

# Enhancing Fairness of Visual Attribute Predictors<sup>\*</sup>

Tobias Hänel<sup>1</sup>, Nishant Kumar<sup>1</sup>, Dmitrij Schlesinger<sup>1</sup>, Mengze Li<sup>2</sup>, Erdem Ünal<sup>1</sup>, Abouzar Eslami<sup>2</sup>, and Stefan Gumhold<sup>1</sup>

<sup>1</sup>Chair for Computer Graphics and Visualization, TU Dresden, Germany

<sup>2</sup>Carl Zeiss Meditec AG, Munich, Germany

Corresponding author: [tobias.haenel@tu-dresden.de](mailto:tobias.haenel@tu-dresden.de)

**Abstract.** The performance of deep neural networks for image recognition tasks such as predicting a smiling face is known to degrade with under-represented classes of sensitive attributes. We address this problem by introducing fairness-aware regularization losses based on batch estimates of Demographic Parity, Equalized Odds, and a novel Intersection-over-Union measure. The experiments performed on facial and medical images from CelebA, UTKFace, and the SIIM-ISIC Melanoma classification challenge show the effectiveness of our proposed fairness losses on bias mitigation as they improve model fairness while maintaining high classification performance. To the best of our knowledge, our work is the first attempt to incorporate these types of losses in an end-to-end training scheme to mitigate biases of visual attribute predictors.

**Keywords:** Algorithmic Fairness · Fair Image Classification · Deep Neural Networks · Visual Attributes · Facial Recognition · Disease Diagnosis

## 1 Introduction

The manifestation of bias is evident in every aspect of our society, from educational institutions [15], to bank credit limits for women [30], to criminal justice for people of color [13]. The core problem is the inability of an individual to make ethically correct objective decisions without being affected by personal opinions. With the advent of recent machine learning (ML) algorithms trained on big data, there is a dramatic shift towards using such algorithms to provide greater discipline to impartial decision-making. However, ML-based algorithms are also prone to making biased decisions [4,27], as the reliability of data-based decision-making is heavily dependent on the data itself. For instance, such models are unfair when the training data is heavily imbalanced towards a particular class of a sensitive attribute such as race [31]. A notable example by [4] shows that by assuming the ML model's target attribute as gender and the sensitive attribute as skin color, the classification error rate is much higher for darker females than for lighter females. A similar concern exists in the medical fraternity, where

---

<sup>\*</sup> Source code is available at <https://github.com/nish03/FVAP>

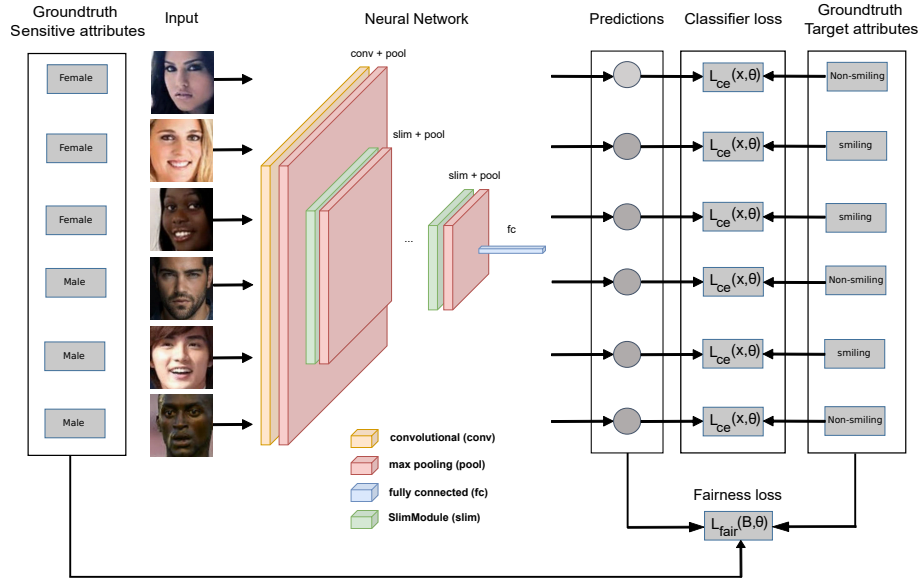
a recent work [22] studied the correlation between the under-representation of darker skin images with the classification error in predicting dermatological disease. Another work [23] showed a sharp decrease in classification performance while diagnosing several types of thoracic diseases for an under-represented gender in the X-ray image training data set. Therefore, it is critical to mitigate these biases in ML-based models for visual recognition tasks, to alleviate ethical concerns while deploying such models in real-world applications.

Recent studies focus on creating balanced data sets [19], or perform data augmentation [36] to remove imbalance with respect to the sensitive attributes. We argue that to make an ML model useful in real-life scenarios, it should achieve algorithmic fairness while still being trained on data sets that consist of real-world biases. In terms of algorithmic fairness, works such as [24,32,45] aim to learn the features in the data that are statistically independent of the sensitive attributes, while [18,2] focus on de-biasing the latent space of a generative model to achieve a fair outcome. We propose that to reduce bias w.r.t. the sensitive attributes, a model must satisfy the fairness notations by learning them during training. Additionally, none of the previous approaches attempted to utilize an inherent IoU-based fairness measure to train an ML model and achieve algorithmic fairness without loss in classification accuracy. Our overall learning scheme is presented in Fig. 1. The contributions of our work are as follows:

- We use classical fairness notations such as Demographic Parity (DP) and Equalized Odds (EO) to define the corresponding fairness loss terms and measure the deviations from the assumptions of the probabilistic independence w.r.t. sensitive attributes. We quantify these deviations by using mean squared error as well as the Kullback-Leibler divergence (KLD) between the learned probability distribution and the best-factorized distribution, leading to the mutual information (MI) between the variables in the learned model.
- We generalize the fairness notations such as DP and EO for categorical variables since the task is usually a multi-class problem in image classification. In the past, such notations were defined for binary variables only.
- We introduce a novel fairness loss based on the Intersection-over-Union (IoU) measure and study its relevancy for achieving fair classification results empirically. Our experiments show that it can simultaneously improve the model fairness and the baseline classification performance when the model is evaluated with fairness metrics.
- We exhaustively evaluate all introduced losses with facial attribute prediction on CelebA [25], age group estimation on UTKFace [49], and disease classification on the SIIM-ISIC Melanoma data set [36]. It was possible for all of these data sets to improve the model fairness with our method.

## 2 Related Work

We discuss current methods that deal with bias mitigation in the data domain and provide an overview of works conducted to achieve fairness-aware facial and medical image recognition systems.



**Fig. 1.** The figure shows our new training procedure that improves the fairness of image classification models w.r.t. to sensitive attributes such as gender, age, and ethnicity. We add a weighted fairness loss to the standard cross-entropy loss during mini-batch gradient descent. It estimates the fairness of the model’s predictions based on the sensitive attributes and the predicted and ground-truth target attributes from the samples within each batch.

**Mitigating bias in the data domain:** The work [33] developed an audit process to highlight the fairness-driven concerns in facial recognition while [45] provided a benchmark for bias removal and highlighted key ethical questions for mitigating bias in vision data sets. Another study [22] estimated skin tones for images with skin disease and showed that the darker skin population is under-represented in the data set. In [11], the authors performed data augmentation to mitigate bias in facial image data sets. Multiple authors [19,28] presented new facial data sets with a balanced number of samples from different classes of sensitive attributes. The works such as [26,7,8] used image generation models to synthesize new facial images and used these images with standard pre-trained classifiers to investigate algorithmic performance bias based on sensitive attributes. Some authors [35,2] used the latent space of the Generative Adversarial Networks (GAN) to de-correlate the target and sensitive attributes, while [42] developed a tool that uses an image data set and its annotations as input and suggests actions to mitigate potential biases. The work [32] learned a GAN-based mapping of the source image to a target fair image representation by removing the semantics of target features such as eyes and lips from the sensitive attribute gender. The issue of unfair predictions due to biased data also exists in few-shot [50] and zero-shot [40] learning where [50] includes a constraint called decision

boundary covariance, enabling independence of target and sensitive attributes while [40] maps the unlabelled target classes to a point in semantic space exterior to the point cluster of labeled source classes.

**Fairness in facial recognition:** A study [37] proposed a classifier that predicts target facial attributes using pre-learned sensitive attributes. Another work [51] calibrated the predicted target class labels to reduce the performance bias while [10] proposed separate fair classifiers for each target class. In [1], they used the cross-entropy loss between the predicted sensitive labels and the uniform distribution, causing the model to be confused and invariant to the sensitive attributes. Many authors [29,9,18,21,14,44] proposed adversarial learning to obtain a fair ML model. A study [12] presented channel-wise attention maps that adapt to diversity in facial features in every demographic group to obtain a fair classifier. A publication [43] used Q-learning to train an agent by rewarding less skew of the inter-class distance between the sensitive attribute such as race. In [17], they used a simple encoder-decoder network, [24,32] used adversarial learning of GAN, and [48] used a style transfer GAN to synthesize facial images independent of sensitive attributes such as gender and race. A work [16] ensured that the generated images from an image generation model have the same error rates for each class of the sensitive attributes. The authors of [47] introduced the False Positive Rate (FPR) as the penalty loss to mitigate bias. The authors in [46] showed that adversarial learning might worsen classification accuracy and fairness performance. They suggested a constraint that standard classification accuracy and fairness measures should be limited to the average of both metrics.

**Fairness in disease diagnosis:** In [23], they showed the problem of gender imbalance in medical imaging data sets, where the classification performance is less for disease diagnosis of the under-represented gender. A study [38] found that True Positive Rate (TPR) disparity for sensitive attributes such as gender, age, and race exists for all classifiers trained on different data sets of chest X-ray images. Another publication [5] highlighted some recent works to mitigate bias via federated learning (FL) [3,20,34]. FL enables multiple clients stationed at different geographical locations to train a single ML model with diversified data sets collaboratively. This method should help overcome the under-representation of individual classes of sensitive attributes in the data, resulting in unbiased models. However, data heterogeneity among the distributed clients remains a challenge for such FL-based models, which might degrade the performance.

In summary, past works did not thoroughly explore the theoretical insights into fairness-based regularization measures. They also did not apply them in experiments on visual data to obtain a model that has high classification performance on target attributes and is unbiased w.r.t. sensitive attributes.

### 3 Method

Let  $x \in \mathcal{X}$  denote an image and  $T = (x_1, x_2, \dots, x_{|T|})$  be the training data consisting of  $|T|$  images. Let  $y_s \in \{1 \dots K_s\}$  and  $y_t \in \{1 \dots K_t\}$  be *sensitive* and

*target* attributes respectively. The latter is our classification target, whereas the former is the one, classification should not depend on. The asterisk will be used to denote the attribute ground truth values, e.g.,  $y_s^*(x)$  means the ground-truth value of the sensitive attribute for an  $x$  from  $T$ . We treat a trainable *classifier* as a conditional probability distribution  $p_\theta(y_t|x)$  parameterized by an unknown parameter  $\theta$  to be learned. In our work, we use Feed Forward networks (FFNs) as classifiers, which means  $\theta$  summarizes network weights.

Next, we consider the joint probability distribution  $p(x, y_t) = p(x) \cdot p_\theta(y_t|x)$ , where  $p(x)$  is some probability distribution, from which training data  $T$  is drawn. Noteworthy, the ground truths can also be considered as random variables, since they are deterministic mappings from  $\mathcal{X}$ , i.e., there exists a unique value  $y_t^*(x)$  as well as a unique value  $y_s^*(x)$  for each  $x \in \mathcal{X}$ . Hence, for the ground truths, we can also consider their statistical properties like joint probability distribution  $p(y_s^*, y_t)$ , independence  $y_t \perp y_s^*$ , or similar.

**Demographic Parity.** The notation reads  $y_t \perp y_s^*$ , i.e. the prediction should not depend on the sensitive attribute. Traditionally (for binary variables) a classifier is said to satisfy demographic parity if

$$p(y_t = 1|y_s^* = 1) = p(y_t = 1) \tag{1}$$

holds. A straightforward generalization to the case of categorical variables is to require the same for all possible values, i.e.

$$p(y_t = a|y_s^* = b) = p(y_t = a), \tag{2}$$

where  $a \in \{1 \dots K_t\}$  and  $b \in \{1 \dots K_s\}$ . The above notation can be used to define a *loss* function, i.e. measure that penalizes the deviation of a given probability distribution  $p_\theta$  from satisfying (2). One possible way is to penalize the sum of squared differences

$$L_{dp}^{l_2}(\theta) = \sum_{ab} [p_\theta(y_t = a|y_s^* = b) - p_\theta(y_t = a)]^2. \tag{3}$$

In fact, during the transition from (2) to (3) we compare the actual joint probability distribution  $p_\theta(y_t, y_s^*)$  to the corresponding factorized (i.e. independent) probability distribution  $p_\theta(y_t) \cdot p_\theta(y_s^*)$  using squared difference, i.e., interpreting probability distributions as vectors to some extent<sup>1</sup>. The squared difference is however not the only way to compare probability distributions. Another option would be e.g., Kullback-Leibler divergence  $D_{KL}(p_\theta(y_t, y_s^*)||p_\theta(y_t) \cdot p_\theta(y_s^*))$  which leads to the mutual information loss

$$\begin{aligned} L_{dp}^{mi}(\theta) &= \sum_{ab} p_\theta(y_t = a, y_s^* = b) \cdot \log \frac{p_\theta(y_t = a, y_s^* = b)}{p_\theta(y_t = a) \cdot p_\theta(y_s^* = b)} = \\ &= H(y_t) + H(y_s^*) - H(y_t, y_s^*), \end{aligned} \tag{4}$$

<sup>1</sup> Strictly speaking, it directly holds only if  $p_\theta(y_s^*)$  is uniform, otherwise (3) corresponds to a squared difference between  $p_\theta(y_t, y_s^*)$  and  $p_\theta(y_t) \cdot p_\theta(y_s^*)$ , where addends are additionally weighted by  $1/p_\theta(y_s^*)^2$ .

where  $H(\cdot)$  denotes the entropy. Note that we derived different losses (3) and (4) from the same independence requirements (2) using different distance measures for probability distributions.

**Equalized Odds.** It is assumed that the predicted target attribute and the ground truth sensitive attribute are *conditionally* independent given a fixed value of the ground truth target, i.e.  $(y_t \perp y_s^*)|y_t^*$ . Hence

$$p(y_t = a|y_t^* = b, y_s^* = c) = p(y_t = a|y_t^* = b) \quad (5)$$

should hold for all triples  $a, b \in \{1 \dots K_t\}$  and  $c \in \{1 \dots K_s\}$ . Again, similarly to the previous case, we consider first the simple quadratic loss

$$L_{eo}^{l_2}(\theta) = \sum_{abc} [p(y_t = a|y_t^* = b, y_s^* = c) - p(y_t = a|y_t^* = b)]^2. \quad (6)$$

It is also possible to measure the deviations of the current model  $p_\theta(y_t, y_t^*, y_s^*)$  from the requirements (5) utilizing the corresponding mutual information as

$$L_{eo}^{mi}(\theta) = \sum_a [H(y_t|y_t^* = a) + H(y_s^*|y_t^* = a) - H(y_t, y_s^*|y_t^* = a)]. \quad (7)$$

**Intersection-Over-Union.** The above losses have a distinct statistical background since they rely on specific independence assumptions. In practice, however, we are often not interested in making some variables independent. Instead, the general goal could be phrased as “the classification performance should be similar for different values of the sensitive attribute”. Hence, the core question is how to measure classifier performance adequately. We argue for the IoU measure because it can appropriately rate performance, especially for unbalanced data. For a target value  $a$ , the corresponding IoU is traditionally defined as

$$\text{IoU}_\theta(a) = \frac{p_\theta(y_t = a \wedge y_t^* = a)}{p_\theta(y_t = a \vee y_t^* = a)}, \quad (8)$$

where  $\wedge$  and  $\vee$  denote logical “and” and “or” respectively. The overall IoU is usually defined by averaging (8) over  $a$ , i.e.

$$\text{IoU}_\theta = \frac{1}{K_t} \sum_a \text{IoU}_\theta(a). \quad (9)$$

For a model  $p_\theta(y_t, y_t^*, y_s^*)$  with target value  $a$  and sensitive value  $b$ , we extend (8) and define  $\text{IoU}_\theta(a, b)$  as

$$\text{IoU}_\theta(a, b) = \frac{p_\theta((y_t = a \wedge y_t^* = a) \wedge y_s^* = b)}{p_\theta((y_t = a \vee y_t^* = a) \wedge y_s^* = b)}. \quad (10)$$

Now given (9),  $\text{IoU}_\theta(b)$  for a particular value  $b$  of the sensitive attribute is

$$\text{IoU}_\theta(b) = \frac{1}{K_t} \sum_a \text{IoU}_\theta(a, b) \quad (11)$$

and the loss penalizes the deviations of these IoU-s from the overall IoU (9) as:

$$L_{iou}(\theta) = \sum_b \left[ \text{IoU}_\theta(b) - \text{IoU}_\theta \right]^2. \quad (12)$$

Note that to define (8) to (12), we again used statistical interpretation of all involved variables, i.e., the joint probability distribution  $p_\theta(y_t, y_t^*, y_s^*)$ . This time however we do not explicitly enforce any independence in contrast to the demographic parity or equalized odds.

**Optimization.** All introduced losses are differentiable w.r.t. unknown parameter  $\theta^2$  because we use probability values for their computation. Hence, we can mix them with other losses simply and conveniently. In particular, if the classifier is e.g., an FFN, we can optimize them using error back-propagation. Second, it should be noted that all losses rely on low-order statistics, i.e., it is only necessary to estimate current  $p_\theta(y_t, y_t^*, y_s^*)$  to compute them. For example, if the involved variables are binary, we only need to estimate 8 values. We assume that they can be reliably estimated from a data mini-batch of a reasonable size instead of computing them over the whole training data. This makes optimizing the introduced losses within commonly used stochastic optimization frameworks possible. To be more specific, the overall loss can be written as

$$\begin{aligned} L(\theta) &= \sum_{x \in T} L_{ce}(x, \theta) + \lambda \cdot L_{fair}(T, \theta) = \\ &= \mathbb{E}_{B \subset T} \sum_{x \in B} L_{ce}(x, \theta) + \lambda \cdot L_{fair}(T, \theta), \end{aligned} \quad (13)$$

where the expectation  $\mathbb{E}$  is over all mini-batches  $B \subset T$  randomly sampled from the training data,  $L_{ce}$  is a “usual” classification loss, e.g., Cross-Entropy,  $\lambda$  is a weighting coefficient, and  $L_{fair}$  is one of the fairness losses introduced above. For the sake of technical convenience, we approximate (13) by

$$L(\theta) = \mathbb{E}_{B \subset T} \left[ \sum_{x \in B} L_{ce}(x, \theta) + \lambda \cdot L_{fair}(B, \theta) \right]. \quad (14)$$

**Impact on the Performance.** We want to show a crucial difference between demographic parity and equalized odds. Imagine a hypothetical “perfect classifier” that always assigns probability 1 to the ground truth label. Hence, the requirements of demographic parity  $y_t \perp y_s^*$  turn into  $y_t^* \perp y_s^*$ . It means that the perfect classifier can satisfy demographic parity only if the ground truth target and the ground truth sensitive attributes are completely uncorrelated, which is hard to expect in practice. It follows from the practical perspective that the classifier performance should decrease when we try to make the classifier fair in the sense of demographic parity.

<sup>2</sup> We also assume that probabilities are differentiable w.r.t. parameters.

On the other side, in our notations, the case of a perfect classifier can be written as

$$p_\theta(y_t, y_t^*, y_s^*) = p_\theta(y_t|y_t^*, y_s^*) \cdot p(y_t^*, y_s^*) = \llbracket y_t = y_t^* \rrbracket \cdot p(y_t^*, y_s^*), \quad (15)$$

where  $\llbracket \cdot \rrbracket$  is 1 if its argument is true. Hence, without loss of generality

$$p_\theta(y_t, y_t^*, y_s^*) = \llbracket y_t = y_t^* \rrbracket \cdot p(y_t^*, y_s^*) = p_\theta(y_t|y_t^*) \cdot p(y_s^*|y_t^*) \cdot p(y_t^*). \quad (16)$$

It means that for a perfect classifier, the predicted target attribute and ground-truth sensitive attribute are conditionally independent, i.e., a perfect classifier automatically satisfies the requirements of equalized odds  $(y_t \perp y_s^*)|y_t^*$ . In practice, if the baseline classifier is already good enough, its performance should not worsen when we try to make the classifier fair w.r.t. equalized odds.

Considering the IoU-loss (12), it is easy to see that it is zero for a perfect classifier, just because all IoU values are equal to one in this case. Hence, as in the case of equalized odds, we do not expect a drop in classifier performance when we try to make it fair w.r.t. the IoU-loss.

**Fairness and Calibration Properties.** We consider a *linear squeezing* operation as follows. Let  $y \in \{0, 1\}$  be a binary variable<sup>3</sup> and  $p(y|x)$  a conditional probability distribution for an input  $x$ . The linear squeezing puts all probability values into the range  $[0.5 - \alpha/2, 0.5 + \alpha/2]$  with  $0 < \alpha < 1$  by applying

$$p'(y|x) = [p(y|x) - 0.5] \cdot \alpha + 0.5. \quad (17)$$

Firstly, this operation does not change the *decision* about  $y$  for a given  $x$ . The decision is made by thresholding  $p(y|x)$  at the 0.5 level, which does not change after applying the linear squeezing. Secondly, it makes the classifier “less confident” about its decision because the output probabilities lie in a narrower range. At the same time, linear squeezing can be understood as mixing the original  $p(y|x)$  and uniform distribution, since (17) can be rewritten as

$$p'(y|x) = p(y|x) \cdot \alpha + 0.5 \cdot (1 - \alpha). \quad (18)$$

Consider now the confusion matrix, i.e.  $p(y_t, y_t^*)$ , obtained by averaging over the training data, and the corresponding IoU-value (9) (for now we do not consider the sensitive attribute). Let us assume evenly distributed ground truth labels for simplicity. So applying (18) to the output probabilities gives

$$p'(y_t, y_t^*) = p(y_t, y_t^*) \cdot \alpha + 0.25 \cdot (1 - \alpha). \quad (19)$$

For  $\alpha$  close to 1, the IoU obtained from  $p'(y_t, y_t^*)$  (i.e., squeezed by (19)) will be close to the original IoU (i.e. obtained from the original non-squeezed  $p(y_t, y_t^*)$ ) for which we assume a rather high value since the classifier is essentially better

<sup>3</sup> We discuss in detail only the case of binary variables and the IoU-loss for simplicity. The situation is similar for other cases.



than random chance. For  $\alpha$  close to zero, the modified IoU converges to  $1/3$ . Hence, IoU differences (i.e. addends in (12)) vanish. To conclude, we can make the IoU-loss (12) alone arbitrarily small just by applying linear squeezing without changing the decision rule.

Note that the squeezing operation with a small  $\alpha$  makes the primary loss, i.e., the cross-entropy, essentially worse since the log-likelihoods of the ground truth labels get smaller. In fact, the fairest classifier is a random choice decision, i.e., which does not depend on the input. It is fair and under-confident but poor in terms of the primary goal, i.e. classification accuracy, and w.r.t. the primary cross-entropy loss. To conclude, adding the IoU-loss to the primary objective (see (14)) pushes the solution towards being less confident. It may be a desired or an undesired behavior depending on whether the baseline classifier is already well-calibrated or not. For over-confident baseline classifiers, employing the IoU-loss should improve calibration properties. The calibration properties may get worse for already well-calibrated or under-confident baseline classifiers.

## 4 Experiments

We validate our contributions on three data sets. The first study of interest in 4.1 concerns the CelebFaces attributes (CelebA) data set [25] which contains more than 200K images of celebrities and manually annotated facial attributes. Secondly in 4.2, we investigate the UTKFace facial image data set [49] which contains over 20K images from a broad range of age groups. In addition to the facial images, we perform experiments with a data set from the SIIM-ISIC Melanoma classification challenge [36] in 4.3 that contains 33k+ skin lesion images. We focus on achieving a balanced target attribute prediction performance that does not depend on the sensitive attribute. We split each data set into a train, validation, and test partition to verify the results of our method. First, we train a baseline model (details in the supplementary material). To improve its fairness, we continue the optimization process by extending the cross-entropy loss with one of the weighted fair losses (see 14) and perform experiments with two different strategies for selecting  $\lambda$  in 4.4.

### 4.1 CelebA

For experiments with the CelebA data set, we use SlimCNN [39], a memory-efficient convolutional neural network, to predict whether a depicted person is smiling or not. To evaluate how our method influences the fairness of this prediction task, we select the binary variables *Male* and *Young* (representing gender and age) as sensitive attributes  $y_s$ . We use the original train, validation, and test partitions in all experiments with this data set.

The results for the experiments with the best  $\lambda$  values are shown in Table 1 (details in the supplementary material). Each row shows the results from an experiment with a particular training loss. The columns list the corresponding prediction accuracy (*Acc*) and all fairness metrics on the validation partition.

**Table 1.** Results of the experiments with manually selected weighing coefficients  $\lambda$  for CelebA facial attribute prediction. The task is to predict the binary target attribute  $y_t = Smiling$ . The experiments #1 to #6 use the sensitive attribute  $y_s = Male$  for the evaluation of the fairness loss terms, while #7 to #12 use the sensitive attribute  $y_s = Young$ . The values in bold are the best results of each evaluation metric.

#	Loss	Acc	$L_{iou}$ (12)	$L_{eo}^{l_2}$ (6)	$L_{eo}^{mi}$ (7)	$L_{dp}^{l_2}$ (3)	$L_{dp}^{mi}$ (4)
1	$L_{ce}$	0.902	$8.73 \times 10^{-4}$	$4.89 \times 10^{-3}$	$5.12 \times 10^{-3}$	$1.77 \times 10^{-2}$	$8.46 \times 10^{-3}$
2	$L_{iou}$	<b>0.903</b>	$7.32 \times 10^{-5}$	$8.59 \times 10^{-4}$	$4.26 \times 10^{-4}$	$2.51 \times 10^{-3}$	$1.20 \times 10^{-3}$
3	$L_{eo}^{l_2}$	0.902	<b><math>1.35 \times 10^{-5}</math></b>	<b><math>1.78 \times 10^{-4}</math></b>	<b><math>7.71 \times 10^{-5}</math></b>	<b><math>1.36 \times 10^{-4}</math></b>	<b><math>6.45 \times 10^{-5}</math></b>
4	$L_{eo}^{mi}$	0.899	$2.37 \times 10^{-5}$	$2.24 \times 10^{-4}$	$1.03 \times 10^{-4}$	$8.40 \times 10^{-4}$	$4.00 \times 10^{-4}$
5	$L_{dp}^{l_2}$	0.899	$4.28 \times 10^{-4}$	$3.75 \times 10^{-3}$	$1.87 \times 10^{-3}$	$1.57 \times 10^{-4}$	$7.43 \times 10^{-5}$
6	$L_{dp}^{mi}$	0.901	$5.28 \times 10^{-4}$	$7.73 \times 10^{-3}$	$3.96 \times 10^{-3}$	$7.15 \times 10^{-4}$	$3.40 \times 10^{-4}$
7	$L_{ce}$	0.901	$1.34 \times 10^{-3}$	$8.28 \times 10^{-3}$	$4.93 \times 10^{-3}$	$1.16 \times 10^{-2}$	$3.48 \times 10^{-3}$
8	$L_{iou}$	0.901	$4.15 \times 10^{-5}$	$9.96 \times 10^{-4}$	$3.06 \times 10^{-4}$	$1.31 \times 10^{-3}$	$3.96 \times 10^{-4}$
9	$L_{eo}^{l_2}$	<b>0.902</b>	<b><math>1.01 \times 10^{-5}</math></b>	<b><math>5.48 \times 10^{-5}</math></b>	<b><math>1.65 \times 10^{-5}</math></b>	<b><math>6.81 \times 10^{-5}</math></b>	<b><math>2.08 \times 10^{-5}</math></b>
10	$L_{eo}^{mi}$	0.901	$1.64 \times 10^{-5}$	$2.57 \times 10^{-4}$	$7.40 \times 10^{-5}$	$4.24 \times 10^{-4}$	$1.29 \times 10^{-4}$
11	$L_{dp}^{l_2}$	0.901	$8.34 \times 10^{-4}$	$6.93 \times 10^{-3}$	$2.18 \times 10^{-3}$	$1.58 \times 10^{-4}$	$4.78 \times 10^{-5}$
12	$L_{dp}^{mi}$	0.901	$5.63 \times 10^{-4}$	$4.26 \times 10^{-3}$	$1.33 \times 10^{-3}$	$1.45 \times 10^{-4}$	$4.44 \times 10^{-5}$

The model fairness improved for all experiments with this data set according to almost all proposed losses. The application of the fairness losses  $L_{iou}$  and the  $L_{eo}^{l_2}$  did not lead to a reduction in the prediction accuracy. The  $L_{eo}^{l_2}$  loss yielded the best fairness improvements according to all proposed fairness losses, while the  $L_{iou}$  loss could even improve the classification performance. However, the model training with other losses slightly decreased the classification accuracy. Furthermore, the  $L_{eo}^{mi}$  loss could improve model fairness w.r.t. to all evaluated metrics, while the DP-based  $L_{dp}^{l_2}$  and  $L_{dp}^{mi}$  losses could only improve their respective fairness losses. Next, applying any fairness loss did not deteriorate the model performance with the sensitive attribute  $y_s = Young$ .

## 4.2 UTKFace

The images in UTKFace have annotations of a binary gender variable (Female, Male), a multi-class categorical ethnicity variable (White, Black, Asian, Indian, and Others), and an integer age variable (0-116 Years). Commonly this data set is used to perform age regression. We derive a categorical age group variable (under 31 Years, between 31-60 Years, over 60 Years) from the original ages as our predicted target attribute.  $y_s = Ethnicity$  and  $y_s = Gender$  represent the sensitive attributes in the experiments. We quantify the performance of the trained model by the accuracy based on data from the validation partition. Preliminary experiments with SlimCNN [39] didn't produce satisfying accuracies. EfficientNet is an alternative convolutional network [41] that can scale the depth, width, and resolution of all filters with a single parameter (we use EfficientNet-B1). Since UTKFace does not have any partitioning information, we split the data

**Table 2.** Quantitative outcomes with manually selected weighting coefficients  $\lambda$  for predicting the multi-class target attribute  $y_t = \textit{Agegroup}$  on the UTKFace facial image data set. Experiments #1 to #6 concern the sensitive attribute  $y_s = \textit{Gender}$  and #7 to #12 cover the sensitive attribute  $y_s = \textit{Ethnicity}$ .

#	Loss	Acc	$L_{iou}$ (12)	$L_{eo}^{l_2}$ (6)	$L_{eo}^{mi}$ (7)	$L_{dp}^{l_2}$ (3)	$L_{dp}^{mi}$ (4)
1	$L_{ce}$	0.847	$1.45 \times 10^{-3}$	$1.66 \times 10^{-1}$	$1.23 \times 10^{-1}$	$6.55 \times 10^{-2}$	$4.21 \times 10^{-2}$
2	$L_{iou}$	0.852	<b><math>3.04 \times 10^{-4}</math></b>	$3.00 \times 10^{-2}$	$2.89 \times 10^{-2}$	$3.77 \times 10^{-2}$	$2.18 \times 10^{-2}$
3	$L_{eo}^{l_2}$	0.856	$3.08 \times 10^{-3}$	$8.11 \times 10^{-2}$	$7.21 \times 10^{-2}$	$4.85 \times 10^{-2}$	$3.10 \times 10^{-2}$
4	$L_{eo}^{mi}$	<b>0.857</b>	$9.82 \times 10^{-4}$	<b><math>2.88 \times 10^{-2}</math></b>	<b><math>2.25 \times 10^{-2}</math></b>	$2.16 \times 10^{-2}$	$1.36 \times 10^{-2}$
5	$L_{dp}^{l_2}$	0.852	$5.61 \times 10^{-4}$	$3.59 \times 10^{-2}$	$3.59 \times 10^{-2}$	$9.39 \times 10^{-3}$	$7.00 \times 10^{-3}$
6	$L_{dp}^{mi}$	0.848	$7.78 \times 10^{-4}$	$9.27 \times 10^{-2}$	$6.40 \times 10^{-2}$	<b><math>5.33 \times 10^{-3}</math></b>	<b><math>4.51 \times 10^{-3}</math></b>
7	$L_{ce}$	0.846	$2.06 \times 10^{-2}$	$3.73 \times 10^{-1}$	$1.81 \times 10^{-1}$	$1.96 \times 10^{-1}$	$5.48 \times 10^{-2}$
8	$L_{iou}$	0.847	$6.62 \times 10^{-3}$	<b><math>9.75 \times 10^{-2}</math></b>	$8.60 \times 10^{-2}$	$1.98 \times 10^{-1}$	$4.53 \times 10^{-2}$
9	$L_{eo}^{l_2}$	0.844	$1.85 \times 10^{-2}$	$2.50 \times 10^{-1}$	$1.53 \times 10^{-1}$	$1.91 \times 10^{-1}$	$4.70 \times 10^{-2}$
10	$L_{eo}^{mi}$	<b>0.857</b>	$1.53 \times 10^{-2}$	$1.12 \times 10^{-1}$	$1.08 \times 10^{-1}$	$1.52 \times 10^{-1}$	$4.21 \times 10^{-2}$
11	$L_{dp}^{l_2}$	<b>0.857</b>	$1.77 \times 10^{-2}$	$1.48 \times 10^{-1}$	$1.62 \times 10^{-1}$	$8.85 \times 10^{-2}$	$3.03 \times 10^{-2}$
12	$L_{dp}^{mi}$	0.854	<b><math>6.61 \times 10^{-3}</math></b>	$1.03 \times 10^{-1}$	<b><math>5.19 \times 10^{-2}</math></b>	<b><math>3.26 \times 10^{-2}</math></b>	<b><math>6.24 \times 10^{-3}</math></b>

set randomly into train, validation, and test partitions which contain 70%, 20%, and 10% of the samples.

Table 2 shows the quantitative results from the experiments with UTKFace. The interpretation of the rows and columns is the same as in Table 1. Again, the model fairness improved w.r.t. to almost all proposed fairness metrics for both sensitive attributes. In addition, applying any fairness loss led to an improvement in the prediction accuracy with the sensitive attribute  $y_s = \textit{Gender}$ . Experiments with the sensitive attribute  $y_s = \textit{Ethnicity}$  also improved the prediction accuracy except when we applied the  $L_{eo}^{l_2}$  loss.

### 4.3 SIIM-ISIC Melanoma Classification

The prediction target attribute in our experiments is a diagnosis, whether a lesion is malignant or benign. Each image has annotations of a binary gender variable (Male, Female) which we use as the sensitive attribute. The performance of the trained model is quantified with the area under the receiver operating curve ( $AUC$ ) on the validation partition, which was the standard evaluation metric in this challenge. We use EfficientNet-B1 as the classification model in all experiments as with UTKFace. As the data set only contains annotations in the original train partition, we used these annotated images and randomly assigned them to train, validation, and test partitions consisting of 70%, 20%, and 10% of the original train samples. We used different transformations to augment the training data, which improved the baseline  $AUC$  score (details in the supplementary material). Since the data set only contains a small fraction of malignant samples, we used the effective number of samples [6] as a class weighting method to deal with this label imbalance. Each sample is assigned a normalized weight  $\alpha_i = \frac{1-\beta}{1-\beta^{n_i}}$  according to the frequency  $n_i$  of the  $i$ -th class

**Table 3.** Experimental results with manually selected weighting coefficients  $\lambda$  for predicting the binary target attribute  $y_t = \textit{Diagnosis}$  from skin lesion images with the SIIM-ISIC Melanoma classification data set. The experiments use the binary variable *Gender* as the sensitive attribute  $y_s$ .

#	Loss	AUC	$L_{iou}$ (12)	$L_{eo}^{l_2}$ (6)	$L_{eo}^{mi}$ (7)	$L_{dp}^{l_2}$ (3)	$L_{dp}^{mi}$ (4)
1	$L_{ce}$	0.829	$1.26 \times 10^{-3}$	$6.22 \times 10^{-2}$	$3.51 \times 10^{-2}$	$3.54 \times 10^{-4}$	$9.27 \times 10^{-4}$
2	$L_{iou}$	0.801	<b><math>5.71 \times 10^{-5}</math></b>	<b><math>3.17 \times 10^{-2}</math></b>	<b><math>1.67 \times 10^{-4}</math></b>	<b><math>1.67 \times 10^{-4}</math></b>	<b><math>1.32 \times 10^{-4}</math></b>
3	$L_{eo}^{l_2}$	<b>0.869</b>	$4.10 \times 10^{-4}$	$8.52 \times 10^{-2}$	$3.95 \times 10^{-2}$	$7.60 \times 10^{-4}$	$1.27 \times 10^{-3}$
4	$L_{eo}^{mi}$	0.854	$7.54 \times 10^{-4}$	$7.71 \times 10^{-2}$	$4.38 \times 10^{-2}$	$7.41 \times 10^{-4}$	$1.43 \times 10^{-3}$
5	$L_{dp}^{l_2}$	0.804	$6.62 \times 10^{-4}$	$8.44 \times 10^{-2}$	$4.86 \times 10^{-2}$	$3.53 \times 10^{-4}$	$8.89 \times 10^{-4}$
6	$L_{dp}^{mi}$	0.835	$3.64 \times 10^{-4}$	$4.71 \times 10^{-2}$	$2.39 \times 10^{-2}$	$4.52 \times 10^{-4}$	$9.43 \times 10^{-4}$


in the train partition. The hyper-parameter  $\beta$  adjusts these weights according to the label distribution in a particular data set, which we set to  $\beta = 0.9998$ . Table 3 shows the results of the experiments with the fine-tuned class-weighting. The application of the EO-based fairness losses  $L_{eo}^{l_2}$  and  $L_{eo}^{mi}$  improved the AUC score considerably. Additionally, incorporation of our novel  $L_{iou}$  based fairness loss helped to improve the fairness of the model w.r.t the baseline for all of the proposed fairness-based evaluation metrics.

It is to be noted that our work is not comparable to closely related approaches. Some works [1,21] propose to remove biases w.r.t. sensitive attributes from the feature representation of the model. Our approach instead focuses on enabling the prediction accuracy of the target attribute to not depend on sensitive attributes. Recent works with similar tasks as ours propose loss functions based on distance [12] or cosine similarity [47] measures while we explore the inherent fairness-driven probabilistic measures as the loss functions in our experimental setup. Such loss functions were not studied before for visual data sets, so a reasonable comparison with methods such as [12,47] is also not possible.


#### 4.4 Effect of $\lambda$ on Fairness vs Accuracy

We studied the effect of the coefficient  $\lambda$  on both model fairness and classification accuracy. Its value depends on the dataset, optimized loss, and whether achieving high accuracy is more relevant for a specific application or fairness. Note that setting  $\lambda = 0$  reverts to the original image classification problem. We obtained the results for the initial experiments by heuristically searching for good weighting coefficients. We observed that starting with small  $\lambda$ -values is beneficial and studied whether the obtained model is fair. One could increase  $\lambda$  until the classification performance significantly decreases. Later, we performed hyper-parameter optimization (HPO) to find  $\lambda$  that leads to the best improvement in the model fairness with maintaining high classification accuracy.

**Heuristic Search** We use the Bessel corrected standard deviation  $\sigma_{IoU}(\lambda) = \sqrt{\frac{1}{K_s-1} \sum_{i=1}^{K_s} (IoU_{\theta}(b_i) - \overline{IoU_{\theta}(b)})^2}$  of the  $IoU_{\theta}(b)$  (10) as a fairness evaluation



**Fig. 2.** Experimental results with the fairness loss weighting coefficient  $\lambda$ . The standard deviation of the  $IoU_{\theta}(b)$  for different sensitive class labels  $b$  is used to quantify the model fairness. The prediction accuracy quantifies the classification performance.



**Fig. 3.** Experimental results with the fairness loss weighting coefficient  $\lambda$  on UTKFace and SIIM-ISIC Melanoma. The standard deviation of the  $IoU_{\theta}(b)$  for different sensitive class labels  $b$  is used to quantify the model fairness (less is fairer). The prediction accuracy quantifies the model performance. The blue lines show the baseline performance and fairness values.

measure and the prediction accuracy as a performance measure. The results for the manual selection of  $\lambda$  on CelebA can be seen in Fig. 2. The  $L_{iou}$  and  $L_{eo}^{mi}$  losses improved the fairness w.r.t. the baseline models (trained with the standard cross-entropy loss  $L_{ce}$ ) for both sensitive attributes. In contrast, the  $L_{eo}^{l2}$  loss and the DP-based losses  $L_{dp}^{mi}$  and  $L_{dp}^{l2}$  improved the fairness only for the sensitive attribute  $y_s = Male$ .

**Hyperparameter Optimization (HPO)** We performed HPO of the weighting coefficient  $\lambda$  for the UTKFace and SIIM-ISIC Melanoma data sets with  $\sigma_{IoU}(\lambda)$  as the minimization objective and searching  $\lambda$  within the range  $[1 \times 10^{-1}, 1 \times 10^3]$ . We use the validation partition for the HPO trials and evaluate the resulting models on the test partition. Due to time and resource constraints, we restrict our experiments to the novel IoU loss  $L_{iou}(\theta)$  (12) with 60 HPO trials on the UTKFace data set and 20 trials on the SIIM-ISIC Melanoma classification data set. The results of these HPO experiments are shown in Fig. 3. We observe a clear trend for experiments on UTKFace with the sensitive at-

tribute  $y_s = \textit{Ethnicity}$  (left figure), where the fairness of the baseline model is relatively low. With an exponentially increasing  $\lambda$ , there is a linear improvement in the model fairness and a linear decrease in the prediction accuracy. However, such a trend does not exist for experiments with UTKFace for the sensitive attribute  $y_s = \textit{Gender}$  (figure in the middle), where the baseline fairness is already good. Here, the model fairness and the prediction accuracy decrease linearly with exponentially increasing  $\lambda$  values. Additionally, there is more variation in the fairness improvements when  $\lambda$  is larger. In these experiments with UTKFace, there is a certain range of  $\lambda$  values (the region with yellow dots) that simultaneously improves the model fairness and prediction accuracy. The baseline fairness is already high for the SIIM-ISIC Melanoma data set with the sensitive attribute  $y_s = \textit{Gender}$  (right figure). Hence, the prediction accuracy and fairness improvement are independent of the  $\lambda$  value, as its effect on these metrics seems random.

## 5 Conclusion

In this work, we presented the theoretical intuition toward obtaining fair image classification models. We implemented various fairness metrics as standardized differentiable loss functions for categorical variables and compared their effectiveness in bias mitigation when compared to our novel IoU loss. Our experiments on publicly available facial and medical image data sets show that the proposed fairness losses do not degrade the classification performance on target attributes and reduce the classification bias at the same time. With this work and the publication of our source code, we provide a tool that encourages further work in this research direction. An interesting future work would be the visualization of relevant regions in the input image space that make the fair model less biased compared to the baseline model, trained with a standalone cross-entropy loss.

**Acknowledgement.** This work primarily received funding from the German Federal Ministry of Education and Research (BMBF) under *Software Campus* (grant 01IS17044) and the Competence Center for Big Data and AI *ScaDS.AI Dresden/Leipzig* (grant 01/S18026A-F). The work also received funding from Deutsche Forschungsgemeinschaft (DFG) (grant 389792660) as part of TRR 248 and the Cluster of Excellence *CeTI* (EXC2050/1, grant 390696704). The authors gratefully acknowledge the Center for Information Services and HPC (ZIH) at TU Dresden for providing computing resources.

## References

1. Alvi, M., Zisserman, A., Nellaker, C.: Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. *ECCV Workshops, Lecture Notes in Computer Science* **11129** (2019)
2. Amini, A., Soleimany, A.P., Schwarting, W., Bhatia, S.N., Rus, D.: Uncovering and mitigating algorithmic bias through learned latent structure. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* pp. 289–295 (2019)

3. Bercea, C.I., Wiestler, B., Ruckert, D., Albarqouni, S.: Feddis: Disentangled federated learning for unsupervised brain pathology segmentation. arXiv preprint arXiv:2103.03705 (2021)
4. Buolamwini, J., Gebru, T.: Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research* **81**, 77–91 (2018)
5. Chen, R.J., Chen, T.Y., Lipkova, J., Wang, J.J., Williamson, D.F., Lu, M.Y., Sahai, S., Mahmood, F.: Algorithm fairness in ai for medicine and healthcare. arXiv preprint arXiv:2110.00603 (2021)
6. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9268–9277 (2019)
7. Denton, E., Hutchinson, B., Mitchell, M., Gebru, T.: Detecting bias with generative counterfactual face attribute augmentation. *CoRR* (2019)
8. Denton, E., Hutchinson, B., Mitchell, M., Gebru, T., Zaldivar, A.: Image counterfactual sensitivity analysis for detecting unintended bias. *CVPR Workshop on Fairness Accountability Transparency and Ethics in Computer Vision* (2019)
9. Dhar, P., Gleason, J., Roy, A., Castillo, C.D., Chellappa, R.: Pass: Protected attribute suppression system for mitigating bias in face recognition. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* pp. 15087–15096 (2021)
10. Dwork, C., Immorlica, N., Kalai, A.T., Leiserson, M.: Decoupled classifiers for group-fair and efficient machine learning. *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, PMLR **81**, 119–133 (2018)
11. Georgopoulos, M., Oldfield, J., Nicolaou, M.A., Panagakis, Y., Pantic, M.: Mitigating demographic bias in facial datasets with style-based multi-attribute transfer. *International Journal of Computer Vision* **129**, 2288–2307 (2021)
12. Gong, S., Liu, X., Jain, A.K.: Mitigating face recognition bias via group adaptive classifier. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 3414–3424 (2021)
13. Hetey, R.C., Eberhardt, J.L.: The numbers don’t speak for themselves: Racial disparities and the persistence of inequality in the criminal justice system. *Current Directions in Psychological Science* **27**(3), 183–187 (2018)
14. Hou, X., Li, Y., Wang, S.: Disentangled representation for age-invariant face recognition: A mutual information minimization perspective. *IEEE/CVF International Conference on Computer Vision (ICCV)* pp. 3672–3681 (2021)
15. Jacobs, J.A.: Gender inequality and higher education. *Annual Review of Sociology* **22**, 153–185 (1996)
16. Jalal, A., Karmalkar, S., Hoffmann, J., Dimakis, A.G., Price, E.: Fairness for image generation with uncertain sensitive attributes. *Proceedings of the 38th International Conference on Machine Learning* (2021)
17. Joo, J., Karkkainen, K.: Gender slopes counterfactual fairness for computer vision models by attribute manipulation. *Proceedings of the 2nd International Workshop on Fairness, Accountability, Transparency and Ethics in Multimedia* (2020)
18. Jung, S., Lee, D., Park, T., Moon, T.: Fair feature distillation for visual recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 12110–12119 (2021)
19. Karkkainen, K., Joo, J.: Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* pp. 1548–1558 (2021)

20. Ke, J., She, Y., Lu, Y.: Style normalization in histology with federated learning. *IEEE 18th International Symposium on Biomedical Imaging (ISBI)* pp. 953–956 (2021)
21. Kim, B., Kim, H., Kim, K., Kim, S., Kim, J.: Learning not to learn: Training deep neural networks with biased data. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 9012–9020 (2019)
22. Kinyanjui, N., Odonga, T., Cintas, C., Codella, N., Panda, R., Sattigeri, P., Varshney, K.R.: Fairness of classifiers across skin tones in dermatology. *Medical Image Computing and Computer Assisted Intervention (MICCAI)* **12266**, 320–329 (2020)
23. Larrazabal, A.J., Nieto, N., Peterson, V., Milone, D.H., Ferrante, E.: Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the National Academy of Sciences* **117**(23), 12592–12594 (2020)
24. Li, P., Zhao, H., Liu, H.: Deep fair clustering for visual learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 9070–9079 (2020)
25. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. *IEEE International Conference on Computer Vision (ICCV)* pp. 3730–3738 (2015)
26. McDuff, D., Song, Y., Kapoor, A., Ma, S.: Characterizing bias in classifiers using generative models. *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (2019)
27. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. *ACM Computing Surveys* **54**(6) (2021)
28. Merler, M., Ratha, N., Feris, R.S., Smith, J.R.: Diversity in faces. *arXiv preprint arXiv:1901.10436* (2019)
29. Morales, A., Fierrez, J., Rodriguez, R.V., Tolosana, R.: Sensitivenets: Learning agnostic representations with application to face images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(6) (2021)
30. Ongena, S., Popov, A.: Gender bias and credit access. *Journal of Money, Credit and Banking* **48** (2016)
31. O’Neil, C.: *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown Publishing Group (2016)
32. Quadrianto, N., Sharmanska, V., Thomas, O.: Discovering fair representations in the data domain. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 8219–8228 (2019)
33. Raji, I.D., Gebru, T., Mitchell, M., Buolamwini, J., Lee, J., Denton, E.: Saving face: Investigating the ethical concerns of facial recognition auditing. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* p. 145–151 (2020)
34. Rajotte, J.F., Mukherjee, S., Robinson, C., et. al.: Reducing bias and increasing utility by federated generative modeling of medical images using a centralized adversary. *Proceedings of the Conference on Information Technology for Social Good* pp. 79–84 (2021)
35. Ramaswamy, V.V., Kim, S.S.Y., Russakovsky, O.: Fair attribute classification through latent space de-biasing. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 9301–9310 (2021)
36. Rotemberg, V., Kurtansky, N., Betz-Stablein, B., et. al.: A patient-centric dataset of images and metadata for identifying melanomas using clinical context. *Scientific Data* **8**(34) (2021)
37. Ryu, H.J., Adam, H., Mitchell, M.: Inclusivefacenet: Improving face attribute detection with race and gender diversity. *Workshop on Fairness, Accountability, and Transparency in Machine Learning* (2018)



38. Seyyed-Kalantari, L., Liu, G., McDermott, M., Chen, I.Y., Ghassemi, M.: Chexclusion: Fairness gaps in deep chest x-ray classifiers. *Pacific Symposium On Biocomputing* **26**, 232–243 (2021)
39. Sharma, A.K., Foroosh, H.: Slim-CNN: A light-weight CNN for face attribute prediction. *15th IEEE International Conference on Automatic Face and Gesture Recognition* pp. 329–335 (2020)
40. Song, J., Shen, C., Yang, Y., Liu, Y., Song, M.: Transductive unbiased embedding for zero-shot learning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* pp. 1024–1033 (2018)
41. Tan, M., Le, Q.: EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the 36th International Conference on Machine Learning* **97**, 6105–6114 (2019)
42. Wang, A., Liu, A., Zhang, R., et. al.: Revise: A tool for measuring and mitigating bias in visual datasets. *International Journal of Computer Vision* (2022)
43. Wang, M., Deng, W.: Mitigate bias in face recognition using skewness-aware reinforcement learning. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 9322–9331 (2020)
44. Wang, M., Deng, W., Jiani Hu, J.P., Tao, X., Huang, Y.: Racial faces in-the-wild: Reducing racial bias by deep unsupervised domain adaptation. *IEEE/CVF International Conference on Computer Vision (ICCV)* pp. 692–702 (2019)
45. Wang, Z., Qinami, K., Karakozis, I.C., et. al.: Towards fairness in visual recognition: Effective strategies for bias mitigation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 8916–8925 (2020)
46. Xu, H., Liu, X., Li, Y., Jain, A., Tang, J.: To be robust or to be fair: Towards fairness in adversarial training. *Proceedings of the 38th International Conference on Machine Learning (PMLR)* **139**, 11492–11501 (2021)
47. Xu, X., Huang, Y., Shen, P., Li, S., Li, J., Huang, F., Li, Y., Cui, Z.: Consistent instance false positive improves fairness in face recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 578–586 (2021)
48. Zhang, B.H., Lemoine, B., Mitchell, M.: Mitigating unwanted biases with adversarial learning. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics and Society* (2018)
49. Zhang, Z., Song, Y., Qi, H.: Age progression/regression by conditional adversarial autoencoder. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 5810–5818 (2017)
50. Zhao, C., Li, C., Li, J., Chen, F.: Fair meta-learning for few-shot classification. *2020 IEEE International Conference on Knowledge Graph (ICKG)* pp. 275–282 (2020)
51. Zhao, J., Wang, T., Yatskar, M., Ordonez, V., Chang, K.W.: Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* pp. 2979–2989 (2017)