

# Fair Visual Recognition in Limited Data Regime using Self-Supervision and Self-Distillation

Pratik Mazumder<sup>1,\*</sup> Pravendra Singh<sup>2,\*</sup> Vinay P. Namboodiri<sup>1,3</sup>

<sup>1</sup>IIT Kanpur, India    <sup>2</sup>IIT Roorkee, India    <sup>3</sup>University of Bath, United Kingdom

pratikm@cse.iitk.ac.in, pravendra.singh@cs.iitr.ac.in, vpn22@bath.ac.uk

## Abstract

*Deep learning models generally learn the biases present in the training data. Researchers have proposed several approaches to mitigate such biases and make the model fair. Bias mitigation techniques assume that a sufficiently large number of training examples are present. However, we observe that if the training data is limited, then the effectiveness of bias mitigation methods is severely degraded. In this paper, we propose a novel approach to address this problem. Specifically, we adapt self-supervision and self-distillation to reduce the impact of biases on the model in this setting. Self-supervision and self-distillation are not used for bias mitigation. However, through this work, we demonstrate for the first time that these techniques are very effective in bias mitigation. We empirically show that our approach can significantly reduce the biases learned by the model. Further, we experimentally demonstrate that our approach is complementary to other bias mitigation strategies. Our approach significantly improves their performance and further reduces the model biases in the limited data regime. Specifically, on the L-CIFAR-10S skewed dataset, our approach significantly reduces the bias score of the baseline model by 78.22% and outperforms it in terms of accuracy by a significant absolute margin of 8.89%. It also significantly reduces the bias score for the state-of-the-art domain independent bias mitigation method by 59.26% and improves its performance by a significant absolute margin of 7.08%.*

## 1. Introduction

Deep learning models have been used to solve several real-world problems, but they need to be trained on a large amount of labeled data. The training data often contains biases related to color, gender, appearance, and others. Standard deep learning training techniques capture such biases, and the resulting models make biased predictions. Researchers have proposed several approaches to mitigate bias in the deep learning models. These approaches assume the

availability of large amounts of training data. In this paper, we empirically demonstrate how a limited training data severely increases the impact of biases on the model predictions. We also demonstrate that the effectiveness of bias mitigation approaches is significantly reduced if the training data is limited. We propose a novel approach for addressing this problem that can also complement other bias mitigation strategies and improve their performance significantly.

The biases in model predictions are generally due to the same biases in the training data. This problem arises when the model finds an unwanted correlation in the data that minimizes the training loss. For example, if in a dataset the majority of the nurses are female, the model may become more inclined towards predicting a woman as a nurse compared to a man since the loss incurred by the model will be minimized. Further, if the dataset contains limited training data, any such biases will get reinforced. Therefore, the model may produce even more biased predictions when trained on limited data.

Some of the popular approaches for bias mitigation are the domain adversarial training [24, 23, 1], Reducing Bias Amplification [31], domain discriminative training [27] and domain independent training [27] approaches. The authors in [27] propose benchmark datasets for bias mitigation by deliberately introducing biases in the form of color/grayscale images, out-of-distribution images, scaled images, and cropped images in some classes. We propose limited data versions of these benchmark datasets [27] and study the performance of bias mitigation methods in the limited data regime.

We propose a novel approach that uses self-supervision and self-distillation to mitigate the bias in the model. Our approach involves training the network on an auxiliary self-supervision task and applying the self-distillation technique to the model. Self-supervision and self-distillation are generally used for improving the representation learning, and generalization ability of the model and are not regarded as bias mitigation approaches. The *novelty* of this work lies in adapting these two techniques for this setting and demonstrating that these techniques are very effective in bias mit-

\*The first two authors contributed equally.

igation (see Sec. 4.1). We empirically show that our approach significantly reduces the impact of biases in the network. Specifically, on the L-CIFAR-10S dataset, our approach significantly reduces the bias score of the baseline model by 78.22%. To the best of our knowledge, this is the *first work* to use such an approach to improve fairness. Further, our approach can be seamlessly integrated with other bias mitigation approaches and can significantly boost their bias mitigation ability in this setting. For example, on the L-CIFAR-10S dataset, our approach significantly reduces the bias score for the strategic sampling, adversarial, domain discriminative, and domain independent bias mitigation approaches by 89.81%, 82.04%, 65.95%, and 59.26%, respectively. We also demonstrate in Sec. 6.6 that our approach improves the model performance even when sufficient training data is present. We perform extensive ablations to validate our approach.

Our contributions can be summarized as follows:

- We empirically show that the effect of biases becomes even more adverse when the training data is limited.
- We empirically demonstrate that the performance of bias mitigation approaches suffers significantly when the available training data is limited.
- We demonstrate for the first time that self-distillation and self-supervision are very effective in bias mitigation. We propose a novel approach that adapts these techniques for mitigating bias in a limited data regime.
- We perform experiments on several benchmark datasets to validate our approach. We empirically show that our approach can augment other bias mitigation methods and significantly improve their performance in the limited data setting.

## 2. Related Works

### 2.1. Bias Mitigation Approaches

Researchers have proposed various approaches for dealing with biases arising from erroneous correlations present in the training dataset [2, 11, 26, 15, 17, 20, 25, 10]. The classes/attributes that are responsible for the spurious correlations are referred to as protected attributes. The methods proposed in [14, 28] aim to mitigate bias in simple linear models. The authors in [31] propose to mitigate bias by updating the inference to match a target distribution. InclusiveFaceNet [22] utilizes demographic information prior to perform better and less biased attribute detection. The domain discriminative approach described in [27] is also based on this approach and creates a new set of classes that incorporates the protected attribute/class. The method in [6] decouples classifiers in order to reduce the effect of biases on the network predictions. The authors in [27] propose the domain independent bias mitigation strategy that uses the same principle. This approach generates separate sets of

predictions for each value of the protected domain attribute to reduce the bias and then combines these predictions. Several researchers have also proposed an adversarial approach of bias mitigation [27, 1, 29, 7, 8] that trains the network to reduce its ability to identify the protected attribute/class.

The authors in [27] propose benchmark datasets for comparing bias mitigation strategies by deliberately introducing biases in the form of color/grayscale images, out-of-distribution images, scaled images, and cropped images in some classes. The work in [27] also compares various bias mitigation approaches on these benchmark datasets.

However, we demonstrate that the biases in model predictions become worse with limited training data. We also empirically show that the effectiveness of bias mitigation approaches significantly diminishes in the case of limited training data.

### 2.2. Self-Supervision

Self-supervision methods use techniques like modifying images to create artificial labels, modeling invariance through contrastive learning, and others [5, 30] to improve representation learning without requiring additional labeled data. The authors in [9] propose to rotate the image and train the network to predict the rotation angle as the self-supervised task. Contrastive learning based self-supervised tasks train the network to bring similar images closer and push dissimilar images away in the feature space. SimCLR [3] applies random data augmentations to produce two different views of the same image and performs contrastive learning using these views, and also uses the negative samples present in the batch. The authors in [4] propose the state-of-the-art SimSiam self-supervision technique. SimSiam also performs contrastive learning on two randomly augmented views of the same image without requiring negative samples or a large batch size. We use SimSiam as an auxiliary task in our approach to mitigate the bias in the network (see Sec. 4.1). We empirically show in Sec. 6.3 that SimSiam outperforms other self-supervised tasks when used as an auxiliary task in our approach.

### 2.3. Self-Distillation

The authors in [13] propose to use knowledge distillation for knowledge transfer from a teacher network to a student network. This process involves training the student network to match the output logits/soft predictions of the teacher network in addition to the training objective of the primary task. When the teacher and student architectures are the same, the knowledge distillation process is referred to as self-distillation. The authors in [19] demonstrate that self-distillation improves the test set performance of the network. The distillation process increases the generalization ability of the network without requiring additional labeled data for training. We use self-distillation in our approach to reduce the impact of biases in the model (see Sec. 4.1).

### 3. Effect of Biased Data on Model Predictions

The authors in [27] propose the CIFAR-10 Skewed (CIFAR-10S) dataset [27] for comparing various bias mitigation strategies. CIFAR-10S retains the 10 classes and 50,000 training images of the CIFAR-10 [16] dataset but modifies the training images to introduce a bias using two domains, i.e., color and grayscale. The color domain has five classes, each containing 95% color and 5% grayscale images (95-5% skew). The remaining five classes belong to the grayscale domain and contain 95% grayscale and 5% color images each (95-5% skew). Therefore, the training data of the color and grayscale domain classes are erroneously skewed towards the colored images and grayscale images, respectively. As a result, the domain can be considered as a protected attribute, and the model performance has to be de-biased with respect to this attribute. For evaluating the model, CIFAR-10S has two separate test datasets that basically are the color and grayscale copies of the entire CIFAR-10 test set, respectively. The CIFAR-10S dataset is designed in such a way that the test accuracy is indicative of the model bias, which helps to easily compare the fairness of different bias mitigation approaches [27].

#### 3.1. L-CIFAR-10S Dataset

We propose a limited data version (L-CIFAR-10S) of the CIFAR-10S dataset, obtained by reducing the number of images per class to 5%. However, we want to preserve the skewness (95-5%) in the dataset. Therefore, we choose the first 5% images from each class while maintaining the same level of skew, i.e., we separately choose the first 5% images from the color and grayscale images in that class. As a result, if a class contained 95% color images and 5% grayscale images or vice-versa, the ratio of color and grayscale images will still remain the same in L-CIFAR-10S. For evaluating the model, we use the same color and grayscale test sets of CIFAR-10S. Therefore, this new L-CIFAR-10S dataset is similar to the full-sized CIFAR-10S dataset except for the number of training images. As a result, similar to CIFAR-10S [27], the test accuracy over L-CIFAR-10S is also indicative of the model bias.

#### 3.2. Effect of Biases with Limited vs. Full Data

When we train the baseline ResNet-18 network on the L-CIFAR-10S dataset, the resulting model achieves 41.04% (see Table 1) classification accuracy on the color test set. This is significantly lower as compared to the 89.0% achieved by training on the CIFAR-10S dataset [27]. If we only consider this result, we may be tricked into assuming that this adverse effect on the baseline model accuracy is only due to the decrease in training data. However, if we convert all the training images in L-CIFAR-10S to grayscale and train the same ResNet-18 model using this data, the resulting model achieves a classification accuracy of 65.57%




Figure 1: Figure shows the confusion matrix of the colored test set predictions of a ResNet-18 model trained on (a) CIFAR-10S (b) L-CIFAR-10S. The confusion matrix indicates that when the training data is limited (L-CIFAR-10S) the biases present in the dataset have a significantly more severe effect on the predictions of the network. The bold labels denote the grayscale domain classes.

on the colored test set (GTCT:GrayTrainColorTest). Similarly, ResNet-18 model trained on an all color version (details are given in the supplementary material) of L-CIFAR-10S achieves 66.53% on the colored test set (CTCT:ColorTrainColorTest). Therefore, the bias in L-CIFAR-10S plays a major role in the significantly lower performance of the baseline (65.57/66.53%→41.04%). This is also evident from the bias score (lower better) of the baseline, which increases from 0.074 (for full training data, i.e., CIFAR-10S) to 0.349 in the limited data setting, which is a massive increase by 371.62%. Similar to the CIFAR-10S dataset, in the L-CIFAR-10S dataset, color domain classes have a majority of colored training images leading to an erroneous correlation in the model. But since the number of images in each class is low, the effect of this erroneous correlation is further magnified in the trained model, and it is even more inclined towards predicting one of the five color domain classes for any colored test image. As pointed out in [27], ideally, the network should perform at least as well as GTCT and approach the performance of CTCT. The gap in the performance of the baseline model w.r.t. GTCT for the L-CIFAR-10S dataset due to the bias in the data is 24.53%, which is significantly greater than the gap of 4% [27] w.r.t. GTCT (93.0%) for the CIFAR-10S dataset (details in the supplementary material). Therefore, with limited training data, the impact of the bias on the model predictions has significantly worsened (as can also be seen in Fig. 1). Our approach trained on L-CIFAR-10S achieves a classification accuracy of 53.31% (see Table 1) on the color test, which is significantly higher than the baseline and reduces this gap from 24.53% to 12.26%. Our approach also reduces the bias score of the baseline by 78.22% (from 0.349 to 0.076). This indicates that the self-supervision and self-distillation used in our approach are very effective in bias mitigation.

#### 3.3. Effect of Biases on Bias Mitigation Approaches with Limited vs. Full Data

We train bias mitigation approaches on the L-CIFAR-10S dataset. The domain discriminative training approach [27] achieves a classification accuracy of 41.47% (see Table

1) on the colored test set of L-CIFAR-10S. Therefore, the gap in its performance compared to the GTCT for the L-CIFAR-10S dataset is 24.1%, which is significantly greater than the gap of 1.8% [27] w.r.t. GTCT (93.0%) for the CIFAR-10S dataset (details in the supplementary material). In fact, the improvement due to the domain discriminative approach is negligible in this setting. Another important observation is that the bias score (lower better) of this approach increases from 0.040 (for full training data, i.e., CIFAR-10S) to 0.232 in the limited data setting, which is a huge increase by 480%. The domain independent training approach [27] achieves a classification accuracy of 59.25% (see Table 1) on the colored test set of L-CIFAR-10S. Therefore, the gap in its performance compared to GTCT for the L-CIFAR-10S dataset is 6.32%, which is significantly greater than the gap of 0.6% [27] w.r.t. GTCT (93.0%) for the CIFAR-10S dataset (details in the supplementary material). The domain independent approach also suffers from a 575% increase in bias score in the limited data setting. Therefore, with limited training data, bias mitigation approaches become significantly less effective. The severe impact of limited training data on the effectiveness of bias mitigation approaches is a *novel finding* of this work. The domain independent method augmented with our approach trained on L-CIFAR-10S achieves a classification accuracy of 66.42% (see Table 1) on the color test set, which outperforms GTCT and is very close to CTCT for the L-CIFAR-10S dataset. Since our approach significantly reduces the bias score of the domain independent method by 59.26% and improves its performance by a significant absolute margin of 7.17%, the effectiveness of our approach in bias mitigation is validated.

## 4. Proposed Approach

### 4.1. Motivation

Our proposed approach, Limited Data Bias Mitigation approach (LDBM), uses self-supervision and self-distillation for bias mitigation. However, self-supervision and self-distillation are used for improving the generalization ability and representation quality of the model. Therefore, these techniques help the model learn useful/meaningful and discriminative features from the training data, thereby minimizing the impact of irrelevant features on the model. Consequently, if the training data suffers from biases due to any unwanted correlations, these techniques will also prevent the model from learning such biases, which is the objective of bias mitigation. In the limited data setting, this problem becomes even more severe, and any biases in the model, including those resulting from unwanted correlations in the training data, will get reinforced. This is why the baseline model and other bias mitigation approaches suffer from a significant impact of biases in this setting. However, due to the reasons mentioned above, our

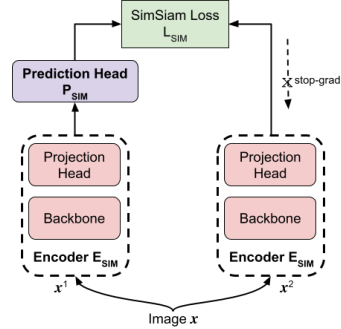


Figure 2: Figure shows the SimSiam auxiliary self-supervised task that trains the network to bring two augmented views of the same image close to each other in the feature space. Please note that, at a time the prediction head  $P_{SIM}$  is applied to the encoder output of either the first view  $x^1$  or the second view  $x^2$  and gradient flow is prevented through the encoder output of the other view [4].

approach is effective in mitigating various types of biases in the baseline model and also significantly improves the effectiveness of the bias mitigation approaches in this setting. The self-supervised auxiliary task and the self-distillation process used in our approach are described as follows.

### 4.2. Training with a Self-Supervised Auxiliary Task

The authors in [4] explain that SimSiam uses a siamese architecture for modeling invariance, which is a focus of representation learning. We employ a self-supervision training technique as an auxiliary task in our approach in order to reduce the biases in the model by improving the representation quality of the network as discussed in Sec. 4.1. In our approach, we use the SimSiam [4] self-supervision task as an auxiliary task. We perform ablations in Sec. 6.3 to validate this choice. SimSiam [4] is a state-of-the-art self-supervised learning technique. It trains the network to bring two slightly different views of the same image closer in the representation space of the network. The two views denote two different sets of data augmentation techniques applied to the same image. The two views of the same image are processed by an encoder ( $E_{SIM}$ ) consisting of a backbone network and a multi-layer perceptron projection head. The encoder shares the same parameter weights between the two views. Let  $x_i^1, x_i^2$  refer to two randomly augmented views of the same image  $x_i$ .

$$z_i^2 = E_{SIM}(x_i^2) \quad (1)$$

Where,  $z_i^2$  refers to the encoder output for  $x_i^2$ .

A multi-layer perceptron based prediction head ( $P_{SIM}$ ) transforms the encoder output of the first view to match the second view (target view). To achieve this objective SimSiam minimizes the negative cosine similarity between the two views (as shown in Fig. 2). The cosine similarity loss is implemented in such a way that the target view encoder output is treated as a constant to prevent gradient flow (stop-gradient) through it [4].

$$p_i^1 = P_{SIM}(E_{SIM}(x_i^1)) \quad (2)$$






Figure 3: Figure shows the self-distillation process from a trained teacher network to the student network.

$$D(p_i^1, z_i^2) = -\frac{p_i^1 \cdot z_i^2}{\|p_i^1\|_2 \cdot \|z_i^2\|_2} \quad (3)$$

Where,  $p_i^1$  refers to the prediction MLP output for  $x_i^1$ .  $\|\cdot\|_2$  refers to  $\ell_2$ -norm.  $D$  is the negative cosine similarity loss.

The authors in [4] also perform a similar operation using  $P_{\text{SIM}}$  to transform the second view to match the first view (target view) in order to achieve a symmetrized loss.

$$L_{\text{SIM}}(x_i^1, x_i^2) = \frac{1}{2}D(p_i^1, z_i^2) + \frac{1}{2}D(p_i^2, z_i^1) \quad (4)$$

Where,  $p_i^2 = P_{\text{SIM}}(E_{\text{SIM}}(x_i^2))$  refers to the prediction MLP output for  $x_i^2$ .  $z_i^1 = E_{\text{SIM}}(x_i^1)$  refers to the encoder output for  $x_i^1$ .  $L_{\text{SIM}}$  is the SimSiam loss.

In our approach, the entire network, apart from the last classification layer, serves as the backbone network. We add a multi-layer perceptron projection head and a multi-layer perceptron based prediction head to our network for the SimSiam self-supervised auxiliary task.

### 4.3. Self Distillation from a Trained Teacher

The self-distillation process trains a student network that has the same architecture as the trained teacher network. The authors in [18] explain how the self-distillation process improves the generalization of the network. We use the self-distillation process to reduce the impact of biases in the model by improving the generalization power of the network as discussed in Sec. 4.1. The self-distillation process also does not require any additional data making it very suitable for our approach.

In the self-distillation process, we minimize the Kullback–Leibler(KL) divergence between the logits/soft probability outputs of the teacher and student networks (as shown in Fig. 3). This process transfers knowledge from the teacher to the student network and also improves the generalization capacity of the student network. The loss function for self-distillation is defined as follows:

$$o_i, o_i^T = \Theta_S(x_i), \Theta_T(x_i) \quad (5)$$

$$L_{\text{SD}}(o_i, o_i^T) = \text{KL-Div}(\sigma(o_i), \sigma(o_i^T)) \quad (6)$$

Where,  $x_i$  denotes a training data point,  $\Theta_S, \Theta_T$  denotes the student and teacher networks, respectively.  $o_i, o_i^T$  denote the logits produced by the student and teacher network for input  $x_i$ .  $\sigma(\cdot)$  denotes the softmax activation function

that transforms the logit  $l_c$  for each class  $c$  into a probability by comparing it with logits  $l_j$  of the remaining classes, i.e.,  $\sigma(l_c) = \frac{\exp^{l_c/\kappa}}{\sum_j \exp^{l_j/\kappa}}$ .  $\kappa$  denotes the temperature hyperparameter [13].  $\text{KL-Div}$  refers to the KL divergence.  $L_{\text{SD}}$  refers to the self distillation loss.

In our proposed limited data bias mitigation approach (LDBM), we train a teacher network  $\Theta_T$  using the cross-entropy loss for the standard classification task and the SimSiam auxiliary self-supervised task loss. The loss function for training the teacher network is defined as follows:

$$L_T = \frac{\sum_{i=1}^N \{L_{\text{CE}}(x_i, y_i) + L_{\text{SIM}}(x_i^1, x_i^2)\}}{N} \quad (7)$$

Where,  $L_{\text{CE}}$  represents the cross-entropy loss,  $L_{\text{SIM}}$  represents the SimSiam loss,  $N$  refers to the number of training samples,  $y_i$  refers to the label of image  $x_i$ , and  $L_T$  refers to the total loss for the teacher network.

Next, we train a student network  $\Theta_S$  that has the same architecture as the teacher network  $\Theta_T$  using the cross-entropy loss for the standard classification task, the SimSiam auxiliary self-supervised task loss, and the self distillation loss from the trained teacher network  $\Theta_T$ . For the self-distillation loss, we minimize the KL divergence between soft predictions of  $\Theta_S$  and  $\Theta_T$  for the same input image. The total loss  $L_S$  for training the student network is defined as follows:

$$L_S = \frac{\sum_{i=1}^N \{L_{\text{CE}}(x_i, y_i) + L_{\text{SIM}}(x_i^1, x_i^2) + L_{\text{SD}}(o_i, o_i^T)\}}{N} \quad (8)$$

After the training is completed, the resulting student network is our final model.

## 5. Experiments

### 5.1. Datasets

We perform the experiments on the limited L-CIFAR-10S skewed dataset described in Sec. 3.1, which is created from the CIFAR-10S dataset [27] with the protected attribute/domain being the color/grayscale image. The authors in [27] also create modified versions of CIFAR-10S by changing the domain attribute. The new domains consist of normal CIFAR-10 images and either a) CIFAR-10S-i: images of the same class from the ImageNet [21] dataset (CIFAR-10S-i) downsampled to  $32 \times 32$ , or, b) CIFAR-10S-c28: images cropped to  $28 \times 28$  and upsampled to  $32 \times 32$ , or, c) CIFAR-10S-d16: images downsampled to  $16 \times 16$  and upsampled to  $32 \times 32$ , or d) CIFAR-10S-d8: images downsampled to  $8 \times 8$  and upsampled to  $32 \times 32$ . All these datasets have the same bias problem as in CIFAR-10S and L-CIFAR-10S. We create the limited data versions of all these datasets, i.e., L-CIFAR-10S-i, L-CIFAR-10S-c28, L-CIFAR-10S-d16, and L-CIFAR-10S-d8, using the same approach as described in Sec. 3.1 to reduce the training data to 5% without changing the level of skew. We perform experiments using the 95-5% skew (described in Sec. 3) for all these datasets. We also perform experiments on the CelebA

multi-labeled dataset after reducing the training set to the first 5% images (L-CelebA). Further details regarding these datasets have been provided in the supplementary material.

## 5.2. Compared Approaches

We perform experiments on the naive baseline method in the limited data setting. It trains using the cross-entropy loss without applying any bias mitigation approaches. We also perform experiments on the strategic sampling, adversarial, domain discriminative, and domain independent approaches described in [27]. The strategic sampling over samples the minority image type in each class, e.g., color images in grayscale domain classes and vice-versa. The adversarial method minimizes the ability of the network to identify the protected attribute/domain, e.g., whether the image is colored or grayscale. The domain discriminative method trains an ND-way classifier where N is the number of image classes and D is the number of domains (values that the protected attribute can take). The domain independent method trains separate N-way classifiers for each domain D. We have provided details regarding these approaches in the supplementary material.

## 5.3. Implementation Details

In order to use SimSiam as an auxiliary task in our approach, we use the ResNet-18 architecture as the backbone network for the L-CIFAR-10S experiments and ResNet-50 architecture (pretrained on ImageNet) as the backbone network for the L-CelebA experiments. For complete details of the training settings refer to the supplementary material.

For the L-CIFAR-10S experiments, we use two metrics also used by [27], i.e., mean per-class per-domain accuracy (primary) and bias amplification [31]. Since the overall test set is fully balanced across the two types of image in each version of L-CIFAR-10S, the mean accuracy reflects the effect of biases on the predictions of the model [27]. For better understanding and completeness w.r.t. to [27], we also provide the mean bias score, defined as follows:

$$\frac{1}{|C|} \sum_{c \in C} \frac{\max(\text{Gr}_c, \text{Col}_c)}{\text{Gr}_c + \text{Col}_c} - 0.5. \quad (9)$$

Where,  $\text{Gr}_c, \text{Col}_c$  refer to the number of grayscale and color images predicted to belong to the class  $c \in C$ .

For the L-CelebA experiments, we perform multi-label classification that generally uses the metric of mean average precision (mAP) across attributes. We use a weighted mAP proposed in [27] to remove the gender bias in the test set. Following [27], we also provide the bias amplification (BA) over the attributes [31]. For further details, refer to the supplementary material.

## 5.4. L-CIFAR-10S Results

Tables 1, 2, 3, 4, 5 compare the performance of different bias mitigation strategies including our proposed LDBM

Model Name	Model	Test Inference	Bias (↓)	Accuracy (% ↑)		
				Color	Gray	Mean
Baseline	N-way	$\arg \max_y P(y x)$	0.349	41.04	42.11	41.58 ± 0.6
<b>w/ LDBM (Ours)</b>	N-way	$\arg \max_y P(y x)$	0.076	53.31	47.63	50.47 ± 0.7
Sampling	N-way	$\arg \max_y P(y x)$	0.265	50.96	52.43	51.7 ± 0.8
<b>w/ LDBM (Ours)</b>	N-way	$\arg \max_y P(y x)$	0.027	56.96	57.15	57.01 ± 0.7
Adversarial	w/ unf. conf.	$\arg \max_y P(y x)$	0.362	40.98	44.24	42.61 ± 1.4
	w/ ∇ rev. proj.	$\arg \max_y P(y x)$	0.324	40.19	45.09	42.64 ± 1.4
<b>w/ LDBM (Ours)</b>	w/ unf. conf.	$\arg \max_y P(y x)$	0.065	54.42	50.99	52.71 ± 1.3
DomainDiscrim	joint ND-way	$\arg \max_y \sum_d P_{tr}(y, d x)$	0.390	38.05	36.72	37.39 ± 0.7
		$\arg \max_y \max_d P_{te}(y, d x)$	0.232	41.38	41.3	41.34 ± 0.9
		$\arg \max_y \sum_d P_{te}(y, d x)$	0.232	41.47	41.33	41.4 ± 0.8
RBA [31]	$y = \mathcal{L}(\sum_d P_{tr}(y, d x))$	0.252	38.21	38.23	38.22 ± 0.9	
<b>w/ LDBM (Ours)</b>	joint ND-way	$\arg \max_y \sum_d P_{te}(y, d x)$	0.079	57.75	53.81	55.78 ± 0.7
DomInd	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.319	50.83	46.25	48.54 ± 0.6
		$\arg \max_y \sum_d s(y, d, x)$	0.027	59.25	58.39	58.64 ± 1.4
<b>w/ LDBM (Ours)</b>	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.333	53.9	50.92	52.41 ± 0.5
		$\arg \max_y \sum_d s(y, d, x)$	<b>0.011</b>	<b>66.42</b>	<b>65.37</b>	<b>65.90 ± 0.6</b>

Table 1: Performance comparison of bias mitigation strategies on L-CIFAR-10S using the ResNet-18 [12].

Model Name	Model	Test Inference	Bias (↓)	Accuracy (% ↑)		
				Color	ImgNet	Mean
Baseline	N-way	$\arg \max_y P(y x)$	0.289	47.94	36.5	42.22 ± 0.5
<b>w/ LDBM (Ours)</b>	N-way	$\arg \max_y P(y x)$	0.091	66.56	49.27	57.92 ± 0.4
Sampling	N-way	$\arg \max_y P(y x)$	0.208	50.16	40.11	45.14 ± 0.5
Adversarial	w/ unf. conf.	$\arg \max_y P(y x)$	0.257	50.32	39.28	44.8 ± 1.2
	w/ ∇ rev. proj.	$\arg \max_y P(y x)$	0.243	49.76	38.29	44.03 ± 1.3
DomainDiscrim	joint ND-way	$\arg \max_y \sum_d P_{tr}(y, d x)$	0.330	48.5	36.79	42.64 ± 0.7
		$\arg \max_y \max_d P_{te}(y, d x)$	0.203	51.94	39.64	45.79 ± 0.8
		$\arg \max_y \sum_d P_{te}(y, d x)$	0.203	52.12	39.73	45.92 ± 0.6
RBA [31]	$y = \mathcal{L}(\sum_d P_{tr}(y, d x))$	0.213	50.88	37.68	44.28 ± 0.6	
DomInd	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.381	42.62	34.30	38.46 ± 0.7
		$\arg \max_y \sum_d s(y, d, x)$	0.038	51.11	42.18	46.64 ± 0.5
<b>w/ LDBM (Ours)</b>	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.364	51.23	39.24	45.24 ± 0.6
		$\arg \max_y \sum_d s(y, d, x)$	<b>0.022</b>	<b>63.55</b>	<b>47.93</b>	<b>55.74 ± 0.6</b>

Table 2: Performance comparison of bias mitigation strategies on L-CIFAR-10S-i using the ResNet-18 [12].

Model Name	Model	Test Inference	Bias (↓)	Accuracy (% ↑)		
				Color	Crop	Mean
Baseline	N-way	$\arg \max_y P(y x)$	0.025	42.47	41.93	42.2 ± 0.6
<b>w/ LDBM (Ours)</b>	N-way	$\arg \max_y P(y x)$	0.016	51.48	53.31	52.39 ± 0.4
Sampling	N-way	$\arg \max_y P(y x)$	0.032	47.75	47.46	47.61 ± 0.6
Adversarial	w/ unf. conf.	$\arg \max_y P(y x)$	0.037	44.87	43.78	44.325 ± 1.2
	w/ ∇ rev. proj.	$\arg \max_y P(y x)$	0.029	45.95	43.96	44.96 ± 1.1
DomainDiscrim	joint ND-way	$\arg \max_y \sum_d P_{tr}(y, d x)$	0.024	39.55	38.67	39.11 ± 0.5
		$\arg \max_y \max_d P_{te}(y, d x)$	0.025	42.71	42.05	42.38 ± 0.6
		$\arg \max_y \sum_d P_{te}(y, d x)$	0.025	42.78	42.08	42.43 ± 0.8
RBA [31]	$y = \mathcal{L}(\sum_d P_{tr}(y, d x))$	0.023	39.58	38.67	39.13 ± 0.6	
DomInd	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.336	47.87	41.59	44.73 ± 0.4
		$\arg \max_y \sum_d s(y, d, x)$	0.035	56.34	55.39	55.86 ± 0.3
<b>w/ LDBM (Ours)</b>	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.322	53.63	53.1	53.37 ± 0.6
		$\arg \max_y \sum_d s(y, d, x)$	<b>0.013</b>	<b>65.14</b>	<b>67.5</b>	<b>66.32 ± 0.4</b>

Table 3: Performance comparison of bias mitigation strategies on L-CIFAR-10S-c28 using the ResNet-18 [12].

approach on L-CIFAR-10S based datasets. The results indicate that the performance of the baselines and all bias mitigation strategies have been significantly affected in the limited data setting (as compared to the full data results reported in [27], Fig. 7). We observe that our proposed LDBM approach significantly improves the performance of the baseline model for all the datasets. Specifically, for L-CIFAR-10S with 95-5% skew, LDBM significantly reduces the bias score of the baseline model by 78.22% and

Model Name	Model	Test Inference	Bias (↓)	Accuracy (% , ↑)	
				Color	Mean
Baseline	N-way	$\arg \max_y P(y x)$	0.146	44.79	36.61 40.7 ± 0.7
<b>w/ LDBM (Ours)</b>	N-way	$\arg \max_y P(y x)$	0.065	49.26	46.89 48.075 ± 0.8
Sampling	N-way	$\arg \max_y P(y x)$	0.167	44.84	41.38 43.11 ± 0.6
Adversarial	w/ unif. conf.	$\arg \max_y P(y x)$	0.136	44.42	38.0 41.21 ± 1.1
	w/ $\nabla$ rev., proj.	$\arg \max_y P(y x)$	0.146	44.18	38.54 41.36 ± 1.4
DomainDiscrim	joint ND-way	$\arg \max_y \sum_d P_{tr}(y, d x)$	0.123	43.33	33.99 38.66 ± 0.7
		$\arg \max_y \max_d P_{te}(y, d x)$	0.102	46.87	35.92 41.39 ± 0.6
		$\arg \max_y \sum_d P_{te}(y, d x)$	0.103	46.94	36.08 41.51 ± 0.6
	RBA [31]	$y = \mathcal{L}(\sum_d P_{tr}(y, d x))$	0.067	43.37	33.99 38.68 ± 0.7
DomInd	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.366	45.51	35.27 40.39 ± 0.7
		$\arg \max_y \sum_d s(y, d, x)$	0.088	56.34	55.39 51.5 ± 0.6
<b>w/ LDBM (Ours)</b>	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.342	49.93	47.85 48.89 ± 0.8
		$\arg \max_y \sum_d s(y, d, x)$	<b>0.026</b>	<b>63.18</b>	<b>59.8</b> <b>61.49 ± 0.6</b>

Table 4: Performance comparison of bias mitigation strategies on L-CIFAR-10S-d16 using the ResNet-18 [12].

Model Name	Model	Test Inference	Bias (↓)	Accuracy (% , ↑)	
				Color	Mean
Baseline	N-way	$\arg \max_y P(y x)$	0.242	41.03	27.74 34.39 ± 0.7
<b>w/ LDBM (Ours)</b>	N-way	$\arg \max_y P(y x)$	0.175	47.89	33.37 40.63 ± 0.6
Sampling	N-way	$\arg \max_y P(y x)$	0.237	40.09	32.78 36.44 ± 0.8
Adversarial	w/ unif. conf.	$\arg \max_y P(y x)$	0.220	41.05	27.21 34.13 ± 1.4
	w/ $\nabla$ rev., proj.	$\arg \max_y P(y x)$	0.237	41.13	25.49 33.31 ± 1.3
DomainDiscrim	joint ND-way	$\arg \max_y \sum_d P_{tr}(y, d x)$	0.219	39.17	25.75 32.46 ± 0.7
		$\arg \max_y \max_d P_{te}(y, d x)$	0.137	42.81	27.28 35.04 ± 0.4
		$\arg \max_y \sum_d P_{te}(y, d x)$	0.136	42.79	27.41 35.1 ± 0.5
	RBA [31]	$y = \mathcal{L}(\sum_d P_{tr}(y, d x))$	0.133	39.25	25.78 32.52 ± 0.7
DomInd	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.411	41.62	28.48 35.05 ± 0.6
		$\arg \max_y \sum_d s(y, d, x)$	0.110	49.09	37.64 43.37 ± 0.3
<b>w/ LDBM (Ours)</b>	N-way per dom.	$\arg \max_y P_{te}(y d^*, x)$	0.403	47.11	33.1 40.11 ± 0.7
		$\arg \max_y \sum_d s(y, d, x)$	<b>0.103</b>	<b>56.19</b>	<b>43.17</b> <b>49.68 ± 0.4</b>

Table 5: Performance comparison of bias mitigation strategies on L-CIFAR-10S-d8 using the ResNet-18 [12].

improves its performance by a significant absolute margin of 8.89% in the mean accuracy. LDBM also significantly reduces the bias score for the state-of-the-art domain independent bias mitigation method by 59.26% and improves its performance by a significant absolute margin of 7.08% in the mean accuracy. In fact, the domain independent training approach with our proposed LDBM achieves a slightly better performance on the color test set than even the model trained on all grayscale version of L-CIFAR-10S (GTCT 65.57%) and very close to the performance of the all color version of L-CIFAR-10S (CTCT 66.53%) as described in Sec. 3.2. LDBM also significantly reduces the bias scores for the strategic sampling, adversarial, and domain discriminative bias mitigation approaches by 89.81%, 82.04%, 65.95%, respectively, and significantly improves their performance by absolute margins of 5.31%, 10.1%, and 14.38%, respectively, on L-CIFAR-10S.

For the other L-CIFAR-10S variants, we use our approach to augment the state-of-the-art domain independent bias mitigation technique. We observe in Table 2 that for the L-CIFAR-10S-i dataset, LDBM significantly reduces the bias score of the baseline model by 68.51% and improves its performance by a significant absolute margin of 15.7%. It also significantly reduces the bias score of the domain independent bias mitigation method by 42.11% and improves its performance by a significant absolute margin of 9.1%.

Model Name	Model	mAP (↑)	BA (↓)
Base	N sigmoids	65.34	0.036
<b>w/ LDBM (Ours)</b>	N sigmoids	66.72	0.017
Adversarial DomainDiscrim	w/uniform conf. [1, 24]	56.82	0.070
	2N sigm, $\sum_d P_{tr}(y, d x)$	64.71	0.020
DomInd	2N sigmoids, $P_{tr}(y d^*, x)$	63.77	0.027
	2N sigm, $\max_d P_{tr}(y d, x)$	64.62	-0.048
	2N sigm, $\sum_d P_{tr}(y d, x)$	65.35	-0.006
	2N sigmoids, $\sum_d s(y, d, x)$	65.73	-0.007
<b>w/ LDBM (Ours)</b>	2N sigmoids, $P_{tr}(y d^*, x)$	66.36	0.024
	2N sigm, $\max_d P_{tr}(y d, x)$	67.92	-0.051
	2N sigm, $\sum_d P_{tr}(y d, x)$	68.49	<b>-0.073</b>
	2N sigmoids, $\sum_d s(y, d, x)$	<b>69.07</b>	-0.019

Table 6: Performance of bias mitigation approaches on the L-CelebA dataset using ResNet-50.




Figure 4: Experimental results showing the test loss (cross-entropy loss) of the network for baseline with or without LDBM trained on L-CIFAR-10S using ResNet-18.




Figure 5: Experimental results showing the decrease in bias amplification scores with the decrease in test loss (improved generalization) of the network when baseline w/ LDBM is trained on L-CIFAR-10S using ResNet-18.

Therefore, our proposed LDBM significantly reduces the effect of biases on the model predictions and significantly improves the performance of bias mitigation strategies.

## 5.5. L-CelebA Results

Tables 6 compares the performance of different bias mitigation strategies, including our proposed LDBM approach on the L-CelebA dataset. We observe that the effectiveness of the bias mitigation strategies diminishes as compared to the full data results reported in [27]. We also observe that our proposed LDBM approach significantly improves the performance of the baseline approach and the domain-independent training approach with respect to average mAP across attributes and the bias amplification score.

## 6. Ablation Experiments

### 6.1. Bias Amplification Score over Training Epochs

Fig. 4 compares the cross-entropy loss on the L-CIFAR-10S test set using the baseline approach and baseline w/ LDBM approach across the training epochs. The results indicate that the baseline w/ LDBM encounters lower test loss than the baseline. Therefore, the generalization power

Self Sup.	Self-Distill	Bias ( $\downarrow$ )	Color Acc. %.( $\downarrow$ )	Gray Acc. %.( $\uparrow$ )	Mean Acc. %.( $\uparrow$ )
$\times$	$\times$	0.027	59.25	58.39	58.82 $\pm$ 0.7
$\times$	$\checkmark$	0.024	61.23	58.96	60.10 $\pm$ 0.6
$\checkmark$	$\times$	0.013	65.20	64.24	64.72 $\pm$ 0.5
$\checkmark$	$\checkmark$	<b>0.011</b>	<b>66.42</b>	<b>65.37</b>	<b>65.90<math>\pm</math>0.6</b>

Table 7: Ablation experiment showing the significance of different components of our approach applied to the domain independent method on L-CIFAR-10S.

Self Supervision	Bias ( $\downarrow$ )	Color Acc. %.( $\downarrow$ )	Gray Acc. %.( $\uparrow$ )	Mean Acc. %.( $\uparrow$ )
Rotation	0.026	59.85	59.41	59.63 $\pm$ 0.8
SimClr	0.017	62.81	62.03	62.42 $\pm$ 0.5
<b>SimSiam</b>	<b>0.011</b>	<b>66.42</b>	<b>65.37</b>	<b>65.90<math>\pm</math>0.6</b>

Table 8: Ablation experiment showing the performance of different self-supervision techniques used as an auxiliary task along with self-distillation in our approach LDBM applied to the domain independent method on L-CIFAR-10S.




Figure 6: Figure shows the performance of the baseline, domain independent and domain independent w/ LDBM methods for different levels of skew in L-CIFAR-10S.

and representation quality of the model obtained using our approach is better than the baseline model. Fig. 5 shows the decrease in the bias amplification score as the model gets trained. Figs. 4, 5 indicate that our LDBM approach reduces the bias score of the model by improving the generalization ability and representation quality of the model.

## 6.2. Significance of Components

The results in Table 7 indicate that both the auxiliary self-supervision and the self-distillation components of our approach are essential to the performance of our approach.

## 6.3. Choice of Self-Supervision

The results in Table 8 indicate that the SimSiam based self-supervised auxiliary task is more effective in our approach as compared to other techniques. The contrastive training approach in the SimSiam and SimClr techniques helps the network to learn better features from the images.

## 6.4. Skew Level vs. Accuracy

The results in Fig. 6 indicate that as the level of skew increases, the performance of the baseline decreases significantly due to an increase in the bias in the data. However, the performance of our approach remains relatively stable and significantly higher than the others.

## 6.5. LDBM complements Bias Mitigation Methods

The results in Table 1 indicate that our proposed LDBM can complement other bias mitigation approaches, and sig-




Figure 7: Figure shows the performance (mean of color & grayscale accuracy) of the baseline, baseline w/ LDBM, domain independent and domain independent w/ LDBM methods on the CIFAR-10S dataset.

nificantly reduces the bias scores of these approaches in the limited data setting. Specifically, it significantly reduces the bias score for the strategic sampling, adversarial, domain discriminative, and domain independent bias mitigation approaches by 89.81%, 82.04%, 65.95%, and 59.26%, respectively on L-CIFAR-10S. It also significantly improves their performance in the limited data setting.

## 6.6. Performance improvement on CIFAR-10S

The results in Fig. 7 indicate that our proposed approach outperforms the baseline method by an absolute margin of 2.55% on the CIFAR-10S dataset, which has a sufficient number of training examples per class. The state-of-the-art domain independent training approach augmented with our proposed LDBM achieves around 1% improvement as compared to 7.08% in the limited data setting. In this setting, since the network is trained on a sufficient amount of data, it is less prone to learning irrelevant features and unwanted correlations from the training data. However, our approach still improves the performance of the network by further reducing the impact of the biases. In fact, the domain independent method w/ LDBM (93.2%) slightly outperforms GTCT (93.0%) on the color test set.

## 7. Conclusion

In this paper, we show that when the number of training examples is limited, the bias in the dataset has a severe impact on the model performance. We also experimentally demonstrate that bias mitigation approaches become less effective with limited training data. We propose a novel approach for bias mitigation. Specifically, we show how self-supervision and self-distillation, which are not used for bias mitigation, can be adapted to significantly reduce the impact of biases on the model, especially in the limited data regime. We empirically demonstrate how our proposed LDBM significantly reduces the bias score of the baseline model and other bias mitigation approaches.

## 8. Acknowledgements

Assistance through SERB grant SERB-CS-2019179 is duly acknowledged.



## References

- [1] Mohsan Alvi, Andrew Zisserman, and Christoffer Nellaker. Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In *ECCV*, 2018.
- [2] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *ECCV*, 2018.
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pages 1597–1607. PMLR, 2020.
- [4] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. *arXiv preprint arXiv:2011.10566*, 2020.
- [5] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 766–774, 2014.
- [6] Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Max Leiserson. Decoupled classifiers for group-fair and efficient machine learning. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, 2018.
- [7] Harrison Edwards and Amos Storkey. Censoring Representations with an Adversary. In *ICLR*, 2016.
- [8] Yaroslav Ganin and Victor Lempitsky. Unsupervised Domain Adaptation by Backpropagation. In *ICML*, 2015.
- [9] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018.
- [10] Aditya Grover, Kristy Choi, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. *arXiv preprint arXiv:1910.12008*, 2019.
- [11] Aditya Grover, Jiaming Song, Ashish Kapoor, Kenneth Tran, Alekh Agarwal, Eric J Horvitz, and Stefano Ermon. Bias correction of learned generative models using likelihood-free importance weighting. In *NeurIPS*, 2019.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016.
- [13] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2014.
- [14] Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. Undoing the Damage of Dataset Bias. In *ECCV*, 2012.
- [15] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *CVPR*, 2019.
- [16] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [17] Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *CVPR*, 2019.
- [18] Hossein Mobahi, Mehrdad Farajtabar, and Peter Bartlett. Self-distillation amplifies regularization in hilbert space. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 3351–3361. Curran Associates, Inc., 2020.
- [19] Hossein Mobahi, Mehrdad Farajtabar, and Peter L. Bartlett. Self-distillation amplifies regularization in hilbert space. In *arXiv preprint arXiv:2002.05715*, 2020.
- [20] Novi Quadrianto, Viktoriia Sharmanska, and Oliver Thomas. Discovering fair representations in the data domain. In *CVPR*, 2019.
- [21] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [22] Hee Jung Ryu, Hartwig Adam, and Margaret Mitchell. Inclusivefacenet: Improving face attribute detection with race and gender diversity. In *Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT/ML)*, 2018.
- [23] Hee Jung Ryu, Margaret Mitchell, and Hartwig Adam. Improving Smiling Detection with Race and Gender Diversity. *arXiv preprint arXiv:1712.00193*, 2017.
- [24] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *CVPR*, 2015.
- [25] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *CVPR*, 2019.
- [26] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *CVPR*, 2019.
- [27] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8919–8928, 2020.
- [28] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning Fair Representations. In *ICML*, 2013.
- [29] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 2018.
- [30] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European Conference on Computer Vision*, pages 649–666. Springer, 2016.
- [31] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men Also Like Shopping: Reducing Gender Bias Amplification using Corpus-level Constraints. In *EMNLP*, 2017.