

Leveling Down in Computer Vision: Pareto Inefficiencies in Fair Deep Classifiers

Dominik Zietlow^{1,2,†,*}, Michael Lohaus^{1,3,†}, Guha Balakrishnan⁴, Matthäus Kleindessner¹,
Francesco Locatello¹, Bernhard Schölkopf^{1,2}, Chris Russell¹

¹Amazon Web Services, Tübingen, ²MPI-IS Tübingen, Germany, ³University of Tübingen, Germany, ⁴Rice University, USA

Abstract

Algorithmic fairness is frequently motivated in terms of a trade-off in which overall performance is decreased so as to improve performance on disadvantaged groups where the algorithm would otherwise be less accurate. Contrary to this, we find that applying existing fairness approaches to computer vision improve fairness by degrading the performance of classifiers across all groups (with increased degradation on the best performing groups).

Extending the bias-variance decomposition for classification to fairness, we theoretically explain why the majority of fairness methods designed for low capacity models should not be used in settings involving high-capacity models, a scenario common to computer vision. We corroborate this analysis with extensive experimental support that shows that many of the fairness heuristics used in computer vision also degrade performance on the most disadvantaged groups. Building on these insights, we propose an adaptive augmentation strategy that, uniquely, of all methods tested, improves performance for the disadvantaged groups.

1. Introduction

High-capacity neural classifiers achieve state-of-the-art performance on most computer vision tasks when assessed by overall test set accuracy. However, researchers have begun examining the unfairness of these models. Here, we use ‘unfairness’ to refer to systematic accuracy differences across protected subgroups defined by human-sensitive attributes like gender and race [7, 8, 36].¹ These differences in accuracy can harm certain population groups, and as a result, numerous strategies for training models that match various measures of accuracy across subgroups have been developed [2, 5, 33, 85]. Typically, such methods quantify unfairness by comparing accuracy-related rates between various groups, for example the *Difference of Equal Opportu-*

nity (DEO) [36] compares groupwise true positive rates.

Many recent computer vision fairness studies are motivated by a fairness-accuracy trade-off originally targeting low-capacity models [1, 2, 15, 22, 47, 56, 85] where high accuracy on better performing (and often larger) groups comes at the cost of lower accuracy on the worse performing (and often smaller) groups. In this case, it is possible to increase fairness by reducing the accuracy on the best-performing group and increasing the accuracy on the worst-performing group, see models A and B in Figure 1.

We revisit this trade-off and show that it does not hold when using high-capacity neural classifiers prevalent in computer vision (see Figure 2). Instead, many fairness methods *degrade the accuracy of networks on all groups*, with a greater degradation occurring for the better performing groups. This increases fairness, but at the cost of producing a worse performing classifier (Figure 1 model C).

The phenomenon of balancing fairness by degrading the performance on the better off groups is referred to as *levelling down* in law and philosophy, where it has received substantial criticism [12, 19, 26, 41, 57]. The behavior we are concerned with is even more extreme than the typical levelling down – rather than just lower the performance on the best performing groups, every group is worse off.

If fairness methods decrease performance for all groups, they are *Pareto Inefficient* with respect to group accuracy and should not be used in contexts where the accuracy of any group is a primary concern.² For example, we find that for classifiers on CelebA [52] regularized by a DEO fairness measure, the increased fairness comes at the cost of degraded performance for every group including the worst performing group (see Figure 2). We attribute this problem to two issues:

¹In this work, we focus on fairness measures that compare accuracy across groups. These are in contrast to measures such as demographic parity [14, 28], which matches the proportion of positive decisions per group.

²Various works in fairness have made use of *Pareto Efficiency*. It has been used both to refer to trading-off global accuracy against notions of fairness [76], and for the notion of trading-off per group accuracies against each other [56]. We only refer to the second case.

[†] Work done during an internship at AWS

* Corresponding author zietld@amazon.de

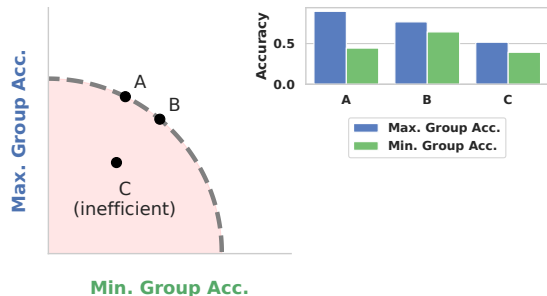


Figure 1. **Pareto curve.** We depict the typical trade-off assumed by most computer vision fairness studies. The accuracies of a family of classifiers on its best- and worst-performing groups form a *Pareto curve* (dotted gray line). Points A and B are maximally efficient configurations that lie on the curve; B is fairer and has a lower accuracy difference between the groups (see bar plot, right). Point C is as fair as B, but is inefficient because it reduces the accuracies of both groups. Applying accuracy-based fairness methods to deep networks tend to result in inefficient configurations like C.

High-capacity classifiers fit training data nearly perfectly: Most methods within the fairness community are designed for low-dimensional data where a classifier cannot fit data from multiple distinct distributions well, even on training data [1, 2, 15, 22, 36, 47, 56, 85]. This is not the case in computer vision, where high-dimensional data and high-capacity models mean that near-zero training error is common [87]. Fairness notions that are aligned with the perfect classifier are therefore trivially satisfied on training data. *Unfortunately, most existing methods do not take this into account and enforce fairness constraints on the training set [2, 10, 22, 24, 25, 27, 44, 47, 53, 56, 61, 82, 85].*

Inappropriate evaluation of fairness methods: Most papers presenting fairness methods report a combination of accuracy and a fairness measure such as DEO, and take a decrease in accuracy and an improved fairness measure as an indication that the method is successfully trading off fairness against accuracy [2, 65, 66, 81, 82, 85]. The crucial question of whether the learned models work better for more disadvantaged groups remains unanswered. Consequently, the choice of metrics may mask a systematic deterioration of classifiers, where accuracy decreases across all groups, and not just in the high-accuracy groups.

To address these limitations, we make three contributions:

1. We revisit the bias-variance trade-off in a simple **decomposition of fairness on test data into training error (bias) and generalization error (variance)**, and observe that in situations where the training error goes to zero, any measure of fairness must be dominated by generalization error. As such, methods not using held-out data for fairness constraints cannot work for the

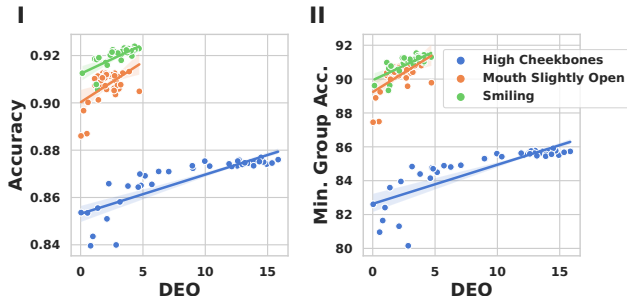


Figure 2. **Accuracy-fairness trade-off.** We trained multiple fair models on CelebA. The protected attribute is “Male” (as annotated in the dataset). Different trade-offs are achieved by varying the strength of a fairness regularizer added to the overall loss [61, 83]. In (I), we plot overall accuracy against the fairness measure DEO (low is fair). As expected, improvements in fairness come with a loss in accuracy. However, this comes at an additional cost; as we see in (II), the worst group accuracy also decreases as DEO improves.

high-capacity classifiers common to computer vision.

2. We perform an extensive **evaluation of existing fairness methods** and show that their reported improvement in fairness metrics is accompanied by **worse performance across all groups**.
3. To confirm our theoretical analysis that **better generalization is key** to improving performance on the most disadvantaged groups, we explore the use of **data augmentation combined with adaptive sampling**. We propose a novel GAN-based augmentation and show that it improves fairness by improving accuracy on the most disadvantaged groups on the CelebA dataset.

2. Related Work

Notions of fairness: We focus on notions of fairness that aim to balance classifier errors across population subgroups, for example, by matching error rates across genders or different racial groups. See [78] for a summary of other forms of fairness, and [79] for when such definitions are inappropriate. One common fairness measure is ‘equal opportunity’ (EO) [36]. This requires a classifier to have equal true positive rates (TPR) on each subgroup. We use \mathcal{A} to refer to the set of protected attributes, such as race or gender, $Y \in \{0, 1\}$ to refer to the ground-truth label of a data-point, and $\hat{Y} \in \{0, 1\}$ the prediction of Y . For two groups $a, a' \in \mathcal{A}$ the violation of equal opportunity is measured by the *difference in equal opportunity* (DEO) defined as

$$|\mathbb{P}(\hat{Y} = 1|Y = 1, A = a) - \mathbb{P}(\hat{Y} = 1|Y = 1, A = a')|. \quad (1)$$

We measure the *difference in equalized odds* (DEOdds) by

$$\sum_y |\mathbb{P}(\hat{Y}=1|Y=y, A=a) - \mathbb{P}(\hat{Y}=1|Y=y, A=a')|. \quad (2)$$

Another fairness notion of recent interest is min-max fairness [22, 55]. In contrast to EO (and the other group fairness notions) min-max fairness does not equalize a statistic across groups, but solely strives to decrease the classification error for the subgroup with the highest error as much as possible. Formally, min-max fairness methods optimize

$$\min_{\hat{Y}} \max_{a \in \mathcal{A}} \mathbb{P}(\hat{Y} \neq Y | A = a). \quad (3)$$

As a measure of min-max fairness we report the minimum group accuracy (*i.e.*, $1 - \text{error}$) and similarly the minimum group TPR. Note that the minimum group TPR is trivially maximized by the constant classifier $\hat{Y} = 1$.

While many variants of these error-matching fairness definitions exist³, the motivation for all measures is broadly the same: to prevent a tyranny of the majority, in which accuracy on smaller groups of people are neglected to achieve higher accuracy on larger groups.

Fairness in computer vision: Following foundational work [13] that identified the systematic failings of face analysis systems on particular racial and gender demographics, algorithmic fairness has attracted growing attention in the computer vision community. Image datasets are known to be biased due to sampling inequalities [3, 64, 75], and human face datasets have been particularly scrutinized [6, 45, 49, 50, 58] as models trained on these data can exhibit systematic failings with respect to attributes protected by the law [46]. Multiple approaches to mitigating dataset bias include collecting more diverse examples [58], using image synthesis to compensate for distribution gaps [6, 50, 72, 74, 82], and resampling [51]. More recently, approaches have been proposed to mitigate the systematic biases of deep learning-based vision models [18, 30, 34, 40, 44, 51, 65, 66, 80, 81].

Outside of computer vision, the most common approaches add additional fairness measures to loss functions [9, 47, 53, 61, 69, 83], enforcing fair representations that are independent of protected attributes [10, 29, 54, 86], and augmenting training data to promote balance [66]. Most of these studies assume an accuracy-fairness tradeoff which we show can lead to suboptimal training or misleading evaluation for deep neural networks.

Active sampling for min-max fairness: While the methods of [22, 55] achieve min-max fairness by means of reweighting and retraining, the recent method of [1] uses

³As a starting point, [79] identified 17 accuracy-based measures excluding min-max fairness.

adaptive sampling and standard SGD updates to minimize a differentiable proxy of (3). This makes the latter method easily applicable to deep neural network training.

3. On Accuracy-based Fairness in Low- and High-capacity Classifiers

This section outlines the challenges peculiar to enforcing fairness in the high-capacity classifiers common to computer vision. In particular, we focus on accuracy-based notions of fairness that aim to match accuracy across groups.

As noted by Wachter et al. [79], any accuracy-based fairness measure⁴ is trivially satisfied by a classifier with zero error. However, as the typical datasets (see [2, 43, 85] for a range of examples) used in the fairness literature are low-dimensional, and with large amounts of label volatility (see Section 3.1), classifiers with zero error do not occur in practice, even on the training data. Training a low-capacity classifier with accuracy-based fairness constraints on the training set is a common approach. Where training error remains high, such approaches can remain effective even if fairness constraints are only enforced over the training data.

This is not the case when enforcing fairness of deep-learning based classifiers on computer vision datasets. Such datasets are empirically *shatterable* [77]. Even if the images in the training set are randomly relabelled, it is possible to learn a classifier with zero error on the training set [87]. In such scenarios, accuracy-based fairness definitions are trivially satisfied over the training set.

In the following subsection, we formalize the gap between fairness on training and held-out data and discuss its implications for computer vision.

3.1. Bias-variance decomposition for classification

It is common in statistics to decompose the error into three terms: irreducible label noise N , a bias B representing how well the regressor can fit the dataset, and a variance V representing additional error induced when generalizing to new data. The standard decomposition was formalized around the squared loss [31], extended to the zero-one loss [48] and generalized to arbitrary losses [23].

We build on the latter formulation, which we summarize in a condensed manner—see [23] for more details. We consider the task of learning $f \in \mathcal{F} : \mathcal{X} \rightarrow \mathcal{Y}$, where $f(x)$ should be a good prediction of the label y of input point x . The quality of a prediction is measured by a loss function $L : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$. The optimal prediction for input x is

$$y_*(x) = \arg \min_{y'} \mathbb{E}_{y|x} [L(y, y')],$$

and the expected loss of model f on x is $\mathbb{E}_{y|x} [L(y, f(x))]$. The conditional expectation corresponds to the fact that,

⁴Wachter et al. refer to these accuracy-based measures as bias-preserving fairness metrics.

in general, the label y is a non-deterministic function of input x . We learn the model f based on a training set $D_n = \{(x_1, y_1), \dots, (x_n, y_n)\}$, and in order to remove the dependency on D_n we consider $\mathbb{E}_{D_n, y|x} [L(y, f(x))]$ as the expected loss of f on x . The latter is the quantity of interest that we want to decompose into bias, variance and noise.

We define the main prediction $y_m(x)$ on x as

$$y_m(x) = \arg \min_{y'} \mathbb{E}_{D_n} [L(f(x), y')].$$

This allows for definition of the intrinsic noise $N(x)$ (expected error induced by label disagreement for a particular datapoint), the bias $B(x)$ (error induced by systematic model imperfection), and the variance $V(x)$ (error difference from the main prediction):

$$B(x) = L(y_*(x), y_m(x)), \quad (4)$$

$$V(x) = \mathbb{E}_{D_n} [L(y_m(x), f(x))], \quad (5)$$

$$N(x) = \mathbb{E}_{y|x} [L(y, y_*(x))]. \quad (6)$$

As [23] shows, for certain loss functions L (including but not limited to squared, zero-one loss, and false negative rate⁵) we can decompose

$$\text{err}_x := \mathbb{E}_{D_n, y|x} [L(y, f(x))] \quad (7)$$

$$= c_1(x)N(x) + B(x) + c_2(x)V(x) \quad (8)$$

for some $c_1(x), c_2(x) \in \mathbb{R}$.

3.2. Expected fairness violations

We make the simplifying assumption that we want to enforce error parity across only two groups A and B . If we define notation for both groups, we use G in place of A or B . We depart slightly from [23] as we consider the per-group error of a classifier trained over both groups rather than the classifier’s overall error. The expected fairness violation for the two-groups case E_{fair} can be defined as

$$E_{\text{fair}} = |\mathbb{E}_{x \in A} [\text{err}_x] - \mathbb{E}_{x \in B} [\text{err}_x]|. \quad (9)$$

With auxiliary definitions

$$B_G = \mathbb{E}_{x \in G} [B(X)], \quad (10)$$

$$V_G = \mathbb{E}_{x \in G} [c_2(x)V(X)], \quad (11)$$

$$N_G = \mathbb{E}_{x \in G} [c_1(x)N(X)], \quad (12)$$

the fairness violation can be rephrased as

$$E_{\text{fair}} = |N_A + B_A + V_A - (N_B + B_B + V_B)|. \quad (13)$$

⁵The use of zero-one loss allows us to consider the bias-variance decomposition of DEOdds, while false negative rate allows us to consider DEO.

3.2.1 Fairness for low-capacity classifiers

For low-capacity models, the variances are strongly dominated by the noises and biases [37]. That is $N_G + B_G \gg V_G$ and the fairness violation can be approximated as

$$E_{\text{fair}} \approx |N_A + B_A - N_B - B_B|. \quad (14)$$

Importantly, $N_G + B_G$ can be directly estimated from the per group training set, and as such methods that match the loss across different groups at training time are likely to work here. Indeed, many methods [2, 36, 85] enforce fairness by explicitly balancing errors on the training set, under the reasonable assumption that these fairness constraints will generalize well to unseen data. The work [25] addressed this by minimizing a combination of unfairness on the training set and an upper bound of the generalization error. However, in practice, the majority of approaches are effective without considering generalization error.

3.2.2 Fairness for high-capacity classifiers

In contrast, for classification tasks in computer vision, the typical behavior is very different. While human labelling of image data is often a noisy process [59], the majority of computer vision datasets do not model label noise and often either only collect one label per datapoint [52], or explicitly denoise the collected labels [21] and only use the most common label per datapoint as the true datapoint.⁶ As such, when we treat dataset acquisition as a stochastic process that can only assign one label per datapoint, eq. (6), which measures the expected loss caused by disagreement between labels, is zero. This point is key to our argument; where multiple disagreeing labels are collected for individual datapoints, it may be possible for methods to improve these forms of fairness without making use of held-out data.

For a computer vision classifier trained to convergence on a fixed dataset, the bias terms also disappear [84]. Here, data is high-dimensional and models have essentially arbitrarily high capacity. Hence, the perfect classifier, which has zero error and predicts the ground-truth label for each datapoint, lies in \mathcal{F} . Hence, $V_G \gg B_G \approx 0$, and the fairness violation is dominated by the generalization error

$$E_{\text{fair}} \approx |V_A - V_B|, \quad (15)$$

i.e., the fairness violation is predominantly determined by the difference of the variances. Others have made related observations, for example, [82] emphasized the importance of data augmentation when reweighting samples, to prevent the gradients going to zero. However, even with augmentations, error on the training set remains much lower than the error on held-out data, and the variance still dominates.

⁶*i.e.*, given an image labeled as cat twice, and dog once, the image is simply treated as being labelled cat.

Given this, it is unsurprising that these methods for enforcing fairness on a training set are not popular in computer vision, and instead research typically centers around heuristics, such as reweighting samples [51], training multiple classifiers on data subsets, and averaging them [70, 71], or data augmentation [50]. With such a wide range of approaches, it is not possible to formally analyze them, beyond saying that they do not use held-out validation data to estimate generalization error, but they all successfully trade-off accuracy against a wide range of fairness metrics.

3.2.3 Rethinking fairness measures in computer vision

We have shown it is not possible to predict per-group error rates on the test set without computing fairness measures on held-out data. Yet the vast majority of methods do not do so, and as demonstrated in Section 5, still consistently improve fairness according to standard metrics such as equalized odds or equal opportunity. How is this possible?

One answer lies in the definitions of equalized opportunity (1) and odds (2) – they are perfectly satisfied by random or constant classifiers. As such, it is possible to increase fairness by decreasing performance on all groups, rather than by rebalancing the errors from one group to another.

Empirically, we find that this is exactly what occurs in practice: improvements to DEO and DEOdds are accompanied by a degradation in the accuracy for all groups, and not a rebalancing of accuracy from one group to another (see Figures 2, 4). While it is not possible to predict the per-group error for unseen data without making use of held-out data, it is always possible to degrade performance by injecting noise into gradients (which may be happening with regularized approaches [61, 83]), augmenting datasets with inappropriate synthetic examples (which might be happening with [66]), or by a wide range of well-motivated and sensible-sounding heuristics, including early stopping.

In light of this issue, when evaluating accuracy-based model fairness, *we recommend choosing a fairness metric that explicitly requires improving performance on the disadvantaged groups*. In our experiments, we measure the accuracy of the worst performing group, *i.e.*, min-max fairness of eq. (3). Another natural choice would be the true positive rate on the worst performing group, but this is easily gamed by classifiers with high false positive rates.

4. Improving accuracy on disadvantaged groups with synthetic data

In the previous section, we presented the importance of evaluating fairness based on the worst-performing group’s accuracy. Next, we consider how to use this insight, along with the result in (15) to better incorporate fairness during the model *training*. Expression (15) shows that fairness violations in high-capacity models are dominated by

the variance of the bias-variance decomposition. Consequently, one way to improve fairness without decreasing the performance for all groups is to decrease the variance of the worst-off group. It is well-known, and experimentally verified [11], that variance typically decreases as the training set increases in size and diversity. Given a fixed dataset, data augmentation is the only way to achieve that. Taking these insights into account, we propose an adaptive augmentation strategy that increases the diversity of samples for the worst-performing group in a classification task.

There are three technical challenges that we must address with our method. How to: (i) decide which group(s) requires augmentation (ii) generate high fidelity in distribution data, and (iii) reliably condition the augmented data on the protected groups and automatically provide target labels. We tackle these challenges by: (i) deploying adaptive sampling strategies using held-out data to prioritize worst-performing groups, (ii) using invertible GAN architectures and latent space traversals to edit images, and (iii) proposing g-SMOTE, a generalization of the synthetic minority over-sampling technique (SMOTE) [16], which produces new labeled training images by traversing GAN latent space.

4.1. Adaptive sampling

At each training iteration, we sample a random batch from one of two datasets: the original training dataset, and an extended training dataset. The extended training set is equal to the original training set at initialization. After each iteration, we determine the worst-performing group using a held-out evaluation dataset, augment a random batch from this group (using g-SMOTE in next section), and add it to the extended set. This approach is inspired by past studies [67, 68] that evaluate the generalization when training on partially augmented data.

We select a batch from the original dataset with probability λ (and $1 - \lambda$ from the extended dataset). The parameter λ allows us to maintain a certain proportion of the original data in the training set, preventing the extended training set from becoming predominantly augmented data. Note that this is more involved than simply balancing the number of elements in each group. Differences in generalization performance can stem from strongly imbalanced group sizes, but also from characteristics of either group that make generalization intrinsically harder.

4.2. Generalized SMOTE: g-SMOTE

Given images from the original dataset, we need a procedure to generative new synthetic images along with attribute labels. To accomplish this, we combine SMOTE [16] with the rich representational and generative capabilities of modern GANs. SMOTE is a simple sampling strategy to overcome imbalanced data that is based on linear feature interpolation between a datapoint and a random datapoint within

Algorithm 1 Adaptive Sampling

- 1: **Inputs:**
 Hyper-parameter $\lambda \in [0, 1]$
 Train dataset $D_{\text{Train}} = \{(x_0, y_0), (x_1, y_1), \dots\}, x_i \in \mathcal{X}, y_i \in \mathcal{Y}$
 Evaluation dataset $D_{\text{Eval}} = \{(x_0^e, y_0^e), (x_1^e, y_1^e), \dots\}, x_i^e \in \mathcal{X}, y_i^e \in \mathcal{Y}$
 Classifier $c_\phi : \mathcal{X} \rightarrow \mathcal{Y}$ (parameterized by ϕ)
 - 2: **Initialize:** $D_{\text{Aug}} := D_{\text{Train}}$
 - 3: **for** $i = 1, \dots, n_{\text{training steps}}$ **do**
 - 4: With probability λ , uniformly sample $(x_i, y_i) \in D_{\text{Train}}$, otherwise sample $(x_i, y_i) \in D_{\text{Aug}}$
 - 5: Update ϕ according to learning objective
 - 6: Determine weakest group based on learning objective and D_{Eval} and augment corresponding $x_{\text{Aug}}, y_{\text{Aug}}$ from that group
 - 7: $D_{\text{Aug}} \leftarrow D_{\text{Aug}} \cup \{(x_{\text{Aug}}, y_{\text{Aug}})\}$
 - 8: **end for**
-

the set of its k nearest neighbors. It has inspired many other approaches [17, 35, 38, 88]. Modern invertible GAN architectures [24, 27, 32, 63] allow one to ‘embed’ images into their latent spaces, which are known to be particularly effective domains to compare and edit images [6, 73]. We propose to use SMOTE in the GAN latent space to generate new diverse synthetic images along with attribute labels.

We propose a generalized variant of SMOTE (g-SMOTE), that extends the classical SMOTE strategy – linear interpolation between a datapoint and a random point among the datapoints m nearest neighbors – to uniform sampling within a k -dimensional simplex formed by k of the m nearest neighbors, to improve data diversity. Given a datapoint and its m nearest neighbors with the same target attribute, we randomly choose k neighbors that span a simplex of dimension $k - 1$ or smaller in the much higher dimensional GAN latent space. Within this simplex, we uniformly sample points in latent space and use the GAN generator to render their images. Our key assumption is that the simplex covers a label-consistent volume in latent space, *i.e.*, that every image stemming from that region shares the same target label. Algorithm S2 describes the data augmentation procedure, Figure S2 illustrates the interpolation mechanism, and Figure 3 shows a real-world example for $k = 3$. See supplementary materials for the effects of varying k and a comparison of SMOTE and g-SMOTE.

5. Experiments

We show that a wide range of methods for fairness in computer vision decrease performance on the worst group despite reporting improved fairness, and that, in contrast, our approach improves performance. As such, we follow existing experimental protocols as closely as possible. We show results on CelebA where all methods are reported to work, we use existing code, and compare each method to its corresponding unfair baseline. We demonstrate that:

Common fairness approaches in computer vision reduce performance for protected groups compared to unfair baselines. As such, we evaluate methods with respect to




Figure 3. Example g-SMOTE augmentations where $k=3$. Given a datapoint (green) and two neighbours (orange), linear interpolations in GAN latent space yield diverse images. Nearest neighbours are chosen to share the target attribute (“eyeglasses”). We give all interpolated images the same attribute label value.

their min-max performance on accuracy and TPR.

Adaptive debiasing with g-SMOTE improves min-max performance. See Section 5.1, and Figure 4 for experimental results. The benefit of adaptive debiasing is emphasized when training on a subset of data, because the GAN allows for effective unsupervised data augmentation.

Adaptive debiasing with g-SMOTE works with cross-sectional groups of multiple protected attributes. See Section 5.2 for experimental results. We show our method works with the four groups defined by the cross product of gender and target attribute binary groups.

g-SMOTE produces better data diversity than popular augmentation strategies. See Section 5.3. When combining g-SMOTE augmentation with popular augmentation methods, such as random flips, crops, rotations and RandAugment, we always observe an improvement in accuracy.

5.1. Evaluating min-max fairness on existing methods and adaptive debiasing

Our experimental setup follows [66] and [82] with full details in Appendix D. We compare against methods to improve fairness, including oversampling, domain discriminative training [70, 71], domain independent models [82], an adversarial approach [4], regularization [61, 83], FairMixup [20] and GAN-based offline dataset debiasing [66].

We summarize results in Figure 4. Some methods use a single multitask model (annotated with *), others use multiple single task models. Note that FairMixup uses a higher resolution (256 px) and a lower capacity model (ResNet18), which leads to a substantially different baseline performance. The only methods to increase the performance on the less accurate groups are g-SMOTE with

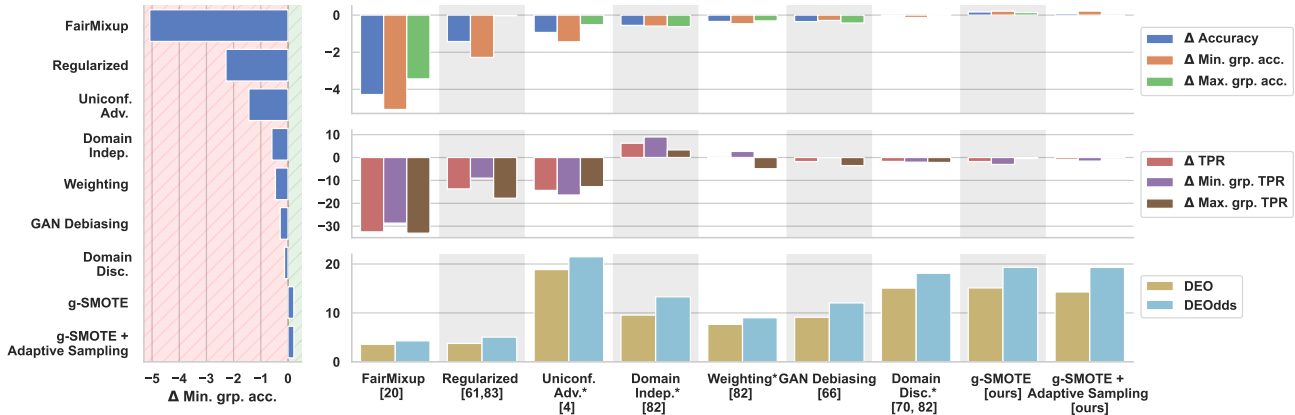


Figure 4. Fairness methods on the CelebA dataset. We report mean scores over 13 labels [66] call *gender independent*. **Left:** Change in minimum group accuracy difference relative to the unconstrained classifier. Only our approaches g-SMOTE and Adaptive g-SMOTE improve accuracy on the worst performing groups. All other methods lower the accuracy. **Top Right:** A decomposition of the change in accuracy into overall change, change in the best performing group, and the worst. **Center Right:** Decomposition of the change in True Positive Rate (TPR). Despite DEO being defined in terms of TPR, all but one of the fairness method decrease the TPR. Note that our g-SMOTE based approaches improve accuracy at the cost of decreasing TPR. **Bottom Right:** plots of DEO and DEOdds. These standard fairness measures track each other closely, with a high score in one corresponding to a high-score in the other. However, they are unpredictable of improvements in TPR or accuracy on any groups. Methods marked with * are based on the code-base of [82], training one model for multi-task classification. Other methods implement one-task classification models.

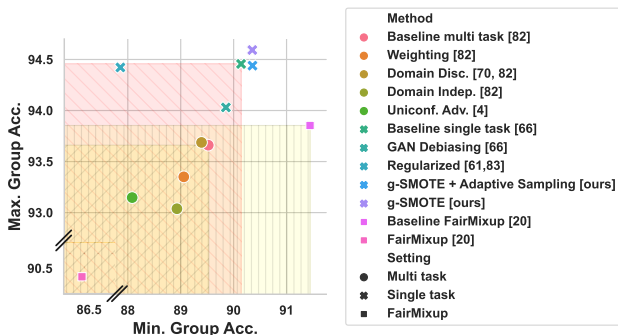


Figure 5. A comparison of fairness methods and unmodified classifiers showing the trade-off behaviour between better and worse performing groups. Results are averaged over 13 labels with *gender independent quality*. The shaded rectangles indicate Pareto-inefficient areas in which neither of the groups gained performance over the unmodified model. Note that FairMixup uses a smaller model and higher resolution. Our methods are the only ones to improve worst group performance over the relevant baseline.

and without adaptive sampling. The baselines lower accuracy across both protected groups. For the offline debiasing approach, we used our own implementation and the same GAN model as used for the g-SMOTE methods.

The argument set out in Section 3.1 suggests that for high capacity models, balancing datasets with adaptive sampling is less important than improving generalization by generating diverse synthetic data. Empirically, we find this to be

the case. Using g-SMOTE without adaptive sampling, we find that performance can be improved for all groups, with only a small decrease in min-group accuracy compared to adaptive sampling, see Table 1. We report the min. group accuracy per attribute in Table S2.

5.2. Adaptive g-SMOTE on multiple groups

Adaptive g-SMOTE is trivially extendable to multiple protected groups. We compare the adaptive g-SMOTE sampling with the regular g-SMOTE augmentation for two protected groups (“male” / not “male”) and four protected groups (“male” / not “male”, target attribute / not target attribute) in Table 1 and Table 2 respectively. The 32 target labels are selected such that each of the four protected groups has at least 0.3% of the total number of datapoints.

This comparison leads to three interesting findings: (1)

2 Protected Groups		Baseline	Adaptive g-SMOTE	g-SMOTE
10k train images	Acc.	89.41	89.57	89.58
	Min. grp. acc.	85.68	85.93	85.91
	DEO	23.18	25.00	24.65
	DEOdds	33.63	35.51	34.64

Table 1. Sampling strategies on CelebA. We report the means over 32 labels. The protected attribute is “male”. The models shown have the highest min. group accuracy over the training period.

g-SMOTE augmentation combined with adaptive sampling yields a significant improvement in min-max fairness when performed on four groups. (2) The performance improve-

4 Protected Groups		Baseline	Adaptive g-SMOTE	g-SMOTE
10k train images	Acc.	85.65	84.91	85.16
	Min. grp. acc.	62.16	64.36	63.54
	DEO	19.41	13.25	19.12
	DEOdds	28.22	19.82	27.75
160k train images	Acc.	87.37	85.77	87.27
	Min. grp. acc.	61.74	68.06	61.84
	DEO	21.99	12.27	21.91
	DEOdds	30.89	19.29	31.24

Table 2. Different sampling strategies on CelebA. Reported are the means over 32 labels. The protected attributes are “male” and the target attribute. The models shown have the highest min. group accuracy over the training period. DEO and DEOdds are only evaluated with respect to the attribute “male”.

ment is even more pronounced when deployed on the full training set (see lower half of Table 2). (3) Optimizing min max fairness with the target attribute as an additional protected attribute substantially improves accuracy-based fairness notions, such as DEO.

5.3. Orthogonality to other augmentation methods

We combine the g-SMOTE augmentations with a variety of existing augmentation methods, such as random crop, random rotation, random flip and RandAugment. For each of the existing augmentation methods, the g-SMOTE augmentation improves accuracy. This indicates that the data diversity generated by the GAN yields a generalization gain that is at least partially orthogonal to the otherwise achieved improvements. Results are reported in Table 3.

	No Augment	Rand Crop	Rand Rot.	Rand Flip	RandAugment
Without g-SMOTE	89.15	89.56	89.66	89.78	90.17
With g-SMOTE	89.63	89.85	89.75	89.86	90.33

Table 3. Min. group accuracy over common augmentations on CelebA with and without AdaptiveSMOTE sampling. We report mean scores over the labels [66] call *gender independent*, using the model with greatest min. group accuracy over the training period.

6. Discussion

The majority of fairness methods examined report improved fairness measures of DEO and DEOdds despite decreasing both accuracy and TPR rates of the most disadvantaged groups. Nonetheless, we have shown that it is possible to improve performance for disadvantaged groups by the targeted generation of synthetic data. For the simple case of these high-capacity models (Table 1), in which protected groups correspond to gender, there is no apparent trade-off, and no need to target particular groups with adaptive sampling. Instead, we can uniformly generate samples and improve performance for everyone. In scenarios where we want high accuracy for groups corresponding to a combination of protected attribute and true label (e.g., males receiving a negative decision) this trade-off still exists. Here,

adaptive sampling gives a substantial improvement in min-group accuracy as well as improving the fairness measures.

Based on our theoretical analysis and experimental findings, we make *two key recommendations* for practitioners:

(1) Evaluate a model using the error of the worst performing group: If accuracy-based notions of fairness are appropriate for a given scenario, *i.e.*, particular groups are being disadvantaged by a high error rate, a better measure for fairness than the difference of error rates is the error over the worst performing group. If the rate is sufficiently low, the method can be safely deployed, otherwise not.

(2) Gather more data for the worst performing groups: We have shown that the key issue preventing fairness in computer vision is an inability to generalize well. However, outside of standardized benchmarks with fixed training sets, generalization can be improved by gathering more data. This leads to a direct approach where one iteratively evaluates on held-out data and gathers additional data on the worst performing group. This is equivalent to algorithm 1 of [1], and under certain constraints is guaranteed to minimize the worst error over all groups. Other works [13] have called for more diverse datasets, but approaching this as a min-max problem tells us how to grow the dataset, and why diversity improves performance even when it leads to a distribution mismatch between training and evaluation data.

Limitations: The analysis set out in Section 3 only holds for the accuracy-based notions of fairness that are satisfied by perfect classifiers. It may be possible to enforce other notions of fairness, such as demographic parity, on maximally accurate classifiers without held-out data. Moreover, a wide range of machine learning scenarios exist for which accuracy-based fairness is inappropriate [79].

Conclusion: The takeaway message from this work should not be that accuracy-based notions of fairness do not work in computer vision, but that if we measure and optimize the right thing, they can be made to work. Our analysis in Section 3 makes it clear that fairness on unseen data is primarily a problem of generalization. As such, three promising directions for improving performance in the most disadvantaged groups lie in Hyperparameter Optimization (HO), Network Architecture Search, and data augmentation, with HO already being used in fairness [42, 62, 76]. We have shown that data augmentation can improve performance on the most disadvantaged groups, but the use of these other techniques to improve worst group performance is a promising direction for future work.

Acknowledgments

CR is a member of the Governance of Emerging Technologies (GET) research programme at the Oxford Internet Institute. He is grateful to members of GET and the Trustworthy Auditing for AI project for informed discussion that helped shape this work.

References

- [1] Jacob Abernethy, Pranjal Awasthi, Matthäus Kleindessner, Jamie Morgenstern, Chris Russell, and Jie Zhang. Active sampling for min-max fairness. *arXiv preprint arXiv:2006.06879*, 2021. 1, 2, 3, 8
- [2] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. A reductions approach to fair classification. In *International Conference on Machine Learning (ICML)*, 2018. 1, 2, 3, 4
- [3] Vitor Albiero, Krishnapriya KS, Kushal Vangara, Kai Zhang, Michael C King, and Kevin W Bowyer. Analysis of gender inequality in face recognition accuracy. In *IEEE/CVF Winter Conference on Applications of Computer Vision Workshops*, 2020. 3
- [4] Mohsan Alvi, Andrew Zisserman, and Christoffer Nellaker. Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In *European Conference on Computer Vision (ECCV) - Workshop*, volume 11129 of *Lecture Notes in Computer Science*, pages 556–572. Springer, 2019. 6
- [5] Pranjal Awasthi, Alex Beutel, Matthäus Kleindessner, Jamie Morgenstern, and Xuezhi Wang. Evaluating fairness of machine learning models under uncertain and incomplete information. In *ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT)*, 2021. 1
- [6] Guha Balakrishnan, Yuanjun Xiong, Wei Xia, and Pietro Perona. Towards causal benchmarking of bias in face analysis algorithms. In *Deep Learning-Based Face Analytics*, pages 327–359. Springer, 2021. 3, 6
- [7] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning*. fairmlbook.org, 2019. <http://www.fairmlbook.org>. 1
- [8] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 2018. 1
- [9] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H. Chi, and Cristos Goodrow. Fairness in recommendation ranking through pairwise comparisons. In *ACM International Conference on Knowledge Discovery and Data Mining (KDD)*, 2019. 3
- [10] Alex Beutel, Jilin Chen, Zhe Zhao, and Ed Huai-hsin Chi. Data decisions and theoretical implications when adversarially learning fair representations. *CoRR*, abs/1707.00075, 2017. 2, 3
- [11] Damien Brain and Geoffrey Webb. On the effect of data set size on bias and variance in classification learning. In *Proceedings of the Fourth Australian Knowledge Acquisition Workshop, University of New South Wales*, pages 117–128, 1999. 5
- [12] Campbell Brown. Giving up levelling down. *Economics and Philosophy*, 19(1):111, 2003. 1
- [13] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, 2018. 3, 8
- [14] Toon Calders and Sicco Verwer. Three naive bayes approaches for discrimination-free classification. *Data mining and knowledge discovery*, 21(2):277–292, 2010. 1
- [15] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. Optimized pre-processing for discrimination prevention. In *Neural Information Processing Systems (NeurIPS)*, 2017. 1, 2
- [16] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002. 5
- [17] Nitesh V. Chawla, Aleksandar Lazarevic, Lawrence O. Hall, and Kevin W. Bowyer. Smoteboost: Improving prediction of the minority class in boosting. In Nada Lavrač, Dragan Gamberger, Ljupčo Todorovski, and Hendrik Blockeel, editors, *Knowledge Discovery in Databases: PKDD 2003*, pages 107–119, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. 6
- [18] Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. In *International Conference on Machine Learning (ICML)*, 2020. 3
- [19] Thomas Christiano and Will Braynen. Inequality, injustice and levelling down. *Ratio*, 21(4):392–420, 2008. 1
- [20] Ching-Yao Chuang and Youssef Mroueh. Fair mixup: Fairness via interpolation. In *International Conference on Learning Representations (ICLR)*, 2021. 6
- [21] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 4
- [22] Emily Diana, Wesley Gill, Michael Kearns, Krishnamurthy Kenchappa, and Aaron Roth. Minimax group fairness: Algorithms and experiments. In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*, 2021. 1, 2, 3
- [23] Pedro Domingos. A unified bias-variance decomposition. In *International Conference on Machine Learning (ICML)*, 2000. 3, 4
- [24] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016. 2, 6
- [25] Michele Donini, Luca Oneto, Shai Ben-David, John Shawe-Taylor, and Massimiliano Pontil. Empirical risk minimization under fairness constraints. In *Neural Information Processing Systems (NeurIPS)*, 2018. 2, 4
- [26] Brett Doran. Reconsidering the levelling-down objection against egalitarianism. *Utilitas*, 13(1):65–85, 2001. 1
- [27] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016. 2, 6
- [28] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Rich Zemel. Fairness through awareness. In *Innovations in Theoretical Computer Science Conference (ITCS)*, 2012. 1

- [29] Harrison Edwards and Amos J. Storkey. Censoring representations with an adversary. In *International Conference on Learning Representations (ICLR)*, 2016. 3
- [30] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 3
- [31] Stuart Geman, Elie Bienenstock, and René Doursat. Neural Networks and the Bias/Variance Dilemma. *Neural Computation*, 4(1):1–58, 01 1992. 3
- [32] Partha Ghosh, Dominik Zietlow, Michael J. Black, Larry S. Davis, and Xiaochen Hu. Invgan: Invertible gans, 2021. 6, 13
- [33] Gabriel Goh, Andrew Cotter, Maya Gupta, and Michael Friedlander. Satisfying real-world goals with dataset constraints. In *Neural Information Processing Systems (NeurIPS)*, 2016. 1
- [34] Aditya Grover, Jiaming Song, Alekh Agarwal, Kenneth Tran, Ashish Kapoor, Eric Horvitz, and Stefano Ermon. Bias correction of learned generative models using likelihood-free importance weighting. *arXiv preprint arXiv:1906.09531*, 2019. 3
- [35] Hui Han, Wen-Yuan Wang, and Bing-Huan Mao. Borderline-smote: A new over-sampling method in imbalanced data sets learning. In De-Shuang Huang, Xiao-Ping Zhang, and Guang-Bin Huang, editors, *Advances in Intelligent Computing*, pages 878–887, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. 6
- [36] Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning. In *Neural Information Processing Systems (NeurIPS)*, 2016. 1, 2, 4
- [37] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning — Data Mining, Inference, and Prediction*. Springer, 2nd edition, 2009. 4
- [38] Haibo He, Yang Bai, Edwardo A. Garcia, and Shutao Li. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pages 1322–1328, 2008. 6
- [39] Kaifeng He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 13
- [40] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 771–787, 2018. 3
- [41] Nils Holtug. Egalitarianism and the levelling down objection. *Analysis*, 58(2):166–174, 1998. 1
- [42] Rashidul Islam, Shimei Pan, and James R Foulds. Can we obtain fairness for free? In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 586–596, 2021. 8
- [43] Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 2012. 3
- [44] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3
- [45] Brendan F Klare, Mark J Burge, Joshua C Klontz, Richard W Vorder Bruegge, and Anil K Jain. Face recognition performance: Role of demographic information. *IEEE Transactions on Information Forensics and Security*, 7(6):1789–1801, 2012. 3
- [46] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein. Discrimination in the age of algorithms. *Journal of Legal Analysis*, 10:113–174, 2018. 3
- [47] Matthäus Kleindessner, Samira Samadi, Muhammad Bilal Zafar, Krishnamurthy Suresh, and Chris Russell. Pairwise fairness for ordinal regression. International Conference on Artificial Intelligence and Statistics (AISTATS), 2022. 1, 2, 3
- [48] Ron Kohavi and David Wolpert. Bias plus variance decomposition for zero-one loss functions. In *ICML*, 1996. 3
- [49] Adam Kortylewski, Bernhard Egger, Andreas Schneider, Thomas Gerig, Andreas Morel-Forster, and Thomas Vetter. Empirically analyzing the effect of dataset biases on deep face recognition systems. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018. 3
- [50] Adam Kortylewski, Bernhard Egger, Andreas Schneider, Thomas Gerig, Andreas Morel-Forster, and Thomas Vetter. Analyzing and reducing the damage of dataset bias to face recognition with synthetic data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019. 3, 5
- [51] Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3, 5
- [52] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015. 1, 4
- [53] Michael Lohaus, Michael Perrot, and Ulrike von Luxburg. Too relaxed to be fair. In *International Conference on Machine Learning (ICML)*, 2020. 2, 3
- [54] David Madras, Elliot Creager, Toniann Pitassi, and Richard S. Zemel. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning (ICML)*, 2018. 3
- [55] Natalia Martinez, Martin Bertran, and Guillermo Sapiro. Minimax Pareto fairness: A multi objective perspective. In *International Conference on Machine Learning (ICML)*, 2020. 3
- [56] Natalia Martinez, Martin Bertrana, and Guillermo Sapiro. Minimax pareto fairness: A multi objective perspective. In *International Conference on Machine Learning (ICML)*, 2020. 1, 2
- [57] Andrew Mason. Egalitarianism and the levelling down objection. *Analysis*, 61(3):246–254, 2001. 1

- [58] Michele Merler, Nalini Ratha, Rogerio S Feris, and John R Smith. Diversity in faces. *arXiv preprint arXiv:1901.10436*, 2019. 3
- [59] Curtis G. Northcutt, Anish Athalye, and Jonas Mueller. Pervasive label errors in test sets destabilize machine learning benchmarks. 2021. *arXiv preprint arXiv:2103.14749*. 4
- [60] Robert Osada, Thomas Funkhouser, Bernard Chazelle, and David Dobkin. Shape distributions. *ACM Transactions on Graphics (TOG)*, 21(4):807–832, 2002. 13
- [61] Manisha Padala and Sujit Gujar. Fnnnc: Achieving fairness through neural networks. In *International Joint Conferences on Artificial Intelligence Organization (IJCAI)*, 2020. 2, 3, 5, 6
- [62] Valerio Perrone, Michele Donini, Muhammad Bilal Zafar, Robin Schmucker, Krishnaram Kenthapadi, and Cédric Archambeau. Fair bayesian optimization. In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*, 2021. 8
- [63] Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6
- [64] Jean Ponce, Tamara L Berg, Mark Everingham, David A Forsyth, Martial Hebert, Svetlana Lazebnik, Marcin Marszałek, Cordelia Schmid, Bryan C Russell, Antonio Torralba, et al. Dataset issues in object recognition. In *Toward category-level object recognition*, pages 29–48. Springer, 2006. 3
- [65] Novi Quadrianto, Viktoriia Sharmanska, and Oliver Thomas. Discovering fair representations in the data domain. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 3
- [66] Vikram V. Ramaswamy, Sunnie S. Y. Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 3, 5, 6, 7, 8, 13, 14
- [67] Suman Ravuri and Oriol Vinyals. Seeing is not necessarily believing: Limitations of biggans for data augmentation. In *ICLR 2019 Workshop LLD*, 2019. 5
- [68] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision (ECCV)*, 2016. 5
- [69] Laurent Risser, Quentin Vincenot, and Jean-Michel Loubes. Tackling algorithmic bias in neural-network classifiers using wasserstein-2 regularization, 2020. 3
- [70] Amélie Royer and Christoph H. Lampert. Classifier adaptation at prediction time. 2015. 5, 6
- [71] Marco Saerens, Patrice Latinne, and Christine Decaestecker. Adjusting the outputs of a classifier to new a priori probabilities: A simple procedure. *Neural computation*, 14:21–41, 02 2002. 5, 6
- [72] Prasanna Sattigeri, Samuel C Hoffman, Vijil Chenthamarakshan, and Kush R Varshney. Fairness gan: Generating datasets with fairness properties using a generative adversarial network. *IBM Journal of Research and Development*, 63(4/5):3–1, 2019. 3
- [73] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of GANs for semantic face editing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6
- [74] Chandan Singh, Guha Balakrishnan, and Pietro Perona. Matched sample selection with gans for mitigating attribute confounding. *arXiv preprint arXiv:2103.13455*, 2021. 3
- [75] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. 3
- [76] Ana Valdivia, Javier Sánchez-Monedero, and Jorge Casillas. How fair can we go in machine learning? assessing the boundaries of accuracy and fairness. *International Journal of Intelligent Systems*, 36(4):1619–1643, 2021. 1, 8
- [77] Vladimir Vapnik. *The nature of statistical learning theory*. Springer, 1995. 3
- [78] Sahil Verma and Julia Rubin. Fairness definitions explained. In *International Workshop on Software Fairness (FairWare)*, 2018. 2
- [79] Sandra Wachter, Brent Mittelstadt, and Chris Russell. Bias preservation in machine learning: the legality of fairness metrics under eu non-discrimination law. *W. Va. L. Rev.*, 123:735, 2020. 2, 3, 8
- [80] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 692–702, 2019. 3
- [81] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5310–5319, 2019. 2, 3
- [82] Zeyu Wang, Klint Qinami, Ioannis Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3, 4, 6, 7, 13
- [83] Michael Wick, Swetasudha Panda, and Jean-Baptiste Tristan. Unlocking fairness: a trade-off revisited. In *Neural Information Processing Systems (NeurIPS)*, 2019. 2, 3, 5, 6, 14
- [84] Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. 2020. *arXiv preprint arXiv:2002.11328*. 4
- [85] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rogriguez, and Krishna P Gummadi. Fairness constraints: Mechanisms for fair classification. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017. 1, 2, 3, 4
- [86] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *International Conference on Machine Learning (ICML)*, 2013. 3
- [87] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning re-

quires rethinking generalization. In *International Conference on Learning Representations (ICLR)*, 2017. [2](#), [3](#)

- [88] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations (ICLR)*, 2018. [6](#)