \rightarrow 1} H_\alpha(P_{Z|X=x}; Q_Z) = H(P_{Z|X=x}; Q_Z)$ [14].

Remark 1: The above derivation is carried for discrete distributions. It still holds for continuous distributions except that for the first inequality in (13), we need to ensure that the conditional differential entropy $h(Z|X=x)$ is non-negative. More specifically, if $P_{Z|X=x}$ is an uncorrelated Gaussian vector (which we use below in the derivation of the Rényi divergence and ensuing experiments), we ensure the non-negativity of the vector's conditional differential entropy by selecting large enough variance components σ_i^2 (see Footnote 4). This also guarantees the non-negativity of the differential Shannon and Rényi cross-entropies: $h(P_{Z|X=x}; Q_Z)$ and $h_\alpha(P_{Z|X=x}; Q_Z)$ with $\alpha < 1$.

Remark 2: If we take $\alpha < 1$ for the Rényi divergence or $\alpha \geq 1$ for the Rényi cross-entropy in (8), we no longer have an upper bound on $I(Z; X)$. But in these cases, the terms in (8) can be considered as approximations to the compression term $I(Z; X)$ that are α -tunable. In other words, as the Rényi divergence is non-decreasing in α and the Rényi cross-entropy is non-increasing in α , choosing values of $\alpha < 1$ or $\alpha \geq 1$, respectively, will result in a smaller approximation for $I(Z; X)$. This allows us to fully control the weight we put on that term compared to the others in the cost function in (7) by adjusting α , being able to adjust it all the way down to zero if desired.

We can similarly leverage the non-negativity of KL divergence to get lower bounds on $I(Z; Y)$ and $I(Z; Y|S)$:

$$I(Z; Y) \geq \mathbb{E}_{P_{Y,Z}} [\log Q_{Y|Z}(Y|Z)] + H(Y), \quad (14)$$

$$I(Z; Y|S) \geq \mathbb{E}_{P_{S,Y,Z}} [\log Q_{Y|S,Z}(Y|S, Z)] + H(Y|S). \quad (15)$$

As the entropy $H(Y)$ and conditional entropy $H(Y|S)$ of the labels do not depend on the parameterization they can be ignored for the optimization.

D. Computing the Bounds

To compute the bounds in practice we use the reparameterization trick [65]. Modeling $P_{Z|X}$ as a density, we let $P_{Z|X} dZ = P_E dE$, where E is a Gaussian (normal) random variable with zero mean and unit variance and $Z = f(X, E)$ is a deterministic function, allowing us to backpropagate gradients and optimize the parameter via gradient descent. We use the data's empirical densities to estimate $P_{X,S}$ and $P_{X,Y,S}$.

Considering a batch $D = \{x_i, s_i, y_i\}_{i=1}^N$ this finally leads to the following RFIB cost function to minimize:

$$J_{\text{RFIB}} = \frac{1}{N} \sum_{i=1}^N \left[F_\alpha(P_{Z|X=x_i}, Q_Z) \right]$$

$$\begin{aligned} & -\beta_1 \mathbb{E}_E [\log(Q_{Y|Z}(y_i | f(x_i, E)))] \\ & -\beta_2 \mathbb{E}_E [\log(Q_{Y|S,Z}(y_i | s_i, f(x_i, E)))] \end{aligned} \quad (16)$$

where we estimate the expectation over E using a single Monte Carlo sample (as in [63]) and

$$F_\alpha(P_{Z|X=x_i}, Q_Z) = \begin{cases} D_\alpha(P_{Z|X=x_i} \| Q_Z) & \text{if } \alpha \geq 1 \\ H_\alpha(P_{Z|X=x_i}; Q_Z) & \text{if } \alpha < 1. \end{cases}$$

We will compare the performance of our RFIB system with both the IB [63] and the *conditional fairness bottleneck* (CFB) [44] schemes. Note that if we set $\alpha = 1$ and $\beta_2 = 0$ in (7), or in the variational loss function (16), we recover the IB system. Furthermore, setting $\alpha = 1$ and forcing $\beta_1 = 0$ (by relaxing its range) yields the CFB loss function. We will therefore use the parameters pairs of $(\alpha = 1, \beta_1 = 0)$ and $(\alpha = 1, \beta_2 = 0)$ as corresponding to the CFB and IB systems, respectively.

E. Derivation of $D_\alpha(P_{Z|X} \| Q_Z)$ for Gaussians

Here we calculate the Rényi divergence term of our cost function when $P_{Z|X}$ and Q_Z are multivariate Gaussian distributions with $P_{Z|X}$ having a diagonal covariance structure while Q_Z being a spherical Gaussian. More specifically, we let

$$\begin{aligned} P_{Z|X} &= \mathcal{N}(Z | \mu_{enc}(X), \text{diag}(\sigma_{enc}^2(X))) \\ Q_Z &= \mathcal{N}(Z | \underline{0}, \gamma^2 I_d), \end{aligned}$$

where μ_{enc} and σ_{enc}^2 are d -dimensional mean and variance vectors (that depend on X), $\text{diag}(\sigma_{enc}^2(X))$ is a $d \times d$ diagonal matrix with the entries of σ_{enc}^2 on the diagonal, $\underline{0}$ is the all-zero vector of size d , γ is a positive scalar, and I_d is the d -dimensional identity matrix. For simplicity of notation, in the rest of the article we write $\mu_{enc}(X)$ as μ_{enc} and $\text{diag}(\sigma_{enc}^2(X))$ as $\sigma_{enc}^2 I_d$.

Starting with the closed-form expression of the Rényi divergence of order α ($\alpha > 0, \alpha \neq 1$) derived in [68] and [69], we have

$$\begin{aligned} D_\alpha(P_{Z|X} \| Q_Z) &= \frac{\alpha}{2} \left(\mu'_{enc} [(\Sigma_\alpha)^*]^{-1} \mu_{enc} \right) \\ &\quad - \frac{1}{2(\alpha-1)} \ln \frac{|(\Sigma_\alpha)^*|}{|\sigma_{enc}^2 I_d|^{1-\alpha} |\gamma^2 I_d|^\alpha} \end{aligned} \quad (17)$$

where μ'_{enc} is the transpose of μ_{enc} ,

$$(\Sigma_\alpha)^* = \alpha \gamma^2 I_d + (1-\alpha) \sigma_{enc}^2 I_d,$$

and

$$\alpha [\sigma_{enc}^2 I_d]^{-1} + (1-\alpha) [\gamma^2 I_d]^{-1}$$

is positive definite. Then

$$\frac{\alpha}{2} \left(\mu'_{enc} [(\Sigma_\alpha)^*]^{-1} \mu_{enc} \right) = \frac{\alpha}{2} \sum_{i=1}^d \frac{\mu_i^2}{\alpha \gamma^2 + (1-\alpha) \sigma_i^2}, \quad (18)$$

where μ_i and σ_i are the i th components of μ_{enc} and σ_{enc} and

$$\frac{1}{2(\alpha-1)} \ln \frac{|(\Sigma_\alpha)^*|}{|\sigma_{enc}^2 I_d|^{1-\alpha} |\gamma^2 I_d|^\alpha}$$

$$\begin{aligned} &= \frac{1}{2(\alpha-1)} \ln \frac{\begin{vmatrix} \alpha \gamma^2 + (1-\alpha) \sigma_1^2 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \alpha \gamma^2 + (1-\alpha) \sigma_d^2 \end{vmatrix}}{\begin{vmatrix} \sigma_1^2 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_d^2 \end{vmatrix}^{1-\alpha} \begin{vmatrix} \gamma^2 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \gamma^2 \end{vmatrix}^\alpha} \\ &= \frac{1}{2(\alpha-1)} \ln \frac{\prod_{i=1}^d [\alpha \gamma^2 + (1-\alpha) \sigma_i^2]}{(\prod_{i=1}^d \sigma_i^2)^{1-\alpha} (\prod_{i=1}^d \gamma^2)^\alpha} \\ &= \frac{1}{2(\alpha-1)} \sum_{i=1}^d \ln \frac{\alpha \gamma^2 + (1-\alpha) \sigma_i^2}{\sigma_i^{2(1-\alpha)} \gamma^{2\alpha}}, \end{aligned} \quad (19)$$

yielding that

$$\begin{aligned} D_\alpha(P_{Z|X} \| Q_Z) &= \frac{\alpha}{2} \sum_{i=1}^d \frac{\mu_i^2}{\alpha \gamma^2 + (1-\alpha) \sigma_i^2} \\ &\quad - \frac{1}{2(\alpha-1)} \sum_{i=1}^d \ln \frac{\alpha \gamma^2 + (1-\alpha) \sigma_i^2}{\sigma_i^{2(1-\alpha)} \gamma^{2\alpha}}. \end{aligned} \quad (20)$$

Since we require the matrix $\alpha [\sigma_{enc}^2 I_d]^{-1} + (1-\alpha) [\gamma^2 I_d]^{-1}$ to be positive definite, the above Rényi divergence expression is valid for

$$\alpha \gamma^2 + (1-\alpha) \sigma_i^2 > 0, \quad i = 1, \dots, d,$$

or equivalently (recalling that $\alpha > 0, \alpha \neq 1$) for

$$\begin{aligned} 0 < \alpha < 1, \sigma_i^2 > 0, \quad i = 1, \dots, d, \\ \alpha > 1, \sigma_i^2 < \frac{\alpha \gamma^2}{\alpha - 1}, \quad i = 1, \dots, d. \end{aligned} \quad (21)$$

We finish this section with a remark.

Remark 3 (Limit as $\alpha \rightarrow 1$): Taking the limit of the Rényi divergence as $\alpha \rightarrow 1$ (or setting $\alpha = 1$ as an extended order) we recover, as expected, the KL divergence expression between Gaussians:

$$\begin{aligned} & \lim_{\alpha \rightarrow 1} D_\alpha(P_{Z|X} \| Q_Z) \\ &= \lim_{\alpha \rightarrow 1} \left[\frac{\alpha}{2} \sum_{i=1}^d \frac{\mu_i^2}{\alpha \gamma^2 + (1-\alpha) \sigma_i^2} \right. \\ &\quad \left. - \frac{1}{2(\alpha-1)} \sum_{i=1}^d \ln \frac{\alpha \gamma^2 + (1-\alpha) \sigma_i^2}{\sigma_i^{2(1-\alpha)} \gamma^{2\alpha}} \right] \\ &= \frac{1}{2} \sum_{i=1}^d \frac{\mu_i^2}{\gamma^2} - \lim_{\alpha \rightarrow 1} \frac{1}{2} \left[\sum_{i=1}^d \frac{\gamma^2 - \sigma_i^2}{\alpha \gamma^2 + (1-\alpha) \sigma_i^2} \right. \\ &\quad \left. + 2 \ln \sigma_i - \ln \gamma^2 \right] \\ &= -\frac{1}{2} \sum_{i=1}^d \left[\ln \sigma_i^2 - \ln \gamma^2 + 1 - \frac{\sigma_i^2}{\gamma^2} - \frac{\mu_i^2}{\gamma^2} \right] \\ &= D_{KL}(P_{Z|X} \| Q_Z). \end{aligned} \quad (22)$$

F. Derivation of $h_\alpha(P_{Z|X} \| Q_Z)$ for Gaussians

Here we calculate the (differential) Rényi cross-entropy term of our cost function when $P_{Z|X}$ and Q_Z are multivariate Gaussian distributions with $P_{Z|X}$ having a diagonal covariance structure while Q_Z being a spherical Gaussian. Again we let

$$P_{Z|X} = \mathcal{N}\left(Z | \mu_{enc}(X), \text{diag}(\sigma_{enc}^2(X))\right)$$

$$Q_Z = \mathcal{N}(Z | \underline{0}, \gamma^2 I_d),$$

where μ_{enc} and σ_{enc}^2 are d -dimensional mean and variance vectors (that depend on X), $\text{diag}(\sigma_{enc}^2(X))$ is a $d \times d$ diagonal matrix with the entries of σ_{enc}^2 on the diagonal, $\underline{0}$ is the all-zero vector of size d , γ is a positive scalar, and I_d is the d -dimensional identity matrix. For simplicity, we write $\mu_{enc}(X)$ as μ_{enc} and $\text{diag}(\sigma_{enc}^2(X))$ as $\sigma_{enc}^2 I_d$.

Starting with the closed-form expression of the Rényi cross-entropy derived in [15], we have

$$h_\alpha(P_{Z|X}; Q_Z)$$

$$= \frac{1}{2-2\alpha} \left(-\ln |A| \left| \sigma_{enc}^2 I_d \right| + (1-\alpha) \ln(2\pi)^d \left| \gamma^2 I_d \right| - G \right) \quad (23)$$

where

$$A = (\sigma_{enc}^2 I_d)^{-1} + (\alpha - 1)(\gamma^2 I_d)^{-1}$$

is positive definite, and

$$G = \mu'_{enc} (\sigma_{enc}^2 I_d)^{-1} \mu_{enc}$$

$$- \mu'_{enc} (\sigma_{enc}^2 I_d)^{-1} A^{-1} (\sigma_{enc}^2 I_d)^{-1} \mu_{enc},$$

with μ'_{enc} denoting the transpose of μ_{enc} . Then

$$-\ln |A| \left| \sigma_{enc}^2 I_d \right|$$

$$= -\ln \begin{vmatrix} \frac{\gamma^2 + \sigma_1^2(\alpha-1)}{\sigma_1^2 \gamma^2} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \frac{\gamma^2 + \sigma_d^2(\alpha-1)}{\sigma_d^2 \gamma^2} \end{vmatrix} \left| \begin{vmatrix} \sigma_1^2 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \sigma_d^2 \end{vmatrix} \right|$$

$$= -\sum_{i=1}^d \ln \left(\frac{\gamma^2 + \sigma_i^2(\alpha-1)}{\gamma^2} \right) \quad (24)$$

and

$$-G = -\sum_{i=1}^d \frac{\mu_i^2}{\sigma_i^2} + \sum_{i=1}^d \frac{\mu_i^2}{\sigma_i^4} \frac{\sigma_i^2 \gamma^2}{\gamma^2 + \sigma_i^2(\alpha-1)}, \quad (25)$$

where μ_i and σ_i are the i th components of μ_{enc} and σ_{enc} , yielding that

$$h_\alpha(P_{Z|X}; Q_Z) = -\frac{1}{2-2\alpha} \sum_{i=1}^d \ln \left(\frac{\gamma^2 + \sigma_i^2(\alpha-1)}{\gamma^2} \right)$$

$$+ \frac{1}{2-2\alpha} (1-\alpha) \gamma^{2d} \ln(2\pi)^d$$

$$- \frac{1}{2-2\alpha} \sum_{i=1}^d \frac{\mu_i^2}{\sigma_i^2} - \frac{\mu_i^2}{\sigma_i^4} \frac{\sigma_i^2 \gamma^2}{\gamma^2 + \sigma_i^2(\alpha-1)}. \quad (26)$$

Since we require the matrix A to be positive definite,⁴ the above Rényi cross-entropy expression is valid for

$$\gamma^2 + \sigma_i^2(\alpha-1) > 0, \quad i = 1, \dots, d,$$

or equivalently (recalling that $\alpha > 0$, $\alpha \neq 1$) for

$$0 < \alpha < 1, \quad \sigma_i^2 < \frac{\gamma^2}{1-\alpha},$$

$$\alpha > 1, \quad \sigma_i^2 > 0, \quad i = 1, \dots, d. \quad (27)$$

IV. EXPERIMENTS

We conduct experiments on three different image datasets: CelebA, FairFace, and EyePACS, and on the tabular Adult and COMPAS datasets.⁵ In this section, we detail the implementation steps of our method, describe the metrics and the datasets we used, and explain how the experiments were performed and present results. We also present a uniform manifold approximation and projection (UMAP) clustering analysis that visualizes the effects of the hyper-parameter α .

A. Data

We next describe the image and tabular datasets (with the target Y taken to be binary). Samples from the image datasets are shown in Fig. 1.

1) *CelebA*: The CelebA dataset [70] contains 202,599 celebrity faces that each have 40 binary attributes. We use age as our prediction target Y , with $Y = 1$ indicating the person is old (according to the perception of the dataset's annotators), and are interested in skin tone as the sensitive attribute S . As this is not included in the dataset we instead use the Individual Topology Angle (ITA) [71] as a proxy, which was found to correlate with the Melanin Index, frequently used in dermatology to classify human skin on the Fitzpatrick scale. As in [19] and [72], we compute ITA via

$$\text{ITA} = \frac{180}{\pi} \arctan \left(\frac{L-50}{b} \right), \quad (28)$$

where L is luminance and b is yellowness in CIE-Lab space. We then binarize ITA where an ITA of ≤ 28 is taken to mean dark skin, matching category thresholds used in [73] and [19].

2) *FairFace*: The FairFace dataset [74] consists of 108,501 face images labeled with race, gender, and age. We consider gender as the target Y and race as sensitive information S . The dataset contains several different categories for race. We include experimental results where we binarize the race labels to Black and non-Black. We also include experiments where S is a categorical variable with three race categories: Black, Latino, and Other, with the Other group consisting of a mix of Indian, East Asian, Southeast Asian, Middle Eastern, and White races.

⁴As stated in Remark 1, we also require $\prod_{i=1}^d \sigma_i^2 \geq \frac{1}{(2\pi e)^d}$ to ensure that $h(Z | X = x) = \frac{1}{2} \log((2\pi e)^d \prod_{i=1}^d \sigma_i^2) \geq 0$. This is satisfied in our simulations as d is large (we use $d = 64$).

⁵The code of our RFIB algorithm and all dataset partitions are available in full detail at the site: <https://github.com/AGronowski/RFIB-Code>




Fig. 1. Examples of images taken respectively from CelebA (top row), FairFace (middle row), and EyePACS (bottom row) datasets. For the EyePACS images, for the left image $(Y, S) = (1, 0)$, for the middle $(Y, S) = (0, 0)$, and for the right $(Y, S) = (0, 1)$.

3) *EyePACS*: The EyePACS dataset [75] is sourced from the Kaggle Diabetic Retinopathy challenge. It consists of 88,692 retinal fundus images of individuals potentially suffering from diabetic retinopathy (DR), an eye disease associated with diabetes that is one of the leading causes of visual impairment worldwide. The dataset contains 5 categories of images based on the severity of the disease, with 0 being completely healthy and 4 being the most severe form of the disease. Similar to [19], we binarize this label into our prediction target Y , with $Y = 1$ corresponding to categories 2-4, considered a positive, referable case for DR and $Y = 0$ corresponding to categories 0-1, considered healthy.

In our experiments we consider skin tone as the sensitive attribute. As with the CelebA dataset, we use ITA as a proxy for skin tone, with S a binary variable that denotes whether or not the ITA of the fundus is ≤ 19 , denoting that the individual has dark skin, as done by [19]. This has the advantage of being significantly easier to determine compared to the expensive and time consuming process of having a clinician manually annotate images. We compare these results with additional experiments we run using a smaller partition of the dataset with the race labels manually added by an ophthalmologist as done in [76], determined based on factors that might correlate with race such as the darkness of pigmentation in the fundus, thickness of blood vessels, and ratio of the optic cup size to optic disk size.

4) *Adult*: The Adult dataset [77], also known as the Census Income dataset, contains 48,842 records of census data from 1994. Each sample contains 15 binary features such as gender, age, or income. We take our prediction target Y to be income, with $Y = 1$ corresponding to an income higher than \$50,000. We take the sensitive attribute S to be gender.

5) *COMPAS*: The ProPublica COMPAS dataset [78] contains 6,172 records of data on criminal defendants, including attributes such as the defendant's gender, age, and race,

and their recidivism outcome within 2 years of their initial screening. We take our sensitive attribute S to be race and our prediction target Y to be the recidivism outcome.

B. Implementation Details

For all experiments, we use an isotropic Gaussian distribution for the encoder with mean and variance learned by a neural network, $P_{Z|X} = \mathcal{N}(Z|\mu_{\text{enc}}, \sigma_{\text{enc}}^2 I_d)$, using the same notation as described in Sections III-E and III-F. Leveraging the reparameterization trick, we compute our representation $Z|X$ as $Z|X = \mu_{\text{enc}} + \sigma_{\text{enc}} \odot E$, where \odot is the element-wise (Hadamard) product and $E \sim \mathcal{N}(0, I_d)$.

We model the approximation of the representation's marginal as a d -dimensional spherical Gaussian, $Q_Z = \mathcal{N}(Z|0, \gamma^2 I_d)$. We calculate the Rényi divergence in (16) between the multivariate Gaussians $P_{Z|X}$ and Q_Z using (20) and calculate the Rényi cross-entropy using (26). For all experiments, we use a value of $\gamma = 1$. We add an additional sigmoid function to the encoder, limiting outputs such that all values of σ_{enc}^2 are < 1 , ensuring the conditions in (21), (27) and Footnote 4 are satisfied.

For binary target Y , we model $Q_{Y|Z}$ and $Q_{Y|Z,S}$ with Bernoulli distributions, $Q_{Y|Z} = \text{Bernoulli}(Y; f(Z))$ and $Q_{Y|Z,S} = \text{Bernoulli}(Y; g(Z, S))$, where⁶ f and g are auxiliary fully connected networks with a sigmoid function at the output, ensuring that $0 < f(Z) < 1$ and $0 < g(Z, S) < 1$. As specified in our description of the datasets above, we only use binary values for Y in our experiments. However, our method can also accommodate categorical (non-binary) values for Y , i.e., $\mathcal{Y} = \{0, \dots, k-1\}$ for integer $k > 2$. In this case, $Q_{Y|Z}$ and $Q_{Y|Z,S}$ can be modeled with Multinoulli distributions, $Q_{Y|Z} = \text{Multinoulli}(Y; f(Z) = (p_0, \dots, p_{k-1}))$ and $Q_{Y|Z,S} = \text{Multinoulli}(Y; g(Z, S) = (p_0, \dots, p_{k-1}))$, where⁷ f and g are auxiliary fully connected networks with a softmax function at the output, ensuring that $\sum_{i=0}^{k-1} p_i = 1$ and $0 < p_i < 1$ for each $i \in \{0, \dots, k-1\}$.

While most prior work on fairness (e.g., [5], [9], [10], [19], [44] among others) consider only binary sensitive attributes, we also use our method with categorical sensitive attributes. When S is binary, we use S directly with (16), while when S is categorical, i.e., $S = \{0, \dots, n-1\}$ for integer $n > 2$, we encode it as a one-hot encoding following standard practice in the literature [79]. We then replace S with its one-hot encoding in our final cost function; i.e., $S = (S_0, \dots, S_{n-1}) =: S_0^{n-1}$, where each S_j , $j \in \{0, \dots, n-1\}$, is a binary random variable and component of the one-hot encoding vector of length n , leading to an RFIB final cost function of

$$J_{\text{RFIB}} = \frac{1}{N} \sum_{i=1}^N \left[F_{\alpha}(P_{Z|X=x_i}, Q_Z) - \beta_1 \mathbb{E}_E \left[\log(Q_{Y|Z}(y_i | f(x_i, E))) \right] \right]$$

⁶The notation $P_U = \text{Bernoulli}(U; p)$ means that U is a Bernoulli random variable with parameter p (i.e., $P(U = 1) = p$).

⁷The notation $P_U = \text{Multinoulli}(U; (p_0, \dots, p_{k-1}))$ means that U is a Multinoulli random variable with parameters p_0, \dots, p_{k-1} (i.e., $P(U = i) = p_i$ for $i \in \{0, \dots, k-1\}$).

$$-\beta_2 \mathbb{E}_E \left[\log \left(Q_{Y|S_0^{n-1}, Z} (y_i | s_{0,i}, \dots, s_{n-1,i}, f(x_i, E)) \right) \right]. \quad (29)$$

The encoder network $P_{Z|X}$ is a ResNet50 [80] classifier pretrained on ImageNet with the final linear layer replaced by two randomly initialized layers with output dimension d equal to the dimension of the representation. The two decoder networks f and g are each two fully connected layers with 100 units and a sigmoid layer.

After creating the representation Z , we use a logistic regression classifier with default settings to predict Y from Z . We evaluate accuracy and fairness on these predictions. Fig. 2 shows the architecture of our model for the case of image input data.

We preprocess images by taking a 128 by 128 pixel center crop of the CelebA images, a 256 pixel by 256 pixel center crop of the EyePACS images, and use the full 224 pixel by 224 pixel images for FairFace. For the tabular Adult and COMPAS datasets, we normalize the input data X to have zero mean and unit variance. We split the data into a training set and test set based on partitions described in detail in Section IV-D. We then further split the training set, randomly choosing 10 % of its data for validation, and use this validation set for early stopping. We train for up to 20 epochs, with early stopping triggering when there is no decrease in validation loss for 5 epochs, using a `min_delta` value of 0.

For all experiments, we train using PyTorch on a NVIDIA GP100 GPU. We use $d = 64$ as the dimension of our representation Z , a batch size of 64, and the Adam optimizer with a learning rate of 0.001. For the experiments on image data with binary sensitive attribute S , we use the Rényi divergence of order α in (16), by varying α linearly from 0 to 1, with $\alpha = 0$ signifying that

$$D_0(P_{Z|X} \| Q_Z) = -\log Q_Z(z : P_{Z|X}(z) > 0) = 0$$

and with $\alpha = 1$ corresponding to $D_\alpha(P_{Z|X} \| Q_Z)$ being given by the KL divergence $D_{KL}(P_{Z|X} \| Q_Z)$; see (11). We also perform some additional experiments with α values > 1 , typically obtaining similar performance behavior as for α values < 1 .

For the experiments on tabular data with binary S and on the image data with categorical S (FairFace), we use the Rényi cross-entropy in (16) and (29), respectively, with its order denoted by $\tilde{\alpha}$ (to differentiate it from the Rényi divergence order α), which we vary from 0 to 1.

C. Metrics

We use the following metrics to evaluate how well the model performs:

1) *Measure of Utility*: We use the overall classification accuracy (later denoted acc):

$$\text{acc} = P(\hat{Y} = Y). \quad (30)$$

2) *Measures of Fairness*: We measure this in multiple ways:

a) using the gap in accuracy (denoted acc_{gap}) between favored and protected subpopulations;

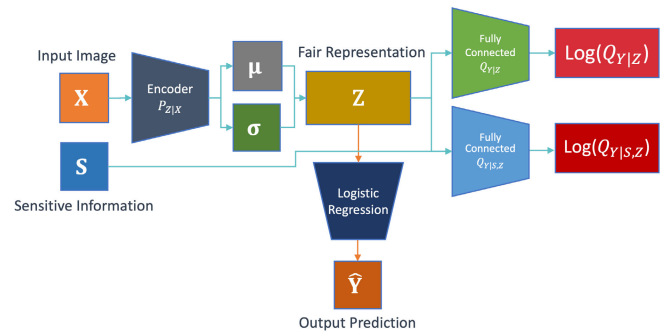


Fig. 2. Architecture of the model is depicted. The input image X is given to the encoder $P_{Z|X}$ that generates the mean μ and standard deviation σ for the distribution of the fair representation Z , $\mathcal{N}(Z|\mu_{\text{enc}}, \sigma_{\text{enc}}^2 I_d)$. Then, Z is given to two fully connected networks, one of which with additional access to the sensitive information S . After training, Z can be used as input to other existing architectures for fair prediction. For our experiments, we use logistic regression to predict Y from Z .

- b) reporting the minimum accuracy across subpopulations (denoted as acc_{min}), which is based on the Rawlsian principle of achieving fairness by maximizing acc_{min} [81];
- c) measuring the adherence to demographic parity via its gap dp_{gap} , following [44];
- d) measuring the adherence using equalized odds via its gap $\text{eqodds}_{\text{gap}}$.

These metrics are defined for only binary sensitive variables in prior work [5], [9], [10], but here we generalize them to also handle non-binary (categorical) sensitive variables where

$$S = \{0, 1, \dots, n-1\}, \quad n \geq 2.$$

We define them as follows:

$$\begin{aligned} \text{acc}_{\text{gap}} &= \frac{1}{n(n-1)} \sum_{i=0}^{n-1} \sum_{\substack{j \neq i \\ j=0}}^{n-1} |P(\hat{Y} = Y | S = i) - P(\hat{Y} = Y | S = j)|, \end{aligned} \quad (31)$$

$$\text{acc}_{\text{min}} = \min_{s \in S} P(\hat{Y} = Y | S = s), \quad (32)$$

$$\begin{aligned} \text{dp}_{\text{gap}} &= \frac{1}{n(n-1)} \sum_{i=0}^{n-1} \sum_{\substack{j \neq i \\ j=0}}^{n-1} |P(\hat{Y} = 1 | S = i) - P(\hat{Y} = 1 | S = j)|, \end{aligned} \quad (33)$$

$$\begin{aligned} \text{eqodds}_{\text{gap}} &= \max_{y \in \{0,1\}} \frac{1}{n(n-1)} \sum_{i=0}^{n-1} \sum_{\substack{j \neq i \\ j=0}}^{n-1} |P(\hat{Y} = 1 | S = i, Y = y) \\ &\quad - P(\hat{Y} = 1 | S = j, Y = y)| \end{aligned} \quad (34)$$

where the last two metrics assume a binary target Y . Note that for the binary sensitive attributes ($n = 2$), these definitions recover the standard definitions from the literature.

3) *Joint Utility-Fairness Measure*: Echoing and comparing with the work in [19], we use a single metric that

jointly captures utility and fairness, the Conjunctive Accuracy Improvement (CAI_λ) measure:

$$\text{CAI}_\lambda = \lambda(\text{acc}_{\text{gap}}^{\text{b}} - \text{acc}_{\text{gap}}^{\text{d}}) + (1 - \lambda)(\text{acc}^{\text{d}} - \text{acc}^{\text{b}}) \quad (35)$$

where $0 \leq \lambda \leq 1$, and acc^{b} and acc^{d} are the accuracy for baseline and debiased algorithms, respectively, while $\text{acc}_{\text{gap}}^{\text{b}}$ and $\text{acc}_{\text{gap}}^{\text{d}}$ are gap in accuracy for the baseline and debiased algorithms. In practice, one can use $\lambda = 0.5$ for an equal balance between utility and fairness or a higher value such as $\lambda = 0.75$ to emphasize fairness.

D. Experiment Details and Results

We perform experiments on the EyePACS, CelebA, and FairFace datasets, considering the challenging case of severe data imbalance where training data is completely missing for one protected subgroup (e.g., diseased dark skin individuals for EyePACS). We also perform additional experiments on the Adult and COMPAS datasets.

1) *Hyper-Parameter Tuning*: For all datasets and experiments, we conduct a hyper-parameter sweep. We use various combinations of α (for Rényi divergence) and $\tilde{\alpha}$ (for Rényi cross-entropy) varied linearly from 0 to 1 (in addition to α values > 1) and β_1 and β_2 varied linearly from 1 to 50 (image datasets) or 1 to 100 (tabular datasets). Results are presented in Tables II-VIII, and are described in detail later in this section.

As mentioned in Section III-D, the RFIB parameter pairs ($\alpha = 1, \beta_1 = 0$) and ($\alpha = 1, \beta_2 = 0$) correspond to the CFB and IB systems, respectively. We therefore compare RFIB with those two methods which are based on the KL divergence and only have a single β hyper-parameter. We do separate comparisons of CFB, IB with RFIB, picking a commonly used value that performs well for CFB, IB and setting RFIB's corresponding hyper-parameter to the same value. We then use grid search to tune the two additional hyper-parameters that our method introduces; to compare with IB we fix a value of β_1 and tune α and β_2 , while to compare with CFB we fix a value of β_2 and then tune α and β_1 .

We implement the IB and CFB methods ourselves and also compare RFIB with two methods used by [19]: adversarial independence, referred to as adversarial debiasing (AD), that minimizes conditional dependence of predictions on sensitive attributes with an adversarial two player game and intelligent augmentation (IA) that generates synthetic data for under-represented populations and performs data augmentation to train a less biased model. To compare with AD and IA, we take results from [19] and report their original CAI scores calculated with respect to their baseline, while for IB and CFB we implement the methods ourselves and calculate CAI scores with respect to results from our own baseline, a ResNet50 network [80].

We recapitulate most acronyms used in Table I.

2) *EyePACS Results*: We predict $Y = \text{DR}$ status while using $S = \text{ITA}$ as the sensitive attribute. We consider a scenario where training data is completely missing for the subgroup of $(Y, S) = (1, 1)$, individuals referable for DR who have dark skin. The goal of our method is for predictions on the

TABLE I
LIST OF ACRONYMS USED

| Acronym | Stands for |
|---------|---|
| AD | Adversarial Debiasing |
| CAI | Conjunctive Accuracy Improvement |
| CFB | Conditional Fairness Bottleneck |
| DI | Disparate Impact |
| DR | Diabetic Retinopathy |
| EO | Equalized Odds |
| IA | Intelligent Augmentation |
| IB | Information Bottleneck |
| ITA | Individual Topology Angle |
| KL | Kullback-Leibler |
| RFIB | Rényi Fair Information Bottleneck |
| UMAP | Uniform Manifold Approximation and Projection |

missing subgroup to be just as accurate as on the group with adequate training data, which is a problem of both fairness and also *domain adaptation*, achieving high performance on a group not present in the original dataset. This scenario matches an important real world problem where data for a minority subgroup such as dark skinned individuals is lacking.

We create a training partition containing both images referable and non-referable for DR of light skin individuals but only non-referable images of dark skin individuals. We use the same partition as in [19] to compare with their method, using a training set that consists of 10,346 light skin images referable for diabetic retinopathy (DR = 1, ITA = 0), 5,173 non-referable light skin images (DR = 0, ITA = 0), and 5,173 non-referable dark skin images (DR = 0, ITA = 1). Then for a fair assessment of our method's performance we evaluate on a balanced test set with an equal number of positive and negative examples for both dark and light skin individuals, with the set containing 2,400 images equally balanced across DR and ITA.

As shown in Table II, our method outperforms all other methods, showing improvements in accuracy and fairness across all metrics. We show a result with a value of $\alpha > 1$ in the bottom row, showing that it is still possible to achieve promising performance with higher α values, although slightly better results for most metrics were achieved with values of $\alpha < 1$. Usual caution should be exercised in interpretations since – despite our aligning with data partitioning in [19] – other variations may exist with [19], [63], and [44] due to non-determinism, parameter setting or other factors. Of interest is that we outperform the baseline in both fairness and accuracy. While somewhat unexpected, as we typically expect a trade-off between fairness and compression, this is potentially due to the compression term $I(Z; X)$ acting as a regularizer and improving generalization performance, as argued by [63].

We perform a second experiment where we use the same networks from before trained using $S = \text{ITA}$ but test on a balanced test set where $S = \text{Ethnicity}$ (defined as in [19]), with labels coming from a human clinician. The test set consists of 400 images equally balanced across DR and ethnicity. We show these experimental results in Table III, where we outperform the other methods across most metrics, including the most important CAI scores. As we achieve similar results using ITA and race labels from a human clinician, we show that ITA can successfully be used as an easily obtained alternative to manual label annotation, further supporting conclusions

TABLE II

RESULTS FOR DEBIASING METHODS ON EYEPACS PREDICTING $Y = \text{DR STATUS}$, TRAINED ON PARTITIONING WITH RESPECT TO $S = \text{ITA}$, AND EVALUATED ON A TEST SET BALANCED ACROSS DR STATUS AND ITA. FOR METRICS WITH AN \uparrow HIGHER IS BETTER WHEREAS FOR \downarrow LOWER IS BETTER. SUBPOPULATION IS THE ONE THAT CORRESPONDS TO THE MINIMUM ACCURACY, WITH (D) INDICATING DARK SKIN AND (L) LIGHT SKIN. METRICS ARE GIVEN AS PERCENTAGES. NA INDICATES THAT RESULTS ARE NOT AVAILABLE DUE TO NOT BEING REPORTED IN THE CITED WORK

| Methods | acc \uparrow | acc _{gap} \downarrow | acc _{min} \uparrow (subpop.) | CAI _{0.5} \uparrow | CAI _{0.75} \uparrow | dp _{gap} \downarrow | eqodds _{gap} \downarrow |
|---|----------------|---------------------------------|--|-------------------------------|--------------------------------|--------------------------------|------------------------------------|
| Baseline (results from [19]) | 70.0 | 3.5 | 68.3 | - | - | NA | NA |
| AD ($\beta = 0.5$) (results from [19]) | 76.12 | 2.41 | 74.92 (L) | 3.61 | 2.35 | NA | NA |
| IA (results from [19]) | 71.5 | 1.5 | 70.16 (D) | 1.75 | 1.875 | NA | NA |
| Baseline (ours) | 73.37 | 8.08 | 69.33 (D) | - | - | 28.25 | 36.33 |
| IB ($\beta_1=30$) [63] | 74.12 | 2.08 | 73.08 (D) | 3.37 | 4.69 | 18.58 | 20.67 |
| CFB ($\beta_2=30$) [44] | 77.83 | 1.66 | 77.0 (L) | 5.44 | 5.93 | 10.83 | 12.5 |
| RFIB (ours) ($\alpha = 0.8, \beta_1 = 36, \beta_2 = 30$) | 79.42 | 0.5 | 79.17 (L) | 6.81 | 7.19 | 16.17 | 16.67 |
| RFIB (ours) ($\alpha = 0.3, \beta_1 = 30, \beta_2 = 50$) | 79.71 | 1.75 | 78.83 (L) | 6.33 | 6.33 | 9.75 | 11.50 |
| RFIB (ours) ($\alpha = 1.8, \beta_1 = 30, \beta_2 = 17$) | 78.35 | 0.25 | 78.25 (L) | 6.41 | 7.12 | 15.58 | 15.83 |

TABLE III

PERFORMANCE RESULTS FOR DEBIASING METHODS ON EYEPACS PREDICTING $Y = \text{DR STATUS}$, TRAINED ON PARTITIONING WITH RESPECT TO $S = \text{ITA}$, AND EVALUATED ON A TEST SET BALANCED ACROSS DR STATUS AND ETHNICITY (DEFINED AS IN [35]). FOR METRICS WITH AN \uparrow HIGHER IS BETTER WHEREAS FOR \downarrow LOWER IS BETTER. SUBPOPULATION IS THE ONE THAT CORRESPONDS TO THE MINIMUM ACCURACY, WITH (W) INDICATING WHITE INDIVIDUALS AND (B) BLACK INDIVIDUALS. METRICS ARE GIVEN AS PERCENTAGES

| Methods | acc \uparrow | acc _{gap} \downarrow | acc _{min} \uparrow (subpop.) | CAI _{0.5} \uparrow | CAI _{0.75} \uparrow | dp _{gap} \downarrow | eqodds _{gap} \downarrow |
|---|----------------|---------------------------------|--|-------------------------------|--------------------------------|--------------------------------|------------------------------------|
| Baseline | 76.00 | 13.00 | 69.50 (B) | - | - | 3.00 | 16.00 |
| IB ($\beta_2=30$) [44] | 78.75 | 3.50 | 77.00 (B) | 6.12 | 7.81 | 0.50 | 4.00 |
| CFB ($\beta_2=30$) [44] | 75.00 | 6.00 | 72.00 (B) | 3.00 | 5.00 | 5.00 | 11.00 |
| RFIB (ours) ($\alpha = 0.8, \beta_1 = 36, \beta_2 = 30$) | 81.5 | 6.00 | 78.50 (B) | 6.25 | 6.63 | 3.00 | 9.00 |
| RFIB (ours) ($\alpha = 0.3, \beta_1 = 30, \beta_2 = 50$) | 81.75 | 4.50 | 79.50 (B) | 7.13 | 7.81 | 11.50 | 16.00 |

TABLE IV

RESULTS FOR DEBIASING METHODS ON CELEBA PREDICTING $Y = \text{AGE}$, TRAINED ON PARTITIONING WITH RESPECT TO $S = \text{ITA}$, AND EVALUATED ON A TEST SET BALANCED ACROSS AGE AND ITA. FOR METRICS WITH AN \uparrow HIGHER IS BETTER WHEREAS FOR \downarrow LOWER IS BETTER. SUBPOPULATION IS THE ONE THAT CORRESPONDS TO THE MINIMUM ACCURACY, WITH (D) INDICATING DARK SKIN AND (L) LIGHT SKIN. METRICS ARE GIVEN AS PERCENTAGES. NA INDICATES THAT RESULTS ARE NOT AVAILABLE DUE TO NOT BEING REPORTED IN THE CITED WORK

| Methods | acc \uparrow | acc _{gap} \downarrow | acc _{min} \uparrow (subpop.) | CAI _{0.5} \uparrow | CAI _{0.75} \uparrow | dp _{gap} \downarrow | eqodds _{gap} \downarrow |
|---|----------------|---------------------------------|--|-------------------------------|--------------------------------|--------------------------------|------------------------------------|
| Baseline (results from [19]) | 74.4 | 13.9 | 67.5 | - | - | NA | NA |
| AD ($\beta = 0.5$) (results from [19]) | 76.45 | 9.6 | 71.65 (D) | 3.17 | 3.75 | NA | NA |
| IA (results from [19]) | 75.29 | 9.18 | 70.7 (D) | 2.8 | 1.56 | NA | NA |
| Baseline (ours) | 70.61 | 16.57 | 62.32 (D) | - | - | 43.82 | 60.4 |
| IB ($\beta_1=30$) [63] | 71.76 | 14.87 | 64.32 (D) | 1.42 | 4.69 | 36.62 | 51.5 |
| CFB ($\beta_2=30$) [44] | 71.77 | 13.25 | 65.15 (D) | 2.24 | 2.78 | 38.3 | 51.55 |
| RFIB (ours) ($\alpha = 0.3, \beta_1 = 30, \beta_2 = 43$) | 75.04 | 10.92 | 69.57 (D) | 5.04 | 5.34 | 29.02 | 39.95 |
| RFIB (ours) ($\alpha = 0.4, \beta_1 = 1, \beta_2 = 30$) | 76.92 | 5.05 | 74.4 (D) | 8.91 | 10.22 | 2.7 | 7.75 |

reached by [19]. These results also demonstrate the ability of our method to perform well in this type of protected factor domain adaptation problem where a different protected factor is used after initial training, which is important in settings where the actual protected factor is not revealed for privacy reasons.

3) *CelebA Results*: We predict age using ITA as the sensitive attribute. Again, we consider the domain adaptation/fairness problem where part of the data is completely missing for a protected subgroup. In this case, older light skin images are missing and we use a training set of 48,000 images,

consisting of 24,000 older dark skin images, 12,000 younger light skin images, and 12,000 older dark skin images. We use a test set of 8,000 images equally balanced across age and ITA, using the same partition as [19] and varying hyper-parameters and comparing methods the same way as for EyePACS. Our results are shown in Table IV, showing our method performs the best across all metrics.

4) *Adult and COMPAS Results*: We predict income on the Adult dataset using gender as the sensitive attribute and predict recidivism outcome on the COMPAS dataset using race as the sensitive attribute. For Adult, we use the train/test partitions

TABLE V

RESULTS FOR DEBIASING METHODS ON ADULT PREDICTING $Y = \text{INCOME}$ WHERE S IS RACE. THE \uparrow INDICATES THAT HIGHER IS BETTER, AND $\tilde{\alpha}$ INDICATES THAT RESULTS WERE OBTAINED USING RÉNYI CROSS-ENTROPY. NA INDICATES THAT RESULTS ARE NOT AVAILABLE DUE TO NOT BEING REPORTED IN THE CITED WORK

| Methods | acc \uparrow | DI \uparrow | EO \uparrow |
|--|----------------|---------------|---------------|
| Fair Classifier (results from [82]) | 83.7 | 0.942 | NA |
| Fair Classifier (results from [82]) | 84.4 | NA | 0.753 |
| FR-Train (results from [83]) | 82.4 | 0.828 | NA |
| FR-Train (results from [83]) | 84.2 | NA | 0.917 |
| RFIB (ours) ($\tilde{\alpha} = .3, \beta_1 = 1, \beta_2 = 84$) | 84.2 | 0.420 | 0.970 |
| RFIB (ours) ($\tilde{\alpha} = .16, \beta_1 = 1, \beta_2 = 47$) | 79.5 | 0.492 | 0.999 |
| RFIB (ours) ($\tilde{\alpha} = .42, \beta_1 = 1, \beta_2 = 0$) | 76.2 | 0.943 | 0.424 |

provided with the dataset, while for COMPAS we randomly split the data using a 70%/30% train/test split. We compare with two methods, Fair Classifier [82] and FR-Train [83], using two fairness metrics defined in those works, Disparate Impact (DP) and Equalized Odds (EO):

$$\text{DI} = \min_{s \in \mathcal{S}} \min_{\hat{s} \neq s} \frac{P(\hat{Y} = 1 | S = \hat{s})}{P(\tilde{Y} = 1 | S = s)}, \quad (36)$$

$$\text{EO} = \min_{y \in \mathcal{Y}} \min_{s \in \mathcal{S}} \min_{\hat{s} \neq s} \frac{P(\hat{Y} = 1 | S = \hat{s}, Y = y)}{P(\tilde{Y} = 1 | S = s, Y = y)}. \quad (37)$$

We show results in Tables V and VI with various combinations of hyper-parameters including where we achieve the highest accuracy, highest DI, and highest EO. As expected, there is typically a trade-off between achieving high accuracy and high performance on the fairness metrics (this is also confirmed in the recent work in [84] based on the AD method). In most cases, higher β values coupled with higher $\tilde{\alpha}$ values result in increased accuracy but worsened fairness (note that the RFIB compression term $I(Z; X)$ is characterized here by the Rényi cross-entropy of order $\tilde{\alpha}$, which is a non-increasing function of $\tilde{\alpha}$).

Compared to Fair Classifier and FR-Train, we suffer a drop in accuracy to achieve a higher DI, but outperform these baselines with respect to EO (the most important fairness metric according to [9]). On the Adult dataset, we achieve higher EO scores than both Fair Classifier and FR-Train while maintaining the same or very comparable accuracy. In cases where we do need to sacrifice accuracy to achieve fairness, we are still able to beat both baselines on at least one metric according to how we tune $\tilde{\alpha}$ and the two β hyper-parameters. In those cases, which method is preferable will depend on the application and the metric that is most suitable for that application.

5) *FairFace Results*: We predict $Y = \text{Gender}$ with $S = \text{Race}$ as the sensitive attribute, testing on a test set balanced across gender and race. As before, we exclude one population subgroup and remove Black females from the training data, matching the common real-world scenario where data is lacking for this group. Unlike the previous experiments, here we can obtain race directly from the dataset rather than using ITA as a proxy for skin tone. We binarize the race labels into

TABLE VI

RESULTS FOR DEBIASING METHODS ON COMPAS PREDICTING $Y = \text{RECIDIVISM OUTCOME}$ WHERE S IS RACE. THE \uparrow INDICATES THAT HIGHER IS BETTER, AND $\tilde{\alpha}$ INDICATES THAT RESULTS WERE OBTAINED USING RÉNYI CROSS-ENTROPY. NA INDICATES THAT RESULTS ARE NOT AVAILABLE DUE TO NOT BEING REPORTED IN THE CITED WORK

| Methods | acc \uparrow | DI \uparrow | EO \uparrow |
|---|----------------|---------------|---------------|
| Fair Classifier (results from [82]) | 67.5 | 0.959 | NA |
| Fair Classifier (results from [82]) | 63.8 | NA | 0.908 |
| FR-Train (results from [83]) | 67.6 | 0.838 | NA |
| FR-Train (results from [83]) | 62.8 | NA | 0.959 |
| RFIB (ours) ($\tilde{\alpha} = .6, \beta_1 = 1, \beta_2 = 53$) | 68.2 | 0.571 | 0.717 |
| RFIB (ours) ($\tilde{\alpha} = .5, \beta_1 = 1, \beta_2 = 6$) | 57.6 | 0.783 | 0.979 |
| RFIB (ours) ($\tilde{\alpha} = .3, \beta_1 = 1, \beta_2 = 2$) | 55.0 | 0.999 | 0.875 |

two groups, a smaller Black group and a larger non-Black group. We use a training set of 16,500 images, consisting of 5,500 male Black images, 5,500 male white images, and 5,500 female white images. We test on a test set of 3,000 images equally balanced across gender and race. Our results, given in Table VII, show that we outperform both the IB and CFB methods both on accuracy and on all fairness metrics.

6) *FairFace Results (Categorical Variables)*: We extend our method to work with non-binary categorical attributes (as described in Section IV-B). We predict $Y = \text{Gender}$ with $S = \text{Race}$ containing three race groups: Black, Hispanic, and Other, with the Other group containing a mix of White, East Asian, Southeast Asian, Middle Eastern, and Indian races. We use a training set of 17,500 images and test set of 7,500 images. Since most prior work on fairness is defined only for binary attributes, we compare RFIB with two non-fairness methods, a ResNet50 baseline and the IB method. We present results in Table VIII, showing that we are able to achieve high accuracy while also achieving improvement across all but one of the fairness metrics, most significantly for the dp_{gap} and $\text{eqodds}_{\text{gap}}$ metrics.

E. UMAP Clustering Analysis and Influence of α

We use UMAP [85] to provide a visual illustration of the effect of the Rényi divergence on the RFIB system performance as it calibrates the compression term by varying its order α . In Fig. 3, we show 2-dimensional UMAP vectors of our representation Z , coloring the points both based on the label Y and sensitive attribute S . Here experiments were done on the EyePACS dataset with $Y = \text{DR}$ and $S = \text{ITA}$. The goal is for the representation to preserve information about Y , allowing the two classes of Y to be easily separated, while removing the sensitive information S and preventing its two classes from being distinguished.

A the value of $\alpha = 0$ corresponds to no compression with $I(Z; X) = 0$, which preserves maximum accuracy but does not help fairness, with the classes of both Y and S being easily separated as shown in Fig. 3. Increasing α gradually results in the different points getting more mixed together (recall that the Rényi divergence is non-decreasing in its order α). While this is seen for both Y and S , the $I(Z; Y)$ and $I(Z; Y|S)$

TABLE VII

RESULTS FOR DEBIASING METHODS ON FAIRFACE PREDICTING $Y = \text{GENDER}$, TRAINED ON PARTITIONING WITH RESPECT TO $S = \text{RACE}$, AND EVALUATED ON A TEST SET BALANCED ACROSS GENDER AND ITA. FOR METRICS WITH AN \uparrow HIGHER IS BETTER WHEREAS FOR \downarrow LOWER IS BETTER. SUBPOPULATION IS THE ONE THAT CORRESPONDS TO THE MINIMUM ACCURACY, WITH (B) INDICATING BLACK. METRICS ARE GIVEN AS PERCENTAGES

| Methods | acc \uparrow | acc _{gap} \downarrow | acc _{min} [↑] (subpop.) | CAI _{0.5} \uparrow | CAI _{0.75} \uparrow | dp _{gap} \downarrow | eqodds _{gap} \downarrow |
|---|----------------|---------------------------------|---|-------------------------------|--------------------------------|--------------------------------|------------------------------------|
| Baseline | 73.97 | 16.47 | 65.73 (B) | - | - | 27.53 | 44.0 |
| IB ($\beta_2=30$) [63] | 73.93 | 14.93 | 68.6 (B) | 0.75 | 1.14 | 30.0 | 44.93 |
| CFB ($\beta_2=30$) [44] | 75.6 | 14.0 | 65.2 (B) | 2.05 | 2.26 | 26.27 | 40.27 |
| RFIB (ours) ($\alpha = 0.2, \beta_1 = 30, \beta_2 = 29$) | 83.6 | 8.53 | 79.33 (B) | 8.78 | 8.36 | 19.47 | 28.0 |
| RFIB (ours) ($\alpha = 0.2, \beta_1 = 1, \beta_2 = 30$) | 82.07 | 7.33 | 78.4 (B) | 8.62 | 8.88 | 19.07 | 26.4 |

TABLE VIII

RESULTS FOR DEBIASING METHODS ON FAIRFACE PREDICTING $Y = \text{GENDER}$ WHERE S IS RACE, AND S IS A CATEGORICAL ATTRIBUTE WITH THREE RACE CATEGORIES. FOR METRICS WITH AN \uparrow HIGHER IS BETTER WHEREAS FOR \downarrow LOWER IS BETTER, AND $\tilde{\alpha}$ INDICATES THAT RESULTS WERE OBTAINED USING RÉNYI CROSS-ENTROPY. SUBPOPULATION IS THE ONE THAT CORRESPONDS TO THE MINIMUM ACCURACY, WITH (B) INDICATING BLACK. METRICS ARE GIVEN AS PERCENTAGES

| Methods | acc \uparrow | acc _{gap} \downarrow | acc _{min} [↑] (subpop.) | CAI _{0.5} \uparrow | CAI _{0.75} \uparrow | dp _{gap} \downarrow | eqodds _{gap} \downarrow |
|---|----------------|---------------------------------|---|-------------------------------|--------------------------------|--------------------------------|------------------------------------|
| Baseline | 83.9 | 6.77 | 76.72 (B) | - | - | 8.17 | 13.18 |
| IB ($\beta_2=30$) [63] | 79.12 | 5.87 | 73.38 (B) | -1.94 | -0.52 | 7.07 | 11.07 |
| RFIB (ours) ($\tilde{\alpha} = 0.2, \beta_1 = 30, \beta_2 = 29$) | 85.24 | 6.74 | 78.01 (B) | 0.68 | 0.35 | 3.88 | 7.84 |

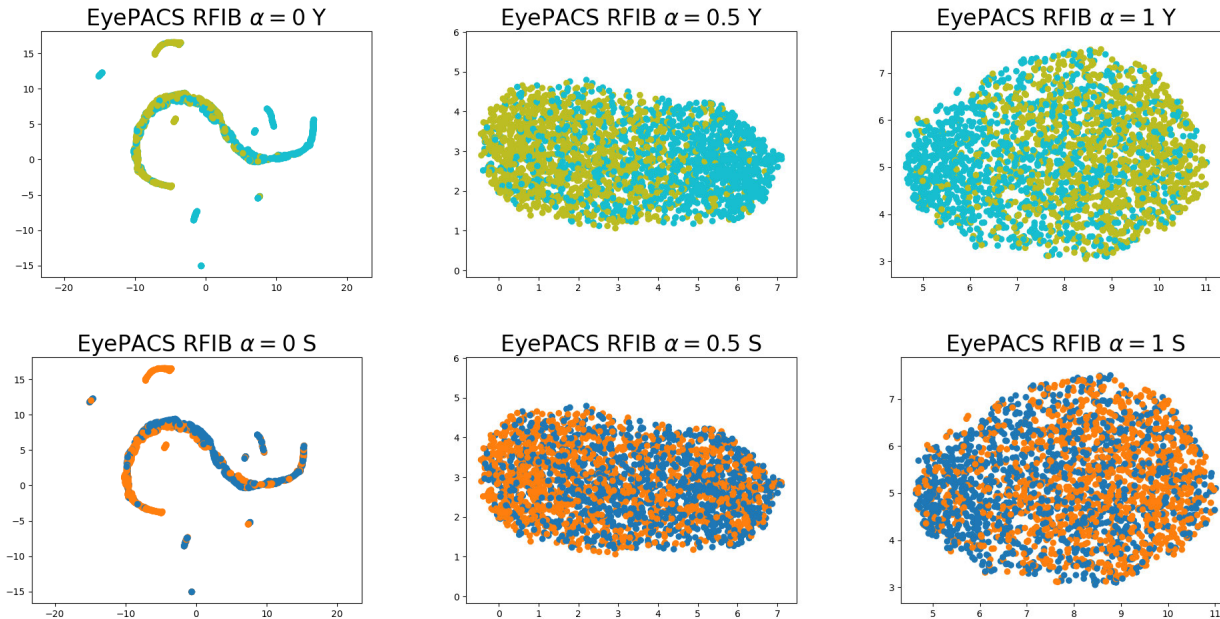


Fig. 3. UMAP dimensionality reduction of Z for the EyePACS dataset using Rényi divergence of order α in (16). The top row shows the two classes of Y (light blue representing the positive examples and light green the negative) while the bottom row shows the two classes of S (orange representing the positive examples and blue the negative). Going from left to right α is increased, resulting in more compression and causing a decrease in separation of the two classes. While separation occurs for both Y and S , it occurs to a greater extent for Y as desired.

terms help ensure that information about Y is still preserved, allowing the classes of Y to still separate fairly well even with higher compression, whereas the classes of S get mixed together as desired.

We note that an intermediate value of α can potentially provide the best compromise between fairness and accuracy, as a more moderate amount of compression can be enough to sufficiently remove sensitive information and further

compression might only harm accuracy. This is illustrated in Fig. 3 where a value of $\alpha = 0.5$ was sufficient to mix together the classes of S while still preserving an obvious separation of the two classes of Y . Further increasing α worsened the separation of Y more than it added additional benefit for S . This is also supported by our experimental results where best overall accuracy-fairness trade-offs were typically obtained for intermediate values of α .

V. CONCLUSION

We propose RFIB, a novel variational IB method based on Rényi measures that offers trade-offs in utility, two fairness objectives, and compactness. Using Rényi divergence and Rényi cross-entropy instead of KL divergence gives a way to control the amount of compression with an additional hyper-parameter. Compared to prior work which incorporates a single definition of fairness, RFIB has the potential benefit of allowing ethicists and policy makers to specify softer and more balanced requirements for fairness that lie between multiple hard requirements, and our work opens the way to future studies expanding on this idea. Experimental results on three different image datasets and two tabular datasets showed RFIB provides benefits vis-a-vis other methods of record including IB, CFB, and other techniques performing augmentation or adversarial debiasing.

Possible directions of future research include using optimization methods other than the weighted sum method we used, such as the ϵ -constraints method [86]. They also include studying different techniques to minimize and maximize the mutual information terms, such as [87], the Mutual Information Neural Estimation algorithm [88] based on the Donsker-Varadhan representation of the Kullback-Leibler divergence [89], or the equivalent algorithm based on a variational representation of the Rényi divergence [90]. In addition, we can further experiment with non-binary (categorical) targets, as outlined in Section IV-B, and study the performance of our method in cases where the sensitive attributes are not explicitly known.

Finally, we note that our method is closely related to the problem of privacy where a privacy mechanism must be designed to release data that has high utility but does not reveal private information (e.g., see [91], [92], [93], [94], [95], [96], [97] and the references therein for various studies on information-theoretic privacy). While similar attributes can be used as sensitive or private variables, for fairness we are concerned with preventing the attribute from affecting classification accuracy, while for privacy we are concerned with preventing an attacker from accessing a user's private data. A very promising future extension of our work is to adapt our method to also remove private information in addition to sensitive information, achieving representations that are both private and fair. Similar to [98], we can then quantify the amount of private information leakage and study theoretical guarantees on privacy for our method.

REFERENCES

- [1] A. Gronowski, W. Paul, F. Alajaji, B. Ghahsifard, and P. Burlina, "Rényi fair information bottleneck for image classification," in *Proc. 17th Can. Workshop Inf. Theory (CWIT)*, Jun. 2022, pp. 1–5. [Online]. Available: <http://cwit.ca/2022/papers/1570796577.pdf>
- [2] N. M. Kinyanjui et al., "Fairness of classifiers across skin tones in dermatology," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 2020, pp. 320–329.
- [3] T. Bolukbasi, K.-W. Chang, J. Y. Zou, V. Saligrama, and A. T. Kalai, "Man is to computer programmer as woman is to homemaker? Debiasing word embeddings," in *Proc. Adv. Neural Inform. Process. Syst.*, 2016, pp. 4349–4357.
- [4] D. Friolo, F. Massacci, C. N. Ngo, and D. Venturi, "Cryptographic and financial fairness," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 3391–3406, 2022.
- [5] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning fair representations," in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 325–333.
- [6] F. Prost, H. Qian, Q. Chen, E. H. Chi, J. Chen, and A. Beutel, "Toward a better trade-off between performance and fairness with kernel-based distribution matching," 2019, *arXiv:1910.11779*.
- [7] S. Caton and C. Haas, "Fairness in machine learning: A survey," 2020, *arXiv:2010.04053*.
- [8] V. Albiero, K. Zhang, M. C. King, and K. W. Bowyer, "Gendered differences in face recognition accuracy explained by hairstyles, makeup, and facial morphology," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 127–137, 2022.
- [9] M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 3315–3323.
- [10] H. Zhao, A. Coston, T. Adel, and G. J. Gordon, "Conditional learning of fair representations," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–17.
- [11] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annu. Allerton Conf. Commun., Control, Comput.*, 1999, pp. 368–377.
- [12] A. Rényi, "On measures of entropy and information," in *Proc. 4th Berkeley Symp. Math. Statist. Probab.*, vol. 1, 1961, pp. 547–561.
- [13] F. Valverde-Albacete and C. Peláez-Moreno, "The case for shifting the Rényi entropy," *Entropy*, vol. 21, no. 1, p. 46, Jan. 2019.
- [14] H. Bhatia, W. Paul, F. Alajaji, B. Ghahsifard, and P. Burlina, "Least k -th-order and Rényi generative adversarial networks," *Neural Comput.*, vol. 33, no. 9, pp. 2473–2510, 2021.
- [15] F. C. Thierrin, F. Alajaji, and T. Linder, "Rényi cross-entropy measures for common distributions and processes with memory," *Entropy*, vol. 24, no. 10, pp. 1–9, 2022.
- [16] G. Pleiss, M. Raghavan, F. Wu, J. Kleinberg, and K. Q. Weinberger, "On fairness and calibration," in *Proc. Adv. Neural Inform. Process. Syst.*, 2017, pp. 5680–5689.
- [17] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 4396–4405.
- [18] A. Grover et al., "Bias correction of learned generative models using likelihood-free importance weighting," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 11058–11070.
- [19] W. Paul, A. Hadzic, N. Joshi, F. Alajaji, and P. Burlina, "TARA: Training and representation alteration for AI fairness and domain generalization," *Neural Comput.*, vol. 34, no. 3, pp. 716–753, Feb. 2022.
- [20] S. Hwang, S. Park, D. Kim, M. Do, and H. Byun, "FairfaceGAN: Fairness-aware facial image-to-image translation," in *Proc. Brit. Mach. Vis. Conf.*, 2020, pp. 1–14.
- [21] P. Sattigeri, S. C. Hoffman, V. Chenthamarakshan, and K. R. Varshney, "Fairness GAN: Generating datasets with fairness properties using a generative adversarial network," *IBM J. Res. Develop.*, vol. 63, pp. 1–9, Jul. 2019.
- [22] N. Quadrianto, V. Sharmanska, and O. Thomas, "Discovering fair representations in the data domain," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 8219–8228.
- [23] C. Wadsworth, F. Vera, and C. Piech, "Achieving fairness through adversarial learning: An application to recidivism prediction," in *Proc. Conf. Fairness, Accountability, Transparency Mach. Learn. (FATML)*, 2018, pp. 1–5.
- [24] B. H. Zhang, B. Lemoine, and M. Mitchell, "Mitigating unwanted biases with adversarial learning," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.*, Dec. 2018, pp. 335–340.
- [25] H. Edwards and A. Storkey, "Censoring representations with an adversary," in *Proc. Int. Conf. Learn. Represent.*, 2016, pp. 1–14.
- [26] A. Beutel, J. Chen, Z. Zhao, and E. H. Chi, "Data decisions and theoretical implications when adversarially learning fair representations," 2017, *arXiv:1707.00075*.
- [27] D. Madras, E. Creager, T. Pitassi, and R. Zemel, "Learning adversarially fair and transferable representations," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 3384–3393.
- [28] P. C. Roy and V. N. Boddeti, "Mitigating information leakage in image representations: A maximum entropy approach," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2581–2589.
- [29] S. Pfohl, B. Marafino, A. Coulet, F. Rodriguez, L. Palaniappan, and N. H. Shah, "Creating fair models of atherosclerotic cardiovascular disease risk," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.*, Jan. 2019, pp. 271–278.

- [30] X. Gitiaux and H. Rangwala, "Learning smooth and fair representations," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2021, pp. 253–261.
- [31] V. Grari, S. Lamprier, and M. Detryniecki, "Fairness-aware neural Rényi minimization for continuous features," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Jul. 2020, pp. 1–8.
- [32] J. Song, P. Kalluri, A. Grover, S. Zhao, and S. Ermon, "Learning controllable fair representations," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, 2019, pp. 2164–2173.
- [33] P. Kairouz, J. Liao, C. Huang, M. Vyas, M. Welfert, and L. Sankar, "Generating fair universal representations using adversarial models," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 1970–1985, 2022.
- [34] A. Kolchinsky, B. D. Tracey, and D. H. Wolpert, "Nonlinear information bottleneck," *Entropy*, vol. 21, no. 12, p. 1181, Nov. 2019.
- [35] I. Fischer, "The conditional entropy bottleneck," *Entropy*, vol. 22, no. 9, p. 999, Sep. 2020.
- [36] A. Makhdoumi, S. Salamatian, N. Fawaz, and M. Médard, "From the information bottleneck to the privacy funnel," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Nov. 2014, pp. 501–505.
- [37] D. Strouse and D. J. Schwab, "The deterministic information bottleneck," *Neural Comput.*, vol. 29, no. 6, pp. 1611–1630, Jun. 2017.
- [38] H. Hsu, S. Asodeh, S. Salamatian, and F. P. Calmon, "Generalizing bottleneck problems," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2018, pp. 531–535.
- [39] S. Asodeh and F. P. Calmon, "Bottleneck problems: An information and estimation-theoretic view," *Entropy*, vol. 22, no. 11, p. 1325, Nov. 2020.
- [40] J.-J. Weng, F. Alajaji, and T. Linder, "An information bottleneck problem with Rényi's entropy," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2021, pp. 2489–2494.
- [41] Z. Goldfeld and Y. Polyanskiy, "The information bottleneck problem and its applications in machine learning," *IEEE J. Sel. Areas Inf. Theory*, vol. 1, no. 1, pp. 19–38, May 2020.
- [42] A. Zaidi, I. Estella-Aguerrí, and S. Shamai (Shitz), "On the information bottleneck problems: Models, connections, applications and information theoretic views," *Entropy*, vol. 22, no. 2, p. 151, Jan. 2020.
- [43] A. Ghassami, S. Khodadadian, and N. Kiyavash, "Fairness in supervised learning: An information theoretic approach," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2018, pp. 176–180.
- [44] B. Rodríguez-Gálvez, R. Thobaben, and M. Skoglund, "A variational approach to privacy and fairness," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Oct. 2021, pp. 1–6.
- [45] A. R. Esposito, M. Gastpar, and I. Issa, "Robust generalization via f-Mutual information," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2020, pp. 2723–2728.
- [46] Y. Li and R. E. Turner, "Rényi divergence variational inference," in *Proc. Adv. Neural Inform. Process. Syst.*, vol. 29, 2016, pp. 1073–1081.
- [47] I. Mironov, "Rényi differential privacy," in *Proc. IEEE 30th Comput. Secur. Found. Symp. (CSF)*, Jun. 2017, pp. 263–275.
- [48] A. Sarraf and Y. Nie, "RGAN: Rényi generative adversarial network," *Social Netw. Comput. Sci.*, vol. 2, no. 1, p. 17, Feb. 2021.
- [49] Y. Pantazis, D. Paul, M. Fasoulakis, Y. Stylianou, and M. A. Katsoulakis, "Cumulant GAN," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 11, pp. 9439–9450, Nov. 2023.
- [50] G. R. Kurri, T. Sypherd, and L. Sankar, "Realizing GANs via a tunable loss function," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Oct. 2021, pp. 1–6.
- [51] G. R. Kurri, M. Welfert, T. Sypherd, and L. Sankar, "A-GAN: Convergence and estimation guarantees," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2022, pp. 276–281.
- [52] M. Welfert, K. Otstot, G. R. Kurri, and L. Sankar, " α_D, α_G -GANs: Addressing GAN training instabilities via dual objectives," 2023, *arXiv:2302.14320*.
- [53] S. Baharlouei, M. Nouiهد, A. Beirami, and M. Razaviyayn, "Rényi fair inference," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–19.
- [54] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan, "A theory of learning from different domains," *Mach. Learn.*, vol. 79, nos. 1–2, pp. 151–175, May 2010.
- [55] Y. Ganin et al., "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2030–2096, May 2015.
- [56] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 97–105.
- [57] N. Zhang, M. Mohri, and J. Hoffman, "Multiple-source adaptation theory and algorithms," *Ann. Math. Artif. Intell.*, vol. 89, nos. 3–4, pp. 237–270, Mar. 2021.
- [58] S. Sagawa, P. W. Koh, T. B. Hashimoto, and P. Liang, "Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 1–19.
- [59] N. Sohoni, J. Dunnmon, G. Angus, A. Gu, and C. Re, "No subclass left behind: Fine-grained robustness in coarse-grained classification problems," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 19339–19352.
- [60] T. Zhang, T. Zhu, K. Gao, W. Zhou, and P. S. Yu, "Balancing learning model privacy, fairness, and accuracy with early stopping criteria," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 9, pp. 5557–5569, Sep. 2023.
- [61] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley Series in Telecommunications and Signal Processing). New York, NY, USA: Wiley-Interscience, 2006.
- [62] F. Alajaji and P.-N. Chen, *An Introduction to Single-User Information Theory*. Singapore: Springer, 2018.
- [63] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, "Deep variational information bottleneck," in *Proc. Int. Conf. Learn. Represent.*, 2017, pp. 1–17.
- [64] F. D. P. Calmon and N. Fawaz, "Privacy against statistical inference," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2012, pp. 1401–1408.
- [65] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. Int. Conf. Learn. Represent.*, 2014, pp. 1–14.
- [66] T. van Erven and P. Harremoës, "Rényi divergence and Kullback–Leibler divergence," *IEEE Trans. Inf. Theory*, vol. 60, no. 7, pp. 3797–3820, Jul. 2014.
- [67] I. Goodfellow et al., "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, 2014, pp. 1–9.
- [68] M. Gil, F. Alajaji, and T. Linder, "Rényi divergence measures for commonly used univariate continuous distributions," *Inf. Sci.*, vol. 249, pp. 124–131, Nov. 2013.
- [69] J. Burbea, "The convexity with respect to Gaussian distributions of divergences of order α ," *Utilitas Mathematica*, vol. 26, pp. 171–192, Nov. 1984.
- [70] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 3730–3738.
- [71] M. Wilkes, C. Y. Wright, J. L. D. Plessis, and A. Reeder, "Fitzpatrick skin type, individual typology angle, and melanin index in an African population: Steps toward universally applicable skin photosensitivity assessments," *JAMA Dermatol.*, vol. 151, no. 8, pp. 902–903, 2015.
- [72] M. Merler, N. Ratha, R. S. Feris, and J. R. Smith, "Diversity in faces," 2019, *arXiv:1901.10436*.
- [73] N. M. Kinyanjui et al., "Estimating skin tone and effects on classification performance in dermatology datasets," in *Proc. NIPS*, 2019, pp. 1–11.
- [74] K. Kärkkäinen and J. Joo, "FairFace: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2021, pp. 1547–1557.
- [75] EyePACS. (2015). *Diabetic Retinopathy Detection*. Data Retrieved From Kaggle. [Online]. Available: <https://www.kaggle.com/agaldran/eyepacs>
- [76] P. Burlina, N. Joshi, W. Paul, K. D. Pacheco, and N. M. Bressler, "Addressing artificial intelligence bias in retinal diagnostics," *Translational Vis. Sci. Technol.*, vol. 10, no. 2, p. 13, Feb. 2021.
- [77] D. Dua and C. Graff. (2017). *UCI Machine Learning Repository*. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [78] W. Dieterich, C. Mendoza, and T. Brennan, "Compas risk scales: Demonstrating accuracy equity and predictive parity," Northpointe, Traverse City, MI, USA, Tech. Rep., Jul. 2016, vol. 7, no. 4.
- [79] C. Mougan, J. M. Alvarez, S. Ruggieri, and S. Staab, "Fairness implications of encoding protected categorical attributes," 2022, *arXiv:2201.11358*.
- [80] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [81] J. Rawls, *Justice as Fairness: A Restatement*. Cambridge, MA, USA: Harvard Univ. Press, 2001.
- [82] J. Cho, G. Hwang, and C. Suh, "A fair classifier using mutual information," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2020, pp. 2521–2526.
- [83] Y. Roh, K. Lee, S. Whang, and C. Suh, "FR-Train: A mutual information-based approach to fair and robust training," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 8147–8157.

- [84] W. Paul, P. Mathew, F. Alajaji, and P. Burlina, "Evaluating trade-offs in computer vision between attribute privacy, fairness and utility," 2023, *arXiv:2302.07917*.
- [85] L. McInnes, J. Healy, N. Saul, and L. Großberger, "UMAP: Uniform manifold approximation and projection," *J. Open Source Softw.*, vol. 3, no. 29, p. 861, Sep. 2018.
- [86] G. Mavrotas, "Effective implementation of the ϵ -constraint method in multi-objective mathematical programming problems," *Appl. Math. Comput.*, vol. 213, no. 2, pp. 455–465, Jul. 2009.
- [87] X. Wang, A. Al-Bashabsheh, C. Zhao, and C. Chan, "Adaptive label smoothing for classifier-based mutual information neural estimation," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2021, pp. 1035–1040.
- [88] M. Ishmael Belghazi et al., "MINE: Mutual information neural estimation," 2018, *arXiv:1801.04062*.
- [89] M. D. Donsker and S. R. S. Varadhan, "Asymptotic evaluation of certain Markov process expectations for large time. IV," *Commun. Pure Appl. Math.*, vol. 36, no. 2, pp. 183–212, Mar. 1983.
- [90] J. Birrell, P. Dupuis, M. A. Katsoulakis, L. Rey-Bellet, and J. Wang, "Variational representations and neural network estimation of Rényi divergences," *SIAM J. Math. Data Sci.*, vol. 3, no. 4, pp. 1093–1116, Jan. 2021.
- [91] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "Utility-privacy tradeoffs in databases: An information-theoretic approach," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 6, pp. 838–852, Jun. 2013.
- [92] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Information extraction under privacy constraints," *Information*, vol. 7, no. 1, p. 15, Mar. 2016.
- [93] S. Asoodeh, M. Diaz, F. Alajaji, and T. Linder, "Estimation efficiency under privacy constraints," *IEEE Trans. Inf. Theory*, vol. 65, no. 3, pp. 1512–1534, Mar. 2019.
- [94] B. Rassouli and D. Gündüz, "Optimal utility-privacy trade-off with total variation distance as a privacy measure," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 594–603, 2020.
- [95] I. Issa, A. B. Wagner, and S. Kamath, "An operational approach to information leakage," *IEEE Trans. Inf. Theory*, vol. 66, no. 3, pp. 1625–1657, Mar. 2020.
- [96] M. Bloch et al., "An overview of information-theoretic security and privacy: Metrics, limits and applications," *IEEE J. Sel. Areas Inf. Theory*, vol. 2, no. 1, pp. 5–22, Mar. 2021.
- [97] A. Zamani, T. J. Oechtering, and M. Skoglund, "Data disclosure with non-zero leakage and non-invertible leakage matrix," *IEEE Trans. Inf. Forensics Security*, vol. 17, pp. 165–179, 2022.
- [98] S. Saeidian, G. Cervia, T. J. Oechtering, and M. Skoglund, "Quantifying membership privacy via information leakage," *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 3096–3108, 2021.



Adam Gronowski received the B.A.Sc. and M.A.Sc. degrees in mathematics and engineering from Queen's University, Canada, in 2020 and 2022, respectively. His research interests include computer vision and machine learning. He received the Canadian Society for Information Theory (CSIT) Best Paper Award in 2022.



William Paul received the B.A. degree in computer science, statistics, and applied mathematics from the University of California at Berkeley, Berkeley, CA, USA, in 2019. His research interests include computer vision, AI robustness, and generative modeling.



Fady Alajaji (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1990 and 1994, respectively. He held a post-doctoral appointment with the Institute for Systems Research, University of Maryland, in 1994. In 1995, he joined the Department of Mathematics and Statistics, Queen's University, Kingston, ON, Canada, where he is currently a Professor of mathematics and engineering. Since 1997, he has also been cross-appointed with the Department of Electrical and Computer Engineering, Queen's University. From 2013 to 2014, he was the Acting Head of the Department of Mathematics and Statistics. From 2003 to 2008 and from 2018 to 2019, he was the Chair of the Queen's Mathematics and Engineering Program. His research interests include information theory, communications, error control coding, data compression, joint source-channel coding, network epidemics, generative adversarial networks, data privacy, and fairness in machine learning. He received the Premiers Research Excellence Award from the Province of Ontario. He was an Associate Editor of Shannon Theory and Information Measures (from 2019 to 2022) of the IEEE TRANSACTIONS ON INFORMATION THEORY. He was an Area Editor of Source-Channel Coding and Signal Processing (from 2008 to 2015) and an Editor of Source and Source-Channel Coding (from 2003 to 2012) of the IEEE TRANSACTIONS ON COMMUNICATIONS. He also served as an organizer and a technical program committee member of several international conferences and workshops.



Bahman Gharesifard (Senior Member, IEEE) received the B.Sc. degree in mechanical engineering and the M.Sc. degree in control and dynamics from Shiraz University, Iran, in 2002 and 2005, respectively, and the Ph.D. degree in mathematics from Queen's University, Canada, in 2009. He held post-doctoral positions with the Department of Mechanical and Aerospace Engineering, University of California, San Diego (from 2009 to 2012), and with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign (from 2012 to 2013). He was an Alexander von Humboldt Research Fellow with the Institute for Systems Theory and Automatic Control, University of Stuttgart, from 2019 to 2020. He has been a Professor with the Department of Mathematics and Statistics, Queen's University, since 2013. He has also been a Professor with the Electrical and Computer Engineering Department, University of California, Los Angeles, since 2021, where he is currently the Area Director of Signals and Systems. His research interests include systems and controls, distributed control, distributed optimization, machine learning, social and economic networks, game theory, geometric control theory, geometric mechanics, and applied Riemannian geometry. He received the 2019 CAIMS-PIMS Early Career Award, jointly awarded by the Canadian Applied and Industrial Math Society and the Pacific Institute for the Mathematical Sciences, the Alexander von Humboldt Foundation Research Fellowship for experienced researchers in 2019, the NSERC Discovery Accelerator in 2019, the SIAG/CST Best SICON Paper Prize in 2021, and the Canadian Society for Information Theory (CSIT) Best Paper Award in 2022. He has served on the Conference Editorial Board for the IEEE Control Systems Society from 2016 to 2020. He is currently an Associate Editor of the IEEE CONTROL SYSTEM LETTERS and IEEE TRANSACTIONS ON NETWORK CONTROL SYSTEMS.



Philippe Burlina (Senior Member, IEEE) received the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland at College Park. He was the Director of software engineering with FileNet (currently IBM). He is currently a Senior Staff Engineer and the Team Lead with Zox (Amazon) and an Associate Research Professor with the Department of Computer Science, Johns Hopkins University, Baltimore, MD. His research interests include computer vision and machine learning.