

Discover and Mitigate Unknown Biases with Debiasing Alternate Networks

Zhiheng Li¹, Anthony Hoogs², and Chenliang Xu¹

¹University of Rochester ²Kitware, Inc.

{zhiheng.li, chenliang.xu}@rochester.edu anthony.hoogs@kitware.com

Abstract. Deep image classifiers have been found to learn biases from datasets. To mitigate the biases, most previous methods require labels of protected attributes (e.g., age, skin tone) as full-supervision, which has two limitations: 1) it is infeasible when the labels are unavailable; 2) they are incapable of mitigating unknown biases—biases that humans do not preconceive. To resolve those problems, we propose Debiasing Alternate Networks (DebiAN), which comprises two networks—a Discoverer and a Classifier. By training in an alternate manner, the discoverer tries to find multiple unknown biases of the classifier without any annotations of biases, and the classifier aims at unlearning the biases identified by the discoverer. While previous works evaluate debiasing results in terms of a single bias, we create Multi-Color MNIST dataset to better benchmark mitigation of multiple biases in a multi-bias setting, which not only reveals the problems in previous methods but also demonstrates the advantage of DebiAN in identifying and mitigating multiple biases simultaneously. We further conduct extensive experiments on real-world datasets, showing that the discoverer in DebiAN can identify unknown biases that may be hard to be found by humans. Regarding debiasing, DebiAN achieves strong bias mitigation performance.

Keywords: Bias Identification, Bias Mitigation, Fairness, Unsupervised Debiasing

1 Introduction

Many studies have verified that AI algorithms learn undesirable biases from the dataset. Some biases provide shortcuts [21] for the network to learn superficial features instead of the intended decision rule causing robustness issues, *e.g.*, static cues for action recognition [7, 14, 45]. Other biases make AI algorithms discriminate against different protected demographic groups such as genders* [3, 29–31, 69, 72, 77] and skin tones [12, 27], leading to serious fairness problems. Therefore, it is imperative to mitigate the biases in AI algorithms. However, most previous bias mitigation methods [4, 65, 71, 75, 77] are supervised methods—requiring annotations of the biases, which has several limitations: First, bias mitigation cannot be performed when labels are not available due to privacy

*In this work, “gender” denotes visually perceived gender, not real gender identity.

concerns. Second, they cannot mitigate *unknown* biases—biases that humans did not preconceive, making the biases impossible to be labeled and mitigated.

Since supervised debiasing methods present many disadvantages, in this work, we focus on a more challenging task—unsupervised debiasing, which mitigates the *unknown* biases in a learned classifier without any annotations. Without loss of generality, we focus on mitigating biases in image classifiers. Solving this problem contains two steps [2,17,40,53,63]: bias identification and bias mitigation.

Due to the absence of bias annotations, the first step is to assign the training samples into different bias groups as the pseudo bias labels, which is challenging since the biases are even unknown. The crux of the problem is to define the unknown bias. Some previous works make strong assumptions about the unknown biases based on empirical observations, such as biases are easier to be learned [53], samples from the same bias group are clustered in feature space [63], which can be tenuous for different datasets or networks. Other works quantify the unknown biases by inversely using the debiasing objective functions [2,17], which can face numerical or convergence problems (more details in Sec. 2). Unlike previous works, we follow an axiomatic principle to define the unknown biases—classifier’s predictions that violate a fairness criterion [16,20,24,26,39,55,66]. Based on this definition, we propose a novel *Equal Opportunity Violation* (EOV) loss to train a *discoverer* network to identify the classifier’s biases. In specific, it shepherds the *discoverer* network to predict bias group assignments such that the *classifier* violates the Equal Opportunity [26,55] fairness criterion (Figs. 1, 2).

Regarding debiasing as the second step, most previous approaches [2,17,63] preprocess the identified biases into pseudo bias labels and resort to other supervised bias mitigation methods [6,57] for debiasing. In contrast, we propose a novel *Reweighted Cross-Entropy* (RCE) loss that leverages soft bias group assignments predicted by the *discoverer* network to mitigate the biases in the *classifier* (Fig. 1). In this way, the *classifier* is guided to meet the Equal Opportunity.

In addition, many previous works [2,17,63] treat bias identification and bias mitigation as two isolated steps. In [2,17], the biases are identified from an undertrained classifier, which is suboptimal since the classifier may learn different biases at different training stages. Consequently, these two-stage methods fail to mitigate other biases learned by the classifier at later training stages. In contrast, we employ an alternate training scheme to carry out bias identification and bias mitigation simultaneously. We jointly update the *discoverer* and *classifier* in an interleaving fashion (Figs. 1 and 2). In this way, the *discoverer* can repetitively inspect multiple biases that the *classifier* learns at the entire training stage.




Fig. 1: DEBIASING ALTERNATE NETWORKS (DebiAN). We alternately train two networks—a *discoverer* and a *classifier*. *Discoverer* actively identifies *classifier*’s unknown biases. At the same time, the *classifier* mitigates the biases identified by the *discoverer*

We integrate our novel losses and training scheme into a unified framework—DEBIASING ALTERNATE NETWORKS (DebiAN), which contains two networks—a *discoverer* D and a *classifier* C (see Fig. 1). We jointly train the two networks in an alternate manner. Supervised by our novel EOv loss, D tries to discover C ’s multiple unknown biases that violate the Equal Opportunity fairness criterion. Trained with our RCE loss, C aims at mitigating multiple biases identified by the *discoverer* D to satisfy Equal Opportunity. After the alternate training, the unknown biases in *classifier* C are mitigated, leading to a fairer and more robust classification model. Besides, when employed with other network explanation methods [59,60,78], the *discoverer* is helpful to interpret the discovered unknown biases, facilitating dataset curators to locate dataset biases [64].

While previous works [7,35,46,53,57] only evaluate debiasing results in terms of a single bias, we create Multi-Color MNIST dataset with two biases in the dataset, which benchmarks debiasing algorithms in the multi-bias setting. Our new dataset surfaces the problems in previous methods (*e.g.*, LfF [53]) and demonstrates the advantage of DebiAN in discovering and mitigating multiple biases. We further conduct extensive experiments to verify the efficacy of DebiAN in real-world image datasets. In the face image domain, DebiAN achieves better gender bias mitigation results on CelebA [50] and bFFHQ [36] datasets. On the gender classification task, DebiAN achieves better debiasing results on CelebA w.r.t. multiple bias attributes. We further show an interesting unknown bias discovered by DebiAN in gender classification—**visible hair area**. Lastly, we show that DebiAN applies to other image domains for broader tasks, such as action recognition and scene classification. Our method not only achieves better debiasing results, but also identifies interesting unknown biases in scene classifiers.

Our contributions are summarized as follows: (1) We propose a novel objective function, *Equal Opportunity Violation* (EOv) loss, for identifying unknown biases of a classifier based on Equal Opportunity. (2) We propose a *Reweighted Cross-Entropy* (RCE) loss to mitigate the discovered unknown biases by leveraging the soft bias group assignments. (3) We create Multi-Color MNIST dataset to benchmark debiasing algorithms in a multi-bias setting. (4) Our DEBIASING ALTERNATE NETWORKS (DebiAN) outperforms previous unsupervised debiasing methods on both synthetic and real-world datasets.

2 Related Work

Bias Identification Most previous works identify *known* biases based on bias labels. In [12], face images are labeled with gender and skin tone to identify the performance gaps across intersectional groups. Balakrishnan *et al.* [8] further synthesize intersectional groups of images and analyze the biases with additional labels. Beyond face images, recent works [51,67] compute the statistics of labels based on the rule mining algorithm [1] or external tools. [38] uses clustering on image embeddings to discover unknown biases. [41,47] discovers *unknown* biases without labels. However, these works rely on GAN [23,33] to synthesize images,

which suffers from image quality issues. In contrast, DebiAN directly classifies real images into different bias attribute groups to discover the unknown biases.

Supervised Debiasing Supervised debiasing methods use bias labels for debiasing. [32] proposes a supervised reweighing method. Wang *et al.* [73] benchmark recent supervised debiasing methods [4,65,75,77]. [18] lets the model be flexibly fair to different attributes during testing. [58] uses disentanglement for debiasing. Singh *et al.* [62] propose a feature splitting approach to mitigate contextual bias. [19,22] use adversarial training to mitigate biases in face recognition.

Known Bias Mitigation with Prior knowledge Without using labels, some works use prior knowledge to mitigate certain known biases. ReBias [7] uses model capacity as the inductive bias to mitigate texture bias and static bias in image and video classification. HEX [68] introduces a texture extractor to mitigate the texture bias. Beyond image classification, RUBi [13] and LearnedMixIn [15] mitigate unimodal bias for visual question answering [5] with prior knowledge.

Unsupervised Debiasing In the field of mitigating unknown biases, Sohoni *et al.* [63] apply clustering on samples in each class and use the clustering assignment as the predicted bias labels, which could be inaccurate due to its unsupervised nature. Li *et al.* [44,45] fix the parameters of feature extractors and focus on mitigating the representation bias. PI [11] uses labels of one known bias attribute to interpolate environments, improving robustness against multiple unknown biases. TOFU [10] infers unstable features from a source task and transfers them to the target task. Concurrent to our work, Bao and Barzilay [9] propose learning to split for detecting unknown biases, which can then be combined with GroupDRO [57] for debiasing. LfF [53] identifies biases by finding easier samples in the training data through training a bias-amplified network supervised by GCE loss [76], which up-weights the samples with smaller loss values and down-weights the samples with larger loss values. In other words, GCE loss does not consider the information of the classifier, *e.g.*, the classifier’s output. Therefore, LfF’s bias-amplified network blindly finds the biases in the data samples instead of the classifier. Unlike LfF, the EOv loss in DebiAN actively identifies biases in the classifier based on the classifier’s predictions, leading to better debiasing performance. Following LfF, BiaSwap [36] uses LfF to discover biases and generate more underrepresented images via style-transfer for training. Other works [2,17,40,70] inversely use the debiasing objective function to maximize an unbounded loss (*e.g.*, gradient norm penalty in IRMv1 [6]) for bias identification, which may encounter numerical or convergence problems. As a comparison, our EOv loss (Eq. (2)) minimizes negative log-likelihood, which is numerically stable and easier to converge.

3 Method

Overview The overview of our proposed DEBIASING ALTERNATE NETWORKS (DebiAN) is shown in Fig. 2. It contains two networks—a *discoverer* D and a *classifier* C . As shown in Fig. 2 (a), the *discoverer* D tries to discover the unknown biases in the *classifier* C by optimizing our proposed EOv loss (\mathcal{L}_{EOv})

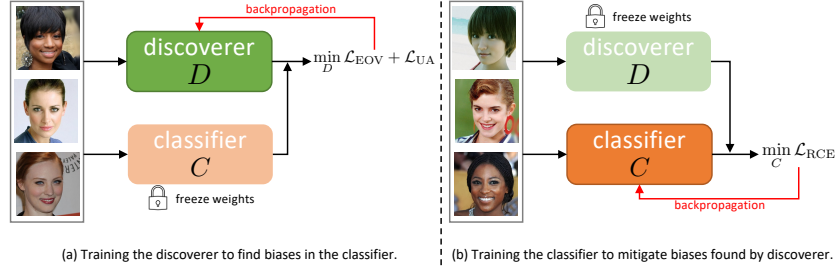


Fig. 2: Overview of DEBIASING ALTERNATE NETWORKS (DebiAN). DebiAN consists of two networks—a *discoverer* D and a *classifier* C . D is trained with \mathcal{L}_{EOV} and \mathcal{L}_{UA} (Sec. 3.1) to find the unknown biases in C . C is optimized with \mathcal{L}_{RCE} (Sec. 3.2) to mitigate the biases identified by D

and UA penalty (\mathcal{L}_{UA}) (Sec. 3.1). As shown in Fig. 2 (b), the *classifier* C ’s goal is to mitigate the biases identified by D via a novel *Reweighted Cross-Entropy* loss (\mathcal{L}_{RCE}) (Sec. 3.2). Lastly, we train the two networks in an alternate manner as the full model for discovering and mitigating the unknown biases (Sec. 3.3).

Background To better explain our motivation for discovering the *unknown* biases (without manual annotations of biases), let us first revisit the traditional approach for identifying *known* biases when labels of biases (*e.g.*, protected attributes) are available, which is illustrated in Fig. 3 (a). The following are given for identifying *known* biases—a well-trained *classifier* C for predicting a target attribute, n testing images $\{\mathbf{I}_i\}_{i=1}^n$, target attribute labels of each image $\{y_i\}_{i=1}^n$, and bias attribute labels $\{b_i\}_{i=1}^n$. We denote the i -th image target attribute as $y_i \in \{1, 2, \dots, K\}$ and K is the number of classes. We consider the bias attribute that is binary or continuously valued (*i.e.*, $b_i \in \{0, 1\}$ or $b_i \in [0, 1]$), such as biological gender (*e.g.*, female and male) and skin tones (*e.g.*, from dark skin tones to light skin tones in Fitzpatrick skin type scale). We leave bias attributes with multi-class values for future works. Then, the given *classifier* C is tested for predicting the target attribute \hat{y}_i for each testing image \mathbf{I}_i . Finally, we check whether the predictions meet a fairness criterion, such as Equal Opportunity [26]:

$$\Pr\{\hat{y} = k \mid b = 0, y = k\} = \Pr\{\hat{y} = k \mid b = 1, y = k\}, \quad (1)$$

where the LHS and RHS are true positive rates (TPR) in negative ($b = 0$) and positive ($b = 1$) bias attribute groups, respectively. $k \in \{1 \dots K\}$ is a target attribute class. Equal Opportunity requires the same TPR across two different bias attribute groups. That is, if the TPR is significantly different in two groups of the bias attribute, we conclude that *classifier* C contains the bias of attribute b because C violates the Equal Opportunity fairness criterion. For example, as shown in Fig. 3 (a), although all images are female, a gender classifier may have a larger TPR for the group of long-hair female images than the group of short-hair female images. Thus the gender classifier is biased against different hair lengths.

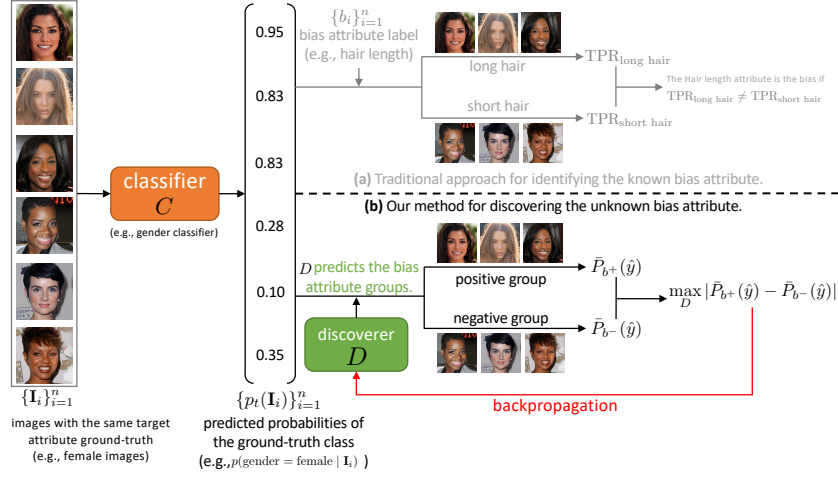


Fig. 3: (a): The traditional approach for identifying the *known* bias attribute (e.g., hair length) by comparing true positive rates (TPR) of the target attribute (e.g., gender) in two groups of bias attributes (e.g., long hair and short hair), where the group assignment of bias attribute is based on the labels of the bias attribute. **(b):** Our method trains a *discoverer* (D) to predict the groups of the unknown bias attribute such that the difference of averaged predicted probabilities on the target attribute (e.g., gender) in two groups are maximized (see Eq. 2)

3.1 Unknown Bias Discovery

As for identifying unknown biases, we do not have the labels to assign images into two groups for comparing TPR since 1) we do not assume images come with bias attribute labels, and 2) the type of bias is even unknown. However, we can compare the difference in TPR for any group assignments based on speculated biases—a significant difference in TPR hints that the Equal Opportunity fairness criterion is most violated (our method mainly focuses on the Equal Opportunity fairness criterion, and we leave other fairness criteria for future work). Motivated by this finding, instead of using labels of bias attribute $\{b_i\}_{i=1}^n$ for group assignment, we train a *discoverer* D to predict the group assignment for each image, i.e., $p(\hat{b} | \mathbf{I}_i) := D(\mathbf{I}_i)$. By optimizing loss functions below, we find the most salient bias of the *classifier* C that violates the Equal Opportunity fairness criterion, which is illustrated in Fig. 3 (b).

Equal Opportunity Violation (EOV) Loss To shepherd the *discoverer* D to find the group assignment where *classifier* C violates the Equal Opportunity fairness criterion, we propose the *Equal Opportunity Violation* (EOV) loss, denoted by \mathcal{L}_{EOV} , as the objective function to train D . For computing \mathcal{L}_{EOV} , we sample a set of n images $\{\mathbf{I}_i\}_{i=1}^n$ with the *same* target attribute labels (i.e., $\forall i, y_i = k$), e.g., all images in Fig. 3 (b) are female. The *classifier* C has been trained for predicting the target attribute y of the images (i.e., $p(\hat{y} | \mathbf{I}_i) := C(\mathbf{I}_i)$). For simplicity, we denote p_t as C 's prediction on images of the ground-truth class

(i.e., $p_t(\mathbf{I}_i) = p(\hat{y} = y_i \mid \mathbf{I}_i)$). Meanwhile, the same set of images $\{\mathbf{I}_i\}$ are fed to the *discoverer* D for predicting the binary bias attribute group assignment: $p(\hat{b} \mid \mathbf{I}_i) := D(\mathbf{I}_i)$. Finally, we define the EOV loss as:

$$\mathcal{L}_{\text{EOV}} = -\log \left(\left| \bar{P}_{b+}(\hat{y}) - \bar{P}_{b-}(\hat{y}) \right| \right), \quad (2)$$

where $\bar{P}_{b+}(\hat{y})$ and $\bar{P}_{b-}(\hat{y})$ are defined by:

$$\begin{aligned} \bar{P}_{b+}(\hat{y}) &= \frac{\sum_{i=1}^n p(\hat{b} = 1 \mid \mathbf{I}_i) p_t(\mathbf{I}_i)}{\sum_{i=1}^n p(\hat{b} = 1 \mid \mathbf{I}_i)}, \\ \bar{P}_{b-}(\hat{y}) &= \frac{\sum_{i=1}^n p(\hat{b} = 0 \mid \mathbf{I}_i) p_t(\mathbf{I}_i)}{\sum_{i=1}^n p(\hat{b} = 0 \mid \mathbf{I}_i)}. \end{aligned} \quad (3)$$

Intuitively, $\bar{P}_{b+}(\hat{y})$ and $\bar{P}_{b-}(\hat{y})$ are the weighted average predicted probabilities of the target attribute in two bias attribute groups, which can be regarded as a relaxation to Equal Opportunity’s true positive rate (Eq. 1) where the predicted probabilities are binarized into predictions with a threshold (e.g., 0.5). Minimizing \mathcal{L}_{EOV} leads D to maximize the discrepancy of averaged predicted probabilities of target attributes in two bias attribute groups (i.e., see $\max_D |\bar{P}_{b+}(\hat{y}) - \bar{P}_{b-}(\hat{y})|$ in Fig. 3), thus finding the bias attribute group assignments where C violates the Equal Opportunity fairness criterion. For example, in Fig. 3 (b), if the gender classifier C is biased against different hair lengths, then by optimizing \mathcal{L}_{EOV} , D can assign the female images into two bias attribute groups (i.e., short hair and long hair) with the predicted bias attribute group assignment probability $p(\hat{b} \mid \mathbf{I}_i)$, such that the difference of averaged predicted probabilities on gender in these two groups is maximized.

Unbalanced Assignment (UA) penalty However, we find that optimizing \mathcal{L}_{EOV} alone may let the *discoverer* D find a trivial solution—assigning all images into one bias attribute group. For example, suppose D assigns all images to the positive bias attribute group (i.e., $\forall_i, p(\hat{b} = 1 \mid \mathbf{I}_i) = 1$). In that case, $\bar{P}_{b-}(\hat{y})$ becomes zero since the negative group contains no images. $\bar{P}_{b+}(\hat{y})$ becomes a large positive number by simply averaging $p_t(\mathbf{I}_i)$ for all of the n images, which can trivially increase $|\bar{P}_{b+}(\hat{y}) - \bar{P}_{b-}(\hat{y})|$, leading to a small \mathcal{L}_{EOV} . To prevent this trivial solution, we propose the *Unbalanced Assignment* (UA) loss denoted by:

$$\mathcal{L}_{\text{UA}} = -\log \left(1 - \frac{1}{n} \left| \sum_{i=1}^n p(\hat{b} = 1 \mid \mathbf{I}_i) - p(\hat{b} = 0 \mid \mathbf{I}_i) \right| \right). \quad (4)$$

Intuitively, minimizing \mathcal{L}_{UA} penalizes the unbalanced assignment that leads to large difference between $\sum_{i=1}^n p(\hat{b} = 1 \mid \mathbf{I}_i)$ and $\sum_{i=1}^n p(\hat{b} = 0 \mid \mathbf{I}_i)$, which can be regarded as the numbers of images assigned into positive and negative bias attribute groups, respectively. Therefore, \mathcal{L}_{EOV} is jointly optimized with \mathcal{L}_{UA} to prevent the trivial solution. We acknowledge a limitation of the UA penalty. Although it resolves the trivial solution, it introduces a trade-off since the bias attribute groups are usually spuriously correlated with the target attribute (e.g., more long-hair females than the short-hair females in the dataset). Hence

encouraging balanced assignments may make the *discoverer* harder to find the correct assignment. However, our ablation study shows that the benefits of using \mathcal{L}_{UA} outweigh its limitations. The results are shown in Sec. 4.1 and Tab. 1.

3.2 Unknown Bias Mitigation by Reweighing

We further mitigate C 's unknown biases identified by D . To this end, we propose a novel *Reweighted Cross-Entropy* loss that adjusts the weight of each image's classification loss. Based on the bias attribute group assignment $p(\hat{b} \mid \mathbf{I}_i)$ predicted by D , we define the weight $\mathcal{W}(\mathbf{I}_i)$ of classification loss for each image \mathbf{I}_i as:

$$\begin{aligned} \mathcal{W}(\mathbf{I}_i) = & \mathbb{1} [\bar{P}_{b+}(\hat{y}) \geq \bar{P}_{b-}(\hat{y})] p(\hat{b} = 0 \mid \mathbf{I}_i) \\ & + \mathbb{1} [\bar{P}_{b+}(\hat{y}) < \bar{P}_{b-}(\hat{y})] p(\hat{b} = 1 \mid \mathbf{I}_i), \end{aligned} \quad (5)$$

where $\mathbb{1}$ is an indicator function. Then, the *Reweighted Cross-Entropy* loss (\mathcal{L}_{RCE}) is defined by:

$$\mathcal{L}_{\text{RCE}} = -\frac{1}{n} \sum_{i=1}^n (1 + \mathcal{W}(\mathbf{I}_i)) \log p_t(\mathbf{I}_i). \quad (6)$$

For example, when C performs better on images from the positive bias attribute group (*i.e.*, $\bar{P}_{b+}(\hat{y}) \geq \bar{P}_{b-}(\hat{y})$), we use $p(\hat{b} = 0 \mid \mathbf{I}_i)$ as the weight, which up-weights the images from the negative bias attribute group, where classifier C is worse-performed. At the same time, it down-weights the images from the positive bias attribute group where C is already better-performed. Adding one to the weight in Eq. (6) lets the loss function degenerate to standard cross-entropy loss when $\mathcal{W}(\mathbf{I}_i) = 0$. By minimizing the *Reweighted Cross-Entropy* loss, C is guided to meet Equal Opportunity.

3.3 Full Model

We summarize the proposed losses in Sec. 3.1 and Sec. 3.2 for the full model of DEBIASING ALTERNATE NETWORKS (DebiAN), which is shown in Fig. 2. When the task is to only discover (*i.e.*, not mitigate) the unknown biases of a given classifier, the classifier's parameters are fixed and we only train the *discoverer* D by minimizing \mathcal{L}_{EOV} (Eq. 2) and \mathcal{L}_{UA} (Eq. 4) on the classifier's training data. When the task is to mitigate the unknown biases, we jointly train two networks in an alternate fashion:

$$\min_D \mathcal{L}_{\text{EOV}} + \mathcal{L}_{\text{UA}}, \quad (7)$$

$$\min_C \mathcal{L}_{\text{RCE}}. \quad (8)$$

In Eq. 7, C 's parameters are fixed, and D is optimized to identify C 's unknown biases where C violates the Equal Opportunity. Through Eq. 8, C is optimized for mitigating the unknown biases discovered by D to satisfy the Equal Opportunity while D 's parameters are frozen. After the alternate training, C 's unknown biases identified by D are mitigated, leading to a fairer and more robust *classifier*. The pseudocode of the complete algorithm is in Appendix A.

4 Experiment

We conduct extensive experiments to verify the efficacy of DebiAN. First, we evaluate the results on our newly created Multi-Color MNIST dataset (Sec. 4.1) in a multi-bias setting. We further conduct experiments on real-world datasets in multiple image domains—face (Sec. 4.2) and other image domains (*e.g.*, scene, action recognition) (Sec. 4.3). More details (*e.g.*, evaluation metrics) are introduced in each subsection. The code and our newly created Multi-Color MNIST dataset are released at <https://github.com/zhihengli-UR/DebiAN>.

Comparison Methods We mainly compare with three unsupervised debiasing methods: 1) LfF [53] uses Generalized Cross-entropy (GCE) loss [76] to train a “biased model” for reweighing the classifier; 2) EIIL [17] identifies the bias groups by optimizing bias group assignment to maximize the IRMv1 [6] objective function. The identified bias groups will serve as pseudo bias labels for other supervised debiasing methods to mitigate the biases. Following [17], IRM [6] is used as the debiasing algorithm for EIIL. 3) PGI [2] follows EIIL to identify the biases by training a small multi-layer perceptron for bias label predictions. Concerning debiasing, PGI minimizes the KL-divergence of the classifier’s predictions across different bias groups. We use the officially released code of LfF, EIIL, and PGI in our experiment. Besides, we also compare with vanilla models, which do not have any debiasing techniques (*i.e.*, only using standard cross-entropy loss for training). On bFFHQ [36] and BAR [53] datasets, we also compare with BiaSwap [36], which follows LfF to identify unknown biases, and then uses style-transfer to generate more underrepresented images for training. Since its code has not been released, we cannot compare DebiAN with BiaSwap on other datasets. All results shown below are the mean results over three random seeds of runs, and we also report the standard deviation as the error bar.

4.1 Experiment on Multi-Color MNIST

Many previous works use synthetic datasets to benchmark bias mitigation performance. For example, Colored MNIST [7,35,46] adds color bias to the original MNIST [42] dataset, where each digit class is spuriously correlated with color (see Fig. 4 (a)). We compare DebiAN with other methods on the Colored MNIST dataset in Appendix E. However, we believe that the single-bias setting is an oversimplification of the real-world scenario where multiple biases may exist. For instance, Lang *et al.* [41] find that gender classifiers are biased with multiple independent bias attributes, including wearing lipsticks, eyebrow thickness, nose width, *etc.* The benchmarking results on such a single-bias synthetic dataset may not help us to design better debiasing algorithms for real-world usage.

To this end, we propose **Multi-Color MNIST** dataset to benchmark debiasing methods under the multi-bias setting. In the training set, each digit class is spuriously correlated with *two* bias attributes—**left color** and **right color** (Fig. 4 (b)). Following the terms used in LfF [53], we call samples that can be correctly predicted with the bias attribute as bias-aligned samples. Samples that cannot be correctly predicted with the bias attribute are called bias-conflicting



Fig. 4: Comparison between (a) previous Colored MNIST [7, 35, 46] with a single color bias and (b) our new Multi-Color MNIST dataset that contains *two* bias attributes—**left color** and **right color**




Fig. 5: Evaluating bias discovery w.r.t. **left color**, **right color** biases throughout the training epochs on Multi-Color MNIST. LfF only finds the more salient **left color** bias (ratio=0.99), whereas DebiAN’s *discoverer* finds both biases at the early training stage. Then accuracies gradually converge to 50% as debiasing is performed in the *classifier*, making the *discoverer* harder to find biases

samples. For example, if most digit “0” images are in red **left color** in the training set, we call them bias-aligned samples w.r.t. **left color** attribute, and we regard digit “0” images in a different **left color** (e.g., yellow) as bias-conflicting samples. Since the dataset contains two bias attributes, there exist images that are bias-aligned w.r.t. to **left color** and bias-conflicting w.r.t. **right color** simultaneously, or vice versa. Following [53], we use the *ratio* of the bias-aligned samples for each bias attribute to indicate how strong the spurious correlation is in the training set. The two ratios for two bias attributes can be different, which is more common in the real-world scenario. The images in the testing set also contain two background colors, but the testing set has a balanced distribution of bias-aligned and bias-conflicting samples w.r.t. each bias attribute.

Evaluation Metrics and Settings Following [53], we report the accuracy results in bias-aligned and bias-conflicting samples on the testing set. Since Multi-Color MNIST contains two bias attributes, we report the four accuracy results in the combination of (bias-aligned, bias-conflicting) \times (**left color**, **right color**), e.g., middle four rows in Tab. 1 for each method. We also report the unbiased accuracy, which averages the four results above. Here, we choose 0.99 as the ratio of bias-aligned samples w.r.t. **left color** and 0.95 as the ratio of bias-aligned samples w.r.t. **right color**. In this way, the **left color** is a *more salient* bias than the **right color**. We report the results of other ratio combinations in Appendix C.3. We strictly use the same set of hyperparameters used in (single) Colored MNIST in LfF. More details are in Appendix B.

Debiasing Results on Multi-Color MNIST The debiasing results are shown in Tab. 1. Except for LfF, all other methods achieve higher accuracy results on **left color** bias-aligned samples (1st and 2nd rows) than **right color** bias-aligned samples (1st and 3rd rows), indicating that most methods are more biased w.r.t. the more salient bias, i.e., **left color** (ratio=0.99) in the

Table 1: Debiasing results on Multi-Color MNIST dataset. The accuracy results in the four combinations of two bias attributes, (*i.e.*, **left color** and **right color**) and (bias-aligned and bias-conflicting) are reported. Unbiased accuracy averages the results over all four combinations. We bold top-2 results and underline lowest results

left color ratio = 0.99	right color ratio = 0.95	vanilla	LfF	EiIL	PGI	w/o \mathcal{L}_{UA} (Ours)	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 \pm 0.0	99.6 \pm 0.5	100.0 \pm 0.0	98.6 \pm 2.3	100.0 \pm 0.0	100.0 \pm 0.0
bias-aligned	bias-conflicting	97.1 \pm 0.5	<u>4.7</u> \pm 0.5	97.2 \pm 1.5	82.6 \pm 19.6	97.2 \pm 0.5	95.6 \pm 0.8
bias-conflicting	bias-aligned	27.5 \pm 3.6	98.6 \pm 0.4	70.8 \pm 4.9	<u>26.6</u> \pm 5.5	71.6 \pm 0.7	76.5 \pm 0.7
bias-conflicting	bias-conflicting	5.2 \pm 0.4	<u>5.1</u> \pm 0.4	10.9 \pm 0.8	9.5 \pm 3.2	13.8 \pm 1.1	16.0 \pm 1.8
unbiased accuracy		57.4 \pm 0.7	<u>52.0</u> \pm 0.1	69.7 \pm 1.0	54.3 \pm 4.0	70.6 \pm 0.3	72.0 \pm 0.8

multi-bias setting. Unlike all other methods, LfF gives abnormal results—high accuracy results (*e.g.*, 99.6, 98.6) for the **right color** bias-aligned samples and low accuracy results (*e.g.*, 4.7, 5.1) for the **right color** bias-conflicting samples. Consequently, LfF achieves the worst unbiased accuracy (52.0). The results indicate that LfF only mitigates the more salient **left color** bias, rendering the classifier to learn the less salient **right color** bias (ratio=0.95). Compared with all other methods, DebiAN achieves better unbiased accuracy results (72.0). More importantly, DebiAN achieves much better debiasing result (16.0) in bias-conflicting samples w.r.t. both **left color** and **right color** attributes, where neither color can provide the shortcut for the correct digit class prediction, demonstrating better debiasing results of DebiAN for mitigating multiple biases simultaneously in the multi-bias setting, which is closer to the real-world scenarios.

Bias Discovery: LfF vs. DebiAN We further evaluate the bias discovery results throughout the entire training epochs, which helps us better understand LfF’s abnormal results and DebiAN’s advantages. We use LfF’s “biased model” and DebiAN’s *discoverer* to predict if a given image is bias-aligned or bias-conflicting w.r.t. a bias attribute (*i.e.*, binary classification, more details in Appendix D.1). We show the accuracy results of bias discovery w.r.t. each bias attribute at the end of each epoch in Fig. 5, which shows that LfF only discovers the more salient **left color** bias attribute (100% accuracy), but completely ignores the less salient **right color** bias (50% accuracy) throughout the entire training stage. It reveals the problem of LfF’s definition of the unknown bias—an attribute in the dataset that is easier, which only holds in the single-bias setting but does not generalize to the multi-bias setting. In contrast, DebiAN uses the principled definition to define the bias—classifier’s predictions that violate equal opportunity, enabling *discoverer* to find both biases accurately at the beginning (it achieves about 60% to 70% accuracy because debiasing is simultaneously performed before the end of the first epoch). At the same time, DebiAN’s alternate training scheme lets the classifier mitigate both biases, making the *discoverer* harder to predict the biases, *e.g.*, accuracies of both bias attributes gradually converge to 50%. More discussions are in Appendix D.4.

Ablation Study on UA penalty We conduct an ablation study to show the effectiveness of Unbalanced Assignment (UA) penalty (Sec. 3.1). Tab. 1 shows

Table 2: Results of mitigating the gender bias of **Blond Hair** classifier on CelebA [50]

	vanilla	LfF	EIIL	PGI	DebiAN (Ours)
Avg Group Acc.	79.8 \pm 0.3	80.9 \pm 1.4	82.0 \pm 1.1	81.6 \pm 0.3	84.0\pm1.4
Worst Group Acc.	37.9 \pm 1.1	43.3 \pm 3.0	46.1 \pm 4.9	40.9 \pm 6.4	52.9\pm4.7

Table 3: Accuracy results on bias-conflicting samples on bFFHQ [36]

vanilla	LfF	PGI	EIIL	BiaSwap	DebiAN
51.03	55.61	55.2 \pm 5.3	59.2 \pm 1.9	58.87	62.8\pm0.6

that \mathcal{L}_{UA} improves the debiasing results (see w/o \mathcal{L}_{UA}). Besides, we also conduct ablation studies on different batch sizes, which are included in Appendix C.2.

4.2 Experiments on Face Image Dataset


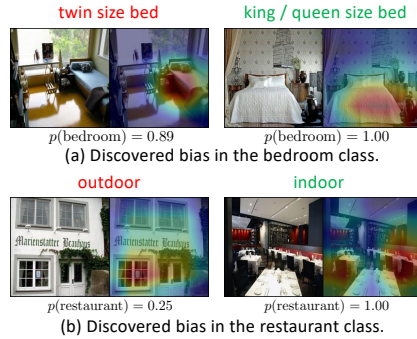
Gender Bias Mitigation In the face image domain, we conduct experiments to evaluate gender bias mitigation results on CelebA [50] dataset, which contains 200K celebrity faces annotated with 40 binary attributes. The dataset has spurious correlations between gender and **Blond Hair**, leading to gender biases when performing hair color classification. We follow most of the settings used in LfF, such as using ResNet-18 [28] as the backbone, using Adam [37] optimizer, *etc.* The only difference is that LfF reports the results on the validation set of CelebA, whereas we use the validation set to select the epoch with the best validation set accuracy (bias labels in the validation set are not used) to report the results on the testing set. All methods (including LfF) are benchmarked under the same setting. We report results in two evaluation metrics: 1) Average Group Accuracy (Avg. Group Acc.), which calculates the unweighted average of accuracies in four groups between target attribute and bias attribute, *i.e.*, (male, female) \times (blond, not blond); 2) Worst Group Accuracy (Worst Group Acc.) [57], which takes the lowest accuracy in the four groups. As shown in Tab. 2, DebiAN achieves better Average and Worst Group accuracy results, which shows that DebiAN can better mitigate gender bias without labels. We also conduct experiments on bFFHQ [36] where the training data contains the spurious correlation between age and gender. We compare DebiAN with other methods of gender bias mitigation. We strictly follow the setting in [36]. We report the age accuracy results on the bias-conflicting samples in the testing set in Tab. 3. The results of vanilla, LfF, and BiaSwap are from [36] and [36] does not provide the error bars. DebiAN achieves the best unsupervised results for mitigating gender bias.

Mitigating Multiple Biases in Gender Classifier The results on Multi-Color MNIST dataset suggest that DebiAN better mitigates multiple biases in the classifier. In the face image domain, a recent study [41] shows that gender classifier is biased by multiple attributes, such as **Heavy Makeup** and **Wearing Lipstick**. Hence, we train gender classifiers on CelebA dataset and evaluate Average Group Accuracy and Worst Group Accuracy w.r.t. these two bias attributes. As shown in Tab. 4, DebiAN achieves better debiasing results w.r.t. both bias attributes, proving that the *discoverer* can find multiple biases in the classifier C during the alternate training, enabling *classifier* to mitigate multiple biases simultaneously.

Identified Unknown Bias in Gender Classifier Gender classifier can have more biases beyond **Wearing Lipstick** and **Heavy Makeup**. For example, Bal-

Table 4: Results of mitigating multiple biases (*i.e.*, Wearing Lipstick and Heavy Makeup) in gender classifier on CelebA dataset

bias attribute	metric	vanilla	LfF	PGI	EIIL	DebiAN (Ours)
Wearing Lipstick	Avg. Group Acc.	86.6 \pm 0.4	87.0 \pm 0.9	86.9 \pm 3.1	86.3 \pm 1.0	88.5\pm1.1
	Worst Group Acc.	53.9 \pm 1.2	55.3 \pm 3.6	56.0 \pm 11.7	52.4 \pm 3.2	61.7\pm4.2
Heavy Makeup	Avg. Group Acc.	85.1 \pm 0.0	85.5 \pm 0.6	85.4 \pm 3.4	84.0 \pm 1.2	87.8\pm1.3
	Worst Group Acc.	45.4 \pm 0.0	46.9 \pm 2.6	46.9 \pm 13.1	40.9 \pm 4.5	56.0\pm5.2

**Fig. 6:** Discovered bias of gender classifier: **visible hair area** based on *discoverer*'s saliency map. $p(\text{female})$ is vanilla classifier's predicted probability of the face is female. In the two groups predicted by D , the visible hair areas are different, where the classifier has different confidences on gender for the same identity**Fig. 7:** Discovered biases in Places [79] dataset. We apply CAM on *discoverer* to generate saliency map. The value $p(\text{bedroom})$ ($p(\text{restaurant})$) is vanilla classifier's predicted probability of the scene image is bedroom (restaurant)

akrishnan *et al.* [8] leverages StyleGAN2 [34] to generate high-quality synthesized images and identify the **hair length** bias of the gender classifier, *e.g.*, longer hair length makes the classifier predict the face as female. Related to their finding, the *discoverer* D in DebiAN identifies an interesting unknown bias: **visible hair area**. We use D to predict the bias attribute group assignment on images in CelebA. To better interpret the bias attribute, we further use the identity labels in CelebA to cluster images with the same identity. Fig. 6 shows that D assigns images of the same identity into two distinct groups based on the visible hair area, which is verified by D 's CAM [78] saliency maps. Strictly speaking, all females in Fig. 6 have long hair. However, due to the hairstyle, pose, or occlusion, visible hair areas differ between the two groups. As a result, the gender classifier has lower predicted probabilities on the female images with smaller visible hair areas. More visualizations are shown in Appendix G.1.

4.3 Experiments on Other Image Domains

Our method is not limited to synthetic and face image domains. Here we conduct experiments on action recognition and scene classification tasks.

Table 5: Results on Biased Action Recognition (BAR) [53] dataset

vanilla	LfF	PGI	EIL	BiaSwap	DebiAN
51.85 \pm 5.92	62.98 \pm 2.76	65.19 \pm 1.32	65.44 \pm 1.17	52.44	69.88\pm2.92

Table 6: Scene classification accuracy results on the *unseen* LSUN [74] dataset

vanilla	LfF	PGI	EIL	DebiAN (Ours)
79.3 \pm 0.3	71.1 \pm 1.0	74.1 \pm 1.9	79.4 \pm 0.2	80.0\pm0.4

Mitigating Place Bias in Action Recognition We conduct experiments on Biased Action Recognition (BAR) dataset [53], an image dataset with the spurious correlation between action and place in the training set. The testing set only contains bias-conflicting samples. Hence, higher accuracy results on the testing set indicate better debiasing results. The accuracy results in Tab. 5 show that DebiAN achieves better debiasing results than other methods.

Improving Cross-dataset Generalization on Scene Classification We conduct experiments on the more challenging scene classification task, where datasets are more complex and may contain multiple unknown biases. The biases in this task are underexplored by previous works partly due to the lack of attribute labels. Due to the absence of attribute labels, we use cross-dataset generalization [64] to evaluate the debiasing results. Concretely, models are trained on Places [79] with ten classes overlapped with LSUN [74] (*e.g.*, bedroom, classroom, *etc.*), and evaluated on the *unseen* LSUN dataset. The results are shown in Tab. 6. DebiAN achieves the best result on the unseen LSUN dataset, showing that DebiAN unlearns the dataset biases [64] in Places to improve the robustness against distributional shifts between different datasets.

Identified Unknown Biases in Scene Classifier DebiAN discovers Places dataset’s unknown biases that humans may not preconceive. In Fig. 7, the *discoverer* separates bedroom and restaurant images based on **size** of beds and **indoor/outdoor**. The vanilla classifier performs worse on bedroom images with twin-size beds and outdoor restaurant images (see more in Appendix G.2).

5 Conclusion

We propose DEBIASING ALTERNATE NETWORKS to discover and mitigate the unknown biases. DebiAN identifies unknown biases that humans may not preconceive and achieves better unsupervised debiasing results. Our Multi-Color MNIST dataset surfaces previous methods’ problems and demonstrates DebiAN’s advantages in the multi-bias setting. Admittedly, our work has some limitations, *e.g.*, DebiAN focuses on binary or continuously valued bias attributes, not multi-class ones. We hope our work can facilitate research on bias discovery and mitigation.

Acknowledgment This work has been partially supported by the National Science Foundation (NSF) under Grant 1764415, 1909912, and 1934962 and by the Center of Excellence in Data Science, an Empire State Development-designated Center of Excellence. The article solely reflects the opinions and conclusions of its authors but not the funding agents.

References

1. Agrawal, R., Srikant, R.: Fast Algorithms for Mining Association Rules in Large Databases. In: International Conference on Very Large Data Bases (1994)
2. Ahmed, F., Bengio, Y., van Seijen, H., Courville, A.: Systematic generalisation with group invariant predictions. In: International Conference on Learning Representations (2021)
3. Albiero, V., K. S., K., Vangara, K., Zhang, K., King, M.C., Bowyer, K.W.: Analysis of Gender Inequality In Face Recognition Accuracy. In: The IEEE Winter Conference on Applications of Computer Vision Workshops (WACVW) (2020)
4. Alvi, M., Zisserman, A., Nellaaker, C.: Turning a Blind Eye: Explicit Removal of Biases and Variation from Deep Neural Network Embeddings. In: The European Conference on Computer Vision Workshop (ECCVW) (2018)
5. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: VQA: Visual Question Answering. In: The IEEE International Conference on Computer Vision (ICCV) (2015)
6. Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D.: Invariant Risk Minimization. arXiv:1907.02893 [cs, stat] (2020)
7. Bahng, H., Chun, S., Yun, S., Choo, J., Oh, S.J.: Learning De-biased Representations with Biased Representations. In: International Conference on Machine Learning (2020)
8. Balakrishnan, G., Xiong, Y., Xia, W., Perona, P.: Towards causal benchmarking of bias in face analysis algorithms. In: The European Conference on Computer Vision (ECCV) (2020)
9. Bao, Y., Barzilay, R.: Learning to Split for Automatic Bias Detection. ArXiv (2022)
10. Bao, Y., Chang, S., Barzilay, D.R.: Learning Stable Classifiers by Transferring Unstable Features. In: International Conference on Machine Learning (2022)
11. Bao, Y., Chang, S., Barzilay, R.: Predict then Interpolate: A Simple Algorithm to Learn Stable Classifiers. In: International Conference on Machine Learning (2021)
12. Buolamwini, J., Gebru, T.: Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In: ACM Conference on Fairness, Accountability, and Transparency (2018)
13. Cadene, R., Dancette, C., Ben younes, H., Cord, M., Parikh, D.: RUBi: Reducing Unimodal Biases for Visual Question Answering. In: Advances in Neural Information Processing Systems (2019)
14. Choi, J., Gao, C., Messou, J.C.E., Huang, J.B.: Why Can't I Dance in the Mall? Learning to Mitigate Scene Bias in Action Recognition. In: Advances in Neural Information Processing Systems (2019)
15. Clark, C., Yatskar, M., Zettlemoyer, L.: Don't Take the Easy Way Out: Ensemble Based Methods for Avoiding Known Dataset Biases. In: Empirical Methods in Natural Language Processing (2019)
16. Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., Huq, A.: Algorithmic Decision Making and the Cost of Fairness. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2017)
17. Creager, E., Jacobsen, J.H., Zemel, R.: Environment Inference for Invariant Learning. In: International Conference on Machine Learning (2021)
18. Creager, E., Madras, D., Jacobsen, J.H., Weis, M., Swersky, K., Pitassi, T., Zemel, R.: Flexibly Fair Representation Learning by Disentanglement. In: International Conference on Machine Learning (2019)

19. Dhar, P., Gleason, J., Roy, A., Castillo, C.D., Chellappa, R.: PASS: Protected Attribute Suppression System for Mitigating Bias in Face Recognition. In: The IEEE International Conference on Computer Vision (ICCV) (2021)
20. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (2012)
21. Geirhos, R., Jacobsen, J.H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., Wichmann, F.A.: Shortcut learning in deep neural networks. *Nature Machine Intelligence* (2020)
22. Gong, S., Liu, X., Jain, A.K.: Jointly De-biasing Face Recognition and Demographic Attribute Estimation. In: The European Conference on Computer Vision (ECCV) (2020)
23. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems* (2014)
24. Grgic-Hlaca, N., Zafar, M.B., Gummadi, K.P., Weller, A.: The case for process fairness in learning: Feature selection for fair decision making. In: *NIPS Symposium on Machine Learning and the Law* (2016)
25. Hanna, A., Denton, E., Smart, A., Smith-Loud, J.: Towards a critical race methodology in algorithmic fairness. In: *Conference on Fairness, Accountability, and Transparency* (2020)
26. Hardt, M., Price, E., Srebro, N.: Equality of Opportunity in Supervised Learning. In: *Advances in Neural Information Processing Systems* (2016)
27. Hazirbas, C., Bitton, J., Dolhansky, B., Pan, J., Gordo, A., Ferrer, C.C.: Towards Measuring Fairness in AI: The Casual Conversations Dataset. *arXiv:2104.02821 [cs]* (2021)
28. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
29. Hendricks, L.A., Burns, K., Saenko, K., Darrell, T., Rohrbach, A.: Women also Snowboard: Overcoming Bias in Captioning Models. In: *The European Conference on Computer Vision (ECCV)* (2018)
30. Jia, S., Meng, T., Zhao, J., Chang, K.W.: Mitigating Gender Bias Amplification in Distribution by Posterior Regularization. In: *Annual Meeting of the Association for Computational Linguistics* (2020)
31. Joo, J., Kärkkäinen, K.: Gender Slopes: Counterfactual Fairness for Computer Vision Models by Attribute Manipulation. In: *International Workshop on Fairness, Accountability, Transparency and Ethics in Multimedia* (2020)
32. Kamiran, F., Calders, T.: Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* (2012)
33. Karras, T., Laine, S., Aila, T.: A Style-Based Generator Architecture for Generative Adversarial Networks. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)
34. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and Improving the Image Quality of StyleGAN. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
35. Kim, B., Kim, H., Kim, K., Kim, S., Kim, J.: Learning Not to Learn: Training Deep Neural Networks With Biased Data. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2019)

36. Kim, E., Lee, J., Choo, J.: BiaSwap: Removing dataset bias with bias-tailored swapping augmentation. In: The IEEE International Conference on Computer Vision (ICCV) (2021)
37. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: International Conference on Learning Representations (2015)
38. Krishnakumar, A., Prabhu, V., Sudhakar, S., Hoffman, J.: UDIS: Unsupervised Discovery of Bias in Deep Visual Recognition Models. In: British Machine Vision Conference, BMVC (2021)
39. Kusner, M.J., Loftus, J., Russell, C., Silva, R.: Counterfactual Fairness. In: Advances in Neural Information Processing Systems (2017)
40. Lahoti, P., Beutel, A., Chen, J., Lee, K., Prost, F., Thain, N., Wang, X., Chi, E.: Fairness without Demographics through Adversarially Reweighted Learning. In: Advances in Neural Information Processing Systems (2020)
41. Lang, O., Gandelsman, Y., Yarom, M., Wald, Y., Elidan, G., Hassidim, A., Freeman, W.T., Isola, P., Globerson, A., Irani, M., Mosseri, I.: Explaining in Style: Training a GAN to explain a classifier in StyleSpace. In: The IEEE International Conference on Computer Vision (ICCV) (2021)
42. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* (1998)
43. Lee, J., Kim, E., Lee, J., Lee, J., Choo, J.: Learning Debaised Representation via Disentangled Feature Augmentation. In: Advances in Neural Information Processing Systems (2021)
44. Li, W., Zhang, P., Zhang, L., Huang, Q., He, X., Lyu, S., Gao, J.: Object-Driven Text-To-Image Synthesis via Adversarial Training. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
45. Li, Y., Li, Y., Vasconcelos, N.: RESOUND: Towards Action Recognition without Representation Bias. In: The European Conference on Computer Vision (ECCV) (2018)
46. Li, Y., Vasconcelos, N.: REPAIR: Removing Representation Bias by Dataset Resampling. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
47. Li, Z., Xu, C.: Discover the Unknown Biased Attribute of an Image Classifier. In: The IEEE International Conference on Computer Vision (ICCV) (2021)
48. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal Loss for Dense Object Detection. In: The IEEE International Conference on Computer Vision (ICCV) (2017)
49. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common Objects in Context. In: The European Conference on Computer Vision (ECCV) (2014)
50. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep Learning Face Attributes in the Wild. In: The IEEE International Conference on Computer Vision (ICCV) (2015)
51. Manjunatha, V., Saini, N., Davis, L.S.: Explicit Bias Discovery in Visual Question Answering Models. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
52. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D., Gebru, T.: Model Cards for Model Reporting. In: ACM Conference on Fairness, Accountability, and Transparency (2019)
53. Nam, J., Cha, H., Ahn, S., Lee, J., Shin, J.: Learning from Failure: Training Debaised Classifier from Biased Classifier. In: Advances in Neural Information Processing Systems (2020)

54. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: *Advances in Neural Information Processing Systems* (2019)
55. Pleiss, G., Raghavan, M., Wu, F., Kleinberg, J., Weinberger, K.Q.: On Fairness and Calibration. In: *Advances in Neural Information Processing Systems* (2017)
56. Reddy, C., Sharma, D., Mehri, S., Romero-Soriano, A., Shabanian, S., Honari, S.: Benchmarking Bias Mitigation Algorithms in Representation Learning through Fairness Metrics. In: *Thirty-Fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)* (2021)
57. Sagawa*, S., Koh*, P.W., Hashimoto, T.B., Liang, P.: Distributionally Robust Neural Networks for Group Shifts: On the Importance of Regularization for Worst-Case Generalization. In: *International Conference on Learning Representations* (2020)
58. Sarhan, M.H., Navab, N., Albarqouni, S.: Fairness by Learning Orthogonal Disentangled Representations. In: *The European Conference on Computer Vision (ECCV)* (2020)
59. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Gradcam: Visual explanations from deep networks via gradient-based localization. In: *The IEEE International Conference on Computer Vision (ICCV)* (2017)
60. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal of Computer Vision* (2019)
61. Shrestha, R., Kafle, K., Kanan, C.: An Investigation of Critical Issues in Bias Mitigation Techniques. In: *The IEEE Winter Conference on Applications of Computer Vision (WACV)* (2022)
62. Singh, K.K., Mahajan, D., Grauman, K., Lee, Y.J., Feiszli, M., Ghadiyaram, D.: Don't Judge an Object by Its Context: Learning to Overcome Contextual Bias. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
63. Sohoni, N.S., Dunnmon, J.A., Angus, G., Gu, A., Ré, C.: No Subclass Left Behind: Fine-Grained Robustness in Coarse-Grained Classification Problems. In: *Advances in Neural Information Processing Systems* (2020)
64. Torralba, A., Efros, A.A.: Unbiased look at dataset bias. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2011)
65. Tzeng, E., Hoffman, J., Darrell, T., Saenko, K.: Simultaneous Deep Transfer Across Domains and Tasks. In: *The IEEE International Conference on Computer Vision (ICCV)* (2015)
66. Verma, S., Rubin, J.: Fairness Definitions Explained. In: *2018 IEEE/ACM International Workshop on Software Fairness (FairWare)* (2018)
67. Wang, A., Narayanan, A., Russakovsky, O.: REVISE: A Tool for Measuring and Mitigating Bias in Image Datasets. In: *The European Conference on Computer Vision (ECCV)* (2020)
68. Wang, H., He, Z., Lipton, Z.C., Xing, E.P.: Learning Robust Representations by Projecting Superficial Statistics Out. In: *International Conference on Learning Representations* (2019)
69. Wang, J., Liu, Y., Wang, X.E.: Are Gender-Neutral Queries Really Gender-Neutral? Mitigating Gender Bias in Image Search. In: *Empirical Methods in Natural Language Processing* (2021)

70. Wang, T., Yue, Z., Huang, J., Sun, Q., Zhang, H.: Self-Supervised Learning Disentangled Group Representation as Feature. In: *Advances in Neural Information Processing Systems* (2021)
71. Wang, X., Ang, M.H., Lee, G.H.: Cascaded Refinement Network for Point Cloud Completion. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
72. Wang, Z., Liu, X., Li, H., Sheng, L., Yan, J., Wang, X., Shao, J.: CAMP: Cross-Modal Adaptive Message Passing for Text-Image Retrieval. In: *The IEEE International Conference on Computer Vision (ICCV)* (2019)
73. Wang, Z., Qinami, K., Karakozis, I.C., Genova, K., Nair, P., Hata, K., Russakovsky, O.: Towards Fairness in Visual Recognition: Effective Strategies for Bias Mitigation. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
74. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop. *arXiv:1506.03365 [cs]* (2016)
75. Zhang, B.H., Lemoine, B., Mitchell, M.: Mitigating Unwanted Biases with Adversarial Learning. In: *AAAI/ACM Conference on AI, Ethics, and Society* (2018)
76. Zhang, Z., Sabuncu, M.: Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels. In: *Advances in Neural Information Processing Systems* (2018)
77. Zhao, J., Wang, T., Yatskar, M., Ordonez, V., Chang, K.W.: Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints. In: *Empirical Methods in Natural Language Processing* (2017)
78. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning Deep Features for Discriminative Localization. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
79. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018)

Appendix

A Pseudocode of DebiAN

We present the pseudocode of DebiAN for two tasks – 1) discover the unknown biases (Alg. 1); 2) mitigate the unknown biases (Alg. 2). To ensure that the sampled images have the same target attribute labels, we select images with the same target attribute label in a batch to compute the loss (line 3-7 in Alg. 1, 2, and line 10-14 in Alg. 2).

Algorithm 1: Discover unknown biases.

Input: C : trained *classifier*, T : number of iterations, K : number of target attribute classes
Output: D : *discoverer*
Data: \mathcal{D} : training set

```

1 for  $t : 1 \dots T$  do
2    $\mathcal{B} := \{(\mathbf{I}_i, y_i)\}_{i=1}^N \sim \mathcal{D}$  // Sample a batch  $\mathcal{B}$  with  $N$  pairs of images
    $\mathbf{I}_i$  and target attribute labels  $y_i$ 
   /* for each target attribute class  $t$  */
3   for  $k : 1 \dots K$  do
4      $\mathcal{B}_k := \{(\mathbf{I}_j, y_j) \mid y_j = k, (\mathbf{I}_j, y_j) \in \mathcal{B}\}_{j=1}^M$  // Select  $M$  pairs from  $\mathcal{B}$ 
       whose labels are  $k$ 
5      $p(\hat{y} \mid \mathbf{I}_j) := C(\mathbf{I}_j), \mathbf{I}_j \in \mathcal{B}_k$  //  $C$  predicts target attribute
6      $p(\hat{b} \mid \mathbf{I}_j) := D(\mathbf{I}_j), \mathbf{I}_j \in \mathcal{B}_k$  //  $D$  predicts bias attribute groups
7      $\mathcal{L}_k = \mathcal{L}_{\text{EOV}} + \mathcal{L}_{\text{UA}}$  // Compute loss on  $\mathcal{B}_t$ .
8   update  $D$  with loss  $1/K \sum_{k=1}^K \mathcal{L}_k$ 

```

B Implementation Details

In DebiAN, the *discoverer* and *classifier* use the same architecture but do not share the parameters.

On Multi-Color MNIST dataset (Sec. 4.1), we follow the same setting used in LfF [53]’s experiment on Colored MNIST. We use Adam [37] optimizer with 10^{-3} learning rate and 256 batch size. We use an MLP with three hidden layers (obtained from the LfF’s official code*). All models are trained for 100 epochs.

In the experiments for gender bias mitigation on CelebA [50] dataset (Sec. 4.2), we follow most of the settings used in LfF. We use ResNet-18 [28] as the network architecture. We use horizontal flip for data augmentation during training. We use Adam optimizer with 10^{-4} learning rate and 256 batch size. All models are trained for 50 epochs. The only difference is that we use CelebA’s validation set

*<https://github.com/alinlab/LfF>

Algorithm 2: Mitigate unknown biases.

Input: T : number of iterations, K : number of target attribute classes
Output: C : classifier, D : discoverer
Data: \mathcal{D} : training set

```

1 for  $t : 1 \dots T$  do
    /* ===== Start: optimize  $C$ , freeze  $D$  ===== */
2    $\mathcal{B} := \{(\mathbf{I}_i, y_i)\}_{i=1}^N \sim \mathcal{D}$  // Sample a batch  $\mathcal{B}$  with  $N$  pairs of images
     $\mathbf{I}_i$  and target attribute labels  $y_i$ 
    /* for each target attribute class  $t$  */
3   for  $k : 1 \dots K$  do
4        $\mathcal{B}_k := \{(\mathbf{I}_j, y_j) \mid y_j = k, (\mathbf{I}_j, y_j) \in \mathcal{B}\}_{j=1}^M$  // Select  $M$  pairs from  $\mathcal{B}$ 
        whose labels are  $k$ 
5        $p(\hat{y} \mid \mathbf{I}_j) := C(\mathbf{I}_j), \mathbf{I}_j \in \mathcal{B}_k$  //  $C$  predicts target attribute
6        $p(\hat{b} \mid \mathbf{I}_j) := D(\mathbf{I}_j), \mathbf{I}_j \in \mathcal{B}_k$  //  $D$  predicts bias attribute groups
7        $\mathcal{L}_k^C = \mathcal{L}_{\text{RCE}}$  // Compute loss on  $\mathcal{B}_k$ 
8   update  $C$  with loss  $1/K \sum_{k=1}^K \mathcal{L}_k^C$ 
    /* ===== End: optimize  $C$ , freeze  $D$  ===== */
    /* ===== Start: optimize  $D$ , freeze  $C$  ===== */
9    $\mathcal{B} := \{(\mathbf{I}_i, y_i)\}_{i=1}^N \sim \mathcal{D}$  // Sample a batch  $\mathcal{B}$  with  $N$  pairs of images
     $\mathbf{I}_i$  and target attribute labels  $y_i$ 
    /* for each target attribute class  $k$  */
10  for  $k : 1 \dots K$  do
11       $\mathcal{B}_k := \{(\mathbf{I}_j, y_j) \mid y_j = k, (\mathbf{I}_j, y_j) \in \mathcal{B}\}_{j=1}^M$  // Select  $M$  pairs from  $\mathcal{B}$ 
        whose labels are  $k$ 
12       $p(\hat{y} \mid \mathbf{I}_j) := C(\mathbf{I}_j), \mathbf{I}_j \in \mathcal{B}_k$  //  $C$  predicts target attribute
13       $p(\hat{b} \mid \mathbf{I}_j) := D(\mathbf{I}_j), \mathbf{I}_j \in \mathcal{B}_k$  //  $D$  predicts bias attribute groups
14       $\mathcal{L}_k^D = \mathcal{L}_{\text{EOV}} + \mathcal{L}_{\text{UA}}$  // Compute loss on  $\mathcal{B}_k$ 
15  update  $D$  with loss  $1/K \sum_{k=1}^K \mathcal{L}_k^D$ 
    /* ===== End: optimize  $D$ , freeze  $C$  ===== */

```

to choose the epoch where models achieve the best validation set accuracy and report the results on the testing set. Note that validation set accuracy does not use any bias attribute labels because unsupervised debiasing should not rely on any labels of bias attributes. LfF directly reports the results at the 50 epoch on the validation set, which is hard to be replicated as reported by other users in their official code GitHub repository[†].

In the experiments for gender bias mitigation on bFFHQ dataset [36], we use the same setting in [36]. We use Adam as the optimizer with 256 batch size. All models are trained for 200 epochs. We use ResNet-18 as the backbone. We notice that Lee *et al.* [43] use a different setting on bFFHQ dataset with StepLR for the learning rate scheduling, which is more complicated than the one in the original paper [36]. Thus, we choose the former one as the setting on bFFHQ dataset.

[†]<https://github.com/alinlab/LfF/issues/2>

Table 7: Ablation study on Unbalanced Assignment (UA) penalty (*i.e.*, \mathcal{L}_{UA}) in mitigating gender bias of **Blond Hair** classifier on CelebA [50] dataset

	w/o \mathcal{L}_{UA} DebiAN	
Avg group Acc.	79.6 \pm 1.7	84.0 \pm 1.4
Worst group Acc.	38.5 \pm 4.7	52.9 \pm 4.7

In the experiment of mitigating multiple biases in gender classifier on CelebA dataset (Sec. 4.2), we choose 64 as the batch size and ResNet-50 as the backbone of classifiers. All models are trained with 50 epochs. We use CelebA’s validation set to choose the epoch that has the best validation set accuracy for each method. We report the results on the testing set. We use Adam as the optimizer with 10^{-4} learning rate.

On Biased Action Recognition (BAR) [53] dataset, we use the setting in [53]. We use Adam as the optimizer with 10^{-4} learning rate. The batch size is 256. We use 224×224 random cropping for data augmentation. All models are trained with 90 epochs.

In the scene classification task, we choose ResNet-18 as the backbone of classifiers and 128 as the batch size. We use Adam as the optimizer with 10^{-4} learning rate. All models are trained only on the Places [79] dataset for 50 epochs. We choose the epoch where the model achieves the best accuracy on Places’s validation set and report the results on LSUN’s [74] validation set.

The code is based on PyTorch [54]. We modify LfF’s code[‡] that generates Colored MNIST to create Multi-Color MNIST dataset.

In implementation, we add $\epsilon = 10^{-6}$ to the denominators of $\bar{P}_{b+}(\hat{y})$ and $\bar{P}_{b-}(\hat{y})$ (Eq. (3)) to avoid zero division.

For the *discoverer* D , we choose two different implementations for different numbers of classes of the target attribute. When the target attribute is binary (*e.g.*, experiments in Sec. 4.2), *i.e.* number of classes is two, D predicts one value for each image, which is the predicted bias attribute group. We denote this implementation as “global” since two classes globally share the predicted bias attribute groups. When the target attribute has $c > 2$ classes, *e.g.*, ten classes in the digit classification, action recognition, and scene classification in Sec. 4.1 and Sec. 4.3, D predicts c values, where each value is the predicted bias attribute group of the corresponding target attribute class. We denote this implementation as “per class” since D predicts bias attribute groups for each target attribute class. We provide an ablation study on this in Appendix C.5.

Table 8: Ablation study on Unbalanced Assignment (UA) penalty (*i.e.*, \mathcal{L}_{UA}) in mitigating multiple biases of gender classifier on CelebA [50] dataset

bias attribute	metric	w/o \mathcal{L}_{UA}	DebiAN
Wearing Lipstick	Avg. Group Acc.	87.7 \pm 0.4	88.5 \pm 1.1
	Worst Group Acc.	58.1 \pm 1.2	61.7 \pm 4.2
Heavy Makeup	Avg. Group Acc.	85.6 \pm 1.2	87.8 \pm 1.3
	Worst Group Acc.	46.9 \pm 5.2	56.0 \pm 5.2

Table 9: Ablation study on different batch sizes on Multi-Color MNIST dataset. We report the accuracy results for images that are both bias-conflicting w.r.t. **left color** and **right color** bias attributes. We also report the unbiased accuracy results. DebiAN achieves better debiasing results under all batch sizes

batch size		vanilla	DebiAN (Ours)
32	both bias-conflicting	8.0 \pm 0.5	18.9 \pm 1.2
	unbiased accuracy	61.7 \pm 1.0	75.0 \pm 0.8
64	both bias-conflicting	8.2 \pm 1.9	18.1 \pm 0.8
	unbiased accuracy	61.7 \pm 1.4	74.2 \pm 0.4
128	both bias-conflicting	5.6 \pm 1.8	17.2 \pm 1.2
	unbiased accuracy	58.7 \pm 2.5	72.1 \pm 0.7
256	both bias-conflicting	5.2 \pm 0.4	16.0 \pm 1.8
	unbiased accuracy	57.4 \pm 0.7	72.0 \pm 0.8
512	both bias-conflicting	4.8 \pm 1.0	12.5 \pm 1.9
	unbiased accuracy	56.1 \pm 1.3	70.1 \pm 1.1

C Ablation Study

C.1 Unbalanced Assignment (UA) penalty

Here we show more ablation study results on the Unbalanced Assignment (UA) penalty on CelebA dataset. The results are shown in Tabs. 7 and 8, which further proves that \mathcal{L}_{UA} can improve fairness results by avoiding the trivial solution—assigning all images into a single bias group (see Unbalanced Assignment (UA) penalty in Sec. 3.1).

C.2 Batch Size

In practice, $\{\mathbf{I}_i\}_{i=1}^n$ (defined in Sec. 3) is a mini-batch of images sampled from the dataset for optimizing the networks. One may have the concern that the sampled batch may not have enough images from different bias groups for the *discoverer* to assign. Therefore, we conduct an ablation study on different batch

[‡]https://github.com/alinlab/LfF/blob/master/make_dataset.py

sizes on Multi-Color MNIST dataset with the same setting introduced in Sec 4.1, where the ratio of the `left color` is 0.99 and `right color` is 0.95. We report the accuracy results for images that are both bias-conflicting w.r.t. `left color` and `right color` bias attributes (see “both bias-conflicting” in Tab. 9). We also report the unbiased accuracy results. The results in Tab. 9 show that DebiAN can achieve better debiasing results under different batch sizes compared with the vanilla model.

C.3 Ablation Study on Different Ratios

We conduct an ablation study on different ratios of bias-aligned samples in Multi-Color MNIST’s training set. We keep the ratio for the `right color` bias attribute to 0.95 and use different ratios for `left color` bias attribute, ranging from 0.995 to 0.95. The results are shown in Tab. 10. DebiAN achieves better unbiased accuracy results and accuracy results on samples that are bias-conflicting w.r.t. both bias attributes. The only exception is the accuracy results of the samples that are bias-conflicting w.r.t. both bias attributes when both ratios are 0.95 (last section in Tab. 10). Both LfF and DebiAN achieve 39.6 accuracy results. However, our method achieves a lower standard deviation (0.2) than LfF (6.9) and achieves much better final unbiased results (81.8 vs. 68.5). We also notice that LfF’s debiasing results have a large standard deviation when the ratios of both bias attributes are 0.95. We provide an explanation in Appendix D.3.

C.4 Alternate Training

We conduct an ablation study on alternate training on the Multi-Color MNIST dataset with the same setting used in Sec. 4.1. To remove the alternate training from DebiAN, we follow EIIL [17] and PGI [2] to train the *discoverer* to identify the unknown biases in a *classifier* trained with one epoch. After training *discoverer*, we fix the parameters of *discoverer* and only train the *classifier* to perform debiasing. The results in Tab. 11 show that alternate training can improve the debiasing results, *e.g.*, higher unbiased accuracy and higher accuracy for samples that are bias-conflicting w.r.t. both bias attributes (4th row), which demonstrates the necessity of alternate training.

C.5 Bias Attribute Groups: Global vs. Per Class

As mentioned in Appendix B, the predicted bias attribute groups from the *discoverer* (D) are shared by both classes in the binary classification task. In the multi-class classification setting (*e.g.*, digit classification, scene classification task, *etc.*), D predicts binary bias group assignments for each class. We justify our implementation choices with the results in Tab. 12.

For the binary age classification on bFFHQ dataset, there is no significant difference between the two implementation choices (*i.e.*, differences are within error bars). Therefore, we choose “global” *discoverer* for the binary classification task due to its simplicity.

However, in the multi-class digit classification on Multi-Color MNIST dataset, we do observe the better result produced by the *discoverer* that predicts bias group assignment for each class (*i.e.*, improvement is greater than the error bar). We suspect that predicting bias attribute groups per class in the binary classification task is redundant because the binary target attribute is spuriously correlated with the binary bias attribute. For example, if the target attribute **age** is spuriously correlated with the bias attribute **gender**, *i.e.*, more young females and old males than old females and young males, then it is not necessary to predict bias attribute group for both genders since both genders share the same bias attribute groups. However, this may not be the case for multi-class settings. For example, in Multi-Color MNIST dataset, each digit class is spuriously correlated with a unique left color, *e.g.*, for bias-aligned samples, digit class 0’s left color is red but digit class 1’s left color is yellow (Fig. 4). In other words, the bias attribute values may not be shared globally across different target attribute classes. Therefore, we choose different numbers of outputs for the *discoverer* under different tasks.

C.6 Alternative Design for Debiasing: $\max_C \mathcal{L}_{\text{EOV}}$

One may consider an alternative design for debiasing—train the *classifier* C to maximize the Equal Opportunity Violation (EOV) loss, or formally $\max_C \mathcal{L}_{\text{EOV}}$. This alternative design, to some degree similar to GAN [23]’s training strategy, may look more “unified” since it lets the *discoverer* D and *classifier* C play the minmax game:

$$\min_C \max_D |\bar{P}_{b+}(\hat{y}) - \bar{P}_{b-}(\hat{y})|, \quad (9)$$

where $\bar{P}_{b+}(\hat{y})$ and $\bar{P}_{b-}(\hat{y})$ are defined in Eq. 3. This alternative design enables C to *directly* meet the Equal Opportunity [26,55] fairness criterion. More concretely, we implement this alternative design of C ’s objectives by the following loss function:

$$\min_C -\log(1 - |\bar{P}_{b+}(\hat{y}) - \bar{P}_{b-}(\hat{y})|) + \text{CE}(p_t(\mathbf{I}_i), y_i), \quad (10)$$

where the first $-\log$ term implements C ’s objective in the minmax game (Eq. (9)) and the second term CE is the standard cross-entropy loss. We conduct an ablation study on the design for debiasing (*i.e.*, playing minmax game vs. RCE loss (Eq. 6)) and the results on Multi-Color MNIST dataset are shown in Tab. 13. The results demonstrate that our RCE loss performs much better than the alternative design. We suspect the reason is that C in DebiAN has two goals – 1) fooling the *discoverer* to achieve fairer results; 2) achieving higher accuracy by optimizing the standard cross-entropy loss, which is different from GAN [23] where the generator only has one goal – fooling the discriminator to achieve better image quality of the synthesized images. Therefore, it is hard to control the balance between the two goals of C in this alternative design. In contrast, our RCE loss can better incorporate the two goals within a single objective function \mathcal{L}_{RCE} , leading to better debiasing results.

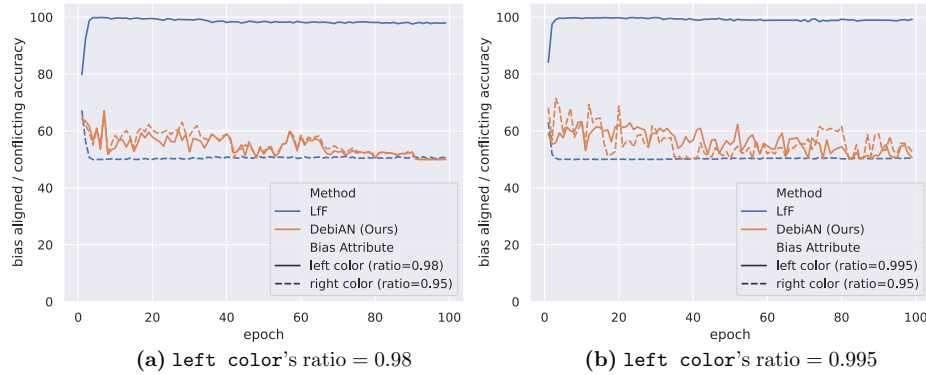


Fig. 8: More bias discovery results w.r.t. **left color** and **right color** biases throughout the training epochs on Multi-Color MNIST dataset under different ratios of bias-aligned samples w.r.t. **left color**. The results are consistent with Fig. 5




Fig. 9: Bias discovery results when **left color** and **right color** are equally salient (both ratios are 0.95). The three plots are results under three different random seeds. In the first two plots, LfF mainly discovers **right color** at the early training stage and gradually discovers both biases. In the third seed, it mainly discovers **left color** bias. The results show that LfF are unstable in bias discovery when two biases are equally salient. In contrast, our DebiAN consistently discovers both biases at the early stage under different random seeds and gradually converges 50% bias discovery accuracy as debiasing is performed in the *classifier*

D Bias Discovery on Multi-Color MNIST

D.1 Implementation Details

To evaluate the bias discovery results, we transform the outputs from LfF’s “biased model” and DebiAN’s *discoverer* into bias-aligned / bias-conflicting prediction by the following approaches.

For LfF, since the biased model is trained to amplify the biases, the biased model’s outputs predict ten colors aligned with digits. Thus, when its predicted color class is the same as the ground-truth digit class, *e.g.*, c -th color for the c -th class, we regard its prediction as bias-aligned. Otherwise, its prediction is bias-conflicting.

For DebiAN, we use *discoverer* and *classifier* to predict each image \mathbf{I}_i ’s predicted bias group assignment $p(\hat{b} \mid \mathbf{I}_i)$ and predicted probability of the ground-

truth class of the target attribute $p_t(\mathbf{I}_i)$ on the entire testing set, respectively. Then, we compute the weighted average predicted probabilities $\bar{P}_{b+}(\hat{y})$ and $\bar{P}_{b-}(\hat{y})$ (see Eq. 3) in two bias groups. If $\bar{P}_{b+}(\hat{y}) \geq \bar{P}_{b-}(\hat{y})$, we use $p(\hat{b} = 1 \mid \mathbf{I}_i)$ as the bias-aligned prediction and $p(\hat{b} = 0 \mid \mathbf{I}_i)$ as the bias-conflicting prediction since the positive bias group has higher weighted average predicted probability, *i.e.*, *classifier* performs better on the positive bias group. If $\bar{P}_{b+}(\hat{y}) < \bar{P}_{b-}(\hat{y})$, we use $p(\hat{b} = 1 \mid \mathbf{I}_i)$ as the bias-conflicting prediction and $p(\hat{b} = 0 \mid \mathbf{I}_i)$ as the bias-aligned prediction.

After obtaining bias-aligned and bias-conflicting predictions, we compute the bias-aligned and bias-conflicting accuracy for **left color** and **right color** biases as follows. Each testing image has two labels—(1) bias-aligned/bias-conflicting w.r.t. left color; (2) bias-aligned/bias-conflicting w.r.t. right color. Thus, the left color (or right color) bias discovery accuracy is computed based on the bias-aligned/bias-conflicting predictions against the left color (or right color) bias-aligned/bias-conflicting labels.

D.2 More Bias Discovery Results under Different Ratios

In the main paper, we evaluate LfF and DebiAN’s bias discovery results when the ratio of **left color** is 0.99 in Fig. 5. Here, we show results under more ratios in Fig. 8, where the ratios of **left color** are 0.98 (Fig. 8 (a)) and 0.995 (Fig. 8 (b)). The results are consistent with Fig. 5—LfF can only discover the more salient **left color** and cannot identify the less salient **right color** bias, whereas DebiAN’s *discoverer* can discover both biases at the early training stage and the bias discovery accuracy gradually converges to 50% when debiasing is performed in the *classifier*.

D.3 Bias Discovery Results under Equally Salient Biases

We further evaluate bias discovery results when **left color** and **right color** biases are equally salient, *i.e.*, ratios of both biases are 0.95. We found that LfF shows more unstable results under different random seeds than in previous settings. Therefore, we show bias discovery results under three different random seeds in Fig. 9. Under the first two random seeds (left and middle plots in Fig. 9), LfF first discovers **right color** bias and gradually discovers both biases. However, under the third random seed (the right plot in Fig. 9), LfF mainly discovers the **left color** bias. Therefore, our Multi-Color MNIST dataset reveals another weakness of LfF—unstable bias discovery results when two biases are equally salient, which also explains LfF’s unstable debiasing results under the equally salient biases (LfF has large error bars, *e.g.*, ± 33.7 and ± 25.9 , of the accuracy results in Tab. 10). In contrast, DebiAN’s bias discovery results are stable—consistent results across different random seeds and under different ratios.

D.4 Detailed Discussion on Bias Discovery

Why LfF’s bias discovery accuracies do not converge to 50%? In Fig. 5 and Fig. 8, LfF’s bias discovery accuracies maintain at about 100% or 50%

throughout the training epochs. One may wonder why it does not converge to 50% as DebiAN does. The reason is that DebiAN discovers biases from the *classifier* (\mathcal{L}_{EOV} is based on *classifier*’s output), whereas LfF identifies biases from the dataset. Concretely, LfF uses the assumption that the bias attribute is easier than the target attribute to define the bias and uses Generalized Cross-entropy (GCE) loss [76] to train a biased model. GCE loss is defined by $-p_t^q \log p_t$, where p_t is bias model’s predicted probability of the ground-truth class and q is a hyperparameter. Intuitively, it up-weights easy examples (*i.e.*, high p_t) with high weight p_t^q and down-weights hard examples (*i.e.*, low p_t) with low weight p_t^q . Therefore, the “biased model” focuses more on the easy examples in the dataset. However, it does not know any biases in the *classifier* (no classifier’s outputs used in GCE). Therefore, whether the classifier is performing debiasing will not affect LfF’s bias discovery, making the bias attribute accuracy stays the same throughout the entire training stage. In contrast, DebiAN’s *discoverer* actively identifies biases in the *classifier*. Therefore, *discoverer*’s bias discovery results will converge to 50% as debiasing is performing in the *classifier*, making the *discoverer* harder to find the biases.

Bias Discovery: EIIL and PGI Here, we discuss the connection and difference between DebiAN and two previous methods—EIIL and PGI. In contrast to LfF that finds biases from the dataset based on the assumption, all EIIL, PGI, and DebiAN actively find biases from the classifier. However, DebiAN differs from EIIL and PGI in terms of the objective function, network architecture, and training scheme. We mainly introduce EIIL because PGI is a follow-up work for EIIL with a difference in the network architecture.

In terms of objective function, EIIL (and PGI) inversely uses the debiasing objective function—IRMv1 [6]. In other words, while minimizing IRMv1 was designed for debiasing in previous works, EIIL maximizes the gradient norm penalty in IRMv1 to identify biases. However, this is suboptimal for two reasons. First, IRMv1 approximates IRM with gradient norm penalty to make it computationally tractable. However, the zero gradient norm may only indicate a local minimum instead of the global minimum. In contrast, DebiAN’s EOv loss uses the principled definition to define the bias—violation of equal opportunity fairness criterion. Second, since the gradient norm does not have an upper bound, maximizing it leads to an optimization problem. As a comparison, DebiAN’s EOv loss minimizes a bounded negative log-likelihood (Eq. 2), which is easier to be optimized.

Regarding network architecture, EIIL does not train any networks but directly optimizes a vector $\mathbf{q} \in \mathbb{R}^N$ for N images in the training set to maximize IRMv1’s gradient norm, which can be regarded as directly optimize the bias group assignments. Since the vector \mathbf{q} is only fitted to the training set, we cannot evaluate EIIL’s bias discovery results on the balanced testing set. PGI uses a slightly different approach by training a small MLP that takes the classifier’s features as the input and predicts the bias group assignments. PGI’s bias discovery objective function is identical with EIIL’s—maximizing IRMv1. We use the trained MLP

to infer the bias-aligned / bias-conflicting on the testing set on Multi-Color MNIST dataset under the setting used in Sec. 4.1 (ratio w.r.t. `left color` is 0.99 and ratio w.r.t. `right color` is 0.95) based on a classifier trained with one epoch (explain in the next paragraph). The bias discovery accuracy results w.r.t. `left color` and `right color` are 50.6 ± 1.9 and 49.9 ± 0.3 . In contrast, DebiAN’s *discoverer* has the same network architecture as the classifier to predict bias group assignments from raw images, which enables *discoverer* to learn its own feature of the bias attribute directly from the raw images. As a result, DebiAN achieves about 60% to 75% accuracy at the first epoch (see Fig. 5). Therefore, we empirically show that DebiAN can discover biases more accurately.

Finally, with respect to the training scheme, by hypothesizing that the classifier learns bias features in the early stage, EIL and PGI use a fixed classifier trained with one epoch to discover biases. However, this might not be the case in the multi-bias setting where multiple biases may be learned at different training stages. In contrast, DebiAN trains *discoverer* and *classifier* in an alternate fashion, enabling *discoverer* to find biases in the *classifier* during the entire training stage. Our ablation study on the alternate training (Appendix C.4) further demonstrates its benefits.

E Results on Colored MNIST (single-bias setting)

We also compare with other methods on Colored MNIST in a single-bias setting. There are two variants of Colored MNIST datasets in previous works—(1) adding colors to the foreground (*i.e.*, digit) [2,35,46]; (2) adding colors to the background [7,56]. Therefore, we conduct experiment on both variants of Colored MNIST dataset. We denote the Colored MNIST with foreground color as “Colored MNIST (foreground)” and the Colored MNIST with background color as “Colored MNIST (background).” Same with the experiment setting on Multi-Color MNIST, we follow the setting used in LfF [53], including using MLP as the network architecture, training with 100 epochs, *etc.* Same with LfF, we report the results under four ratios of bias-aligned samples in the training set—0.995, 0.99, 0.98, and 0.95. In terms of evaluation metrics, we follow LfF to report the accuracy results on bias-conflicting samples and unbiased accuracy. We additionally report the accuracy results on bias-aligned samples.

We report LfF’s reported results on Colored MNIST (foreground). Besides, we also replicate LfF’s results via their officially released code. We cannot replicate their reported results. This issue is also reported in the official code’s GitHub repository by other people[§].

The results on Colored MNIST (foreground color) are in Tab. 14. Although LfF achieves better bias-conflicting accuracy results, it also achieves much lower bias-aligned accuracy, revealing that LfF’s reweighing method overly focuses on the bias-conflicting samples than the bias-aligned samples. As a result, LfF achieves low unbiased accuracy. Overall, DebiAN achieves comparable or slightly lower unbiased accuracy results.

[§]<https://github.com/alinalab/LfF/issues/1>

The results on Colored MNIST (background color) are in Tab. 15. Similar to the results on Colored MNIST (foreground color), LfF achieves better bias-conflicting accuracy but low bias-aligned accuracy and unbiased accuracy. Overall, DebiAN achieves comparable or slightly better unbiased accuracy results.

We also notice that PGI achieves inconsistent results on Colored MNIST (foreground color) and Colored MNIST (background color). While PGI achieves very good results on Colored MNIST (foreground color), *e.g.*, top-1 unbiased accuracy when ratio = 0.995 and ratio = 0.98 in Tab. 14, it achieves bad results on Colored MNIST (background color), *e.g.* the lowest unbiased accuracy results in Tab. 15. We suspect the reason is that PGI is overly sensitive to hyperparameters, *e.g.*, the coefficient of the KL-divergence for debiasing. In contrast, our method achieves good results across two Colored MNIST variants without tuning or changing any hyperparameters.

Finally, we restate that Colored MNIST is a single-bias setting, which may not be the case in real-world scenarios where multiple biases exist. Therefore, we regard that we should focus more on our new Multi-Color dataset to evaluate the debiasing results w.r.t. multiple biases (Tab. 10), where DebiAN achieves better debiasing results.

F More results on Mitigating Multiple Biases in Gender Classification

In Tab. 4, we show better DebiAN’s better debiasing results w.r.t. **Wearing Lipstick** and **Heavy Makeup**. To demonstrate that DebiAN’s better debiasing results w.r.t. more bias attributes, we evaluate on Transects [8] dataset (mentioned on L567 in the main paper). Transects dataset contains high-quality face images synthesized by StyleGAN2 [34], which are also balanced w.r.t. multiple biases such as **Hair Length** and **Skin Color**. The dataset does not contain training split because it is designed as a testing set to identify biases in gender classifier. We use gender classifiers trained on CelebA dataset (same setting used in Sec. 4.2) to evaluate their debiasing performance w.r.t. **Hair Length** and **Skin Color** bias attributes on Transects dataset. The results are shown in Tab. 16, which demonstrates DebiAN’s better capability in mitigating multiple biases simultaneously. Furthermore, the better results w.r.t. **Hair Length** can also reflect that DebiAN’s *discoverer* identifies **visible hair area** bias attribute (Fig. 6 and Fig. 10).

G More Qualitative Results

G.1 More Examples of Discovered Biases on Face Images

While Fig. 5 shows the discovered **visible hair area** bias of gender classifier on female images, we further show the male image examples in Fig. 10. *D* focuses on the **visible hair area** to separate images into “small visible hair area” and “large visible hair area” groups, which is consistent with the female examples shown in Fig. 5.

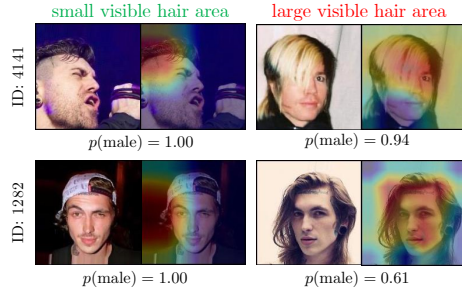


Fig. 10: Discovered bias of gender classifier: **visible hair area**. D 's CAM saliency map is paired with each image. The vanilla gender classifier train on CelebA dataset performs worse for males with larger visible hair area

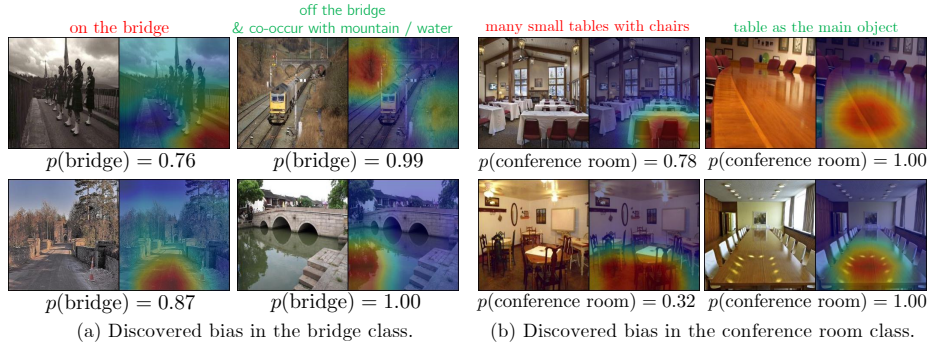


Fig. 11: Discovered biases of vanilla scene classifiers. D 's CAM is paired with the image

G.2 More Examples of Discovered Biases on Scene Images

We show more examples of discovered biases on the scene images in Fig. 11. For bridge images, D predicts two bias groups. When the photos are taken on the bridge, the vanilla classifier performs worse. In comparison, the vanilla classifier performs better when the photos are taken off the bridge with some correlated backgrounds, such as mountains or water. For conference room images, the vanilla classifier performs better when the table is the major object in the scene. However, it performs worse when the conference room images have many tables, and the tables do not occlude the chairs.

H Discussion

H.1 Is DebiAN a hard negative method?

No. Hard negative methods focus on addressing the imbalanced problem by letting the classifier focus on hard misclassified examples and pay less attention to easy examples. A seminal hard negative method is focal loss [48], which reweights the standard cross-entropy loss $(-\log(p_t))$, where p_t is the predicted




Fig. 12: Vanilla’s cross-entropy loss and DebiAN’s RCE loss in (a) blond hair classification and (b) gender classification

probability of the ground-truth class) to $-\alpha_t(1-p_t)^\gamma \log(p_t)$, where α_t and γ are hyperparameters. Intuitively, it uses classifier’s predicted probability to reweigh the cross-entropy loss, where hard examples (*i.e.*, low p_t) are up-weighted with high $\alpha_t(1-p_t)^\gamma$ weights and easy examples (*i.e.*, high p_t) are down-weighted with low $\alpha_t(1-p_t)^\gamma$ weights. Different from focal loss, our RCE loss (Eq. 6) uses *discoverer*’s predicted bias group assignments to reweigh the cross-entropy loss, and the *discoverer* is trained with \mathcal{L}_{EOV} to differentiate *classifier*’s p_t on examples from the same target class where the Equal Opportunity is violated. Therefore, DebiAN is not a hard negative method because we do not use easy or hard samples (*i.e.*, low or high p_t) to perform reweighing, but rather use samples’ estimated bias group assignments to perform debiasing.

We also compare with focal loss on Multi-Color MNIST dataset. We use $\alpha = 0.25$ and $\gamma = 2.0$ as they perform the best in [48]. The results are shown in Tab. 17, where focal loss’s results are even worse than the vanilla model. The results prove that hard negative methods are not well-suited for debiasing. Since DebiAN is different from hard negative methods by using estimated bias group assignments to mitigate biases, our method achieves much better debiasing results.

H.2 Is EOv loss simply doing clustering?

No. EOv loss is used to train *discoverer* to classify different bias groups values. Therefore, instead of simply clustering *classifier*’s prediction, EOv loss guide the *discoverer* to do a classification for the bias group assignment. The results in Fig. 5, Fig. 8, and Fig. 9 show that *discoverer* can accurately classify if the samples on the testing set (unseen during training) are bias-aligned or bias-conflicting, demonstrating that EOv loss guides the *discoverer* to do classification based on different bias groups values and it can generalize to testing set’s images.

H.3 Will *classifier* achieve 100% accuracy such that *discoverer* cannot predict bias group assignments?

No. First, note that *discoverer*'s EOV loss is based on *classifier*'s predicted probabilities instead of thresholded hard predictions. Therefore, 100% accuracy does not indicate that *discoverer* cannot predict the bias group assignments. Second, we show the vanilla model's cross-entropy loss and DebiAN's RCE loss on blond hair classification and gender classification tasks in Fig. 12. The results show that the losses do not completely converge to zero, which proves that there always exist samples in the training set that *classifier* does not achieve 1.0 predicted probabilities. Therefore, it still leaves the room for *discoverer* to predict bias group assignments based on *classifier*'s different predictions on different samples.

H.4 What if the mini-batch only contains the samples from a single bias group?

It mainly happens under two conditions—(1) very strong spurious correlation; (2) small mini-batch size. For the first case, our ablation study on the ratios on Multi-Color MNIST dataset (Appendix D.2) shows that DebiAN achieves better debiasing results even when the ratio of `left color` is 0.995 (*i.e.*, very strong spurious correlation). For the second case, our ablation study on different batch sizes (Appendix C.2) shows that our method still achieves strong debiasing results when the batch size is small.

H.5 RCE loss compared with previous reweighing-based methods

LfF and focal loss are two previous reweighing-based methods for unsupervised debiasing. At a high level, LfF, focal loss, and DebiAN's RCE loss all target at up-weighting worse performed samples and down-weighting better-performed samples. The difference is how to compute the weight. LfF uses the ratio of cross-entropy loss between the biased model and the classifier to compute weights. Focal loss, as a hard negative method (see Appendix H.1), directly uses classifier's predicted probabilities to compute the weights. Different from previous methods, RCE loss uses *discoverer*'s predicted bias group assignment to compute the weights. Compared with LfF and focal, DebiAN achieves better debiasing results (Tab. 1-6 and Tab. 17).

H.6 Evaluation on Discovered Unknown Biases

In Fig. 6, Fig. 7, Fig. 10, and Fig. 11, we show some interesting unknown biases that human may not preconceive via saliency map. One may wonder if there exist other approaches to evaluate the results. First, it is hard to directly quantify the findings due to lack of annotations of the discovered bias attributes, *e.g.*, CelebA does not have attribute annotations or segmentation annotations of `visible hair area`. Using other datasets (*e.g.*, COCO [49]) with more attribute

or segmentation ground-truth may not help since the discovered unknown biases may still be out of the annotations. Second, UDIS [38], a recent bias discovery method, also uses saliency maps to interpret the bias. We believe that using saliency maps is an established evaluation protocol in this task. Third, although it is hard to evaluate bias discovery in real-world dataset, our evaluation of bias discovery on Multi-Color MNIST (Fig. 5, Fig. 8, and Fig. 9) has shown that DebiAN achieves strong bias discovery results. Finally, we believe that our better debiasing results w.r.t. **Hair Length** bias attribute on the Transects dataset can also indirectly prove that *discoverer* identifies **visible hair area** bias (see Appendix F).

H.7 Why not add two colors to the foreground in Multi-Color MNIST?

The reason is that foreground digits are not always well aligned to the center of the images. If we assign two colors to the foreground digit based on whether the foreground is on the left or right, we may encounter cases where the digit is mainly on the right and only has a tiny area on the left, *e.g.* an italic digit “1.” Thus, we choose to add colors to the background.

H.8 Difference between Multi-Color MNIST and Biased MNIST

Shrestha *et al.* [61] recently proposed the Biased MNIST dataset, which contains seven biases. However, all seven biases in Biased MNIST share the same bias-aligned ratio (*i.e.*, 0.7). In contrast, our Multi-color MNIST contains two biases that are in different bias-aligned ratios, which we believe is more common in real-world scenarios and can better reveal the failure modes of existing debiasing methods. For example, while LfF performs the best in the Biased MNIST benchmark, our Multi-Color MNIST dataset reveals that LfF can only discover the more salient bias—the bias with a larger bias-align ratio.

H.9 Why evaluate on scene classification task?

First, we regard that scene classification as a core vision task on par with object classification. Second, while many previous debiasing works create datasets [36,53] that contain a single bias (*e.g.*, artificially introducing the spurious correlation w.r.t. a single bias), we believe that the classical cross-dataset generalization evaluation approach [64] does not have the single-bias assumption. The subgroup distribution w.r.t. multiple biases may vary across different datasets, which is closer to the real-world setting.

H.10 Limitations and Future Directions

We list some limitations that DebiAN has not fully resolved. First, we only assume that the bias attribute is binary or continuously valued from 0 to 1 (*i.e.*, two

bias attribute groups). Future works can focus on extending DebiAN to discover and mitigate unknown biases with more than two groups. Second, DebiAN can only discover the biases caused by spurious correlation rather than lack of coverage. For example, suppose a face image dataset only contains long-hair female images and does not contain any short-hair female images. In that case, DebiAN cannot discover the `hair length` bias attribute because the *discoverer* does not have samples to categorize female images into two groups in terms of the `hair length` bias attribute. Finally, in terms of interpreting the discovered biases, DebiAN’s approach, using the saliency maps on real-world images, is not as easy as interpreting biases from synthesized counterfactual images [41,47]. Future works can further explore better interpreting the discovered unknown biases on real-world images.

H.11 Potential Negative Social Impact

One potential negative social impact is that DebiAN’s discovered biases could be used as a way to choose real-world images as the adversarial images to attack visual models in some safety-critical domains, *e.g.*, self-driving cars. Therefore, we encourage the defender to use DebiAN to mitigate the biases as the defense strategy.

Since our bias discovery approach relies on the fairness criterion based on equations, *e.g.*, equal true positive rates among two groups, our method cannot identify the biases that a fairness criterion cannot capture, *e.g.*, discrimination against the historically disadvantaged group. To mitigate this issue, we include a model card [52] in the released code to clarify that our method’s intended use case is discovering and mitigating biases that violate the equal opportunity fairness criterion [26], and the model’s out-of-scope use case is identifying or mitigating other biases that cannot be captured by the equation of a fairness criterion, *e.g.*, discrimination against the historically disadvantaged group [25].

Table 10: Ablation study on the ratios on Multi-Color MNIST dataset. Top-1 accuracy results are bolded, and the lowest accuracy results are underlined

left color ratio = 0.995	right color ratio = 0.95	vanilla	LfF	EIIL	PGI	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	<u>96.3</u> ± 0.5	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
bias-aligned	bias-conflicting	98.7 ± 0.6	<u>7.6</u> ± 1.0	98.4 ± 0.2	92.2 ± 11.1	98.1 ± 0.4
bias-conflicting	bias-aligned	<u>6.5</u> ± 1.0	96.5 ± 1.6	46.8 ± 0.5	27.7 ± 14.2	55.4 ± 2.1
bias-conflicting	bias-conflicting	<u>2.0</u> ± 0.4	5.9 ± 1.3	7.4 ± 0.1	6.3 ± 2.4	9.2 ± 0.8
unbiased		51.8 ± 0.2	<u>51.6</u> ± 0.6	63.2 ± 0.1	56.5 ± 4.3	65.7 ± 0.7
left color ratio = 0.99	right color ratio = 0.95	vanilla	LfF	EIIL	PGI	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	99.6 ± 0.5	100.0 ± 0.0	98.6 ± 2.3	100.0 ± 0.0
bias-aligned	bias-conflicting	97.1 ± 0.5	<u>4.7</u> ± 0.5	97.2 ± 1.5	82.6 ± 19.6	95.6 ± 0.8
bias-conflicting	bias-aligned	27.5 ± 3.6	98.6 ± 0.4	70.8 ± 4.9	<u>26.6</u> ± 5.5	76.5 ± 0.7
bias-conflicting	bias-conflicting	5.2 ± 0.4	<u>5.1</u> ± 0.4	10.9 ± 0.8	9.5 ± 3.2	16.0 ± 1.8
unbiased		57.4 ± 0.7	<u>52.0</u> ± 0.1	69.7 ± 1.0	54.3 ± 4.0	72.0 ± 0.8
left color ratio = 0.98	right color ratio = 0.95	vanilla	LfF	EIIL	PGI	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	99.0 ± 1.7	100.0 ± 0.0	<u>89.0</u> ± 19.0	100.0 ± 0.0
bias-aligned	bias-conflicting	96.6 ± 1.2	<u>9.7</u> ± 0.7	96.0 ± 0.3	78.6 ± 32.4	97.1 ± 0.8
bias-conflicting	bias-aligned	<u>64.4</u> ± 2.3	98.3 ± 0.9	84.0 ± 0.6	69.5 ± 27.7	85.1 ± 3.4
bias-conflicting	bias-conflicting	12.4 ± 1.1	<u>11.5</u> ± 1.1	16.0 ± 1.7	16.4 ± 1.1	19.4 ± 1.3
unbiased		68.3 ± 1.4	<u>54.6</u> ± 0.5	74.0 ± 0.5	63.42 ± 19.3	75.4 ± 0.9
left color ratio = 0.95	right color ratio = 0.95	vanilla	LfF	EIIL	PGI	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	93.4 ± 5.8	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
bias-aligned	bias-conflicting	91.1 ± 2.3	<u>71.1</u> ± 33.7	92.7 ± 0.5	76.5 ± 17.8	94.7 ± 0.9
bias-conflicting	bias-aligned	87.0 ± 3.7	<u>69.8</u> ± 25.9	90.0 ± 1.1	74.4 ± 17.7	92.7 ± 1.3
bias-conflicting	bias-conflicting	26.0 ± 1.3	39.6 ± 6.9	34.7 ± 3.3	<u>15.8</u> ± 4.7	39.6 ± 0.2
unbiased		76.0 ± 1.6	68.5 ± 3.2	79.3 ± 0.7	<u>66.7</u> ± 10.0	81.8 ± 0.6

Table 11: Ablation study on alternate training on Multi-Color MNIST dataset

left color ratio = 0.99	right color ratio = 0.95	w/o alternate training	DebiAN
bias-aligned	bias-aligned	100.0 ± 0.0	100.0 ± 0.0
bias-aligned	bias-conflicting	97.3 ± 0.1	95.6 ± 0.8
bias-conflicting	bias-aligned	74.0 ± 0.3	76.5 ± 0.7
bias-conflicting	bias-conflicting	11.5 ± 0.7	16.0 ± 1.8
unbiased accuracy		70.7 ± 0.1	72.0 ± 0.8

Table 12: Results of ablation study on *discoverer*’s outputs. Bolded methods are used to report results in the main paper. For binary classification (*i.e.*, age classification on bFFHQ dataset) task, there is no significant difference whether or not the *discoverer* predicts bias attribute groups per class or globally. When it comes to the multi-class digit classification on Multi-Color MNIST dataset, predicting bias attribute groups per class has better accuracy results on the images that are bias-conflicting w.r.t. both **left color** and **right color** bias attributes

dataset	#classes	per class	DebiAN (global)
bFFHQ	2	62.80 \pm 0.60	62.87 \pm 0.61
dataset	#classes	global	DebiAN (per class)
Multi-Color MNIST	10	13.5 \pm 0.1	16.0 \pm 1.8

Table 13: Results of an alternative design for debiasing in DebiAN – D and C play the minmax game (see DebiAN (minmax)) on Multi-Color MNIST dataset. Our RCE loss significantly outperforms the alternative design

left color ratio = 0.99	right color ratio = 0.95	DebiAN (minmax)	DebiAN (RCE)
bias-aligned	bias-aligned	97.3 \pm 4.6	100.0 \pm 0.0
bias-aligned	bias-conflicting	95.7 \pm 0.8	95.6 \pm 0.8
bias-conflicting	bias-aligned	64.5 \pm 2.0	76.5 \pm 0.7
bias-conflicting	bias-conflicting	7.7 \pm 1.9	16.0 \pm 1.8
unbiased accuracy		66.3 \pm 0.9	72.0 \pm 0.8

Table 14: Results on Colored MNIST (foreground color) under different ratios of bias-aligned samples in the training set. We bold top-1 results (except LfF’s results reported in the original paper since they cannot be replicated by their officially released code) and underline the lowest results based on the mean value. Although LfF achieves better bias-conflicting accuracy, it achieves lower bias-aligned accuracy, resulting in low unbiased accuracy. Overall, DebiAN achieves comparable or slightly lower unbiased accuracy results compared with other methods

ratio		vanilla	LfF (reported in the paper)	LfF (replicate via official code)	EIIL	PGI	DebiAN (Ours)
0.995	bias-aligned	100.00 \pm 0.00	-	54.13 \pm 6.33	99.90 \pm 0.01	99.86 \pm 0.11	100.00 \pm 0.00
	bias-conflicting	7.92 \pm 4.68	63.49 \pm 1.94	57.11 \pm 6.22	24.63 \pm 0.37	27.01 \pm 5.49	24.83 \pm 1.83
	unbiased accuracy	53.96 \pm 2.34	63.39 \pm 1.97	55.62 \pm 6.26	62.27 \pm 1.85	64.44 \pm 2.78	62.41 \pm 0.91
0.99	bias-aligned	99.97 \pm 0.05	-	61.96 \pm 3.29	99.81 \pm 0.19	99.86 \pm 0.07	99.86 \pm 0.05
	bias-conflicting	18.73 \pm 2.78	74.19 \pm 0.94	67.20 \pm 4.58	40.71 \pm 1.93	41.88 \pm 0.99	43.33 \pm 0.86
	unbiased accuracy	59.35 \pm 1.40	74.01 \pm 2.21	64.58 \pm 2.22	70.26 \pm 1.06	70.87 \pm 0.53	71.60 \pm 0.44
0.98	bias-aligned	99.80 \pm 0.16	-	73.11 \pm 5.41	99.77 \pm 0.14	99.73 \pm 0.11	99.76 \pm 0.05
	bias-conflicting	39.23 \pm 1.63	80.67 \pm 0.56	78.23 \pm 1.56	54.44 \pm 1.26	57.46 \pm 1.13	55.46 \pm 0.71
	unbiased accuracy	69.52 \pm 0.90	80.48 \pm 0.45	75.67 \pm 2.95	77.10 \pm 0.55	78.60 \pm 0.55	77.61 \pm 0.37
0.95	bias-aligned	99.60 \pm 0.17	-	71.67 \pm 0.68	99.61 \pm 0.25	99.37 \pm 0.31	99.70 \pm 0.17
	bias-conflicting	70.99 \pm 2.45	85.77 \pm 0.66	82.37 \pm 1.49	73.01 \pm 1.03	70.63 \pm 2.24	73.04 \pm 2.20
	unbiased accuracy	85.30 \pm 1.30	85.39 \pm 0.94	77.02 \pm 4.11	86.31 \pm 0.46	85.00 \pm 1.05	86.37 \pm 1.10

Table 15: Results on Colored MNIST (background color) under different ratios of bias-aligned samples in the training set. We bold top-1 results and underline the lowest results based on the mean value. Similar to the results in Colored MNIST (foreground color) (Tab. 14), although LfF achieves better bias-conflicting accuracy, it achieves lower bias-aligned accuracy, resulting in low unbiased accuracy. Overall, DebiAN achieves comparable or slightly higher unbiased accuracy results compared with other methods

ratio		vanilla	LfF	EIIL	PGI	DebiAN (Ours)
0.995	bias-aligned	99.97 \pm 0.05	<u>42.98\pm1.67</u>	100.00\pm0.00	100.00\pm0.00	100.00\pm0.00
	bias-conflicting	<u>1.98\pm0.35</u>	61.37\pm1.69	14.16 \pm 4.89	9.95 \pm 3.28	20.94 \pm 3.92
	unbiased accuracy	<u>50.98\pm0.18</u>	52.17 \pm 1.62	57.08 \pm 2.44	54.98 \pm 1.64	60.47\pm1.96
0.99	bias-aligned	100.00\pm0.00	<u>52.24\pm1.84</u>	99.97 \pm 0.05	99.26 \pm 1.10	100.00\pm0.00
	bias-conflicting	<u>5.43\pm0.32</u>	67.94\pm0.88	35.71 \pm 1.78	12.89 \pm 1.62	42.57 \pm 0.36
	unbiased accuracy	<u>52.75\pm0.16</u>	60.09 \pm 1.26	67.84 \pm 0.90	56.07 \pm 0.81	71.29\pm0.18
0.98	bias-aligned	99.97\pm0.05	<u>71.90\pm4.86</u>	99.90 \pm 0.09	97.83 \pm 3.54	99.90 \pm 0.00
	bias-conflicting	21.84 \pm 5.39	80.28\pm0.98	51.25 \pm 2.86	<u>18.06\pm5.47</u>	53.41 \pm 1.21
	unbiased accuracy	60.91 \pm 2.71	76.09 \pm 2.91	75.58 \pm 1.46	<u>59.75\pm4.02</u>	76.66\pm0.60
0.95	bias-aligned	99.87 \pm 0.15	<u>69.71\pm7.07</u>	99.90 \pm 0.16	96.78 \pm 4.90	99.90\pm0.01
	bias-conflicting	55.14 \pm 2.06	81.67\pm4.54	69.17 \pm 3.15	<u>33.90\pm15.55</u>	70.70 \pm 0.76
	unbiased accuracy	77.51 \pm 1.09	75.69 \pm 5.50	84.53 \pm 1.54	<u>65.34\pm9.80</u>	85.30\pm0.37

Table 16: Average group accuracy results of gender classification on Transects [8] dataset. All models are trained on CelebA dataset and evaluated on Transects w.r.t. two bias attributes—**Hair Length** and **Skin Color**. DebiAN achieves better results, which demonstrates that DebiAN better mitigate multiple biases simultaneously in the real-world multi-bias setting. Besides, it also reflects that DebiAN discovers **visible hair area** bias attribute to achieve better debiasing results w.r.t. **Hair Length** bias attribute

bias attribute	vanilla	LfF	EIIL	PGI	DebiAN (Ours)
Hair Length	55.1 \pm 5.8	54.7 \pm 2.9	54.0 \pm 0.4	56.2 \pm 1.3	60.5\pm1.7
Skin Color	53.5 \pm 5.3	53.3 \pm 2.9	53.1 \pm 0.08	57.4 \pm 0.3	60.1\pm1.2

Table 17: Comparing with focal loss [48], a hard negative method, on Multi-Color MNIST dataset

left color skew = 0.995	right color skew = 0.95	vanilla	focal	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
bias-aligned	bias-conflicting	98.7 ± 0.6	<u>97.9</u> ± 0.9	98.1 ± 0.4
bias-conflicting	bias-aligned	6.5 ± 1.0	<u>0.4</u> ± 0.2	55.4 ± 2.1
bias-conflicting	bias-conflicting	2.0 ± 0.4	<u>1.2</u> ± 0.3	9.2 ± 0.8
unbiased		51.8 ± 0.2	<u>49.2</u> ± 0.2	65.7 ± 0.7
left color skew = 0.99	right color skew = 0.95	vanilla	focal	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
bias-aligned	bias-conflicting	97.1 ± 0.5	95.7 ± 0.6	<u>95.6</u> ± 0.8
bias-conflicting	bias-aligned	27.5 ± 3.6	<u>3.3</u> ± 2.0	76.5 ± 0.7
bias-conflicting	bias-conflicting	5.2 ± 0.4	<u>2.4</u> ± 0.3	16.0 ± 1.8
unbiased		57.4 ± 0.7	<u>50.3</u> ± 0.4	72.0 ± 0.8
left color skew = 0.98	right color skew = 0.95	vanilla	focal	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
bias-aligned	bias-conflicting	96.6 ± 1.2	<u>85.1</u> ± 2.1	97.1 ± 0.8
bias-conflicting	bias-aligned	64.4 ± 2.3	<u>15.9</u> ± 4.4	85.1 ± 3.4
bias-conflicting	bias-conflicting	12.4 ± 1.1	<u>6.0</u> ± 0.1	19.4 ± 1.3
unbiased		68.3 ± 1.4	<u>51.7</u> ± 0.8	75.4 ± 0.9
left color skew = 0.95	right color skew = 0.95	vanilla	focal	DebiAN (Ours)
bias-aligned	bias-aligned	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
bias-aligned	bias-conflicting	91.1 ± 2.3	<u>63.7</u> ± 3.8	94.7 ± 0.9
bias-conflicting	bias-aligned	87.0 ± 3.7	<u>54.4</u> ± 4.1	92.7 ± 1.3
bias-conflicting	bias-conflicting	26.0 ± 1.3	<u>11.3</u> ± 0.1	39.6 ± 0.2
unbiased		76.0 ± 1.6	<u>57.3</u> ± 1.2	81.8 ± 0.6