




Constructing Balance from Imbalance for Long-tailed Image Recognition

Yue Xu¹^{*}, Yong-Lu Li^{1,2}^{*}, Jiefeng Li¹, and Cewu Lu¹^{**}

¹ Shanghai Jiao Tong University

² Hong Kong University of Science and Technology
{silicxuyue,yonglu.li,ljf.likit,lucewu}@sjtu.edu.cn

Abstract. Long-tailed image recognition presents massive challenges to deep learning systems since the imbalance between majority (head) classes and minority (tail) classes severely skews the data-driven deep neural networks. Previous methods tackle with data imbalance from the viewpoints of data distribution, feature space, and model design, etc. In this work, instead of directly learning a recognition model, we suggest confronting the bottleneck of head-to-tail bias before classifier learning, from the previously omitted perspective of **balancing label space**. To alleviate the head-to-tail bias, we propose a concise paradigm by progressively adjusting label space and dividing the head classes and tail classes, *dynamically constructing balance from imbalance* to facilitate the classification. With flexible data filtering and label space mapping, we can easily **embed** our approach to most classification models, especially the decoupled training methods. Besides, we find the separability of head-tail classes varies among different features with different inductive biases. Hence, our proposed model also provides a *feature evaluation* method and paves the way for long-tailed **feature** learning. Extensive experiments show that our method can boost the performance of state-of-the-arts of different types on widely-used benchmarks. **Code is available at <https://github.com/silicx/DLSA>.**

Keywords: Image Classification, Long-Tailed Recognition, Normalizing Flows

1 Introduction

Deep learning shows its superiority in various computer vision tasks [53,32,17], especially in balanced data scenarios. Though, real-world data is usually severely imbalanced, following a long-tailed distribution [71,55,34,35], *i.e.*, very few frequent classes take up the majority of data (head) while most classes are infrequent (tail). The highly biased data skews classifier learning and leads to performance drop on tail classes. As shown in Fig. 1, most errors stem from the

* The first two authors contribute equally.

** Cewu Lu is the corresponding author, member of Qing Yuan Research Institute and Shanghai Qi Zhi institute.

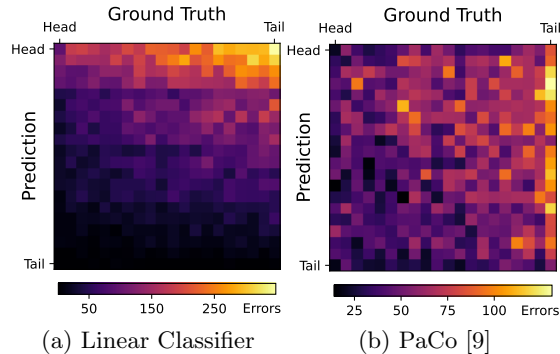


Fig. 1: Confusion matrices of models on ImageNet-LT [38] test set, indicating the severe head-tail bias and that tail classes are particularly prone to confusion with head classes (high density at right-top). **The correct samples are omitted for clarity.** The classes are ordered by their frequency and merged into 20 bins.

head-to-tail bias and a large number of tail samples are misclassified as head classes, even with the very recent long-tail learning technique [9].

Many approaches have been proposed to re-balance long-tail learning by balancing the data distribution [24,15,1,2,6,33], balancing the output logits [41,67], balancing the training losses [19,20,60,3,46,50,54], balancing the feature space [63,26,58,25,9] or with balanced training strategy [27,70,22]. However, as it is the definition of class labels to blame for the long-tailed data distribution, there are very few works tackling long-tailed recognition from the perspective of balancing the *label space*. Samuel *et al* [48] decompose the class labels into semantic class descriptors and exploit its familiarity effect, which is commonly used in zero-shot learning. Wu *et al* [61] reorganize the class labels into a tree hierarchy by realistic taxonomy. Both methods partly alleviate the label imbalance but are restricted by the *class setting* (semantics, hierarchy). Here, we want to dynamically adjust the label space according to the imbalance of *realistic data distribution* to fit long-tail learning.

We can speculate from Fig. 1 that if the head and tail classes are ideally separated, the tail-to-head errors and the complexity of imbalanced learning can be significantly reduced. Thus, we conduct a simple probing to analyze the effect of head-tail separation in Fig. 2 (detailed in supplementary materials: Supp. Sec. 1). We separate the test samples into two groups (head: classes with >50 samples; tail: the rest classes) before individually classifying the groups. Fig. 2 shows that accurate head-tail separation can help long-tail learning. In light of this, we incorporate a **Dynamic Label Space Adjustment (DLSA)** method as shown in Fig. 3. We propose to first confront the *bottleneck* of head-to-tail bias and deal with long-tailed recognition in a **two-stage** fashion: first adjust the label space and separate the head and tail samples, and then apply classification for the divided groups respectively. In virtue of the *inductive bias* of deep learning




Fig. 2: The performance curve of linear classifier on ImageNet-LT [38] with different head-tail label space for imbalanced data and divide separation accuracy (50%: **random** separation; 100%: **ideal** separation). The separation model divides the classes into head and tail, and then we classify within the two individual groups. **Better** separation leads to a **higher** overall accuracy.




Fig. 3: Demonstration of our dynamic label space adjustment (DLSA). H_i stands for a head class and T_j is a tail class. We re-define head and tail classes into two groups progressively to reduce head-to-tail bias, seeking balance from the new data sub-space. DLSA not only separates head and tail but also **re-defines** the label space for the convenience of classifiers. So some of the head samples may be tail classes in new label space due to the \mathcal{L}_{bal} constraint.

feature models (*e.g.*, pretrained backbone), we could define the new label space as the initial clusters of features and approach *head-tail separation*. Moreover, we hope the clusters in the new label space are balanced and contain only a few classes. We formulate these assumptions on the latent space as: 1) **Head-Tail Separability**: the head and tail are separated during the adjustment; 2) **Cluster Balancedness**: the clusters have balanced sizes; 3) **Purity**: samples in each cluster are pure, *i.e.*, belong to as few classes as possible.

Specifically, we propose a *plug-and-play* module which can be embedded in most two-stage learning methods [27]. We use Gaussian mixture to model the pretrained features and produce clusters. Then the samples are divided into two groups and classified independently, where head and tail classes are desired to be separated. In practice, multiple modules can be linked to a cascade model to progressively separate head and tail classes and reduce the imbalance bias. Experiments show that the dataset partition and label re-assignment can effectively alleviate label bias. On top of that, we find backbone models with various inductive biases have different head-tail class separability. With DLSA, we can qualitatively evaluate the feature learning of the backbones from the perspective of head-to-tail bias, which can facilitate long-tailed recognition in practice.

Our main contributions are: 1) Proposing a dynamic label space adjustment paradigm for long-tail learning. 2) Our method is plug-and-play and can boost the performances of the state-of-the-arts on widely-adopted benchmarks. 3) Our method can also act as an evaluation of feature model selection to guide long-tailed feature learning and recognition.

2 Related Work

2.1 Long-Tailed Recognition

Real-world data are long-tailed distributed, which skews machine learning models. Numerous tasks face the challenge of long-tailed data, including object detection [14], attribute recognition [36] and action understanding [4,37]. There is increasing literature trying to alleviate the bias brought by imbalanced data.

Some small adjustments to the model components can alleviate imbalanced learning. The most intuitive approach is re-balance the **data distribution**, either by over-sampling the minority [1,2], or under-sampling of majority [24,15,1]. Data **augmentation** and **generation** [6,33,7] also flattens the long-tailed distribution and helps tail learning. Re-balancing by **loss function** adjust the significance of samples on each class [19,20,60,3,46,50,54]. Instead of weighting the losses, some methods balance the **output logits** [41,29,67] after training.

There are also methods to modify the whole training process. One of the paths is **knowledge transfer** from head classes to tail classes [38,64,27,70,22,30]. Liu *et al* [38] enhance the feature of minority classes with a memory module. Zhang *et al* [64] distill the tail-centric teacher models into the general student network to facilitate tail learning while keeping the head performance. Some works involve specific balanced **training strategy** for imbalanced learning. Kang *et al* [27] exploit the two-stage approach by first training a feature extractor and then re-train a balanced classifier. Zhou *et al* [70] combine uniform sampling and reversed sampling in a curriculum learning fashion. Jamal *et al* [22] estimate class weights with meta-learning to modulate classification loss.

Recently, self-supervised learning is applied in long-tailed recognition for **feature space** re-balancing, among which contrastive learning is the most trendy technique [42,5,16]. Yang and Xu [63] first use self-supervised pretraining to overcome long-tailed label bias. Kang *et al* [26] systematically compare contrastive learning with traditional supervised learning and find that the contrastive model learns a more balanced feature space which alleviates the bias brought by data imbalance. Here, our DLSA provides another strong proof for this conclusion (Sec. 4.4). Wang *et al* [58] combine contrastive feature learning and cross-entropy-based classifier training with a curriculum learning strategy. Jiang *et al* [25] exploit the contrast of the network and its pruned competitor to overcome the forgetting of the minority class. Cui *et al* [9] propose parametric contrastive learning to rebalance the sampling process in contrastive learning.

However, very few works deal with long-tail bias from the perspective of balancing **label space**. Samuel *et al* [48] incorporate semantic class descriptors and enhance minority learning with the familiarity effect of the descriptors. Wu *et al* [61] reorganize the label space into tree hierarchy with semantics. In this paper, we propose DLSA to filter the balance subset from imbalanced data.

2.2 Normalizing Flows

Normalizing flows [47,31,43,12] are a family of invertible networks and are widely adopted in density estimation and generative models. Specifically, a normalizing

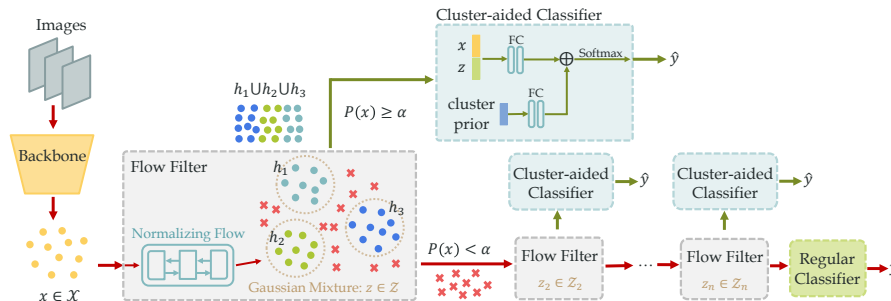


Fig. 4: Method overview. The features x are extracted by pretrained backbones and sent to **Gaussian Mixture Flow Filter**. The Flow Filter maps the samples to a Gaussian mixture latent space \mathcal{Z} with three constraints: **cluster balancedness**, **purity**, and **head-tail separability** to balance the label space and separate head and tail classes. Then we re-group the data according to the likelihood of each sample. The samples that are more conforming to the Gaussian mixture (dots) are classified by **Cluster-aided Classifier** with cluster prior information. The samples with low likelihoods (red crosses) are *progressively* solved by Flow Filters or classified by a regular classifier at the last layer.

flow g is a cascaded transforms data $x \in \mathcal{X}$ to representation u in latent space \mathcal{U} : $u = g^{-1}(x)$. By assuming a tractable latent distribution, we can recover the data distribution with the *change of variable formula*:

$$P(x) = P(u) \cdot |\mathbf{J}(g^{-1}(x))|, \quad (1)$$

where $\mathbf{J}(\cdot)$ is the Jacobian determinant of a transformation and is a crucial concern in the design of normalizing flows due to its high complexity. Rezende and Mohamed [47] propose some cascaded normalizing flows and their basic blocks have linear-time Jacobian computation. Real NVP [12] updates only part of the input vector in the normalizing flow block with a simple bijection, which is more capable of modeling high-dimension data distribution. IAF [31] and MAF [43] are more general than Real NVP, incorporating autoregressive transformations. In practical applications of density estimation, the latent distribution \mathcal{U} is usually unit Gaussian for simplicity. Some works [21] exploit Gaussian mixture to explicitly capture the cluster structure of data. In this work, we also use a unit Gaussian mixture to simultaneously transform data and clusters.

3 Approach

3.1 Preliminaries

Long-tailed recognition is to learn a model that predict labels $y = 1, 2, \dots, \mathcal{C}$ of the data sample $x \in \mathcal{X}$, while the training sample numbers of each class

$n_i = |\{x|y = i, x \in X_{train}\}|$ is highly imbalanced. The imbalance of data are measured by imbalance factor $\beta = \max(n_i)/\min(n_i)$ and the β of typical long-tailed datasets ranges from 10 to 1000.

Since the **head-tail separation** is a well-known key for long-tail learning, our idea is to separate head and tail samples recursively and classify in a divide-and-conquer manner. So the longtailedness is reduced in each partition of data and, in this sense, we construct *balance from imbalance*. We first divide head and tail classes, by re-defining a more balanced label space $h = 1, 2, \dots, K$. As discussed, if we can ideally map the data and labels to a balanced space and separate head and tail classes, it will be easier to learn classifiers. However, it is impractical to achieve an *absolutely* balanced space. So we propose to *progressively* map and cluster samples, and filter out the samples that are well subject to the balanced distributions. During the clustering, we force the model to learn head-tail class separability. In this paradigm, the complexity of long-tail learning is partially *transferred* to the label space learning and latent space mapping process, while these can be more easily settled with the help of the inductive biases of pretrained features and the transformation ability of normalizing flows [47].

Backbone Models. Different inductive biases exist in the various deep learning models and training schemes. Cross-entropy-based supervised learning imposes class separability on deep feature learning. As previously studied [27], for deep supervised learning, the data imbalance is mainly detrimental to the classifiers while the feature learning suffers less. Self-supervised contrastive learning methods [42,5] introduce instance-level transformation invariance to deep models, leading to a more balanced feature since there are no labels involved. Further supervised contrastive learning [63,26] combines the two aspects and incorporates both instance-level multi-view invariance and class-wise contrast. We compare the main-stream pretrained models with different inductive biases in Sec. 4.4 and supervised contrastive learning shows its superiority over the rest.

Normalizing Flow. As introduced in Sec. 2.2, normalizing flows are invertible transformations and are suitable for density estimation and distribution mapping. We utilize normalizing flows to simultaneously map the data x to representations z in a more balanced latent space and estimate the new class label h . The mathematical beauty of normalizing flow enables the maximum likelihood estimation and the learning of data mapping. We also design multiple constraints for the normalizing flows to re-balance the latent space.

3.2 Overview

Fig. 4 depicts the pipeline of our method. We apply our DLSA in feature space and the initial features are extracted by pretrained self-supervised models. There are mainly two types of modules arranged in a cascaded manner. **Gaussian Mixture Flow Filter** (Sec. 3.3) clusters the input data in latent space and assigns new cluster labels. Then, the samples are divided into two groups. The well-clustered samples are filtered out to mitigate the head-to-tail bias and learning complexity, and are sent to a dedicated **Cluster-aided Classifier** (Sec. 3.4)

which exploits the cluster information in classification. The rest outliers are either forwarded to the next normalizing flow or regularly classified.

Specifically, the input features are transformed by normalizing flows to a Gaussian mixture distribution to obtain cluster labels. For better separability in the long-tailed scenario, besides the maximum likelihood objective of Gaussian mixture, the flow model is constrained by three additional **objectives**: 1) Head-tail Separability: head and tail classes should be separated; 2) Cluster Balancedness: the sizes of latent clusters should be as balanced as possible; 3) Purity: the samples in each cluster should belong to as few classes as possible.

3.3 Gaussian Mixture Flow Filter

Gaussian Mixture Flow Filter modules transform data to a desired latent distribution, perform clustering and filter the samples. Without loss of generality, we take the first Flow Filter (magnified in Fig. 4) as an example.

The filter module incorporates a normalizing flow model $g_\theta^{-1} : \mathcal{X} \rightarrow \mathcal{Z}$ with trainable parameter θ , mapping data samples $x \in \mathcal{X}$ to latent distribution $z \in \mathcal{Z}$. To cluster the input samples into K clusters, we assume the latent distribution $P(z)$ is a Gaussian mixture and each component corresponds to a cluster:

$$P(z) = \sum_{k=1}^K P(h = k)P(z|h = k), \quad (2)$$

where h is the random variable of cluster index. $P(z|h = k) = \mathcal{N}(z|\mu_k, I)$ is a Gaussian probability density with mean vectors μ_k and identity covariance matrix. The μ_k are randomly sampled from a normal distribution and fixed during training [21]. A proper assumption on *prior probability* $P(h)$ is required for a tractable optimization methods of NF with GMM. Since $P(h)$ in GMM determines the “size” of a Gaussian component, with the **balancedness** assumptions, the prior probability $P(h)$ can be set to uniform distribution, thus:

$$P(z) = \frac{1}{K} \sum_{k=1}^K P(z|h = k), \quad (3)$$

and the prediction of sample x is given by Bayes’ theorem:

$$P(k|x) = \frac{P(z|h = k)}{\sum_{k'=1}^K P(z|h = k')}, \quad \hat{h} = \operatorname{argmax}_k P(k|x). \quad (4)$$

We train the normalizing flow model with end-to-end gradient descent optimization with the following objectives.

Maximum Likelihood Loss and Head-tail Separability. We propose a weighted maximum likelihood loss to simultaneously train the cluster model and learn head-tail class separation. With the elegant invertible property of

normalizing flow and change of variable formula Eq. 1, we can painlessly obtain the Gaussian mixture likelihood of data sample x :

$$L(\theta; x) = \sum_{k=1}^K \frac{1}{K} \mathcal{N}(g_{\theta}^{-1}(x) | \mu_k, I) |\mathbf{J}(g_{\theta}^{-1}(x))|. \quad (5)$$

The maximum likelihood estimation (MLE) of parameter θ will produce a clustering model. To enhance the separability of head and tail classes of the Flow Filter, we incorporate *sample weighting* and impose higher weights on tail samples than head samples. In DLSA, samples in class i are weighted by:

$$\omega(i) = \frac{n_i^{-q}}{\sum_j n_j^{-q}}, \quad (6)$$

where n_i is the training sample size of class i . q is a positive number usually ranging from 1 to 2, hence the tail samples receive larger weights. The weighted $L(\theta; x)$ will impose higher likelihoods on tail samples than head samples, therefore tail classes are more likely to be filtered out by Flow Filter and separated with head classes. The weighted maximum likelihood loss is constructed with the negative log-likelihood of a data batch \mathcal{B} :

$$\mathcal{L}_{MLE}(\mathcal{B}) = - \sum_{x \in \mathcal{B}} \omega(y) \log L(\theta; x). \quad (7)$$

Training the flow model with a single MLE constraint may result in trivial solutions or a latent distribution that is not suitable for long-tailed recognition, *e.g.*, g_{θ}^{-1} maps all samples to one cluster (Fig. 5 and 6). To avoid *model collapse* and *degeneration*, we introduce two more objectives for the flow filter.

Cluster Balancedness Loss. Though the prior distribution $P(h)$ has been set to uniform, the posterior after observing data samples X :

$$P(h = k|X) = \sum_{x \in X} P(x|X)P(h = k|x) = \sum_{x \in X} \frac{1}{|X|} P(h = k|x), \quad (8)$$

can still be imbalanced; if so, the actual cluster sizes would be biased or even severely long-tailed distributed. To promote cluster balance, we introduce

$$\mathcal{L}_{bal}(X) = \mathbb{E}_h [\log P(h|X)], \quad (9)$$

a cluster balancedness loss which is the negative entropy of the posterior. It is also equivalent to the KL divergence between $P(k|X)$ and a discrete uniform distribution, whose simple proof is shown in Supp. Sec. 2.1. So minimizing the loss will force the cluster sizes to be more even. As shown in Fig. 5, balancedness loss can significantly reduce the unbalancedness of the cluster sizes.

In practice, the optimization of \mathcal{L}_{bal} can be unstable, especially with mini-batch based optimizer [28]. For each mini-batch \mathcal{B}_t at t^{th} iteration, $\hat{p} = P(h = k|\mathcal{B}_t)$ is an unbiased estimator of the posterior $P(h = k|X)$. But small batch



Fig. 5: Cluster sizes of DLSA w/ or w/o \mathcal{L}_{bal} on ImageNet-LT [27].




Fig. 6: Cluster purity of DLSA w/ or w/o \mathcal{L}_{pure} on ImageNet-LT [27].

size can lead to large variance of \hat{p} . Therefore, we propose to exploit the *history information* and incorporate a momentum-based estimator: $\tilde{p}_t = \eta\tilde{p}_{t-1} + (1 - \eta)P(h = k|\mathcal{B}_t)$, where η is decay factor. \tilde{p}_{t-1} is the estimation of last batch and we set \tilde{p}_0 to 0. Thus, the bias-corrected estimator $\frac{\tilde{p}_t}{1 - \eta^t}$ is an unbiased and more efficient estimator than \hat{p} , and is more stable during mini-batch training. The proof of its unbiasedness and efficiency is shown in Supp. Sec. 2.2.

Purity Loss. We define a scalar “purity” $Purity(k)$ as the sample proportion of the largest class in the cluster k . To enhance the class purity in each clusters, we repeatedly randomly sample a pair of samples x_i, x_j from two different classes i, j and suppressing their similarity in cluster predictions:

$$\mathcal{L}_{pure}(x_i, x_j) = \sum_{k=1}^K P(h = k|x_i) \log P(h = k|x_j). \quad (10)$$

By minimizing the purity loss, we increase the cross-entropy between $P(h|x_i)$ and $P(h|x_j)$, thus pushing the two samples farther. Fig. 6 indicates purity loss effectively increases the purity of the clusters. Different sampling strategies can be applied to the purity loss. Our experiments show that the best practice is sampling with equal probability for each class.

Overall, the total loss of Flow Filter module is:

$$\mathcal{L}_{total} = \mathcal{L}_{MLE} + \lambda_{bal}\mathcal{L}_{bal} + \lambda_{pure}\mathcal{L}_{pure}. \quad (11)$$

Among the constraints, \mathcal{L}_{MLE} is for learning the data distribution and its sample-wise weights make the model lean to head samples, enabling head-tail separation with sample likelihood $P(x)$. \mathcal{L}_{pure} ensures that each cluster in GMM belongs to only one class, so we can use cluster information to enhance classification. \mathcal{L}_{bal} reduces the longtailedness in clusters and prevents the model from trivial solutions, *e.g.*, most samples go to one cluster. \mathcal{L}_{pure} and \mathcal{L}_{bal} ensure the training stability and enhance the separability learned by \mathcal{L}_{MLE} .

Filtering The Flow Filter implements the divide step in the DLSA pipeline. After training, some samples do not belong to any cluster and become outliers of the Gaussian mixture. We use a likelihood threshold α to separate the well-clustered samples and outliers. As shown in Fig. 4, the samples with higher

confidence ($P(x) \leq \alpha$, green and blue) are closer to the cluster centers and the rest ($P(x) > \alpha$, red) are outliers of any clusters so they are sent to next adjustment module to further decrease the training difficulty. Accordingly, the samples are divided into two groups to separate head and tail classes.

In practice, we set the threshold α to a quantile of $P(x)$ on training data, so that a certain proportion of data will be filtered out.

3.4 Cluster-Aided Classifier

The high-confidence samples filtered out by the Flow Filter are sent to a Cluster-aided Classifier. Under ideal conditions, each cluster h in the Flow Filter contains only one class y , so the prediction can be directly obtained by a simple label-class mapping. However, the learned Gaussian mixture is rather noisy and besides a majority class y in one cluster, there are some samples from other classes. To compensate for these noise samples, we introduce a softer label-class mapping method. As demonstrated in Fig. 4, for each sample x_i , we first compute cluster prior $P(y|h = \hat{k}_i)$, which is the class frequency of training samples belonging to the sample’s predicted cluster \hat{k}_i . The cluster prior vector and the concatenation of feature x_i and latent representation z_i are forwarded to two independent fully-connected layers (FC). Their outputs are added and we take its Softmax as output probability: $P(y|x, z) = \text{Softmax} \left(FC[x, z] + FC[P(y|h = \hat{k}_i)] \right)$.

4 Experiment

4.1 Datasets

We evaluate our approach on three main-stream benchmarks: ImageNet-LT [38], Places-LT [38], and iNaturalist18 [56]. ImageNet-LT [38] and Places-LT [38] are long-tailed subsets of ImageNet [11] and Places-365 [69], with 1,000 and 365 categories and about 186 K and 106 K images respectively. Both datasets have Pareto distributed train sets and *balanced test sets* and the imbalanced factor β is 256 and 996. iNaturalist18 [56] is a fine-grained image classification dataset which is naturally highly imbalanced ($\beta = 500$), with 8,142 categories and over 437 K images. Following [38], we report overall accuracy on all datasets, and Many-shot (classes with over 100 images), Medium-shot (classes with 20-100 images), Few-shot (classes with fewer than 20 images) accuracy. We also use Matthews correlation coefficient (MCC) [40] and normalized mutual information (NMI) [10] to measure the performance in long-tailed scenario.

4.2 Baselines

We embed and evaluate our models in multiple two-stage [27] long-tail learning methods, with different trending feature learning and classifier learning. The involved methods cover various long-tail learning strategies, including data rebalancing, loss balancing, feature space balancing, and decision boundary adjusting.

Feature Learning. (1) **Supervised Cross-Entropy (CE)**: a baseline backbone trained with vanilla cross-entropy loss. (2) **Class-balance Sampler (CBS)** [52]: a balanced data sampler where each class has equal sampling probability. (3) **PaCo** [9]: supervised contrastive learning model based on MoCo [16] It manages to alleviate the sampling bias in contrastive learning. PaCo is currently the state-of-the-art approach on the two mentioned benchmarks. Following the original implementation, it is integrated with Balanced Softmax [46]. We reproduce the two-stage training version of PaCo and RandAug [8] is not applied in classifier learning for a fair comparison with other features.

Classifier Learning. (1) **BalSoftmax** [46]: an adjusted Softmax-cross-entropy loss for imbalanced learning, by adding $\log(n_j)$ to logit of class j during training, where n_j is the training sample size of class j . (2) **Class-balance Sampler (CBS)** [52]: same as CBS in Feature Learning. (3) **Cosine classifier (Cosine)** [44,13]: predicting according to the cosine distance of features x and class embeddings w_i . Or namely, normalize the features x and weight vector w_i of each class i of a linear classifier: $\hat{y} = \operatorname{argmin}_i \{\cos\langle w_i, x \rangle\}$ (4) **M2M** [30] transfers head information to tail by resampling strategies. (5) **RIDE** [59] trains multiple diversified expert classifier and dynamically select the best expert for each sample. We use RIDE with two experts.

4.3 Implementation Details

We use 2-layered MAF [43] as the Flow Filter model. The Gaussian mixture centers are randomly sampled from $\mathcal{N}(0, 0.05^2)$ and the details of variance selection please refer to Supp. Sec. 3. The model is trained with a SGD optimizer and the decay rate for balancedness loss momentum $\eta = 0.7$. By default, the weight q in the \mathcal{L}_{MLE} is 2.0, and 30% samples are filtered to Cluster-aided Classifiers in each division step. **We reproduce the baselines with decoupled strategy** [27]. The detailed hyper-parameters of each dataset are listed in Supp. Sec. 5.

4.4 Results

We evaluate the methods on some baselines in Tab. 1, 2 and 3. Supp. Sec. 6 shows the detailed results on iNaturalist18. Our method promotes the accuracy of all the baselines. On state-of-art method PaCo [9]+BalSoftmax [46], we achieve accuracy gain of 1.4%, 1.2%, 1.0% on the datasets respectively. Especially, on few-shot classes, the method brings over 1.5% accuracy bonus. Though iNaturalist18 is a large-scale, realistic, and severely imbalanced dataset, DLSA still brings 1% improvement with PaCo representations. DLSA also brings about 1% and 0.5% gain on MCC and NMI score.

Comparison of classifiers. In Tab. 1 and 2, we compare different classifier learning models on same pretrained features (last 10 rows of Tab. 1 and rows beginning with “CE+” in Tab. 2). On each dataset, all the classifier methods use a same trained Flow Filter model since we incorporate two-stage training in DLSA and the learning of Flow Filters only depends on the feature models. DLSA brings comparable improvement on different classifiers (1.2%, 1.4%, 1.2%,

Table 1: Results on ImageNet-LT [38] with ResNet-50 [17] backbone.

Feature	Classifier	Overall	Many	Medium	Few	MCC	MNI
CE	BalSoftmax [46]	42.8	54.1	39.4	23.2	42.7	70.1
	BalSoftmax [46] + DLSA	43.9 (+1.1)	54.5	41.0	24.0	43.8	70.3
CBS [52]	BalSoftmax [46]	42.2	55.8	38.3	17.6	42.2	69.9
	BalSoftmax [46] + DLSA	43.1 (+0.9)	55.3	40.2	18.9	43.1	70.2
PaCo [9]	CBS [52]	54.4	61.7	52.0	42.5	54.4	74.5
	CBS [52] + DLSA	55.6 (+1.2)	62.9	52.7	45.1	55.6	74.9
	BalSoftmax [46]	54.9	67.0	50.1	38.0	54.9	74.9
	BalSoftmax [46] + DLSA	56.3 (+1.4)	67.2	52.1	40.2	56.1	75.4
	Cosine [44,13]	55.7	64.9	53.0	39.5	55.7	75.2
	Cosine [44,13] + DLSA	56.9 (+1.2)	64.6	54.9	41.8	56.8	75.7
	M2M [30]	55.8	67.3	52.1	36.5	55.8	75.3
	M2M [30] + DLSA	56.7 (+0.9)	68.0	52.8	38.2	56.6	75.7
	RIDE [59]	56.5	67.3	53.3	37.3	56.5	75.6
	RIDE [59] + DLSA	57.5 (+1.0)	67.8	54.5	38.8	57.5	75.9

0.9%, 1.0% on ImageNet-LT; 1.2%, 1.2%, 1.0% on Places-LT). The performance improvement of our model mainly depends on the initial structure of the features.

Feature model comparison and evaluation.

With DLSA, models with different feature models have comparable overall accuracy improvement (1.1%, 0.9%, 1.4% with BalSoftmax classifier on ImageNet-LT), but the improvement on few-shot classes differs (0.8%, 1.3%, 2.2% with BalSoftmax classifier). This is due to the different *head-tail separability* of pretrained features with different inductive biases. Thus, we compare the head-tail separation accuracy of DLSA on the three feature models. We analyze the samples filtered out by Flow Filters and compute the accuracy of these samples being tail classes. As discussed in Sec. 1, an ideal separation with 100% accuracy will significantly boost the long-tailed recognition. On PaCo+BalSoftmax, 100% accurate head-tail separation brings $\sim 10\%$ improvement. Fig. 7 shows the accuracy of head-tail classes separation of models with different features. At all three Flow Filter layers, features of PaCo show superior separability than feature models pretrained with CE or CBS, thus PaCo has more potential of enhancing tail recognition. We also evaluate the DLSA on CLIP [45] features of ImageNet-LT. It achieves **65.2%** overall accuracy and over **75%** separation accuracy, surpassing all baselines, showing the **potential** of DLSA in the latest feature spaces. These also strongly support the conclusion of [26], that contrastive learning generates a more balanced feature space and enhances generalization capability. The positive correlation between




Fig. 7: The accuracy of head-tail separation of DLSA on different feature models. The threshold of head/tail is 50.

Table 2: Results on Places-LT [38] with ImageNet [11]-pretrained ResNet-152 [17].

Feature	Overall	Many	Medium	Few	MCC	NMI
Joint (baseline) [27]	30.2	45.7	27.3	8.2	-	-
cRT [27]	36.7	42.0	37.6	24.9	-	-
τ -norm [27]	37.9	37.8	40.7	31.8	-	-
LWS [27]	37.6	40.6	39.1	28.6	-	-
MetaDA [23]	37.1	-	-	-	-	-
LFME [62]	36.2	39.3	39.6	24.2	-	-
FeatAug [7]	36.4	42.8	37.5	22.7	-	-
BALMS [46]	38.7	41.2	39.8	31.6	-	-
RSG [57]	39.3	41.9	41.4	32.0	-	-
DisAlign [65]	39.3	40.4	42.4	30.1	-	-
CE+BalSoftmax [46]	37.8	40.0	40.2	28.5	37.7	58.1
CE+BalSoftmax [46]+ DLSA	39.0 (+1.2)	42.0	39.8	31.3	38.8	58.6
CE+Cosine [44,13]	37.0	39.2	37.7	31.4	36.8	57.8
CE+Cosine [44,13]+ DLSA	38.2 (+1.2)	39.7	39.2	33.3	38.1	58.3
CE+RIDE [59]	41.2	44.4	43.2	31.1	41.1	59.7
CE+RIDE [59]+ DLSA	42.2 (+1.0)	45.8	43.0	33.7	42.0	60.2
PaCo+BalSoftmax [9]	40.9	44.5	42.7	30.4	40.8	59.6
PaCo+BalSoftmax [9]+ DLSA	42.1 (+1.2)	44.4	44.6	32.3	42.0	59.7




Fig. 8: Sample separation process of PaCo [9]+BalSoftmax [46] on ImageNet-LT [38] test set. The **red** and **blue** points are **head** and **tail** samples respectively. The grey boxes are Flow Filters and blue boxes are Cluster-aided Classifiers.

head-tail separability and performance gain indicates that our proposed DLSA is a practical tool for feature analysis and evaluation with long-tailed data.

4.5 Visualization

We visualize our cascade division process in Fig. 8 via t-SNE [39]. We take the PaCo+BalSoftmax as the example. Our method can effectively divide the samples of head and tail classes gradually. The samples filtered out to Cluster-aided Classifiers (in the blue rectangles) are mainly tail samples, while the samples passed to the next Flow Filters (in the grey rectangles) are mixed.

4.6 Ablation Study

We conduct ablation studies and look into the model components on the test set of ImageNet-LT [38] with PaCo [9]+BalSoftmax [46] model, from the following

Table 3: Results on iNaturalist-18 [56] with ResNet-50 [17].

Feature	Overall
LWS [27]	69.5
MetaDA [23]	67.6
Deep-RTC [61]	64.0
Remix [6]	70.5
FeatAug [7]	65.9
smDRAGON [49]	69.1
KCL [26]	68.6
MiSLAS [68]	71.6
DisAlign [65]	70.6
Hybrid-PSC [58]	70.4
FSR [66]	65.5
DRO-LT [51]	69.7
DIVE [18]	71.7
PaCo [9]	71.8
PaCo [9] + Ours	72.8 (+1.0)

Table 4: Ablation study on ImageNet-LT [38] with PaCo [9]+BalSoftmax [46] model and ResNet-50 [17] backbone.

Method	Overall	Many	Medium	Few
Full model	56.3	67.2	52.1	40.2
w/o \mathcal{L}_{MLE}	54.7	66.1	50.8	36.3
w/o \mathcal{L}_{bal}	55.2	65.8	51.2	38.9
w/o \mathcal{L}_{pure}	55.4	66.1	51.5	39.2
300 clusters	55.5	66.1	51.3	39.3
1000 clusters	55.1	66.5	50.3	39.6
2 blocks	55.4	66.3	51.2	39.3
4 blocks	55.8	66.5	51.7	39.2

aspects. The results are reported in Tab. 4. The extended ablation studies on Places-LT are shown in Supp. Sec. 7.

Objectives. We evaluate the three objectives \mathcal{L}_{MLE} , \mathcal{L}_{bal} , \mathcal{L}_{pure} by removing one of them. Removing any loss leads to a significant performance drop, which indicates the necessity of all objectives. Among these losses, w/o \mathcal{L}_{pure} shows the least degradation since the pretrained backbones have moderately learned intra-class similarity and \mathcal{L}_{pure} aim to explicitly enhance the compactness.

Cluster number. Models with less/more clusters perform worse than default 500 clusters. Larger cluster number results in slow training and inference too.

Filter block number. Models with more Flow Filter blocks can separate head and tail classes more finely. But excessive division operations lead to few training samples for the clustering model and classifiers. Due to this trade-off, the default model with 3 blocks outperforms that with 2 or 4 blocks.

5 Conclusions

In this paper, we propose to confront the head-to-tail bias by re-balancing the label space and separating head and tail classes. We present a plug-and-play module DLSA, which automatically adjusts the data distribution and constructs new label space to facilitate the recognition. We embed DLSA in various types of long-tailed recognition state-of-the-arts and boost their performances. We observe that DLSA is also capable of evaluating the different feature learning models. Our future work may extend to combining our paradigm with end-to-end models and more tasks, *e.g.*, object detection and segmentation.

Acknowledgement This work was supported by the National Key R&D Program of China (No. 2021ZD0110700), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), Shanghai Qi Zhi Institute, and SHEITC (2018-RGZN-02046).

References

1. Buda, M., Maki, A., Mazurowski, M.A.: A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks* **106**, 249–259 (2018)
2. Byrd, J., Lipton, Z.: What is the effect of importance weighting in deep learning? In: *International Conference on Machine Learning*. pp. 872–881. PMLR (2019)
3. Cao, K., Wei, C., Gaidon, A., Arechiga, N., Ma, T.: Learning imbalanced datasets with label-distribution-aware margin loss. *arXiv preprint arXiv:1906.07413* (2019)
4. Chao, Y.W., Wang, Z., He, Y., Wang, J., Deng, J.: Hico: A benchmark for recognizing human-object interactions in images. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1017–1025 (2015)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International conference on machine learning*. pp. 1597–1607. PMLR (2020)
6. Chou, H.P., Chang, S.C., Pan, J.Y., Wei, W., Juan, D.C.: Remix: Rebalanced mixup. In: *European Conference on Computer Vision*. pp. 95–110. Springer (2020)
7. Chu, P., Bian, X., Liu, S., Ling, H.: Feature space augmentation for long-tailed data. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*. pp. 694–710. Springer (2020)
8. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. pp. 702–703 (2020)
9. Cui, J., Zhong, Z., Liu, S., Yu, B., Jia, J.: Parametric contrastive learning. *arXiv preprint arXiv:2107.12028* (2021)
10. Danon, L., Diaz-Guilera, A., Duch, J., Arenas, A.: Comparing community structure identification. *Journal of statistical mechanics: Theory and experiment* **2005**(09), P09008 (2005)
11. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
12. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real nvp. *arXiv preprint arXiv:1605.08803* (2016)
13. Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4367–4375 (2018)
14. Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5356–5364 (2019)
15. He, H., Garcia, E.A.: Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering* **21**(9), 1263–1284 (2009)
16. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722* (2019)
17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
18. He, Y.Y., Wu, J., Wei, X.S.: Distilling virtual examples for long-tailed recognition. *arXiv preprint arXiv:2103.15042* (2021)
19. Huang, C., Li, Y., Loy, C.C., Tang, X.: Learning deep representation for imbalanced classification. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5375–5384 (2016)

20. Huang, C., Li, Y., Loy, C.C., Tang, X.: Deep imbalanced learning for face recognition and attribute prediction. *IEEE transactions on pattern analysis and machine intelligence* **42**(11), 2781–2794 (2019)
21. Izmailov, P., Kirichenko, P., Finzi, M., Wilson, A.G.: Semi-supervised learning with normalizing flows. In: *International Conference on Machine Learning*. pp. 4615–4630. PMLR (2020)
22. Jamal, M.A., Brown, M., Yang, M.H., Wang, L., Gong, B.: Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7610–7619 (2020)
23. Jamal, M.A., Brown, M., Yang, M.H., Wang, L., Gong, B.: Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7610–7619 (2020)
24. Japkowicz, N., Stephen, S.: The class imbalance problem: A systematic study. *Intelligent data analysis* **6**(5), 429–449 (2002)
25. Jiang, Z., Chen, T., Mortazavi, B., Wang, Z.: Self-damaging contrastive learning. *arXiv preprint arXiv:2106.02990* (2021)
26. Kang, B., Li, Y., Xie, S., Yuan, Z., Feng, J.: Exploring balanced feature spaces for representation learning. In: *International Conference on Learning Representations* (2020)
27. Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217* (2019)
28. Kiefer, J., Wolfowitz, J.: Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics* pp. 462–466 (1952)
29. Kim, B., Kim, J.: Adjusting decision boundary for class imbalanced learning. *IEEE Access* **8**, 81674–81685 (2020)
30. Kim, J., Jeong, J., Shin, J.: M2m: Imbalanced classification via major-to-minor translation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13896–13905 (2020)
31. Kingma, D.P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., Welling, M.: Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems* **29**, 4743–4751 (2016)
32. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **25**, 1097–1105 (2012)
33. Li, S., Gong, K., Liu, C.H., Wang, Y., Qiao, F., Cheng, X.: Metasaug: Meta semantic augmentation for long-tailed visual recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5212–5221 (2021)
34. Li, Y.L., Liu, X., Wu, X., Li, Y., Qiu, Z., Xu, L., Xu, Y., Fang, H.S., Lu, C.: Hake: A knowledge engine foundation for human activity understanding. *arXiv preprint arXiv:2202.06851* (2022)
35. Li, Y.L., Xu, L., Liu, X., Huang, X., Xu, Y., Wang, S., Fang, H.S., Ma, Z., Chen, M., Lu, C.: Pastanet: Toward human activity knowledge engine. In: *CVPR* (2020)
36. Li, Y.L., Xu, Y., Mao, X., Lu, C.: Symmetry and group in attribute-object compositions. In: *CVPR* (2020)
37. Li, Y.L., Zhou, S., Huang, X., Xu, L., Ma, Z., Fang, H.S., Wang, Y., Lu, C.: Transferable interactiveness knowledge for human-object interaction detection. In: *CVPR* (2019)

38. Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2537–2546 (2019)
39. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(11) (2008)
40. Matthews, B.W.: Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure* **405**(2), 442–451 (1975)
41. Menon, A.K., Jayasumana, S., Rawat, A.S., Jain, H., Veit, A., Kumar, S.: Long-tail learning via logit adjustment. arXiv preprint arXiv:2007.07314 (2020)
42. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
43. Papamakarios, G., Pavlakou, T., Murray, I.: Masked autoregressive flow for density estimation. arXiv preprint arXiv:1705.07057 (2017)
44. Qi, H., Brown, M., Lowe, D.G.: Low-shot learning with imprinted weights. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5822–5830 (2018)
45. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
46. Ren, J., Yu, C., Sheng, S., Ma, X., Zhao, H., Yi, S., Li, H.: Balanced meta-softmax for long-tailed visual recognition. arXiv preprint arXiv:2007.10740 (2020)
47. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: International conference on machine learning. pp. 1530–1538. PMLR (2015)
48. Samuel, D., Atzmon, Y., Chechik, G.: From generalized zero-shot learning to long-tail with class descriptors. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 286–295 (2021)
49. Samuel, D., Atzmon, Y., Chechik, G.: From generalized zero-shot learning to long-tail with class descriptors. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 286–295 (2021)
50. Samuel, D., Chechik, G.: Distributional robustness loss for long-tail learning. arXiv preprint arXiv:2104.03066 (2021)
51. Samuel, D., Chechik, G.: Distributional robustness loss for long-tail learning. arXiv preprint arXiv:2104.03066 (2021)
52. Shen, L., Lin, Z., Huang, Q.: Relay backpropagation for effective learning of deep convolutional neural networks. In: European conference on computer vision. pp. 467–482. Springer (2016)
53. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
54. Sinha, S., Ohashi, H., Nakamura, K.: Class-wise difficulty-balanced loss for solving class-imbalance. In: Proceedings of the Asian Conference on Computer Vision (2020)
55. Spain, M., Perona, P.: Measuring and predicting importance of objects in our visual world (2007)
56. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8769–8778 (2018)

57. Wang, J., Lukasiewicz, T., Hu, X., Cai, J., Xu, Z.: Rsg: A simple but effective module for learning imbalanced datasets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3784–3793 (2021)
58. Wang, P., Han, K., Wei, X.S., Zhang, L., Wang, L.: Contrastive learning based hybrid networks for long-tailed image classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 943–952 (2021)
59. Wang, X., Lian, L., Miao, Z., Liu, Z., Yu, S.X.: Long-tailed recognition by routing diverse distribution-aware experts. arXiv preprint arXiv:2010.01809 (2020)
60. Wang, Y.X., Ramanan, D., Hebert, M.: Learning to model the tail. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. pp. 7032–7042 (2017)
61. Wu, T.Y., Morgado, P., Wang, P., Ho, C.H., Vasconcelos, N.: Solving long-tailed recognition with deep realistic taxonomic classifier. In: European Conference on Computer Vision. pp. 171–189. Springer (2020)
62. Xiang, L., Ding, G., Han, J.: Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In: European Conference on Computer Vision. pp. 247–263. Springer (2020)
63. Yang, Y., Xu, Z.: Rethinking the value of labels for improving class-imbalanced learning. arXiv preprint arXiv:2006.07529 (2020)
64. Zhang, S., Chen, C., Hu, X., Peng, S.: Balanced knowledge distillation for long-tailed learning. arXiv preprint arXiv:2104.10510 (2021)
65. Zhang, S., Li, Z., Yan, S., He, X., Sun, J.: Distribution alignment: A unified framework for long-tail visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2361–2370 (2021)
66. Zhang, Z., Pfister, T.: Learning fast sample re-weighting without reward data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 725–734 (2021)
67. Zhong, Z., Cui, J., Liu, S., Jia, J.: Improving calibration for long-tailed recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16489–16498 (2021)
68. Zhong, Z., Cui, J., Liu, S., Jia, J.: Improving calibration for long-tailed recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16489–16498 (2021)
69. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence* **40**(6), 1452–1464 (2017)
70. Zhou, B., Cui, Q., Wei, X.S., Chen, Z.M.: 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. pp. 9719–9728 (2020)
71. Zipf, G.K.: The psycho-biology of language: An introduction to dynamic philology. Routledge (2013)