

# Combating Unknown Bias with Effective Bias-Conflicting Scoring and Gradient Alignment

Bowen Zhao<sup>1</sup>, Chen Chen<sup>2</sup>, Qian-Wei Wang<sup>1,3</sup>, Anfeng He<sup>2</sup>, Shu-Tao Xia<sup>1,3</sup>\*

<sup>1</sup>Tsinghua University

<sup>2</sup>Tencent TEG AI

<sup>3</sup>Peng Cheng Laboratory

zbw18@mails.tsinghua.edu.cn, chen1634chen@gmail.com, xiazt@sz.tsinghua.edu.cn

## Abstract

Models notoriously suffer from dataset biases which are detrimental to robustness and generalization. The identify-emphasize paradigm shows a promising effect in dealing with unknown biases. However, we find that it is still plagued by two challenges: A, the quality of the identified bias-conflicting samples is far from satisfactory; B, the emphasizing strategies just yield suboptimal performance. In this work, for challenge A, we propose an effective bias-conflicting scoring method to boost the identification accuracy with two practical strategies — peer-picking and epoch-ensemble. For challenge B, we point out that the gradient contribution statistics can be a reliable indicator to inspect whether the optimization is dominated by bias-aligned samples. Then, we propose gradient alignment, which employs gradient statistics to balance the contributions of the mined bias-aligned and bias-conflicting samples dynamically throughout the learning process, forcing models to leverage intrinsic features to make fair decisions. Experiments are conducted on multiple datasets in various settings, demonstrating that the proposed solution can alleviate the impact of unknown biases and achieve state-of-the-art performance.

## 1 Introduction

Deep Neural Networks (DNNs) have achieved significant progress in various visual tasks. DNNs tend to learn **intended** decision rules to accomplish target tasks commonly. However, they may follow **unintended** decision rules based on the easy-to-learn shortcuts to “achieve” target goals in some scenarios (Bahng et al. 2020). For instance, when training a model to classify digits on Colored MNIST (Kim et al. 2019), in which the images of each class are mainly dyed by one pre-defined color respectively (e.g., most ‘0’ are red, ‘1’ are yellow), the intended decision rules classify images based on the shape of digits, while the unintended decision rules utilize color information instead. Following Nam et al. (2020), sample  $x$  that can be “correctly” classified by unintended decision rules is denoted as a **bias-aligned** sample  $\underline{x}$  (e.g., red ‘0’) and vice versa a **bias-conflicting** sample  $\bar{x}$  (e.g., green ‘0’).

There are also many similar scenarios in real-world, e.g., an animal-centric image set may be biased by the habitats in

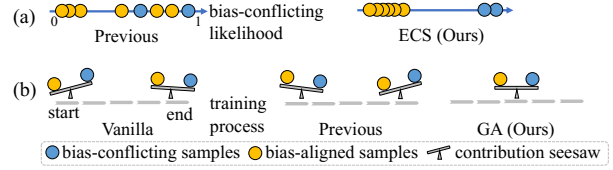


Figure 1: (a) Effective bias-Conflicting Scoring helps identify real bias-conflicting samples in stage I. (b) Gradient Alignment balances contributions from the mined bias-aligned and bias-conflicting samples throughout training, enforcing models to focus on intrinsic features in stage II.

backgrounds, and a human-centric set can be biased by gender or racial information. Models blinded by biased datasets usually perform poorly in mismatched distribution (e.g., a red ‘8’ may be incorrectly categorized as ‘0’). Worse still, models with racial or gender bias, *etc.* can cause severe negative social impacts. Moreover, in most real-world problems, the bias information (both bias type and precise labels of bias attribute) is unknown, making debiasing more challenging. Therefore, combating unknown biases is urgently demanded when deploying AI systems in realistic applications.

One major issue that leads to biased models is that the training objective (e.g., vanilla empirical risk minimization) can be accomplished via only unintended decision rules (Sagawa et al. 2020b). Accordingly, some studies (Nam et al. 2020; Kim et al. 2021) try to identify and emphasize the bias-conflicting samples. Nevertheless, we find that the low identification accuracy and the suboptimal emphasizing strategies impair the debiasing effect. In this work, we build an enhanced two-stage debiasing scheme to combat unknown dataset biases. In stage I, we present an Effective bias-Conflicting Scoring (ECS) function to mine bias-conflicting samples. On top of the off-the-shelf method, we propose a peer-picking mechanism to consciously pursue seriously biased auxiliary models and employ epoch-ensemble to obtain more accurate and stable scores. In stage II, we propose Gradient Alignment (GA), which balances the gradient contributions across the mined bias-aligned and bias-conflicting samples to prevent models from being biased. The gradient information is served as an indicator to down-weight (up-weight) the mined bias-aligned (bias-conflicting) samples in order to achieve dynamic balance throughout op-

\*Corresponding authors: Chen Chen and Shu-Tao Xia.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

timization. The effects of ECS and GA are illustrated in Fig. 1. Compared to other debiasing techniques, the proposed solution (i) does not rely on comprehensive bias annotations or a pre-defined bias type; (ii) does not require disentangled representations, which may fail in complex scenarios where disentangled features are hard to extract; (iii) does not introduce any data augmentations, avoiding additional training complexity such as in generative models; (iv) does not involve any modification of model architectures, making it easy to be applied to other networks. (v) significantly improves the debiasing performance.

The main contributions are: **(i)** To combat unknown dataset biases, we present an enhanced two-stage approach (illustrated in Fig. 2). We propose an effective bias-conflicting scoring algorithm equipped with peer-picking and epoch-ensemble in stage I (Sec. 3.1), and gradient alignment in stage II (Sec. 3.2). **(ii)** Broad experiments on commonly used datasets are conducted to compare several debiasing methods in a fair manner (overall, we train more than 650 models), among which the proposed method achieves state-of-the-art performance (Sec. 4.1 and Sec. 4.2). **(iii)** We undertake comprehensive analysis of the efficacy of each component, hyper-parameters, *etc.* (Sec. 4.3).

## 2 Related Work

**Combating biases with known types and labels.** Many debiasing approaches require explicit bias types and bias labels for each training sample. A large group of strategies aims at disentangling spurious and intrinsic features (Moyer et al. 2018). For example, EnD (Tartaglione, Barbano, and Grangetto 2021) designs regularizers to disentangle representations with the same bias label and entangle features with the same target label; BiasCon (Hong and Yang 2021) pulls samples with the same target label but different bias labels closer in the feature space based on contrastive learning; and some other studies learn disentangled representation by mutual information minimization (Zhu et al. 2021; Kim et al. 2019; Ragonesi et al. 2021). Another classic approach is to reweigh/resample training samples based on sample number or loss of different explicit groups (Li, Li, and Vasconcelos 2018; Sagawa et al. 2020b; Li and Vasconcelos 2019), or even to synthesize samples (Agarwal, Shetty, and Fritz 2020). Besides, Sagawa et al. (2020a) and Goel et al. (2021) intend to improve the worst-group performance through group distributionally robust optimization (Goh and Sim 2010) and Cycle-GAN (Zhu et al. 2017) based data augmentation, respectively. Furthermore, IRM (Arjovsky et al. 2019) is designed to learn a representation that performs well in all environments; domain-independent classifiers are introduced by Wang et al. (2020) to accomplish target tasks in each known bias situation.

**Combating biases with known types.** To alleviate expensive bias annotation costs, some bias-tailored methods relax the demands by requiring only the bias types (Geirhos et al. 2018). Bahng et al. (2020) elaborately design specific networks based on the bias types to obtain biased representations on purpose (*e.g.*, using 2D CNNs to extract static bias in action recognition). Then, the debiased representation is learned by encouraging it to be independent of the biased

one. Wang et al. (2019) try to project the model’s representation onto the subspace orthogonal to the texture-biased representation. SoftCon (Hong and Yang 2021) serves as an extension of BiasCon to handle the cases where only the bias type is available. In addition, the ensemble approach that consists of a bias-type customized biased model and a debiased model is employed in natural language processing as well (He, Zha, and Wang 2019; Clark, Yatskar, and Zettlemoyer 2019; Cadene et al. 2019; Utama, Moosavi, and Gurevych 2020a; Clark, Yatskar, and Zettlemoyer 2020).

**Combating unknown biases.** Despite the effectiveness of the methodologies described above, the assumptions limit their applications, as manually discovering bias types heavily relies on experts’ knowledge and labeling bias attributes for each training sample is even more laborious. As a result, recent studies (Le Bras et al. 2020; Kim, Ghorbani, and Zou 2019; Hashimoto et al. 2018) try to obtain debiased models with unknown biases, which are more realistic. Nam et al. (2020) mine bias-conflicting samples with generalized cross entropy (GCE) loss (Zhang and Sabuncu 2018) and emphasize them using a designed weight assignment function. Kim et al. (2021) further synthesize diverse bias-conflicting samples via feature-level data augmentation, whereas Kim, Lee, and Choo (2021) directly generate them with SwapAE (Park et al. 2020). RNF (Du et al. 2021) uses the neutralized representations from samples with the same target label but different bias labels (generated by GCE-based biased models, the version that accesses real bias labels is called RNF-GT) to train the classification head alone. Besides GCE loss, feature clustering (Sohoni et al. 2020), early-stopping (Liu et al. 2021), forgettable examples (Yaghoobzadeh et al. 2021) and limited network capacity (Sanh et al. 2020; Utama, Moosavi, and Gurevych 2020b) are involved to identify bias-conflicting samples. Furthermore, Creager, Jacobsen, and Zemel (2021) and Lahoti et al. (2020) alternatively infer dataset partitions and enhance domain-invariant feature learning by min-max adversarial training. In addition to the identify-emphasize paradigm, Pezeshki et al. (2020) introduces a novel regularization method for decoupling feature learning dynamics in order to improve model robustness.

## 3 Methodology

### 3.1 Effective Bias-Conflicting Scoring

Due to the explicit bias information is not available, we try to describe how likely input  $x$  is a bias-conflicting sample via the **bias-conflicting (b-c) score**:  $s(x, y) \in [0, 1]$ ,  $y \in \{1, 2, \dots, C\}$  stands for the target label. A larger  $s(x, y)$  indicates that  $x$  is harder to be recognized via unintended decision rules. As models are prone to fitting shortcuts, previous studies (Kim, Lee, and Choo 2021; Liu et al. 2021) resort model’s output probability on target class to define  $s(x, y)$  as  $1 - p(y|\hat{f}(x))$ , where  $p(k|\hat{f}(x)) = \frac{e^{\hat{f}(x)[k]}}{\sum_{k'=1}^C e^{\hat{f}(x)[k']}}$ ,  $\hat{f}$  is an auxiliary biased model and  $\hat{f}(x)[k]$  denotes the  $k^{th}$  index of logits  $\hat{f}(x)$ . Despite this, over-parameterized networks tend to “memorize” all samples, resulting in low scores for the real bias-conflicting samples as well. To avoid it, we propose the following two strategies. The whole scoring framework

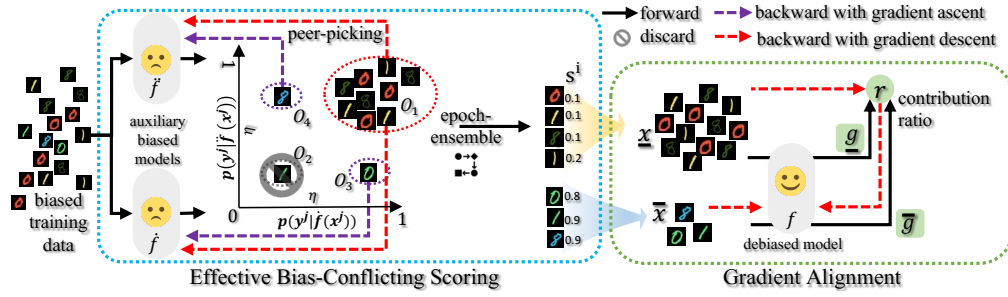


Figure 2: Our debiasing scheme. Stage I: training auxiliary biased models  $\hat{f}, \tilde{f}$  with peer-picking and epoch-ensemble to score the likelihood that a sample is bias-conflicting (Sec. 3.1). Stage II: learning debiased model  $f$  with gradient alignment (Sec. 3.2). A dashed arrow starting from a sample cluster indicates that the model is updated with gradients from these samples.

---

**Algorithm 1:** Effective bias-Conflicting Scoring (ECS)

---

**Input:**  $\mathcal{D}=\{(x^i, y^i)\}_{i=1}^N$ ; initial models  $\hat{f}^0, \tilde{f}^0$  and b-c scores  $\{s^i \leftarrow 0\}_{i=1}^N$ ; loss function  $\ell$ ; threshold  $\eta$ .

```

1 for  $t = 0$  to  $T - 1$  do
2    $\mathcal{B}=\{(x^j, y^j)\}_{j=1}^B \leftarrow \text{Batch}(\mathcal{D})$  // batch size  $B$ 
3    $\{p(y^j|\hat{f}^t(x^j)), \{p(y^j|\tilde{f}^t(x^j))\} \leftarrow \text{Forward}(\mathcal{B}, \hat{f}^t, \tilde{f}^t)$ ;
4    $\hat{l}^t \leftarrow 0$ ;  $\tilde{l}^t \leftarrow 0$ ; // initialize loss
5   for  $j = 1$  to  $B$  do
6     if  $p(y^j|\hat{f}^t(x^j)) > \eta$  and  $p(y^j|\tilde{f}^t(x^j)) > \eta$  then
7        $\hat{l}^t += \ell(\hat{f}^t(x^j), y^j)$ ;  $\tilde{l}^t += \ell(\tilde{f}^t(x^j), y^j)$ ;
8     else if  $p(y^j|\hat{f}^t(x^j)) > \eta$  and  $p(y^j|\tilde{f}^t(x^j)) \leq \eta$  then
9        $\hat{l}^t -= \ell(\hat{f}^t(x^j), y^j)$ ;
10    else if  $p(y^j|\hat{f}^t(x^j)) \leq \eta$  and  $p(y^j|\tilde{f}^t(x^j)) > \eta$  then
11       $\tilde{l}^t -= \ell(\tilde{f}^t(x^j), y^j)$ ;
12     $\hat{f}^{t+1} \leftarrow \text{Update}(\hat{f}^t, \frac{\hat{l}^t}{B})$ ;  $\tilde{f}^{t+1} \leftarrow \text{Update}(\tilde{f}^t, \frac{\tilde{l}^t}{B})$ ;
13    if  $(t + 1) \% T' = 0$  then
14      for  $i = 1$  to  $N$  do
15         $s^i += \frac{T'}{T} [1 - \frac{p(y^i|\hat{f}^{t+1}(x^i)) + p(y^i|\tilde{f}^{t+1}(x^i))}{2}]$ 

```

**Output:** the estimated b-c scores  $\{s^i\}_{i=1}^N$ .

---

is summarized in Alg. 1 (noting that the “for” loop is used for better clarification, it can be avoided in practice).

**Training auxiliary biased models with peer-picking.** Deliberately amplifying the auxiliary model’s bias seems to be a promising strategy for better scoring (Nam et al. 2020), as heavily biased models can assign high b-c scores to bias-conflicting samples. We achieve this by **confident-picking** — only picking samples with confident predictions (which are more like bias-aligned samples) to update auxiliary models. Nonetheless, a few bias-conflicting samples can still be overfitted and the memorization will be strengthened with continuous training. Thus, with the assist of **peer model**, we propose **peer-picking**, a co-training-like (Han et al. 2018) paradigm, to train auxiliary biased models (more auxiliary models can get slightly better results but raise costs).

Our method maintains two auxiliary biased models  $\hat{f}$  and  $\tilde{f}$  simultaneously (identical structure). Considering a training set  $\mathcal{D} = \{(x^i, y^i)\}_{i=1}^N$  with  $B$  samples in each batch, with a threshold  $\eta \in (0, 1)$ , each model divides samples into confident and unconfident groups relying on the output prob-

abilities on target classes. Consequently, four clusters are formed as shown in Fig. 2. For the red cluster ( $\mathcal{O}_1$ ), since two models are confident on them, it is reasonable to believe that they are indeed bias-aligned samples, therefore we pick up them to update model via gradient descent as usual (L7, 12 of Alg. 1). While the gray cluster ( $\mathcal{O}_2$ ), on which both two models are unconfident, will be discarded outright as they might be bias-conflicting samples. The remaining purple clusters ( $\mathcal{O}_3$  and  $\mathcal{O}_4$ ) indicate that some samples may be bias-conflicting, but they are memorized by one of auxiliary models. Inspired by the work for handling noisy labels (Han et al. 2020), we endeavor to force the corresponding model to forget the memorized suspicious samples via gradient ascent (L9, 11, 12). We average the output results of the two heavily biased models  $\hat{f}$  and  $\tilde{f}$  to obtain b-c scores (L15).

**Collecting results with epoch-ensemble.** During the early stage of training, b-c scores  $\{s^i\}$  ( $s^i := s(x^i, y^i)$ ) of real bias-conflicting samples are usually higher than those of bias-aligned ones, while the scores may be indistinguishable at the end of training due to overfitting. Unfortunately, selecting an optimal moment for scoring is strenuous. To avoid tedious hyper-parameter tuning, we collect results every  $T'$  iterations (typically every epoch in practice, *i.e.*,  $T' = \lfloor \frac{N}{B} \rfloor$ ) and adopt the ensemble averages of multiple results as the final b-c scores (L15). We find that the ensemble can alleviate the randomness of a specific checkpoint and achieve superior results without using tricks like early-stopping.

### 3.2 Gradients Alignment

Then, we attempt to train the debiased model  $f$ . We focus on an important precondition of the presence of biased models: the training objective can be achieved through unintended decision rules. To avoid it, one should develop a new learning objective that cannot be accomplished by these rules. The most straightforward inspiration is the use of plain reweighting (Rew) to intentionally rebalance sample contributions from different domains (Sagawa et al. 2020b):

$$\mathcal{R}_{\text{Rew}} = \sum_{i=1}^N \frac{\bar{N}}{\gamma \cdot \underline{N}} \cdot \ell(f(\underline{x}^i), y^i) + \sum_{j=1}^{\bar{N}} \ell(f(\bar{x}^j), y^j), \quad (1)$$

where  $\bar{N}$  and  $\underline{N}$  are the number of bias-conflicting and bias-aligned samples respectively,  $\gamma \in (0, \infty)$  is a reserved hyper-parameter to conveniently adjust the tendency: when

$\gamma \rightarrow 0$ , models intend to exploit bias-aligned samples more and when  $\gamma \rightarrow \infty$ , the behavior is reversed. As depicted in Fig. 3, assisted with reweighting, unbiased accuracy sky-rockets in the beginning, indicating that the model tends to learn intrinsic features in the first few epochs, while declines gradually, manifesting that the model is biased progressively (adjusting  $\gamma$  can not reverse the tendency).

The above results show that the static ratio between  $\bar{N}$  and  $\underline{N}$  is not a good indicator to show how balanced the training is, as the influence of samples can fluctuate during training. Accordingly, we are inspired to directly choose gradient statistics as a metric to indicate whether the training is overwhelmed by bias-aligned samples. Let us revisit the commonly used cross-entropy loss:  $\ell(f(x), y) = -\sum_{k=1}^C \mathbb{I}_{k=y} \log p(k|f(x))$ . For a sample, the gradient on logits  $f(x)$  is given by  $\nabla_{f(x)} \ell(f(x), y) = [\frac{\partial \ell(f(x), y)}{\partial f(x)[1]}, \frac{\partial \ell(f(x), y)}{\partial f(x)[2]}, \dots, \frac{\partial \ell(f(x), y)}{\partial f(x)[C]}]^T$ . We define the current gradient contribution of sample  $x$  as  $g(x, y|f) = \|\nabla_{f(x)} \ell(f(x), y)\|_1$ , which can be efficiently calculated by  $\sum_{k=1}^C |\frac{\partial \ell(f(x), y)}{\partial f(x)[k]}| = 2|\frac{\partial \ell(f(x), y)}{\partial f(x)[y]}| = 2 - 2p(y|f(x))$ . Assuming within the  $t^{\text{th}}$  iteration ( $t \in [0, T-1]$ ), the batch is composed of  $\underline{B}^t$  bias-aligned and  $\bar{B}^t$  bias-conflicting samples ( $B$  in total,  $\underline{B}^t \gg \bar{B}^t$  under our concerned circumstance). The accumulated gradient contributions generated by bias-aligned samples are denoted as  $\underline{g}^t = \sum_{i=1}^{\underline{B}^t} g(\underline{x}^i, y^i|f^t)$ , similarly for the contributions of bias-conflicting samples:  $\bar{g}^t$ .

We present the statistics of  $\{\bar{g}^t\}_{t=0}^{T-1}$  and  $\{\underline{g}^t\}_{t=0}^{T-1}$  when learning with standard ERM (Vanilla), Eq. (1) (Rew) in Fig. 4. For vanilla training, we find the gradient contributions of bias-aligned samples overwhelm that of bias-conflicting samples at the beginning, thus the model becomes biased towards spurious correlations rapidly. Even though at the late stage, the gap in gradient contributions shrinks, it is hard to rectify the already biased model. For Rew, we find the contributions of bias-conflicting and bias-aligned samples are relatively close at the beginning (compared to those under Vanilla), thus both of them can be well learned. Nonetheless, the bias-conflicting samples are memorized soon due to their small quantity, and the gradient contributions from the bias-conflicting samples become smaller than that of the bias-aligned samples gradually, leading to biased models step by step.

The above phenomena are well consistent with the accuracy curves in Fig. 3, indicating that the gradient statistics can be a useful ‘‘barometer’’ to reflect the optimization process. Therefore, the core idea of gradient alignment is to rebalance bias-aligned and bias-conflicting samples according to their currently produced gradient contributions. Within the  $t^{\text{th}}$  iteration, We define the contribution ratio  $r^t$  as:

$$r^t = \frac{\bar{g}^t}{\gamma \cdot \underline{g}^t} = \frac{\sum_{j=1}^{\bar{B}^t} [1 - p(y^j|f^t(\bar{x}^j))]}{\gamma \cdot \sum_{i=1}^{\underline{B}^t} [1 - p(y^i|f^t(\underline{x}^i))]}, \quad (2)$$

where  $\gamma$  plays a similar role as in Rew. Then, with  $r^t$ , we rescale the gradient contributions derived from bias-aligned samples to achieve alignment with that from bias-conflicting ones, which can be simply implemented by reweighting the

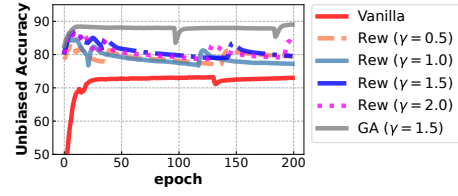


Figure 3: Unbiased accuracy (%) on Colored MNIST.

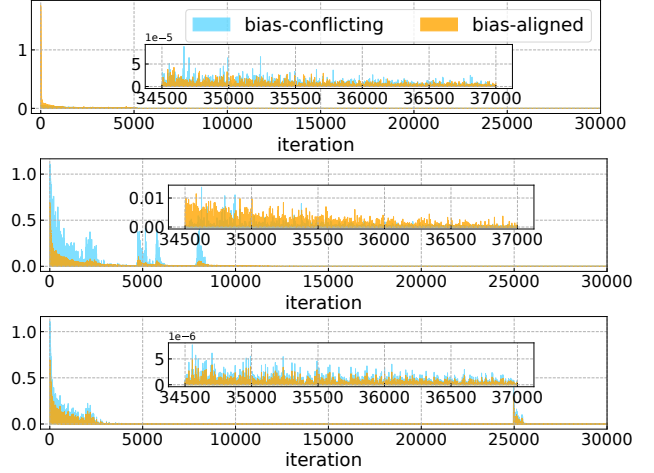


Figure 4: Statistics of  $\{\bar{g}^t\}_{t=0}^{T-1}$  and  $\{\underline{g}^t\}_{t=0}^{T-1}$ . Vanilla (top), Rew (middle), GA (bottom). Results in the late stage are enlarged and shown in each figure.

learning objective for the  $t^{\text{th}}$  iteration:

$$\mathcal{R}_{GA}^t = \sum_{i=1}^{\underline{B}^t} r^t \cdot \ell(f^t(\underline{x}^i), y^i) + \sum_{j=1}^{\bar{B}^t} \ell(f^t(\bar{x}^j), y^j), \quad (3)$$

i.e., the modulation weight is adaptively calibrated in each iteration. As shown in Eq. (2) and (3), GA only needs negligible computational extra cost ( $1 \times$  forward and backward as usual, only increases the cost of computing  $r^t$ ). As shown in Fig. 4, GA can dynamically balance the contributions throughout the whole training process. Correspondingly, it obtains optimal and stable predictions as demonstrated in Fig. 3 and multiple other challenging datasets in Sec. 4. Noting that as bias-conflicting samples are exceedingly scarce, it is unrealistic to ensure that every class can be sampled in one batch, thus all classes share the same ratio in our design.

To handle unknown biases, we simply utilize the estimated b-c score  $\{s^i\}_{i=1}^N$  and a threshold  $\tau$  to assign input  $x$  as bias-conflicting ( $s(x, y) \geq \tau$ ) or bias-aligned ( $s(x, y) < \tau$ ) here. For clarity, GA with the pseudo annotations (bias-conflicting or bias-aligned) produced by ECS will be denoted as ‘ECS+GA’ (similarly, ‘ECS+ $\Delta$ ’ represents combining ECS with method  $\Delta$ ).

## 4 Experiments

### 4.1 Experimental Settings

**Datasets.** We mainly conduct experiments on five benchmark datasets. For Colored MNIST (C-MNIST), the task is

$\rho$	Colored MNIST				Corrupted CIFAR10 <sup>1</sup>				Corrupted CIFAR10 <sup>2</sup>				B-Birds	B-CelebA
	95%	98%	99%	99.5%	95%	98%	99%	99.5%	95%	98%	99%	99.5%	95%	99%
Vanilla	85.7	73.6	60.7	45.4	44.9	30.4	22.4	17.9	42.7	27.2	20.6	17.4	77.1	77.4
Focal	86.7	75.8	62.4	45.9	45.5	30.7	22.9	17.8	41.9	26.9	21.0	17.0	78.6	78.1
GEORGE	87.0	76.2	62.4	46.4	44.6	29.5	21.8	17.9	44.2	27.3	20.7	17.7	79.3	78.2
LfF	88.2	86.7	80.3	73.2	59.6	50.4	42.9	34.6	58.5	49.0	42.2	33.4	80.4	84.4
DFA	89.8	86.9	81.8	74.1	58.2	50.0	41.8	35.6	58.6	48.7	41.5	35.2	79.5	84.3
SD	86.7	73.9	59.7	42.4	43.1	28.6	21.6	17.7	41.4	27.0	20.0	17.5	76.8	77.8
ECS+Rew	91.8	88.6	84.2	78.9	58.5	47.5	38.6	33.4	61.4	53.2	47.4	40.3	82.7	88.3
ECS+ERew	91.0	87.5	81.4	71.3	59.8	47.9	38.5	30.2	62.2	51.1	41.4	25.9	84.9	80.5
ECS+PoE	80.2	75.4	64.4	50.0	54.4	48.7	<b>45.6</b>	<b>42.7</b>	47.9	40.3	36.8	<b>42.4</b>	85.8	81.1
ECS+GA	<b>92.1</b>	<b>89.5</b>	<b>86.4</b>	<b>79.9</b>	<b>61.0</b>	<b>51.7</b>	42.6	35.0	<b>64.1</b>	<b>57.0</b>	<b>50.0</b>	41.8	<b>86.1</b>	<b>89.5</b>
<sup>†</sup> REBIAS	85.5	74.0	61.1	44.5	44.8	29.9	22.4	17.7	41.5	27.0	20.6	17.9	77.5	78.1
<sup>†</sup> Rew	91.5	87.9	83.8	77.6	59.1	48.9	40.4	33.4	61.1	53.1	46.9	41.2	86.0	90.7
<sup>†</sup> RNF-GT	84.3	75.9	66.3	59.1	52.1	39.1	30.6	22.2	50.3	34.9	27.9	19.8	81.2	85.1
<sup>†</sup> BiasCon	90.9	86.7	83.0	79.0	59.0	48.6	39.0	32.4	60.0	49.9	43.0	37.4	84.1	90.4
<sup>†</sup> GA	92.4	89.1	85.7	80.4	61.5	52.9	43.5	33.9	64.5	56.9	51.1	43.6	87.9	92.3

Table 1: Overall unbiased accuracy (%) over three runs. Best results with unknown biases are shown in bold. <sup>†</sup> indicates that the method requires prior knowledge regarding bias.

to recognize digits (0-9), in which the images of each target class are dyed by the corresponding color with probability  $\rho \in \{95\%, 98\%, 99\%, 99.5\%\}$  and by other colors with probability  $1 - \rho$  (a higher  $\rho$  indicates more severe biases). Similarly, for Corrupted CIFAR10, each object class in it holds a spurious correlation with a corruption type. Two sets of corruption protocols are utilized, leading to two biased datasets (Nam et al. 2020): Corrupted CIFAR10<sup>1</sup> and CIFAR10<sup>2</sup> (C-CIFAR10<sup>1</sup>, C-CIFAR10<sup>2</sup>) with  $\rho \in \{95\%, 98\%, 99\%, 99.5\%\}$ . In Biased Waterbirds (B-Birds), “waterbirds” and “landbirds” are highly correlated with “wet” and “dry” habitats (95% bias-aligned samples, *i.e.*,  $\rho=95\%$ ). Consequently, the task aiming to distinguish images as “waterbird” or “landbird” can be influenced by background. To focus on the debiasing problem, we balance the number of images per class. In Biased CelebA (B-CelebA), blond hair is predominantly found on women, whereas non-blond hair mostly appears on men ( $\rho=99\%$ ). When the goal is to classify the hair color as “blond” or “non-blond”, the information of gender (“male” or “female” in this dataset) can be served as a shortcut (Nam et al. 2020).

**Compared methods.** We choose various methods for comparison: standard ERM (Vanilla), Focal loss (Lin et al. 2017), plain reweighting (Sagawa et al. 2020b) (Rew, ECS+Rew), REBIAS (Bahng et al. 2020), BiasCon (Hong and Yang 2021), RNF-GT (Du et al. 2021), GEORGE (Sohoni et al. 2020), LfF (Nam et al. 2020), DFA (Kim et al. 2021), SD (Pezeshki et al. 2020), ERew (Clark, Yatskar, and Zettlemoyer 2019) and PoE (Clark, Yatskar, and Zettlemoyer 2019) (ECS+ERew, ECS+PoE)<sup>1</sup>. Among them, REBIAS requires bias types; Rew, BiasCon, RNF-GT, and GA are performed with real conflicting/aligned annotations.

**Evaluation metrics.** We mainly report the overall unbiased accuracy, alongside the fairness performance in terms of DP

<sup>1</sup>As the auxiliary biased models used in ERew and PoE are designed for NLP tasks, here, we combine our ECS with them.

and EqOdd (Reddy et al. 2021).

**Implementation.** The studies for the previous debiasing approaches are usually conducted with varying network architectures and training schedules. We run the representative methods with identical configurations to make fair comparisons. We use an MLP with three hidden layers (each hidden layer comprises 100 hidden units) for C-MNIST, except for the biased models in REBIAS (using CNN). ResNet-20 (He et al. 2016) is employed for C-CIFAR10<sup>1</sup> and C-CIFAR10<sup>2</sup>. ResNet-18 is utilized for B-Birds and B-CelebA.

## 4.2 Main Results

**The proposed method achieves better performance than others.** The overall unbiased accuracy is reported in Tab. 1. Vanilla models commonly fail to produce acceptable results on unbiased test sets, and the phenomenon is aggravated as  $\rho$  goes larger. Different debiasing methods moderate bias propagation with varying degrees of capability. When compared to other SOTA methods, the proposed approach achieves competitive results on C-CIFAR10<sup>1</sup> and noticeable improvements on other datasets across most values of  $\rho$ . For instance, the vanilla model trained on C-CIFAR10<sup>2</sup> ( $\rho=99\%$ ) only achieves 20.6% unbiased accuracy, indicating that the model is heavily biased. While, ECS+GA leads to 50.0% accuracy, and exceeds other prevailing debiasing methods by 3%-30%. When applied to the real-world dataset B-CelebA, the proposed scheme also shows superior results, demonstrating that it can effectively deal with subtle actual biases. Though the main purpose of this work is to combat unknown biases, we find GA also achieves better performance compared to the corresponding competitors when the prior information is available.

**Plain reweighting is an important baseline.** We find Rew (and ECS+Rew) can achieve surprising results compared with recent SOTA methods, while it is overlooked by some studies. The results also indicate that explicitly balancing the

	Colored MNIST			Corrupted CIFAR-10 <sup>1</sup>			Corrupted CIFAR-10 <sup>2</sup>			Biased Waterbirds			Biased CelebA		
	best $\uparrow$	last $\uparrow$	$\Delta$ $\downarrow$	best $\uparrow$	last $\uparrow$	$\Delta$ $\downarrow$	best $\uparrow$	last $\uparrow$	$\Delta$ $\downarrow$	best $\uparrow$	last $\uparrow$	$\Delta$ $\downarrow$	best $\uparrow$	last $\uparrow$	$\Delta$ $\downarrow$
LfF	86.7	75.1	11.6	50.4	49.4	1.0	49.0	47.3	1.7	80.4	76.9	3.5	84.4	61.0	23.4
DFA	86.9	81.2	5.8	50.0	47.9	2.1	48.7	46.0	2.7	79.5	74.5	5.0	84.3	73.2	11.0
ECS+Rew	88.6	79.6	9.0	47.5	42.6	5.0	53.2	48.8	4.4	82.7	77.2	5.5	88.3	80.5	7.8
ECS+GA	<b>89.5</b>	<b>88.9</b>	<b>0.6</b>	<b>51.7</b>	<b>50.8</b>	<b>0.9</b>	<b>57.0</b>	<b>55.4</b>	<b>1.6</b>	<b>86.1</b>	<b>85.5</b>	<b>0.6</b>	<b>89.5</b>	<b>87.4</b>	<b>2.1</b>
<sup>†</sup> Rew	87.9	79.2	8.7	48.9	45.3	3.6	53.1	47.5	5.6	86.0	78.8	7.3	90.7	82.8	7.8
<sup>†</sup> BiasCon	86.7	79.4	7.3	48.6	40.9	7.7	49.9	42.0	7.9	84.1	78.5	5.7	90.4	75.2	15.2
<sup>†</sup> GA	89.1	88.6	0.4	52.9	49.9	2.9	56.9	55.8	1.1	87.9	87.7	0.2	92.3	91.8	0.4

Table 2: Overall accuracy of the best epoch, the last epoch, and their difference (%). Results on Colored MNIST and Corrupted CIFAR-10 are under  $\rho = 98\%$ .

contributions is extremely important and effective.

**Early-stopping is not necessary for GA to select models.** Plain reweighting requires strong regularizations such as early-stopping to produce satisfactory results (Byrd and Lipton 2019; Sagawa et al. 2020a), implying that the results are not stable. Due to the nature of combating unknown biases, the unbiased validation set is not available, thus recent studies choose to report the best results among epochs (Nam et al. 2020; Kim et al. 2021) for convenient comparison. We follow this evaluation protocol in Tab. 1. However, in the absence of prior knowledge, deciding when to stop can be troublesome, thus some results in Tab. 1 are excessively optimistic. We claim that if the network can attain dynamic balance throughout the training phase, such early-stopping may not be necessary. We further provide the last epoch results in Tab. 2 to validate it. We find that some methods suffer from serious performance degradation. On the contrary, GA achieves steady results (with the same and fair training configurations). In other words, our method shows superiority under two model selection strategies simultaneously.

**The proposed method has strong performance on fairness metrics as well.** As shown in Tab. 3, our method also obtains significant improvement in terms of DP and EqOdd. These results further demonstrate that the proposed method is capable of balancing bias-aligned and bias-conflicting samples, as well as producing superior and impartial results.

### 4.3 Analysis and Discussion

**ECS shows superior ability to mine bias-conflicting samples.** We separately verify the effectiveness of each component of ECS on C-MNIST ( $\rho=98\%$ ) and B-CelebA. A good bias-conflicting scoring method should prompt superior precision-recall curves for the mined bias-conflicting samples, *i.e.*, give real bias-conflicting (aligned) samples high (low) scores. Therefore, we provide the average precision (AP) in Tab. 4. When comparing #0, #4, #5, and #6, we observe that epoch-ensemble, confident-picking and peer model all can improve the scoring method. In addition, as shown in Tab. 1, ECS+GA achieves results similar to GA with the help of ECS; ERew, PoE, and Rew combined with ECS also successfully alleviate biases to some extent, demonstrating that the proposed ECS is feasible, robust, and can be adopted in a variety of debiasing approaches.

We further compare the methods: #1 collecting results

	DP $\uparrow$ EqOdd $\uparrow$			DP $\uparrow$ EqOdd $\uparrow$	
	DP $\uparrow$	EqOdd $\uparrow$		DP $\uparrow$	EqOdd $\uparrow$
Vanilla	0.57	0.57	Vanilla	0.43	0.43
LfF	0.63	0.61	LfF	<b>0.80</b>	0.76
DFA	0.55	0.55	DFA	0.69	0.76
ECS+Rew	0.61	0.60	ECS+Rew	0.59	0.68
ECS+GA	<b>0.99</b>	<b>0.99</b>	ECS+GA	0.73	<b>0.91</b>

Table 3: Performance in terms of DP and EqOdd on Biased Waterbirds (left) and Biased CelebA (right).

with early-stopping in JTT (Liu et al. 2021), #2 training auxiliary biased model with GCE loss (and #3 collecting results with epoch-ensemble on top of it). When comparing #1 and #4, both early-stopping and epoch-ensemble can reduce the overfitting to bias-conflicting samples when training biased models, yielding more accurate scoring results. However, early-stopping is laborious to tune (Liu et al. 2021), whereas epoch-ensemble is more straightforward and robust. From #2 and #3, we see that epoch-ensemble can also enhance other strategies. Comparing #3 and #5, GCE loss is helpful, while confident-picking gains better results. Noting that though co-training with peer model raises some costs, it is not computationally complex and can yield significant benefits (#6), and even without peer model, #5 still outperform previous ways. Peer models are expected to better prevent bias-conflicting samples from affecting the training, so we can get better auxiliary biased models. Though the only difference between peer models is initialization in our experiments, as DNNs are highly nonconvex, different initializations can lead to different local optimal (Han et al. 2018).

**Why does GA outperform counterparts?** Focal loss, LfF, DFA, and ERew just reweight a sample with the information from itself (individual information), different from them, GA, as well as Rew, use global information within one batch to obtain modulation weight. Correspondingly, the methods based on individual sample information can not maintain the contribution balance between bias-aligned and bias-conflicting samples, which is crucial for this problem as presented in Sec. 4.2. Different from the static rebalance method Rew, we propose a dynamic rebalance training strategy with aligned gradient contributions throughout the learning process, which enforces models to dive into intrinsic features instead of spurious correlations. Learning with

	Epoch- Ensemble	Confident- Picking	Peer Model	C-MNIST	B-CelebA
#0 VM				27.0	13.3
#1 ES				45.6	47.9
#2 GCE				37.0	27.8
#3 GCE + EE	✓			89.3	52.1
#4 ECS	✓			53.8	46.5
#5 ECS	✓	✓		95.0	61.5
#6 ECS	✓	✓	✓	<b>98.8</b>	<b>67.6</b>

Table 4: Average Precision (%) of the mined bias-conflicting samples. VM represents scoring with vanilla model.

w.r.t. left color bias	a	a	c	c	Avg.
w.r.t. right color bias	a	c	a	c	
LfF <sup>§</sup>	99.6	4.7	<b>98.6</b>	5.1	52.0
PGI <sup>§</sup>	98.6	82.6	26.6	9.5	54.3
EIL <sup>§</sup>	<b>100.0</b>	<b>97.2</b>	70.8	10.9	69.7
DebiAN <sup>§</sup>	<b>100.0</b>	95.6	76.5	16.0	72.0
ECS+GA	<b>100.0</b>	89.7	96.1	<b>24.3</b>	<b>77.5</b>

Table 5: Accuracies (%) on four test groups of Multi-Color MNIST. <sup>§</sup>states the reported results from DebiAN. a: aligned; c: conflicting.

GA, as demonstrated in Fig. 3 and Tab. 2, produces improved results with no degradation. The impact of GA on the learning trajectory presented in Fig. 4 also shows that GA can schedule the optimization processes and take full advantage of the potential of different samples. Moreover, unlike the methods for class imbalance (Cui et al. 2019; Tan et al. 2021; Zhao et al. 2020), we try to rebalance the contributions of implicit groups rather than explicit categories.

**The method can manage multiple biases.** Most debiasing studies (Nam et al. 2020; Kim et al. 2021) only discussed single bias. However, there may be multiple biases, which are more difficult to analyze. To study the multiple biases, we adopt the Multi-Color MNIST dataset (Li, Hoogs, and Xu 2022) which holds two bias attributes: left color ( $\rho=99\%$ ) and right color ( $\rho=95\%$ ). In such training sets, though it seems more intricate to group a sample as bias-aligned or bias-conflicting (as a sample can be aligned or conflicting w.r.t. left color bias or right color bias separately), we still simply train debiased models with GA based on the b-c scores obtained via ECS. We evaluate ECS+GA on four test groups separately and present them in Tab. 5. We find the proposed method also can manage the multi-bias situation.

**The introduced hyper-parameters are principled and insensitive.** Though the hyper-parameters are critical for methods aimed at combating unknown biases, recent studies did not include an analysis for them. Here, we present the ablation studies on C-MNIST ( $\rho=98\%$ ) for the hyper-parameters ( $\eta$ ,  $\tau$ ,  $\gamma$ ) in our method as shown in Fig. 5. We find that the method performs well under a wide range of hyper-parameters. The determination of  $\eta$  is related to the number of categories and the difficulty of tasks, *e.g.*, 0.5

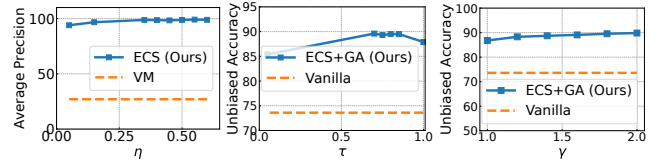


Figure 5: Ablation on thresholds  $\eta$ ,  $\tau$  and balance ratio  $\gamma$ .

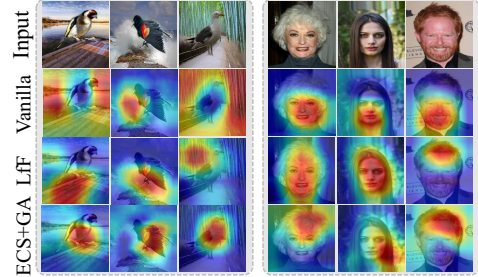


Figure 6: Activation maps on Waterbirds and CelebA.

for C-MNIST, 0.1 for C-CIFAR10<sup>1</sup> and C-CIFAR10<sup>2</sup> (10-class classification tasks), 0.9 for B-Birds and B-CelebA (2-class) here. To ensure the precision first,  $\tau$  is typically set to 0.8 for all experiments. For the balance ratio  $\gamma$ , though the results are reported with  $\gamma = 1.6$  for all settings on C-MNIST, C-CIFAR10<sup>1</sup> and C-CIFAR10<sup>2</sup>, 1.0 for B-Birds and B-CelebA, the proposed method is not sensitive to  $\gamma \in [1.0, 2.0]$ , which is reasonable as  $\gamma$  in such region makes the contributions from bias-conflicting samples close to that from bias-aligned samples.

**The proposed method makes decisions based on intrinsic features.** We visualize the activation maps via CAM (Zhou et al. 2016) in Fig. 6. Vanilla models usually activate regions related to biases when making predictions, *e.g.*, the background in B-Birds, the gender characteristics (faces, beards) in B-CelebA. LfF can focus attention on key areas in some situations, but there are still some deviations. Meanwhile, the proposed ECS+GA mostly utilizes compact essential features to make decisions.

## 5 Conclusions

Biased models can cause poor out-of-distribution performance and even negative social impacts. In this paper, we focus on combating unknown biases which is urgently required in realistic applications, and propose an enhanced two-stage debiasing method. In the first stage, an effective bias-conflicting scoring approach containing peer-picking and epoch-ensemble is proposed. Then we derive a new learning objective with the idea of gradient alignment in the second stage, which dynamically balances the gradient contributions from the mined bias-conflicting and bias-aligned samples throughout the learning process. Extensive experiments on synthetic and real-world datasets reveal that the proposed solution outperforms previous methods.

## Acknowledgments

This work is supported in part by the National Natural Science Foundation of China under Grant 62171248, the R&D Program of Shenzhen under Grant JCYJ20220818101012025, the PCNL KEY project (PCL2021A07), and Shenzhen Science and Technology Innovation Commission (Research Center for Computer Network (Shenzhen) Ministry of Education).

## References

- Agarwal, V.; Shetty, R.; and Fritz, M. 2020. Towards causal vqa: Revealing and reducing spurious correlations by invariant and co-variant semantic editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9690–9698.
- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*.
- Bahng, H.; Chun, S.; Yun, S.; Choo, J.; and Oh, S. J. 2020. Learning de-biased representations with biased representations. In *International Conference on Machine Learning*, 528–539. PMLR.
- Byrd, J.; and Lipton, Z. 2019. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning*, 872–881. PMLR.
- Cadene, R.; Dancette, C.; Cord, M.; Parikh, D.; et al. 2019. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems*, 32: 841–852.
- Clark, C.; Yatskar, M.; and Zettlemoyer, L. 2019. Don’t Take the Easy Way Out: Ensemble Based Methods for Avoiding Known Dataset Biases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4069–4082.
- Clark, C.; Yatskar, M.; and Zettlemoyer, L. 2020. Learning to Model and Ignore Dataset Bias with Mixed Capacity Ensembles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, 3031–3045.
- Creager, E.; Jacobsen, J.-H.; and Zemel, R. 2021. Environment inference for invariant learning. In *International Conference on Machine Learning*, 2189–2200. PMLR.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9268–9277.
- Du, M.; Mukherjee, S.; Wang, G.; Tang, R.; Awadallah, A.; and Hu, X. 2021. Fairness via Representation Neutralization. *Advances in Neural Information Processing Systems*, 34.
- Geirhos, R.; Rubisch, P.; Michaelis, C.; Bethge, M.; Wichmann, F. A.; and Brendel, W. 2018. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*.
- Goel, K.; Gu, A.; Li, Y.; and Re, C. 2021. Model Patching: Closing the Subgroup Performance Gap with Data Augmentation. In *International Conference on Learning Representations*.
- Goh, J.; and Sim, M. 2010. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1): 902–917.
- Han, B.; Niu, G.; Yu, X.; Yao, Q.; Xu, M.; Tsang, I.; and Sugiyama, M. 2020. SIGUA: Forgetting May Make Learning with Noisy Labels More Robust. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 4006–4016. PMLR.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I. W.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*.
- Hashimoto, T.; Srivastava, M.; Namkoong, H.; and Liang, P. 2018. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, 1929–1938. PMLR.
- He, H.; Zha, S.; and Wang, H. 2019. Unlearn Dataset Bias in Natural Language Inference by Fitting the Residual. In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, 132–142.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hong, Y.; and Yang, E. 2021. Unbiased Classification through Bias-Contrastive and Bias-Balanced Learning. *Advances in Neural Information Processing Systems*, 34.
- Kim, B.; Kim, H.; Kim, K.; Kim, S.; and Kim, J. 2019. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9012–9020.
- Kim, E.; Lee, J.; and Choo, J. 2021. BiasSwap: Removing dataset bias with bias-tailored swapping augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14992–15001.
- Kim, E.; Lee, J.; Lee, J.; Lee, J.; and Choo, J. 2021. Learning Debiased Representation via Disentangled Feature Augmentation. *arXiv preprint arXiv:2107.01372*.
- Kim, M. P.; Ghorbani, A.; and Zou, J. 2019. Multiaccuracy: Black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 247–254.
- Lahoti, P.; Beutel, A.; Chen, J.; Lee, K.; Prost, F.; Thain, N.; Wang, X.; and Chi, E. H. 2020. Fairness without demographics through adversarially reweighted learning. *arXiv preprint arXiv:2006.13114*.
- Le Bras, R.; Swayamdipta, S.; Bhagavatula, C.; Zellers, R.; Peters, M.; Sabharwal, A.; and Choi, Y. 2020. Adversarial filters of dataset biases. In *International Conference on Machine Learning*, 1078–1088. PMLR.
- Li, Y.; Li, Y.; and Vasconcelos, N. 2018. Resound: Towards action recognition without representation bias. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 513–528.
- Li, Y.; and Vasconcelos, N. 2019. Repair: Removing representation bias by dataset resampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9572–9581.
- Li, Z.; Hoogs, A.; and Xu, C. 2022. Discover and Mitigate Unknown Biases with Debiasing Alternate Networks. In *The European Conference on Computer Vision (ECCV)*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, E. Z.; Haghighi, B.; Chen, A. S.; Raghunathan, A.; Koh, P. W.; Sagawa, S.; Liang, P.; and Finn, C. 2021. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning*, 6781–6792. PMLR.
- Moyer, D.; Gao, S.; Breckelmanns, R.; Galstyan, A.; and Ver Steeg, G. 2018. Invariant Representations without Adversarial Training. *Advances in Neural Information Processing Systems*, 31: 9084–9093.

- Nam, J.; Cha, H.; Ahn, S.; Lee, J.; and Shin, J. 2020. Learning from Failure: Training Debiased Classifier from Biased Classifier. In *Advances in Neural Information Processing Systems*.
- Park, T.; Zhu, J.-Y.; Wang, O.; Lu, J.; Shechtman, E.; Efros, A. A.; and Zhang, R. 2020. Swapping autoencoder for deep image manipulation. *arXiv preprint arXiv:2007.00653*.
- Pezeshki, M.; Kaba, S.-O.; Bengio, Y.; Courville, A.; Precup, D.; and Lajoie, G. 2020. Gradient starvation: A learning proclivity in neural networks. *arXiv preprint arXiv:2011.09468*.
- Ragonesi, R.; Volpi, R.; Cavazza, J.; and Murino, V. 2021. Learning unbiased representations via mutual information backpropagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2729–2738.
- Reddy, C.; Sharma, D.; Mehri, S.; Romero-Soriano, A.; Shabanian, S.; and Honari, S. 2021. Benchmarking Bias Mitigation Algorithms in Representation Learning through Fairness Metrics. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2020a. Distributionally Robust Neural Networks. In *International Conference on Learning Representations*.
- Sagawa, S.; Raghunathan, A.; Koh, P. W.; and Liang, P. 2020b. An investigation of why overparameterization exacerbates spurious correlations. In *International Conference on Machine Learning*, 8346–8356. PMLR.
- Sanh, V.; Wolf, T.; Belinkov, Y.; and Rush, A. M. 2020. Learning from others’ mistakes: Avoiding dataset biases without modeling them. In *International Conference on Learning Representations*.
- Sohoni, N.; Dunnmon, J.; Angus, G.; Gu, A.; and Ré, C. 2020. No Subclass Left Behind: Fine-Grained Robustness in Coarse-Grained Classification Problems. *Advances in Neural Information Processing Systems*, 33.
- Tan, J.; Lu, X.; Zhang, G.; Yin, C.; and Li, Q. 2021. Equalization loss v2: A new gradient balance approach for long-tailed object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1685–1694.
- Tartaglione, E.; Barbano, C. A.; and Grangetto, M. 2021. EnD: Entangling and Disentangling deep representations for bias correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13508–13517.
- Utama, P. A.; Moosavi, N. S.; and Gurevych, I. 2020a. Mind the Trade-off: Debiasing NLU Models without Degrading the In-distribution Performance. In *Proceedings of the 58th Conference of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Utama, P. A.; Moosavi, N. S.; and Gurevych, I. 2020b. Towards Debiasing NLU Models from Unknown Biases. In *EMNLP*.
- Wang, H.; He, Z.; Lipton, Z. L.; and Xing, E. P. 2019. Learning Robust Representations by Projecting Superficial Statistics Out. In *International Conference on Learning Representations*.
- Wang, Z.; Qinami, K.; Karakozis, I. C.; Genova, K.; Nair, P.; Hata, K.; and Russakovsky, O. 2020. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8919–8928.
- Yaghoobzadeh, Y.; Mehri, S.; des Combes, R. T.; Hazen, T. J.; and Sordoni, A. 2021. Increasing Robustness to Spurious Correlations using Forgettable Examples. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 3319–3332.
- Zhang, Z.; and Sabuncu, M. R. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. In *32nd Conference on Neural Information Processing Systems (NeurIPS)*.
- Zhao, B.; Xiao, X.; Gan, G.; Zhang, B.; and Xia, S.-T. 2020. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13208–13217.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.
- Zhu, W.; Zheng, H.; Liao, H.; Li, W.; and Luo, J. 2021. Learning Bias-Invariant Representation by Cross-Sample Mutual Information Minimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15002–15012.