Identifying Spurious Biases Early in Training through the Lens of Simplicity Bias

Yu Yang
University of California, Los Angeles

Eric Gan
University of California, Los Angeles

Gintare Karolina Dziugaite
Google Deepmind

Baharan Mirzasoleiman
University of California, Los Angeles

Abstract

Neural networks trained with (stochastic) gradient descent have an inductive bias towards learning simpler solutions. This makes them highly prone to learning *spurious correlations* in the training data, that may not hold at test time. In this work, we provide the first theoretical analysis of the effect of simplicity bias on learning spurious correlations. Notably, we show that examples with spurious features are *provably* separable based on the model's output *early in training*. We further illustrate that if spurious features have a small enough noise-to-signal ratio, the network's output on majority of examples is almost exclusively determined by the spurious features, leading to poor *worst-group* test accuracy. Finally, we propose Spare, which identifies spurious correlations early in training, and utilizes importance sampling to alleviate their effect. Empirically, we demonstrate that Spare outperforms state-of-the-art methods by up to 21.1% in worst-group accuracy, while being up to 12x faster. We also show that Spare is a highly effective but lightweight method to *discover spurious correlations*. Code is available at https://github.com/BigML-CS-UCLA/SPARE.

1 INTRODUCTION

The *simplicity bias* of gradient-based training algorithms towards learning simpler solutions has been suggested as a key factor for the superior generalization performance of overparameterized neural networks (Hermann and Lampinen, 2020; Hu et al., 2020; Nakkiran et al., 2019; Neyshabur et al., 2014; Pezeshki et al., 2021; Shah et al., 2020). At the same time, it is conjectured to make neural networks vulnerable to learning *spurious correlations* frequently found in real-world datasets (Sagawa et al., 2019; Sohoni et al., 2020). Neural networks trained with gradient-based methods can exclusively rely on simple *spurious features* that are highly correlated with a class in the training data but are not predictive of the class in general, and remain invariant to the predictive but more complex *core features* (Shah et al., 2020). This results in a poor *worst-group test accuracy* on groups of examples where the spurious correlations do not hold (Shah et al., 2020; Teney et al., 2022). For example, in an image classification task, if the majority of images of a 'bird' appear on a 'sky' background, the classifier learns the sky instead of bird, and misclassifies birds that do not appear in the sky at test time.

An effective way to mitigate a spurious correlation and improve the worst-group test accuracy is to upweight examples that do not contain the spurious feature during training (Sagawa et al., 2019). However, inspecting all training examples to find such examples becomes prohibitive in real-world datasets. This has motivated a growing body of work on group inference: separating majority groups exhibiting spurious correlation with a class, from minority groups without the spurious correlation. Such methods first train a neural network with gradient methods to learn the spurious correlation. Then, they rely on model's misclassification (Liu et al., 2021), loss (Creager et al., 2021), or representations (Sohoni et al., 2020; Ahmed et al., 2020) at a certain point during training, as an indicative of minority examples. The time of group inference and how much to upweight the minority groups are heavily tuned based on a group-labeled validation data (Sohoni et al., 2020; Liu et al.,

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

2021; Creager et al., 2021; Ahmed et al., 2020).

Despite their success on simple benchmark datasets, we show that such methods suffer from several issues: (1) they often misidentify minority examples as majority and mistakenly downweight them; then, (2) to counteract the spurious correlation they need to heavily upweight their small inferred minority group. This magnifies milder spurious correlations that may exist in the minority group (Li et al., 2023) and harms the performance; as (3) there is no theoretical guideline for finding the time of group inference and group weights, such methods rely on extensive hyperparameter tuning. This limits their applicability and scalability.

In this work, we make several theoretical and empirical contributions towards addressing the above issues. First, we prove that the simplicity bias of gradient descent can be leveraged to identify spurious correlations. We analyze a two-layer fully connected neural network trained with SGD, and leverage recent results showing its early-time learning dynamics can be mimicked by training a linear model on the inputs (Hu et al., 2020). We show that the contribution of a spurious feature to the network output in the initial training phase increases linearly with the amount of spurious correlation. Thus, minority and majority groups can be *provably* separated based on the model's output, *early in training*. This enables more accurate identification of minorities, and limits the range of group inference to the first few training epochs, without extensive hyperparameter tuning.

Next, we show that once the initial linear model converges, if the noise-to-signal ratio of a spurious feature is lower than that of the core feature in a class, the network will not learn the core features of the majority groups. This explains prior empirical observations (Shah et al., 2020), by revealing *when and why* neural networks trained with gradient almost exclusively rely on spurious features and remain invariant to the predictive but more complex core features. To the best of our knowledge, this is the first analysis of the effect of SGD's simplicity bias on learning spurious vs core features.

Finally, we propose an efficient and lightweight method, Spare (SePArate early and REsample), that clusters model's output early in training, and leverage importance sampling based on inverse cluster sizes to mitigate spurious correlations. This results in a superior worst-group accuracy on more challenging tasks, without increasing the training time, or requiring extensive hyperparameter tuning. Unlike existing methods, Spare can operate without a group-labeled validation data, which allows it to *discover unknown spurious correlations.*

Our extensive experiments confirm that Spare achieves up to 42.9% higher worst-group accuracy over state-of-the-art on most commonly used benchmarks, including CMNIST (Alain et al., 2015) (with multiple minority groups), Waterbirds (Sagawa et al., 2019), CelebA (Liu et al., 2015) and UrbanCars (Li et al., 2023) (with multiple spurious correlations) while being up to 12x faster. On CMNIST, Spare performs well across varying noise-to-signal ratios, whereas other state-of-the-art methods struggle. Applied to Restricted ImageNet, a dataset without known spurious correlations or group-labeled validation set available for hyperparameter tuning, identifies the spurious correlation much more effectively than the state-of-the-art group inference methods and improves the model's accuracy on minority groups by up to 23.2% higher than them after robust training.

2 RELATED WORK

Mitigating Spurious Correlations With Group Inference. To mitigate spurious correlations, group inference methods typically fall into three categories: (1) end-of-training clustering, (2) Environment Invariance Maximization, and (3) misclassification methods. GEORGE (Sohoni et al., 2020) utilizes end-of-training clustering by first training a model using Empirical Risk Minimization (ERM). The clustered feature representations are then used to train a robust model with Group Distributionally Robust Optimization (GDRO) (Sagawa et al., 2019). Methods use Environment Invariance Maximization, EIIL (Creager et al., 2021) and PGI (Ahmed et al., 2020), train an initial model using ERM and then partition the data to maximize the Invariant Risk Minimization (IRM) objective (Arjovsky and Bottou, 2017). Then, EIIL trains the robust model with GDRO, whereas PGI minimizes the KL divergence of softmaxed logits for same-class samples across groups. The misclassification approach, exemplified by JTT (Liu et al., 2021) trains an ERM model for some epochs and identifies misclassified examples. The training set is then upsampled with these examples, and a robust model is trained using ERM on this augmented set.

State-of-the-art methods commonly misidentify minority examples as the majority and struggle to mitigate spurious correlations. They heavily rely on a group-labeled validation for hyperparameter tuning, and often increase training time during group inference or robust training. In contrast, guided by theory, Spare can accurately separate groups with spurious features in the first few epochs, eliminating the need for extensive hyperparameter tuning and achieving better performance on minority groups without extending training time.

Simplicity Bias. Simplicity bias of SGD in learning simpler functions before complex ones (Hermann and Lampinen, 2020; Hu et al., 2020; Nakkiran et al., 2019;

Neyshabur et al., 2014; Pezeshki et al., 2021; Shah et al., 2020) is empirically observed in various network architectures, including MobileNetV2, ResNet50, and DenseNet121 (Sandler et al., 2018; He et al., 2016; Shah et al., 2020). Hu et al. (2020) formally proved that initial learning dynamics of a two-layer FC network can be mimicked by a linear model and extended this to multi-layer FC and convolutional networks. While simplicity bias helps overparameterized networks generalize well, it is also *conjectured* to favor simpler features over complex ones, even if they are less predictive (Shah et al., 2020; Teney et al., 2022). However, the exact notion of the simplicity of features and the mechanism by which they are learned remain poorly understood except in certain simplistic settings (Nagarajan et al., 2020; Shah et al., 2020). In this work, we build on Hu et al. (2020) to rigorously specify the conditions and mechanism of learning spurious features in a two-layer FC network.

3 PROBLEM FORMULATION

Let $\mathcal{D}=\{(\boldsymbol{x}_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$ be n training data with features $\boldsymbol{x}_i \in \mathbb{R}^d$, and labels $y_i \in C = \{1,-1\}$. For simplicity, we consider binary classification with ℓ_2 loss, but our analysis generalizes to multi-class classification with CE loss, and other model architectures, as we will confirm experimentally.

Features & Groups. We assume every class $c \in C$ has a *core* feature $\boldsymbol{v}_c$, which is the invariant feature of the class that appears in both training and test sets. Besides, there is a set of *spurious* features $\mathcal{A}$ that are shared between classes and are present in both training and test sets, but may not have a spurious correlation with the labels at test time. For example, in the CMNIST dataset containing images of colored handwritten digits (Fig. 1), the digit is the core feature, and its color is the spurious feature. Assuming w.l.o.g. that all $\boldsymbol{v}_c, \boldsymbol{v}_s \in \mathbb{R}^d$ are orthogonal vectors, the feature vector of every example $\boldsymbol{x}_i$ in class c can be written as $\boldsymbol{x}_i = \boldsymbol{v}_c + \boldsymbol{v}_s + \boldsymbol{\xi}_i$, where $\boldsymbol{v}_s \in \mathcal{A}$, and each $\boldsymbol{\xi}_i$ is a noise vector drawn i.i.d. from $\mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_\xi)$. Noise-to-signal (NSR) ratio of a feature is defined as its variance over magnitude, i.e., $R_. = \sigma_./\|\boldsymbol{v}_.\|$. Features with smaller NSR are simpler to learn. Training examples can be partitioned into groups $g_{c,s}$ based on the combinations of their core and spurious features $(\boldsymbol{v}_c, \boldsymbol{v}_s)$.

Neural Network & Training. We consider a two-layer FC neural network with m hidden neurons:

$$f(\boldsymbol{x}; \boldsymbol{W}, \boldsymbol{z}) = \frac{1}{\sqrt{m}} \sum_{r=1}^{m} z_r \phi(\boldsymbol{w}_r^T \boldsymbol{x}/\sqrt{d}) = \frac{1}{\sqrt{m}} \boldsymbol{z}^T \phi(\boldsymbol{W}\boldsymbol{x}/\sqrt{d}),$$

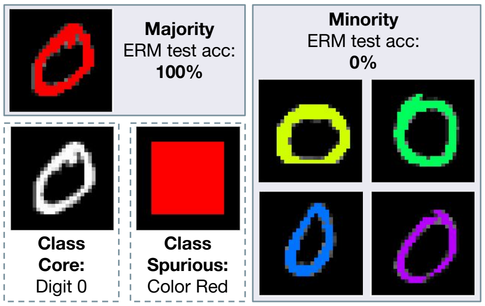

Figure 1: Colored MNIST as an example of datasets containing spurious correlations. Each digit is a class; the majority of digits in a class have a particular color, and the remaining digits are in other colors. Models trained with ERM learn to rely on spurious features (colors) instead of the core feature (digits) and thus do not perform well on groups of examples where the spurious correlation does not hold.

where $\boldsymbol{x} \in \mathbb{R}^d$ is the input, $\boldsymbol{W} = [\boldsymbol{w}_1, \cdots, \boldsymbol{w}_m]^T \in \mathbb{R}^{m\times d}$ is the weight matrix in the first layer, and $\boldsymbol{z} = [z_1, \cdots, z_m]^T \in \mathbb{R}^m$ is the weight vector in the second layer. Here $\phi : \mathbb{R} \to \mathbb{R}$ is a smooth or piecewise linear activation function (including ReLU, Leaky ReLU, Erf, Tanh, Sigmoid, Softplus, etc.) that acts entry-wise on vectors or matrices. We consider the following ℓ_2 training loss:

$$\mathcal{L}(\boldsymbol{W}, \boldsymbol{z}) = \frac{1}{2n} \sum_{i=1}^{n} (f(\boldsymbol{x}_i; \boldsymbol{W}, \boldsymbol{z}) - y_i)^2. \quad (1)$$

We train the network by applying gradient descent on the loss (1) starting from random initialization[1]:

$$\boldsymbol{W}_{t+1} = \boldsymbol{W}_t - \eta \nabla_{\boldsymbol{W}} \mathcal{L}(\boldsymbol{W}_t, \boldsymbol{z}_t), \quad (2)$$
$$\boldsymbol{z}_{t+1} = \boldsymbol{z}_t - \eta \nabla_{\boldsymbol{z}} \mathcal{L}(\boldsymbol{W}_t, \boldsymbol{z}_t), \quad (3)$$

Worst-group Error. We quantify the performance of the model based on its highest test error across groups $\mathcal{G} = \{g_{c,s}\}_{c,s}$ in all classes. Formally, *worst-group* test error is defined as:

$$\text{Err}_{wg} = \max_{g\in\mathcal{G}} \mathop{\mathbb{E}}_{(\boldsymbol{x}_i, y_i)\in g} [y_i \neq y_f(\boldsymbol{x}_i; \boldsymbol{W}, \boldsymbol{z})], \quad (4)$$

where $y_f(\boldsymbol{x}_i; \boldsymbol{W}, \boldsymbol{z})$ is the label predicted by the model. That is, Err_{wg} measures the highest fraction of examples that are incorrectly classified across all groups.

[1] Detailed assumptions on the activations, and initialization can be found in Sec. A

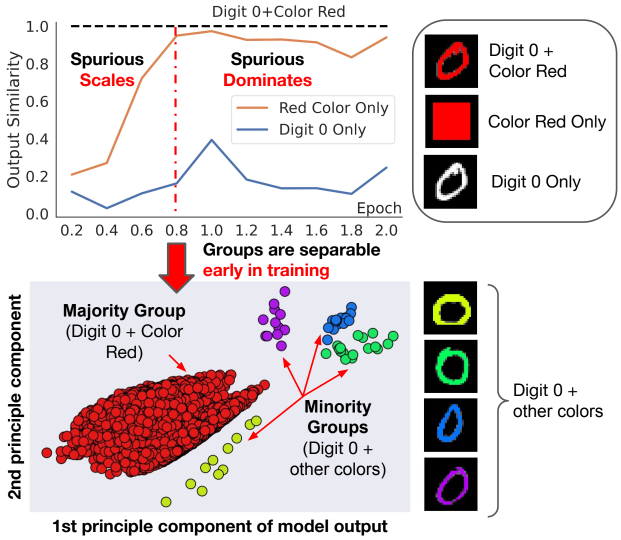

Figure 2: Training LeNet-5 on Colored MNIST. **Top**: Up to epoch 2, the network output is almost exclusively indicated by the color red (spurious feature in the majority group). **Bottom**: Majority and minority groups are separable based on the network output, e.g. via clustering. Minority groups that have a spurious feature in majority groups of other classes (yellow, purple, blue, green) are also separable from each other. Similar results on Waterbirds are shown in Fig. 6.

4 INVESTIGATING SPURIOUS FEATURE LEARNING IN NEURAL NETWORKS

We start by investigating how spurious features are learned during training a two-layer fully-connected neural network. Our analysis reveals two phases in early-time learning. First, in the initial training iterations, the contribution of a spurious feature to the network output increases linearly with the amount of the spurious correlation. Interestingly, if the majority group is sufficiently large, majority and minority groups are separable at this phase by the network output. Second, if the noise-to-signal ratio of the spurious feature of the majority group is smaller than that of the core feature, the network's output on the majority of examples in the class will be almost exclusively determined by the spurious feature and will remain mostly invariant to the core feature. Next, we will discuss the two phases in detail.

4.1 Spurious Features are Learned in the Initial Training Iterations

We start by analyzing the effect of spurious features on learning dynamics of a two-layer FC neural network trained with gradient descent in the initial training iterations. With the data model $\boldsymbol{x}_i = \boldsymbol{v}_c + \boldsymbol{v}_s + \boldsymbol{\xi}_i$ defined in Sec. 3, the following theorem shows that if a majority group is sufficiently large, contribution of its spurious feature to the model's output is magnified by the network at every step early in training.

Theorem 4.1. *Let $\alpha \in (0, \frac{1}{4})$ be a fixed constant. Suppose the number of training samples n and the network width m satisfy $n \gtrsim d^{1+\alpha}$ and $m \gtrsim d^{1+\alpha}$. Let n_c be the number of examples in class c, and $n_{c,s} = |g_{c,s}|$ be the size of group $g_{c,s}$ with label c and spurious feature $\boldsymbol{v}_s \in \mathcal{A}$. Then, under the setting of Sec. 3 there exist a constant $\nu_1 > 0$, such that with high probability, for all $0 \le t \le \nu_1 \cdot \sqrt{\frac{d^{1-\alpha}}{\eta}}$, the contribution of the core and spurious features to the network output can be quantified as follows:*

$$f(\boldsymbol{v}_c; \boldsymbol{W}_t, \boldsymbol{z}_t) = \frac{2\eta\zeta^2 c\|\boldsymbol{v}_c\|^2 t}{d}\left(\frac{n_c}{n} \pm O(d^{-\Omega(\alpha)})\right), \quad (5)$$

$$f(\boldsymbol{v}_s; \boldsymbol{W}_t, \boldsymbol{z}_t) = \frac{2\eta\zeta^2 c\|\boldsymbol{v}_s\|^2 t}{d}\Bigg(\frac{n_{c,s} - n_{c',s}}{n} \quad (6)$$

$$\pm O(d^{-\Omega(\alpha)})\Bigg),$$

where $c' = C \setminus c$, and ζ is the expected gradient of activation functions at random initialization.

The proof for Theorem 4.1 is detailed in Sec. B.2. At a high level, as the model is nearly linear in the initial $\nu_1 \cdot \frac{d \log d}{\eta}$ iterations, the contribution of the spurious feature $\boldsymbol{v}_s$ to the network output grows almost linearly with $(n_{c,s} - n_{c',s})\|\boldsymbol{v}_s\|^2$, at *every iteration* in the initial phase of training. Here, $n_{c,s} - n_{c',s}$ represents the correlation between the spurious feature and label c. When $n_{c,s} \gg n_{c',s}$, the spurious feature exists almost exclusively in the majority group of class c, and thus has a high correlation only with class c. In this case, if the magnitude of the spurious feature is significant, the contribution of the spurious feature to the model's output grows very rapidly, early in training. In particular, if $(n_{c,s} - n_{c',s})\|\boldsymbol{v}_s\|^2 \gg n_c\|\boldsymbol{v}_c\|^2$, the model's output is increasingly determined by the spurious feature, but not the core feature.

Remember from Sec. 3 that every example consists of a core and a spurious feature. As the effect of spurious features of the majority groups is amplified in the network output, the model's output will differ for examples in the majority and minority groups. The following corollary shows that the majority and minority groups are separable based on the network's output early in training. Notably, multiple minority groups with spurious features contained in majority groups of other classes are also separable.

Corollary 4.2 (**Separability of majority and minority groups**). *Suppose that for all classes, a major-*

ity group has at least K examples and a minority group has at most k examples. Then, under the assumptions of Theorem 4.1, examples in the majority and minority groups are nearly separable with high probability based on the model's output, early in training. That is, for all $0 \leq t \leq \nu_1 \cdot \sqrt{\frac{d^{1-\alpha}}{\eta}}$, with high probability, the following holds for at least $1 - O(d^{-\Omega(\alpha)})$ fraction of the training examples $\boldsymbol{x}_i$ in group $g_{c,s}$:

If $g_{c,s}$ is in a majority group in class $c = 1$:

$$f(\boldsymbol{x}_i;\boldsymbol{W}_t,z_t) \geq \frac{2\eta\zeta^2 t}{d}\Bigg(\frac{\|\boldsymbol{v}_s\|^2(K-k)}{n} + \xi \pm O(d^{-\Omega(\alpha)})\Bigg) + \rho(t,\phi,\Sigma), \tag{7}$$

If $g_{c,s}$ is in a minority group in class $c = 1$, but $g_{c',s}$ is a majority group in class $c' = -1$:

$$f(\boldsymbol{x}_i;\boldsymbol{W}_t,z_t) \leq \frac{2\eta\zeta^2 t}{d}\Bigg(-\frac{\|\boldsymbol{v}_s\|^2(K-k)}{n} + \xi \pm O(d^{-\Omega(\alpha)})\Bigg) + \rho(t,\phi,\Sigma), \tag{8}$$

where ρ is constant for all examples in the same class, $\xi \sim \mathcal{N}(0,\kappa)$ with $\kappa = \frac{1}{n}(\sum_c n_c^2\sigma_c^2\|\boldsymbol{v}_c\|^2)^{1/2} + \frac{1}{n}(\sum_s(n_{c,s} - n_{c',s})^2\sigma_s^2\|\boldsymbol{v}_s\|^2)^{1/2}$ is the total effect of noise on the model.

Analogous statements holds for the class $c = -1$ by changing the sign and direction of the inequality.

The proof can be found in Sec. B.2. Corollary 4.2 shows that when the majority group is considerably larger than the minority groups ($K \gg k$), the prediction of examples in the majority group move toward their label considerably faster, due to the contribution of the spurious feature. Hence, majority and minority groups can be separated from each other, early in training. Importantly, multiple minority groups can be also separated from each other, if their spurious feature exists in majority groups of other classes. Note that $K > k + |\xi|$ is the minimum requirement for the separation to happen. Separation is more significant when $K \gg k$ and when $\|\boldsymbol{v}_s\|$ is significant.

4.2 Network Relies on Simple Spurious Features for Majority of Examples

Next, we analyze the second phase in early-time learning of a two-layer neural network. In particular, we show that if the noise-to-signal ratio of the spurious feature of the majority group of class c, i.e., $R_s = \sigma_s/\|\boldsymbol{v}_s\|$ is smaller than that of the core feature $R_c = \sigma_c/\|\boldsymbol{v}_c\|$, then the neural network's output is almost exclusively determined by the spurious feature and remain invariant to the core feature at $T = \nu_2 \cdot \frac{d\log d}{\eta}$, even though the core feature is more predictive of the class.

Theorem 4.3. *Under the assumptions of Theorem 4.1, if the classes are balanced, and the total size of the minority groups in class c is small, i.e., $O(n^{1-\gamma})$ for some $\gamma > 0$, then there exists a constant $\nu_2 > 0$ such that at $T = \nu_2 \cdot \frac{d\log d}{\eta}$, for an example $\boldsymbol{x}_i$ in a majority group $g_{c,s}$, the contribution of the core feature to the model's output is at most:*

$$|f(\boldsymbol{v}_c;\boldsymbol{W}_T,z_T)| \leq \left(\frac{\sqrt{2}R_s}{R_c} + O(n^{-\gamma} + d^{-\Omega(\alpha)})\right). \tag{9}$$

In particular if $\min\{R_c, 1\} \gg R_s$, then the model's output is mostly indicated by the spurious feature instead of the core feature:

$$|f(\boldsymbol{v}_s;\boldsymbol{W}_T,z_T)| \gg |f(\boldsymbol{v}_c;\boldsymbol{W}_T,z_T)|. \tag{10}$$

The proof can be found in Sec. B.3. Theorem 4.3 shows that at $T = \nu_2 \cdot \frac{d\log d}{\eta}$, the contribution of the core feature to the network's output is at most proportional to R_s/R_c. Hence, if $R_s \ll R_c$, the network almost exclusively relies on the spurious feature of the majority group instead of the core feature.

We note that our results in Theorem 4.1, Corollary 4.2, and Theorem 4.3 generalize to more than two classes and hold if the classes are imbalanced, as we will confirm by our experiments. Similar results can be shown for multi-layer fully connected and convolutional networks, following (Hu et al., 2020).

4.3 Separability of Majority and Minority Groups in The Two Early Training Phases

The initial phase of training iterations, when $0 \leq t \leq \nu_1 \cdot \sqrt{\frac{d^{1-\alpha}}{\eta}}$, is characterized by approximately linear change in the loss. Corollary 4.2 shows that in Phase 1, majority and minority groups are separable based on the network output, if spurious feature is strongly correlated with label and has a higher magnitude than the core. Phase 2 happens when $T = \nu_2 \cdot \frac{d\log d}{\eta}$ and marks the point where the approximate linear model converges to its optimal parameters, and the network starts learning higher-order (non-linear) functions. Theorem 4.3 shows that in Phase 2, majority and minority groups are separable based on the network output, if the noise-to-signal ratio of spurious feature (R_s) is smaller than core (R_c).

The above discussion implies that one can separate majority and minority groups in Phase 1 or Phase 2, as long as the corresponding conditions are met.

Visualization of Theoretical Results. We empirically illustrate the above results during early-time

Algorithm 1 SePArate early and REsample (Spare)

```
Input: Network $f(., \boldsymbol{W})$, data $\mathcal{D} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^n$, loss function $\mathcal{L}$, iteration numbers $T_N, T_{init}$.
Output: Model $f$ trained without bias
  Stage 1: Early Bias Identification
  for $t = 0, \cdots, T_{init}$ do
    $\boldsymbol{W}_{t+1} \leftarrow \boldsymbol{W}_t - \eta \nabla \mathcal{L}(\boldsymbol{W}_t; \mathcal{D})$
  end for
  for every class $c \in C$ with examples $V_c$ do
    Identify $\lambda$, # of clusters $k$ via Silhouette analysis
    Cluster $V_c$ into $\{V_{c,j}\}_{j=1}^k$ based on $f(\boldsymbol{x}_i; \boldsymbol{W}_t)$
    Weight every $\boldsymbol{x}_i \in V_{c,j}$ by $w_i = 1/|V_{c,j}|$, $p_i = w_i^\lambda / \sum_i w_i^\lambda$
  end for
  Stage 2: Learning without Bias
  for $t = 0, \cdots, T_N$ do
    Sample a mini-batch $\mathcal{M}_t = \{(\boldsymbol{x}_i, y_i)\}_i$ with probabilities $p_i$
    $\boldsymbol{W}_{t+1} = \boldsymbol{W}_t - \eta \nabla \mathcal{L}(\boldsymbol{W}_t; \mathcal{M}_t)$.
  end for
```

training of LeNet-5 (LeCun et al., 1998) on the Colored MNIST (Alain et al., 2015) dataset containing colored handwritten digits. Fig. 2 shows that the prediction of the network on the majority group is almost exclusively indicated by the color of the majority group, confirming Theorem 4.3. The bottom of Fig. 2 shows that the majority and minority groups are separable based on the network output, confirming Corollary 4.2.

Strong Spurious Correlations Make the Network Invariant to Core Features of Majority Groups. Finally, note that by only learning the spurious feature, the neural network can shrink the training loss on the majority of examples in class c to nearly zero and correctly classify them. Hence, the contribution of the spurious feature of the majority group of class c to the model's output is retained throughout the training. On the other hand, if minority groups are small, higher complexity functions that appear later in training overfit the minority groups, as observed by (Sagawa et al., 2020). This results in a small training error but a poor worst-group generalization performance on the minorities.

5 SPARE: ELIMINATING SPURIOUS BIAS EARLY IN TRAINING

Based on the theoretical foundations outlined in Sec. 4, we develop a principled pipeline, Spare, to discover and mitigate spurious correlations *early in training.* The pseudocode is illustrated in Alg. 1.

Discovering Spurious Correlations: Separating the Groups Early in Training. Corollary 4.2 shows that majority and minority groups are separable based on the network's output. To identify the majority and minority groups, we cluster examples V_c in every class $c \in C$ based on the output of the network, during the first few epochs. We tune T_{init} for maximum recall of Spare's clusters against the validation set groups in the first 1-2 epochs (discussed in Sec. E.2). We determine the number of clusters via silhouette analysis (Rousseeuw, 1987), a technique that assesses the cohesion and separation of clusters by evaluating how close each point in one cluster is to points in the neighboring clusters. In doing so, we can separate majority and minority groups in each class of examples with different spurious features. Any clustering algorithm such as k-means or k-median clustering (Mirzasoleiman et al., 2013, 2015) can be applied to separate groups in large data.

Mitigation after Discovery: Balancing Groups via Importance Sampling. To alleviate the spurious correlations and enable effective learning of the core features, we employ an importance sampling method on examples in each class to upsample examples in the smaller clusters and downsample examples in the larger clusters. To do so, we assign every example $i \in V_{c,j}$ a weight given by the size of the cluster it belongs to, i.e., $w_i = 1/|V_{c,j}|$. Then we sample examples in every mini-batch with probabilities equal to $p_i = w_i^\lambda / \sum_i w_i^\lambda$, where λ can be determined based on the average silhouette score of clusters in each class, *without further tuning.* Our importance sampling method does not increase the size of the training data but only changes the data distribution. Hence, it does not increase the training time.

6 EXPERIMENTS

In this section, we first confirm that Spare outperforms state-of-the-art baselines in identifying and mitigating spurious correlations across multiple curated benchmark datasets. Most notably, Spare excels on UrbanCars, a challenging dataset with multiple spurious correlations within each class. Then, we demonstrate that Spare effectively discovers and mitigates naturally occurring spurious correlations early in training on Restricted ImageNet—a realistic dataset *without known spurious correlations.*

6.1 Mitigating Curated Spurious Correlations in Benchmark Datasets

First, we evaluate the effectiveness of Spare in alleviating spurious correlations on spurious benchmarks. The reported results are averaged over three runs with different model initializations.

Table 1: Worst-group and average accuracy (%) of training with SPARE vs. state-of-the-art algorithms. SPARE achieves a superior performance much faster that existing methods. CB, GB indicate balancing classes and groups. Range for training cost encompasses all datasets, and accounts for (1) training the reference model for group inference and (2) number of training examples involved in robust training (excluding tuning cost). Baseline results for CMNIST, UrbanCars are from the benchmarks (Zhang et al., 2022; Li et al., 2023). ♦ and △ indicate using group-labeled validation for tuning group inference, and robust training. (E#) shows the early group inference epoch for SPARE. (◊) shows SPARE doesn't heavily rely on validation set (◊). We couldn't successfully run CnC on UrbanCars and DFR on CMNIST.

	Grp label required	Train cost	**CMNIST** (1 Spurious × **5 Classes**)		**Waterbirds** (1 Spurious × 2 Classes)		**CelebA** (1 Spurious × 2 Classes)		**UrbanCars** (**2 Spurious** × 2 Classes)	
			Worst-group	Average	Worst-group	Average	Worst-group	Average	Worst-group	Average
ERM	−−	1x	$0.0_{\pm0.0}$	$20.1_{\pm0.2}$	$62.6_{\pm0.3}$	$97.3_{\pm1.0}$	$47.7_{\pm2.1}$	$94.9_{\pm0.3}$	28.4	97.6
CB	−−	1x	$0.0_{\pm0.0}$	$23.7_{\pm3.1}$	$62.8_{\pm1.6}$	$97.1_{\pm0.1}$	$46.1_{\pm1.5}$	$95.2_{\pm0.4}$	33.7	98.1
GEORGE (Sohoni et al., 2020)	−−	2x	$76.4_{\pm2.3}$	$89.5_{\pm0.3}$	$76.2_{\pm2.0}$	$95.7_{\pm0.5}$	$54.9_{\pm1.9}$	$94.6_{\pm0.2}$	35.2	97.9
PGI (Ahmed et al., 2020)	♦−	1x	$73.5_{\pm1.8}$	$88.5_{\pm1.4}$	$79.5_{\pm1.9}$	$95.5_{\pm0.8}$	$85.3_{\pm0.3}$	$87.3_{\pm0.1}$	34.0	95.7
EIIL (Creager et al., 2021)	♦−	1x	$72.8_{\pm\mathbf{6.8}}$	$90.7_{\pm0.9}$	$83.5_{\pm2.8}$	$94.2_{\pm1.3}$	$81.7_{\pm0.8}$	$85.7_{\pm0.1}$	50.6	95.5
LfF (Nam et al., 2020)	♦△	2x	$0.0_{\pm0.0}$	$25.0_{\pm0.5}$	$78.0_{N/A}$	$91.2_{N/A}$	$77.2_{N/A}$	$85.1_{N/A}$	34.0	97.2
JTT (Liu et al., 2021)	♦△	5x-6x	$74.5_{\pm\mathbf{2.4}}$	$90.2_{\pm0.8}$	$83.1_{\pm\mathbf{3.5}}$	$90.6_{\pm0.3}$	$81.5_{\pm1.7}$	$88.1_{\pm0.3}$	55.8	95.9
CnC (Zhang et al., 2022)	♦△	2x-12x	$77.4_{\pm3.0}$	$90.9_{\pm0.6}$	$88.5_{\pm0.3}$	$90.9_{\pm0.1}$	$88.8_{\pm0.9}$	$89.9_{\pm0.5}$	-	-
SPARE (−2nd best)	◊−	**1x**	**(E2) 83.0**$_{\pm1.7}$ **(+5.6)**	$91.8_{\pm0.7}$	**(E2) 91.6**$_{\pm\mathbf{0.8}}$ **(+3.1)**	$96.2_{\pm0.6}$	**(E1) 90.3**$_{\pm\mathbf{0.3}}$ **(+1.5)**	$91.1_{\pm0.1}$	**(E2) 76.9**$_{\pm1.8}$ **(+21.1)**	$96.6_{\pm0.5}$
DFR (Kirichenko et al., 2023)	train sub	1x	-	-	$90.4_{\pm1.5}$	$94.1_{\pm0.5}$	$80.1_{\pm1.1}$	$89.7_{\pm0.4}$	44.5	89.7
GB	train full	1x	$82.2_{\pm1.0}$	$91.7_{\pm0.6}$	$86.3_{\pm0.3}$	$93.0_{\pm1.5}$	$85.0_{\pm1.1}$	$92.7_{\pm0.1}$	73.9	92.2
GDRO (Sagawa et al., 2019)	train full	1x	$78.5_{\pm4.5}$	$90.6_{\pm0.1}$	$89.9_{\pm0.6}$	$92.0_{\pm0.6}$	$88.9_{\pm1.3}$	$93.9_{\pm0.1}$	75.2	91.6

Benchmark Datasets & Models. (1) CMNIST (Alain et al., 2015) contains colored handwritten digits derived from MNIST (LeCun et al., 1998). We follow the challenging 5-class setting in Zhang et al. (2022) where every two digits form one class and 99.5% of training examples in each class are spuriously correlated with a distinct color. We use a 5-layer CNN (LeNet-5 (LeCun et al., 1998)) for CMNIST. (2) Waterbirds (Sagawa et al., 2019) contains two classes (landbird vs. waterbird) and the background (land or water) is the spurious feature. Majority groups are (waterbird, water) and (landbird, land). (3) CelebA (Liu et al., 2015) is a face datasets containing two classes (blond vs dark) and gender (male or female) as the spurious feature (Sagawa et al., 2019). Majority groups are (blond, female) and (non-blond, male). (4) UrbanCars (Li et al., 2023) is a challenging task containing two classes (urban vs. country cars) and *two spurious features*: (1) background (BG): urban vs. country and (2) co-occurring object (CoObj): fireplug and stop sign vs cow and horse. For Waterbirds, CelebA, and UrbanCars, we follow the standard settings in previous work to train a ResNet-50 model (He et al., 2016) pretrained on ImageNet provided by the Pytorch library (Paszke et al., 2019). More details about the datasets and the experimental settings are in Sec. E.

Baselines. We compare SPARE with the state-of-the-art methods for eliminating spurious correlations in Tab. 1, in terms of both worst-group and average accuracy. We use adjusted average accuracy for Waterbirds, i.e., the average accuracy over groups weighted by their size. This is consistent with prior work, and is done because the validation and test sets are group-balanced while the training set is skewed. For UrbanCars, our average accuracy corresponds to the ID accuracy in the original paper (Li et al., 2023), and our worst-group accuracy is computed by adding the largest of BG/CoObj/BG+CoObj gap to the ID accuracy. GB (Group Balancing) and GDRO (Sagawa et al., 2019) use the group label of all training examples. DFR (Kirichenko et al., 2023) uses group-balanced data drawn from either validation (DFR$^{Tr}_{Tr}$) or training (DFR$^{Val}_{Tr}$) data. We considered DFR$^{tr}_{tr}$, which trains on group-balanced training data, for a fair comparison with baselines that only use the validation set to tune. The rest of the methods infer the group labels without using such information.

SPARE Outperforms SOTA Algorithms, Especially When Multiple Spurious Correlations Exist. Tab. 1 shows that SPARE consistently outperforms the best baselines, on worst-group and average accuracy. On UrbanCars, where baselines have been shown to amplify one spurious when trying to mitigate the other (Li et al., 2023), SPARE achieves a 21.1% higher accuracy than the next best state-of-the-art method that does not rely on ground-truth group labels during training. Notably, SPARE performs comparably to those that *use* the group information, and even achieves a better worst-group accuracy on CMNIST, CelebA, and UrbanCars, and has a comparable worst-group but higher average accuracy on the Waterbirds. As group labels are unavailable in real-world datasets, methods that do not rely on group labels are more practical. Moreover, group labels can sometimes

Method	BG (↑)	CoObj (↑)	BG+CoObj (↑)
JTT (E1)	-8.1	-13.3	-40.1
EIIL (E1)	-4.2	-24.7	-44.9
JTT (E2)	-23.3	-5.3	-52.1
EIIL (E2)	-21.5	-6.8	-49.6
Spare (E2)	**-5.3**	**-3.1**	**-8.9**

Table 2: UrbanCars with two spurious correlations (BG and CoObj). SOTA methods show "whack-a-mole" behavior (Li et al., 2023): mitigating one spurious correlation amplifies the other, Spare finds minority groups more accurately and does not exhibit whack-a-mole.

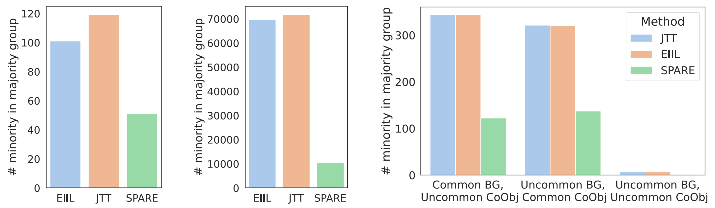

(a) Waterbirds (b) CelebA (c) UrbanCars

Figure 3: Number of minority examples inferred as majority. JTT and EIIL infer many minority examples as majority and mistakenly downweight them. Spare identifies minority groups more accurately, and correctly upweights them.

be inaccurate, so group-inference methods like Spare can better identify the groups. This aligns with EIIL's observation that inferred groups can be more informative than underlying group labels (Creager et al., 2021). Among such methods, Spare has a superior performance and easily scales to large datasets.

Spare Does Not Exhibit Whack-a-mole Behavior when Data Contains Multiple Spurious Correlations. Tab. 2 shows the accuracy of different groups in UrbanCars datasets with two spurious correlations (BG and CoObj). We see that JTT and EIIL drop the accuracy on BG or CoObj when trying to improve the accuracy on the other. This is because they find a smaller minority group, and hence need to upweight it heavily to mitigate the spurious correlation. In doing so, they amplify the other subtle spurious correlation in the minority group. In contrast, Spare finds minority groups more accurately and mitigates the spurious correlation without introducing a new one. Notably, Spare achieves 31.2% better accuracy than the best baseline on the smallest group (BG+CoObj).

Spare is Much Faster and Easier to Tune. As our theoretical results narrow the range for inference time to the initial epochs (often 1 or 2 in practice), it can be tuned easily for Spare, while others need to search over a wide range from epoch 1 to 60. In addition to the time saved for hyperparameter tuning, Spare has up to 12x lower computational cost (k-means & total wall-clock runtimes are reported in Tab. 6 and Tab. 7) compared to the state-of-the-art.

6.2 Why Does Spare Work Better?

Spare Finds Minority Groups More Accurately. Fig. 3 explains the superior performance of Spare over the state-of-the-art: JTT (Liu et al., 2021) and EIIL (Creager et al., 2021) mistakenly group many minority examples into majority groups and thus mistakenly downweight them. The minority samples include, by definition, all instances where spurious correlations are not present. In contrast, by finding the minority groups more accurately, Spare effectively upweights them to mitigate the spurious correlations and improve the worst-group accuracy. We also compare the worst-group and average accuracy of models trained with GDRO and JTT's upsampling method applied to groups inferred by Spare in Tab. 3. Comparing to Tab. 1, we see that both methods obtain a better worst-group accuracy using Spare's groups. In particular, training on Spare's groups with GDRO outperform George and EIIL that use GDRO by 11.4%, 4.1%. Similarly, training on Spare's groups with JTT's upsampling outperforms JTT by 4.5%, further confirming that Spare finds minorities more accurately.

Table 3: Spare's importance sampling is more effective in improving the worst-group accuracy than GDRO and JTT's upsampling, when applied to Spare's groups.

Groups	Robust training	Worst-group	Avg Acc
Spare	JTT	86.2 ± 3.6	92.0 ± 0.8
Spare	GDRO(/George/EIIL)	87.6 ± 0.8	89.4 ± 1.3
Spare	Spare	**91.6 ± 0.8**	**96.2 ± 0.6**

Spare's Importance Sampling is More Effective and Efficient. Next, we compare the worst-group and average accuracy of models trained with GDRO, JTT's upsampling, and Spare's importance sampling, applied to groups inferred by Spare. Tab. 3 shows that Spare's importance sampling is more effective in improving the worst-group accuracy and outperforms GDRO by 4% and JTT's upsampling by 5.4%. Note that both methods require tuning based on group-labeled validation data, and upsampling drastically increases the training time. On the other hand, Spare's importance sampling does not require any hyperparameter tuning or increase the training time.

6.3 Discovering Natural Spurious Correlations in Restricted ImageNet

Next, we show the applicability of Spare to discover and mitigate spurious correlations in Restricted Ima-

Table 4: Discovering & mitigating spurious correlations in Restricted ImageNet. Spare infers groups more accurately and improves both insect and frog minority accuracy by 1.2% and 11.5% respectively, with only a minor drop in total accuracy. Note that the group with worst-group accuracy changes during the training.

Method	Minority Recall	Test Avg Acc	Insect Min Acc	Frog Min Acc	W-G Acc
ERM	-	96.0%	91.7%	80.8%	80.8%
CB	-	95.9%	**93.7**% ↑	80.8%–	80.8%–
EIIL	78.8%	93.1%	88.3% ↓	69.2% ↓	69.2% ↓
JTT	82.6%	92.8%	75.0% ↓	**92.3**% ↑	75.0% ↓
GEORGE	85.4%	94.8%	89.4% ↓	80.8%–	80.8%–
Spare	**92.8**%	95.4%	**92.9**% ↑	**92.3**% ↑	**92.3**% ↑

geNet (Taghanaki et al., 2021), a 9-superclass subset of ImageNet. Here, we train ResNet-50 from scratch. See G for more details on the dataset and experiment.

Spare Discovers Spurious Correlations We observe Spare clusters during the initial training epochs. Inspecting the clusters with the highest fraction of misclassified examples to another class, we find that many frog images are misclassified as insects. Fig. 4a shows examples from the two groups Spare finds for the Insect class at epoch 8, where clusters with spurious features are visually evident [2]. GradCAM reveals an obvious spurious correlation between "green leaf" and the insect class that is maintained until the end of the training, as illustrated in Fig. 4b. We also observe a large gap between the confidence of examples in the two groups. This indicates that the model has learned the spurious feature early in training.

Spare Discovers Spurious Correlations without Relying on Group-labeled Validation Data. State-of-the-art group inference baselines heavily rely on a group-labeled validation set to identify the time of group inference during training with ERM. While Spare can also benefit from a group-labeled validation, this is not essential. In fact, our theoretical results reveal that the range for inference time should be within the initial epochs. Thus we only visually inspect a few initial epochs (2, 4, 6, 8) to verify the spurious correlations. This sets Spare apart as a more generally applicable method for discovering and mitigating spurious correlations, even in the absence of a group-labeled validation set.

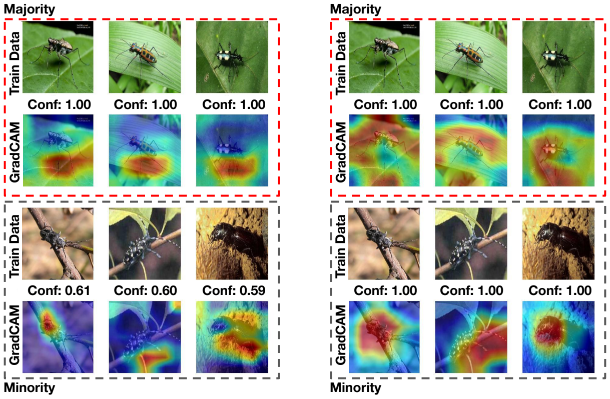

(a) ERM epoch 8: model learns the spurious feature "green leaf" in Insect class.

(b) ERM end of training: model keep relying on "green leaf", instead of Insect.

Figure 4: Spare-discovered spurious correlation between "green leaf" & "insect" in Restricted ImageNet.

Spare Achieves State-of-the-art Accuracy on Minority Groups. Based on the spurious correlations Spare discovered, we manually labeled the background of both training and test data for the *insect* and *frog* classes, and these group labels to tune the baseline group inference methods. Tab. 4 shows Spare separates the insect majority group with the spurious correlation better than other group inference methods and improves both insect and frog minority accuracy by 1.2% and 11.5% respectively, with only a minor drop in total accuracy. Note that group with worst-group accuracy changes during the training. CB only improves insect minority accuracy. JTT decreases the model's accuracy on the insect minority a lot while improving the frog minority. EIIL decreases both minority and total accuracy as it finds the least majority. Unlike the baselines, Spare effectively balances groups, mitigating spurious correlations.

7 CONCLUSION

We studied the effect of simplicity bias of SGD on learning spurious features. Specifically, we analyzed a two-layer fully-connected neural network and showed that spurious features can be identified early in training based on model output. If spurious features have a low noise-to-signal ratio, they dominate the network's output, overshadowing core features. Based on the above theoretical insights, we proposed Spare, which separates majority and minority groups by clustering the model output early in training. Then, it applies importance sampling based on the cluster sizes to make the groups relatively balanced. It outperforms state-of-the-art methods in worst-group accuracy on benchmark datasets and scales well to large-scale applications. Importantly, it can deal with multiple spurious correlations, minimizes the need for hyperparameter tuning, and can discover spurious correlations in realistic scenarios like Restricted ImageNet, early in training.

[2] Since the model is not pretrained, it is expected that the spurious clusters form slightly later. For pretrained models, spurious clusters form very early, as shown in Tab. 1

Acknowledgements

This research was partially supported by the National Science Foundation CAREER Award 2146492, Amazon PhD fellowship, and Google Cloud research credits.

References

F. Ahmed, Y. Bengio, H. van Seijen, and A. Courville. Systematic generalisation with group invariant predictions. In *International Conference on Learning Representations*, 2020.

G. Alain, A. Lamb, C. Sankar, A. Courville, and Y. Bengio. Variance reduction in sgd by distributed importance sampling. *arXiv preprint arXiv:1511.06481*, 2015.

M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862*, 2017.

A. Badanidiyuru, B. Mirzasoleiman, A. Karbasi, and A. Krause. Streaming submodular maximization: Massive data summarization on the fly. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 671–680, 2014.

H. Cramér. *Mathematical methods of statistics*, volume 43. Princeton university press, 1999.

E. Creager, J.-H. Jacobsen, and R. Zemel. Environment inference for invariant learning. In *International Conference on Machine Learning*, pages 2189–2200. PMLR, 2021.

P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.

A. Gupta, P. Dollar, and R. Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

K. Hermann and A. Lampinen. What shapes feature representations? exploring datasets, architectures, and training. *Advances in Neural Information Processing Systems*, 33:9995–10006, 2020.

W. Hu, L. Xiao, B. Adlam, and J. Pennington. The surprising simplicity of the early-time learning dynamics of neural networks. *Advances in Neural Information Processing Systems*, 33:17116–17128, 2020.

P. Kirichenko, P. Izmailov, and A. G. Wilson. Last layer re-training is sufficient for robustness to spurious correlations. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=Zb6c8A-Fghk.

J. Krause, M. Stark, J. Deng, and L. Fei-Fei. 3d object representations for fine-grained categorization. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 554–561, 2013. doi: 10.1109/ICCVW.2013.77.

Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Z. Li, I. Evtimov, A. Gordo, C. Hazirbas, T. Hassner, C. C. Ferrer, C. Xu, and M. Ibrahim. A whac-a-mole dilemma: Shortcuts come in multiples where mitigating one amplifies others. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20071–20082, 2023.

E. Z. Liu, B. Haghgoo, A. S. Chen, A. Raghunathan, P. W. Koh, S. Sagawa, P. Liang, and C. Finn. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning*, pages 6781–6792. PMLR, 2021.

Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause. Distributed submodular maximization: Identifying representative elements in massive data. In *Advances in Neural Information Processing Systems*, pages 2049–2057, 2013.

B. Mirzasoleiman, A. Badanidiyuru, A. Karbasi, J. Vondrák, and A. Krause. Lazier than lazy greedy. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

V. Nagarajan, A. Andreassen, and B. Neyshabur. Understanding the failure modes of out-of-distribution generalization. *arXiv preprint arXiv:2010.15775*, 2020.

P. Nakkiran, G. Kaplun, D. Kalimeris, T. Yang, B. L. Edelman, F. Zhang, and B. Barak. Sgd on neural networks learns functions of increasing complexity. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 3496–3506, 2019.

J. Nam, H. Cha, S. Ahn, J. Lee, and J. Shin. Learning from failure: De-biasing classifier from biased classifier. *Advances in Neural Information Processing Systems*, 33:20673–20684, 2020.

J. Nam, J. Kim, J. Lee, and J. Shin. Spread spurious attribute: Improving worst-group accuracy with spurious attribute estimation. In *International Conference on Learning Representations*, 2021.

B. Neyshabur, R. Tomioka, and N. Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.

A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. 2019.

M. Pezeshki, O. Kaba, Y. Bengio, A. C. Courville, D. Precup, and G. Lajoie. Gradient starvation: A learning proclivity in neural networks. *Advances in Neural Information Processing Systems*, 34:1256–1272, 2021.

P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20: 53–65, 1987.

S. Sagawa, P. W. Koh, T. B. Hashimoto, and P. Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2019.

S. Sagawa, A. Raghunathan, P. W. Koh, and P. Liang. An investigation of why overparameterization exacerbates spurious correlations. In *International Conference on Machine Learning*, pages 8346–8356. PMLR, 2020.

M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.

R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.

H. Shah, K. Tamuly, A. Raghunathan, P. Jain, and P. Netrapalli. The pitfalls of simplicity bias in neural networks. *Advances in Neural Information Processing Systems*, 33:9573–9585, 2020.

N. Sohoni, J. Dunnmon, G. Angus, A. Gu, and C. Ré. No subclass left behind: Fine-grained robustness in coarse-grained classification problems. *Advances in Neural Information Processing Systems*, 33:19339–19352, 2020.

S. A. Taghanaki, K. Choi, A. H. Khasahmadi, and A. Goyal. Robust representation learning via perceptual similarity metrics. In *International Conference on Machine Learning*, pages 10043–10053. PMLR, 2021.

D. Teney, E. Abbasnejad, S. Lucey, and A. Van den Hengel. Evading the simplicity bias: Training a diverse set of models discovers solutions with superior ood generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16761–16772, 2022.

D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=SyxAb30cY7.

C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.

L. A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.

M. Zhang, N. S. Sohoni, H. R. Zhang, C. Finn, and C. Ré. Correct-n-contrast: A contrastive approach for improving robustness to spurious correlations. *arXiv preprint arXiv:2203.01517*, 2022.

B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2018. doi: 10.1109/TPAMI.2017.2723009.

Checklist

1. For all models and algorithms presented, check if you include:
   (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
   (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
   (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
   (a) Statements of the full set of assumptions of all theoretical results. [Yes]
   (b) Complete proofs of all theoretical results. [Yes]
   (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
   (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
   (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
   (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
   (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
   (a) Citations of the creator If your work uses existing assets. [Yes]
   (b) The license information of the assets, if applicable. [Yes]
   (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
   (d) Information about consent from data providers/curators. [Not Applicable]
   (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
   (a) The full text of instructions given to participants and screenshots. [Not Applicable]
   (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
   (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]