

Project-Probe-Aggregate: Efficient Fine-Tuning for Group Robustness

Beier Zhu^{1,2} Jiequan Cui¹ Hanwang Zhang¹ Chi Zhang^{2†}

¹Nanyang Technological University, ²Westlake University

{beier.zhu, jiequan.cui, hanwangzhang}@ntu.edu.sg, chizhang@westlake.edu.cn

Abstract

While image-text foundation models have succeeded across diverse downstream tasks, they still face challenges in the presence of spurious correlations between the input and label. To address this issue, we propose a simple three-step approach—Project-Probe-Aggregate (PPA)—that enables parameter-efficient fine-tuning for foundation models without relying on group annotations. Building upon the failure-based debiasing scheme, our method, PPA, improves its two key components: minority samples identification and the robust training algorithm. Specifically, we first train biased classifiers by projecting image features onto the nullspace of class proxies from text encoders. Next, we infer group labels using the biased classifier and probe group targets with prior correction. Finally, we aggregate group weights of each class to produce the debiased classifier. Our theoretical analysis shows that our PPA enhances minority group identification and is Bayes optimal for minimizing the balanced group error, mitigating spurious correlations. Extensive experimental results confirm the effectiveness of our PPA: it outperforms the state-of-the-art by an average worst-group accuracy while requiring less than 0.01% tunable parameters without training group labels.

1. Introduction

Image-text foundation models [12, 31, 40]—large models pre-trained on web-scale data—are becoming increasingly prevalent in real-world deployments. However, adapting these models to downstream tasks, whether through zero-shot transfer or fine-tuning, remains challenging due to fundamental issues of group robustness: they achieve low average test error but incur high risk on certain groups of examples [3, 44, 47, 48]. This issue becomes particularly pronounced in the presence of spurious correlations [8], where a classifier relies on non-causal relationships that happen to align with class labels in training but break down during deployment. For example, consider the case where “water-

Method	Worst-group Recall (%)	Worst-group Precision (%)
CA [48]	46.4	15.6
JTT [21]	78.6	24.1
PPA (ours)	80.4	44.3

Table 1. The precision and recall of the worst-group examples (*i.e.*, the group with lowest validation accuracy) identified by biased models on the Waterbird training set [32] using CLIP ResNet-50. Our PPA is more susceptible to spurious features, achieving higher worst-group recall and precision compared to previous methods.

birds” predominantly appear in images with a “water” background. When trained on such a dataset, neural networks often rely on the background to recognize the objects.

As a result, several fine-tuning methods have emerged to improve group robustness in foundation models [3, 29, 44, 47, 48]. Unlike conventional robust training frameworks, which typically re-train the entire model [21, 22, 26, 27] (which is computationally expensive) and/or require group labels [1, 32] (which demands additional annotation costs), these fine-tuning methods prioritize parameter and annotation efficiency. They address spurious correlations by optimizing only a small subset of parameters without relying on training group labels. For instance, Zhang and Ré [48] fine-tune a lightweight adapter and You *et al.* [47] propose to optimize projection layers to improve distributional robustness. In this paper, we build on these ideas by focusing on efficient fine-tuning to enhance group robustness, *i.e.*, training linear probes.

At a high level, in the absence of group annotations, a common strategy [5, 19, 21, 22, 26, 48, 49] involves first identifying minority groups by training a *biased model* that easily overfits to spurious features. Subsequently, a *debiased model* is trained using the inferred group labels. To improve the identification of minority groups, prior work proposes various strategies: [22, 26] employ generalized cross-entropy loss to amplify the model bias, JTT [21] simply trains a model with empirical risk minimization (ERM) and tunes hyperparameters to exacerbate bias, CA [48] leverages the correctness of zero-shot pre-

This work was done during Beier Zhu’s visit at AGI Lab, Westlake University. † denotes corresponding author.

dictionaries and [5] uses low-capacity networks to capture spurious features. In this paper, we leverage the pre-trained knowledge from the image-text foundation models to effectively induce bias, further enhancing the identification of minority groups. Specifically, we **project** out class proxies, provided by the text encoders, from the input image features and use only the remaining information to train the biased classifiers. Intuitively, removing class information forces the model to rely more heavily on spurious features for predictions. We provide theoretical analysis to support this intuition in Proposition 1. Empirically, in Tab. 1, we evaluate the quality of the inferred groups in the training datasets obtained by our method, and compare it to prior approaches [21, 48]. Specifically, we measure precision—the fraction of examples within the identified minor groups that belong to the worst-performing group, and recall—the fraction of examples from the worst-performing group that are captured by the identified minor groups. As expected, our method achieves higher precision and recall, demonstrating superior group identification.

Given the inferred groups, typical debiased learning algorithms up-weight/up-sample minority group samples to enhance performance on these challenging instances [5, 10, 17, 21, 26]. However, determining the optimal weights requires careful hyper-parameter tuning and lacks theoretical guidance. Since the goal of debiased training is to minimize the Balanced Group Error (BGE) [22, 36], we propose a methodology designed to achieve this objective directly. Rather than predicting class labels, our method is trained to **probe** for group targets by compensating group priors into the softmax cross-entropy during training. Afterward, we **aggregate** group weights of each class to produce the final debiased classifier. We dub our overall debiasing framework Project-Probe-Aggregate (PPA), and demonstrate that it is *Bayes* optimal for minimizing BGE in Proposition 2.

In summary, our contributions are:

- We project out the class proxies offered by the text encoder and train the biased model on the residual image features, with theoretical analysis showing this approach increases reliance on spurious features, improving minority group identification.
- We propose a novel debiasing method that is consistent for minimizing the balanced group error. The debiased classifier predicts group labels compensated by group priors and uses weight-space aggregation after fine-tuning to avoid group inference overhead.
- We confirm the effectiveness of our proposal across five benchmarks with spurious correlations. With less than 0.01% of the foundation model parameters, our approach PPA improves the worst group accuracy over the state-of-the-art without training group annotations. Beyond linear probes, we also adapt other parameter-efficient fine-tuning paradigms (e.g., prompt tuning and adapters) to

show the versatility of the proposed method.

2. Related Work

Improving group robustness. There have been numerous works aimed at improving group robustness, which broadly fall into two categories based on whether group labels are available during training. (i) *Group supervised methods* focus on balancing the contribution of each group during the training phase, such as through re-sampling/re-weighting [10, 11, 13, 17, 24, 39], or robust optimization [1, 32, 34]. However, this approach requires additional annotation costs for group labels. In this paper, we focus on (ii) *group unsupervised methods*, which do not assume the availability of group labels during training [2, 14, 15, 18, 35, 42, 43, 46]. Without group labels, a common strategy is to first train an easily biased model, and then infer groups based on its biased predictions. A robust model is subsequently trained using the inferred group labels through techniques such as importance weighting [21, 26, 30], robust optimization [4, 27], or by learning similar representations for groups within the same class [49]. In this paper, we leverage the general knowledge from image-text foundation models to train a biased classifier by excluding the class proxy information generated by the CLIP text encoder. Note that classifier in [39] also predicts group labels, our approach differs in *two key aspects*: 1) We further incorporate logit offsets during training, which compensates for group imbalances and leads to lower balanced error. 2) Unlike [39], our method *does not* require access to ground-truth group labels.

Efficient robust fine-tuning for foundation models. The advent of image-text foundation models has led to several efforts to develop efficient group-robust fine-tuning methods [3, 6, 16, 17, 28, 44, 47, 48, 50, 55]. Kirichenko *et al.* [17] demonstrate that last-layer retraining significantly improves group robustness in pretrained models. Zhang and Ré [48] use contrastive objectives to train adapters, bringing same-class embeddings closer together. Chuang *et al.* [3] reduce bias in image-text models by removing biased directions in text embeddings. Other works [44, 47] focus on lightweight fine-tuning by optimizing only the projection layers of the vision branch for distributional robustness. Our approach builds on efficient fine-tuning by training only a linear classification head for group targets and further introduces weight-space ensembling after fine-tuning to eliminate the overhead of group inference.

3. Preliminaries

3.1. Setup

Considering a classification problem with instances $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d$ and labels $y \in \mathcal{Y} = [K]$. Each data point (\mathbf{x}, y) has an input attribute $a(\mathbf{x}) \in \mathcal{A}$ that is spuriously correlated

with the target y . We form $|\mathcal{G}| = |\mathcal{A}| \times |\mathcal{Y}|$ groups, where each group label $g \in \mathcal{G}$ is the combination of the attribute a and the class label y , i.e., $g := (a, y)$. Given a training dataset $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$ drawn from some unknown distribution \mathbb{P} , and let \mathbb{P}_g be the distribution conditioned on g for any $g \in \mathcal{G}$. We assume that spurious attribute annotations are *not* available for the samples in the training set. Our goal is to train a model $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^K$ that outputs prediction score, such that it achieves low expected error:

$$\mathcal{R} = \mathbb{E}_{\mathbf{x}, y \sim \mathbb{P}}[y \neq \operatorname{argmax}_{y' \in \mathcal{Y}} f_\theta(\mathbf{x})_{y'}] \quad (1)$$

and, more importantly, exhibits *group robust*, i.e., achieving low worst-group error:

$$\mathcal{R}_{\text{wg}} = \max_{g \in \mathcal{G}} \mathbb{E}_{\mathbf{x}, y \sim \mathbb{P}_g}[y \neq \operatorname{argmax}_{y' \in \mathcal{Y}} f_\theta(\mathbf{x})_{y'}] \quad (2)$$

For example, in the Waterbirds [32] dataset, the task is to classify birds $y \in \{\text{waterbird}, \text{landbird}\}$, with the background serving as a spurious attribute ($a \in \{\text{water}, \text{land}\}$). A notable pattern is that about 95% of waterbird images have a water background, causing models to rely on the “water” background to predict “waterbird”. This reliance leads to poor performance on minority groups, such as landbirds with water backgrounds $g = (\text{water}, \text{landbird})$. Unlike domain generalization or out-of-distribution evaluation, all groups are present in the training, validation, and test splits. However, the training groups are imbalanced, which can result in poor group robustness on the test set.

3.2. Common Strategy for Improving Unsupervised Group Robustness

At a high level, typical methods without access to group labels first train a *biased model* and then use its biased predictions to infer the groups. Subsequently, a *robust model* is trained using the inferred group labels.¹

Taking JTT [21] as an example, it identifies the misclassified samples from ERM model f_{erm} as the error set \mathcal{E} :

$$\mathcal{E} = \{(\mathbf{x}_i, y_i) \text{ s.t. } f_{\text{erm}}(\mathbf{x}_i) \neq y_i\}, \quad (3)$$

then trains a robust model by upweighting the samples in \mathcal{E} :

$$\mathcal{L}_{\text{JTT}}(\theta, \mathcal{E}) = \lambda \sum_{(\mathbf{x}, y) \in \mathcal{E}} \ell(\mathbf{x}, y; \theta) + \sum_{(\mathbf{x}, y) \notin \mathcal{E}} \ell(\mathbf{x}, y; \theta), \quad (4)$$

where $\lambda \in \mathbb{R}_+$ is a hyperparameter. The core idea behind JTT is that ERM models tend to fit groups with easily learnable spurious correlations. The samples in \mathcal{E} are primarily from minority groups where these correlations fail. Upweighting such challenging samples improves the model’s performance on the worst-performing groups.

¹It is also known as the failure-based debiasing scheme. [26]

Remark. Improving group robustness without access to group labels hinges on identifying minority groups by training a biased model that overfits on spurious features. Prior work amplifies this bias through various strategies: LfF [26] employs generalized cross-entropy (GCE) loss to bias the first-stage model, JTT [21] tunes hyperparameters to induce bias, and [5] uses low-capacity networks to capture spurious features. In this paper, we enhance bias by projecting out the class proxies from the input features, using only the remaining information for classification. In Sec. 5.1, we prove that this operation increases susceptibility to spurious features, aiding the identification of minority groups.

Another key factor is the robust learning algorithm given the pseudo groups. Prior methods address this through importance weighting [21, 26], robust optimization [4, 27], or aligning representations within the same class [49]. Our approach, PPA, predicts pseudo group labels and applies a group prior offset to correct for imbalance. In Sec. 5.2, we prove that our training algorithm effectively mitigates spurious correlations.

4. Methods

Tuning entire pre-trained models is both time-consuming and computationally costly. Recent work [17] shows that simple last-layer retraining performs well on spurious correlation benchmarks. In our approach, the input \mathbf{x} represents image features encoded by CLIP. We freeze the backbone and learn linear classifiers for efficient fine-tuning.

We now present PPA, a simple three-step approach following the common strategy in Sec. 3.2. First, we train a biased classifier by removing the class proxies generated by the CLIP text encoder. Next, we group samples by classification correctness and train a robust model to predict group labels using our proposed group logit adjustment. Finally, we aggregate group weights to produce the final classifier.

Step 1: Project out class proxies. We compute the matrix $Z = [\mathbf{z}_1, \dots, \mathbf{z}_K]^\top \in \mathbb{R}^{K \times d}$ whose rows are the text embeddings of the K class names. Specifically, each \mathbf{z}_j is derived from a prompt like “a photo of a [CLASS].” with the class token [CLASS] replaced by the j -th class name. Let $\Pi \in \mathbb{R}^{d \times d}$ be the projection operator onto the null space of Z , given by:

$$\Pi = I - Z^\top(ZZ^\top)^{-1}Z \quad (5)$$

With Π , the biased model $f_b : \mathcal{X} \rightarrow \mathbb{R}^K$ is formulated as a linear classifier parameterized by a matrix $W_b \in \mathbb{R}^{K \times d}$:

$$f_b(\mathbf{x}) = W_b \Pi \mathbf{x} \quad (6)$$

To account for potential class imbalance², we apply the logit-adjustment loss [25, 54] to ensure f_b focuses on spu-

²Datasets like Waterbirds [32] and CelebA [23] are class imbalanced.

rious features rather than class priors:

$$\ell_{\text{la}}(y, f_b(\mathbf{x})) = -\ln \frac{\exp(f_b(\mathbf{x}) + \ln \pi_y)}{\sum_{y' \in \mathcal{Y}} \exp(f_b(\mathbf{x}) + \ln \pi_{y'})}, \quad (7)$$

where π_y is the prior of class y . In Proposition 1, we show that our biased model is more easily influenced by spurious correlation.

Step 2: Probe with group target. We follow the heuristic of JTT [21] to identify minority groups that f_b misclassifies:

$$\hat{a}(\mathbf{x}) = \mathbb{1}[y \neq \operatorname{argmax}_{y' \in \mathcal{Y}} f_b(\mathbf{x})_{y'}], \quad (8)$$

where $\mathbb{1}[\cdot]$ is the indicator function. Then, each training sample is augmented as (\mathbf{x}, y, \hat{a}) , with the pseudo group defined as the combination of y and \hat{a} : $\hat{g} := (y, \hat{a}) \in \hat{\mathcal{G}}$.

Our approach PPA uses a group scorer, implemented via linear probing, to predict pseudo group labels. Specifically, the group score $h_d(\mathbf{x}) = W_d \mathbf{x} : \mathcal{X} \rightarrow \mathbb{R}^{|\hat{\mathcal{G}}|}$ is parameterized by $W_d \in \mathbb{R}^{|\hat{\mathcal{G}}| \times d}$. Let $\hat{\beta}$ denote the group priors, we propose the group logit adjustment loss to achieve balanced group performance.:

$$\ell_{\text{gla}}(\hat{g}, h_d(\mathbf{x})) = -\ln \frac{\exp(h_d(\mathbf{x}) + \tau \cdot \ln \hat{\beta}_{\hat{g}})}{\sum_{g' \in \hat{\mathcal{G}}} \exp(h_d(\mathbf{x}) + \tau \cdot \ln \hat{\beta}_{g'})}, \quad (9)$$

where $\tau \in \mathbb{R}_+$ is a hyper-parameter.

Step 3: Aggregate weights. The final debiased classifier $f_d(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}^K$ is constructed by aggregating the weights corresponding to each class:

$$f_d(\mathbf{x})_y = \mathbf{w}_y^\top \mathbf{x}, \text{ where } \mathbf{w}_y^\top = \sum_{g \in \hat{\mathcal{G}}(y)} W_{d,g}. \quad (10)$$

Here, $W_{d,g}$ represents the weight vector in W_d associated with group g . This weight-space aggregation transforms group classification into class prediction, eliminating the need for group inference overhead.

In Proposition 2, we prove that the debiased model $f_d(\mathbf{x})$ is the *Bayes* optimal classifier for minimizing the balanced group error. Beyond linear classifiers, we also employ other parameter-efficient fine-tuning paradigms (e.g., prompt tuning and adapters) to implement our debiased classifier $f_d(\mathbf{x})$, as detailed in Sec. 6.5. The full pipeline is outlined in Algorithm 1.

5. Theoretical Analysis

5.1. Removal of Class Proxies Amplifies Model Bias

Intuitively, removing class information forces the model f_b to rely more on spurious features for predicting class labels. To support this intuition, we provide theoretical evidence in a linear regression setup. Let \mathbf{c} denote the core features

Algorithm 1 Pipeline of PPA

- 1: **Given:** Training dataset $\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$, class proxies $\{\mathbf{z}_j\}_{j=1}^K$.
 - Step 1: Project**
 - 2: Compute the projection matrix Π using Eq. (5).
 - 3: Build $f_b(\mathbf{x}) = W_b \Pi \mathbf{x}$ and optimize W_b using Eq. (7).
 - Step 2: Probe**
 - 4: Identify pseudo groups $\{\hat{g}\}$ via Eq. (8).
 - 5: Build $h_d(\mathbf{x}) = W_d \mathbf{x}$ and optimize W_d using Eq. (9).
 - Step 3: Aggregate**
 - 6: Construct $f_d(\mathbf{x})$ by aggregating weights of W_d , as in Eq. (10).
 - 7: **Return** $f_d(\mathbf{x})$.
-

which are stable for predicting the target y and s be a spurious feature. The spurious feature is correlated with the target in the training set, but this correlation may not hold during testing. Suppose we observe n features, stacked as $C = [\mathbf{c}_1, \dots, \mathbf{c}_n]^\top \in \mathbb{R}^{N \times d}$ and $\mathbf{s} = [s_1, \dots, s_n]^\top \in \mathbb{R}^N$. We are interested in the contribution of the spurious feature s across the following two models:

- **Full model:** Linear regression on the core features \mathbf{c} and the spurious feature s :

$$\mathbf{y} = C\boldsymbol{\beta} + \gamma\mathbf{s} + \boldsymbol{\varepsilon}, \quad (11)$$

where $\boldsymbol{\beta}$ and γ are the weights associated with the core features and the spurious feature, respectively. $\boldsymbol{\varepsilon}$ is a noise term with an expected value of 0.

- **Projected model:** Applying the projection matrix Π to C to obtain $\tilde{C} = C\Pi$, followed by linear regression on $\tilde{\mathbf{c}}$ and s :

$$\mathbf{y} = \tilde{C}\tilde{\boldsymbol{\beta}} + \gamma'\mathbf{s} + \boldsymbol{\varepsilon}', \quad (12)$$

where $\tilde{\boldsymbol{\beta}}$ and γ' are the weights for the projected core features and the spurious feature, respectively.

In Eq. (11), the core features C can be decomposed into the remaining part \tilde{C} and the projected-out part C_o .

$$C = C\Pi + C(I - \Pi) = \tilde{C} + C_o. \quad (13)$$

Let $\mathbf{y}_o = C_o\boldsymbol{\beta}$ denote the contribution of the projected-out core features in the full model. Define $M = I - \tilde{C}(\tilde{C}^\top \tilde{C})^{-1} \tilde{C}^\top$, $\mathbf{r}_{\mathbf{y}_o} = M\mathbf{y}_o$ and $\mathbf{r}_s = M\mathbf{s}$. The following proposition states that projecting out core features can make the model more susceptible to spurious feature (proof in Appendix A.1).

Proposition 1. *The weight of the spurious feature after projection is*

$$\gamma' = \gamma + \frac{\mathbf{r}_s^\top \mathbf{r}_{\mathbf{y}_o}}{\mathbf{r}_s^\top \mathbf{r}_s}. \quad (14)$$

Remark. Since the denominator $\mathbf{r}_s^\top \mathbf{r}_s$ is non-negative, the model relies more on spurious feature, i.e., $\gamma' > \gamma$, if

$\mathbf{r}_s^\top \mathbf{r}_{y_o} > 0$. In other words, if the spurious feature s is positively correlated with the contribution of projected-out core features \mathbf{c}_o in the space of M , the weight of the spurious feature will increase. For instance, consider the Waterbirds dataset [32], the spurious feature *water* background is positively correlated with the core feature *waterbird* (about 95% of *waterbird* images have a *water* background). If we remove the contribution of *waterbird* by projection, the weight of the spurious feature *water* will increase.

5.2. Group Classification and Aggregation Mitigate Spurious Correlation

The goal of training a debiased model to avoid spurious correlation is to minimize the Balanced Group Error (BGE) rates [22]:

$$\text{BGE}(f) = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \mathbb{E}_{\mathbf{x}|g} [y \neq \arg\max_{y' \in \mathcal{Y}} f(\mathbf{x})_{y'}] \quad (15)$$

A natural question is: what is the *Bayes* optimal classifier for this problem, i.e., $f^* \in \arg \min_f \text{BGE}(f)$. Suppose the underlying group-probabilities $\mathbb{P}(g|\mathbf{x}) \propto \exp(h^*(\mathbf{x})_g)$ for unknown scorer: $h^* : \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{G}|}$. The *Bayes* optimal classifier f^* is derived from the following Proposition (proof in Appendix A.2):

Proposition 2. Let $\mathcal{G}(y)$ denote the set of groups with class label y , i.e., $\mathcal{G}(y) := \{g = (y', a) \in \mathcal{G} | y' = y\}$. Let β denote the group priors, i.e., $\beta_g = \mathbb{P}(g)$. The prediction :

$$\arg \max_{y \in \mathcal{Y}} f^*(\mathbf{x})_y = \arg \max_{y \in \mathcal{Y}} \sum_{g \in \mathcal{G}(y)} (h(\mathbf{x}) - \ln \beta)_g \quad (16)$$

is *Bayes optimal for the problem in Eq. (15)*.

Remark. We translate the unknown distributional class scores based on group logits $h(\mathbf{x})$ and group priors β . To obtain f^* , we can learn the ERM solution for group logits $h(\mathbf{x})$. Then, we can minus the logarithm of the group prior $\ln \beta$ and aggregate the logits that belongs to $\mathcal{G}(y)$ to yield the prediction for each class y . In practice, instead predict Eq. (47), we introduce a hyper-parameter $\tau \in \mathbb{R}_+$:

$$\arg \max_{y \in \mathcal{Y}} f^*(\mathbf{x})_y = \arg \max_{y \in \mathcal{Y}} \sum_{g \in \mathcal{G}(y)} (h(\mathbf{x}) - \tau \cdot \ln \beta)_g \quad (17)$$

Inspired by Menon *et al.* [25], we incorporate logit adjustment into the softmax cross-entropy ℓ_{ce} . To do so, we directly enforce group priors offset while learning the group classifier $h_d(\mathbf{x})$. We refer to this as the group logit adjustment loss:

$$\ell_{gl_a}(\hat{g}, h_d(\mathbf{x})) = \ell_{ce}(\hat{g}, h_d(\mathbf{x}) + \tau \cdot \ln \hat{\beta}), \quad (18)$$

which is equivalent to Eq. (9). Now, the estimated debiased classifier $f_d(\mathbf{x})$ becomes:

$$f_d(\mathbf{x})_y = \sum_{g \in \hat{\mathcal{G}}(y)} h_d(\mathbf{x})_g = \sum_{g \in \hat{\mathcal{G}}(y)} W_{d,g} \mathbf{x} \quad (19)$$

Since h_d is linear, summing over the output-space is equivalent to aggregating in weight-space, eliminating the overhead of group inference.

6. Experiments

In this section, we evaluate the effectiveness of our method across five benchmarks: Waterbirds [32], CelebA [23], MetaShift [20], BAR [26], and Living-17 [33]. For brevity, we provide an overview of our experimental setup here; further details can be found in Appendix B.

6.1. Datasets

Waterbirds [32] was created by overlaying bird images from the Caltech-UCSD Birds [38] dataset onto background scenes from the Places [51] dataset. The target classes are bird types ($\mathcal{Y} = \{\text{waterbird}, \text{landbird}\}$), while the spurious attribute is the scene type ($\mathcal{A} = \{\text{water}, \text{land}\}$). Approximately 95% of the waterbird training samples are associated with a water background.

CelebA [23] is a real-world dataset comprising 200K celebrity portrait images. The objective is to classify hair color ($\mathcal{Y} = \{\text{not blond}, \text{blond}\}$), with gender ($\mathcal{A} = \{\text{male}, \text{female}\}$) as a spurious attribute. Notably, over 94% of the blond-haired training images feature women.

MetaShift [20]: The classification objective of MetaShift is to distinguish between cats and dogs, with background type ($\mathcal{A} = \{\text{indoor}, \text{outdoor}\}$) as a spurious attribute. Following [45], we use a pre-processed version of MetaShift. In the training set, an inherent bias emerges: cats predominantly appear in indoor scenes, while dogs are more frequently seen in outdoor settings.

BAR [26] is a real-world dataset designed to classify six actions that are each biased toward specific locations. In the original training set, the six action-location pairs are {climbing, rock wall}, {diving, underwater}, {fishing, water surface}, {racing, paved track}, {throwing, playing Field}, and {vaulting, sky}. However, the test set features different locations than the training set. We adapted the dataset for a group robustness setting by incorporating 5% of the test set images into the training set.

Living-17 [33]: Each class in the Living-17 represents a broad animal category that encompasses several fine-grained groups. The groups in the training and testing sets can exhibit noticeable visual differences, e.g., the bear class includes images of sloth bear and ice bear. Following [48], 5% of the images in each testing group are added to the training groups.

6.2. Baselines

We compare our PPA against 14 baselines: (1) Zero-shot prompting [31], which matches image features with classification weights by extending class names into pre-

Table 2. Evaluation of methods for improving group robustness of CLIP models across the Waterbirds, CelebA, and MetaShift benchmarks. Best worst-group accuracy (WGA) of the methods without group labels are in **bold**.

Group labels in train sets?	Method	CLIP ResNet-50						CLIP ViT-L/14					
		Waterbirds		CelebA		MetaShift		Waterbirds		CelebA		MetaShift	
		WGA	Avg	WGA	Avg	WGA	Avg	WGA	Avg	WGA	Avg	WGA	Avg
✓	GroupDRO [32]	75.1	83.8	84.1	89.5	83.2	87.3	90.8	96.4	88.3	91.2	93.9	97.4
	S-CS [44]	77.5	83.2	75.2	80.4	81.2	89.8	89.1	95.7	86.1	89.3	92.3	97.1
	S-CL [44]	75.2	86.0	75.6	80.4	81.5	88.8	89.9	96.0	87.8	90.5	93.1	96.9
	DFR [17]	73.2	83.8	80.0	92.8	83.1	88.3	89.7	97.8	85.6	90.8	92.3	97.0
✗	Zero-Shot (ZS) [31]	54.2	92.4	55.0	88.0	86.2	95.4	26.5	88.2	27.0	85.9	93.2	96.2
	Group Prompt ZS [31]	46.4	91.7	53.4	73.5	84.6	95.2	25.4	85.8	66.9	83.1	93.9	96.7
	ERM [37]	7.9	93.5	11.9	94.7	75.4	94.4	65.9	97.6	28.3	94.7	84.6	96.7
	WiSE-FT [41]	49.8	91.0	85.6	88.6	86.2	95.4	65.9	97.6	80.0	87.4	93.9	97.2
	Orth-Cali [3]	74.0	78.7	82.2	84.4	86.2	94.8	68.8	84.5	76.1	86.2	92.7	96.2
	AFR [30]	48.4	89.3	53.4	94.3	76.9	86.8	73.4	88.2	70.0	85.2	90.3	97.1
	JTT [21]	61.7	90.6	60.2	79.9	78.5	89.4	83.6	97.3	75.6	93.3	91.2	94.2
	CnC [49]	61.2	87.1	63.9	90.3	78.3	87.1	84.5	97.5	79.2	89.3	92.2	94.7
	CA [48]	83.7	89.4	90.0	90.7	77.9	85.5	86.9	96.2	84.6	90.4	91.3	93.4
	CFR [47]	76.9	77.6	73.7	81.1	81.5	89.5	88.2	96.8	84.8	87.8	93.7	95.5
	PPA (ours)	84.3	88.3	91.1	92.1	90.8	94.7	87.2	94.6	90.4	91.0	94.8	96.8

defined prompts, *e.g.*, “a photo of a *waterbird*”. (2) Group-informed prompting [31], which incorporates group descriptions into the prompts, *e.g.*, “a photo of a *waterbird* on *land*”. (3) ERM [37], a linear classifier trained using cross-entropy loss, (4) WiSE-FT [41], which interpolates zero-shot and ERM models in weight-space, (5) Orth-Cali [3], a debiased zero-shot model that removes biased direction in the text embedding, (6) AFR [30], retrains the last layer of an ERM model with a weighted loss that emphasizes examples where the ERM model performs poorly, (7) JTT [21], (see details in Sec. 3.2), (8) CnC [49], which builds upon JTT [21], but it uses a contrastive loss to train the debiased model, (9) CA [48], trains contrastive adapters by learning similar representations across different groups, (10) CFR [47], refines the representation by aligning them with the class centroids via contrastive loss. We also compare with three group-supervised methods – (11) GroupDRO [32], (12) DFR [17], (13) S-CS [44] and (14) S-SL [44].

6.3. Training and Evaluation Details

For each method, we evaluate both its worst-group accuracy and average (in-distribution) accuracy on the test sets. We use CLIP ResNet-50 and CLIP ViT-L/14 as primary models, while also reporting results for CLIP ResNet-101, ViT-B/32, and ViT-B/16. Following prior work [44, 47, 48], we freeze the model backbones and train all models using the SGD optimizer for 100 epochs. The hyperparameter τ is set to {1.2, 1.0, 0.9, 1.0, 1.0} for Waterbirds, CelebA, MetaShift, Living-17, and BAR, respectively. Model selection is based on the highest worst-group validation ac-

Table 3. Experimental results on Living-17 and BAR benchmarks.

Method	CLIP ResNet-50			CLIP ViT-L/14		
	Living-17		BAR	Living-17		BAR
	WGA	Avg	WGA	WGA	Avg	WGA
ZS [31]	38.0	83.2	74.4	48.0	92.3	81.7
ERM [37]	53.3	90.8	54.9	84.0	98.6	89.0
JTT [21]	44.0	86.4	58.5	78.7	97.3	90.2
CA [48]	62.0	90.9	-	80.0	97.5	-
PPA (ours)	65.0	92.5	76.8	85.0	98.1	91.5

curacy, except for BAR, which lacks a validation set; for BAR, we evaluate the checkpoint from the last epoch. Additionally, since the majority groups in training set do not appear in testing, we measure only the worst-group accuracy for BAR. Further implementation details are provided in Appendix B.

6.4. Main Results

Tab. 2 presents the worst-group accuracy (WGA) and average (in-distribution) accuracy for all methods on Waterbirds [32], CelebA [23] and MetaShift [20] datasets. Compared to other group-unsupervised debiasing approaches, our PPA achieves superior performance with less than 0.1% trainable parameters (CLIP ResNet-50 has 25,557,032 parameters, while our model $f_d(\cdot)$ has only 4,096 trainable parameters). Specifically, PPA based on CLIP ResNet-50 outperforms the previous state-of-the-art, CFR [47], with improvements of {7.4%, 17.4%, 9.3%} on Water-

Table 4. PPA consistently improves group robustness for across different efficient fine-tuning paradigms using CLIP ResNet-50. The results using CLIP ViT-L/14 are reported in Appendix B.3.

Method	Waterbirds		CelebA		MetaShift	
	WGA	Avg	WGA	Avg	WGA	Avg
Linear Probe + ERM	7.9	93.5	11.9	94.7	75.4	94.4
Linear Probe + PPA	84.3	88.3	91.1	92.1	90.8	94.7
CoOp + ERM	61.7	96.2	35.0	94.8	75.4	92.8
CoOp + PPA	83.0	87.6	90.0	92.2	89.0	93.3
Adapter + ERM	61.4	96.7	60.6	94.4	84.6	93.9
Adapter + PPA	78.2	91.8	91.5	92.5	87.7	93.4

Table 5. PPA consistently improves group robustness on Waterbirds across different architectures.

Method	CLIP RN-101		CLIP ViT-B/32		CLIP ViT-B/16	
	WGA	Avg	WGA	Avg	WGA	Avg
Zero-shot	33.6	90.0	47.0	88.8	34.0	88.1
ERM LP	62.9	96.5	56.8	95.6	66.5	96.4
PPA	83.6	89.7	83.5	89.1	84.4	89.9

birds, CelebA, and MetaShift, respectively. Moreover, PPA is competitive with group-supervised methods that rely on group annotations. Notably, our improvements in worst-group accuracy come at the cost of a modest reduction in in-distribution accuracy, with an average drop of 1.8 points compared to ERM models using CLIP ViT-L/14. In Tab. 3, we report results for the Living-17 [33] and BAR [26] datasets. We again observe that PPA achieves the highest worst-group performance. Notably, for the BAR dataset, where model selection is not performed and only the last checkpoint is used (due to the absence of a validation set), PPA still attains the highest worst-group accuracy.

6.5. Further Analysis and Ablation Studies

We conduct comprehensive studies to (1) confirm the versatility of our proposed method—specifically, its compatibility with other efficient fine-tuning paradigms and model architectures—and (2) to understand the contribution of each component to the overall improvement.

Extending PPA to other parameter-efficient fine-tuning paradigms. In this paper, we adopt linear probes for efficient fine-tuning of foundation models. We also explore the applicability of PPA to other efficient fine-tuning methods, such as prompt tuning [52, 53] (CoOp) and adapter-based tuning [7, 9, 48] (small bottleneck MLPs). Specifically, we replace the linear probe $h_d(\mathbf{x}) = W_d \mathbf{x}$ in Step 2 of Algorithm 1 with the following models:

- Prompt tuning (CoOp). The prompt for group i is designed as:

$$\mathbf{t}_i = [V_1^i][V_2^i] \dots [V_M^i][\text{CLASS}], \quad (20)$$

Table 6. **Main component analysis.** We present the worst group accuracies using CLIP ResNet-50. “**Proj.**” and “**GLA**” stands for the projection operation in Eq. (6) and group logit adjustment loss in Eq. (9), respectively. “**GT**” means we use the ground-truth group labels for training debiased models.

	Proj.	GLA	GT	Waterbirds	CelebA	MetaShift
(a)				7.9	11.9	75.4
(b)	✓			54.4	29.4	86.2
(c)		✓		81.6	70.0	89.2
(d)	✓	✓		84.3	91.1	90.8
(e)		✓	✓	86.8	91.5	91.3




Figure 1. The precision and recall of the worst-group examples on Waterbirds and CelebA using CLIP ResNet-50.

where $[V_m^i]$ is a vector with the same dimensionality as the word embedding and serves as a parameter to be optimized. We follow CoOp to set $M = 16$. $[\text{CLASS}]$ represents the word embedding of the class name. Let $\Phi(\cdot)$ denote the CLIP text encoder. The group scorer for group i is then defined as:

$$h_d^{\text{pt}}(\mathbf{x}; \mathbf{t})_i = \Phi(\mathbf{t}_i)^\top \mathbf{x} / \tau_{\text{clip}}, \quad (21)$$

where τ_{clip} is the temperature.

- Adapter-based tuning. We adopt [7] to transform the visual embeddings $\mathbf{x} \in \mathbb{R}^D$ and the text embedding $Z \in \mathbb{R}^{K \times D}$ using simple 2-layer bottleneck MLPs. Let H denote the feature the hidden-layer dimension. We follow [48] to set $H = 128$. With ReLU function $\sigma(\cdot)$, visual adapter weights $W_v^1 \in \mathbb{R}^{H \times D}$, $W_v^2 \in \mathbb{R}^{D \times H}$ and textual adapter weights $W_t^1 \in \mathbb{R}^{H \times D}$, $W_t^2 \in \mathbb{R}^{D \times H}$, the input \mathbf{x} and text embeddings Z are adapted as:

$$\tilde{\mathbf{x}} = W_v^2 \sigma(W_v^1 \mathbf{x}), \quad \tilde{Z}^\top = W_t^2 \sigma(W_t^1 Z^\top). \quad (22)$$

We then use the normalized adapted embeddings and the temperature τ_{clip} to output the final scorer:

$$h_d^{\text{at}}(\mathbf{x})_i = \tilde{\mathbf{z}}^\top \tilde{\mathbf{x}} / \tau_{\text{clip}} \quad (23)$$

Tab. 4 demonstrates consistent improvements in worst-group accuracy over the ERM-trained CLIP ResNet-50 model across three benchmarks, highlighting the versatility




Figure 2. Study on the effect of τ using CLIP ResNet-50. The optimal worst group accuracy is typically observed around $\tau = 1.0$.

of our PPA across various efficient fine-tuning paradigms. For instance, on the CelebA dataset, our PPA achieves worst-group accuracy improvements of 55.0% with CoOp and 30.9% with Adapter. Additionally, integrating PPA with these three fine-tuning methods yields comparable performance. However, linear probing is more computationally efficient than prompt tuning, as the latter requires both forward and backward passes through the large text encoder to update the prompts. It is also more storage-efficient than Adapter-based methods, which introduce additional training parameters, *i.e.*, θ and ϕ .

Study on model versatility. In Tab. 5, we implement our PPA to various backbone architectures, including CLIP ResNet-101, CLIP ViT-B/32 and CLIP ViT-B/16. We observe consistent performance gains on the worst group accuracy, indicating the model versatility of our PPA.

Main component analysis. In Tab. 6, we provide an ablation study to assess the impact of the projection operation (Proj.) and group logit adjustment (GLA) loss on worst-group accuracy using CLIP ResNet-50. Row (a) represents the ERM-trained linear probing model; Row (b) probes group targets without prior correction; and Row (c) omits the projection operation during training. Row (d) shows results from our proposed method (PPA), while Row (e) replaces estimated group labels with ground-truth labels. Comparing Rows (b) and (d), we observe that adjusting the margin significantly improves worst-group accuracy. In the comparison of Rows (c) and (d), projecting out class proxies enhances group identification, resulting in better worst-group performance. Finally, Row (e) demonstrates the oracle performance using ground-truth labels, revealing that PPA approximates the maximum achievable accuracy, underscoring the importance of accurate pseudo-group labels.

Study on the quality of estimated group labels. In the absence of training group labels, we rely on the quality of pseudo group labels inferred from the biased model $f_b(\mathbf{x})$. To evaluate the quality of identified minority groups within the training set, we measure two metrics: *precision*, defined as the fraction of examples within the identified minority groups that belong to the worst-performing group, and *recall*, the fraction of examples from the worst-performing

group captured by the identified minority groups. Here, the worst-performing group is defined as the group on which the ERM model achieves the lowest validation accuracy. In Fig. 1, we compare our method with two baselines: (1) ERM, which uses standard empirical risk minimization to train biased models, and (2) CA [48], which leverages zero-shot models to identify minority groups. Our method, PPA, achieves significantly higher recall and precision than both ERM and CA. This improvement is attributed to our biased model f_b , which relies more heavily on spurious features to make predictions.

Study on the hyper-parameter τ . We examine the effect of $\tau \in \mathbb{R}_+$ in Eq. (9), as illustrated in Fig. 2 (more details are in Appendix B.6). Specifically, we vary τ from 0.0 to 2.0 and report both the worst-group accuracy and average accuracy on the Waterbirds, CelebA, and MetaShift datasets using CLIP ResNet-50. A larger τ promotes a greater margin for the minority groups, while $\tau = 0$ reduces to the standard cross-entropy loss. Our findings reveal that (1) average accuracy decreases monotonically as τ increases, and (2) the optimal worst-group accuracy is generally achieved around $\tau = 1.0$. For example, the optimal values of τ are 1.2, 1.0 and 0.9 for Waterbirds, CelebA and MetaShift, respectively. This aligns with our theoretical analysis, which suggests that the *Bayes* optimal classifier for minimizing the balanced group error is attained when $\tau = 1$.

7. Conclusion

In this paper, we introduce Project-Probe-Aggregate (PPA), a novel, parameter-efficient fine-tuning framework for enhancing the group robustness of image-text foundation models without relying on training group labels. Our approach improves the identification of minority groups by projecting image features onto the null space of class proxies. Additionally, we design a group logit adjustment loss to minimize balanced group error and mitigate spurious correlations. Through theoretical analysis and empirical evaluations across diverse benchmarks, fine-tuning paradigms and model architectures, we demonstrate that the proposed method effectively enhances robust group performance of foundation models.

Acknowledgment

This research is supported by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore.

References

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. [1](#), [2](#)
- [2] Saeid Asgari, Aliasghar Khani, Fereshte Khani, Ali Gholami, Linh Tran, Ali Mahdavi Amiri, and Ghassan Hamarneh. Masktune: Mitigating spurious correlations by forcing to explore. In *NeurIPS*, 2022. [2](#)
- [3] Ching-Yao Chuang, Varun Jampani, Yuanzhen Li, Antonio Torralba, and Stefanie Jegelka. Debiasing vision-language models via biased prompts. *arXiv preprint arXiv:2302.00070*, 2023. [1](#), [2](#), [6](#)
- [4] Elliot Creager, Jörn-Henrik Jacobsen, and Richard Zemel. Environment inference for invariant learning. In *ICML*, 2021. [2](#), [3](#)
- [5] Nikolay Dageev, Brett D Roads, Xiaoliang Luo, Daniel N Barry, Kaustubh R Patil, and Bradley C Love. A too-good-to-be-true prior to reduce shortcut reliance. *Pattern recognition letters*, 166:164–171, 2023. [1](#), [2](#), [3](#)
- [6] Mateo Espinosa Zarlenga, Swami Sankaranarayanan, Jerone TA Andrews, Zohreh Shams, Mateja Jamnik, and Alice Xiang. Efficient bias mitigation without privileged information. In *ECCV*, 2024. [2](#)
- [7] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *IJCV*, 2024. [7](#)
- [8] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. [1](#)
- [9] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *ICML*, 2019. [7](#)
- [10] Badr Youbi Idrissi, Martin Arjovsky, Mohammad Pezeshki, and David Lopez-Paz. Simple data balancing achieves competitive worst-group-accuracy. In *Conference on Causal Learning and Reasoning*, 2022. [2](#)
- [11] Pavel Izmailov, Polina Kirichenko, Nate Gruver, and Andrew G Wilson. On feature learning in the presence of spurious correlations. In *NeurIPS*, 2022. [2](#)
- [12] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. [1](#)
- [13] Sangwon Jung, Sanghyuk Chun, and Taesup Moon. Learning fair classifiers with partially annotated group labels. In *CVPR*, 2022. [2](#)
- [14] Eungyeup Kim, Jihyeon Lee, and Jaegul Choo. Biaswap: Removing dataset bias with bias-tailored swapping augmentation. In *ICCV*, 2021. [2](#)
- [15] Nayeong Kim, Sehyun Hwang, Sungsoo Ahn, Jaesik Park, and Suha Kwak. Learning debiased classifier with biased committee. In *NeurIPS*, 2022. [2](#)
- [16] Younghyun Kim, Sangwoo Mo, Minkyu Kim, Kyungmin Lee, Jaeho Lee, and Jinwoo Shin. Discovering and mitigating visual biases through keyword explanation. In *CVPR*, 2024. [2](#), [13](#)
- [17] Polina Kirichenko, Pavel Izmailov, and Andrew Gordon Wilson. Last layer re-training is sufficient for robustness to spurious correlations. In *ICML*, 2022. [2](#), [3](#), [6](#)
- [18] Jungsoo Lee, Eungyeup Kim, Juyoung Lee, Jihyeon Lee, and Jaegul Choo. Learning debiased representation via disentangled feature augmentation. In *NeurIPS*, 2021. [2](#)
- [19] Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *CVPR*, 2019. [1](#)
- [20] Weixin Liang and James Zou. Metashift: A dataset of datasets for evaluating contextual distribution shifts and training conflicts. In *ICLR*, 2022. [5](#), [6](#)
- [21] Evan Z Liu, Behzad Haghighi, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training group information. In *ICML*, 2021. [1](#), [2](#), [3](#), [4](#), [6](#)
- [22] Sheng Liu, Xu Zhang, Nitesh Sekhar, Yue Wu, Prateek Singhal, and Carlos Fernandez-Granda. Avoiding spurious correlations via logit correction. In *ICLR*, 2023. [1](#), [2](#), [5](#)
- [23] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. [3](#), [5](#), [6](#)
- [24] Aditya Krishna Menon, Ankit Singh Rawat, and Sanjiv Kumar. Overparameterisation and worst-case generalisation: friend or foe? In *ICLR*, 2020. [2](#)
- [25] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *ICLR*, 2021. [3](#), [5](#)
- [26] Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: De-biasing classifier from biased classifier. In *NeurIPS*, 2020. [1](#), [2](#), [3](#), [5](#), [7](#)
- [27] Junhyun Nam, Jaehyung Kim, Jaeho Lee, and Jinwoo Shin. Spread spurious attribute: Improving worst-group accuracy with spurious attribute estimation. In *ICLR*, 2022. [1](#), [2](#), [3](#)
- [28] Hoang Phan, Andrew Gordon Wilson, and Qi Lei. Controllable prompt tuning for balancing group distributional robustness. *arXiv preprint arXiv:2403.02695*, 2024. [2](#)
- [29] Hoang Phan, Andrew Gordon Wilson, and Qi Lei. Controllable prompt tuning for balancing group distributional robustness. In *ICML*, 2024. [1](#)
- [30] Shikai Qiu, Andres Potapczynski, Pavel Izmailov, and Andrew Gordon Wilson. Simple and fast group robustness by automatic feature reweighting. In *International Conference on Machine Learning*, pages 28448–28467. PMLR, 2023. [2](#), [6](#)

- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 5, 6
- [32] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks. In *ICLR*, 2019. 1, 2, 3, 5, 6
- [33] Shibani Santurkar, Dimitris Tsipras, and Aleksander Madry. Breeds: Benchmarks for subpopulation shift. In *ICLR*, 2021. 5, 7
- [34] Saeid A Taghanaki, Kristy Choi, Amir Hosein Khasahmadi, and Anirudh Goyal. Robust representation learning via perceptual similarity metrics. In *ICML*, 2021. 2
- [35] Enzo Tartaglione, Carlo Alberto Barbano, and Marco Grangetto. End: Entangling and disentangling deep representations for bias correction. In *CVPR*, 2021. 2
- [36] Christos Tsirigotis, Joao Monteiro, Pau Rodriguez, David Vazquez, and Aaron C Courville. Group robust classification without any group information. *NeurIPS*, 2024. 2
- [37] Vladimir Vapnik. Principles of risk minimization for learning theory. In *NeurIPS*, 1991. 6
- [38] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 5
- [39] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *CVPR*, 2020. 2
- [40] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. In *ICLR*, 2022. 1
- [41] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *CVPR*, 2022. 6
- [42] Shirley Wu, Mert Yuksekgonul, Linjun Zhang, and James Zou. Discover and cure: Concept-aware mitigation of spurious correlation. In *ICML*, 2023. 2
- [43] Yadollah Yaghoobzadeh, Soroush Mehri, Remi Tachet, Timothy J Hazen, and Alessandro Sordani. Increasing robustness to spurious correlations using forgettable examples. *arXiv preprint arXiv:1911.03861*, 2019. 2
- [44] Yu Yang, Besmira Nushi, Hamid Palangi, and Baharan Mirzasoleiman. Mitigating spurious correlations in multi-modal models during fine-tuning. In *ICML*, 2023. 1, 2, 6
- [45] Yuzhe Yang, Haoran Zhang, Dina Katabi, and Marzyeh Ghassemi. Change is hard: A closer look at subpopulation shift. *arXiv preprint arXiv:2302.12254*, 2023. 5
- [46] Yao-Yuan Yang, Chi-Ning Chou, and Kamalika Chaudhuri. Understanding rare spurious correlations in neural networks. *arXiv preprint arXiv:2202.05189*, 2022. 2
- [47] Chenyu You, Yifei Min, Weicheng Dai, Jasjeet S Sekhon, Lawrence Staib, and James S Duncan. Calibrating multi-modal representations: A pursuit of group robustness without annotations. In *CVPR*, 2024. 1, 2, 6
- [48] Michael Zhang and Christopher Ré. Contrastive adapters for foundation model group robustness. In *NeurIPS*, 2022. 1, 2, 5, 6, 7, 8
- [49] Michael Zhang, Nimit S Sohoni, Hongyang R Zhang, Chelsea Finn, and Christopher Re. Correct-n-contrast: a contrastive approach for improving robustness to spurious correlations. In *ICML*, 2022. 1, 2, 3, 6
- [50] Yuhui Zhang, Jeff Z HaoChen, Shih-Cheng Huang, Kuan-Chieh Wang, James Zou, and Serena Yeung. Diagnosing and rectifying vision models using language. *arXiv preprint arXiv:2302.04269*, 2023. 2
- [51] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE TPAMI*, 40(6):1452–1464, 2017. 5
- [52] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 2022. 7
- [53] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. In *ICCV*, 2023. 7
- [54] Beier Zhu, Kaihua Tang, Qianru Sun, and Hanwang Zhang. Generalized logit adjustment: Calibrating fine-tuned models by removing label bias in foundation models. In *NeurIPS*, 2023. 3
- [55] Beier Zhu, Jiequan Cui, and Hanwang Zhang. Robust fine-tuning of zero-shot models via variance reduction. In *NeurIPS*, 2024. 2