




FairNN - Conjoint Learning of Fair Representations for Fair Decisions

Tongxin Hu¹, Vasileios Iosifidis², Wentong Liao ¹, Hang Zhang¹,
Michael Ying Yang ³, Eirini Ntoutsi ², and Bodo Rosenhahn¹

¹Institut für Informationsverarbeitung, Leibniz University of Hanover, Germany

²Leibniz University Hanover & L3S Research Center, Germany

³Scene Understanding Group, University of Twente, The Netherlands

¹{second_name}@tnt.uni-hannover.de, ²{second_name}@l3s.de,

³michael.yang@utwente.nl

Abstract. In this paper, we propose *FairNN* a neural network that performs joint feature representation and classification for fairness-aware learning. Our approach optimizes a multi-objective loss function which (a) learns a fair representation by suppressing protected attributes (b) maintains the information content by minimizing the reconstruction loss and (c) allows for solving a classification task in a fair manner by minimizing the classification error and respecting the equalized odds-based fairness regularizer. Our experiments on a variety of datasets demonstrate that such a joint approach is superior to separate treatment of unfairness in representation learning or supervised learning. Additionally, our regularizers can be adaptively weighted to balance the different components of the loss function, thus allowing for a very general framework for conjoint fair representation learning and decision making.

Keywords: Fairness, Bias, Neural Networks, Auto-encoders

1 Introduction

The wide usage of AI-based systems, mostly powered nowadays by data and machine learning algorithms, in areas of high societal impact raises a lot of concerns regarding accountability, fairness, and transparency [25] of their decisions. Such systems can become discriminatory towards groups of people or individuals based on *protected attributes* like gender, race, religious beliefs etc, as it has been already showcased in a variety of cases [9,3,5]. For example, [3] shows that Google’s ad-targeting system was displaying more highly paid jobs to men than to women, thus making discriminatory decisions based on gender. Such incidents call for methods that explicitly target bias and discrimination in AI-systems, while maintaining their predictive power. The ever increased interest in this area is already reflected in the large, given the recency of the field, body of literature on fairness-aware learning and responsible AI, in general (see [23] for a recent survey).

However, despite the large number of methods and approaches for fairness-aware machine learning proposed thus far, most of these approaches refer to supervised learning upon a given feature representation. Some approaches that target fair representation learning also exist, e.g., [28] but they focus on learning a fair lower dimensional representation of the data which can be used either as a standalone result (e.g., for visualization purposes), or as an input to some other learning task (e.g., for learning a classifier upon the reduced representation) [26]. Only a few approaches exist that jointly target fairness in both representation learning *and* supervised learning, e.g., [6,19,22,24].

In this work we argue that a joint tackling of fairness in the machine learning pipeline (data \rightarrow algorithm \rightarrow model) is superior to the separate treatment of unfairness in representation- or supervised- learning. This is because bias-related corrections in representation learning do not guarantee that a model derived from the corrected data will be fair. Instead, the learning algorithm might still pick up certain data peculiarities that lead to discriminatory outcomes. Therefore, a joint *goal-oriented* consideration in the pipeline is much more effective, as also demonstrated in our experimental results. To this end, we aim for a fair representation learning that preserves as much as possible the original data while obfuscating information on the protected attribute so decisions based on the protected attribute in the latent space are not possible. Additionally, the learned representation should structure itself in such a fashion, that a task-goal, such as a classification task, can still be appropriately solved.

Our contributions can be summarized as follows:

- We propose a neural network that learns a fair representation and a fair classifier *jointly* in an end-to-end manner.
- The contribution of the different components during training can be adjusted, leading to a very flexible and competitive framework.
- Our experiments demonstrate that *FairNN* with a goal-oriented fair representation is superior to a plain fair classifier without explicit representation constraints as well as to a standard fair representation learner without an explicit classification goal.
- The source code is available ¹.

The rest of the paper is organized as follows: Related work is summarized in Section 2. Necessary background is provided in Section 3. Our joint goal-oriented approach to fairness-aware learning is introduced in Section 4. Experimental results are presented in Section 5. Finally, Section 6 concludes our work and identifies interesting directions for future research.

2 Related work

The domain of fairness-aware machine learning can be categorized into pre-processing, in-processing and post-processing approaches to fairness depending

¹ [git@github.com:wtliao/FairNN.git](https://github.com:wtliao/FairNN.git)

on whether they focus on mitigating discrimination at the data, algorithms or model output, respectively.

Mitigating fairness in supervised learning: *Pre-processing approaches to fairness* assume that there exist encoded (e.g., societal) biases in the data which they try to eliminate before “feeding” the data to some learning algorithm. For example, [14] proposes instance re-weighting, label swapping, and data augmentation to eliminate discrimination in the input data. Similar ideas, but for the online scenario, were proposed by [13]. Data augmentation has also been used in [11] in order to force the model so as to learn efficiently all the population segments. In [10] a bagging schema is proposed to equalize the data distributions for the different population segments. In [2] a probabilistic framework for discrimination-preventing preprocessing in supervised learning is introduced with the goal to preserve the utility of the data for the learning task while controlling the correlation between the protected attributes and class and minimizing instance distortion. *In-processing approaches to fairness* aim to explicitly consider fairness into the learning algorithm by constraining or regularizing the model during the training phase. It comprises the most popular category to fairness mitigation, which however depends on the algorithm per se. For example, in [29] the authors tweak the objective function of the linear SVM and Logistic Regression models by inserting convex-concave fairness-related constraints (they use equalized odds as fairness measure). In [15], a fairness-aware splitting criterion for decision trees is proposed that evaluates not only the splitting quality w.r.t. the class but also the discrimination effect of a potential split. The work is extended in [30] for online learning, using Hoeffding Trees as the underlying model. In [12] the authors aim to eliminate discrimination in sequential learning scenarios (in particular, boosting) by dynamically adapting the data distributions over the training rounds using a cumulative version of equalized odds. In [17] it is assumed that there exist latent fair class labels (non-observable) which are estimated via an iterative process. Finally, *post-processing approaches to fairness* work directly at the output of a model and change its outcomes until a chosen fairness notion is satisfied. For example, [7] shifts the decision boundary of AdaBoost w.r.t a protected attribute until statistical parity is achieved. In [8] different thresholds are introduced for different population segments to enforce equal error rates. In [16] the predictions of probabilistic classifiers and ensemble models for instances close to the decision boundary are altered until statistical parity is fulfilled. Our *FairNN* belongs to the category of in-processing approaches as the objective function of the NN is altered to account for fairness. In contrast to the majority of the previous approaches however, our method comprises a *joint* approach for fair-feature representation- and classifier-learning.

Fair representation learning approaches: Fair representation learning aims to learn a transformation to a lower dimensional space where the protected and non-protected groups are indistinguishable. In [28] the authors propose Fair-PCA, an extension of PCA, that forces similar reconstruction errors between protected and non-protected groups. In [18], the Variational Fair Auto Encoder is proposed that is able to also learn fair non-linear functions, which can be

used after as input to other learning models. Our *FairNN* also derives non-linear transformations via autoencoders, however on the contrary to [18], we don't only focus on fair-representation learning but rather on joint representation-and classifier-learning. In [27] an approach for learning individually fair representations is proposed using an end-to-end model with autoencoders. On the contrary, our *FairNN* aims at learning representations that are fair for each group (i.e., protected and non-protected).

Closer to our work are the joint approaches [6,19] that aim at both fair representation- and classifier-learning. In [6,19] instead of using some constraining to reduce the dependencies on the sensitive attribute in the latent space (e.g., by minimizing KL-divergence as in our *FairNN*), they train an adversary classifier to discriminate between the protected and non-protected groups. In particular, in [6] they optimize for statistical parity, whereas [19] extends the idea for more fairness measures. It is not clear in what circumstances a constraint-based approach or an adversary one should be preferred [6], but we include [19] in our experimental analysis.

3 Basic concepts and definitions

Let $A = \{A_1, \dots, A_d\}$ be a d -dimensional feature space of mixed attribute types. We assume the existence of a protected attribute $S \in A$, e.g., $S = gender$. We assume S is binary: $S = \{s, \bar{s}\}$, with s denoting the protected group (e.g., $s = female$), and \bar{s} the non-protected group e.g., $\bar{s} = male$. An instance $X \in A_1 \times A_2 \cdots \times A_n$ is a d -dimensional feature vector representing an object in the vector space A . Each instance is assigned a label $c \in C$ by some unknown target function $g: A \rightarrow C$. For simplicity, we assume the class attribute is also binary, i.e., $C = \{+, -\}$. We use the notation s_+ (s_-), \bar{s}_+ (\bar{s}_-) to denote the protected and non-protected group for the positive (negative, respectively) class.

The target function $g()$ is unknown, instead a training set $D = \{(X_i, c)\}$ of i.i.d. instances drawn from the joint attribute-class space $A \times C$ is available and can be used for approximating $g()$. The goal of fairness-aware supervised learning is to approximate $g()$ via a mapping function $f()$ that does not only map correctly future unseen instances of the population from A into C , but also mitigates discriminatory outcomes. The former aspect corresponds to the typical objective of supervised learning achieved through empirical risk minimization. The latter aspect is evaluated in terms of some fairness measure (c.f. Section 2).

3.1 Formalizing fairness

In this work, we employ Equalized Odds [8] (shortly *Eq.Odds*) as our fairness measure. *Eq.Odds* accounts for the percentage difference among protected and non-protected groups in the model's outcomes. In particular, let δFPR (δFNR) be the difference in false positive rates (false negative rates, respectively) between

the protected and non-protected groups, defined as follows:

$$\begin{aligned}\delta FPR &= P(c \neq \hat{c}|\bar{s}_-) - P(c \neq \hat{c}|s_-) \\ \delta FNR &= P(c \neq \hat{c}|\bar{s}_+) - P(c \neq \hat{c}|s_+)\end{aligned}\tag{1}$$

where \hat{c} are the predicted labels. The goal of *Eq.Odds* is to minimize both differences:

$$Eq.Odds = |\delta FPR| + |\delta FNR|\tag{2}$$

where *Eq.Odds* $\in [0, 2]$, with 0 indicating no discrimination and 2 indicating maximum discrimination.

Eq.Odds has become quite popular among recent state-of-the-art fairness-aware methods [12,8,29,17,19]. In contrast to the well-known statistical parity [14], which uses only the positive predicted outcomes without the aid of true labels, or equal opportunity [8], which accounts only for the false negative difference among s and \bar{s} , *Eq.Odds* is able to locate discriminatory outcomes for both classes. Furthermore, statistical parity is prone to favor groups by discriminating on specific individuals [4].

3.2 Auto-encoders

An auto-encoder (AE) is an unsupervised neural network that learns an approximation of the identity function such that the output of the network is similar to its input. A reduced/compressed representation is learned by placing constraints in the structure of the network, e.g. by using a bottleneck layer.

In this work, we consider mixed attribute type data of numerical and nominal attributes. Reconstructing the numerical attributes could be considered as a regression task, so we use the *Mean Square Error* as the loss function for numerical attributes. Since for the nominal attributes there is no order among their values, reconstructing their values could be considered as a classification task, so we use the *Cross Entropy* as the loss function for nominal attributes. We assume there exist K numerical and N nominal features, such that $K + N = d$. We combine the feature-type specific loss functions in the overall objective function of the auto-encoder as follows (we compute the loss per batch of B instances):

$$L(X, \hat{X}) = \frac{1}{B} \sum_{b=1}^B \left(\sum_{k=1}^K (X_{b,k} - \hat{X}_{b,k})^2 - \sum_{j=1}^N \sum_{l_j=1}^{M_j} X_{b,l_j} \log(p_{b,l_j}) \right)\tag{3}$$

where X is the original instance, \hat{X} is the reconstructed instance and X_{ji} is the value of instance j in dimension i . The first term of the above equation refers to the loss of numerical attributes: $X_{b,k}$, $\hat{X}_{b,k}$ denotes the original and reconstructed data of numerical attributes, respectively. The second term of the the above equation refers to the loss of nominal attributes. For each nominal attribute j , l_j represents the class label and M_j the number of values of the feature. For the j -th nominal attribute in instance b , X_{b,l_j} has the binary value (positive or negative) which indicates if the class label l_j is the correct classification, p_{b,l_j} represents the predicted probability of class l_j .

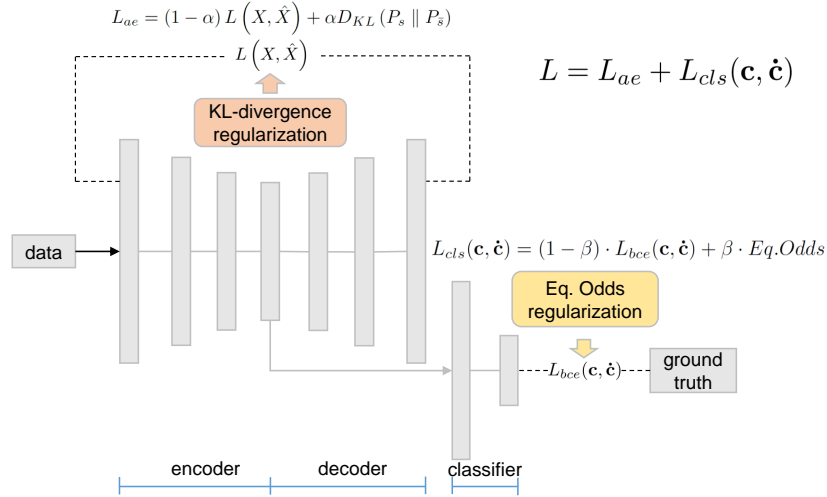


Fig. 1. An overview of *FairNN* that jointly learns a fair representation and a fair mapping function for classification. The auto-encoder (left part) is responsible for representation learning; the KL-divergence constraint forces the representation to be fair. The loss function of the classifier (right part) is tweaked towards fairness through the Eq.Odds regularization. Both aspects are reflected in the joint objective

4 *FairNN*

In this section, we introduce our proposed method, namely *FairNN* that jointly learns a fair representation and a fair mapping function for classification. An overview of our approach is depicted in Fig. 1. The architecture consists of two parts, an auto-encoder block aiming at learning a fair latent representation of the data (left) and a classification block aiming at learning a fair classifier (right). We explicitly consider fairness in the representation learning by adding an additional constraint to the latent space of the auto-encoder in order to obfuscate the information on the protected attribute (Section 4.1). Likewise, we explicitly consider fairness in the classification part by adding an additional constraint to the loss function based on the Equalized Odds fairness notion (Eq. 2) (Section 4.2). We consider these aspects *jointly* and optimize a multi-loss objective function that balances the importance of the different components in training (Section 4.3).

4.1 Fair representation learning via KL-divergence regularization

In order to learn fair feature transformations for the protected and non-protected groups, *KL divergence* is added to the loss function to train the auto-encoder, which constrains the learned features of different groups to have similar distribution properties. With this constraint, the auto-encoder is trained to mix up the

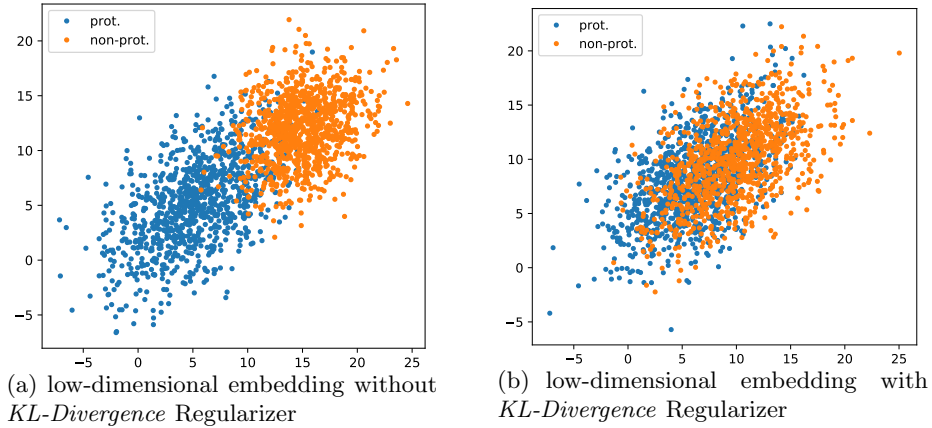


Fig. 2. Effect of the *KL-Divergence* Regularizer in (fair) representation learning

protected attribute information and meanwhile to maintain good reconstruction ability. In practice, we use the *KL divergence* as an additional regularization in the objective function. Based on the values of protected attributes, we divide the data points into protected group s and non-protected group \bar{s} . Without loss of generality, we assume their distribution in the latent space as d -dimensional normal distributions with means $\mu_s, \mu_{\bar{s}}$ and covariance matrices $\Sigma_s, \Sigma_{\bar{s}}$ respectively. Then, the *KL divergence* between the their distributions is given as:

$$D_{KL}(P_s \parallel P_{\bar{s}}) = \frac{1}{2} \left(\log \frac{\det(\Sigma_{\bar{s}})}{\det(\Sigma_s)} - d + \text{tr}(\Sigma_{\bar{s}}^{-1} \Sigma_s) + (\mu_{\bar{s}} - \mu_s)^T \Sigma_{\bar{s}}^{-1} (\mu_{\bar{s}} - \mu_s) \right) \quad (4)$$

where, $\det(\Sigma)$ is the determinant of the covariance matrix Σ , and $\text{tr}(\cdot)$ is the trace of the matrix, which is the sum of elements on the main diagonal of the matrix. With the *KL-Divergence* Regularization, the original reconstruction loss function of the auto-encoder (c.f., Eq. 3) is rewritten as:

$$L_{ae} = (1 - \alpha) L(X, \hat{X}) + \alpha D_{KL}(P_s \parallel P_{\bar{s}}) \quad (5)$$

where $\alpha \in [0, 1)$, is a coefficient for balancing the two terms.

Fig. 2 demonstrates the impact of our KL-divergence regularizer, as distribution of data points in a low-dimensional feature space, in contrast to a transformation that has been learned without KL-Divergence regularization. The protected and non-protected groups are denoted in blue and orange respectively. Fig. 2(a) shows that the data points belonging to different groups are easy to be separated in the latent space with direct implications to fairness. The regularizer mixes-up the distributions of the two groups making it hard to predict the protected attribute, c.f., Fig. 2(b).

4.2 Fair classifier learning via equalized odds regularization

The classifier is an MLP with two FC layers followed by *Relu* activation. The output is a scalar that is squashed by the sigmoid function between 0 and 1 for our binary classification task. The *Binary Cross Entropy* is used as loss function to train the classifier as follows:

$$L_{bce}(\mathbf{c}, \hat{\mathbf{c}}) = -\frac{1}{B} \sum_{n=1}^b ((c_b \log(\hat{c}_b) + (1 - c_b) \log(1 - \hat{c}_b))) \quad (6)$$

where c_b is the true label and \hat{c}_b is the predicted probability of the data point b having the label c_b .

Our goal is to improve the fairness performance without losing the classification performance. This motivates us to add an additional fairness measurement as a regularization term in the objective function. As we mentioned before, among different fairness measurements, *Equalized Odds* does not only consider the predicted outcome but also compares it to the actual outcome recorded in the dataset. It considers both the samples with actual positive labels and also those with negative labels. Therefore, *Equalized Odds* (Eq.Odds) is used as the constraint term and added to the classification loss Eq. (6):

$$L_{cls}(\mathbf{c}, \hat{\mathbf{c}}) = (1 - \beta) \cdot L_{bce}(\mathbf{c}, \hat{\mathbf{c}}) + \beta \cdot Eq.Odds \quad (7)$$

where $\beta \in [0, 1)$, is a balancing coefficient between the classification loss L_{bce} and the *Eq.Odds* fairness regularization.

4.3 Fair Representation and Classifier-learning via Joint Optimization

By combining the two parts of our network, which are the auto-encoder (Eq. 5) and classifier loss (Eq. 7), the acquired multi-loss function can be expressed as:

$$L = L_{ae} + L_{cls}(\mathbf{c}, \hat{\mathbf{c}}). \quad (8)$$

It is known that neural networks can easily be over-parameterized and tend to overfit, given limited training data. The additional constraints in our architecture, together with the auto-encoder component enforces better generalization, as demonstrated in our experiments (Section 5). We implemented *FairNN* in the Python framework using PyTorch.

5 Experiments

We evaluate the predictive and fairness performance of *FairNN* and compare the results with recent state-of-the-art methods. Additionally, we perform several ablation studies to demonstrate the importance of each component in our proposed framework. Accuracy and balanced accuracy are reported for evaluating the *predictive* performance and *Equalized Odds* for *fairness* performance. Since *Equalized Odds* reports the difference between two groups and we also want to maintain the predictive performance for both groups, we also report the actual *TPR* and *TNR* of both groups.

Table 1. An overview of the datasets.

	#Instances	#Attributes	protected attribute	protected group	class ratio (+:-)	positive class
Adult Census	45,175	14	<i>gender</i>	<i>female</i>	1:3.03	<i>>50K</i>
Bank Marketing	40,004	16	<i>marital status</i>	<i>married</i>	1:7.57	<i>yes</i>

5.1 Experimental setup

5.1.1 Datasets We evaluate our method on two real-world datasets, summarized in Tab. 1:


- **Adult Census Income Dataset**[1] is extracted from the 1994 American Census Database. The task is to predict whether a person’s income is over 50K a year. People with label *>50K* belong to the positive class. $S = \textit{gender}$ is considered as the protected attribute, $s = \textit{female}$ the protected group and $\bar{s} = \textit{male}$ the non-protected group.
- **Bank Marketing Dataset**[21] is collected from a Portuguese bank that focuses on selling long-term deposits over the phone. The task is to predict whether a client will make a deposit subscription. We take $S = \textit{marital status}$ as the protected attribute, $s = \textit{married}$ the protected group and $\bar{s} = \textit{single/divorced}$ as the non-protected group.

5.1.2 Experimental settings The nominal attributes are encoded to one-hot vector and max-normalization is applied to the numerical attributes to ensure the values are in $[0, 1]$. In the auto-encoder block, both the encoder and decoder have three fully-connected linear layers and each is followed by a *ReLU* activation. Following the evaluation setup in [29,17,12], 50% of the data is used for training in which 20% of them are used for validation, and the other 50% is for testing. All experiments are evaluated using 10 random splits. We train the auto-encoder and classifier simultaneously by minimizing the objective function Eq. 8. For training, we use the Adam optimization method, with batch size $B = 512$ and a learning rate 0.002. In order to get the best $\alpha - \beta$ combination (see Eq. (5) and (7)), grid search is operated within $\alpha \in [0.4, 0.5, 0.6, 0.7, 0.8, 0.9]$ and $\beta \in [0.1, 0.2, 0.3, 0.4, 0.5]$. Finally, $\alpha = 0.9, \beta = 0.2$ for the Adult Census Income Dataset and $\alpha = 0.8, \beta = 0.4$ for the Bank Marketing Dataset are selected.

To further improve the performance, the well known preferential sampling [14] is applied. The samples are ranked according to their classification scores ascendingly. Centered on classification score 0.5, K samples whose scores > 0.5 are duplicated while K samples whose scores < 0.5 are skipped. K is computed based on the size of sensitive attributes and labels.

5.2 Comparison with other Methods

We compared our approach with the recently proposed state-of-the-art in-processing approaches which mainly aim to minimize *Eq.Odds*.



(a) Comparison on Adult Census Income Dataset. (b) Comparison on Bank Marketing Dataset.

Fig. 3. Comparison with the state-of-the-art methods on Adult Census Income dataset and Bank Marketing dataset. For fairness measurement *Eq.Odds*, lower values are better; For others, higher are better.

AdaFair [12]: a boosting model which assigns fairness related weights in each boosting round by observing the cumulative fairness behavior of the ensemble.

LAFTR [19]: a holistic approach that learns a latent fair representation using an encoder/decoder and an adversary (where the encoder/decoder seek to minimize the adversary’s objective), and at the same time trains a fair classifier on the latent space.

FairPCA-SVM [28]: aims to find a low dimensional representation of the original data while maintaining similar fidelity for two groups. We project the data to the Fair PCA space and use SVM for binary classification.

PCA-SVM: Similar to FairPCA-SVM, we project the data to the PCA space and use an SVM classifier. This is only a naive baseline method for comparison.

EO-Network [20]: A two-layer neural network, with *Eq.Odds* as a constraint in the loss function. This can be seen as our model without the auto-encoder part.

Krasanakis et al. [17]: In this work, the authors assume the existence of a latent fair class distribution, which they approximate through the CULEP model by re-estimating the instance weights iteratively.

Zafar et al. [29]: In this work, the authors formulate fairness as a set of convex-concave constraints which are embedded in the objective function of a logistic regression model.

The experimental results from different methods on two datasets are depicted in Fig. 3, detailed discussion on each dataset follows hereafter. The results of [17,29,12] are taken from [12].

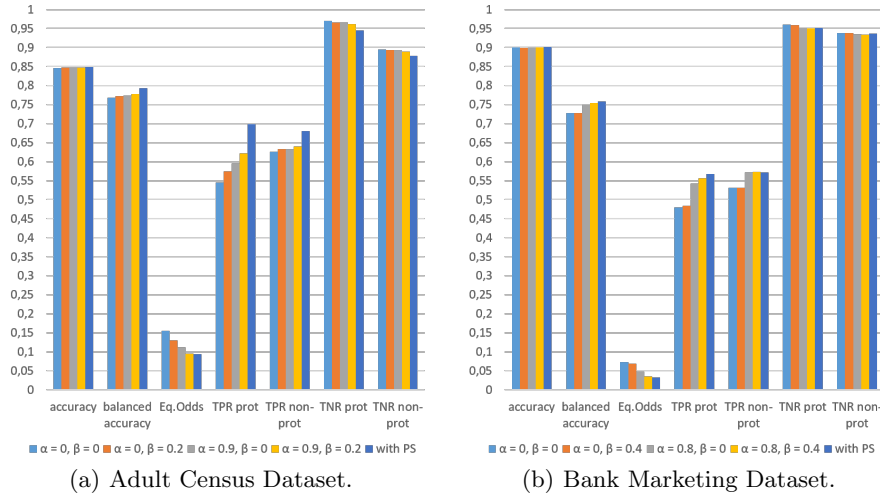


Fig. 4. Ablation Study (for *Eq.Odds* lower values are better - for the rest, higher values are better.)

5.2.1 Adult Census Income Fig. 3(a) displays the baselines, state-of-the-art and our final experimental results on the Adult Census Dataset. Our method achieves the highest accuracy and balanced accuracy rates. The lowest *Eq.Odds* is achieved by Krasanakis et al. However, its TPRs for both *protected*(TPR prot) and *non-protected*(TPR non-prot) groups are much lower than the other methods (the lowest TPR prot and the second-lowest TPR non-prot). Fair-PCA aims to learn a fair feature representation in the low-dimensional space. But the learned representation may be unsuited for the binary classification task. It achieves fairer decision-making (lower *Eq.Odds*) comparing to PCA yet performs worse compared to our method. The comparison of our method with EO-Network demonstrates an 8% decrease in *Eq.Odds* and 14% improvement in TPR prot, revealing the effectiveness of generating low-dimensional features. Similar to our approach, LAFTR also leverages the joint-learning thought, but ours is more effective comparing to theirs: balanced accuracy is 5% higher and *Eq.Odds* 3% lower. Our method also brings a significant increase in TPR prot(18% higher). The superior performance from our method indicates that our method is able to learn the fair representation. It balances the balanced accuracy and *Eq. Odds* well.

5.2.2 Bank In Fig. 3(b), we report experimental results on the Bank Marketing Dataset. Due to the class imbalance problem, both PCA-SVM and FairPCA-SVM perform poorly on this dataset. They output all zeros for the binary classification task which result in balanced accuracy 0.5, TPR prot and TPR non-prot are 0. In EO-Network, the weight parameter of the *Eq.Odds* constraint is the

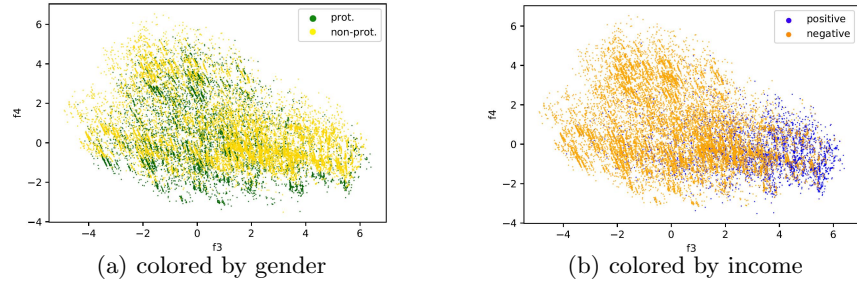


Fig. 5. Visualization of learned features colored by (a) gender and (b) income.

same as used in our method i.e., 0.4. LAFTR reaches the lowest *Eq. Odds* result but its TPRs for both groups are also the lowest. There is a minor difference in *Eq. Odds* between LAFTR and our method, yet ours achieves the much higher balanced accuracy rate, TPR prot, and TPR non-prot. Compared to Zafar et al. and Krasanakis et al., our method reports a higher balanced accuracy rate, higher TPRs for both groups and also the comparable *Eq. Odds*. It proves that our method maintains classification performance while achieving fairness.

5.3 Ablation study

We perform ablation studies to evaluate how different parts influence the predictive and fairness performance of our method. In Fig. 4, $\alpha = 0$ represents the outcome without *KL-Divergence* regularization and $\beta = 0$ without *Eq. Odds* regularization respectively. Fig. 4(a) demonstrates the ablation study on the Adult Census Income Dataset and Fig. 4(b) the Bank Marketing Dataset. We can see that, integrating only the *KL-Divergence* regularization is more effective than integrating *Eq. Odds* regularization only (comparing the second and third bars in Fig. 4(a)). Applying both regularizations further improves the performance (the fourth bars in Fig. 4(a)). Preferential sampling further improves the TPR prot and TPR non-prot while almost not affecting *Eq. Odds*. The ablation study results on the Bank Dataset (as shown in Fig. 4(b)) display a similar tendency yet Preferential sampling does not bring much improvement on TPRs.

5.4 Feature visualization

To better understand what kind of features are learned from the auto-encoder part, we visualize the extracted features by randomly selecting 2 dimensions of the 10 dimensional latent space and color them according to the protected attribute (Fig. 5(a)) and by the label (Fig. 5(b)) respectively. Fig. 5(a) illustrates that the protected attribute information is mixed up in the latent space, which indicates that the fair representation is learned. Fig. 5(b) shows that the label information is distinguishable. The learned representation is not only fair but also suitable for the binary classification task which follows afterwards.

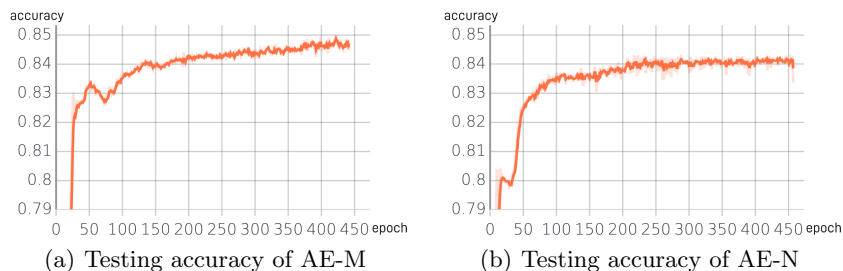


Fig. 6. Comparison of auto-encoder with MSE+Cross-Entropy loss (left) and with normal MSE loss function (right).

5.5 The effect of the multi-loss function on the accuracy

In this experiment we evaluate the effect of the multi-loss function on the accuracy and compare the auto-encoder with MSE loss + Cross Entropy loss (we call the network AE-M), to an auto-encoder with normal MSE loss (AE-N). We set $\alpha = 0, \beta = 0$, which means to ignore the equations 5 and 7.

By observing the testing accuracy shown in Fig. 6(a) and Fig. 6(b), we can conclude that AE-M does not perform worse but even achieves a slightly better predictive performance (testing accuracy is 0.76% higher than AE-N).

6 Conclusion

We proposed *FairNN*, a neural network approach for fairness-aware learning that jointly learns a feature representation and classification model. The neural network consists of two parts, an autoencoder component for fair representation learning and a classification component for fair decision making. Our approach optimizes a multi-objective loss function which (a) learns fair representation by suppressing protected attributes (b) maintains the information content by minimizing the reconstruction loss and (c) allows for fair classification by minimizing the classification error and respecting the equalized odds-based fairness regularizer. Our experiments demonstrate that such a joint approach is superior to a separate treatment of unfairness in representation learning or classifier learning. Our method achieves the highest accuracy and balanced accuracy rates. All components are important as demonstrated by the ablation study.

Note that our architecture contains a branch of an auto-encoder which allows unsupervised learning. Thus, our framework is suited for semi-supervised learning with sparsely labeled data. We will elaborate on this aspect in future works.

Acknowledgements

The work is supported by BIAS (*Bias and Discrimination in Big Data and Algorithmic Processing. Philosophical Assessments, Legal Dimensions, and Technical*

Solutions) a project funded by the Volkswagen Foundation within the initiative *AI and the Society of the Future* for which the last authors are Principal Investigators.

References

1. Bache, K., Lichman, M.: UCI machine learning repository (2013) [9](#)
2. Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K.N., Varshney, K.R.: Optimized pre-processing for discrimination prevention. In: NeurIPS. pp. 3992–4001 (2017) [3](#)
3. Datta, A., Tschantz, M.C., Datta, A.: Automated experiments on ad privacy settings. *Privacy Enhancing Technologies* **2015**(1), 92–112 (2015) [1](#)
4. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.S.: Fairness through awareness. In: *Innovations in Theoretical Computer Science 2012*, Cambridge, MA, USA, January 8–10, 2012. pp. 214–226 (2012) [5](#)
5. Edelman, B.G., Luca, M.: Digital discrimination: The case of airbnb. com (2014) [1](#)
6. Edwards, H., Storkey, A.J.: Censoring representations with an adversary. In: Bengio, Y., LeCun, Y. (eds.) *4th International Conference on Learning Representations, ICLR 2016*, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings (2016) [2](#), [4](#)
7. Fish, B., Kun, J., Lelkes, Á.D.: A confidence-based approach for balancing fairness and accuracy. In: Venkatasubramanian, S.C., Jr., W.M. (eds.) *Proceedings of the 2016 SIAM International Conference on Data Mining*, Miami, Florida, USA, May 5–7, 2016. pp. 144–152. SIAM (2016) [3](#)
8. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016*, December 5–10, 2016, Barcelona, Spain. pp. 3315–3323 (2016) [3](#), [4](#), [5](#)
9. Ingold, D., Soper, S.: Amazon doesn’t consider the race of its customers. should it. Bloomberg, April (2016) [1](#)
10. Iosifidis, V., Fetahu, B., Ntoutsis, E.: Fae: A fairness-aware ensemble framework. In: *2019 IEEE International Conference on Big Data (Big Data)*. pp. 1375–1380. IEEE (2019) [3](#)
11. Iosifidis, V., Ntoutsis, E.: Dealing with bias via data augmentation in supervised learning scenarios. Jo Bates Paul D. Clough Robert Jäschke p. 24 (2018) [3](#)
12. Iosifidis, V., Ntoutsis, E.: Adafair: Cumulative fairness adaptive boosting. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019*, Beijing, China, November 3–7, 2019. pp. 781–790 (2019) [3](#), [5](#), [9](#), [10](#)
13. Iosifidis, V., Tran, T.N.H., Ntoutsis, E.: Fairness-enhancing interventions in stream classification. In: *Database and Expert Systems Applications - 30th International Conference, DEXA 2019*, Linz, Austria, August 26–29, 2019, Proceedings, Part I. pp. 261–276 (2019) [3](#)
14. Kamiran, F., Calders, T.: Data preprocessing techniques for classification without discrimination. *Knowl. Inf. Syst.* **33**(1), 1–33 (2011) [3](#), [5](#), [9](#)
15. Kamiran, F., Calders, T., Pechenizkiy, M.: Discrimination aware decision tree learning. In: *ICDM*. pp. 869–874. IEEE Computer Society (2010) [3](#)

16. Kamiran, F., Mansha, S., Karim, A., Zhang, X.: Exploiting reject option in classification for social discrimination control. *Inf. Sci.* **425**, 18–33 (2018) [3](#)
17. Krasanakis, E., Xioufis, E.S., Papadopoulos, S., Kompatsiaris, Y.: Adaptive sensitive reweighting to mitigate bias in fairness-aware classification. In: Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23–27, 2018. pp. 853–862 (2018) [3](#), [5](#), [9](#), [10](#)
18. Louizos, C., Swersky, K., Li, Y., Welling, M., Zemel, R.S.: The variational fair autoencoder. In: Bengio, Y., LeCun, Y. (eds.) 4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings (2016) [3](#), [4](#)
19. Madras, D., Creager, E., Pitassi, T., Zemel, R.: Learning adversarially fair and transferable representations. arXiv preprint arXiv:1802.06309 (2018) [2](#), [4](#), [5](#), [10](#)
20. Manisha, P., Gujar, S.: A neural network framework for fair classifier. arXiv preprint arXiv:1811.00247 (2018) [10](#)
21. Moro, S., Cortez, P., Rita, P.: A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems* **62**, 22–31 (2014) [9](#)
22. Navarin, N., Oneto, L., Donini, M.: Learning deep fair graph neural networks (2020) [2](#)
23. Ntoutsi, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejd, W., Vidal, M.E., Ruggieri, S., Turini, F., Papadopoulos, S., Krasanakis, E., Kompatsiaris, I., Kinder-Kurlanda, K., Wagner, C., Karimi, F., Fernandez, M., Alani, H., Berendt, B., Kruegel, T., Heinze, C., Broelemann, K., Kasneci, G., Tiropanis, T., Staab, S.: Bias in data-driven artificial intelligence systems - an introductory survey. *WIREs Data Mining and Knowledge Discovery* (2020) [1](#)
24. Oneto, L., Donini, M., Maurer, A., Pontil, M.: Learning fair and transferable representations. arXiv preprint arXiv:1906.10673 (2019) [2](#)
25. of the President, U.S.E.O., Podesta, J.: Big data: Seizing opportunities, preserving values. White House, Executive Office of the President (2014) [1](#)
26. Rudolph, M., Wandt, B., Rosenhahn, B.: Structuring autoencoders. In: Third International Workshop on “Robust Subspace Learning and Applications in Computer Vision” (ICCV) (Aug 2019) [2](#)
27. Ruoss, A., Balunovic, M., Fischer, M., Vechev, M.T.: Learning certified individually fair representations. *CoRR* **abs/2002.10312** (2020) [4](#)
28. Samadi, S., Tantipongpipat, U.T., Morgenstern, J.H., Singh, M., Vempala, S.S.: The price of fair PCA: one extra dimension. In: Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3–8 December 2018, Montréal, Canada. pp. 10999–11010 (2018) [2](#), [3](#), [10](#)
29. Zafar, M.B., Valera, I., Gomez-Rodriguez, M., Gummadi, K.P.: Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In: Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3–7, 2017. pp. 1171–1180 (2017) [3](#), [5](#), [9](#), [10](#)
30. Zhang, W., Ntoutsi, E.: FAHT: an adaptive fairness-aware decision tree classifier. In: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10–16, 2019. pp. 1480–1486 (2019) [3](#)