

Disentangling Label Distribution for Long-tailed Visual Recognition

Youngkyu Hong* Seungju Han* Kwanghee Choi* Seokjun Seo Beomsu Kim Buru Chang†
Hyperconnect

{youngkyu.hong, seungju.han, kwanghee.choi, seokjun.seo, beomsu.kim, buru.chang}@hpcnt.com

Abstract

The current evaluation protocol of long-tailed visual recognition trains the classification model on the long-tailed source label distribution and evaluates its performance on the uniform target label distribution. Such protocol has questionable practicality since the target may also be long-tailed. Therefore, we formulate long-tailed visual recognition as a label shift problem where the target and source label distributions are different. One of the significant hurdles in dealing with the label shift problem is the entanglement between the source label distribution and the model prediction. In this paper, we focus on disentangling the source label distribution from the model prediction. We first introduce a simple but overlooked baseline method that matches the target label distribution by post-processing the model prediction trained by the cross-entropy loss and the Softmax function. Although this method surpasses state-of-the-art methods on benchmark datasets, it can be further improved by directly disentangling the source label distribution from the model prediction in the training phase. Thus, we propose a novel method, LAbel distribution DisEntangling (LADE) loss based on the optimal bound of Donsker-Varadhan representation. LADE achieves state-of-the-art performance on benchmark datasets such as CIFAR-100-LT, Places-LT, ImageNet-LT, and iNaturalist 2018. Moreover, LADE outperforms existing methods on various shifted target label distributions, showing the general adaptability of our proposed method.

1. Introduction

Based on large-scale datasets such as ImageNet [50], COCO [34], and Places [64], deep neural networks have achieved significant progress in various visual recognition tasks, including classification [22, 52], object detection [47, 18], and segmentation [48]. In contrast to these relatively balanced datasets, real-world data often exhibit

*Equal contribution. Author ordering determined by coin flip.

†Corresponding author.

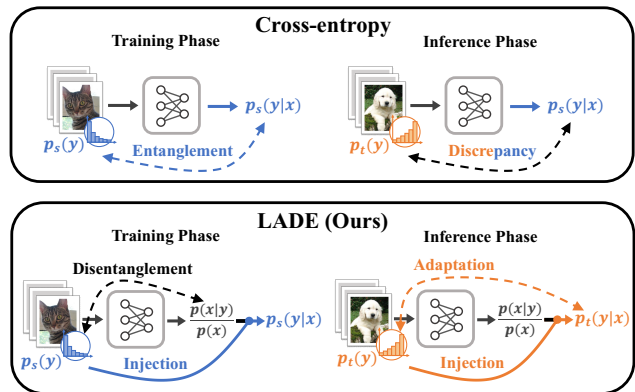


Figure 1: A comparison between the cross-entropy loss and our proposed LADE loss in long-tailed visual recognition. After training with the cross-entropy loss, the model prediction gets entangled with the source label distribution $p_s(y)$, which causes a discrepancy with the target label distribution $p_t(y)$ during the inference phase. Our proposed LADE disentangles $p_s(y)$ from the model prediction so that it can adapt to the arbitrary target probability by injecting $p_t(y)$.

long-tailed distribution where head (major) classes occupy most of the data, while tail (minor) classes have a handful of samples [58, 38]. Unfortunately, the performance of state-of-the-art classification models degrades on datasets following the long-tailed distribution [7, 21, 60].

To tackle this problem, many long-tailed visual recognition methods [7, 21, 25, 8, 51, 62, 9] have been proposed. These methods compare their effectiveness by (1) training on the long-tailed source label distribution $p_s(y)$ and (2) evaluating on the uniform target label distribution $p_t(y)$. However, we argue that this evaluation protocol is often impractical as it is natural to assume that $p_t(y)$ could be the arbitrary distribution such as uniform distribution [50] and long-tailed distribution [17, 10].

From this perspective, we are motivated to explore a new method that adapts the model to the arbitrary $p_t(y)$. In this paper, we borrow the concept of the label distribution shift problems [16, 36, 56] to the long-tailed visual recognition




Figure 2: The average probability for each class calculated on the balanced test set. The model is ResNet-32 [22] trained on CIFAR-100-LT [9], which has uniform target label distribution $p_t(y)$ while the source dataset has long-tailed distribution $p_s(y)$. (a) depicts the average probability when trained with the cross-entropy loss, which shows average probability correlates with $p_s(y)$, resulting in the discrepancy with $p_t(y)$. (b) depicts the average probability when trained and inferred with LADE, which shows the ability of LADE on adapting to $p_t(y)$.

task.

However, it is problematic to directly use the model prediction $p(y|x; \theta)$ which is fitted to the source probability $p_s(y|x)$, as the target probability $p_t(y|x)$ is shifted from $p_s(y)$ to $p_t(y)$ (Figure 1). Figure 2a shows the entanglement between the model prediction and $p_s(y)$ when the model is trained by the cross-entropy (CE) loss and the Softmax function. To alleviate this problem, we focus on disentangling $p_s(y)$ from the model outputs so that the shifted target label distribution $p_t(y)$ can be injected to estimate the target probability.

We shed light on a simple yet strong baseline, called **Post-Compensated Softmax (PC Softmax)** that post-processes the model prediction to disentangle $p_s(y)$ from $p(y|x; \theta)$ and then incorporate $p_t(y)$ to the disentangled model output probability. Despite the simplicity of the method, PC Softmax outperforms state-of-the-art methods in long-tailed visual recognition (will be described in Section 4). Although this observation demonstrates the effectiveness of the disentanglement in the inference phase, PC Softmax can be further improved by directly disentangling $p_s(y)$ in the training phase.

Thus, we propose a novel method, **L**abel distribution **D**isEntangling (**LADE**) loss. LADE utilizes the Donsker-Varadhan (DV) representation [15] to directly disentangle $p_s(y)$ from $p(y|x; \theta)$. Figure 2b shows that LADE disentangles $p_s(y)$ from $p(y|x; \theta)$. We claim that the disentanglement in the training phase shows even better performance on adapting to arbitrary target label distributions.

We conduct several experiments to compare our proposed method with existing long-tailed visual recognition methods, and show that LADE achieves state-of-the-art

performance on benchmark datasets such as CIFAR-100-LT [30], Places-LT [38], ImageNet-LT [38], and iNaturalist 2018 [57]. Moreover, we demonstrate that the classification model trained with LADE can cope with arbitrary $p_t(y)$ by evaluating the performance on datasets with various shifted $p_t(y)$. We further show that our proposed LADE can also be effective in terms of confidence calibration. Our contributions in this paper are summarized as follows:

- We introduce a simple yet strong baseline method, PC Softmax, which outperforms state-of-the-art methods in long-tailed visual recognition benchmark datasets.
- We propose a novel loss called LADE that directly disentangles the source label distribution in the training phase so that the model effectively adapts to arbitrary target label distributions.
- We show that LADE achieves state-of-the-art performance in long-tailed visual recognition on various target label distributions.

2. Related work

2.1. Long-tailed visual recognition

Most long-tailed visual recognition methods can be divided into two strategies: modifying the data sampler to balance the class frequency during optimization [7, 21, 25, 8, 51], and modifying the class-wise weights of the classification loss to increase the importance of tail classes in terms of empirical risk minimization [1, 65, 37, 40, 25, 21, 7].

Both strategies suffer from under-representation of head classes or memorization of tail classes [62, 9]. To overcome these problems, [23, 28, 9, 63] introduce strategies for preventing deteriorated representation learning caused by re-balancing. [60, 38] utilize knowledge from head classes to learn tail classes. [29, 59] augment tail class samples while preserving the diversity of dataset. [24] applies domain adaptation on learning tail class representation.

Recent approaches introduce advanced re-balancing methods for better accommodation of tail class samples. [13] calculates the effective number of samples per each class to re-balance the loss. [9] enforces a greater margin from the decision boundary for tail classes. [55] disentangles feature learning from the confounding effect in momentum by backdoor adjustment.

2.2. Label distribution shift

In this paper, we cope with long-tailed visual recognition as one of the label distribution shift problems. We summarize recently proposed studies in label distribution shift problems that assume $p_s(y) \neq p_t(y)$ and $p_s(x|y) = p_t(x|y)$. [36] estimates the degree of label shift using a black box predictor. [16] extends the work of [36] with

the expectation-maximization algorithm. [2] introduces a domain adaptation to handle label distribution shifts by estimating importance weights. [46] adjusts the logits before applying the Softmax function by the frequency of each class considering the uniform target label distribution.

2.3. Donsker-Varadhan representation

Donsker-Varadhan (DV) representation [15] is the dual variational representation of Kullback-Leibler (KL) divergence [32]. It is proven that the optimal bound of the DV representation is the log-likelihood ratio of two distributions of the KL divergence [3, 4]. The usefulness of the DV representation has been broadly shown in the area including mutual information estimation [4, 35, 45] or generative models [4, 44]. However, [54, 12, 41] have pointed out the instability of directly using the DV representation. To avoid the issue, we use the regularized DV representation from [12] to approximate network logits as the log-likelihood ratio $\log(p(x|y)/p(x))$. To the best of our knowledge, this is the first attempt to utilize the optimal bound inside the DV representation in the long-tailed visual recognition.

3. Method

3.1. Preliminaries

We start by revisiting the most common loss for training the Softmax regression (also known as the multinomial logistic regression) model [5], namely the CE loss:

$$p(y|x; \theta) = \frac{e^{f_\theta(x)[y]}}{\sum_c e^{f_\theta(x)[c]}} \quad (1)$$

$$\mathcal{L}_{CE}(f_\theta(x), y) = -\log(p(y|x; \theta)), \quad (2)$$

where x is the input image and y is the target label, $p_s(x, y)$ and $p_t(x, y)$ are the source (train) and target (test) data distributions, and $f_\theta(x)[y]$ is the logit of class y of the model. The Softmax regression model estimates $p_s(y|x)$ and works well *when the source and the target label distribution are the same*, i.e. $p_s(y) = p_t(y)$.

However, we focus on the *label shift* problem [16, 36] where the target label distribution is shifted from the source label distribution, i.e. $p_s(x|y) = p_t(x|y)$ but $p_s(y) \neq p_t(y)$. Since the model prediction estimates $p_s(y|x)$, it cannot be used to predict the shifted distribution. This is due to the strong coupling between $p_s(y|x)$ and $p_s(y)$, as justified from the Bayes' rule:

$$p_s(y|x) = \frac{p_s(y)p_s(x|y)}{p_s(x)} = \frac{p_s(y)p_s(x|y)}{\sum_c p_s(c)p_s(x|c)}. \quad (3)$$

3.2. PC Softmax: Post-Compensated Softmax

The straightforward way to handle the label distribution shift is by replacing $p_s(y)$ with $p_t(y)$. We introduce a *post-compensation* (PC) strategy that modifies the logit in the inference phase:

Definition 3.1 (Post-Compensation Strategy) *The Post-Compensation strategy modifies model logits as follows:*

$$f_\theta^{PC}(x)[y] = f_\theta(x)[y] - \log p_s(y) + \log p_t(y) \quad (4)$$

where $p_s(y)$ is the distribution which the model logits are entangled with, and $p_t(y)$ is the target distribution that the model tries to incorporate with.

Note that the concept of the PC strategy is not entirely new since [40, 7, 26] previously covered as a different form of multiplying $p_t(y)/p_s(y)$ to the output probability. However, our PC strategy does not violate the categorical probability assumption, i.e. $\sum_c p_t(y = c|x) = 1$.

We apply the PC strategy to the Softmax regression model, which we call Post-Compensated Softmax (PC Softmax). For the Softmax regression model, the PC strategy is the proper adjustment for estimating target data distribution.

Theorem 1 (Post-Compensated Softmax). Let $p_s(x, y)$ and $p_t(x, y)$ be the source and target data distributions, respectively. If $f_\theta(x)[y]$ is the logit of class y from the Softmax regression model estimating $p_s(y|x)$, then the estimation of $p_t(y|x)$ is formulated as:

$$p_t(y|x; \theta) = \frac{\frac{p_t(y)}{p_s(y)} \cdot e^{f_\theta(x)[y]}}{\sum_c \frac{p_t(c)}{p_s(c)} \cdot e^{f_\theta(x)[c]}} \quad (5)$$

$$= \frac{e^{(f_\theta(x)[y] - \log p_s(y) + \log p_t(y))}}{\sum_c e^{(f_\theta(x)[c] - \log p_s(c) + \log p_t(c))}} \quad (6)$$

$$= \frac{e^{f_\theta^{PC}(x)[y]}}{\sum_c e^{f_\theta^{PC}(x)[c]}}. \quad (7)$$

Proof. See the Supplementary Material.

We emphasize that PC Softmax becomes a strong baseline that surpasses previous state-of-the-art long-tailed visual recognition methods. However, recent literature does not consider this as a baseline.

PC Softmax can also be viewed as an extension of Balanced Softmax [46], which modifies the Softmax function to accommodate the uniform target label distribution in the training phase. In contrast, PC Softmax modifies the model logits in the inference phase to match the arbitrary target label distribution $p_t(y)$.

3.3. LADER: Label distribution DisEntangling Regularizer

Performance gain from the PC strategy shows the efficacy of disentangling the source label distribution. However, the PC strategy does not involve the disentanglement in the training phase, which we claim as the ingredient for better adaptability to arbitrary target label distributions. To achieve this, we design a new modeling objective that works

as a substitute for $p_s(y|x)$. We derive the new objective in two steps: (1) detaching $p_s(y)$ from $p_s(y|x)$, which results in $p_s(x|y)/p_s(x)$, and (2) replacing $p_s(y)$ in $p_s(x)$ with the uniform prior $p_u(y)$, i.e. $p_u(y = c) = 1/C$, where C is the total number of classes.

Finally, the modeling objective for the model logits is:

$$f_\theta(x)[y] = \log \frac{p_u(x|y)}{p_u(x)}. \quad (8)$$

We utilize the optimal form of the regularized Donsker-Varadhan (DV) representation [15, 3, 4, 12] to model the log-likelihood ratio above explicitly.

Theorem 2 (Optimal form of the regularized DV representation). Let \mathbb{P}, \mathbb{Q} be arbitrary distributions with $\text{supp}(\mathbb{P}) \subseteq \text{supp}(\mathbb{Q})$. Suppose for every function $T : \Omega \rightarrow \mathbb{R}$ on some domain Ω , the function T that minimizes the regularized DV representation is the log-likelihood ratio of \mathbb{P} and \mathbb{Q} :

$$\log \frac{d\mathbb{P}}{d\mathbb{Q}} = \arg \max_{T: \Omega \rightarrow \mathbb{R}} (\mathbb{E}_{\mathbb{P}}[T] - \log(\mathbb{E}_{\mathbb{Q}}[e^T]) - \lambda(\log(\mathbb{E}_{\mathbb{Q}}[e^T]))^2), \quad (9)$$

for any $\lambda \in \mathbb{R}^+$ when the expectations are finite.

Proof. See Subsection 7.2 from [12].

By plugging $\mathbb{P} = p_u(x|y)$ and $\mathbb{Q} = p_u(x)$ into Equation 9 and choosing the function family of $T : \Omega \rightarrow \mathbb{R}$ to be parametrized by the logits of the deep neural network, the optimal $f_\theta(x)[y]$ approaches to the target objective in Equation 8:

$$\begin{aligned} \log \frac{p_u(x|y)}{p_u(x)} &\geq \arg \max_{f_\theta} (\mathbb{E}_{x \sim p_u(x|y)} [f_\theta(x)[y]] \\ &\quad - \log \mathbb{E}_{x \sim p_u(x)} [e^{f_\theta(x)[y]}]) \\ &\quad - \lambda(\log(\mathbb{E}_{x \sim p_u(x)} [e^{f_\theta(x)[y]}]))^2. \end{aligned} \quad (10)$$

Since the exact estimations of the expectation with respect to $p_u(x|y)$ and $p_u(x)$ are intractable, we use the Monte Carlo approximation [49] using a single batch:

$$\mathbb{E}_{x \sim p_u(x|c)} [f_\theta(x)[c]] \approx \frac{1}{N_c} \sum_{i=1}^N \mathbb{1}_{y_i=c} \cdot f_\theta(x_i)[c] \quad (11)$$

$$\mathbb{E}_{x \sim p_u(x)} [e^{f_\theta(x)[c]}] = \mathbb{E}_{(x,y) \sim p_s(x,y)} \left[\frac{p_u(y)}{p_s(y)} e^{f_\theta(x)[c]} \right] \quad (12)$$

$$\approx \frac{1}{N} \sum_{i=1}^N \frac{p_u(y_i)}{p_s(y_i)} \cdot e^{f_\theta(x_i)[c]}, \quad (13)$$

where x_i and y_i are i -th sample and label, respectively, N is the total number of samples, and N_c is the number of samples for class c . In Equation 12, importance sampling [27]

is used to approximate the expectation with respect to $p_u(x)$ using samples from $p_s(x)$:

$$\frac{p_u(x)}{p_s(x)} = \frac{\sum_c p_u(x|c)p_u(c)}{\sum_c p_s(x|c)p_s(c)} = \frac{p_u(y)}{p_s(y)}, \quad (14)$$

for the sample label pair $(x, y) \sim p_s(x, y)$, where we assume $p_s(x|c) = 0$ for $c \neq y$.

Finally, we derive a novel loss that regularizes the logits to approach Equation 8 by applying Equation 11, 12, and 13 to Equation 10:

Definition 3.2 (LADER) For a single batch of sample-label pairs (x_i, y_i) with $i = 1, \dots, N$, LAbel distribution DisEntangling Regularizer (LADER) is defined as follows:

$$\begin{aligned} \mathcal{L}_{LADERc} &= -\frac{1}{N_c} \sum_{i=1}^N \mathbb{1}_{y_i=c} \cdot f_\theta(x_i)[c] \\ &\quad + \log\left(\frac{1}{N} \sum_{i=1}^N \frac{p_u(y_i)}{p_s(y_i)} \cdot e^{f_\theta(x_i)[c]}\right) \\ &\quad + \lambda\left(\log\left(\frac{1}{N} \sum_{i=1}^N \frac{p_u(y_i)}{p_s(y_i)} \cdot e^{f_\theta(x_i)[c]}\right)\right)^2 \end{aligned} \quad (15)$$

$$\mathcal{L}_{LADER} = \sum_{c \in \mathbb{S}} \alpha_c \cdot \mathcal{L}_{LADERc}, \quad (16)$$

with nonnegative hyperparameters $\lambda, \alpha_1, \dots, \alpha_C$, where C is total number of classes, N_c is the number of samples of class c and \mathbb{S} is the set of classes existing inside the batch.

Empirically, we find out that regularizing the head classes more strongly than the tail classes is more effective. Thus, we apply $\alpha_c = p_s(y = c)$ as the weight for the regularizer of class c , \mathcal{L}_{LADERc} in Equation 16.

3.4. Deriving the conditional probability from disentangled logits

LADER regularizes the logits to be $\log(p_u(x|y)/p_u(x))$ to ensure the logits are explicitly disentangled from the source label distribution $p_s(y)$. To estimate the conditional probability $p_t(y|x)$ of the arbitrary data distribution $p_t(x, y)$ from the regularized logits, we use the modified Softmax function derived from the Bayes' rule with the assumption of $p_t(x|y) = p_u(x|y)$:

$$\begin{aligned} p_t(y|x; \theta) &= \frac{p_t(y)p_t(x|y; \theta)}{\sum_c p_t(c)p_t(x|c; \theta)} \\ &= \frac{p_t(y)p_u(x|y; \theta)}{\sum_c p_t(c)p_u(x|c; \theta)} = \frac{p_t(y) \cdot e^{f_\theta(x)[y]}}{\sum_c p_t(c) \cdot e^{f_\theta(x)[c]}}. \end{aligned} \quad (17)$$

Similar to this, we can estimate $p_s(y|x)$ by swapping $p_t(y)$ of Equation 17 with $p_s(y)$, so that $p_s(y|x; \theta)$ can be optimized by the CE loss. Thus, we can combine LADER with the CE loss as our final loss for training:

Definition 3.3 (LADE) *Label distribution DisEntangling (LADE) loss is defined as follows:*

$$\mathcal{L}_{LADE-CE}(f_{\theta}(x), y) = -\log(p_s(y|x; \theta)) \quad (18)$$

$$= -\log\left(\frac{p_s(y) \cdot e^{f_{\theta}(x)[y]}}{\sum_c p_s(c) \cdot e^{f_{\theta}(x)[c]}}\right) \quad (19)$$

$$\mathcal{L}_{LADE}(f_{\theta}(x), y) = \mathcal{L}_{LADE-CE}(f_{\theta}(x), y) + \alpha \cdot \mathcal{L}_{LADER}(f_{\theta}(x), y), \quad (20)$$

where α is a nonnegative hyperparameter, which determines the regularization strength of \mathcal{L}_{LADER} .

Note that Balanced Softmax [46] is equivalent to LADE with $\alpha = 0$, but LADE is derived from an entirely different perspective of directly regularizing the logits. Furthermore, Balanced Softmax only covers the uniform target label distribution, while our method is designed to cover arbitrary target label distributions without re-training.

In the inference phase, we inject the target label distribution as in Equation 17.

4. Experiments

We compare PC Softmax and LADE with current state-of-the-art methods. First, we evaluate the performance on the uniform target label distribution, which is the prevalent evaluation scheme of long-tailed visual recognition. Then, we assess the performance on variously shifted target label distributions. Finally, we conduct further analysis to show that LADE successfully disentangles the source label distribution and improves confidence calibration. We provide source codes¹ of LADE for the reproduction of the experiments conducted in this paper. Details of the hyperparameter tuning process and the results of the ablation test are reported in the Supplementary Material.

4.1. Experimental setup

Long-tailed dataset We follow the common evaluation protocol [38, 9, 63, 55] in long-tailed visual recognition, which trains classification models on the long-tailed source label distribution and evaluates their performance on the uniform target label distribution. We use four benchmark datasets with at least 100 classes to simulate the real-world long-tailed data distribution: CIFAR-100-LT [9], Places-LT [38], ImageNet-LT [38], and iNaturalist 2018 [57]. We define the *imbalance ratio* as N_{max}/N_{min} , where N is the number of samples in each class. CIFAR-100-LT has three variants with controllable data imbalance ratios 10, 50, and 100. The details of datasets are summarized in Table 1.

¹<https://github.com/hyperconnect/LADE>

Table 1: The details of the training set of long-tailed datasets.

Dataset	# of classes	# of samples	Imbalance ratio
CIFAR-100-LT	100	50K	{10, 50, 100}
Places-LT	365	62.5K	996
ImageNet-LT	1K	186K	256
iNaturalist 2018	8K	437K	500

Comparison with other methods. We compare PC Softmax and LADE with three categories of methods:

- **Baseline methods.** For our baseline, we use Softmax (Equation 1), Focal loss (Focal) [33], OLTR [38], CB-Focal [13], and LDAM [9].
- **Two-stage training.** To demonstrate our method’s efficiency and effectiveness, we compare our method with two-staged state-of-the-art methods that employ a fine-tuning strategy. LDAM+DRW [9] applies a fine-tuning step with loss re-weighting. Decouple [28] re-balances the classifier during the fine-tuning stage.
- **Other state-of-the-art methods.** BBN [63], Causal Norm [55], and Balanced Softmax [46] are recently proposed state-of-the-art methods on long-tail visual recognition. BBN uses an extra additional network branch to deal with an imbalanced training set. Causal Norm utilizes backdoor adjustment to remove indirect causal effect caused by imbalanced source label distribution.

Evaluation Protocol. We report evaluation results using top-1 accuracy. Following [38], for ImageNet-LT and Places-LT, we categorize the classes into three groups depending on the number of samples of each class and further report each group’s evaluation results. The three groups are defined as follows: *Many* covering classes with > 100 images, *Medium* covering classes with ≥ 20 and ≤ 100 images, *Few* covering classes with < 20 images.

4.2. Results on balanced test label distribution

Evaluation results on CIFAR-100-LT, Places-LT, ImageNet-LT, and iNaturalist 2018 are shown in Table 2, 3, 4, and 5, respectively. All the datasets have a uniform target label distribution. PC Softmax shows comparable or better results than the previous state-of-the-art results on benchmark datasets. This result is quite surprising considering the simplicity of PC Softmax. Our proposed method, LADE, achieves even better performance in long-tailed visual recognition on all four benchmark datasets, advancing the state-of-the-art even further.

Table 2: Top-1 accuracy on CIFAR-100-LT with different imbalance ratios. Rows with † denote results directly borrowed from [55]. We use the same backbone network with [55].

Dataset	CIFAR-100 LT		
	100	50	10
Focal Loss†	38.4	44.3	55.8
LDAM†	42.0	46.6	58.7
BBN†	42.6	47.0	59.1
Causal Norm†	44.1	50.3	59.6
Balanced Softmax	45.1	49.9	61.6
Softmax	41.0	45.5	59.0
PC Softmax	45.3	49.5	61.2
LADE	45.4	50.5	61.7

Table 3: The performances on Places-LT [38], starting from an ImageNet pre-trained ResNet-152. Rows with † denote results directly borrowed from [28].

Method	Many	Medium	Few	All
Focal Loss†	41.1	34.8	22.4	34.6
OLTR†	44.7	37.0	25.3	35.9
Decouple- τ -norm†	37.8	40.7	31.8	37.9
Decouple-LWS†	40.6	39.1	28.6	37.6
Causal Norm	23.8	35.8	40.4	32.4
Balanced Softmax	42.0	39.3	30.5	38.6
Softmax	46.4	27.9	12.5	31.5
PC Softmax	43.0	39.1	29.6	38.7
LADE	42.8	39.0	31.2	38.8

CIFAR-100-LT Table 2 shows the evaluation results on CIFAR-100-LT. As shown in the table, in CIFAR-100-LT, LADE outperforms all the baselines over all the imbalance ratios. PC Softmax also shows better performance than other methods except for Balanced Softmax and our proposed LADE.

Places-LT We further evaluate PC Softmax and LADE on Places-LT, and Table 3 shows the experimental results. LADE achieves a new state-of-the-art of 38.8% top-1 overall accuracy, without using a two-stage training as in Decouple- τ -norm [28]. PC Softmax shows yet another promising result by surpassing the previous state-of-the-art, while Softmax offers poor results. This result is quite impressive since both models are the same, and the only difference occurs in the inference phase.

ImageNet-LT We conduct experiments on ImageNet-LT to demonstrate the effectiveness of LADE in the large-scale dataset. We observe that the model is under-fitting at 90 epochs when using LADE. Previous works [28, 63] train model for longer epochs to deal with under-fitting. Thus, we

Table 4: The performances on ImageNet-LT [38]. Rows with § denote results directly borrowed from [55].

Method	Many	Medium	Few	All
<i>90 epochs</i>				
Focal Loss§	64.3	37.1	8.2	43.7
OLTR§	51.0	40.8	20.8	41.9
Decouple-cRT§	61.8	46.2	27.4	49.6
Decouple- τ -norm§	59.1	46.9	30.7	49.4
Decouple-LWS§	60.2	47.2	30.3	49.9
Causal Norm§	62.7	48.8	31.6	51.8
Balanced Softmax	62.2	48.8	29.8	51.4
Softmax	65.1	35.7	6.6	43.1
PC Softmax	60.4	46.7	23.8	48.9
LADE	62.3	49.3	31.2	51.9
<i>180 epochs</i>				
Causal Norm	65.2	47.7	29.8	52.0
Balanced Softmax	63.6	48.4	32.9	52.1
Softmax	68.1	41.9	14.4	48.2
PC Softmax	63.9	49.1	34.3	52.8
LADE	65.1	48.9	33.4	53.0

Table 5: Top-1 accuracy over all classes on iNaturalist 2018. Rows with † denote results directly borrowed from [28] and * denotes the result directly borrowed from [63].

Method	Top-1 Accuracy
CB-Focal†	61.1
LDAM†	64.6
LDAM+DRW†	68.0
Decouple- τ -norm†	69.3
Decouple-LWS†	69.5
BBN*	69.6
Causal Norm	63.9
Balanced Softmax	69.8
Softmax	65.0
PC Softmax	69.3
LADE	70.0

also report the evaluation results at both 90 and 180 epochs, respectively, and this is different from [55] where they train the baseline methods during 90 epochs. Table 4 presents the performance of our method on ImageNet-LT. LADE yields 53.0% top-1 overall accuracy with 180 epochs, which is better than the previous state-of-the-art, including Causal Norm trained with 180 epochs. PC Softmax also shows a favorable result, 52.8%, where it also outperforms the previous state-of-the-art-results. Besides, LADE still achieves the best result compared to the methods when LADE and the methods are trained for 90 epochs.

iNaturalist 2018 To show the scalability of LADE on a large-scale dataset, we evaluate our methods in the real-world long-tailed dataset, iNaturalist 2018. Since iNaturalist 2018 does not contain a validation set, we train the model

Table 6: Top-1 accuracy over all classes on test time shifted ImageNet-LT. All models are trained for 180 epochs.

Dataset	Forward					Uniform	Backward				
Imbalance ratio	50	25	10	5	2	1	2	5	10	25	50
Causal Norm	64.1	62.5	60.1	57.8	54.6	52.0	49.3	45.8	43.4	40.4	38.4
Balanced Softmax	62.5	60.9	58.8	57.0	54.4	52.1	49.6	46.5	44.1	41.4	39.7
Softmax	66.3	63.9	60.4	57.1	52.3	48.2	44.2	38.9	35.0	30.5	27.9
PC Causal Norm	66.7	64.3	60.9	58.1	54.6	52.0	49.8	47.9	47.0	46.7	46.7
PC Balanced Softmax	65.5	63.1	59.9	57.3	54.3	52.1	50.2	48.8	48.3	48.5	49.0
PC Softmax	66.6	63.9	60.6	58.1	55.0	52.8	51.0	49.3	48.8	48.5	49.0
LADE	67.4	64.8	61.3	58.6	55.2	53.0	51.2	49.8	49.2	49.3	50.0

with LADE for 200 epochs and report the test accuracy at 200 epochs, following [28]. Table 5 shows the top-1 accuracy over all classes on iNaturalist 2018. LADE reaches the best accuracy 70.0% among the other methods, even without any branch structure of fine-tuning as BBN [63], nor a two-stage training scheme as [28]. Further, LADE surpasses PC Softmax with a large gap, +0.7%, where PC Softmax still shows a competitive result compared to other methods. This result indicates that PC Softmax is effective for small datasets but performance worsens for larger datasets, while LADE scales well on large datasets as well.

4.3. Results on variously shifted test label distributions

Test sets are rarely well-balanced in real-world scenarios. To simulate and compare the performance of various state-of-the-art methods in the wild, we propose a more realistic evaluation protocol. We first train the model on the long-tailed source label distribution. Then we examine the performance on a range of target label distributions, from distributions that resemble the source label distributions to radically different distributions, similar to [2, 36].

We choose the ImageNet-LT from Section 4.2 as the source dataset. The test set of ImageNet-LT is uniformly distributed, and each class has 50 samples; hence the maximum imbalance ratio is 50. Similar to constructing the CIFAR-LT training dataset [13, 9], we additionally design two types of test datasets. Let us assume ImageNet-LT classes are sorted by descending values of the number of samples per class. Then the shifted test dataset is defined as follows: (1) **Forward**. $n_j = N \cdot \mu^{(j-1)/C}$. As the imbalance ratio increases, it becomes similar to the source label distribution. (2) **Backward**. $n_j = N \cdot \mu^{(C-j)/C}$. The order is flipped so that it gets more different as the imbalance ratio increases. Here μ is the imbalance ratio, N is the number of samples per class in the original ImageNet-LT test set ($= 50$), C is the number of classes, j is the class index and $1 \leq j \leq C$, and n_j is the number of samples in class j for the shifted test set.

We compare LADE with Softmax, Balanced Softmax,

and Causal Norm on this evaluation protocol. For a fair comparison, we apply our PC strategy to state-of-the-art methods: $\log p_t(y) - \log p_u(y)$ is added to the logits before applying the softmax function for Balanced Softmax, and Causal Norm, as both methods target the uninformative prior. Theorem 1 is used for the Softmax instead.

Table 6 shows the top-1 accuracy in the range of test datasets between the Forward- and Backward-type datasets. Our PC strategy shows consistent performance gain, which indicates the benefits of plug-and-play target label distributions. Moreover, LADE outperforms all the other methods in every imbalance settings, and the performance gap between LADE and PC Softmax gets wider as the dataset gets more imbalanced. These results demonstrate the general adaptability of our proposed method on real-world scenarios, where the target label distribution is variously shifted.

4.4. Further analysis

Visualization of the logit values In this subsection, we visualize the logit values of each class in order to demonstrate the effect of LADE. By disentangling the source label distribution with LADE as described in Section 3.3, the logit value $f_\theta(x)[y]$ should converge to $\log C$ for the positive samples:

$$f_\theta(x)[y] = \log \frac{p_u(x|y)}{p_u(x)} = \log \frac{p_u(y|x)}{p_u(y)} = \log C, \quad (21)$$

where we assume perfectly separable case, i.e. $p_u(y|x) = 1$ for $(x, y) \sim p_u(x, y)$ and C is the number of classes.

Figure 3 shows how the logits are distributed for each class. The hyperparameter α represents the regularization strength for LADER. As α increases, the logit values gradually converge to the theoretical value $y = \log C = \log 100$ (dotted line in the figure), reconfirming Theorem 2. This result indicates that LADER successfully regularizes the logit values as we intended.

Confidence calibration Previous literature claims that the confidence of the neural network classifier does not represent its true accuracy [20]. We regard a classifier is well-




Figure 3: ResNet-32 model logits correspond to each class, where the model is trained on CIFAR-100-LT with an imbalance ratio of 100. For each class c , *positive samples* denote the sample corresponds to class c , and *negative samples* denote the other. The colored area denotes the variance of logit values, while the line indicates the mean.




Figure 4: Reliability diagrams of ResNeXt-50-32x4d [61] on ImageNet-LT. The average confidence of the model trained by LADE nearly matches its accuracy.

calibrated when its predictive probability, $\max_y p(y|x)$, represents the true probability [20]. For example, when a calibrated classifier predicts the label y with predictive probability 0.75, it has a 75% chance of being correct. It will be catastrophic if we cannot trust the confidence of the neural network in the domain of medical diagnosis [11] and self-driving car [6].

[42] suspect the endlessly growing logit values induced by a combination of naive CE loss and the Softmax function as the culprit of over-confidence. Since LADER regularizes the logit size, we expect that using LADE prevents the model from being over-confident. Through experiments, we observe that LADE improves calibration. Figure 4 shows the reliability diagrams with 20-bins. Using expected calibration error (ECE) [43], we quantitatively measure the miscalibration rate of the model trained on ImageNet-LT. We compare our LADE with PC Softmax and the current state-of-the-art methods, Causal Norm [55], and Balanced Softmax [46]. Results show that LADE produces a more calibrated classifier than other methods, with the ECE of 0.0346, which confirms our expectation.

5. Conclusion

In this paper, we suggest that disentangling the source label distribution from the model prediction is useful for long-tailed visual recognition. To disentangle the source label distribution, we first introduce a simple yet strong baseline, called PC Softmax, that matches the target label distribution by post-processing the model prediction trained by the cross-entropy loss and the Softmax function. We further propose a novel loss, LADE, that directly disentangles the source label distribution in the training phase based on the optimal bound of Donsker-Varadhan representation. Experiment results demonstrate that PC Softmax and our proposed LADE outperform state-of-the-art long-tailed visual recognition methods on real-world benchmark datasets. Furthermore, LADE achieves state-of-the-art performance on various shifted target label distributions. Lastly, further experiments show that our proposed LADE is also effective in terms of confidence calibration. We plan to extend our research to other vision domain problems that suffer from long-tailed distributions, such as object detection and segmentation.

References

- [1] Rehan Akbani, Stephen Kwek, and Nathalie Japkowicz. Applying support vector machines to imbalanced datasets. In *European conference on machine learning*, pages 39–50. Springer, 2004.
- [2] Kamyar Azizzadenesheli, Anqi Liu, Fanny Yang, and Animesh Anandkumar. Regularized learning for domain adaptation under label shifts. In *International Conference on Learning Representations*, 2019.
- [3] Arindam Banerjee. On bayesian bounds. In *Proceedings of the 23rd international conference on Machine learning*, pages 81–88, 2006.
- [4] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. Mutual information neural estimation. In *International Conference on Machine Learning*, pages 531–540, 2018.
- [5] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [6] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseem Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- [7] 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.
- [8] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning*, pages 872–881, 2019.
- [9] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Advances in Neural Information Processing Systems*, pages 1567–1578, 2019.
- [10] João Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *CoRR*, abs/1808.01340, 2018.
- [11] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. Intelligible models for health-care: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1721–1730, 2015.
- [12] Kwanghee Choi and Siyeong Lee. Regularized mutual information neural estimation. *arXiv preprint arXiv:2011.07932*, 2020.
- [13] 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 Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019.
- [14] 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.
- [15] MD Donsker and SRS Varadhan. Large deviations for stationary gaussian processes. *Communications in Mathematical Physics*, 97(1-2):187–210, 1985.
- [16] Saurabh Garg, Yifan Wu, Sivaraman Balakrishnan, and Zachary C. Lipton. A unified view of label shift estimation. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [17] Krzysztof J. Geras, Stacey Wolfson, S. Gene Kim, Linda Moy, and Kyunghyun Cho. High-resolution breast cancer screening with multi-view deep convolutional neural networks. *CoRR*, abs/1703.07047, 2017.
- [18] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [19] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- [20] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, pages 1321–1330, 2017.
- [21] Haibo He and Edwardo A Garcia. Learning from imbalanced data. *IEEE Transactions on knowledge and data engineering*, 21(9):1263–1284, 2009.
- [22] 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.
- [23] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5375–5384, 2016.
- [24] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. 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*, pages 7610–7619, 2020.
- [25] Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5):429–449, 2002.
- [26] Justin M Johnson and Taghi M Khoshgoftaar. Survey on deep learning with class imbalance. *Journal of Big Data*, 6(1):27, 2019.
- [27] Herman Kahn. Use of different monte carlo sampling techniques. 1955.
- [28] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *International Conference on Learning Representations*, 2019.
- [29] Jaehyung Kim, Jongheon Jeong, and Jinwoo Shin. M2m: Imbalanced classification via major-to-minor translation. In

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13896–13905, 2020.
- [30] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [31] Meelis Kull, Miquel Perello Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [32] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951.
- [33] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [35] Xiao Lin, Indranil Sur, Samuel A Nastase, Ajay Divakaran, Uri Hasson, and Mohamed R Amer. Data-efficient mutual information neural estimator. *arXiv preprint arXiv:1905.03319*, 2019.
- [36] Zachary C. Lipton, Yu-Xiang Wang, and Alexander J. Smola. Detecting and correcting for label shift with black box predictors. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 3128–3136. PMLR, 2018.
- [37] Xu-Ying Liu and Zhi-Hua Zhou. The influence of class imbalance on cost-sensitive learning: An empirical study. In *Sixth International Conference on Data Mining (ICDM’06)*, pages 970–974. IEEE, 2006.
- [38] 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 Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019.
- [39] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [40] Dragos Margineantu. When does imbalanced data require more than cost-sensitive learning. In *Proceedings of the AAAI’2000 Workshop on Learning from Imbalanced Data Sets*, pages 47–50, 2000.
- [41] David McAllester and Karl Stratos. Formal limitations on the measurement of mutual information. In *International Conference on Artificial Intelligence and Statistics*, pages 875–884, 2020.
- [42] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? In *Advances in Neural Information Processing Systems*, pages 4694–4703, 2019.
- [43] Mahdi Pakdaman Naeini, Gregory F Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence*, volume 2015, page 2901. NIH Public Access, 2015.
- [44] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pages 271–279, 2016.
- [45] Ben Poole, Sherjil Ozair, Aaron van den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *ICML*, 2019.
- [46] Jiawei Ren, Cunjun Yu, shunan sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4175–4186. Curran Associates, Inc., 2020.
- [47] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [48] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [49] Reuven Y Rubinstein and Dirk P Kroese. *Simulation and the Monte Carlo method*, volume 10. John Wiley & Sons, 2016.
- [50] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [51] Li Shen, Zhouchen Lin, and Qingming Huang. Relay back-propagation for effective learning of deep convolutional neural networks. In *European conference on computer vision*, pages 467–482. Springer, 2016.
- [52] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [53] Jasper Snoek, Yaniv Ovadia, Emily Fertig, Balaji Lakshminarayanan, Sebastian Nowozin, D. Sculley, Joshua V. Dillon, Jie Ren, and Zachary Nado. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13969–13980, 2019.

- [54] Jiaming Song and Stefano Ermon. Understanding the limitations of variational mutual information estimators. In *International Conference on Learning Representations*, 2019.
- [55] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *Advances in Neural Information Processing Systems*, 33, 2020.
- [56] Junjiao Tian, Yen-Cheng Liu, Nathaniel Glaser, Yen-Chang Hsu, and Zsolt Kira. Posterior re-calibration for imbalanced datasets. *Advances in Neural Information Processing Systems*, 33, 2020.
- [57] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8769–8778, 2018.
- [58] Grant Van Horn and Pietro Perona. The devil is in the tails: Fine-grained classification in the wild. *arXiv preprint arXiv:1709.01450*, 2017.
- [59] Xinyue Wang, Yilin Lyu, and Liping Jing. Deep generative model for robust imbalance classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14124–14133, 2020.
- [60] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *Advances in Neural Information Processing Systems*, pages 7029–7039, 2017.
- [61] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [62] Yuzhe Yang and Zhi Xu. Rethinking the value of labels for improving class-imbalanced learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- [63] 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.
- [64] 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.
- [65] Zhi-Hua Zhou and Xu-Ying Liu. Training cost-sensitive neural networks with methods addressing the class imbalance problem. *IEEE Transactions on knowledge and data engineering*, 18(1):63–77, 2005.