

Invariant Representations with Stochastically Quantized Neural Networks

Mattia Cerrato^{1*}, Marius Köppel^{2*}, Roberto Esposito³, Stefan Kramer¹

¹ Institute of Computer Science, Johannes Gutenberg-Universität Mainz, Germany

² Institute for Nuclear Physics, Johannes Gutenberg-Universität Mainz, Germany

³ Computer Science Department, Università degli Studi di Torino, Italy

{mcerrato, mkoepfel}@uni-mainz.de, roberto.esposito@unito.it, kramer@informatik.uni-mainz.de

Abstract

Representation learning algorithms offer the opportunity to learn invariant representations of the input data with regard to nuisance factors. Many authors have leveraged such strategies to learn fair representations, i.e., vectors where information about sensitive attributes is removed. These methods are attractive as they may be interpreted as minimizing the mutual information between a neural layer’s activations and a sensitive attribute. However, the theoretical grounding of such methods relies either on the computation of infinitely accurate adversaries or on minimizing a variational upper bound of a mutual information estimate. In this paper, we propose a methodology for direct computation of the mutual information between neurons in a layer and a sensitive attribute. We employ stochastically-activated binary neural networks, which lets us treat neurons as random variables. Our method is therefore able to minimize an upper bound on the mutual information between the neural representations and a sensitive attribute. We show that this method compares favorably with the state of the art in fair representation learning and that the learned representations display a higher level of invariance compared to full-precision neural networks.

1 Introduction

Representation learning algorithms based on neural networks are being employed extensively in information retrieval and data mining applications. The social impact of what the general public refers to as “AI” is now a topic of much discussion, with regulators in the EU even putting forward legal proposals which would require practitioners to “[...] minimize the risk of unfair biases embedded in the model [...]” (Commission 2021). Such proposals refer to biases with regard to individual characteristics which are protected by the law, such as gender and ethnicity. The concern is that models trained on biased data might then learn those biases (Barocas, Hardt, and Narayanan 2019), therefore perpetuating historical discrimination against certain groups of individuals. Machine learning methodologies designed to avoid these situations are often said to be “group-fair”.

One possible approach to the group fairness problem is fair representation learning. Fair representation learning is

a set of techniques which learns new representations of the original data where information about protected characteristics such as gender or ethnicity has been removed. Numerous authors have employed neural networks as the base learning algorithm in this task (Zemel et al. 2013; Xie et al. 2017; Madras et al. 2018). The core concept in fair representation learning is to remove sensitive information from a neural network’s representations. If one is able to guarantee that the network has no information about, for instance, an individual’s gender identity, it follows that the decisions undertaken by the model are independent of it. One possible formalization for the above desideratum relies on computing the mutual information between the ℓ -th neural layer T^ℓ and a sensitive attribute S representing the protected information. If the mutual information $I(T^\ell; S)$ equals 0, for instance, one would obtain group-invariant representations and therefore unbiased decisions which do not rely on S . Estimating mutual information over high-dimensional distributions is however highly challenging in general. In practice, an upper bound on $I(T^\ell; S)$ is sufficient for minimization purposes. Previous research has employed post-training quantization (Tishby and Zaslavsky 2015), variational approximations of the encoding distribution (Moyer et al. 2018; Belghazi et al. 2018) and adversarial learning (Madras et al. 2018; Xie et al. 2017). Also, the mutual information $I(T^\ell; S)$ is ill-defined in neural networks as they are deterministic functions of the input data. Implying that $I(X; T^\ell)$ is vacuous, it does not depend on the network parameters (Goldfeld et al. 2019).

In this work, we provide an alternative approach to fair representation learning by employing stochastically-quantized neural networks. These networks have a low-precision activation (1 bit) which may be interpreted as a random variable following the Bernoulli distribution. We show a sketch of these networks in Figure 1. Thus, we are able to obtain a finite, exact value for the mutual information $I(T_i^\ell; S)$ between a neuron i in any layer ℓ and the sensitive attribute, without relying on variational approximation (Moyer et al. 2018), adversarial learning (Madras et al. 2018) or adding noise to the representations (Goldfeld et al. 2019). We then employ this exact value to upper bound the overall mutual information in layer ℓ , $I(T^\ell; S)$, obtaining an optimization objective which leads to invariant representations. Furthermore, we discuss how to compute $I(T^\ell; S)$ in stochastically activated binary neural networks via den-

*These authors contributed equally.

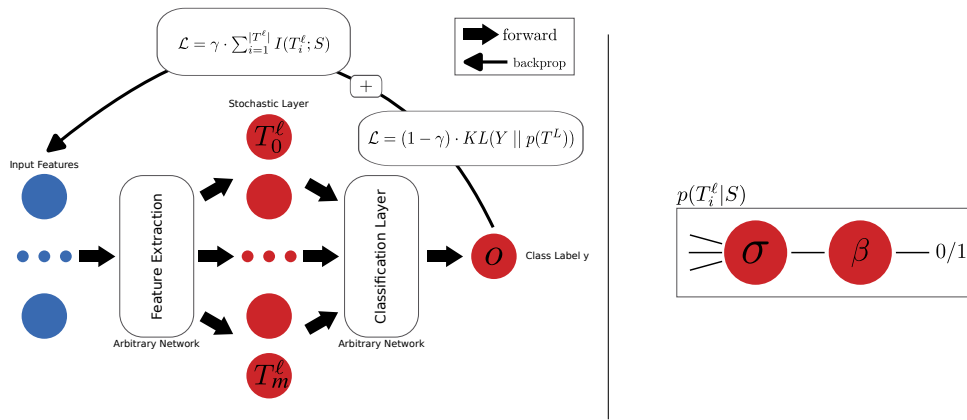


Figure 1: **Left:** sketch of a Stochastically Quantized Neural Network. The stochastic layer T^ℓ , holding the quantized neurons, is shown with $T_0^\ell \dots T_m^\ell$. During the forward step the input features are extracted via feature extraction layers before they enter the stochastic layer. After the stochastic layer further classification layers are used to classify the class label y . During the backpropagation the loss function is evaluated e.g. via binary cross entropy for the class label y and by calculating the mutual information of the stochastic layer T^ℓ and the sensitive attribute S . The influence of the two loss functions is controlled via the parameter γ . The feature extraction and classification layers may be chosen to be full-precision. **Right:** sketch of a stochastically quantized neuron. The neuron is sigmoid-activated, but the sigmoid output is employed as the parameter for a Bernoulli distribution, which we then sample from. This enables the interpretation of T_i^ℓ as a random variable and T^ℓ as a random vector, in turn allowing us to compute $I(T^\ell; S)$.

sity estimation. Our experimentation proves that quantized, low-precision models are able to obtain group-invariant representations in a fair classification setting and avoid training set biases in image data. Our contributions can be summarized as follows: I. We show how to compute the mutual information $I(T_i^\ell; S)$ between a stochastic binary neuron and the sensitive attribute S . II. We show how to use this value to upper bound the mutual information for the whole layer $I(T^\ell; S)$, which is a natural objective for invariant representation learning. III. We employ density estimation to compute the mutual information $I(T^\ell; S)$. IV. We perform experiments on three fair classification datasets and one invariant representation learning dataset, showing that our low-precision model is competitive with the state of the art.

2 Related Work

Algorithmic fairness has attracted significant attention from the academic community and general public in recent years, thanks in no small part to the ProPublica/COMPAS debate (Angwin et al. 2016; Rudin, Wang, and Coker 2020). To the best of our knowledge, however, the first contribution in this area dates back to 1996, when Friedman and Nisselbaum (Friedman and Nissenbaum 1996) contributed that automatic decision systems need to be developed with particular attention to systemic discrimination and moral reasoning. The need to tackle automatic discrimination is also part of EU-level law in the GDPR, Recital 71 in particular (Malgieri 2020).

One possible way to tackle these issues is to remove the information about the “nuisance factor” S from the data X by employing fair representation learning. Fair representation learning techniques learn a projection of X into a latent feature space Z where all information about S has been re-

moved. One seminal contribution in this area is due to Zemel et al. (Zemel et al. 2013). Since then, neural networks have been extensively used in this space. Some proposals (Xie et al. 2017; Madras et al. 2018) employ adversarial learning, a technique due to Ganin et al. (Ganin et al. 2016) in which two networks are pitted against one another in predicting Y and removing information about S . Another line of work (Louizos et al. 2016; Moyer et al. 2018) employs variational inference to approximate the intractable distribution $p(Z | X)$. A combination of architecture design (Louizos et al. 2016) and information-theoretic loss functions (Moyer et al. 2018; Gretton et al. 2012) may then be employed to encourage invariance of the neural representations with regard to S . Our proposal differs in that we employ a stochastically quantized neural network for which it is possible to compute the mutual information between the i -th neuron at layer ℓ and the sensitive attribute S . This lets us avoid using approximations of the target distribution $p(T^\ell, S)$ and provides a more stable training objective for representation invariance compared to adversarial training. More recently, neural architectures have been proposed for other fairness-related settings such as fair ranking (Zehlike and Castillo 2019; Cerrato et al. 2020; Narasimhan et al. 2020) and fair recourse (Sharma et al. 2021). Our experimental comparison will focus, however, on the fair/invariant classification setting so to enable a comparison with other methods in this area which bound or approximate information measures in neural representations.

3 Method

Our contribution deals with learning fair (group-invariant) representations in a principled way by employing stochastically quantized neural layers. In this section, we provide a

theoretical motivation for our work by contextualizing it in an information-theoretic framework similar to the one introduced in the “Information Bottleneck” literature (Tishby and Zaslavsky 2015; Goldfeld et al. 2019). As previously mentioned, the work done so far in this space has approximated measuring the information theoretic quantities in neural networks either via post-training quantization (Tishby and Zaslavsky 2015), adversarial bounding (Madras et al. 2018), variational approximations (Moyer et al. 2018) or the addition of stochastic noise in the representations (Goldfeld et al. 2019; Cerrato, Esposito, and Li Puma 2020). These approximations are necessary as the mutual information is an ill-defined concept in deterministic neural nets (Goldfeld and Polyanskiy 2020). Our approach is instead to employ stochastically quantized neural layers to exactly compute the mutual information between a sensitive attribute S and any neuron in the layer.

Invariant Representations and Mutual Information

A feedforward neural network with L layers may be formalized as a sequence of “layer functions” ϕ^ℓ which compute the neural activations given an input $x \in \mathbb{R}^{d_0}$:

$$\phi^\ell(x) = \sigma(A^\ell \phi^{\ell-1}(x) + b^\ell), \quad \ell = 1, \dots, L \quad (1)$$

$$\phi^0(x) = x, \quad (2)$$

where $A_\ell \in \mathbb{R}^{d_\ell \times d_{\ell-1}}$ is a weight matrix, $b \in \mathbb{R}^{d_\ell}$ is a bias vector, σ is an activation function, and d_ℓ is the size of the ℓ -th layer. We now define neural representations as applications of ϕ to the random variable X which follows the empirical distribution of the x samples:

$$T^\ell := \phi^\ell(X), \quad \ell = 1, \dots, L$$

We note that, in a supervised learning setting, the last layer T^L is a reproduction of Y . Representation invariance may be formalized as an information-theoretic objective in which the representations of the ℓ -th neural layer T^ℓ display minimal mutual information with regard to a sensitive attribute S :

$$I(T^\ell; S) \leq \alpha, \quad (3)$$

with $\alpha \approx 0$ being a threshold where the information can be considered minimal, and $I(T^\ell; S)$ is the mutual information between T^ℓ and S . Mutual information is an attractive measure in invariant representation learning, as two random variables are statistically independent if and only if $I(A; B) = 0$. Thus, one may obtain S -invariant representations by minimizing the mutual information between S and T^ℓ , therefore certifying that no information about the sensitive attribute is contained in the representations. Minimizing this objective might, however, remove all information about the original data X and, possibly, the labels Y . To avoid this issue, one might want to guarantee that some information about X is preserved. This changes the setting to a problem which closely resembles the Information Bottleneck problem (Tishby and Zaslavsky 2015) in which the task is to minimize $I(T^\ell; S)$ s.t. $I(T^\ell; X) \geq \beta$, where β is a positive real number. However, constrained optimization is

highly challenging in neural networks: in practice, previous work in this area has focused on, e.g., minimizing a reconstruction loss as a surrogate for the $I(T^\ell; X) \geq \beta$ constraint (Madras et al. 2018).

At the same time, computing $I(T^\ell; S)$ is very challenging. Since the distribution of T^ℓ is non-parametric and in general unknown, one would need to resort to density estimation to approximate it from data samples. However, the number of needed samples scales exponentially with the dimensionality of T^ℓ (Paninski 2003). Furthermore, mutual information is in general ill-defined in neural networks: if X is a random variable representing the empirical data distribution and f is a deterministic injective function, the mutual information $I(X; f(X))$ is either constant (when X is discrete) or infinite (when X is continuous) (Goldfeld et al. 2019). As activation functions in neural networks such as sigmoid and tanh are injective and real-valued, $I(T^\ell; S)$ is infinite. When the ReLU activation function is employed, the mutual information is finite but still vacuous – i.e. independent of the network’s parameters¹.

Previous work in fair representation learning has circumvented this issue by employing different techniques. One possible approach is to perform density estimation by grouping the real-valued activations into a finite number of bins (Tishby and Zaslavsky 2015). This approach, however, returns different values for the mutual information depending on the number of bins (Tishby and Zaslavsky 2015) and obtains the actual, “true” mutual information value only as the number of bins approaches infinity (Goldfeld and Polyanskiy 2020). Some authors have instead relied on variational approximations of the intractable distribution $p(T^\ell|X)$, which leads to a tractable upper bound (Moyer et al. 2018). Lastly, it is possible to bound the mutual information term $I(T^\ell; S)$ with the loss of an adversary network which tries to predict S from T^ℓ (Madras et al. 2018; Ganin et al. 2016; Xie et al. 2017). The adversary is, however, usually chosen to be a deterministic network, and thus this result suffers from the issues described in this section.

Our approach is instead to employ stochastically-quantized neural layers, a technique used in binary networks (Courbariaux and Bengio 2016), enabling the interpretation of neural activations as (discrete) random variables. This approach has multiple benefits: It does not rely on variational approximations or adversarial training; it lets us treat T^ℓ as a random variable, avoiding the infinite mutual information issue described above; lastly, it avoids adding noise to the representations as a way to obtain stochasticity (Goldfeld et al. 2019).

Mutual Information Computation via Bernoulli activations

Our methodology relies on stochastically quantized neural layers in which the activations are stochastic. More specifically, we employ *binary* neural layers in which either the

¹Proper contextualization of this result requires some additional preliminary results, which we will avoid reporting here due to space constraints. We refer the interested reader to Goldfeld and Polyanskiy (Goldfeld and Polyanskiy 2020).

weights or the activations have 1-bit precision. In a binary network, T_i^ℓ may be computed either deterministically or stochastically from the activations of the previous layer $T^{\ell-1}$ and the learnable weights and biases w connecting the two. For the sake of presentation, we report in the following the formula for deterministic computation of T^ℓ :

$$T_i^\ell = \mathbb{I} \left[\frac{1}{|T^{\ell-1}|} \sum_{i=1}^{|T^{\ell-1}|} w_i T_i^{\ell-1} \geq 0.5 \right],$$

where \mathbb{I} is the indicator function. In this paper, however, we employ stochastic quantization of binary neurons, which we compute by sampling from $\mathcal{B}(\theta_i^\ell)$, where:

$$\theta_i^\ell = \sigma \left(\frac{1}{|T^{\ell-1}|} \sum_{i=1}^{|T^{\ell-1}|} w_i T_i^{\ell-1} \right), \quad (4)$$

where σ is the sigmoid function. We then sample from this distribution to compute the actual activation T_i^ℓ . Thus, T_i^ℓ may be interpreted as a random variable following the Bernoulli distribution: $T_i^\ell \sim \mathcal{B}(\theta_i^\ell)$. The entropy of a random variable following the Bernoulli distribution has a closed, analytical form:

$$H(T_i^\ell) = -(1 - \theta_i^\ell) \cdot \log_2(1 - \theta_i^\ell) - \theta_i^\ell \cdot \log_2(\theta_i^\ell). \quad (5)$$

Recalling that the definition of mutual information may be rewritten in terms of entropy as $I(T_i^\ell; S) = H(T_i^\ell) - H(T_i^\ell | S)$, we then note that it is possible to compute it exactly for a sensitive attribute S and a neuron T_i^ℓ . The conditional entropy $H(T_i^\ell | S = s)$ may be computed by selecting those representations $T_i^\ell = \phi(x)$ for which it is true that $S = s$. We then consider the whole layer as a stochastic random vector $T^\ell = [T_1^\ell, T_2^\ell, \dots, T_m^\ell]$. We show in the following lemma that $\sum_{i=1}^{|T^\ell|} I(T_i^\ell; S) \geq I(T^\ell; S)$. Thus, minimizing $\sum_{i=1}^{|T^\ell|} I(T_i^\ell; S)$ will minimize $I(T^\ell; S)$.

Lemma. *Let T^ℓ be a random vector of a given layer in the network, S a random variable representing the sensitive attribute, $|T^\ell|$ the number of neurons in T^ℓ and $TC(T^\ell; S)$ the informativeness (Gao et al. 2019). Then, the mutual information $I(T^\ell; S)$ is minimized if $\sum_{i=1}^{|T^\ell|} I(T_i^\ell; S)$ is minimized.*

Proof. We recall the definition of total correlation (Watanabe 1960a) for a random vector $A = (a_1, \dots, a_n)$ having the probability density function $p(A)$:

$$\begin{aligned} TC(A) &= \sum_{i=1}^n H(a_i) - H(A) \\ &= KL[p(A) || p(a_1)p(a_2)\dots p(a_n)] \end{aligned} \quad (6)$$

We then define the informativeness of the sensitive attribute and a layer T^ℓ in the network as:

$$TC(T^\ell; S) = TC(T^\ell) - TC(T^\ell | S), \quad (7)$$

where $TC(T^\ell)$ is the total correlation (Watanabe 1960b) or multi-information (Studený and Vejnarová 1998) of T^ℓ and

$TC(T^\ell | S)$ the conditional total correlation of T^ℓ given S . One can rewrite the total correlation for T^ℓ and S as entropy terms following Equation 6. Similarly, the conditional total correlation may be reformulated as:

$$TC(T^\ell | S) = \sum_{i=1}^{|T^\ell|} H(T_i^\ell | S) - H(T^\ell | S). \quad (8)$$

Following the derivation of Gao et al. (Gao et al. 2019) Equation 7 can be rewritten, using Equation 6 and Equation 8, to:

$$\begin{aligned} TC(T^\ell; S) &= \sum_{i=1}^{|T^\ell|} H(T_i^\ell) - H(T^\ell) \\ &\quad - \sum_{i=1}^{|T^\ell|} H(T_i^\ell | S) + H(T^\ell | S) \\ &= \sum_{i=1}^{|T^\ell|} I(T_i^\ell; S) - \underbrace{(H(T^\ell) - H(T^\ell | S))}_{I(T^\ell; S)}. \end{aligned}$$

Since $TC(T^\ell; S) \geq 0$, $I(T_i^\ell; S) \geq 0$ and $I(T^\ell; S) \geq 0$ it follows that:

$$\sum_{i=1}^{|T^\ell|} I(T_i^\ell; S) \geq I(T^\ell; S). \quad (9)$$

□

Thus, given a stochastically quantized binary neural layer, we can impose representation invariance with respect to S by stochastic gradient descent over a loss function \mathcal{L} that directly incorporates $I(T_i^\ell; S)$:

$$\mathcal{L} = \gamma \cdot \sum_{i=1}^{|T^\ell|} I(T_i^\ell; S) + (1 - \gamma) \cdot KL(Y || p(T^L)), \quad (10)$$

where KL is the Kullback-Leibler divergence and γ is a trade-off parameter weighting the importance of representation invariance and accuracy on Y . It is worthwhile to mention that only a single layer ℓ needs to be stochastically quantized for the mutual information term to be computable. We explored the performance of both a “hybrid” network with a mix of binary-precision and full-precision layers and a fully binary-precision network (see supplementary material). Nevertheless, we found in our experiments that the “hybrid” method has the strongest performers in all cases.

We also note that stochastically quantized binary neural networks are a natural fit for density estimation techniques. As the network is stochastic and quantized by default, it is possible to estimate the conditional distribution $p(T^\ell | S)$ from samples while avoiding the issues with post-training quantization and ill-definedness described earlier in this section and in the literature (Goldfeld et al. 2019; Goldfeld and Polyanskiy 2020). We do not give specifics about this method in this section, as it relies on standard counting techniques. However, we give a full description of this strategy

in the supplementary material. We test this method in the next section, where we dubbed it `BinaryMI` to contrast it with the methodology described previously in this section, which we will refer to as `BinaryBernoulli`.

4 Experiments

Our experimental setting focuses on analyzing the performance of the methodology presented in this paper in fair classification and invariant representation learning settings. Our experimentation aims to answer the following questions:

Q1. Is the present methodology able to learn fair models which avoid discrimination? **A1.** Yes. We analyze our models’ accuracy and fairness by measuring the area under curve (AUC) and their disparate impact / disparate mistreatment. We compare them with full-precision neural networks trained for fair classification and see that our approach is able to find strong accuracy/fairness tradeoffs under the assumption that both are equally important. We also observe that supervised classifiers trained with the learned representations and the sensitive data are unable to generalize to the test set, as their accuracy is close to random guessing performance. Furthermore, we show that our proposed models are able to remove nuisance factors from image datasets.

Q2. Is the present methodology *stable*, i.e. is it able to learn different tradeoffs between accuracy and fairness? **A2.** Yes. Our method is able to explore the fairness-accuracy tradeoff with a higher degree of stability than previously proposed neural models when changing the tradeoff parameter γ . We observe high positive correlation between higher γ and higher fairness for both `BinaryBernoulli` and `BinaryMI`.

Datasets

COMPAS. This dataset was released by ProPublica (Angwin et al. 2016). The ground truth is whether an individual committed a crime in the following two years. The sensitive attribute is the individual’s ethnicity. Machine learning models trained on this dataset may display disparate mistreatment (Zafar et al. 2017), thus our objective is minimizing an equal opportunity metric, the Group-dependent Pairwise Accuracy, while maximizing accuracy (AUC).

Adult. This dataset is part of the UCI repository (Dua and Graff 2017). The ground truth is whether an individual’s annual salary is over 50K\$ per year or not (Kohavi 1996). This dataset has been shown to be biased against gender (Louizos et al. 2016; Zemel et al. 2013).

Bank marketing. In this dataset, the classification goal is whether an individual will subscribe a term deposit. Models trained on this dataset may display both disparate impact and disparate mistreatment with regard to age, more specifically on whether individuals are under 25 and over 65 years of age.

Biased-MNIST. This is an image dataset based on the well-known MNIST Handwritten Digits database in which the background has been modified so to display a color bias (Bahng et al. 2020). More specifically, a nuisance factor C is introduced which is highly correlated with the ground truth

Y and whose values represent the background color in the training set. Ten different colors are pre-selected for each value of $Y = \{0\dots9\}$ and inserted as a background in the training images with high probability ($p = 0.99, 0.995$). The test images, on the other hand, have background color chosen at random. We show samples from the training and test data in the supplementary material. Thus, the background color/nuisance factor provides a very strong training bias. The simplest strategy for a model to achieve high accuracy on the training set is to overfit the background color. Therefore, models that are unable to learn invariant representations and decisions will inevitably overfit the training set (Bahng et al. 2020).

Metrics

Group-dependent Pairwise Accuracy. We employ this to test the disparate mistreatment of our models. Its introduction is due to Narashiman et al. (Narasimhan et al. 2020). We report its definition in the following.

Let G_1, \dots, G_K be a set of K groups such that every instance inside the dataset \mathcal{D} belongs to one of these groups. The *group-dependent pairwise accuracy* $A_{G_i > G_j}$ is then defined as the accuracy of a classifier on instances which are labeled as positives and belong to group G_i and instances labeled as negatives which belong to group G_j . Group-dependent pairwise accuracy is then defined as $|\hat{A}_{G_i > G_j} - A_{G_j > G_i}|$, and should be close to zero. The rationale for this metric is that the false positive and false negative rates should be equalized across groups if possible, similarly to the notions of equality of opportunity (Hardt, Price, and Srebro 2016) or disparate mistreatment (Zafar et al. 2017). In the following, we call the Group-dependent Pairwise Accuracy *GPA*.

Area under Discrimination Curve (AUDC) We take the discrimination as a measure of disparate impact (Zemel et al. 2013), which is given by:

$$y\text{Discrim} = \left| \frac{\sum_{n:s_n=1}^n \hat{y}_n}{\sum_{n:s_n=1}^n 1} - \frac{\sum_{n:s_n=0}^n \hat{y}_n}{\sum_{n:s_n=0}^n 1} \right|,$$

where $n : s_n = 1$ denotes that the n -th example has a value of s equal to 1. We then generalize this metric in a similar fashion to how accuracy may be generalized to obtain a classifier’s area under the curve (AUC): We evaluate the measure above for different classification thresholds and then compute the area under this curve. We employed 100 equispaced thresholds in our experiments. In the following, we will refer to this measure as AUDC (area under the discrimination curve) as done elsewhere in the literature (Cer-rato et al. 2022). Contrary to AUC, lower values are better.

Experimental setup

We split all datasets into 3 internal and 3 external folds. On the 3 internal folds, we employ a Bayesian optimization technique to find the best hyperparameters for our model. A summary of our models’ best hyperparameters can be found in the supplementary material. As our interest is to obtain models which are both fair and accurate, we employ Bayesian optimization to maximize the sum of the

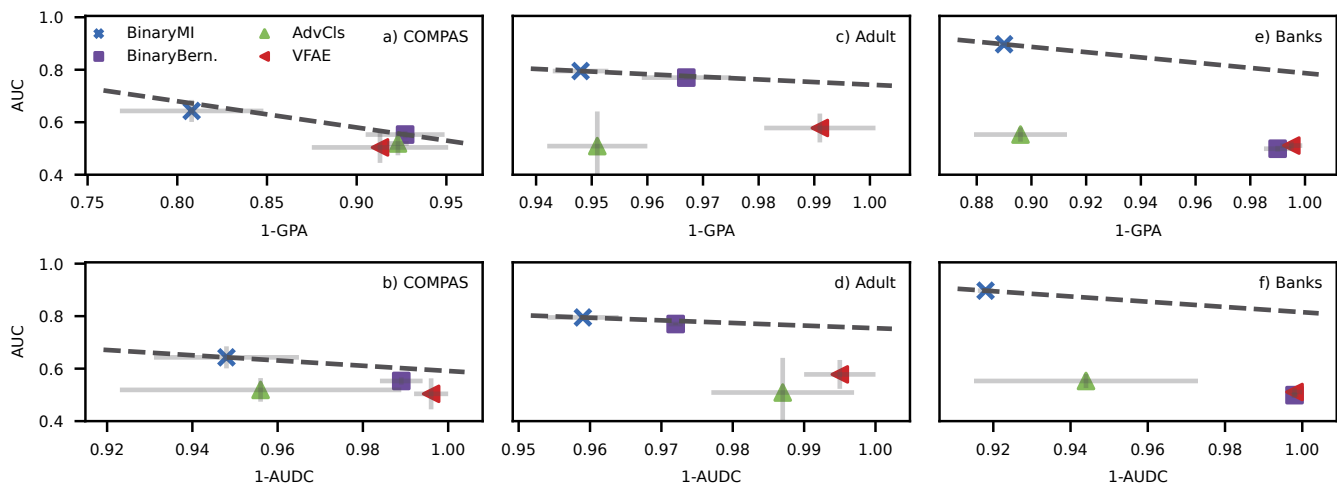


Figure 2: Experiment results for our Bernoulli entropy model (BinaryBernoulli), our density estimation model (BinaryMI), an adversarial classifier (AdvCls), and a fair variational autoencoder (VFAE). The dotted line represents the line of equivalent fairness/accuracy tradeoffs and goes through the model closest (under the L_1 norm) to perfect accuracy and perfect fairness. We show on top our best performing models for the AUC/1-GPA tradeoff, and on the bottom the best performing models for AUC/1-AUDC.

models’ AUC, 1-GPA and 1-AUDC. We set the maximum number of iterations to 200. The best hyperparameter setting found this way is then evaluated on the 3 external folds and reported. We relied on the Weights & Biases platform for an implementation of Bayesian optimization and overall experiment tracking (Biewald 2020). On the fairness datasets, we compare with an adversarial classifier (AdvCls in the Figures) trained as described by Xie et al. (Xie et al. 2017) and a variational fair autoencoder (VFAE) (Louizos et al. 2016). We obtained publicly available implementations for these models and optimized the hyperparameters with the same strategy we employed for our models. We report our model relying on the Bernoulli entropy computation as BinaryBernoulli in the figures, while the model relying on density estimation to compute the joint $p(T^\ell)$ is referred to as BinaryMI.

Fair and Invariant Classification

We report plots analyzing the accuracy/fairness tradeoff of our models trained for classification in Figure 2. We take AUC as our measure of accuracy and both 1-GPA and 1-AUDC as the fairness metric. The ideal model in the fair classification setting displays maximal AUC and little to no GPA/AUDC and would appear at the very top right in Figure 2. This result is not attainable on the datasets we consider, as there is usually some correlation between S and Y , which prevents us to obtain perfectly fair and accurate decisions. Thus, one needs to consider possible accuracy/fairness tradeoffs. We assume a balanced accuracy/fairness tradeoff and consider as the “best” model the one closest to (1,1) under the L_1 norm. We then show all equivalent tradeoff points as a dotted line. We see that on both COMPAS and Adult, our BinaryMI model is able to find a stronger tradeoff than the competitors, with

BinaryBernoulli very close to an equivalent tradeoff. The same may be said for the third dataset, Banks, in which however our two models find very different tradeoffs, with BinaryBernoulli preferring an almost perfectly fair result and BinaryMI finding a very accurate model. Compared to adversarial learning and variational inference, our models either find the best tradeoff or lie closest to the best tradeoff line. We also analyze the accuracy/fairness tradeoff of the representations learned by our models. We extract neural activations from all the networks considered at the penultimate layer $T^{\ell-1}$, which we stochastically quantized in all experiments. We then report in Figure 4 in the supplementary material the performance of a Random Forest algorithm with 1000 base estimators trained to predict the sensitive attribute S associated with each representation. We observe that BinaryMI is able to find the best tradeoff between informativeness on Y (bottom row) and invariance to S . Our best BinaryBernoulli model representations, which are comparatively not very invariant on COMPAS, performed strongly in this regard on both Adult and Banks.

Stability and Complexity Analysis

It is important that fair classification models are able to find different tradeoffs between fairness and accuracy depending on the application requirements. This tradeoff may be regulated, in practical terms, via a parameter γ which weights the importance of accuracy and invariance in the loss function. This idea is common, to the best of our knowledge, to all the fair representation learning algorithms developed so far. We explore how the performance of our model changes with γ in Figure 3, where we report the performance of a BinaryBernoulli model trained on the COMPAS dataset, an adversarial classifier and a fair variational autoencoder for comparison. What we observe in Figure 3 is




Figure 3: 1-GPA vs. AUC plot for BinaryBernoulli, AdvCls and VFAE while varying the γ parameter. The values are regressed with smoothed lines using splines fitted on the result for each method separately. The color of the spline segments represents the value of γ for the specific run. The data used for these experiments was the COMPAS dataset.

that our model displays a relatively stable performance. As γ grows, so does the fairness of the model in terms of 1-GPA (linear correlation $\rho = 0.924$). The model is highly sensitive to values of γ between 0.01 and 0.5. The same trend, but reversed, may be observed for AUC ($\rho = -0.836$). We observed similar trends for BinaryMI, with correlations of 0.781 and -0.760 for 1-GPA and AUC, respectively. Thus, we reason that our proposal may be employed under different fairness requirements with minimal changes (a tweak of the γ parameter). The adversarial classifier, on the other hand, displays little correlation between γ and its performance. While the best models are on par or almost on par with BinaryBernoulli, this happens for arbitrary values of γ . The adversarial model does not seem to be able to explore the fairness/accuracy tradeoff quite as well, and the effect of the γ parameter is unpredictable. We posit that this behavior may be due to the difficulty of striking a balance between the predictive power of the two sub-networks which predict Y and S alternatively (Xie et al. 2017). This is a well-known issue for generative adversarial models which pit different networks against each other (Chu, Minami, and Fukumizu 2020). We also note that the variational approximation-based model (VFAE) struggles to come to an accurate result, whereas it mostly takes fair decisions. The correlation between γ , which controls the strength of the Maximum Mean Discrepancy regularization in this model, and AUC/1-GPA is also quite low (0.173 and -0.168 respectively).

Biased-MNIST

We report our method’s performance on the Biased-MNIST dataset in Table 1. We also report results from Bahng et al. (Bahng et al. 2020) as a comparison. To enable this comparison, we experimented with the same setup as the authors’ by training our model for 80 epochs. As usual we selected our best hyperparameters with a Bayesian optimization strategy

Table 1: Accuracies for the biased MNIST experiments. Results for other methodologies as reported by Bahng et al. (Bahng et al. 2020). We report results for two bias levels $\rho = 0.99$ and 0.995. A higher value of ρ implies a higher chance of a biased sample in the training set.

ρ	Vanilla	ReBias	LearnedMixIn	RUBi	BinaryMI	BinaryBernoulli
0.995	72.1	76.0	78.2	90.4	89.08	90.64
0.990	89.1	88.1	88.3	93.6	88.54	96.02

employing an Alexnet-like structure (Krizhevsky, Sutskever, and Hinton 2012) by alternating (binary) convolutional layers and max-pooling layers. We report the full information for the best hyperparameters we found in the supplementary material. We then tested with two different bias levels, i.e. the probability of a training sample displaying a specific color bias. We observe that BinaryBernoulli is the strongest performer on both bias levels, even when compared with other full-precision strategies. BinaryMI has a comparable performance to a “vanilla” convolutional network when the probability of training bias is $\rho = 0.99$. However, it displays better scaling to the higher bias level than the baseline method. In this dataset, we see that our methodology is also a strong performer when removing biases from image data is necessary for classification accuracy.

5 Conclusion and Future Work

In this paper we proposed a methodology to compute the mutual information between a stochastically activated neuron and a sensitive attribute. We then generalized this methodology into two different strategies to compute the mutual information between a layer of neural activations and a sensitive attribute. Both our strategies perform strongly on both fair classification datasets and invariant image classification. Furthermore, our methodology displays high stability to changes of the accuracy/fairness tradeoff parameter γ , especially when compared to adversarial learning (Xie et al. 2017; Madras et al. 2018). A possible direction for further development is to employ the methodologies discussed in this paper to revisit the debate on the information bottleneck problem introduced by Tishby et al. (Tishby and Zaslavsky 2015). As our models are partly stochastically quantized, they naturally lend themselves to mutual information computation, avoiding many of the common issues in estimating information measures in neural networks (Goldfeld et al. 2019). While estimating the conditional probability $p(T^\ell | S)$ does seem to scale to very wide networks (i.e. layers with many neurons), it is possible to repeat the sampling (the stochastic quantization) as many times as needed, which could alleviate this issue. Furthermore, we would like to test the capabilities of the presented models in domain adaptation scenarios, where adversarial models are still used extensively (Ganin et al. 2016; Madras et al. 2018).

6 Acknowledgements

We thank Alexander Segner for useful discussions throughout the development of this paper. The paper was partially supported by the “TOPML: Trading Off Non-Functional Properties of Machine Learning” project funded by the Carl-Zeiss-Stiftung in the Förderprogramm “Durchbrüche”, identifying code P2021-02-014. This work has been partially supported by the European PILOT project and by the HPC4AI project (Aldinucci et al. 2018).

References

- Aldinucci, M.; Rabellino, S.; Pironi, M.; Spiga, F.; Viviani, P.; Drocco, M.; et al. 2018. HPC4AI, an AI-on-demand federated platform endeavour. In *ACM Computing Frontiers*. Ischia, Italy.
- Angwin, J.; Larson, J.; Mattu, S.; and Kirchner, L. 2016. Machine Bias. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>. Accessed: 2023-04-24.
- Bahng, H.; Chun, S.; Yun, S.; Choo, J.; and Oh, S. J. 2020. Learning de-biased representations with biased representations. In *International Conference on Machine Learning*, 528–539. PMLR.
- Barocas, S.; Hardt, M.; and Narayanan, A. 2019. Fairness and Machine Learning. <http://www.fairmlbook.org>. Accessed: 2023-04-24.
- Belghazi, M. I.; Baratin, A.; Rajeshwar, S.; Ozair, S.; Bengio, Y.; Courville, A.; and Hjelm, D. 2018. Mutual information neural estimation. In *International Conference on Machine Learning*, 531–540. PMLR.
- Biewald, L. 2020. Experiment Tracking with Weights and Biases. <https://www.wandb.com/>. Accessed: 2023-04-24.
- Cerrato, M.; Coronel, A. V.; Köppel, M.; Segner, A.; Esposito, R.; and Kramer, S. 2022. Fair Interpretable Representation Learning with Correction Vectors.
- Cerrato, M.; Esposito, R.; and Li Puma, L. 2020. Constraining Deep Representations with a Noise Module for Fair Classification. In *ACM SAC*.
- Cerrato, M.; Köppel, M.; Segner, A.; Esposito, R.; and Kramer, S. 2020. Fair pairwise learning to rank. *DSAA*.
- Chu, C.; Minami, K.; and Fukumizu, K. 2020. Smoothness and Stability in GANs. In *International Conference on Learning Representations*.
- Commission, E. 2021. Regulatory framework proposal on artificial intelligence. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>. Accessed: 2023-04-24.
- Courbariaux, M.; and Bengio, Y. 2016. BinaryNet: Training Deep Neural Networks with Weights and Activations Constrained to +1 or -1. *CoRR*, abs/1602.02830.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>. Accessed: 2023-04-24.
- Friedman, B.; and Nissenbaum, H. 1996. Bias in computer systems. *ACM Transactions on Information Systems (TOIS)*, 14(3): 330–347.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; et al. 2016. Domain-adversarial training of neural networks. *JMLR*, 17.
- Gao, S.; Brekermans, R.; Steeg, G. V.; and Galstyan, A. G. 2019. Auto-Encoding Total Correlation Explanation. In *AIS-TATS*.
- Goldfeld, Z.; and Polyanskiy, Y. 2020. The information bottleneck problem and its applications in machine learning. *IEEE Journal on Selected Areas in Information Theory*, 1(1): 19–38.
- Goldfeld, Z.; van den Berg, E.; Greenewald, K. H.; Melnyk, I.; Nguyen, N.; Kingsbury, B.; and Polyanskiy, Y. 2019. Estimating Information Flow in Deep Neural Networks. In *ICML*.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(3): 723–773.
- Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29.
- Kohavi, R. 1996. Scaling up the Accuracy of Naive-Bayes Classifiers: A Decision-Tree Hybrid. *KDD*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *NeurIPS*.
- Louizos, C.; Swersky, K.; Li, Y.; Welling, M.; and Zemel, R. 2016. The Variational Fair Autoencoder. *ICLR*.
- Madras, D.; Creager, E.; Pitassi, T.; and Zemel, R. S. 2018. Learning Adversarially Fair and Transferable Representations. *CoRR*, abs/1802.06309.
- Malgieri, G. 2020. The Concept of Fairness in the GDPR: A Linguistic and Contextual Interpretation. *FAT**.
- Moyer, D.; Gao, S.; Brekermans, R.; Galstyan, A.; and Ver Steeg, G. 2018. Invariant representations without adversarial training. *NeurIPS*.
- Narasimhan, H.; Cotter, A.; Gupta, M.; and Wang, S. 2020. Pairwise Fairness for Ranking and Regression. *AAAI*.
- Paninski, L. 2003. Estimation of entropy and mutual information. *Neural computation*, 15(6): 1191–1253.
- Rudin, C.; Wang, C.; and Coker, B. 2020. The Age of Secrecy and Unfairness in Recidivism Prediction. *Harvard Data Science Review*, 2(1). <https://hdsr.mitpress.mit.edu/pub/7z10o269>.
- Sharma, S.; Gee, A. H.; Paydarfar, D.; and Ghosh, J. 2021. FaiR-N: Fair and Robust Neural Networks for Structured Data. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery.
- Studený, M.; and Vejnarová, J. 1998. The Multiinformation Function as a Tool for Measuring Stochastic Dependence. *Learning in Graphical Models*, 261–297.
- Tishby, N.; and Zaslavsky, N. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, 1–5. IEEE.
- Watanabe, S. 1960a. Information Theoretical Analysis of Multivariate Correlation. *IBM Journal of Research and Development*, 4(1): 66–82.

Watanabe, S. 1960b. Information Theoretical Analysis of Multivariate Correlation. *IBM Journal of Research and Development*, 4(1): 66–82.

Xie, Q.; Dai, Z.; Du, Y.; Hovy, E.; and Neubig, G. 2017. Controllable Invariance through Adversarial Feature Learning. *NeurIPS*.

Zafar, M. B.; Valera, I.; Gomez Rodriguez, M.; et al. 2017. Fairness beyond disparate treatment & disparate impact. *WWW*.

Zehlike, M.; and Castillo, C. 2019. Reducing disparate exposure in ranking: A learning to rank approach. *WWW*.

Zemel, R.; Wu, Y.; Swersky, K.; Pitassi, T.; and Dwork, C. 2013. Learning fair representations. *ICML*.