

Rényi Fair Information Bottleneck for Image Classification

Adam Gronowski
Dept. Mathematics and Statistics
Queen's University
Kingston, ON, Canada
adam.gronowski@queensu.ca

William Paul
Applied Physics Laboratory
Johns Hopkins University
Laurel, MD, United States
william.paul@jhuapl.edu

Fady Alajaji
Dept. Mathematics and Statistics
Queen's University
Kingston, ON, Canada
fa@queensu.ca

Bahman Ghahesifard
Electrical & Computer Engineering Dept.
University of California, Los Angeles
Los Angeles, CA, United States
ghahesifard@ucla.edu

Philippe Burlina
Applied Physics Laboratory
Johns Hopkins University
Laurel, MD, United States
philippe.burlina@jhuapl.edu

Abstract—We develop a novel method for ensuring fairness in machine learning which we term as the Rényi Fair Information Bottleneck (RFIB). We consider two different fairness constraints - demographic parity and equalized odds - for learning fair representations and derive a loss function via a variational approach that uses Rényi's divergence with its tunable parameter α and that takes into account the triple constraints of utility, fairness, and compactness of representation. We then evaluate the performance of our method for image classification using the EyePACS medical imaging dataset, showing it outperforms competing state of the art techniques with performance measured using a variety of compound utility/fairness metrics, including accuracy gap and Rawls' minimal accuracy.

I. INTRODUCTION

The problem of fairness in machine learning is to obtain accurate predictions of a target of interest while remaining free of bias due to sensitive information such as gender, race, age, or other similar attributes. Representing input data as random variable $X \in \mathcal{X}$, prediction target as random variable $Y \in \mathcal{Y}$, and sensitive information as random variable $S \in \mathcal{S}$, the goal is to predict Y from X in a way that is uninfluenced by S .

One solution to this problem is to learn *fair representations*, finding an intermediate representation $Z \in \mathcal{Z}$ that can then be used instead of X with existing machine learning architectures to make predictions, as done by [1], [2]. The new representation Z must simultaneously preserve information from X relevant to predicting Y while removing sensitive information that could lead to bias.

Learning fair representations can be done both by adversarial methods, such as [3]–[5], and through variational approaches such as [6], [7] that maximize and minimize different mutual information terms to ensure the representation is both expressive and fair. These approaches are related to the *information bottleneck (IB)* method [8] and its variational approximations [9]–[11] which aim to find a representation

that maximally compresses X while remaining informative about Y .

We develop a related variational IB method for fairness using Rényi divergence due to the extra degree of freedom provided by its tunable α parameter, rather than the more commonly used Kullback-Leibler divergence. We learn representations that fulfill three criteria: the representation is compact, expressive about Y in order to improve classification accuracy, and free from information that can lead to bias. Unlike most prior work that focuses on a single definition of fairness, we consider how to jointly address and balance two of arguably the most important measures for fairness, demographic parity and equalized odds, with possible benefits for ethicists and policy makers to translate policies into actual engineering systems.

Recent work based on Rényi information measures (and its variants) include an IB problem under a Rényi entropy complexity constraint [12], bounding the generalization error in learning algorithms [13], Rényi divergence variational inference [14], Rényi differential privacy [15] and the analysis and development of deep generative adversarial networks [16]–[19].

The rest of the paper is organized as follows. In Section II we derive a cost function for our method to be minimized and then in Section III present experimental results using our method to classify retinal fundus images from the EyePACS dataset. Finally, we conclude with a discussion in Section IV.

II. RÉNYI FAIR INFORMATION BOTTLENECK

We adopt a variational approach which we call *Rényi Fair Information Bottleneck* (RFIB) to encode data X into a new representation Z that can be used to draw inferences about Y free from bias due to S . In light of our model, we assume that the Markov chain $(Y, S) \rightarrow X \rightarrow Z$ holds. To simplify notation, 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 related to, but a stronger constraint than, equality of opportunity, 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 , i.e.,

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

while for equalized odds the goal is to achieve this independence by conditioning on the actual target Y , i.e.,

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

B. Lagrangian Formulation

To encourage equalized odds, we minimize $I(Z; S|Y)$; i.e., we minimize the average amount of information that Z has about S given Y . To both obtain good classification accuracy and help promote demographic parity, we maximize $I(Z; Y|S)$. Maximizing mutual information between Z and Y ensures the representation will be expressive about its target while the conditioning on S ensures that Z does not keep information shared by S , encouraging demographic parity.

In addition, we minimize $I(Z; X|S, Y)$, a compression term similar to one from the IB problem [9]. 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. Finally, we maximize the utility term $I(Z; Y)$; this optimization, similar to the IB problem, solely ensures the representation is maximally expressive of the target Y .

Combining these 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

$$\begin{aligned} \mathcal{L} = & I(Z; S|Y) + I(Z; X|S, Y) - \lambda_1 I(Z; Y) \\ & - \lambda_2 I(Z; Y|S), \end{aligned} \quad (3)$$

where λ_1 and λ_2 are hyperparameters. Reworking this Lagrangian, we have that:

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

$$= I(Z; X) - (\lambda_1 + 1)I(Z; Y) - \lambda_2 I(Z; Y|S), \quad (4)$$

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

$$\mathcal{L} = I(Z; X) - \beta_1 I(Z; Y) - \beta_2 I(Z; Y|S), \quad (5)$$

where $\beta_1 = \lambda_1 + 1$ and $\beta_2 = \lambda_2$. This simpler Lagrangian is easier to compute while maintaining similar properties to 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. The two hyperparameters β_1 and β_2 control trade-offs between accuracy and fairness, with higher β values corresponding to a higher priority on accuracy and lower β values giving more influence to the compression term $I(Z; X)$ that discards unwanted information, potentially improving fairness at the expense of accuracy. 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 in (5), 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 distributions $P_{Y|S,Z}$, $P_{Y|Z}$, and P_Z so we replace them with variational approximations $Q_{Y|S,Z}$, $Q_{Y|Z}$ and Q_Z , respectively. We next derive an upper bound for $I(Z; X)$ with the novel use of Rényi divergence:

$$\begin{aligned} I(Z; X) &= \sum_{(z,x) \in \mathcal{Z} \times \mathcal{X}} P_{Z,X}(z, x) \log \frac{P_{Z|X}(z|x)}{P_Z(z)} \\ &= \sum_{(z,x) \in \mathcal{Z} \times \mathcal{X}} P_{Z,X}(z, x) \log P_{Z|X}(z|x) \\ &\quad - D_{KL}(P_Z || Q_Z) - \sum_{z \in \mathcal{Z}} P_Z(z) \log Q_Z(z) \\ &\leq \sum_{(z,x) \in \mathcal{Z} \times \mathcal{X}} P_{Z,X}(z, x) \log \frac{P_{Z|X}(z|x)}{Q_Z(z)} \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}_{P_X} D_{KL}(P_{Z|X} \| Q_Z) \\
&\leq \mathbb{E}_{P_X} D_\alpha(P_{Z|X} \| Q_Z), \tag{6}
\end{aligned}$$

for $\alpha > 1$. The first inequality follows from the non-negativity of Kullback-Leibler (KL) divergence, similar to [2], [6], [9]. For the final step, we take the Rényi divergence $D_\alpha(\cdot \| \cdot)$ of order α (e.g., see [20]), rather than the KL divergence as typically done in the literature, where

$$D_\alpha(P \| Q) = \frac{1}{\alpha - 1} \log \left(\sum_{x \in \mathcal{X}} P(x)^\alpha Q(x)^{1-\alpha} \right) \tag{7}$$

for $\alpha > 0$, $\alpha \neq 1$ and distributions P and Q with common support \mathcal{X} .¹ Using Rényi divergence gives an extra degree of freedom and allows more control over the compression term $I(X; Z)$. As the Rényi divergence is non-decreasing with α , a higher α will more strongly force the distribution $P_{Z|X}$ closer to Q_Z , resulting in more compression.

The upper bound in (6) holds for $\alpha > 1$ since D_α is non-decreasing in α and $\lim_{\alpha \rightarrow 1} D_\alpha(P \| Q) = D_{KL}(P \| Q)$.² When $\alpha < 1$, then $\mathbb{E}_{P_X} D_\alpha(P_{Z|X} \| Q_Z)$ is no longer an upper bound on $I(Z; X)$; but it can be considered as a potentially useful approximation that is tunable by varying α .

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

$$I(Z; Y) \geq \mathbb{E}_{P_{Y,Z}} [\log Q_{Y|Z}(Y|Z)] + H(Y), \tag{8}$$

$$I(Z; Y|S) \geq \mathbb{E}_{P_{S,Y,Z}} [\log Q_{Y|S,Z}(Y|S, Z)] + H(Y|S). \tag{9}$$

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 [21]. Modeling $P_{Z|X}$ as a density, we let $P_{Z|X} dZ = P_E dE$ where E is a random variable 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:

$$\begin{aligned}
J_{\text{RFIB}} = & \frac{1}{N} \sum_{i=1}^N \left[D_\alpha(P_{Z|X=x_i} \| Q_Z) \right. \\
& - \beta_1 \mathbb{E}_E [\log(Q_{Y|Z}(y_i | f(x_i, E)))] \\
& \left. - \beta_2 \mathbb{E}_E [\log(Q_{Y|S,Z}(y_i | s_i, f(x_i, E)))] \right], \tag{10}
\end{aligned}$$

where we estimate the expectation over E using a single Monte Carlo sample.

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

²For simplicity and by the continuity property of D_α , we define its extended orders at $\alpha = 1$ and $\alpha = 0$ [20] as $D_1(P \| Q) := D_{KL}(P \| Q)$ and $D_0(P \| Q) := \lim_{\alpha \rightarrow 0} D_\alpha(P \| Q) = -\log Q(x : P(x) > 0)$, which is equal to 0 when P and Q share a common support.

We note that depending on the choice of α , β_1 , and β_2 , from our method we can recover both the *IB* [9] and *conditional fairness bottleneck* (CFB) [6] schemes to which we compare our results. Letting $\alpha = 1$ and $\beta_2 = 0$ corresponds to IB, while setting $\alpha = 1$ and $\beta_1 = 0$ corresponds to CFB.

III. EXPERIMENTS

We present experimental results on the EyePACS dataset of retinal images.

A. Data

The EyePACS dataset [22] 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 [23], we binarize this label into our prediction target Y , with $Y = 1$ corresponding to categories 1-4, considered a positive, referable case for DR, and $Y = 0$ corresponding to category 0, a healthy eye with no disease.

We are interested in skin tone for our sensitive variable S with $S = 0$ representing light skin and $S = 1$ dark. However, as skin tone is not included in the dataset, we instead use the Individual Topology Angle (ITA) [24] 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 [23], [25], we compute ITA via

$$\text{ITA} = \frac{180}{\pi} \arctan \left(\frac{L - 50}{b} \right) \tag{11}$$

where L is luminescence and b is ‘‘yellowness’’ in CIE-Lab space. We then binarize ITA where an ITA of ≤ 19 is taken to mean dark skin, as done in [23], [26]. Using ITA as a proxy for skin tone has the advantage of being significantly easier to determine compared to the potential issues arising in having a clinician manually annotate images as done in [27]. Sample images from the EyePACS dataset are shown in Fig. 1.




Fig. 1. Examples of retinal fundus images from the EyePACS dataset. For the left image $(Y, S) = (1, 0)$, for the middle $(Y, S) = (0, 0)$, and for the right $(Y, S) = (0, 1)$.

B. 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*).

2) *Measures of Fairness*: We measure this in multiple ways: a) using the gap in accuracy (denoted acc_{gap}) between favored and protected subpopulations; and 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} [28]. Also echoing [6], we measure: c) the adherence to demographic parity via its gap dp_{gap} ; and adherence to d) equalized odds via its gap $eqodds_{\text{gap}}$. The latter two metrics are respectively given by:

$$dp_{\text{gap}} = |P(\hat{Y} = 1|S = 0) - P(\hat{Y} = 1|S = 1)| \quad (12)$$

and

$$eqodds_{\text{gap}} = \max_{y \in \{0,1\}} |P(\hat{Y} = 1|S = 0, Y = y) - P(\hat{Y} = 1|S = 1, Y = y)|. \quad (13)$$

3) *Joint Utility-Fairness Measure*: Echoing and comparing with the work in [23], we use a single metric that jointly captures utility and fairness, the Conjunctive Accuracy Improvement (CAI_{λ}) measure:

$$CAI_{\lambda} = \lambda(acc_{\text{gap}}^b - acc_{\text{gap}}^d) + (1 - \lambda)(acc^d - acc^b) \quad (14)$$

where $0 \leq \lambda \leq 1$, and acc^b and acc^d are the accuracy for baseline and debiased algorithms, respectively, while acc_{gap}^b and acc_{gap}^d are gap in accuracy for the baseline and debiased algorithms. In practice, one uses either $\lambda = 0.5$ for equal balance between utility and fairness or $\lambda = 0.75$ to emphasize fairness.

C. Additional Implementation Details

We use an isotropic Gaussian distribution for the encoder with mean and variance learned by a neural network, $P_{Z|X} = \mathcal{N}(Z|\mu_{\text{enc}}(X), \sigma_{\text{enc}}^2(X)I_d)$, where the mean and variance are two outputs of the encoder, d is the dimension of Z , and I_d is the d -dimensional identity matrix. The representation is then computed as $Z = \mu_{\text{enc}}(X) + \sigma_{\text{enc}}\mathcal{N}(\underline{0}, I_d)$, where $\underline{0}$ is the all-zero vector of size d .

We model the approximation of the representation’s marginal as a d -dimensional spherical Gaussian, $Q_Z = \mathcal{N}(Z|\underline{0}, I_d)$. As a result, we calculate the Rényi divergence in (10) between the multivariate Gaussians $P_{Z|X}$ and Q_Z using the closed-form expression derived in [29], [30].

Finally, as we only use binary values for Y , we model $Q_{Y|Z}$ with Bernoulli distributions, $Q_{Y|Z} = \text{Bernoulli}(Y; f(Z))$ and $Q_{Y|Z,S} = \text{Bernoulli}(Y; g(Z, S))$ where f and g are auxiliary fully connected networks.

We use ResNet50 as the encoder network while f and g consist of two linear layers followed by 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.

D. Results

We predict $Y = \text{DR Status}$ while using $S = \text{ITA}$ as the sensitive attribute. We consider the case of an extreme data imbalance where training data is completely missing

for one protected subgroup (dark skin individuals) and for a specific value of $Y = 1$ (DR-referable individuals). 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. The goal is for predictions on the missing subgroup to be just as accurate as on the group with adequate training data. This is a problem of both fairness and also *domain adaptation*, and matches an important real world problem where data for dark skin individuals is lacking compared to light skin individuals.

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. We use the same partition as in [23] to compare with their method. For their method, we report their original CAI scores calculated with respect to their baseline whereas we calculate our CAI scores with respect to results from our own baseline, a ResNet50 network.

Hyper-parameter tuning: We use various combinations of hyperparameters β_1 and β_2 varied linearly from 1 to 50 and α varied linearly from 0 to 1, where $\alpha = 1$ signifies KL divergence instead of Rényi divergence. As values of $\alpha = 1$ and $\beta_2 = 0$ correspond to the IB method, to compare it with our method we first find an intermediate value of β_1 that performs well for the IB method, then fix our value of β_1 to the same value and vary α and β_2 . We compare our method to the CFB in the same way, fixing a value of β_2 that we use for both the CFB and our method, and then varying α and β_1 . We implement these methods ourselves and also compare to two methods with results taken from [23]: adversarial independence (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 underrepresented populations and performs data augmentation to train a less biased model.

As seen in Table I, our method mostly outperforms all other methods, showing improvements in accuracy and fairness across nearly all metrics. Usual caution should be exercised in interpretations since – despite our aligning with data partitioning in [23] – other variations may exist with [6], [9], [23] due to non-determinism, parameter setting or other factors.

IV. DISCUSSION

This study proposed RFIB, a new variational approach that encodes trade-offs between different principles of fairness along with utility and compactness. Compared to prior work which incorporates a single definition of fairness, our approach has the potential benefit of allowing ethicists and policy makers to specify softer and more balanced fairness requirements that may lie between multiple hard fairness requirements. Our work opens the way to more future studies expanding on this idea.

In additional experiments we evaluated the performance of our method when the fairness model for EyePACS was developed vis-a-vis the ITA protected factor but was then tested on a dataset where the protected factor was race, as

TABLE I
RESULTS FOR DEBIASING METHODS ON EYEPACS

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 (from [23])	70.0	3.5	68.3	-	-	NA	NA
AD ($\beta = 0.5$) ([23])	76.12	2.41	74.92 (L)	3.61	2.35	NA	NA
IA ([23])	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 ($\beta_1=30$) ([9])	74.12	2.08	73.08 (D)	3.37	4.69	18.58	20.67
CFB ($\beta_2=30$) ([6])	77.83	1.66	77.0 (L)	5.84	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) ($\alpha = 0.8, \beta_1 = 30, \beta_2 = 50$)	79.17	0.83	78.75 (L)	6.54	6.91	17.0	17.83
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

We predict $Y = DR$ Status, trained on partitioning with respect to $S = ITA$, 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.

annotated by clinicians. The results demonstrate the ability of our method to perform well in this type of protected factor domain adaptation problem, where our approach outperformed all other methods. This is particularly important in a setting where the actual protected factor is not revealed for privacy reasons. Regarding addressing methods that are both fair and private, the type of information blinding approach pursued here may have implications for models that protect against attribute inference attack, another possible future direction.

Overall, evaluation of the method showed benefits vis-a-vis methods of record including IB, CFB, and other augmentation or adversarial debiasing techniques, with the potential to further improve results by combining our method with other post-processing, pre-processing and data intervention methods.

REFERENCES

- [1] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning fair representations," in *International Conference on Machine Learning*, 2013, pp. 325–333.
- [2] J. Song, P. Kalluri, A. Grover, S. Zhao, and S. Ermon, "Learning controllable fair representations," in *Twenty-Second International Conference on Artificial Intelligence and Statistics*, 2019, pp. 2164–2173.
- [3] A. Beutel, J. Chen, Z. Zhao, and E. H. Chi, "Data decisions and theoretical implications when adversarially learning fair representations," *arXiv:1707.00075*, 2017.
- [4] H. Edwards and A. Storkey, "Censoring representations with an adversary," *arXiv:1511.05897*, 2015.
- [5] D. Madras, E. Creager, T. Pitassi, and R. Zemel, "Learning adversarially fair and transferable representations," in *International Conference on Machine Learning*. PMLR, 2018, pp. 3384–3393.
- [6] B. Rodríguez-Gálvez, R. Thobaben, and M. Skoglund, "A variational approach to privacy and fairness," in *Proceedings of the IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–6.
- [7] A. Ghassami, S. Khodadadian, and N. Kiyavash, "Fairness in supervised learning: An information theoretic approach," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2018, pp. 176–180.
- [8] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," in *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, 1999, p. 368–377.
- [9] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, "Deep variational information bottleneck," in *Proc. 5th Int. Conf. Learning Representations (ICLR)*, 2017, pp. 1–17.
- [10] I. Fischer, "The conditional entropy bottleneck," *Entropy*, vol. 22, no. 9, p. 999, 2020.
- [11] A. Kolchinsky, B. D. Tracey, and D. H. Wolpert, "Nonlinear information bottleneck," *Entropy*, vol. 21, no. 12, p. 1181, 2019.
- [12] J.-J. Weng, F. Alajaji, and T. Linder, "An information bottleneck problem with Rényi's entropy," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2021, pp. 2489–2494.
- [13] A. R. Esposito, M. Gastpar, and I. Issa, "Robust generalization via α -mutual information," in *Proceedings of the International Zurich Seminar on Information and Communication*, 2020, pp. 96–100. [Online]. Available: <https://doi.org/10.3929/ethz-b-000403224>
- [14] Y. Li and R. E. Turner, "Rényi divergence variational inference," in *Advances in Neural Information Processing Systems*, vol. 29, 2016, pp. 1073–1081.
- [15] I. Mironov, "Rényi differential privacy," in *IEEE 30th Computer Security Foundations Symposium (CSF)*, 2017, pp. 263–275.
- [16] H. Bhatia, W. Paul, F. Alajaji, B. Ghahserifard, and P. Burlina, "Least k th-order and Rényi generative adversarial networks," *Neural Computation*, vol. 33, no. 9, pp. 2473–2510, 2021.
- [17] A. Sarraf and Y. Nie, "RGAN: Rényi generative adversarial network," *SN Computer Science*, vol. 2, no. 1, p. 17, 2021.
- [18] Y. Pantazis, D. Paul, M. Fasoulakis, Y. Stylianou, and M. Katsoulakis, "Cumulant GAN," *arXiv:2006.06625*, 2020.
- [19] G. R. Kurri, T. Sypherd, and L. Sankar, "Realizing GANs via a tunable loss function," in *Proceedings of the IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–6.
- [20] T. van Erven and P. Harremoës, "Rényi divergence and Kullback-Leibler divergence," *IEEE Transactions on Information Theory*, vol. 60, no. 7, pp. 3797 – 3820, 2014.
- [21] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," *arXiv:1312.6114*, 2013.
- [22] EyePACS, "Diabetic retinopathy detection," Data retrieved from kaggle, <https://www.kaggle.com/agaldran/eyepacs>, 2015.
- [23] W. Paul, A. Hadzic, N. Joshi, F. Alajaji, and P. Burlina, "TARA: Training and representation alteration for AI fairness and domain generalization," *Neural Computation*, vol. 34, no. 3, pp. 716–753, 2022.
- [24] M. Wilkes, C. Y. Wright, J. L. du 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 dermatology*, vol. 151, no. 8, pp. 902–903, 2015.
- [25] M. Merler, N. Ratha, R. S. Feris, and J. R. Smith, "Diversity in faces," *arXiv:1901.10436*, 2019.
- [26] N. M. Kinyanjui, T. Odonga, C. Cintas, N. C. Codella, R. Panda, P. Sattigeri, and K. R. Varshney, "Estimating skin tone and effects on classification performance in dermatology datasets," *arXiv:1910.13268*, 2019.
- [27] P. Burlina, N. Joshi, W. Paul, K. D. Pacheco, and N. M. Bressler, "Addressing artificial intelligence bias in retinal diagnostics," *Translational Vision Science & Technology*, vol. 10, no. 2, pp. 1–13, 2021.
- [28] J. Rawls, *Justice as fairness: A restatement*. Harvard University Press, 2001.
- [29] M. Gil, F. Alajaji, and T. Linder, "Rényi divergence measures for commonly used univariate continuous distributions," *Information Sciences*, vol. 249, pp. 124–131, 2013.
- [30] J. Burbea, "The convexity with respect to Gaussian distributions of divergences of order α ," *Utilitas Mathematica*, vol. 26, pp. 171–192, 1984.