

MDCS: More Diverse Experts with Consistency Self-distillation for Long-tailed Recognition

Qihao Zhao^{1,2}, Chen Jiang¹, Wei Hu¹, Fan Zhang^{1*}, Jun Liu²

¹ Beijing University of Chemical Technology, China

² Singapore University of Technology and Design, Singapore

{zhaoqh, jiangchen, huwei, zhangf}@mail.buct.edu.cn, jun.liu@sutd.edu.sg

Abstract

Recently, multi-expert methods have led to significant improvements in long-tail recognition (LTR). We summarize two aspects that need further enhancement to contribute to LTR boosting: (1) More diverse experts; (2) Lower model variance. However, the previous methods didn't handle them well. To this end, we propose More Diverse experts with Consistency Self-distillation (MDCS) to bridge the gap left by earlier methods. Our MDCS approach consists of two core components: Diversity Loss (DL) and Consistency Self-distillation (CS). In detail, DL promotes diversity among experts by controlling their focus on different categories. To reduce the model variance, we employ KL divergence to distill the richer knowledge of weakly augmented instances for the experts' self-distillation. In particular, we design Confident Instance Sampling (CIS) to select the correctly classified instances for CS to avoid biased/noisy knowledge. In the analysis and ablation study, we demonstrate that our method compared with previous work can effectively increase the diversity of experts, significantly reduce the variance of the model, and improve recognition accuracy. Moreover, the roles of our DL and CS are mutually reinforcing and coupled: the diversity of experts benefits from the CS, and the CS cannot achieve remarkable results without the DL. Experiments show our MDCS outperforms the state-of-the-art by 1% ~ 2% on five popular long-tailed benchmarks, including CIFAR10-LT, CIFAR100-LT, ImageNet-LT, Places-LT, and iNaturalist 2018. The code is available at <https://github.com/fistyee/MDCS>

1. Introduction

Deep learning has achieved remarkable progress in a range of computer vision (CV) tasks, such as image recog-

*The Corresponding author is with the College of Information Science and Technology and the Interdisciplinary Research Center for Artificial Intelligence, Beijing University of Chemical Technology, China

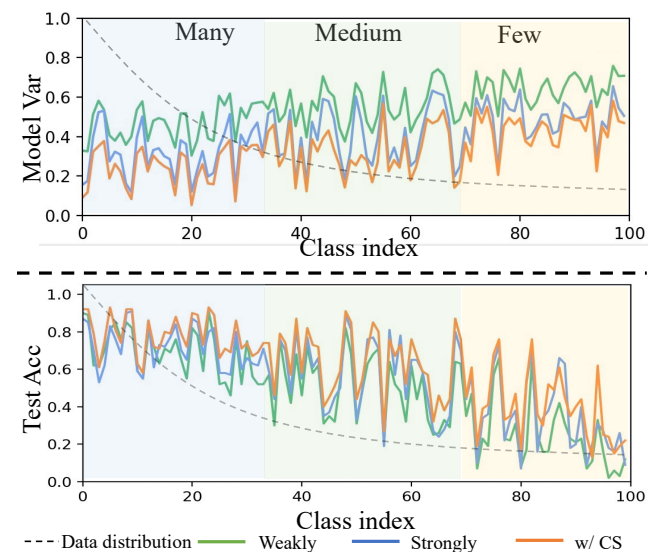


Figure 1. We evaluate a ResNet-32 model based on Balanced Softmax [41, 35] with weakly/strongly augmentation. All experiments were performed with an Imbalanced Factor (IF) of 100 on the CIFAR100-LT dataset. **Top**: model variance [51]. The model trained with weakly augmented instances has a higher variance, whereas the model trained with strongly augmented instances is better than weak augmentation. **Bottom**: Test accuracy. In the case of training on weakly/strongly augmented instances, the model supervised with one-hot labels presents lower accuracy. In contrast, our CS transfers richer knowledge from weakly augmented instances, preventing the model from overfitting instances as well as reducing the model variance and improving recognition accuracy. It indicates that the prediction of the weakly augmented model could provide richer supervision knowledge to the strongly augmented instances better than its one-hot label.

niton [17, 14, 64, 30, 63], object detection [45, 12], and action recognition [46]. Despite advances in deep technologies and computational capability, great success is also highly dependent on well-designed large datasets such as ImageNet [13] and Places [67], where each category has

sufficient and roughly balanced training samples. However, real-world data tends to be long-tailed over semantic categories [60]: a few categories contain many instances (called head categories), while most categories contain only a few instances (called tail categories). Long-tailed recognition (LTR) is challenging because it needs to deal not only with the numerous small data learning problems of the tail categories but also with the extremely unbalanced classification of all categories. Deep models trained with such long-tailed data are usually biased toward head categories on balanced testing data and perform poorly on tail categories.

To address this challenge, many approaches have explored long-tail recognition in order to learn well-performing models from long-tailed data, such as class re-balancing/re-weighting [3, 4, 24, 29, 11, 6, 54, 35], decoupling learning [23] and contrastive learning [56, 22, 49, 68, 10]. Recently, long-tailed recognition methods employing multi-expert ensemble learning [53, 51, 5, 59, 27] have achieved state-of-the-art (SOTA) performance. We summarize two key aspects of these approaches that need further improvement for boosting LTR. (1) Diverse experts focus on different aspects, maximizing the expertise of each [5, 51]. More diversity can support experts in improving LTR. (2) There is a heavy model variance in the prediction of the model, especially for the tail category. So, reducing model variance is essential for LTR. Previous multi-expert methods [53, 51, 5, 59, 27] focused on the above two aspects but did not handle them well. RIDE [51] utilizes a loss to moderate diversity, yet individual experts focus primarily on head categories. ACE and SADE [5, 51, 59] focus on the diverse experts, which learn classification knowledge from sub-categories or dominant categories. However, The "tail category experts" of these methods can greatly suppress head category performance while focusing on the tail categories. Furthermore, these multi-expert methods all employ an ensemble method to reduce the final variance while ignoring the variance of each expert. Among them, NCL [27] introduces strong data augmentation [8] that provides a better generalization of the model. However, there is still a high risk of model variance in its one-hot label supervision for strongly augmented instances. To this end, we design a novel method, namely **More Diverse experts with Consistency Self-distillation (MDCS)**, for long-tailed recognition.

Our proposed MDCS contains two key components, Diversity Loss (DL) and Consistency Self-distillation (CS). Our DL contains an adjustable distribution weight, to cater to the diversity of each expert. By adjusting the distribution weight, each expert tends to recognize different categories, such as Many-shot categories, Medium-shot categories, and Few-shot categories. It is a simple yet effective method for increasing diversity and significantly improving recognition accuracy over previous methods (discussed in Sec.

4). To reduce the model variance and avoid model over-fitting instances, we look forward to providing each expert with a richer form of supervision when learning strongly augmented samples. The label-smoothing regularization [47, 37] is a straightforward way, and further MiSLAS [65] proposes label-aware smoothing for long-tailed recognition. However, the proportion of label-smoothing assignments of these methods is still instance-agnostic, and more reasonable label assignment principles remain to be explored. To this end, we design CS for each expert, which distills richer instance knowledge from predictions of weakly augmented data to regularize strongly augmented instances. Especially for a mini-batch instance, we propose Confident Instance Sampling (CIS) to select the correctly classified instances for consistency self-distillation. In this way, our proposed CIS can prevent CS from introducing biased/noisy knowledge. As illustrated in Fig. 1, the model trained with strong augmentation method [8] could reduce the model variance [51] compared with the model trained with weakly augmented instances (e.g., flipped, cropped). However, our CS trains the model on strongly augmented instances and is supervised by "soft labels" from the predictions of weakly augmented instances, leading to lower model variance and higher recognition accuracy. These "soft labels", produced by prediction on weakly augmented representation, contain more knowledge than their one-hot labels. In addition, the roles of our DL and CS are mutually reinforcing and coupled: (1) Our CS is designed for each expert, which increases the diversity and recognition accuracy of a single expert, and ultimately benefits the ensemble model. (2) Without the DL, the CS cannot achieve remarkable results as the model is biased towards head categories (discussed in Sec. 6).

In the experiments, our proposed MDCS model outperforms state-of-the-art (SOTA) methods by a significant margin on five commonly used benchmark datasets. For instance, on CIFAR-100-LT with an imbalance factor of 100, our approach achieves an accuracy of 56.1%. Similarly, on ImageNet-LT with ResNeXt-50, our model achieves an accuracy of 61.8%, while on iNaturalist 2018 with ResNet-50, we achieve an accuracy of 75.6%.

2. Related Work

Long-tailed Visual Recognition. Conventional methods to alleviate the long-tailed problem are to design re-balancing paradigms that consist of re-sampling and re-weighting. Re-sampling methods, which over-sample tail classes or under-sample head classes, aim to achieve a more balanced data distribution. Re-sampling by simply over-sampling minority classes [3, 4, 38] and under-sampling by abandoning data for dominant classes [21, 16, 3]. However, over-sampling duplicated tailed samples might lead to the over-fitting of minority classes [3]. Under-sampling poten-

tially loses head class information and certainly impairs the generalization ability of the DNNs. Re-weighting methods [24, 29, 11, 6, 54, 35, 42, 55] assign weights to different classes by loss modification or logits adjustment. However, some researchers observed that re-balancing methods will hurt representation learning, and decoupling representation with classifiers will lead to better features. Therefore, two-stage learning was proposed which first trains the model with original data and then fine-tunes the classifier with class-balanced data [33]. Transfer learning is another way to tackle the long-tailed problem, aiming to transfer knowledge learned from majority classes to minority classes. But knowledge transfer methods often need carefully designed structures such as memory bank [31, 50, 32]. More recently, many works try to improve the performance of long-tailed visual recognition by using contrastive learning (CL) strategy [56, 22, 49, 68, 10, 28]. For example, PaCo [10] introduces a set of class-wise learnable center to overcome bias on high-frequency classes of basic supervised contrastive learning (SCL).

Ensemble-based methods, which use multiple experts with aggregation methods, are receiving more and more attention due to their effectiveness on long-tailed recognition. LFME [53] trains different experts with different parts of the dataset and distills the knowledge from these experts to a student model. RIDE [51] optimizes experts jointly with distribution-aware diversity loss and trains a router to handle hard samples. SADE [58] proposed test-time experts aggregating method to handle unknown test class distributions. Recently proposed NCL [27] uses mutual distillation allowing every expert to learn knowledge from others. They still have shortcomings in terms of expert diversity and model variance.

Knowledge Distillation. Knowledge distillation(KD) [2, 18, 39] was proposed for the purpose of model compression by transferring knowledge learned from a powerful teacher model to a student model. KD is performed by supervising the student model with soft labels generated by the teacher model, which also provides better generalization to the student model. KD has gradually evolved from an offline process [40, 18, 39], where the teacher model has trained ahead of time, to an online process [7, 15, 62], where teacher model and student model are trained simultaneously. Unlike offline or online KD, self-distillation [57] assumes that one model can be its own teacher, where the teacher model and student model are identical.

Consistency regularization. Consistency regularization has played a very important role in semi-supervised learning, which was first proposed by Bachman [1] and popularized by Sajjadi [43] and Laine [26]. Consistency regularization utilizes unlabeled data by assuming the model should output the same result when the inputs are similar. Specifically, given two different formations of perturbation

of a training sample, the gap in the output is treated as a loss to train the model. There are various ways to generate perturbed input [36, 48]. A common method is employing two different formations of data augmentation on the same image [44]. FixMatch [44] computes an artificial label for each unlabeled sample by computing the model’s predicted class distribution given a weakly-augmented version. Unlike the above methods, our proposed consistency self-distillation first designs without extra hyper-parameters, combines consistency mechanisms that transfer the richer knowledge of weakly augmented instances to provide more supervision, and employs confidence instance sampling to remove biased/noisy knowledge. Benefiting from these well-designed components, our method effectively reduces the model variance and improves generalization ability.

3. Method

The proposed MDCS consists of two parts, Diversity Loss(DL) and Consistency Self-distillation(CS), aiming to improve the diversity of experts and reduce model variance, respectively. In the following part, we first introduce the preliminaries of long-tailed recognition. Then, we elaborate on our proposed DL and CS. Finally, we show the overall loss of the training process.

3.1. Preliminaries

Long-tail identification attempts to learn a well-represented classification model from a training dataset with a long-tailed class distribution. Formally, let $\mathbb{D}_s = \{(x_i, y_i) | 1 \leq i \leq n_s\}$ be a training set, which x_i is the i -th training sample and $y_i \in \{0, 1\}^C$ is its corresponding one-hot label over C classes. The test set $\mathbb{D}_t = \{(x_i, y_i) | 1 \leq i \leq n_t\}$ is defined in the similar way. Let n_j denote the number of training samples for class j , and let $N = \sum_{j=1}^C n_j$ be the total number of training samples. Without loss of generality, we assume that the classes are decreasingly ordered, i.e., if $i < j$, then $n_i \geq n_j$. Additionally, an imbalanced dataset has significant differences in the class instance numbers, $n_i \gg n_j$.

3.2. Diversity Loss

For training diversity experts, one intuitive approach is to train different experts to focus on different sub-categories. We propose our Diversity Softmax defined as:

$$p(x; \theta) = \frac{n_k^\lambda \exp(v^k)}{\sum_{c=1}^C n_c^\lambda \exp(v^c)}, \lambda \in (-\infty, \infty), \quad (1)$$

where v^k is the class- k output of the model $f(x_i; \theta)$ with parameter θ , and n_k is the number of training samples for category k . The diversity softmax function maps a model’s

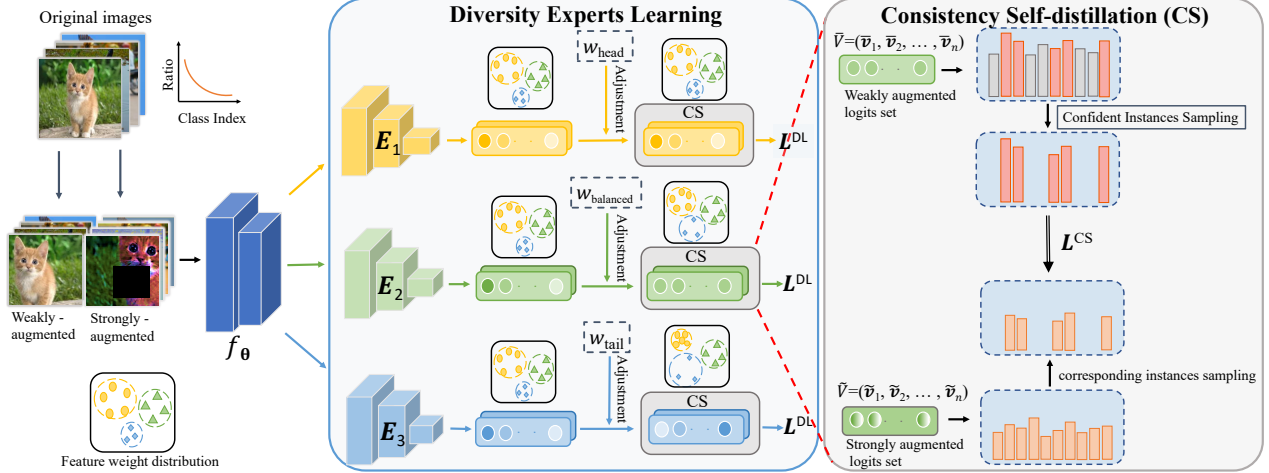


Figure 2. Our method consists of two core components: (i) Diversity Loss (DL), which trains diversity experts; and (ii) Consistency Self-distillation (CS), which reduces the variance of each expert. Firstly, original images with weak and strong augmentation as inputs for shared backbone experts. Secondly, our DL controls experts’ focus on different categories with an adjustable distribution weight (W) to learn different feature weight distributions. Finally, our proposed Consistency Self-distillation distills the richer knowledge from the prediction of weakly augmented instances to address overfitting and reduce model variance. Meanwhile, our proposed Confident Instances Sampling prevents CS from introducing biased knowledge.

class- k output v^k to the probability $p(x; \theta)$. More importantly, the introduced λ acts as a weight distribution parameter for logit adjustment. Fortunately, in our experiments, we discover that it has the effect of generating a reversed weight of long-tailed distribution when $\lambda > 1$. Then, we can employ it to train an expert to improve the accuracy of minority categories with original long-tailed data distribution. Similarly, when $\lambda < 0$, it has the effect of aggravating the imbalance of original long-tailed data, which makes the expert pay more attention to the head categories. When the λ is set to $(0, 1)$, it can weaken the influence of long-tailed distribution [41, 35]. Notably, the aggravation is sensible because our intention is to improve the diversity of all experts in **all** categories. The experiments in Sec. 6 also demonstrate the effect of λ for simulating weight distribution.

With diversity softmax, we propose our Diversity Loss (DL) for diversity experts learning. The DL is defined as:

$$\mathcal{L}^{DL} = \frac{1}{\|\mathbb{D}\|} \sum_{x_i \in \mathbb{D}} -y_i \log \sigma(f(x_i; \theta) + w), \quad (2)$$

where the $\sigma(\cdot)$ is the standard softmax function and w is:

$$w = \lambda \log N^C, \quad (3)$$

where N^C is a list consisting of the number of training samples for each category. In the default setting, we employ DL for training three experts, namely E_1 , E_2 , and E_3 , focusing on Many-shot classes, Medium-shot classes, and Few-shot classes respectively. Visually, Fig. 2 illustrates a multi-expert model with a shared backbone f_θ and three experts

trained with DL. The only difference between the Diverse Loss for different experts is the distribution weight w , such as w_{head} for E_1 , $w_{balance}$ for E_2 , and w_{tail} for E_3 . In general, the structure of our diversity experts learning can have any number of experts $E_\mu (\mu = 1, 2, 3, \dots)$. The effect of the number of experts is shown in the later section 6.

3.3. Consistency Self-distillation

Overall view. In this section, we propose an elegant Consistency Self-distillation (CS) approach to tackle the model variance problem. Our CS method distills richer knowledge from a normal image to a distorted version of the same image. As demonstrated in the left part of Fig. 2, we first construct an original image x_i to two different views denoted as \bar{x}_i and \tilde{x}_i by a weak augmentation (e.g. crop, flip) and a strong argumentation (e.g. RandAug [8]). For different Expert E_μ , we employ diversity softmax to conduct probabilities $\bar{p}(x_i; \theta)$ and $\tilde{p}(x_i; \theta)$ for given (\bar{x}_i, \tilde{x}_i) :

$$\bar{p}(x_i; \theta) = \frac{n_k^\lambda \exp(\bar{v}_i^k / T)}{\sum_{c=1}^C n_c^\lambda \exp(\bar{v}_i^c / T)}, \quad (4)$$

$$\tilde{p}(x_i; \theta) = \frac{n_k^\lambda \exp(\tilde{v}_i^k / T)}{\sum_{c=1}^C n_c^\lambda \exp(\tilde{v}_i^c / T)}, \quad (5)$$

where T is a temperature (a higher T produces a softer probability distribution over categories[18]). Then, our proposed CS employs the Kullback-Leibler(KL) divergence to

perform self-distillation for an instance, which can be formulated as:

$$\mathcal{L}^{CS} = KL(p(\bar{x}_i; \theta) || p(\tilde{x}_i; \theta)). \quad (6)$$

Confident Instance Sampling As our diversity experts specialize in certain categories and may perform poorly in other categories. To prevent CS distills from all instances introducing biased knowledge, we only distill the instances which are correctly classified. Thus, we define Confident Instance set contain all correctly classified instances as:

$$\mathbb{D}_{CI} = \{x_i \in \mathbb{D} | \text{argmax}(p(\bar{x}_i; \theta)) == y_i\}, \quad (7)$$

where y_i is the ground-truth label of instance x_i . Furthermore, we re-formulate the loss of CS with CIS as:

$$\mathcal{L}^{CS} = \frac{1}{|\mathbb{D}_{CI}|} \sum_{x_i \in \mathbb{D}_{CI}} KL(p(\bar{x}_i; \theta) || p(\tilde{x}_i; \theta)) \quad (8)$$

3.4. Model Training

The overall loss in our proposed MDCS consists of two parts, the Loss \mathcal{L}^{DL} of diversity loss and the Loss \mathcal{L}^{CS} of consistency self-distillation with CIS. Finally, we denote the set of Expert as E and formulate the overall loss as:

$$\mathcal{L}_{all} = \sum_{E_\mu \in E} (\mathcal{L}_\mu^{DL} + \alpha \mathcal{L}_\mu^{CS}) \quad (9)$$

where α is a hyperparameter to adjust the weight of Consistency Self-distillation. In addition, we conduct the effect of parameter α in Sec. 6.

4. Method Analysis

4.1. More Diverse Experts

Definition of diversity. According to our empirical analysis, more diverse experts could contribute to the improvement of long-tailed recognition. However, the previous works [59, 51] don't present a measure of diversity. Here, we propose a measure called the diversity factor (σ), defined for a model containing M experts as:

$$\sigma = \bigcup_{\mu=1}^M S_\mu \quad (10)$$

where S_μ is all the correctly classified samples in the test set by Expert E_μ . The S_μ can define as:

$$S_\mu = \{\text{argmax}(p(x_i; \theta_\mu)) == y_i, (x_i, y_i) \in \mathbb{D}_t\} \quad (11)$$

The bigger σ represents greater diversity for the ensemble model.

Model	RIDE [51]							
	ImageNet-LT				CIFAR100-LT			
	Many	Med.	Few	All	Many	Med.	Few	All
E1 Acc	64.3	49.0	31.9	52.6	63.5	44.8	20.3	44.0
E2 Acc	64.7	49.4	31.2	52.8	63.1	44.7	20.2	43.8
E3 Acc	64.3	48.9	31.8	52.5	63.9	45.1	20.5	44.3
Ensemble Acc	68.0	52.9	35.1	56.3	67.4	49.5	23.7	48.0
Ensemble (σ)	76.6	62.9	51.8	60.2	75.8	61.5	37.5	53.1

Model	SADE [59]							
	ImageNet-LT				CIFAR100-LT			
	Many	Med.	Few	All	Many	Med.	Few	All
E1 Acc	68.8	43.7	17.2	49.8	67.6	36.3	6.8	38.4
E2 Acc	65.5	50.5	33.3	53.9	61.2	44.7	23.5	44.2
E3 Acc	43.4	48.6	53.9	47.3	14.0	27.6	41.2	25.8
Ensemble Acc	67.0	56.7	42.6	58.8	61.6	50.5	33.9	49.4
Ensemble (σ)	78.3	62.4	49.3	61.4	75.5	57.1	46.5	59.8

Model	MDCS (ours)							
	ImageNet-LT				CIFAR100-LT			
	Many	Med.	Few	All	Many	Med.	Few	All
E1 Acc	71.9	40.8	12.1	48.9	75.2	37.3	4.1	40.6
E2 Acc	68.2	54.1	36.8	57.1	66.4	51.7	31.4	50.8
E3 Acc	51.8	56.5	58.8	55.7	23.9	37.8	48.2	36.0
Ensemble Acc	72.6	58.1	44.3	61.8	72.4	57.8	35.0	56.1
Ensemble (σ)	81.2	64.6	53.4	65.3	81.8	63.2	55.2	66.9

Table 1. Recognition accuracy (%) and diversity factor (%) of each expert and ensemble model in Many-shot, Medium-shot, and Few-shot categories. The experiment is conducted on CIFAR100-LT with IF = 100. The results show our method outperforms SOTA in terms of both expert diversity and a single expert's accuracy, which are critical to the performance of the ensemble model.

Comparison with SOTA methods. The recognition accuracy and diversity factor results are shown in Table 1, where we compare our result with the prior art, such as RIDE and SADE. The RIDE [51] aims to improve diversity through KL-divergence between experts. However, simply maximizing the KL divergence between experts cannot lead to good diversity and accuracy. SADE is limited to only generating different inversely long-tailed data distributions by adjusting the hyper-parameter in the inverse softmax loss [59] and the accuracy and diversity of Many or Medium-shot is severely inhibited for the "tail category expert," E3. The results show a significant advantage of our method in terms of both diversity and accuracy. The

E1 trained with our DL shows significant improvement in diversity and accuracy in Many-shot categories and similar results in medium-shot categories and Few-shot categories. The E2 and E3 in MDCS also show great strengths in all three shots, which demonstrates the effectiveness of our DL.

The effect of λ for diversity. Table 2 shows different λ used in diversity loss to increase the model diversity. With λ all set to 0, the experts focus on head or tail categories, which gives the model poor diversity. With λ set to $\{1, 1, 1\}$, experts focus on average different categories, the model could get better diversity. With λ set to $\{0, 1, 2\}$ and $\{-0.5, 1, 2.5\}$, experts focus on different categories, the model gets best diversity.

λ	Many	Med.	Few	All
$\{0, 0, 0\}$	85.3	59.2	22.4	55.4
$\{1, 1, 1\}$	80.6	66.8	47.4	64.7
$\{2, 2, 2\}$	54.2	58.4	62.3	58.2
$\{0, 0, 1\}$	84.4	63.9	37.6	61.8
$\{1, 2, 2\}$	70.5	63.3	60.8	64.9
$\{0, 1, 2\}$	80.7	64.4	51.2	65.4
$\{-0.5, 1, 2.5\}$	81.8	63.2	55.2	66.9

Table 2. The effect λ for the three-expert model on the CIFAR100-LT (IF = 100). Different combinations of λ affect the diversity of the model.

4.2. Lower Model Variance

Model variance is the degree of variation in the predictions produced by the same model on different training datasets. With high model variance, the model may perform very differently on different training data, which may indicate that the model overfitted the training data and thus performs poorly in generalizing to unseen data. For n random data sets $\mathbb{B}_{(1)}, \dots, \mathbb{B}_{(m)}$, the k -th models trained on $\mathbb{B}_{(k)}$ will predict $y_{(k)}$ for instance x . The mean predicted value of these models is \bar{y} , which is denote:

$$\bar{y} = \frac{1}{m} \sum_{k=1}^m y_{(k)}, \quad (12)$$

and the model variance denotes:

$$\mathbf{Var}(x, f) = \frac{1}{m} \sum_{k=1}^m (y_{(k)} - \bar{y})^2. \quad (13)$$

To establish a benchmark for model variance, we compare our approach against three baseline methods: cRT [23], RIDE [51] and RIDE with label smoothing (LS) [47]. These metrics are evaluated using twenty independently trained models, trained on CIFAR100-LT with 300 samples for class 0 (IF = 100) [51]. In Table 3, compare with cRT,

RIDE and RIDE with LS, our model has better accuracy performance as well as lower model variance. It also suggests that our approach has better generalization than using ensemble model and label smoothing regularization to reduce model variance.

Method	cRT	RIDE	RIDE + LS	MDCS (Ours)
Var	0.50	0.42	0.41	0.36
Acc	36.4	40.5	41.3	46.1

Table 3. Comparison of mean accuracy and variance of baselines and our MDCS based on CIFAR100-LT. The experiment settings follow RIDE [51].

We also conduct experiments to show the effect of our proposed method on the model variance. Table 4 shows that the model with our consistency self-distillation (CS) effect reduces the model variance of each expert for the Many-, Medium-, and Few-shot subsets.

method	Many-shot	Medium-shot	Few-shot	All
w/o CS	0.28	0.42	0.49	0.40
w/ CS	0.24	0.38	0.46	0.36

Table 4. The effect of CS on the model variance for the three-expert model on the CIFAR100-LT (IF = 100).

5. Experiments

In this section, we perform experiments on five widely used datasets in long-tailed recognition, including CIFAR100/10-LT [25], ImageNet-LT [31], Places-LT [31], and iNaturalist 2018 [20]. After that, we conduct ablation experiments on the CIFAR100-LT and ImageNet-LT datasets to gain further insights.

5.1. Dataset

CIFAR100/10-LT. CIFAR100/10-LT is the long-tailed version of CIFAR100/10 [25]. CIFAR-100/10 contains 50,000 images for training and 10,000 images for the validation of size 32×32 with 100/10 classes. Following [51, 58], we use the same long-tailed version for a fair comparison. The imbalanced factor (IF) β is defined by $\beta = N_{max}/N_{min}$, and this reflects the degree of imbalance in the data. The imbalance factors used in the experiment are set to 100 and 50.

ImageNet-LT and Places-LT. ImageNet-LT and Places-LT are the long-tailed versions of the dataset ImageNet-2012 [13] and the large-scale scene classification dataset Places [67] proposed by Liu [31]. We follow their work by conducting the same dataset by sampling subsets following the Pareto distribution with the power value $\gamma = 6$. Overall, ImageNet-LT has 115.8K images from one thousand categories with an imbalanced factor $\beta = 1280/5$. Places-LT

contains 184.5K images from 365 categories with imbalanced factor $\beta = 4980/5$.

iNaturalist 2018. iNaturalist [20] is a large-scale real-world dataset for long-tailed recognition, which suffers from extremely imbalanced distribution. It contains 437.5K training images and 24.4K validation images from 8142 categories. In addition, the fine-grained problem makes it more challenging [52]. Moreover, we follow the works [27, 59] to divide classes into Many-shot (with more than 100 images), Medium-shot (with 20 - 100 images), and Few-shot (with less than 20 images) parts and report the results on each part.

5.2. Implementation Details

Ensemble method. The final ensemble is average across the experts.

Architecture and settings. We use the same setup for all the baselines and our method. Specifically, following previous work [51, 27, 58], we employ ResNet-32 for CIFAR100/10-LT, ResNeXt-50/ResNet-50 for ImageNet-LT, ResNet-152 for Places-LT and ResNet-50 for iNaturalist 2018 as backbones, respectively. Moreover, we adopt the cosine classifier for prediction on all datasets. If not specified, we use the SGD optimizer with a momentum of 0.9 and set the initial learning rate as 0.1 with linear decay. We set $\lambda = \{-0.5, 1, 2.5\}$ and $\alpha = 0.6$ in our method for all benchmarks. The results of our comparison method are taken from their original paper and our results are averaged over three experiments. More implementation details about epochs we have marked in the comparison table and the hyper-parameter statistics are reported in Appendix.

Augmentation. Our proposed CS utilizes weakly-augmented view and strongly-augmented view to conduct self-distillation. On the CIFAR10/100-LT dataset, the weak augmentation includes crop, horizontal flip, and rotation. The strong augmentation uses CIFAR10Policy besides the basic augmentation. For ImageNet-LT, Places-LT, and iNaturalist, we use cropping, horizontal flipping, rotation, and ColorJitter as weak augmentation. For a fair comparison, we utilize RandAug [8] as strong augmentation for ImageNet-LT and iNaturalist 2018. We add RandomGrayscale and Gaussian Blur to the basic data augmentation composing strong augmentation for Places-LT following previous work [27].

5.3. Comparisons with SOTA on Benchmarks

Long-Tailed CIFAR-100 and CIFAR-10. The comparison results between MDCS and other methods on long-tailed CIFAR datasets are shown in Table 5. We conduct experiments on CIFAR100-LT and CIFAR10-LT with imbalance factors of 100 and 50. Additionally, for fairness, we compare results for 200 epochs and 400 epochs respectively. Our MDCS significantly outperforms the previous

Method	CIFAR100-LT		CIFAR10-LT	
	100	50	100	50
Imbalance Factor				
200 epochs				
CB Focal loss [6]	39.6	45.1	74.5	79.2
LDAM+DRW [6]	42.0	46.6	77.1	81.0
BBN[66]	42.5	47.0	79.8	81.1
LFME[53]	42.3	-	-	-
CAM[61]	47.8	51.7	80.0	83.6
Logit Adj.[35]	43.9	-	77.7	-
Hybrid-SC[49]	46.7	51.8	81.4	85.3
RIDE[51]	49.1	-	-	-
ResLT[9]	48.2	52.7	82.4	85.3
SADE[58]	49.8	53.9	-	-
BCL†[68]	51.9	56.5	84.3	87.2
MDCS†(Ours)	53.2	57.2	85.8	89.4
400 epochs				
ACE[5]	49.6	51.9	81.4	84.9
BSCE†[41]	50.6	55.0	84.0	85.8
PaCo†[10]	52.0	56.0	-	-
SADE†[58]	52.2	57.3	-	-
NCL†[27]	54.2	58.2	85.5	87.3
MDCS†(Ours)	56.1	60.1	87.2	88.3

Table 5. Comparisons on CIFAR100-LT and CIFAR10-LT datasets with the IF of 100 and 50. †denotes models trained with RandAugment[8].

method on all groups, including 56.1% on the CIFAR100-LT dataset with an IF of 100 when trained for 400 epochs.

ImageNet-LT, Places-LT, and iNaturalist 2018. Table 6, 7, and 8 list the Top-1 accuracy of SOTA methods utilizing different backbones on ImageNet-LT, Places-LT, and iNaturalist 2018, respectively. We report the overall Top-1 accuracy as well as the Top-1 accuracy on Many-shot, Medium-shot, and Few-shot groups for Place-LT, and iNaturalist 2018. For fair comparisons, we report the accuracy results at different epochs and these results are from their origin papers. Compared with prior arts, such as PaCo, BCL, NCL, and SADE, our proposed MDCS achieves SOTA performance in the same setting.

6. Ablation Study and Further Analysis

Simulating weight distributions by λ . To enrich the diversity of each expert, the Diversity Loss employs the hyper-parameter λ to simulate different weight distributions for each expert’s training. Fig. 3 indicates how different λ can affect the accuracy of Many-shot, medium-shot, and few-shot categories and overall accuracy. When λ increases, the accuracy of Many-shot categories decreases while the accuracy of few-shot categories increases, which demonstrates the ability to simulate different weight distributions. Besides, when λ gets high enough, the accuracy of

Method	multi-experts	ResNet-50	ResNeXt-50
180 epochs			
LADE[19]		-	53.0
BBN[66]	✓	48.3	49.3
PaCo†[10]		-	56.0
BCL†[68]		-	57.1
SADE[58]	✓	-	58.8
MDCS†(Ours)	✓	59.3	60.2
400 epochs			
ACE[5]	✓	54.7	56.6
τ -norm[23]		54.5	56.0
BSCE†[41]		55.0	56.2
PaCo†[10]		57.0	58.2
NCL†[27]	✓	59.5	60.5
MDCS†(Ours)	✓	60.7	61.8

Table 6. Comparisons on ImageNet-LT. †denotes models trained with RandAug [8].

Method	Many	Medium	Few	All
OLTR[31]	44.7	37.0	25.3	35.9
τ -norm[23]	37.8	40.7	31.8	37.9
ResLT[9]	39.8	43.6	31.4	39.8
MiSLAS[65]	39.6	43.3	36.1	40.4
BSCE†[41]	-	-	-	40.2
PaCo†[10]	36.1	47.9	35.3	41.2
NCL†[27]	-	-	-	41.8
MDCS†(Ours)	43.1	42.9	36.3	42.4

Table 7. Comparisons on Places-LT, starting from an ImageNet pre-trained ResNet-152 provided by Torchvision [34]. †denotes models trained with RandAug [8].

few-shot classes will decrease. This is due to the few-shot group having 30 categories on CIFAR100-LT, this categorization is not fine-grained enough to cater to the effect of the extremely inversely long-tailed distribution generated.

Influence of data augmentations. The RandAug [8] is widely employed as its strong generalization for long-tailed recognition [27, 68, 10]. In this subsection, we conduct different augmentations on training samples to evaluate the effectiveness of weak-strong consistency self-distillation. The results are shown in Table 9, where we compare weak-strong distillation with weak-weak distillation and strong-strong distillation. The results of weak-strong consistency self-distillation are better than strong-strong distillation, which demonstrates that the excellence of our structure does not depend entirely on RandAugment.

Impact of a different number of experts. Our proposed MDCS is a multi-expert framework. The num-

Method	Many	Medium	Few	All
100 epochs				
BBN[66]	49.4	70.8	65.3	66.3
τ -norm[23]	65.6	65.3	65.9	65.2
BCL†[68]	-	-	-	71.8
MDCS†(Ours)	71.8	73.1	72.4	72.5
200 epochs				
CE	68.1	41.5	14.0	48.2
RIDE[51]	70.5	73.7	73.3	73.2
SADE[58]	74.5	72.5	73.0	72.9
400 epochs				
ACE[5]	-	-	-	72.9
BSCE†[41]	72.3	72.6	71.7	71.8
PaCo†[10]	70.3	73.2	73.6	73.2
SADE†[58]	75.5	73.7	75.1	74.5
NCL†[27]	72.7	75.6	74.5	74.9
MDCS†(Ours)	76.5	75.5	75.2	75.6

Table 8. Comparisons on iNaturalist 2018. †denotes models trained with RandAugment[8].




Figure 3. Experiments on the effect of λ to simulate weight distributions. The results demonstrate λ has the ability to simulate different weight distributions focusing on Many-shot, Medium-shot, and Few-shot.

ber of experts can be easily extended by adjusting λ for different experts. We conduct experiments to demonstrate the power of multiple experts. As shown in Fig. 4 (a), the performance of MDCS tends to get better when the number of experts increases. The λ for the one-expert model is $\{1\}$, for the two-expert model is $\{-0.5, 2.5\}$, for the three-expert model is $\{-0.5, 1, 2.5\}$, for the four-expert model is $\{-0.5, 0, 1, 2.5\}$, for the

View1	View2	Top1-Acc
Weak augmentation	Weak augmentation	50.7
Strong augmentation	Strong augmentation	54.6
Weak augmentation	Strong augmentation	56.1

Table 9. Comparisons of training the model with weak-weak augmentation, strong-strong augmentation, and weak-strong augmentation. Experiments are conducted on CIFAR100-LT with IF = 100.

five-expert model is $\{-0.5, 0, 1, 2, 2.5\}$, for the six-expert model is $\{-1, -0.5, 0, 2, 2.5, 3\}$, and for the seven-expert model is $\{-1, -0.5, 0, 1, 2, 2.5, 3\}$.

The rule for setting λ . The ensemble models are not sensitive to the hyper-parameter λ within a reasonable range, so we can easily choose λ just to spread across this range. When the number of branches of experts increases, we first average divide experts into the three groups to set λ . For the head group, the $\lambda \in [-1, 0.5]$, for the balanced group, the $\lambda \in (0.5, 1.5)$, for the tail group, the $\lambda \in [1.5, 3]$. when the values of λ for different experts fell within the above three ranges, the multi-expert model exhibits effective performance improvements. For example, we set $\{-0.5, 1, 2.5\}$ for three experts and $\{-1, -0.5, 0, 1, 2, 2.5, 3\}$ for seven experts.

Influence of loss weight α . The α is an adjusted loss weight of Consistency Self-distillation to control the contribution of the CS part in total loss. To find an appropriate for α , A series of values experimented on the CIFAR100-LT dataset. As shown in Fig. 4 (b), the best performance is achieved when $\alpha = 0.6$. The best result means a balance of supervised learning and self-distillation.

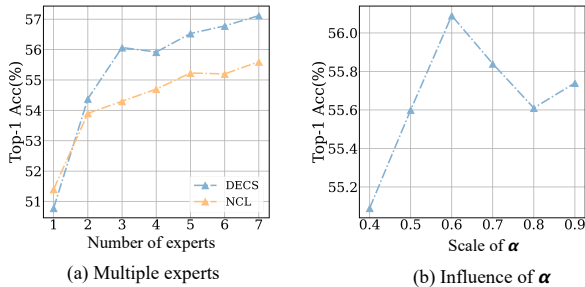


Figure 4. (a) Comparison of using different number of experts with NCL [27]. We report the performance of all categories. When the number of experts increases, the model’s performance also tends to improve. (b) the loss of weight of Consistency Self-distillation. The best result is achieved when $\alpha = 0.6$.

Ablation studies on all components. In this subsection, we conduct detailed ablation studies on the CIFAR100-LT dataset to analyze every component of our MDCS. As shown in Table 10, we evaluate the proposed components including Diverse Loss (DL), weak-strong augmen-

DL	w/RandAug	CS	Accuracy(%)
✗	✗	✗	47.8
✗	✓	✓	48.5
✓	✗	✗	50.7
✓	✓	✗	54.1
✓	✓	✓	56.1

Table 10. Ablation study on the CIFAR100-LT dataset with an IF of 100. The DL indicates Diverse Loss. w/RandAug means weak-strong augmentation. The ✗ in w/RandAug means conducting weak-weak augmentation.

tation(w/RandAug), Consistency Self-distillation (CS), respectively. The ✗ in DL means we use normal Softmax to conduct experiments and the ✗ in w/RandAug means we employ weak-weak augmentation. As shown in Table, our proposed Diversity Loss improves the performance by 2.9%. It is a core component because, without our DL, the other components are less effective at improving performance. MDCS Employing weak-strong augmentation can improve performance from 50.7% to 54.1%, which proves the strength of RandAug [8, 27]. Eventually, when conducting our proposed CS, the performance is significantly further improved, from 54.6% to 56.1%.

7. Conclusion

In this paper, we propose a novel method, MDCS, to cater to the diversity and variance of multi-expert, leading to improved long-tailed recognition accuracy. Our MDCS contains two core components: (1) diversity loss (DL), which can effectively enhance the diversity of experts. (2) consistency self-distillation (CS), which is a novel self-distillation method for reducing the model variance. Furthermore, we propose confident instance sampling in CS to ensure unbiased knowledge. In analyses and ablation studies, we analyze the effectiveness of our proposed core components through experimental results. Moreover, the roles of our DL and CS are mutually reinforcing and coupled. Experimental evidence shows that our MDCS achieves significant performance over the SOTA methods on five popular benchmarks, including 56.1% (+1.9%) accuracy on CIFAR100-LT with an IF 100, 61.8% (+1.3%) accuracy on ImageNet-LT with ResNeXt-50, and 75.6% (+0.7%) accuracy on iNaturalist 2018 with ResNet-50.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant No.62271034, and in part by the Fundamental Research Funds for the Central Universities under Grant XK2020-03.

References

- [1] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. *Advances in neural information processing systems*, 27, 2014.
- [2] Cristian Bucilua, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541, 2006.
- [3] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- [4] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning*, pages 872–881. PMLR, 2019.
- [5] Jiarui Cai, Yizhou Wang, and Jenq-Neng Hwang. Ace: Ally complementary experts for solving long-tailed recognition in one-shot. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 112–121, 2021.
- [6] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *arXiv preprint arXiv:1906.07413*, 2019.
- [7] Defang Chen, Jian-Ping Mei, Can Wang, Yan Feng, and Chun Chen. Online knowledge distillation with diverse peers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3430–3437, 2020.
- [8] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020.
- [9] Jiequan Cui, Shu Liu, Zhuotao Tian, Zhisheng Zhong, and Jiaya Jia. Reslt: Residual learning for long-tailed recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [10] Jiequan Cui, Zhisheng Zhong, Shu Liu, Bei Yu, and Jiaya Jia. Parametric contrastive learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 715–724, 2021.
- [11] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019.
- [12] Xiyang Dai, Yinpeng Chen, Jianwei Yang, Pengchuan Zhang, Lu Yuan, and Lei Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2988–2997, 2021.
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [14] 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. *arXiv preprint arXiv:2010.11929*, 2020.
- [15] Qiushan Guo, Xinjiang Wang, Yichao Wu, Zhipeng Yu, Ding Liang, Xiaolin Hu, and Ping Luo. Online knowledge distillation via collaborative learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11020–11029, 2020.
- [16] Haibo He and Eduardo A Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.
- [19] Youngkyu Hong, Seungju Han, Kwanghee Choi, Seokjun Seo, Beomsu Kim, and Buru Chang. Disentangling label distribution for long-tailed visual recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6626–6636, 2021.
- [20] Grant Van Horn, Oisin Mac Aodha, Yang Song, Alexander Shepard, Hartwig Adam, Pietro Perona, and Serge J. Belongie. The inaturalist challenge 2017 dataset. *CoRR*, abs/1707.06642, 2017.
- [21] Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5):429–449, 2002.
- [22] Bingyi Kang, Yu Li, Sa Xie, Zehuan Yuan, and Jiashi Feng. Exploring balanced feature spaces for representation learning. In *International Conference on Learning Representations*, 2020.
- [23] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019.
- [24] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems*, 29(8):3573–3587, 2017.
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [26] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.
- [27] Jun Li, Zichang Tan, Jun Wan, Zhen Lei, and Guodong Guo. Nested collaborative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6949–6958, 2022.
- [28] Tianhong Li, Peng Cao, Yuan Yuan, Lijie Fan, Yuzhe Yang, Rogerio S Feris, Piotr Indyk, and Dina Katabi. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6928, 2022.

- [29] Tsung-Yi Lin, Priyal Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021.
- [31] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019.
- [32] Alexander Long, Wei Yin, Thalaiyasingam Ajanthan, Vu Nguyen, Pulak Purkait, Ravi Garg, Alan Blair, Chunhua Shen, and Anton van den Hengel. Retrieval augmented classification for long-tail visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6959–6969, June 2022.
- [33] Eran Malach and Shai Shalev-Shwartz. Decoupling “when to update” from “how to update”. In *Advances in Neural Information Processing Systems*, pages 960–970, 2017.
- [34] Sébastien Marcel and Yann Rodriguez. Torchvision the machine-vision package of torch. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1485–1488, 2010.
- [35] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*, 2020.
- [36] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- [37] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances in neural information processing systems*, 32, 2019.
- [38] Seulki Park, Youngkyu Hong, Byeongho Heo, Sangdoon Yun, and Jin Young Choi. The majority can help the minority: Context-rich minority oversampling for long-tailed classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6887–6896, June 2022.
- [39] Nikolaos Passalis and Anastasios Tefas. Learning deep representations with probabilistic knowledge transfer. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 268–284, 2018.
- [40] Baoyun Peng, Xiao Jin, Jiaheng Liu, Dongsheng Li, Yichao Wu, Yu Liu, Shunfeng Zhou, and Zhaoning Zhang. Correlation congruence for knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5007–5016, 2019.
- [41] Jiawei Ren, Cunjun Yu, Shunan Sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. *arXiv preprint arXiv:2007.10740*, 2020.
- [42] Jiawei Ren, Cunjun Yu, Shunan Sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. *arXiv preprint arXiv:2007.10740*, 2020.
- [43] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems*, 29, 2016.
- [44] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020.
- [45] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14454–14463, 2021.
- [46] Zehua Sun, Qihong Ke, Hossein Rahmani, Mohammed Bennamoun, Gang Wang, and Jun Liu. Human action recognition from various data modalities: A review. *IEEE transactions on pattern analysis and machine intelligence*, 2022.
- [47] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [48] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017.
- [49] Peng Wang, Kai Han, Xiu-Shen Wei, Lei Zhang, and Lei Wang. Contrastive learning based hybrid networks for long-tailed image classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 943–952, 2021.
- [50] Peng Wang, Kai Han, Xiu-Shen Wei, Lei Zhang, and Lei Wang. Contrastive learning based hybrid networks for long-tailed image classification. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 943–952, 2021.
- [51] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella X Yu. Long-tailed recognition by routing diverse distribution-aware experts. *arXiv preprint arXiv:2010.01809*, 2020.
- [52] Xiu-Shen Wei, Peng Wang, Lingqiao Liu, Chunhua Shen, and Jianxin Wu. Piecewise classifier mappings: Learning fine-grained learners for novel categories with few examples. *IEEE Transactions on Image Processing*, 28(12):6116–6125, 2019.
- [53] Liuyu Xiang, Guiguang Ding, and Jungong Han. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *European Conference on Computer Vision*, pages 247–263. Springer, 2020.

- [54] Cihang Xie and Alan Yuille. Intriguing properties of adversarial training at scale. *arXiv preprint arXiv:1906.03787*, 2019.
- [55] Yue Xu, Yong-Lu Li, Jiefeng Li, and Cewu Lu. Constructing balance from imbalance for long-tailed image recognition. In *European Conference on Computer Vision*, pages 38–56. Springer, 2022.
- [56] Yuzhe Yang and Zhi Xu. Rethinking the value of labels for improving class-imbalanced learning. *Advances in neural information processing systems*, 33:19290–19301, 2020.
- [57] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3713–3722, 2019.
- [58] Yifan Zhang, Bryan Hooi, Lanqing Hong, and Jiashi Feng. Test-agnostic long-tailed recognition by test-time aggregating diverse experts with self-supervision. *arXiv preprint arXiv:2107.09249*, 2021.
- [59] Yifan Zhang, Bryan Hooi, Lanqing Hong, and Jiashi Feng. Self-supervised aggregation of diverse experts for test-agnostic long-tailed recognition. *Advances in Neural Information Processing Systems*, 3, 2022.
- [60] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *arXiv preprint arXiv:2110.04596*, 2021.
- [61] Yongshun Zhang, Xiu-Shen Wei, Boyan Zhou, and Jianxin Wu. Bag of tricks for long-tailed visual recognition with deep convolutional neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 3447–3455, 2021.
- [62] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4320–4328, 2018.
- [63] QiHao Zhao, Wei Hu, Yangyu Huang, and Fan Zhang. P-diff+: Improving learning classifier with noisy labels by noisy negative learning loss. *Neural Networks*, 144:1–10, 2021.
- [64] Qihao Zhao, Yangyu Huang, Wei Hu, Fan Zhang, and Jun Liu. Mixpro: Data augmentation with maskmix and progressive attention labeling for vision transformer. In *The Eleventh International Conference on Learning Representations*, 2022.
- [65] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16489–16498, June 2021.
- [66] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9719–9728, 2020.
- [67] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017.
- [68] Jianggang Zhu, Zheng Wang, Jingjing Chen, Yi-Ping Phoebe Chen, and Yu-Gang Jiang. Balanced contrastive learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6908–6917, 2022.