

Classes Are Not Equal: An Empirical Study on Image Recognition Fairness

Jiequan Cui¹ Beier Zhu¹ Xin Wen² Xiaojuan Qi² Bei Yu³ Hanwang Zhang¹

Nanyang Technological University¹ The University of Hong Kong² The Chinese University of Hong Kong³

Abstract

In this paper, we present an empirical study on image recognition unfairness, *i.e.*, extreme class accuracy disparity on balanced data like ImageNet. We demonstrate that classes are not equal and unfairness is prevalent for image classification models across various datasets, network architectures, and model capacities. Moreover, several intriguing properties of fairness are identified. First, the unfairness lies in problematic representation rather than classifier bias distinguished from long-tailed recognition. Second, with the proposed concept of Model Prediction Bias, we investigate the origins of problematic representation during training optimization. Our findings reveal that models tend to exhibit greater prediction biases for classes that are more challenging to recognize. It means that more other classes will be confused with harder classes. Then the False Positives (FPs) will dominate the learning in optimization, thus leading to their poor accuracy. Further, we conclude that data augmentation and representation learning algorithms improve overall performance by promoting fairness to some degree in image classification.

1. Introduction

In the past decade, significant advancements have been made in image recognition. Researchers have diligently explored a multitude of techniques to continually enhance the recognition capabilities of deep models, such as data augmentations [7, 46, 47], the evolution of model architectures [17, 18, 25, 41], and representation learning [13, 19, 20, 34]. However, it's worth noting that most of these endeavors primarily focus on achieving state-of-the-art overall accuracy, often overlooking the goal of consistent performance across all interested classes.

Recently, long-tailed recognition has garnered significant attention within the realms of computer vision and machine learning, due to its growing relevance in real-world applications. In practical scenarios, data often exhibits a long-tailed distribution, *i.e.*, a few classes occupy plenty of data while most of the classes only have a few sam-

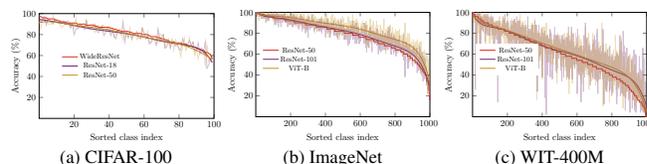


Figure 1. **The unfairness is prevalent in image classification across various datasets, network architectures, and model capacities.** We sort classes by the performance of WideResNet-34-10 on CIFAR-100 and ResNet-50 on ImageNet. For CLIP models, we sort classes by the zero-shot performance on ImageNet of the CLIP ResNet-50 model. Note that data rebalancing is considered in the collection of WIT-400M [37].

ples. Models trained on long-tailed data show extremely imbalanced performance. Especially, the accuracies of the low-frequency classes are pretty poor. However, as shown in Figure 1, we identify that a significant accuracy disparity can occur even on balanced datasets such as ImageNet [15]. *The best class achieves 100% top-1 accuracy while the worst class only achieves 16% top-1 accuracy with a ResNet-50 model on ImageNet.* This phenomenon suggests that factors beyond class frequency contribute to imbalanced performance. We hope that our findings will draw further attention from the research community to this pressing issue.

In this paper, we empirically investigate the image classification fairness issue, *i.e.*, *extreme accuracy disparity among classes*, on 8 balanced datasets including CIFAR and ImageNet. Our analysis also encompasses vision-language models, *i.e.*, CLIP [37] and the stable diffusion model [38]. With a range of network architectures (CNNs and vision transformers), model capacities, and datasets, we identify this as a widespread issue in image classification models.

Representation or classifier bias? Inspired by long-tailed recognition, we investigate the root of unfairness regarding classifier bias including ℓ_2 -norm of classifier weight and class separation angles. Challenging the conventional belief that larger separation angles correlate with higher accuracy, our findings reveal an intriguing contrast: *classes with lower performance exhibit larger separation angles.* Confirmed with k -nearest neighbors (k -NN) algorithm and ETF classifier, we conclude that the problematic representation serves as a crucial factor contributing to the unfairness.

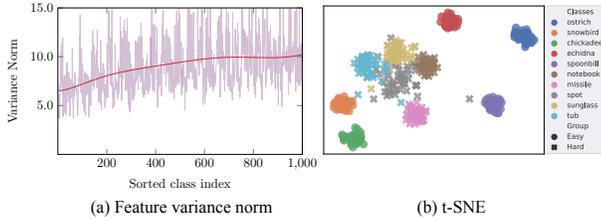


Figure 2. **Data diversity imbalance.** (a) Feature variance ℓ_2 -norm. (b) t-SNE visualization on ImageNet. We sort classes by the performance of ResNet-50 on ImageNet.

Optimization challenges for problematic representation.

We examine the underlying causes of problematic representation within the optimization process. With the proposed concept of *Model Prediction Bias*, our observations indicate a striking trend: *The model displays higher prediction bias on classes that are more challenging to recognize.* This trend exactly contradicts the principles of long-tailed recognition. It means that the harder the class, the more other classes will be confused with it, leading to False Positives (FPs) overwhelming True Positives (TPs) learning in the training optimization and thus their poor accuracy.

Meanwhile, as shown in Figure 2(a), the increasing curve of “the ℓ_2 -norm of the class mean feature variance” implies that the feature distribution of hard classes is more diverse than that of easy classes. The t-SNE visualization in Figure 2(b) also indicates that the features of easy classes are more compact and separable than those of hard classes. This phenomenon suggests that data associated with hard classes encompasses greater diversity, covering more complex scenarios and resulting in overlaps with other classes in high probability. Such overlaps can lead to class confusion in optimization. This is coherent with the analysis of model prediction bias.

Finally, we study the techniques to improve fairness. Data augmentations and representation learning algorithms promote overall performance, usually ignoring per-class accuracy. With sound experiments, we observe that there are more performance gains in hard classes than in easy ones, leading to better fairness. A combination of data augmentations and representation learning tricks can boost fairness along with improved overall accuracy.

Our key contributions are as follows.

- We identify that the unfairness, *i.e.*, extreme class accuracy disparity, is a general problem in image classification even on balanced data like ImageNet [15].
- We study several intriguing properties of the unfairness phenomenon to understand it better. First, the extreme class accuracy disparity comes from problematic representation instead of biased classifiers. Second, with analysis of the model prediction bias, the harder the class is, the more other classes will be confused with it, thus leading to hard optimization and poor accuracy.

- We observe that data augmentations and representation learning algorithms achieve better fairness with improved overall accuracy while re-weighting sacrifices the accuracy of easy classes.

2. Related Work

Fairness. While unfairness is typically defined in terms of disparities related to sensitive attributes, like gender, race, disabilities, and sexual or political orientation, it is a growing concept. In federated learning [28], fairness can be a measurement of the degree of uniformity in performance across federated client devices. For group distributionally robust optimization (DRO), spurious correlations are studied with group labels on the Waterbirds [40]. In this paper, the unfairness is defined as the class accuracy disparities for image classification models.

Imbalanced Learning. Re-weighting [14] and re-sampling [21] are two classical methods to deal with data imbalance. Since Cao et al. [4] and Kang et al. [22] observe that re-weighting and re-sampling can hurt the learned representation and thus degrade performance, two-stage methods [48, 49] become popular, *i.e.*, decoupling representation and classifier learning. Wang et al. [42] validate the effectiveness of model ensembling on long-tailed data. Cui et al. [11] propose the residual mechanism for imbalanced learning. After that, methods based on representation learning [10, 13] achieve new state-of-the-art results.

Data Augmentation and Representation Learning. Data augmentations [7, 8, 16, 46, 47] can significantly improve the generalization ability of models taking nearly no additional cost. Single-image augmentations [7, 8, 16] do image-processing operations on each image independently. Cross-image augmentations [46, 47] mix images and their labels simultaneously. Representation learning is a fundamental task. Contrastive loss [6, 19] measures the similarities of sample pairs in feature space, and positive sample pairs are encouraged to have similar representations. Recently, masked modeling transferred from the neural language processing (NLP) community has made great progress in learning good representation for images, like [1, 20], benefiting a wide spectrum downstream tasks.

Data augmentations and representation learning techniques both improve feature quality for recognition, thus promoting overall accuracy. However, per-class accuracy is usually ignored and we bridge the gap in this paper.

3. Fairness Issue Widely Exists

To showcase the widespread prevalence of fairness issues in image classification models, we conduct experiments on 8 datasets including CIFAR-100 [24] and ImageNet [39] with various network architectures including convolutional neu-

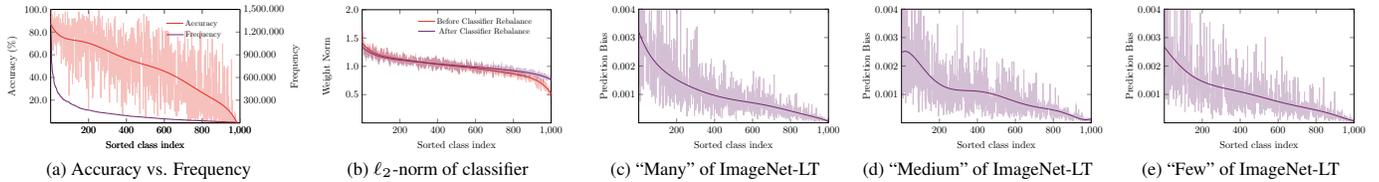


Figure 3. **Analysis on ImageNet-LT.** (a) Data distribution and per-class accuracy with ResNet-50 on ImageNet-LT; (b) The ℓ_2 -norm of classifier weights before/after classifier rebalancing; (c), (d), and (e) evaluate the prediction bias on “Many”, “Medium”, and “Few” classes data separately. We sort classes by the number of samples in the classes on ImageNet-LT.

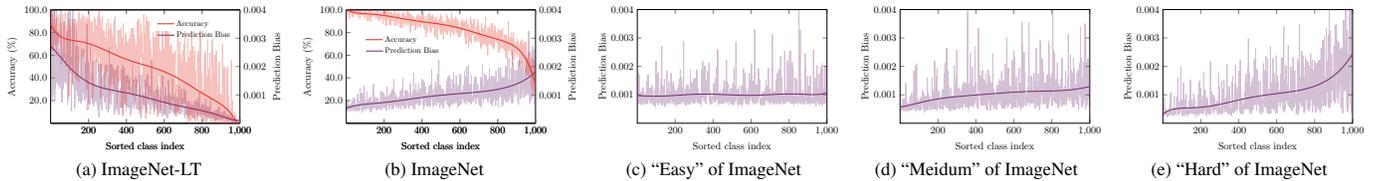


Figure 4. **Model prediction bias.** (a) and (b) show the relationship between class accuracy and prediction bias. ImageNet-LT and ImageNet appear contrary conclusions. The details are discussed in Section 4.2. (c), (d), and (e) evaluate the prediction bias on “Easy”, “Medium”, and “Hard” classes data separately. ResNet-50 is used for ImageNet-LT while ViT-B is adopted on ImageNet. Classes are sorted by their frequency on ImageNet-LT. On ImageNet, classes are sorted by the performance of ResNet-50.

Table 1. **Top-1 accuracy on ImageNet and CIFAR-100.** CLIP models trained on large-scale image-text pairs are also included.

Method	Easy	Medium	Hard	All	Best	Worst
ImageNet						
ResNet-50	93.1	81.1	59.4	77.8	100.0	16.0
ResNet-101	94.1	82.6	61.5	79.4	100.0	16.0
ViT-B	95.4	86.6	68.9	83.6	100.0	20.0
CIFAR-100						
ResNet-50	89.8	81.4	68.5	79.8	98.0	51.0
WideResNet-34-10	92.2	83.2	70.1	81.7	98.0	57.0
CLIP model on ImageNet						
ResNet-50	84.3	61.8	33.4	59.8	100.0	0.0
ResNet-101	83.6	63.1	40.3	62.3	100.0	0.0

ral networks (CNNs) and vision transformers along with different model capacities. CLIP [37] and the Stable Diffusion [38] models trained on large-scale image-text pair data are also considered in our study. Please refer to the supplementary file for analysis of more datasets.

Training on CIFAR-100. The CIFAR-100 dataset has 50,000 training images and 10,000 testing images with 100 classes. Each class has 500 training images and 100 testing images. Following previous work [9], standard data augmentation including random crops with 4 pixels of padding and random horizontal flip is performed for data preprocessing in the course of training.

We use WideResNet34-10, and ResNet-50 as our backbones and follow the training schedule in [12, 43]. We train the models 200 epochs with a cosine learning rate strategy and a batch size of 128 on 4 GPUs. The initial learning rate is set to 0.2. A weight decay of $5e-4$ is adopted. We use the

SGD optimizer with a momentum of 0.9 in training.

Training on ImageNet. ImageNet [39] is one of the most challenging datasets for classification, which consists of 1.2 million images for training and 50K images for validation over 1,000 classes. It is worth noting that ImageNet is a carefully curated dataset. Data is balanced over all classes and each class has nearly 1000 training samples.

We use CNNs (ResNet-50 and ResNet-101) and vision transformers (ViT-B) as our backbones. For CNNs, we follow the training schedule from [46]. Models are trained in 300 epochs with a cosine learning rate strategy and a batch size of 256 on 8 GPUs. The initial learning rate and the weight decay are set to 0.1 and $1e-4$ separately. The SGD optimizer with a momentum of 0.9 is used.

For training with ViT-B, we follow the training schedule in [20], *i.e.*, firstly pre-train it with masked modeling technique and then fine-tune it 100 epochs. Our implementation is built on their open-sourced code.

Evaluation. In the long-tailed recognition field, “Many”, “Medium”, and “Few” classes are defined for evaluation according to the number of samples in the classes. Similarly, to quantify the unfairness, we define “Hard”, “Medium”, and “Easy” classes with a well-trained reference model. ResNet50 and WideResNet34-10 are adopted as the reference models for ImageNet and CIFAR-100 separately. The number of classes in “Hard”, “Medium”, and “Easy” are evenly divided according to the performance of the reference model. For CLIP [37] models, we directly evaluate their pre-trained model on the ImageNet validation set and sort the classes with the performance of the ResNet-50 CLIP model.

Results Summary. We summarize the experimental results in Figure 1 and Table 1. As plotted in Figure 1, from easy to hard classes, the accuracy significantly decreases for models trained on CIFAR-100, ImageNet, and WIT-400M. The phenomenon is consistent across various backbones including CNNs and vision transformers. Especially, the best class “ostrich” achieves **100%** top-1 accuracy while the worst class “screen” only obtains **16%** top-1 accuracy on ImageNet with ResNet-50, which demonstrates a large performance disparity on balanced data.

As shown in Table 1, “Easy” classes can get much better performance than “Hard” classes although training samples are evenly distributed among all the classes. For the CLIP ResNet-50 model, the average accuracy of “Easy” classes is 50% higher than that of “Hard” classes. The huge accuracy gap between the best (**100%**) and the worst (**0%**) classes indicates a surprising performance imbalance on balanced datasets. It is worth noting that data rebalancing strategies have already been applied in the collection of WIT-400M. Analysis of the other five datasets is included in the supplementary file.

With these experimental analyses, we conclude that the fairness issue in image classification is a general problem across various datasets, network architectures, and model capacities. Further, the similar trends on the same datasets with various models demonstrate that the unfairness highly depends on training data distribution. We expect that our study could raise more attention from the community to this issue and facilitate progress.

Discussion. When trained with long-tailed data, models tend to perform well on high-frequency classes but show significantly reduced accuracy on low-frequency ones. This pattern suggests to us that class frequency can be the primary driver of the extremely imbalanced performance. Although several works start to consider hard example learning, most works [4, 11, 14, 22] in long-tailed recognition still focus on this aspect. However, *our research reveals that significant performance disparities (100% accuracy on the best class while only 16% on the worst class on ImageNet) also occur in balanced datasets. This indicates the presence of other significant factors that influence per-class accuracy, beyond just class frequency.*

4. Analysis on Image Recognition Fairness

As shown in Figure 1 and Figure 3 (a), extreme performance disparity happens both in long-tailed recognition and unfairness. We delve into the properties of fairness in image classification by comparing it with long-tailed recognition. The discussion on whether bias in representation or classifier leads to unfairness is presented in Section 4.1. Following that, we conduct an in-depth examination of the underlying optimization challenges for fairness in Section 4.2.

4.1. Representation or Classifier Bias?

Inspired by long-tailed recognition, we diagnose the classifier bias for unfairness in terms of their weights ℓ_2 -norm and class separation angles. Then we conclude that the problematic representation results in unfairness.

ℓ_2 -norm of Classifier Weights. Previous works [4, 22, 49] study the classifier bias in long-tailed learning and observe ℓ_2 -norm of classifier weights matters. Kang et al. [22] propose a two-stage pipeline: uniform data sampling for representation learning, and then class-balanced re-sampling for classifier learning with previously learned representation fixed. With the method, we train a ResNet-50 model on ImageNet-LT [30], which obeys the long-tailed distribution shown in Figure 3(a).

With the uniform sampling strategy for representation learning, the ℓ_2 -norm of the learned classifier weights positively correlates to class accuracy and class frequency as plotted in Figure 3(a) and Figure 3(b). After classifier rebalancing using the re-sampling strategy, the norm of classifier weights becomes more uniform, which is visualized in Figure 3(b), and the accuracies of low-frequency classes are promoted.

Interestingly, although the models suffering from unfairness display extremely imbalanced accuracy varying from 100% to 16%, we notice that the ℓ_2 -norm of classifier weights is already balanced for models trained on ImageNet shown in Figure 5(a), implying that the ℓ_2 -norm of classifier weights is not the core factor for unfairness.

Definition 1 (Simplex Equiangular Tight Frame). A collection of vectors $\mathbf{m}_i \in \mathbb{R}^d, i = 1, 2, \dots, K, d \leq K - 1$, is said to be a simplex equiangular tight frame if:

$$\mathbf{M} = \sqrt{\frac{K}{K-1}} \mathbf{U} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T \right), \quad (1)$$

where $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_K] \in \mathbb{R}^{d \times K}$ allows a rotation and satisfies $\mathbf{U}^T \mathbf{U} = \mathbf{I}_K$, \mathbf{I}_K is the identity matrix, and $\mathbf{1}_K$ is an all-ones vector.

All vectors in a simplex ETF have an equal ℓ_2 -norm and the same pair-wise angle, *i.e.*, $\mathbf{m}_i^T \mathbf{m}_j = -\frac{1}{K-1}$ if $i \neq j$.

Maximal Equiangular Separation. The simplex equiangular tight frame (ETF) in Definition 1 is introduced by Pappayan et al. [35]. Then several works [45, 50] explore the ETF classifier for long-tailed recognition. ETF classifier enjoys the same pair-wise angle for any two different classes and allows low-frequency classes to have the maximal equiangular separation.

We examine if bias in unfairness comes from the smaller separation angles between “Hard” and other classes by conducting experiments on ImageNet. The mean angle for class

i , denoted by $\text{CMA}(i)$, is defined as follows:

$$\text{CMA}(i) = \frac{1}{C} \sum_j^C \frac{\mathbf{w}_i \cdot \mathbf{w}_j}{\|\mathbf{w}_i\|_2 \cdot \|\mathbf{w}_j\|_2}, \quad (2)$$

where \mathbf{w}_i is the classifier weight for class i , and C is the number of classes.

As shown in Figure 5(b), most of the ‘‘Hard’’ classes get smaller CMA values than ‘‘Easy’’ ones and thus they even possess larger separation angles with other classes, which is quite contrary to our common belief: *larger separation angles usually indicate a higher accuracy*.

We hypothesize that the problematic representations lead to the contrary. When samples in two classes look similar or have overlapping scenarios, the model can be confused. During optimization, the cross-entropy loss will enlarge the separation angles between the two classes and meanwhile encourage their representations more distinguishable from each other. However, if there is little effect on the representations, the two classes still suffer from poor accuracy despite large separation angles between them.

Problematic Representation in Unfairness. We verify the hypothesis with the following empirical studies:

- Evaluate the model with k -NN algorithm and examine the existence of unfairness.
- Train models with a fixed ETF classifier and examine the existence of unfairness.

For the k -NN algorithm, with k nearest neighbors in training data, the prediction is determined on the basis of a majority vote, *i.e.*, the label most frequently represented around its neighbors. Following DINO [5], k is set to 20 for evaluation. ETF classifier satisfies the maximal equiangular separation principle. We follow the implementation in previous works [45, 50]. As shown in Figure 5(c) and Figure 5(d), the class performance still exhibits extreme imbalance, confirming that problematic representation is the core issue for unfairness.

Discussion. This study highlights the fundamental distinctions between unfairness and long-tailed recognition. Biased classifiers can exacerbate performance disparities in long-tailed recognition. Conversely, unfairness in balanced data stems from problematic representation. It provides valuable insights indicating that effective representation learning algorithms are needed to explicitly improve feature quality for the recognition of ‘‘Hard’’ classes. Building on these findings, we delve into strategies to enhance fairness in Section 5.

4.2. Model Prediction Bias

Models trained on balanced datasets such as ImageNet exhibit problematic representations, leading to unfairness. However, the underlying cause of this issue in terms of

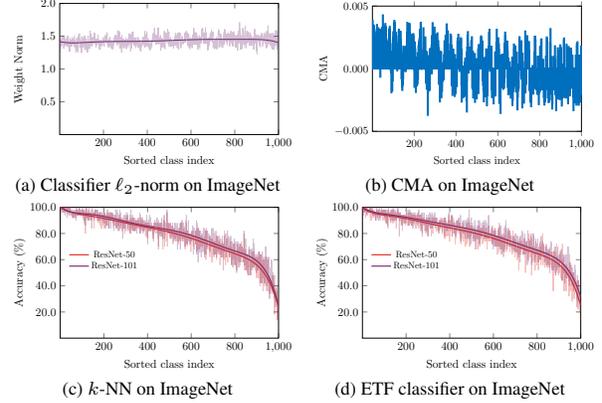


Figure 5. **Analysis of fairness on bias from representation and classifier.** (a) The ℓ_2 -norm of classifier weights on ImageNet; (b) Class mean angles defined in Eq. (2); (c) Unfairness exists with ResNet-50/101 and k -NN algorithm; (d) Unfairness exists with ResNet-50/101 and the ETF classifier; Classes are sorted by the performance of ResNet-50 on ImageNet.

optimization remains a mystery. In this section, we are prompted to explore this question with the introduced concept of *Model Prediction Bias* in Definition 2.

Definition 2 (Model Prediction Bias). For a well-calibrated [26] classification model \mathcal{M} , $x_i \in X$ is a input sample. Then the prediction bias for \mathcal{M} is calculated on X with the following formulation:

$$\mathbf{G}_n = \sum_{x_i \in X} \sigma(\mathcal{M}(x_i)) \odot (\mathbf{1} - \mathbf{y}_i), \quad (3)$$

$$\mathbf{G}_p = \sum_{x_i \in X} \sigma(\mathcal{M}(x_i)) \odot \mathbf{y}_i, \quad (4)$$

where σ is the softmax activation. $\mathcal{M}(x_i) \in \mathbb{R}^C$ is the logits output. \mathbf{y}_i is the one-hot label. C is the number of classes. Normalization can be applied to \mathbf{G}_n or \mathbf{G}_p .

\mathcal{M} is unbiased when X is balanced and \mathbf{G}_n is uniform. \mathbf{G}_p with the normalization indicates the class accuracy for the well-calibrated model \mathcal{M} .

High Accuracy Partially from Prediction Bias in LT.

With the concept of *Model Prediction Bias*, we calculate the prediction bias vector \mathbf{G}_n of the model trained on ImageNet-LT. As shown in Figure 3(b), the ℓ_2 -norm of classifier weights for high-frequency classes is much larger than that for low-frequency ones. This leads to the model being biased toward high-frequency classes on all the data (‘‘Many’’, ‘‘Medium’’, and ‘‘Few’’) when giving predictions, as demonstrated in Figs 3(c), 3(d), 3(e) and 4(a). Thus, the model prediction bias results in a much better performance of high-frequency classes than low-frequency classes. In a word, in long-tailed recognition, the high performance of high-frequency classes benefits from the prediction bias of the trained model to some degree.

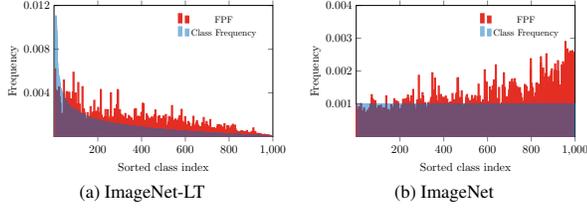


Figure 6. **Class frequency vs. False Positive Frequency (FPF).** Classes are sorted by the number of samples in the class for ImageNet-LT and the performance of ResNet-50 on ImageNet.

Low Accuracy from Prediction Bias for Unfairness. On the contrary, as shown in Figure 4(b), the model suffering from unfairness tends to be biased towards “Hard” classes which are with poor accuracy. In particular, *the harder the class, the higher the prediction bias the model exhibits.* We further calculate the prediction bias \mathbf{G}_n on “Easy”, “Medium”, and “Hard” subsets data individually. Illustrated by Figure 4(c), we observe that the model has little prediction bias with “Easy” subset data. On “Medium” subset data, the slight tendency of prediction bias towards “Medium” and “Hard” classes appears. Figure 4(e) indicates an obvious prediction bias towards “Hard” classes of the model when evaluating on the “Hard” subset data, which implies that the unfairness mainly comes from the confusion among “Hard” classes and thus results in their pretty low performance.

High Accuracy or Low Accuracy from Prediction Bias?

The conclusion appears to be contradictory:

- *In long-tail learning, higher prediction biases are found in “high-frequency” classes which enjoy high accuracies. In particular, all classes are biased towards high-frequency classes as shown in Figs. 3(c), 3(d), and 3(e).*
- *On balanced data, higher prediction biases are observed in “Hard” classes which suffer from low accuracies. Specifically, only “Hard” and “Medium” classes are biased towards “Hard” classes in Figs. 4(c), 4(d), and 4(e).*

We detail the reasons behind the phenomenon by introducing the following Lemma 1.

Lemma 1 (Gradients Convergence Condition). When the training of model \mathcal{M} converges, the gradients with respect to parameters will be $\mathbf{0}$. Then, the following equation is established.

$$\mathbf{G}_n + \mathbf{G}_p = \mathbf{CF}, \quad (5)$$

where $\mathbf{CF}[i]$ represents the number of samples in class i .

Proof. With the rule of backpropagation, the gradients with sample x_i on $\mathcal{M}(x_i)$ is what follows,

$$\frac{\partial \mathcal{L}}{\partial \mathcal{M}(x_i)} = \sigma(\mathcal{M}(x_i)) - \mathbf{y}_i. \quad (6)$$

The gradients integrated on the whole training data should be $\mathbf{0}$ for the gradients convergence of training. Then we

derive the equation:

$$\sum_{x_i \in X} \sigma(\mathcal{M}(x_i)) - \mathbf{y}_i = \mathbf{0}. \quad (7)$$

Expand Eq. (7) with Eq. (3), and Eq. (4),

$$\sum_{x_i \in X} \sigma(\mathcal{M}(x_i)) \odot (\mathbf{1} - \mathbf{y}_i) + \sum_{x_i \in X} \sigma(\mathcal{M}(x_i)) \odot \mathbf{y}_i = \sum_{x_i \in X} \mathbf{y}_i. \quad (8)$$

Eq. (8) is equal to Eq. (5), thus completing the proof.

In Lemma 1, we theoretically prove the relations among the prediction bias \mathbf{G}_n , the class performance \mathbf{G}_p , and the class frequency \mathbf{CF} . As demonstrated by Eq. (5), for long-tailed recognition, the large number of samples (\mathbf{CF}) still allows their good performance (\mathbf{G}_p) for high-frequency classes, although the model shows higher prediction bias (\mathbf{G}_n) towards them. Meanwhile, in terms of the unfairness on balanced data, the higher prediction bias (\mathbf{G}_n) toward “Hard” classes necessarily results in their low accuracy (\mathbf{G}_p) because the class frequency is uniform (\mathbf{CF}).

Intuitively, with *Model Prediction Bias* in Definition 2, \mathbf{G}_n can be regarded as the false positive frequency (FPF) and the class frequency \mathbf{CF} is our expected true positive frequency (TPF). During training, we improve the score of the ground-truth label and decrease the scores for other classes with cross-entropy loss, which means we minimize the FPF and expect the TPF to match the class frequency (\mathbf{CF}). To achieve better accuracy for one specific class, the training should be dominated by true positives (TPs) instead of false positives (FPs) learning. Otherwise, FPs will overwhelm the learning of TPs, leading to hard optimization and resulting in low performance.

As shown in Figure 6(a), in long-tailed recognition, the model displays higher \mathbf{G}_n towards high-frequency classes. However, TPs still dominate the training due to their high \mathbf{CF} . On the contrary, illustrated by Figure 6(b), ImageNet enjoys a uniform \mathbf{CF} and the high \mathbf{G}_n has overwhelmed their TPs learning for “Hard” classes, resulting in their hard optimization and poor accuracy.

The Unfairness Can Stem from Data Bias. We have analyzed the reason for the unfairness phenomenon in terms of training optimization. From the perspective of input data, bias in the datasets can lead to unfairness.

Class co-occurrence happens in the ImageNet [2]. For example, “keyboard” and “monitor” frequently appear in conjunction. Thus, single-label image annotations in ImageNet can lead to confusion and hard optimization. The ImageNet validation set is relabeled as ImageNet-ReaL where multiple labels can be assigned per image. Evaluating on ImageNet-ReaL, the accuracy of the worst class improves from 16% to 45% but is still far behind the best class 100%,

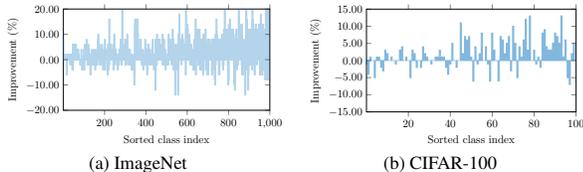


Figure 7. **Improvements of per class with data augmentations on ImageNet and CIFAR-100.** (a) Trained with Autoaug, Mixup, and CutMix on ImageNet. (b) Trained with CutMix on CIFAR-100. ViT-B backbone is used on ImageNet while WideResNet-34-10 is used on CIFAR-100.

demonstrating the existence of other potential factors leading to the imbalance.

Moreover, Some class objects occur in much more complex scenes than others. Example classes in “Easy” and “Hard” subsets on ImageNet are shown in Figures 8 and 9 of the supplementary file. We calculate “the ℓ_2 -norm of the class mean feature variance” with a well-trained ResNet-50 model on ImageNet. As shown in Figure 2, the feature distribution of “Hard” classes has a higher variance than that of “Easy” classes. It demonstrates that “Hard” classes cover much more complex scenarios which can have some overlaps with other classes. In this case, without enough training data for the “Hard” classes, the model can overfit some scenario background and lead to prediction confusion with other classes that share similar scenarios.

4.3. Discussion

As illustrated in Section 3, unfairness is a general problem in image classification models. In this section, we both empirically and theoretically uncover the mystery behind the unfairness from the perspective of gradient optimization during training. Moreover, regarding dataset bias, we observe that features of “Hard” classes usually have much higher variances than “Easy” classes as shown by Figure 2, implying that data for “Hard” classes enjoys larger diversity and possesses lower sample density. From this point of view, the unfairness can be considered as a general imbalance problem, *i.e.*, *data diversity imbalance*. Besides, although we usually assume classes are independent, the relations among classes (*e.g.*, class co-occurrence) will make the *data diversity imbalance* heavier. To our knowledge, there are few efforts to explicitly solve the *data diversity imbalance*. We expect our study to raise more attention from the community to the problem.

5. Improving Fairness

In this section, we investigate how regular training techniques, like data augmentation and representation learning, affect fairness in image classification. Previous work including Mixup [47], CutMix [46], AutoAug [7], RandAug [8], and self-supervised pertaining [20] boosts overall performance. However, fairness is rarely discussed and ex-

plored. We bridge the gap in this work.

5.1. Data Augmentation

Mixup and CutMix. Mixup [47] combines image-label pairs (x_i, y_i) and (x_j, y_j) with a uniform sampled λ ,

$$x_{mix} = \lambda * x_i + (1 - \lambda) * x_j, \quad (9)$$

$$y_{mix} = \lambda * y_i + (1 - \lambda) * y_j, \quad (10)$$

where $\lambda \in [0, 1]$.

CutMix [46] extends Mixup and randomly exchanges an image region in x_i with another one in image x_j . Then the mixed label is determined by mixing y_i and y_j with λ which depends on the area of x_i region in the mixed image.

AutoAug and RandAug. AutoAug [7] is one data augmentation strategy that is automatically learned from the data. It is learned from a pre-defined search space including various image process operations, like translation, rotation, or shearing. RandAug [8] selects image transformations in a well-designed search space with a uniform probability.

Training and Evaluation. We use the same training pipeline as described in Section 3. For evaluation, we calculate the improvements for each class between models with and without using data augmentation strategies. We also calculate the improvements of the “Head”, “Medium”, and “Easy” subsets separately. WideResNet-34-10 is used on CIFAR-100 while ViT-B is adopted for ImageNet.

Results and Analysis. Table 3 shows the improvements on the “Hard”, “Medium”, and “Easy” subsets on ImageNet and CIFAR-100 datasets with data augmentations. With the experimental results, the consistent observation is that improvements of “Hard” classes are much larger than that of “Easy” classes. We also plot per-class improvement on CIFAR-100 and ImageNet. As shown in Figure 7, most of the “Hard” classes usually enjoy higher performance improvement than “Easy” classes after adopting data augmentation. It indicates that data augmentation strategies can improve the performance fairness of image classification models. Moreover, data augmentation strategies can be complementary to each other. Combining all of these techniques, the “Hard” classes obtain more performance gains than applying only a single one.

5.2. Representation Learning

Contrastive Learning. We examine the effects of contrastive learning on fairness in two cases: 1) with the pre-trained weights by MoCo [19] as initialization; 2) using the trained model by GPaco [13]. For fair comparisons, we use the same training setting in GPaco [13] for baseline models. ResNet-50 backbone is used in experiments.

Masked Modeling. Besides contrastive learning, we empirically study how the masked modeling pre-training meth-

Table 2. **Fairness improvements on ImageNet with contrastive learning and masked modeling.** “*” represents the model is trained from scratch in 300 epochs.

Method	Easy	Medium	Hard	All
Contrastive Learning (ResNet-50)				
w/o pre-training	93.4	82.0	62.0	79.1
w/ pre-training	93.4(+0.0)	82.1(+0.1)	62.2(+0.2)	79.2
GPaCo [13]	93.8(+0.4)	82.7(+0.7)	62.6(+0.6)	79.7
Masked Modeling (ViT-B)				
w/o pre-training	92.2	79.6	59.6	77.1
w/o pre-training*	94.3	84.3	64.8	81.1
w/ pertaining	95.4(+1.1)	86.6(+2.3)	68.9(+4.1)	83.6
Contrastive Learning and Masked Modeling (ViT-B)				
GPaCo [13]	95.5(+1.2)	87.2(+2.9)	69.1(+4.3)	84.0

ods [1, 20, 44] affect fairness in image recognition. We follow the same training pipeline as MAE [20] and the ViT-B backbone is adopted in experiments. We also demonstrate that fairness can be boosted by combining contrastive learning and masked modeling pretraining.

Results. We list the results in Table 2. Illustrated by the empirical study, both contrastive learning and masked pertaining can achieve better performance gains on “Hard” classes than “Easy” classes. Further, we also observe that the masked pertaining [20] is more conducive to fairness than contrastive learning. Nevertheless, the two techniques are complementary, and better fairness can be achieved by making them work together.

5.3. Re-weighting and Other Methods

Re-weighting plays an important role in long-tailed recognition [14, 31]. It assigns larger weights for low-frequency classes and smaller weights for high-frequency ones and expects the low-frequency classes to achieve better performance. We transfer the re-weighting method to address the fairness issue. With a validation set, we calculate the effective frequency (EF) for each class as follows. For class c ,

$$EF(c) = \frac{1}{|X_c|} \sum_{x_i \in X_c} \sigma(\mathcal{M}(x_i))[c], \quad (11)$$

where X_c is a subset of the validation set X . Its samples belong to class c . σ is the softmax activation.

With the derived effective frequency (EF), we conduct experiments on ImageNet using [14, 31]. Besides, we investigate the debiased method LfF [32] and worst-group performance optimization method JTT [29] on CIFAR-100.

Results. As shown in Table 4, re-weighting methods can improve the accuracy of “Hard” classes at the cost of some performance degradation on “Easy” and “Medium” classes. However, it fails to promote overall accuracy and even achieves inferior performance than the baseline, which is different from data augmentations and representation

Table 3. **Fairness improvements on ImageNet and CIFAR-100 with data augmentation strategies.**

Method	Easy	Medium	Hard	All
ImageNet				
ViT-B	94.0	84.1	64.2	80.8
+AutoAug	94.9(+0.9)	85.9(+1.8)	66.2(+2.0)	82.4
+Mixup	94.5(+0.5)	85.3(+1.2)	66.4(+2.2)	82.0
+CutMix	94.4(+0.4)	84.6(+0.5)	65.3(+1.1)	81.4
+All	95.4(+1.4)	86.6(+2.5)	68.9(+4.7)	83.6
CIFAR-100				
WideResNet-34-10	92.2	83.2	70.1	81.7
+AutoAug	91.7(-0.5)	83.5(+0.3)	72.7(+2.6)	82.5
+Mixup	91.3(-0.9)	83.5(+0.3)	71.5(+1.4)	82.0
+CutMix	92.2(+0.0)	85.0(+1.8)	73.9(+3.8)	83.6
+Cutout	92.2(+0.0)	84.2(+1.0)	72.5(+2.4)	82.9

Table 4. **Re-weighting and other methods on fairness.**

Method	Easy	Medium	Hard	All
ImageNet				
Baseline (ResNet-50)	93.1	81.1	59.4	77.8
Re-weighting [31]	91.7(-1.4)	79.8(-1.3)	61.4(+2.0)	77.6
Re-weighting [14]	91.6(-1.5)	79.6(-1.5)	61.3(+1.9)	77.5
CIFAR-100				
Baseline (WideResNet34)	92.2	83.2	70.1	81.7
Re-weighting [31]	91.5(-0.7)	83.2(+0.0)	70.4(+0.3)	81.6
LfF [32]	90.4(-1.8)	81.5(-1.7)	67.7(-2.4)	79.7
JtT [29]	91.3(-0.9)	82.4(-0.8)	69.6(-0.5)	81.0

learning techniques. Moreover, we observe that methods for worst-group performance or classifier debiasing cannot handle complex data distributions, like CIFAR-100.

6. Conclusion and Future Work

In this paper, we investigate the image recognition fairness issue, *i.e.*, extreme class performance disparity on balanced data. This unfairness phenomenon widely exists in image classification models. We identify several properties of the fairness issue by comparing it with long-tailed recognition: (1) The unfairness comes from problematic representation instead of classifier bias; (2) The harder the class, the more other classes can be confused with it, resulting in their low performance. Finally, we verify that data augmentation and representation learning tricks can improve fairness with better overall performance.

With our empirical study, representation learning is promising for improving fairness. However, it hasn’t been explored to explicitly promote fairness learning. Designing proper data augmentation strategies or self-supervised pre-training techniques specifically for fairness is an interesting direction. We hope the study will attract more researchers to solve this problem.

Acknowledgements. This research was supported by the funding MOE AcRF Tier 2 (MOE2019-T2-2-062).

References

- [1] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. In *ICLR*, 2022. 2, 8
- [2] Lucas Beyer, Olivier J Hénaff, Alexander Kolesnikov, Xi-aohua Zhai, and Aäron van den Oord. Are we done with imagenet? *arXiv preprint arXiv:2006.07159*, 2020. 6
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *ECCV*, 2014. 12
- [4] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, 2019. 2, 4
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 5
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020. 2
- [7] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. In *CVPR*, 2019. 1, 2, 7
- [8] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPR workshops*, pages 702–703, 2020. 2, 7
- [9] Jiequan Cui, Pengguang Chen, Ruiyu Li, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Fast and practical neural architecture search. In *ICCV*, pages 6509–6518, 2019. 3
- [10] Jiequan Cui, Zhisheng Zhong, Shu Liu, Bei Yu, and Jiaya Jia. Parametric contrastive learning. In *ICCV*, pages 715–724, 2021. 2
- [11] Jiequan Cui, Shu Liu, Zhuotao Tian, Zhisheng Zhong, and Jiaya Jia. Reslt: Residual learning for long-tailed recognition. *IEEE TPAMI*, pages 3695–3706, 2023. 2, 4
- [12] Jiequan Cui, Zhuotao Tian, Zhisheng Zhong, Xiaojuan Qi, Bei Yu, and Hanwang Zhang. Decoupled kullback-leibler divergence loss. *arXiv preprint arXiv:2305.13948*, 2023. 3
- [13] Jiequan Cui, Zhisheng Zhong, Zhuotao Tian, Shu Liu, Bei Yu, and Jiaya Jia. Generalized parametric contrastive learning. *IEEE TPAMI*, pages 1–12, 2023. 1, 2, 7, 8
- [14] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, pages 9268–9277, 2019. 2, 4, 8
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 1, 2
- [16] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 2
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 1
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 1, 2, 7
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 1, 2, 3, 7, 8
- [21] Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, pages 429–449, 2002. 2
- [22] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *ICLR*, 2020. 2, 4
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCV*, pages 554–561, 2013. 12
- [24] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 2
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 1
- [26] Ananya Kumar, Tengyu Ma, Percy Liang, and Aditi Raghunathan. Calibrated ensembles can mitigate accuracy trade-offs under distribution shift. In *UAI*, pages 1041–1051, 2022. 5
- [27] Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your diffusion model is secretly a zero-shot classifier. *arXiv preprint arXiv:2303.16203*, 2023. 11
- [28] Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. Fair resource allocation in federated learning. In *ICLR*, 2020. 2
- [29] Evan Z Liu, Behzad Haghgoo, 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*, pages 6781–6792. PMLR, 2021. 8
- [30] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, pages 2537–2546, 2019. 4
- [31] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *ICLR*, 2021. 8
- [32] Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: De-biasing classifier from biased classifier. In *NeurIPS*, pages 20673–20684, 2020. 8
- [33] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pages 722–729, 2008. 12

- [34] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [1](#)
- [35] Vardan Papayan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *PNAS*, 117(40):24652–24663, 2020. [4](#)
- [36] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *CVPR*, pages 3498–3505, 2012. [11](#)
- [37] 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*, pages 8748–8763, 2021. [1](#), [3](#)
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. [1](#), [3](#), [11](#)
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, pages 211–252, 2015. [2](#), [3](#)
- [40] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks. In *ICLR*, 2019. [2](#)
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [1](#)
- [42] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella Yu. Long-tailed recognition by routing diverse distribution-aware experts. In *ICLR*, 2021. [2](#)
- [43] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *NeurIPS*, pages 2958–2969, 2020. [3](#)
- [44] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, pages 9653–9663, 2022. [8](#)
- [45] Yibo Yang, Shixiang Chen, Xiangtai Li, Liang Xie, Zhouchen Lin, and Dacheng Tao. Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network? *NeurIPS*, pages 37991–38002, 2022. [4](#), [5](#)
- [46] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, pages 6023–6032, 2019. [1](#), [2](#), [3](#), [7](#)
- [47] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. [1](#), [2](#), [7](#)
- [48] Songyang Zhang, Zeming Li, Shipeng Yan, Xuming He, and Jian Sun. Distribution alignment: A unified framework for long-tail visual recognition. In *CVPR*, pages 2361–2370, 2021. [2](#)
- [49] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In *CVPR*, pages 16489–16498, 2021. [2](#), [4](#)
- [50] Zhisheng Zhong, Jiequan Cui, Yibo Yang, Xiaoyang Wu, Xiaojuan Qi, Xiangyu Zhang, and Jiaya Jia. Understanding imbalanced semantic segmentation through neural collapse. In *CVPR*, pages 19550–19560, 2023. [4](#), [5](#)