

Unsupervised Learning of Debaised Representations with Pseudo-Attributes

Seonguk Seo¹ Joon-Young Lee³ Bohyung Han^{1,2}

¹ECE & ¹ASRI & ^{1,2}IPAI, Seoul National University ³Adobe Research

{seonguk, bhhan}@snu.ac.kr jolee@adobe.com

Abstract

Dataset bias is a critical challenge in machine learning since it often leads to a negative impact on a model due to the unintended decision rules captured by spurious correlations. Although existing works often handle this issue based on human supervision, the availability of the proper annotations is impractical and even unrealistic. To better tackle the limitation, we propose a simple but effective unsupervised debiasing technique. Specifically, we first identify pseudo-attributes based on the results from clustering performed in the feature embedding space even without an explicit bias attribute supervision. Then, we employ a novel cluster-wise reweighting scheme to learn debaised representation; the proposed method prevents minority groups from being discounted for minimizing the overall loss, which is desirable for worst-case generalization. The extensive experiments demonstrate the outstanding performance of our approach on multiple standard benchmarks, even achieving the competitive accuracy to the supervised counterpart. The source code is available at our project page¹.

1. Introduction

Deep neural networks have achieved impressive performance by minimizing the average loss on training datasets. Although we typically adopt the empirical risk minimization framework as a training objective, it is sometimes problematic due to dataset bias leading to significant degradation of worst-case generalization performance as discussed in [2, 12, 18, 37, 38]. This is because models do not always learn what we expect, but, to the contrary, rather capture unintended decision rules from spurious correlations. For example, on the Colored MNIST dataset [1, 21, 25], where each digit is highly correlated to a certain color, a network often learns the color patterns in images, not the digit information, while ignoring few conflicting samples. Such an unintended rule works well on most of the training samples but incurs unexpected worst-case errors for the examples in

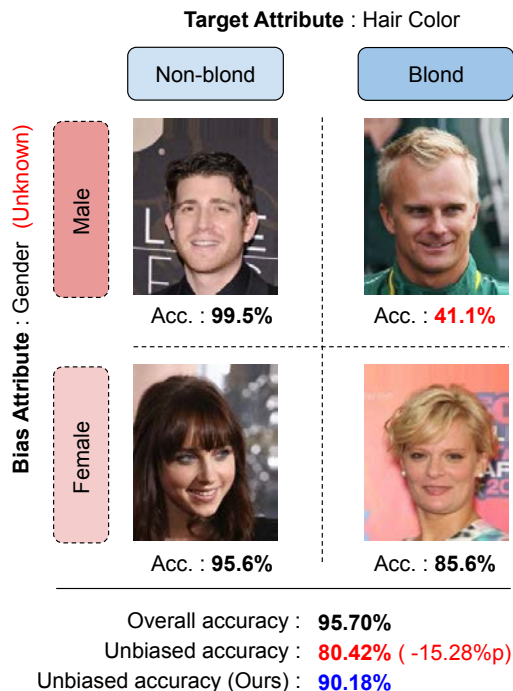


Figure 1. Representative examples on the CelebA dataset for the problem that we focus on. Since most of the people with blond hair are women, *hair color* attribute has a spurious correlation with *gender* attribute. Thus, when trained to classify the *hair color*, a network captures the unintended decision rule using *gender*, leading to poor worst-group and unbiased accuracies, despite its high overall accuracy. Our model aims to learn debaised representation, which gives better worst-group and unbiased accuracies, especially when the bias information is unavailable.

minority groups, which makes the model unable to generalize on test environments with distribution shifts or robustness constraints. Figure 1 illustrates the problem that we mainly deal with in this paper.

To mitigate the bias issue, learning debaised representations from a biased dataset has received growing attention from machine learning community [1, 3, 14, 21, 26, 28, 34]. However, most previous works rely on the explicit supervision or prior knowledge under the assumption of the pres-

¹<https://github.com/skynbe/pseudo-attributes>

ence of dataset bias. This setting is problematic because it is challenging to identify what kinds of bias exist and which attributes involve spurious correlations without thorough analysis of model and dataset. Note that, even with the information about dataset bias, the relevant annotations over all training examples is challenging. Contrary to the supervised approaches, [25,30] tackle a more challenging setting, where the bias information is unavailable, via failure-based learning or subgroup-based penalizing.

This paper presents a simple but effective unsupervised debiasing technique via feature clustering and cluster reweighting. We first observe that the examples with the same label for a certain attribute other than the target attribute tend to have similar representations in the feature space by the model trained sufficiently. Based on this observation, we estimate bias *pseudo-attributes* in an unsupervised manner from the clustering results within each class. To exploit the bias pseudo-attributes for learning debiased representations, we introduce a reweighting scheme for the corresponding clusters, where each cluster has an importance weight depending on its size and task-specific accuracy. This strategy encourages the minority clusters to participate in the optimization process actively, which is critical to improving worst-group generalization. Despite its simplicity, our method turns out to be effective for debiasing without any supervision of bias information; it is even comparable to the supervised debiasing method. The main contributions of our work are summarized below:

- We propose a simple but effective unsupervised debiasing approach, which requires no explicit supervision about spurious correlations across attributes.
- We introduce a technique to learn debiased representations by identifying bias pseudo-attributes via clustering and reweighting the corresponding clusters based on both their size and target loss.
- We provide extensive experimental results and achieve outstanding performance in terms of unbiased and worst-group accuracies, which are even as competitive as supervised debiasing methods.

The rest of this paper is organized as follows. We review the prior research in Section 2. Section 3 presents the proposed framework for learning debiased representations, and Section 4 demonstrates its effectiveness with extensive empirical analysis. We conclude our paper in Section 5.

2. Related Work

2.1. Bias in computer vision tasks

Real-world datasets are inevitably biased due to their insufficiently controlled collection process, and consequently,

deep neural networks often capture unintended correlations between the true labels and the spuriously correlated ones. Measuring and mitigating the potential risks posed by dataset or algorithmic bias has been extensively investigated in various computer vision tasks [3,6,16,21,31,33,36]. For example, VQA models frequently exploits statistical regularities between answer occurrences and patterns of questions while ignoring visual information [3,7]. Semantic segmentation models typically take advantage of scene context for pixel-level predictions of semantic labels [6]. To prevent using undesirable correlation in biased datasets, existing approaches often rely on human supervision for bias annotations and present several technical directions such as data augmentation [13,39], model ensembles [3,7], and statistical regularizations [1]. Such supervised debiasing techniques have been applied to various computer vision tasks by exploiting known application-specific bias information, including uni-modality of the dataset in visual question answering [3], stereotype textures in image recognition [13], temporal invariance in action recognition [20], and demographic information in facial image recognition [26,35,39].

2.2. Handling distribution shifts

Distribution shift has recently emerged as a critical challenge in machine learning, where the goal of the optimization is to learn a robust model in a test environment with a different distribution. Distributionally robust optimization (DRO) [2,9,11,17] has been proposed to improve the worst-case generalization performance over a family of target distributions, and has provided theoretical background of Group DRO [26] and its variation [30]. However, the objective of DRO often leads to an overly conservative model and results in performance degeneration on unseen environments [10,15]. To relax the constraints for the uncertainty set of test distributions, some approaches pose additional assumptions. For instance, the dataset consists of multiple groups with the shared properties and the uncertainty set is represented by a mixture of these groups. This assumption is also used in robust federated learning [19,24], algorithmic fairness [5,8,37], and domain generalization [4,18,29]. Our framework also takes advantage of this assumption but does not rely on the supervision of group information.

2.3. Debiasing via loss-based reweighting

There exist several generic debiasing techniques via sample reweighting based on observed task-specific losses under the supervised environment [26] or the unsupervised setting [25,30]. Group DRO [26] exploits the group information specified by the bias attributes and aims to improve the worst-group generalization performance. On the other hand, Nam *et al.* [25] employ the difference between the generalized and standard cross-entropy loss to capture the bias-alignment for sample reweighting while Sohoni *et*

al. [30] estimate subclass labels via clustering and utilize the information for distributionally robust optimization to mitigate hidden stratification. Although the unsupervised approaches work well in small and artificial datasets such as MNIST, their performance improvement becomes marginal in real-world datasets including CelebA. Our framework also belongs to unsupervised methods that do not rely on the bias information to learn debiased representations.

3. Method

This section presents our debiasing technique via bias pseudo-attribute estimation and sample reweighting.

3.1. Preliminaries

Let an example \mathbf{x} be associated with a set of m attributes $\mathcal{A} := \{a_1, \dots, a_m\}$. The goal of our model is to predict a target attribute $a_t \in \mathcal{A}$ by estimating the intended causation $p(a_t|\mathbf{x})$, which does not involve any undesirable correlation to other latent attributes, *i.e.*, $p(a_t|\mathbf{x}) = p(a_t|\mathbf{x}, a_i), \forall a_i \in \mathcal{A} - \{a_t\}$. On the other hand, spurious correlation indicates strong coincidence between two attributes $a_i, a_j \in \mathcal{A}$; the conditional entropy $H(a_i|a_j)$ is close to zero and there exists no causal relationship between them. A machine learning algorithm is considered biased if a certain attribute $a_b \in \mathcal{A}$ has a spurious correlation with the target attribute a_t and affects the prediction, *i.e.*, $p(a_t|\mathbf{x}) \neq p(a_t|\mathbf{x}, a_b)$. Our approach performs debiasing by estimating groups in the dataset without supervision, where the group is defined by a pair of target and bias attributes, *e.g.*, $g = (a_t, a_b)$.

3.2. Observation

If a bias attribute is highly correlated to a target attribute while being easy to learn, the model may ignore few conflicting examples and learn its decision rule based on the bias attributes with spurious correlations to maximize accuracy [25, 27]. To prevent this undesirable situation, a simple group upweighting or resampling strategies [27] are known to be effective while they work poorly in realistic scenarios, where the bias information is unknown during training.

To overcome this challenge, from our intuition, we analyze the feature semantics over the target and bias attributes. We first naively train a base model on the CelebA dataset to classify *hair color*, and visualize the representation of the examples after convergence with a sufficient number of epochs ($T = 100$). We select *gender* as a bias attribute, but do not utilize any information of the bias attribute during training. It turns out that, even without using the bias information during training, the examples drawn from certain groups, which are given by a combination of hair color and gender attribute values in this case, *e.g.*, (male, non-blonde) and (female, blonde), are located closely in the feature space. This observation implies that it is possible to identify bias pseudo-attributes by taking advantage of the

embedding results even without attribute supervisions. Our unsupervised debiasing framework is based on the capability to identify the bias pseudo-attributes via clustering.

3.3. Formulation

Suppose that training examples, (\mathbf{x}, y) , are drawn from a certain distribution \tilde{P} . Given a loss function $\ell(\cdot)$ and model parameters θ , the objective of the empirical risk minimization (ERM) is to optimize the following expected loss:

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, y) \sim P} [\ell((\mathbf{x}, y); \theta)], \quad (1)$$

where P is the empirical distribution over training data that approximates the true distribution \tilde{P} . Although ERM generally works well, it tends to ignore examples in minority groups that conflict with bias attributes and implicitly assumes the consistency of the underlying distributions for training and testing data. Consequently, the approach often leads to high unbiased and worst-group test error [9, 26].

Several distributionally robust optimization (DRO) techniques [2, 9, 17] can be employed to tackle the dataset bias and distribution shift problems and maximize unbiased generalization accuracy. They consider a particular uncertainty set \mathcal{Q}_P , which is close to the training distribution P , *e.g.*, $\mathcal{Q}_P = \{Q : D_f[Q||P] \leq \delta\}$, where $D_f[\cdot||\cdot]$ indicates an f -divergence function². To minimize the worst-case loss over the uncertainty distribution set \mathcal{Q}_P , DRO optimizes

$$\min_{\theta} \left\{ \mathbb{R}_{\mathcal{Q}_P}(\theta) := \sup_{Q \in \mathcal{Q}_P} \mathbb{E}_{(\mathbf{x}, y) \sim Q} [\ell((\mathbf{x}, y); \theta)] \right\}. \quad (2)$$

However, this objective is overly pessimistic and makes the model consider the implausible worst cases [10, 15].

The group distributionally robust optimization, referred to as group DRO [26] creates more realistic sets of possible test distributions by leveraging the prior knowledge of group information. They assume the training distribution P is a mixture of G groups, P_g , which is given by

$$P_G = \sum_{g \in \mathcal{G}} c_g P_g, \quad c \in \Delta_G \quad (3)$$

where $\mathcal{G} = \{1, \dots, G\}$ and Δ_G is a $(G - 1)$ -dimensional simplex. Then the uncertainty set \mathcal{Q}_P is defined by a set of all possible mixtures of these groups, *i.e.*, $\mathcal{Q}_{P_G} = \{\sum_{g \in \mathcal{G}} c_g P_g : c \in \Delta_G\}$. Because \mathcal{Q}_{P_G} is a simplex, its optimum is achieved at a vertex, thus minimizing the worst-case risk of \mathcal{Q}_{P_G} is equivalent to

$$\min_{\theta} \left\{ \mathbb{R}_G(\theta) := \max_{g \in \mathcal{G}} \mathbb{E}_{(\mathbf{x}, y) \sim P_g} [\ell((\mathbf{x}, y); \theta)] \right\}. \quad (4)$$

²Let P and Q be probability distributions over a space Ω , then f -divergence is $D_f(P||Q) = \int_{\Omega} f\left(\frac{dP}{dQ}\right) dQ$.

Algorithm 1: Debiasing with bias pseudo-attribute

- 1 **Require:** step size η_θ , momentum m , training steps T , batch size B , the number of clusters K
- 2 **Base model:**
- 3 **Initialize** $\tilde{\theta}$
- 4 **for** $t = 1, \dots, T$ **do**
- 5 Sample $(\mathbf{x}_i, y_i) \sim P$ for $i = 1, \dots, B$;
- 6 $\tilde{\theta} \leftarrow \tilde{\theta} - \eta_\theta \sum_{i=1}^B \nabla \ell((\mathbf{x}_i, y_i); \tilde{\theta})$;
- 7 **end**
- 8 **for** $k = 1, \dots, K$ **do**
- 9 $P_k = \{(\mathbf{x}_n, y_n) \mid h((\mathbf{x}_n, y_n); \tilde{\theta}) = k \text{ for all } n\}$;
- 10 $N_k = |P_k|$;
- 11 **end**
- 12 **Target model:**
- 13 **Initialize** θ and ω_k for $k = 1, \dots, K$
- 14 **for** $t = 1, \dots, T$ **do**
- 15 $\omega_k \leftarrow (1 - m)\omega_k + \frac{m}{N_k} \mathbb{E}_{(\mathbf{x}, y) \sim P_k} [\ell((\mathbf{x}, y); \theta)]$
- 16 for $k = 1, \dots, K$;
- 17 Sample $(\mathbf{x}_i, y_i) \sim P$ for $i = 1, \dots, B$;
- 18 $\alpha_i = \omega_{h((\mathbf{x}_i, y_i); \tilde{\theta})}$;
- 19 $\bar{\alpha}_i = \alpha_i / \sum_{i=1}^B \alpha_i$;
- 20 $\theta \leftarrow \theta - \eta_\theta \sum_{i=1}^B \bar{\alpha}_i \nabla \ell((\mathbf{x}_i, y_i); \theta)$;
- 21 **end**

Different from the group DRO setting, we do not know the group assignment for each training example. Instead, we use the bias pseudo-attribute information, obtained by any clustering algorithm in the feature embedding space, to define groups. Note that the clustering is performed with the representations given by the base model trained without debiasing, which is parameterized by $\tilde{\theta}$. Our goal is to alleviate dataset bias and maximize unbiased accuracy, and we need to consider all groups fairly for optimization. To this end, we assign a proper importance weight, ω_k , to the k^{th} cluster, where $k \in \mathcal{K} = \{1, \dots, K\}$, and the final objective of our framework is given by minimizing a weighted empirical risk as follows:

$$\min_{\theta} \left\{ \mathbb{R}_{\mathcal{K}}(\theta) := \mathbb{E}_{(\mathbf{x}, y) \sim P} \left[\omega_{h((\mathbf{x}, y); \tilde{\theta})} \ell((\mathbf{x}, y); \theta) \right] \right\}, \quad (5)$$

where $h((\mathbf{x}, y); \tilde{\theta})$ denotes the cluster membership of an example (\mathbf{x}, y) . The details of the weight assignment method will be discussed next.

3.4. Sample weighting with bias pseudo-attributes

Based on our observation described in Section 3.2, we first cluster training examples in each class on the feature

embedding space defined by the base model optimized sufficiently, e.g., for 100 epochs using the standard classification loss. We suppose that each cluster corresponds to a bias pseudo-attribute. Among all clusters, we focus on the examples in the minority clusters, especially when they have high average losses. A common failure case in the presence of dataset bias is incurred by ignoring specific subpopulation groups for minimizing the overall training loss, and minority clusters are prone to be ignored due to their sizes. The problematic cases among the clusters are the ones that contain many bias-conflicting examples, having high losses, and thus result in poor worst-case errors. If the minority clusters consist of mostly bias-aligned samples, they will apparently achieve high classification accuracy.

Therefore, to handle dataset bias issue, we should consider both scale and average difficulty (loss) of each cluster, unlike group DRO [26] and George [30], which focus only on the average loss. We calculate the importance weight of each cluster by our reweighting scheme to train the target model, which is given by

$$\begin{aligned} \omega_k &= \frac{\mathbb{E}_{(\mathbf{x}, y) \sim P_k} [\ell((\mathbf{x}, y); \theta)]}{N_k} \\ &= \frac{\mathbb{E}_{(\mathbf{x}, y) \sim P} [\ell((\mathbf{x}, y); \theta) \mid h((\mathbf{x}, y); \tilde{\theta}) = k]}{\sum_i \mathbb{1}(h((\mathbf{x}_i, y_i); \tilde{\theta}) = k)}, \quad (6) \end{aligned}$$

where θ and $\tilde{\theta}$ indicates the parameters of the final and base models, respectively, $h(\cdot, \cdot)$ is a cluster membership function, and $\mathbb{1}(\cdot)$ is an indicator function. Note that P_k denotes the sample distribution of the k^{th} cluster and N_k is the number of samples in the k^{th} cluster, where $k \in \mathcal{K} = \{1, \dots, K\}$.

3.5. Algorithm procedure

Algorithm 1 presents the optimization procedure of the proposed framework. We first naïvely train a baseline network (line 4-7), parameterized by $\tilde{\theta}$. Then, we cluster all training examples based on the features extracted from the network to obtain the membership distribution P_k and the size of cluster N_k (line 8-11). Based on the cluster assignments, we calculate the importance weight of each cluster ω_k using the target model, parameterized by θ , where the weight is updated by exponential moving average at each iteration (line 15). We finally use the normalized importance weight of each sample $\bar{\alpha}_i$ over a mini-batch to train the target model (line 18-19).

4. Experiments

4.1. Dataset

CelebA [22] is a large-scale face dataset for face image recognition, containing 40 attributes for each image. Following the previous works [25, 26], we set *hair color* and

Table 1. Unbiased and worst-group results in the existence of spurious correlation between target and bias attributes on the test split of the CelebA dataset. LfF* denotes a variant of LfF [25], which fine-tuned only the classification layer of a trained baseline model, for additional comparison to ours. Bold and underline fonts indicate the first and second place among the compared approaches, respectively. All experimental results are the average of three runs.

Target	Bias	Unbiased accuracy (%)					Worst-group accuracy (%)				
		Unsupervised				Supervised	Unsupervised				Supervised
		Base	LfF	LfF*	BPA (ours)	Group DRO	Base	LfF	LfF*	BPA (ours)	Group DRO
Blond Hair	Gender	80.42	59.46	84.89	<u>90.18</u>	91.39	41.02	34.23	57.96	<u>82.54</u>	87.86
Heavy Makeup	Gender	71.19	56.34	71.85	<u>73.78</u>	<u>72.70</u>	17.35	30.81	23.87	39.84	21.36
Pale Skin	Gender	71.50	78.69	75.23	<u>90.06</u>	90.55	36.64	57.38	43.26	88.60	<u>87.68</u>
Wearing Lipstick	Gender	73.90	53.79	73.84	78.28	<u>78.26</u>	31.38	25.52	31.92	46.52	<u>46.08</u>
Young	Gender	78.19	45.99	79.58	<u>82.27</u>	82.40	52.79	0.34	57.79	<u>74.33</u>	76.29
Double Chin	Gender	64.61	65.46	68.47	<u>82.92</u>	83.19	21.33	28.19	28.24	<u>67.78</u>	72.94
Chubby	Gender	67.42	60.03	71.56	83.88	<u>81.90</u>	24.30	7.60	34.09	<u>72.32</u>	72.64
Wearing Hat	Gender	93.53	84.56	94.81	<u>96.80</u>	96.84	85.12	69.06	88.31	94.94	<u>94.67</u>
Oval Face	Gender	62.70	57.64	62.30	<u>67.18</u>	<u>65.40</u>	29.15	7.40	36.00	<u>55.78</u>	56.84
Pointy Nose	Gender	62.10	42.20	63.83	<u>68.90</u>	70.71	25.80	1.05	38.04	<u>52.48</u>	63.76
Straight Hair	Gender	70.28	39.57	72.84	79.18	<u>77.04</u>	47.82	1.95	58.53	72.09	<u>66.10</u>
Blurry	Gender	73.05	76.70	77.52	88.93	<u>87.05</u>	45.68	43.81	52.35	84.10	<u>82.06</u>
Narrow Eyes	Gender	63.18	68.53	67.77	<u>76.39</u>	76.72	27.01	31.81	38.53	73.24	<u>71.47</u>
Arched Eyebrows	Gender	69.72	56.17	71.87	<u>74.77</u>	78.30	34.76	26.21	44.97	<u>54.36</u>	69.44
Bags Under Eyes	Gender	69.47	44.61	71.86	77.84	<u>75.88</u>	41.65	0.06	49.10	<u>62.55</u>	63.34
Bangs	Gender	89.04	41.41	89.04	<u>93.94</u>	94.45	76.91	3.18	82.37	92.21	<u>92.12</u>
Big Lips	Gender	60.87	46.74	62.15	66.50	<u>63.70</u>	30.85	31.44	38.54	56.99	<u>47.55</u>
No Beard	Gender	73.11	60.12	73.13	79.58	<u>77.86</u>	13.30	11.92	20.00	40.00	<u>36.70</u>
Receding Hairline	Gender	69.72	70.57	74.58	<u>84.95</u>	85.15	35.69	32.10	45.53	<u>79.11</u>	79.12
Wavy Hair	Gender	73.10	48.00	74.53	79.89	<u>79.65</u>	38.01	0.06	45.24	<u>65.74</u>	66.79
Wearing Earrings	Gender	72.17	59.35	74.17	84.57	<u>83.50</u>	26.26	0.10	32.95	<u>72.81</u>	75.24
Wearing Necklace	Gender	55.04	58.64	57.21	68.96	<u>62.89</u>	2.72	0.22	6.67	41.93	<u>24.34</u>
Average	Gender	72.67	58.65	74.87	81.74	<u>80.87</u>	39.91	21.91	47.88	69.84	<u>69.68</u>

Table 2. Unbiased and worst-group results on the Waterbirds dataset.

Target	Bias	Unbiased accuracy (%)					Worst-group accuracy (%)				
		Unsupervised				Supervised	Unsupervised				Supervised
		Base	LfF	LfF*	BPA (ours)	Group DRO	Base	LfF	LfF*	BPA (ours)	Group DRO
Object	Place	84.63	85.48	84.57	<u>87.05</u>	88.99	62.39	68.02	61.68	<u>71.39</u>	80.82
Place	Object	87.99	85.77	85.05	<u>88.44</u>	89.20	73.34	62.37	60.00	<u>79.16</u>	85.27

Table 3. Unbiased accuracy (%) on the valid split of the Colored-MNIST dataset.

Target	Bias	Unsupervised			Supervised
		Baseline	LfF	BPA (ours)	Group DRO
Digit	Color	74.48	85.15	<u>85.26</u>	85.88
Color	Digit	99.95	<u>99.91</u>	99.82	98.96

heavy makeup as the target attribute. Note that gender attribute is spuriously correlated to the two attributes and employed as the bias attribute for worst-group accuracy evaluation in our experiment. For more comprehensive results, we also consider the other 32 attributes as the target attributes.

Waterbirds [26] is a synthesized dataset with 4,795 training examples, created by combining birds photographs from the Caltech-UCSD Birds-200-2011 (CUB) dataset [32] and

the background images in the Places dataset [40]. There exist two attributes in the dataset; one is the type of a bird, {waterbird, landbird}, which is the target attribute, and the other is the background place, {water, land}.

The Colored-MNIST dataset [1, 21, 25] is an extension of MNIST with the color attributes, where each digit is highly correlated to a certain color. There are 60K training examples and 10K test images, where the ratio of bias-aligned samples³ is 95%. We follow the protocol employed in [25] for the experiment.

4.2. Implementation details

For CelebA and Waterbirds, we use a ResNet-18 as our backbone network, which is pretrained on ImageNet. We

³It denotes the samples that can be correctly classified by using the bias attribute (color).

train both the base and target models using the Adam optimizer with a learning rate of 1×10^{-4} , a batch size of 256, and a weight decay rate of 0.01. For the Colored MNIST dataset, we adopt a multi-layer perceptron with three hidden layers, each of which has 100 hidden units. We also employ the same Adam optimizer with a learning rate of 1×10^{-3} . We train the models for 100 epochs for all experiments, and decay the learning rate using the cosine annealing [23].

For clustering, we extract features from a separately trained base network with the standard classification loss and perform k -means clustering with $K = 8$ in all experiments. The cluster weight of the k^{th} cluster, ω_k , is updated by exponential moving average at each iteration with a momentum m of 0.3. All the results reported in our paper are obtained from the average of three runs.

4.3. Evaluation protocol

To evaluate the unbiased accuracy with an imbalanced evaluation set, we measure the accuracy of each group $g = (a_t, a_b)$, defined by a pair of target and bias attribute values. We finally report the average accuracy of each group and worst-group accuracy among all groups as in [25].

4.4. Results

We present our main results on the standard benchmarks including CelebA, WaterBirds, and Colored-MNIST. In the rest of this section, our method is denoted by BPA, which stands for bias pseudo-attribute.

CelebA Before evaluating our frameworks, we first thoroughly analyze the CelebA dataset in terms of algorithmic bias among the attributes. There are a total of 40 attributes in the CelebA dataset. The bias attribute is fixed to *gender*, and we analyze its relation to the attributes of the target candidates given by the rest of 39 attributes. We accept the target attributes if the smallest group given by its combinations with the bias attribute in the test split contains at least 10 examples for statistical stability⁴. We suppose that the spurious correlation exists between target and bias attributes when a baseline model gives a large performance gap between its overall accuracy and unbiased accuracy (e.g., $>5\%$ points). We found that 26 out of 32 attributes have spurious correlation to *gender*, and report the results for the attributes. See our supplementary files for more detailed analysis.

Table 1 presents the experimental results of the proposed algorithm on the CelebA dataset, in comparison to the existing methods as well as the baseline model. Our model significantly outperforms the baseline and LfF [25] for all target attributes in terms of both unbiased and worst-group accuracies. Note that our model is almost competitive to

⁴The removed target attributes are *5 o'clock shadow*, *bald*, *rosy cheeks*, *sideburns*, *goatee*, *mustache*, and *wearing necktie*.

a supervised approach, Group DRO [26], without the explicit bias information. On the other hand, we observe that training the model with LfF deteriorates performance even compared to the baseline. This is because it fixes the feature extractor and only trains its classification layer at the end⁵. To conduct a meaningful comparison with the stable version of LfF, we first train the baseline model used in our experiment for 100 epochs and then fine-tune the classification layer only using the LfF algorithm; this revised version is referred to as LfF* in the rest of this section. Although the performance of LfF* is stable, the improvement by debiasing is still limited compared to Group DRO and our approach. Additional experimental results for other bias attributes are provided in our supplementary documents.

Waterbirds We also evaluate our model on the Waterbirds dataset and present the results in Table 2. As in the CelebA dataset, our model achieves the best accuracies among the unsupervised methods in terms of both the unbiased and worst-group accuracies, and presents comparable results to the supervised method [26]. Our successful results on Waterbirds imply that the proposed method is robust to small-scale datasets as well.

Colored-MNIST Table 3 shows that our model achieves consistent accuracy in the digit classification with color bias. Besides, the color classification performance, where the algorithmic bias does not exist, turns out to be also competitive to the baseline, while the supervised approach is not good at this setting.

4.5. Analysis

Results with no algorithmic bias We test our algorithm on unbiased datasets to make sure that it is dependable on the cases without algorithmic bias. The unbiased setting is defined by the configuration that a baseline model involves a marginal difference between its overall accuracy and unbiased accuracy (e.g., $< 5\%$ points). Similar to Table 1, we identify a subset of target attributes in CelebA, which is not spuriously correlated to *gender*; there exist 6 out of 32 attributes. Table 4 illustrates the results of the 6 target attributes, where the accuracy of our approach is most consistent among the 4 methods. This implies that our framework can be incorporated into the existing recognition models directly, without knowing the presence of dataset bias. Note that color classification with digit bias on Colored-MNIST or background place classification with object bias on Waterbirds are also qualified as unbiased settings, where our model gives consistent results.

Multiple bias attributes Thanks to the unsupervised nature of our method, we can simply evaluate our model on

⁵<https://github.com/alinelab/LfF>

Table 4. Unbiased and worst-group accuracies on the CelebA dataset with the target attributes, where the algorithmic bias does not exist.

Target	Bias	Unbiased accuracy (%)					Worst-group accuracy (%)				
		Unsupervised				Supervised	Unsupervised				Supervised
		Base	LfF	LfF*	BPA (ours)	Group DRO	Base	LfF	LfF*	BPA (ours)	Group DRO
Attractive	Gender	76.05	30.18	75.97	<u>77.90</u>	78.35	63.61	6.09	64.78	<u>65.20</u>	66.30
Smiling	Gender	91.66	74.62	91.20	92.08	<u>91.64</u>	88.49	60.09	<u>88.65</u>	90.06	88.48
Mouth Open	Gender	93.10	81.85	92.96	<u>93.45</u>	93.64	91.52	66.92	<u>92.44</u>	92.27	91.69
High Cheekbones	Gender	83.44	48.40	83.70	84.93	<u>84.52</u>	70.49	7.92	73.56	78.56	<u>78.37</u>
Eyeglasses	Gender	98.20	85.47	98.38	<u>98.39</u>	98.65	96.24	76.89	96.85	<u>97.22</u>	97.64
Black Hair	Gender	84.92	61.00	85.19	<u>86.57</u>	86.76	75.47	22.04	75.69	81.28	<u>80.67</u>
Average	Gender	87.90	63.59	87.92	<u>88.89</u>	88.93	80.97	39.51	82.29	84.10	<u>83.86</u>

Table 5. Unbiased accuracy (%) with multiple bias attributes. For each target, our model only requires a single model, while Group DRO [26] should train separate models depending on the bias set.

Target	Blond Hair				Blurry			
	Baseline	LfF*	BPA (ours)	Supervised Group DRO	Baseline	LfF*	BPA (ours)	Supervised Group DRO
Gender	80.42	84.89	<u>90.18</u>	91.39	73.05	76.70	88.93	<u>87.05</u>
Gender, Heavy Makeup	83.64	<u>88.82</u>	91.90	81.09	75.37	<u>79.55</u>	89.09	72.17
Gender, Wearing Lipstick	80.34	84.13	91.63	<u>85.93</u>	79.88	<u>83.21</u>	89.79	79.88
Gender, Young	78.39	81.21	89.05	<u>87.96</u>	72.97	77.77	88.66	<u>85.39</u>
Gender, No Beard	79.50	82.51	89.92	<u>85.01</u>	78.91	79.84	84.06	<u>81.07</u>
Gender, Wearing Necklace	79.25	81.03	92.62	<u>92.26</u>	71.80	78.07	89.60	<u>85.33</u>
Gender, Big Nose	81.18	84.10	<u>90.58</u>	90.83	71.89	77.11	88.57	<u>87.11</u>
Gender, Smiling	79.75	82.91	<u>89.85</u>	91.73	73.31	78.04	89.32	<u>87.87</u>
Average	80.29 ±1.71	83.53 ±2.64	90.79 ±1.29	<u>87.83 ±4.10</u>	74.88 ±3.32	79.08 ±2.06	88.44 ±1.98	<u>82.69 ±5.50</u>

Table 6. Ablation results on our importance weighting scheme on CelebA with *blond hair* and *gender* for target and bias attributes, respectively, in terms of unbiased and worst-group accuracies (%).

Scale	Loss	Unbiased Acc.	Worst-group Acc.
		80.42	41.02
	✓	83.86	57.44
✓		89.08	76.55
✓	✓	90.18	82.54

multi-bias scenarios, where multiple bias attributes exist in the dataset, without modification. Table 5 presents the unbiased results with multiple bias attributes using our method and Group DRO [26], where we additionally report the average and standard deviation over all bias sets to compare the overall effectiveness and robustness. We also present the results of LfF*, a variant of LfF [25], introduced in Table 1. When trained on multiple bias attributes, the accuracy of Group DRO is sensitive to bias sets while our method achieves stable and superior results for a variety of sets. Note also that our model is applicable to any bias sets without additional fine-tuning while a supervised method should separately train their model for each set.

Ablative results on importance weighting We perform the ablative experiments on the CelebA dataset to analyze

Table 7. Unbiased accuracy (%) with bias-unspecified settings. The results are the average of a set of unbiased accuracies, each of which adopts one of the 25 unspecified attributes jointly with the specified bias attribute, *gender*, as the bias attributes to define groups.

Target	Baseline	BPA (ours)	Group DRO
Blond Hair	79.13 ± 2.72	90.16 ± 3.19	90.82 ± 2.76
Heavy Makeup	70.26 ± 3.84	73.52 ± 3.86	71.57 ± 4.33
Young	77.56 ± 1.80	81.31 ± 2.31	80.56 ± 2.09
Double Chin	62.56 ± 2.22	81.71 ± 3.61	78.46 ± 3.37
Chubby	67.80 ± 2.77	82.36 ± 3.55	79.91 ± 3.43
Wearing Hat	90.80 ± 4.01	95.11 ± 3.10	94.71 ± 3.76
Oval Face	61.77 ± 1.80	66.63 ± 1.63	65.35 ± 1.45
Pointy Nose	63.96 ± 1.42	70.53 ± 1.52	70.16 ± 1.45

the effectiveness of our cluster weighting strategy. Table 6 presents the results when only one of the scale N_k and the average loss $\mathbb{E}_{(\mathbf{x}, y) \sim P_k} [\ell((\mathbf{x}, y); \theta)]$, respectively, are taken into account to compute ω_k for the k^{th} cluster in Eq. (6). Note that our ablative model with the loss factor only is closely related to [30]. Table 6 also shows that combining both terms plays a crucial role for learning debiased representations while the scale factor is clearly more important.

Unspecified group shifts To verify the robustness in another realistic scenario, we test unspecified group shifts,




Figure 2. Sensitivity analysis on the number of clusters in our framework on the CelebA dataset.

where the group information at test time is not fully provided during training. The bias attribute specified during training, which is exploited by group DRO, is fixed to *gender*. To evaluate the performance in this setting, we measure a set of unbiased accuracies corresponding to the combinations of the specified bias attribute, *gender*, and each of 25 unspecified bias attributes except the target attribute⁶. Note that the unspecified bias attributes are not introduced during training but used to define groups at test time. Table 7 clearly shows that our model outperforms Group DRO in the bias-unspecified setting, where we only report the average and standard deviation of the set of unbiased accuracies due to space limitation. This implies that, although Group DRO handles group shifts well within the simplex of the specified group distributions, it suffers from worst-case generalization for unspecified group shifts. Note that the proposed approach is free from the issue because it does not use any information about bias for training.

Sensitivity analysis on the number of clusters We conduct ablation study on the number of groups for clustering on the feature embedding space to obtain bias pseudo-attributes on the CelebA dataset. We set *gender* as the bias attribute and evaluate the unbiased accuracies for several target attributes. Figure 2 presents that the results are stable over a wide range of the number of clusters and the accuracy is saturated when $K \geq 4$.

Feature visualization Figure 3 illustrates the t-SNE results of instance embeddings for the baseline model (left) and ours (right) on the CelebA dataset for the *blond hair* attribute classification, where we visualize only negative (*blond hair = false*) examples for an effective visualization. Blue and orange colors indicate values—female and male, respectively—of the bias attribute, *gender*. We observe that our model helps to mix the examples in different groups within the same class, which is desirable for debiasing.




Figure 3. The t-SNE plots of feature embeddings using baseline (left) and ours (right) trained to classify *hair color*. We visualize the distribution of samples that have the same target value (*blond hair = false*). Blue and orange colors denote different gender values, female and male, respectively. Our framework helps mix samples which have the same target but different bias attribute values.

5. Conclusion

We presented a generic unsupervised debiasing framework using pseudo-attribute. We observed that the examples sampled from the same groups are closely located in the feature embedding space. Based on our empirical observation, we claim that it is possible to identify the pseudo-attributes by taking advantage of the embedding results even without the attribute supervision. Inspired by this fact, we introduced a novel cluster-based weighting strategy for learning debiased representations. We demonstrated the effectiveness of our method on multiple standard benchmarks, which is even as competitive as the supervised debiasing method, group DRO. We also conducted a thorough analysis of our framework in many realistic scenarios, where our model provides substantial gains consistently.

Potential societal impact and limitation Machine learning models typically focus on performance improvement unconditionally. Hence, it is often exposed to the risk caused by dataset and/or algorithmic bias, which need to be carefully addressed for enhancing reliability and robust of models. This research contributes to a positive societal impact from this point of view. Although the proposed algorithm turns out to be effective for bias identification, there may be blind spots due to unexplored types of bias. Therefore, we believe that the identification of hidden and unobservable biases without prior knowledge is a promising research direction.

Acknowledgments This work was partly supported by the IITP grants [2021-0-02068, Artificial Intelligence Innovation Hub; 2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)] and the National Research Foundation of Korea (NRF) grant [2022R1A2C3012210] funded by the Korean government (MSIT) and by Samsung Electronics Co., Ltd. [IO210917-08957-01].

⁶There exist 26 (out of 39) valid attributes as described in Section 4.4.

References

- [1] Hyojin Bahng, Sanghyuk Chun, Sangdoon Yun, Jaegul Choo, and Seong Joon Oh. Learning de-biased representations with biased representations. In *ICML*, 2020. 1, 2, 5
- [2] Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenare, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013. 1, 2, 3
- [3] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. Rubi: Reducing unimodal biases for visual question answering. In *NeurIPS*, 2019. 1, 2
- [4] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *CVPR*, 2019. 2
- [5] Alexandra Chouldechova and Aaron Roth. The frontiers of fairness in machine learning. *arXiv preprint arXiv:1810.08810*, 2018. 2
- [6] Sanghyeok Chu, Dongwan Kim, and Bohyung Han. Learning debiased and disentangled representations for semantic segmentation. In *NeurIPS*, 2021. 2
- [7] Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. In *EMNLP*, 2019. 2
- [8] Michele Donini, Luca Oneto, Shai Ben-David, John S Shawe-Taylor, and Massimiliano Pontil. Empirical risk minimization under fairness constraints. In *NIPS*, 2018. 2
- [9] John Duchi, Peter Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *arXiv preprint arXiv:1610.03425*, 2016. 2, 3
- [10] John Duchi, Tatsunori Hashimoto, and Hongseok Namkoong. Distributionally robust losses for latent covariate mixtures. *arXiv preprint arXiv:2007.13982*, 2020. 2, 3
- [11] Rui Gao, Xi Chen, and Anton J Kleywegt. Wasserstein distributional robustness and regularization in statistical learning. *arXiv preprint arXiv:1712.06050*, 2017. 2
- [12] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *arXiv preprint arXiv:2004.07780*, 2020. 1
- [13] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2019. 2
- [14] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. Women also snowboard: Overcoming bias in captioning models. In *ECCV*, 2018. 1
- [15] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. Does distributionally robust supervised learning give robust classifiers? In *ICML*, 2018. 2, 3
- [16] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017. 2
- [17] Johannes Kirschner, Ilija Bogunovic, Stefanie Jegelka, and Andreas Krause. Distributionally robust bayesian optimization. In *AISTATS*, 2020. 2, 3
- [18] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, 2017. 1, 2
- [19] Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*, 2019. 2
- [20] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *ECCV*, 2018. 2
- [21] Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *CVPR*, 2019. 1, 2, 5
- [22] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 4
- [23] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 6
- [24] Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh. Agnostic federated learning. *arXiv preprint arXiv:1902.00146*, 2019. 2
- [25] Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: Training debiased classifier from biased classifier. In *NeurIPS*, 2020. 1, 2, 3, 4, 5, 6, 7
- [26] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In *ICLR*, 2020. 1, 2, 3, 4, 5, 6, 7
- [27] Shiori Sagawa, Aditi Raghunathan, Pang Wei Koh, and Percy Liang. An investigation of why overparameterization exacerbates spurious correlations. In *ICML*, 2020. 3
- [28] Seonguk Seo, Joon-Young Lee, and Bohyung Han. Information-theoretic bias reduction via causal view of spurious correlation. In *AAAI*, 2022. 1
- [29] Seonguk Seo, Yumin Suh, Dongwan Kim, Geeho Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. In *ECCV*, 2020. 2
- [30] Nimit Sohoni, Jared Dunnmon, Geoffrey Angus, Albert Gu, and Christopher Ré. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. In *NeurIPS*, 2020. 2, 3, 4, 7
- [31] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR*, 2011. 2
- [32] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 5
- [33] Angelina Wang, Arvind Narayanan, and Olga Russakovsky. REVISE: A tool for measuring and mitigating bias in visual datasets. In *ECCV*, 2020. 2
- [34] Haohan Wang, Zexue He, Zachary L. Lipton, and Eric P. Xing. Learning robust representations by projecting superficial statistics out. In *ICLR*, 2019. 1
- [35] 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 *ICCV*, 2019. [2](#)
- [36] 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 *ICCV*, 2019. [2](#)
- [37] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *ICML*, 2013. [1](#), [2](#)
- [38] Marvin Zhang, Henrik Marklund, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: A meta-learning approach for tackling group shift. *arXiv preprint arXiv:2007.02931*, 2020. [1](#)
- [39] Yi Zhang and Jitao Sang. Towards accuracy-fairness paradox: Adversarial example-based data augmentation for visual debiasing. *arXiv preprint arXiv:2007.13632*, 2020. [2](#)
- [40] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *TPAMI*, 40(6):1452–1464, 2017. [5](#)