

# No One Left Behind: Improving the Worst Categories in Long-Tailed Learning

Yingxiao Du, Jianxin Wu\*

State Key Laboratory for Novel Software Technology, Nanjing University  
Nanjing, China, 210023

duyx@lamda.nju.edu.cn, wujx2001@gmail.com

## Abstract

Unlike the case when using a balanced training dataset, the per-class recall (i.e., accuracy) of neural networks trained with an imbalanced dataset are known to vary a lot from category to category. The convention in long-tailed recognition is to manually split all categories into three subsets and report the average accuracy within each subset. We argue that under such an evaluation setting, some categories are inevitably sacrificed. On one hand, focusing on the average accuracy on a balanced test set incurs little penalty even if some worst performing categories have zero accuracy. On the other hand, classes in the “Few” subset do not necessarily perform worse than those in the “Many” or “Medium” subsets. We therefore advocate to focus more on improving the lowest recall among all categories and the harmonic mean of all recall values. Specifically, we propose a simple plug-in method that is applicable to a wide range of methods. By simply re-training the classifier of an existing pre-trained model with our proposed loss function and using an optional ensemble trick that combines the predictions of the two classifiers, we achieve a more uniform distribution of recall values across categories, which leads to a higher harmonic mean accuracy while the (arithmetic) average accuracy is still high. The effectiveness of our method is justified on widely used benchmark datasets.

## 1. Introduction

Various gaps exist when adapting image recognition techniques that are developed in the lab to industrial applications. The most noteworthy one is perhaps the difference between training datasets. Most training datasets used in academic research [6, 16] are balanced with respect to the number of images per class. This should not be taken for granted because datasets used in real-world applications are more likely to be imbalanced. Training deep models

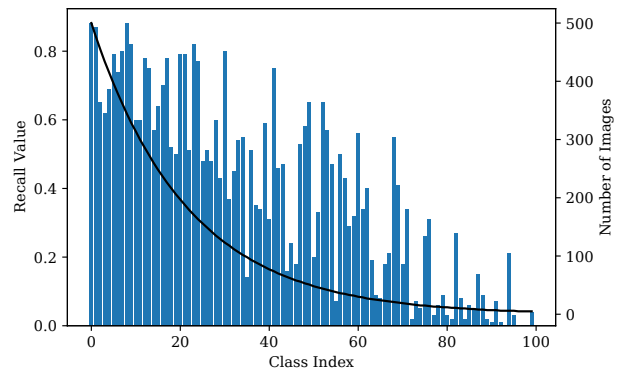


Figure 1. The per-class recall of models trained on the imbalanced CIFAR100 (with imbalance ratio 100). Per-class recall value varies a lot from category to category. Moreover, it is not necessarily true that all categories in the “Few” subset have lower accuracy than those in the “Many” or “Medium” subsets.

Method	Mean Accuracy	Lowest Recall
BSCE [22]	42.24	3.00
DiVE [12]	45.11	2.00
MiSLAS [30]	47.05	5.00
RIDE [26]	48.64	2.00
PaCo [4]	51.24	5.00

Table 1. The lowest per-class recall of various state-of-the-art methods on the imbalanced CIFAR100 (with imbalance ratio 100). Although there are rapid improvements over the mean accuracy, the lowest per-class recall remains very low.

on these datasets is not trivial as models are known to perform poorly on such datasets. Long-tailed recognition is a research field that aims at tackling this challenge.

By plotting the per-class recall (i.e., per-class accuracy) of the model trained on the imbalanced CIFAR00 dataset (with imbalance ratio 100) in Fig. 1, we find that the recall varies dramatically from category to category. We argue that in real world applications, *all categories are equally*

\*J. Wu is the corresponding author. This research was partly supported by the National Natural Science Foundation of China under Grant 62276123 and Grant 61921006.

*important and no one should be left behind.* We also find that despite the rapid developments in this community, no obvious improvement over the lowest per-class recall is witnessed in the past few years [4, 12, 22, 26, 30], as is shown in Tab. 1. The convention in long-tailed recognition research is to split the classes into three subsets based on the number of training images. The accuracy within each subset is often reported along with the overall accuracy. While this evaluation scheme seems reasonable at the first glance, we argue it is potentially problematic. First of all, computing the average recall within each subset is way too coarse, making it impossible to reflect whether some classes are completely sacrificed but covered up by other “easy” tail classes. What’s more, we find it is not necessarily true that classes in the “Few” category all have lower recall than classes in the other two subsets, as is shown in Fig. 1. Therefore, focusing on improving the mean accuracy alone is not enough, especially in real-world applications.

In this paper, we propose a novel method to make sure that no category is left behind. Specifically, we argue that although mean accuracy is widely used in image classification as an optimization objective, due to the fact that different classes have very different recall in long-tailed recognition, it is not the most suitable objective as it incurs little penalty even if some categories have very small per-class accuracy values (e.g., close to 0). Hence, to improve even the worst-performing categories, we believe the harmonic mean of per-class recall would be a better objective. Since harmonic mean is very sensitive to small numbers, further improving classes that already have high recall brings little benefits to the harmonic mean, which makes sure no single class will be left behind. Also, now all classes are treated equally, forcing us to improve classes that have low recall no matter which subset it belongs to.

However, it is difficult to directly minimize the harmonic mean. We therefore propose a novel loss function that maximizes the geometric mean instead, which can be viewed as a surrogate. Our method serves as a simple plug-in that can be used together with both baseline and various state-of-the-art methods. We also propose an ensemble trick that uses the pre-trained and fine-tuned models together to make predictions during inference with nearly no extra cost. We are able to yield a more uniform distribution of recall across categories (i.e., no one left behind), which achieves higher harmonic mean of the recall while the (arithmetic) mean accuracy remains high, too.

In summary, our work has the following contributions:

- We are the first to emphasize the importance of the correct recognition of all categories in long-tailed recognition.
- We propose a novel method that aims at increasing the harmonic mean of per-class recall as well as an ensemble

trick that combines two existing models together during inference with nearly no extra cost.

- We experimented on three widely used benchmark datasets, which justify the effectiveness of our method in terms of both overall and worst per-class accuracy.

## 2. Related Work

Long-tailed recognition is a research field that aims at training models using imbalanced datasets. Since it is widely believed that the learning procedure is dominated by head classes, most existing works focus explicitly on improving the recognition of tail classes.

### 2.1. Re-sampling and Re-weighting

Re-sampling and re-weighting are two classical approaches. Re-sampling methods aim at balancing the training set by either over-sampling the tail classes [3, 9, 21, 24] or under-sampling the head classes [7, 10]. Some methods also transfer statistics from the major classes to minor classes to obtain a balanced dataset [15, 19]. One disadvantage of re-sampling methods is that it may lead to either over-fitting of tail classes or under-fitting of head classes. Re-weighting methods, on the other hand, give each instance a weight based on its true label when computing the loss [2, 22, 23]. The major drawback of re-weighting is that it makes the loss function hard to optimize.

### 2.2. Two-Stage Decoupling

Due to the shortcomings of re-sampling and re-weighting, various other methods are proposed recently. For example, two-stage methods like [14, 18, 30] propose to decouple the learning of features and classifiers and achieve impressive results. Our method is similar to them in the sense that we also re-train a classifier given an existing model. But our motivation is to design a simple and flexible method that can be applied to a wide variety of existing models, and we propose a novel loss function, while two-stage methods like cRT [14] usually use the conventional cross entropy loss in the fine-tuning stage.

### 2.3. Hybrid and Multiple Heads

Hybrid methods, like [26, 31], make use of multiple heads to improve the recognition of different classes, which is similar to our proposed ensemble trick. But these methods require a joint training of multiple heads, together with a complex routing module that dynamically determines the head to use during inference, which increases the complexity of the model and makes it hard to use in practice. Our method, on the other hand, is simpler and requires no extra training nor complex modules to combine the predictions, which is different from them.

## 2.4. Miscellaneous

There are many approaches that use knowledge distillation [12, 18] and contrastive learning [4, 17, 25, 28, 33] and achieve descent results. There are also methods that perform logit adjustment [20, 29], which formulate long-tailed recognition as a label distribution shift problem. Recently, there is also one work that tries to solve the problem by performing some regularization on the classifier’s weight [1]. They also visualize the recall of each class in their work and find the model is biased towards common classes, which is similar to our motivation. However, they do not emphasize the importance of correctly recognizing all classes, which is different from us.

## 3. Method

We will present our method in this section. First, we will introduce the notations used. Then we will describe the major framework of our method, which includes our proposed loss function and a simple ensemble trick.

### 3.1. Notations

The neural network can be represented as a non-linear function  $\mathcal{F}$ . Given the  $i$ -th training image  $(\mathbf{x}^i, y^i)$  from the imbalanced dataset that has  $N$  images, where  $\mathbf{x}^i \in \mathbb{R}^{H \times W \times D}$ , a forward propagation through the network yields logit  $\mathbf{o}^i = \mathcal{F}(\mathbf{x}^i)$ . For a multi-class classification problem with  $C$  categories, a softmax activation function is usually applied and the result can be represented as  $\tilde{\mathbf{p}}^i = \text{softmax}(\mathbf{o}^i)$ , where  $\tilde{\mathbf{p}}^i \in \mathbb{R}^C$ .

### 3.2. A Simple Plug-In Method

Over the past few years, various methods [4, 14, 22] have been proposed to tackle the challenges in long tailed recognition. Although our evaluation results in Tab. 1 reveal that no obvious improvements have been made regarding the worst-performing category, many of their design choices are critical and effective in long-tailed recognition. Therefore, in order to make full use of existing advances, we do not aim at building up a whole new training framework, but rather propose a simple plug-in method that can be applied to a wide range of existing methods.

Our method is simple. Given a model trained using either one of existing methods, we re-initialize the last FC layer and re-train it using our proposed loss function while the backbone is frozen. We borrow this idea from many two-stage decoupling methods [14, 30], which is effective and flexible, making our method widely applicable. However, unlike normal decoupling methods that use the conventional cross entropy loss in the second stage, we propose a novel loss function to improve the worst-performing categories. Our proposed loss function, called GML (Geometric

Mean Loss), is defined as

$$\mathcal{L}_{\text{GML}} = -\frac{1}{C} \sum_{c=1}^C \log \bar{\mathbf{p}}_c, \quad (1)$$

where  $\bar{\mathbf{p}}_c$  is computed as

$$\bar{\mathbf{p}}_c = \frac{1}{N_c} \sum_{i=1}^{N_c} \tilde{\mathbf{p}}_{y^i}^i, \quad (2)$$

which is the average of  $\tilde{\mathbf{p}}_{y^i}^i$  across all training samples belonging to class  $c$  in this mini-batch, and  $N_c$  is the number of examples from this category in the mini-batch. Inspired by [22], we perform re-weighting when we compute  $\tilde{\mathbf{p}}^i$ :

$$\tilde{\mathbf{p}}_j^i = \frac{N_j \exp(o_j^i)}{\sum_{c=1}^C N_c \exp(o_c^i)}. \quad (3)$$

### 3.3. Why GML Works?

The definition of mean accuracy is

$$\text{Acc} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}(y^i = \arg \max \tilde{\mathbf{p}}^i), \quad (4)$$

where  $\mathbb{1}(\cdot)$  is the indicator function. On a balanced test set, it is the same as the arithmetic mean of per-class recall. We argue that maximizing the mean accuracy inevitably ignores some categories because arithmetic mean incurs little punishment to categories that have very small recall values. On the other hand, harmonic mean, defined as

$$\text{HM}(x_1, \dots, x_n) = \frac{n}{\frac{1}{x_1} + \dots + \frac{1}{x_n}}, \quad (5)$$

incurs strong punishment to small values. Therefore, we believe it can be a better alternative to optimize for. However, harmonic mean is defined using reciprocal, making it hard and numeric unstable to be optimized. To this extent, we propose to maximize the geometric mean of per-class recall, defined as

$$\text{GM}(x_1, \dots, x_n) = \sqrt[n]{x_1 \times \dots \times x_n}. \quad (6)$$

As an illustration,  $\frac{0.01+0.99}{2} = 0.5$ ,  $\text{HM}(0.01, 0.99) = 0.02$ , and  $\text{GM}(0.01, 0.99) = 0.10$ . Both harmonic and geometric means are heavily affected by the small value (0.01), while the arithmetic mean is much less sensitive to it.

Specifically, given the per-class recall as  $r_1, \dots, r_C$ , we transform their geometric mean using a simple logarithm transformation:

$$\sqrt[r_1 \dots r_C] = \exp\left(\log(r_1 \dots r_C)^{1/C}\right) \quad (7)$$

$$= \exp\left(\frac{1}{C} \log(r_1 \dots r_C)\right) \quad (8)$$

$$= \exp\left(\frac{1}{C} \sum_{i=1}^C \log r_i\right). \quad (9)$$

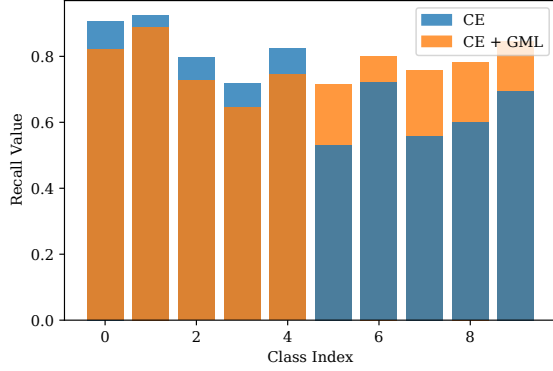


Figure 2. Bar plot of per-class recall on the imbalanced CIFAR10 (with imbalance ratio 100) before (‘CE’) and after (‘CE+GML’) the fine-tuning. For the fine-tuned model, the recall of the first 5 classes dropped while the recall of the latter 5 classes increased, motivating us to combine the prediction of both models.

Since recall is computed by computing the mean of a series of indicator functions and these indicator functions only take value 0 or 1 and is non-differentiable, we use  $\tilde{p}_c$  as a surrogate for  $r_c$  during training, resulting in

$$\sqrt[r_1 \dots r_C] \propto \exp(-\mathcal{L}_{\text{GML}}). \quad (10)$$

So, by minimizing  $\mathcal{L}_{\text{GML}}$ , we are effectively maximizing the geometric mean of per-class recall, thus being able to improve on the worst-performing categories.

### 3.4. Combining the Strength of Both Worlds

By fine-tuning the pre-trained model with our proposed loss function, indeed we are able to obtain a more uniform distribution of recall values across categories. As is shown in Fig. 2, the recall of classes that perform worse before are improved while the recall of classes that perform well before are dropped. Motivated by these dynamics, we then propose a simple ensemble trick that aims at combining the strength of both worlds.

Our idea is simple again. Since we re-initialize the classifier during fine-tuning and the backbone is kept frozen, the old classifier is still usable during inference. Let these two models be  $\mathcal{F}_{\text{old}}$  and  $\mathcal{F}_{\text{new}}$ , respectively, we then make predictions by

$$\tilde{\mathbf{p}}_{\text{ensembled}} = \frac{\tilde{\mathbf{p}}_{\text{new}} + \tilde{\mathbf{p}}_{\text{old}}}{2}. \quad (11)$$

In practice, since the calibration of some models can be poor [8, 30] (e.g., the baseline using cross entropy), we need to manually calibrate the prediction before combining them. For simplicity, we choose temperature scaling to calibrate the model [8] and use two temperature variables  $t_{\text{new}}$  and

$t_{\text{old}}$  when computing  $\tilde{\mathbf{p}}_{\text{new}}$  and  $\tilde{\mathbf{p}}_{\text{old}}$  separately:

$$\tilde{\mathbf{p}}_{\text{new}} = \text{softmax} \left( \frac{\mathbf{o}_{\text{new}}}{t_{\text{new}}} \right), \quad (12)$$

$$\tilde{\mathbf{p}}_{\text{old}} = \text{softmax} \left( \frac{\mathbf{o}_{\text{old}}}{t_{\text{old}}} \right). \quad (13)$$

One big advantage of our proposed ensemble trick is that it brings nearly no extra cost. Unlike many multi-head [31] or multi-expert [26] methods, we do not need extra training as we simply use the classifier that comes with the pre-trained model. During inference, the only extra cost is the fully connected layer to compute  $\tilde{\mathbf{p}}_{\text{old}}$ , which is negligible. Also, no complex routing rules are needed.

The two temperature variables are the only two hyper-parameters introduced by our method, and they can be easily tuned as no training is required. Plus, as we will later show in the ablation studies, most values work perfectly fine so actually no painful tuning is needed. What’s more, this ensemble trick is completely optional, our method works well enough even without it.

### 3.5. Training Pipeline

Our method is very similar to the two-stage decoupling methods [14, 30]. We design it in this way on purpose because we want it to be as flexible and simple as possible. During experiments, we found that our method is orthogonal to many design choices in state-of-the-art methods and is able to improve the worse-case performance when applied to various different pre-trained models. The overall training pipeline of our method consists of three stages:

1. **Pre-training stage:** Obtain the pre-trained model from scratch using any one of the existing methods.
2. **Fine-tuning stage:** Freeze the backbone and re-train the classifier using our proposed loss function.
3. **An optional ensemble stage:** Calibrate the prediction of two classifiers and combine them additively.

The overall training procedure is summarized in Algorithm 1.

## 4. Experiments

In order to validate the effectiveness of our method, we have conducted experiments on three widely used benchmark datasets. In this section, we will present our experimental results. First, we will introduce the datasets used. Then we will describe various implementation details and evaluation metrics used. After that, we will compare our method with various other methods. Finally, we will present the results of some ablation studies.

Dataset	Number of Classes	# Training Images	# Test Images	Imbalance Ratio
CIFAR100-LT [16]	100	10,847	10,000	100
ImageNet-LT [6, 19]	1,000	115,846	50,000	256
Places-LT [32]	365	62,500	36,500	996

Table 2. Statistics of three imbalanced datasets used in our experiments.

---

**Algorithm 1** The overall training procedure

---

**Input:** Training images  $x$  and their labels  $y$ .

- 1: Randomly initialize the network and train it from scratch to obtain the pre-trained model.
  - 2: Use the pre-trained model to perform initialization.
  - 3: Freeze the backbone and re-initialize the classifier.
  - 4: Fine-tune the model using the loss function defined as Eq. (1) for a few epochs.
  - 5: During inference, combine the results from both classifiers as defined in Eq. (11).
- 

#### 4.1. Datasets

We use three widely used long-tailed image recognition datasets: CIFAR100-LT, ImageNet-LT and Places-LT. Some statistics about these three datasets are summarized in Tab. 2. Since CIFAR100 [16], ImageNet [6] and Places [32] are all balanced datasets, we follow previous works [19, 31] to down-sample their original training set. Details about the construction process can be found in the supplementary material.

#### 4.2. Implementation Details

Following previous works [19, 31], we use ResNet-32 [11], ResNeXt-50 [27] and ResNet-152 as the backbone network for CIFAR100-LT, ImageNet-LT and Places-LT, respectively. For the pre-training stage, if there exists released checkpoints, we use them directly. If not, we use the officially released codes and try to reproduce the results in our best effort without changing their original settings. This includes stuffs like hyper-parameters and data augmentations used. For the fine-tuning stage, detailed training settings can be found in the supplementary material.

#### 4.3. Evaluation Metrics and Comparison Methods

The convention in long-tailed recognition research is to split the classes into three subsets based on the number of training images [14, 19]. Typically, “Head” denotes classes that have more than 100 images, “Few” denotes classes that have less than 20 images and all other classes are denoted as “Medium”. The accuracy within each subset is often reported along with the overall accuracy to justify the claim that the performance improvement mainly comes from tail classes. And according to our observation in Sec. 1, we ar-

gue that such an evaluation scheme can be problematic and we propose to first compute the recall within each category and then compute their harmonic mean. In our experiments, since our method optimizes geometric mean of per-class recall as a surrogate, we also report the geometric mean. Besides, we report the overall mean accuracy as well.

We compare our methods with both baseline methods like simple cross entropy loss and various state-of-the-arts methods [4, 12, 22, 26]. Many of these methods have very different training settings and strictly speaking, they are not directly comparable. But since our method is a simple plug-in method, the main focus is to quantify the improvements brought by our method compared to its respective baseline result.

#### 4.4. Main Comparisons

We validate our method on three benchmark datasets and present the results on each dataset here separately.

##### 4.4.1 CIFAR100-LT

Tab. 3 shows the experimental results on the CIFAR100-LT dataset. CIFAR100-LT has three different versions with different imbalance ratios and here we only consider the one with imbalance ratio 100. As we can see from the table, although various different methods have been proposed in recent years and there are rapid improvements on the overall accuracy, the performance of the worst-performing category is still relatively poor. The geometric and harmonic mean of per-class recall also justify our claim. We apply our method to both the baseline method CE and state-of-the-art method PaCo [4] on this dataset. When our method is applied to CE and PaCo, we are able to improve the lowest recall from 0.00 and 5.00 to 6.00 and 9.00 respectively. And if we further apply our proposed ensemble trick, we are able to improve the lowest recall to 15.00, achieving state-of-the-art performance in both geometric mean and harmonic mean while keeping the overall accuracy roughly unchanged.

It is worth noting that even the simplest CE+GML outperforms all existing methods in terms of lowest recall, that is, in terms of “no one left behind”.

##### 4.4.2 ImageNet-LT and Places-LT

Tab. 4 shows the result on ImageNet-LT. The overall observation is similar to that on CIFAR100-LT, except that



Methods	Accuracy	Geometric Mean	Harmonic Mean	Lowest Recall
CE	38.74	21.03*	2.05*	0.00
BSCE [22]	42.24	35.16	26.24	3.00
LDAM [2]	43.51	33.61	21.49	3.00
LADE [13]	44.39	39.05	32.58	5.00
DiVE [12]	45.11	37.08	25.18	2.00
MiSLAS [30]	47.05	40.16	30.93	5.00
RIDE [26]	48.64	38.71	23.86	2.00
PaCo [4]	51.24	45.29	36.42	5.00
CE + GML	41.06	36.59	31.26	6.00
PaCo + GML	50.53	45.47	39.20	9.00
PaCo + GML (Ensemble)	49.82	<b>45.70</b>	<b>41.02</b>	<b>15.00</b>

Table 3. Results on the CIFAR100-LT dataset with imbalance ratio 100. Numbers with \* are computed by substituting zero elements with a small number ( $10^{-3}$ ) or else the geometric and harmonic mean will all be zero.

Methods	Accuracy	Geometric Mean	Harmonic Mean	Lowest Recall
CE	43.90	23.25*	1.25*	0.00
BSCE [22]	50.48	42.32*	13.74*	0.00
cRT [14]	49.64	41.35*	13.82*	0.00
DiVE [12]	53.63	45.49*	12.76*	0.00
RIDE [26]	55.69	47.56*	17.32*	0.00
PaCo [4]	58.53	51.32*	21.79*	0.00
CE + GML	45.61	39.82	31.67	2.00
PaCo + GML	55.57	50.67	43.71	2.00
PaCo + GML (Ensemble)	57.22	<b>52.29</b>	<b>44.80</b>	<b>2.00</b>

Table 4. Results on the ImageNet-LT dataset. Numbers with \* are computed by substituting zero elements with a small number ( $10^{-3}$ ) or else the geometric and harmonic mean will all be zero.

since this dataset is larger and harder, *the lowest recall of all current methods are zero*, making the harmonic mean of per-class recall very low, even when a small number substitutes zero per-class recalls. Our method, on the other hand, greatly improves the harmonic mean by successfully improving the lowest recall value. All categories have per-class recall higher than 0 when our method is used, even in the simplest CE+GML.

Tab. 5 shows the result on Places-LT. Current state-of-the-art on this dataset is MiSLAS [30], so we apply our method upon it and the simple baseline CE. Again, our method successfully improves the lowest recall, as well as the harmonic mean of per-class recall. And when the ensemble trick is further applied, we are able to achieve new state-of-the-art results.

#### 4.4.3 Universality

Since we want our method to be a universal plug-in, in this subsection we present the results of the experiment of adding our method on top of various different methods. The experiment is conducted on CIFAR100-LT, and the re-

sult is shown in Tab. 6. Note that here we do not perform any ensemble. As we can see, our method is applicable to various different methods, and is able to achieve consistent improvements. However, we do observed that there exists some methods, like DiVE [12], where we failed to achieve obvious improvements. More details about this can be found when we discuss limitations in Sec. 5.

#### 4.5. Ablation Studies

For ablation studies, if not otherwise specified, all experiments are conducted on CIFAR100-LT (with imbalance ratio 100) under the default training setting.

##### 4.5.1 Effects of Two Temperature Variables in the Ensemble Stage

When we design our method, we'd like to keep the number of new hyper-parameters as less as possible in order for it to be readily applicable in real-world applications. The only two hyper-parameters introduced by our method are in the optional ensemble stage. Here we study the effects of choosing different values of them. Specifically, we change

Methods	Accuracy	Geometric Mean	Harmonic Mean	Lowest Recall
CE	28.71	12.11*	0.73*	0.00
BSCE [22]	37.18	29.30*	5.64*	0.00
PaCo [4]	40.45	27.88*	2.53*	0.00
MiSLAS [30]	40.48	35.53	28.97	3.00
CE + GML	36.82	29.30*	8.00*	0.00
MiSLAS + GML	39.90	35.41	29.98	5.00
MiSLAS + GML (Ensemble)	40.37	<b>35.92</b>	<b>30.51</b>	<b>5.00</b>

Table 5. Results on the Places-LT dataset. Numbers with \* are computed by substituting zero elements with a small number ( $10^{-3}$ ) or else the geometric and harmonic mean will all be zero. Note that to plug GML into MiSLAS, since MiSLAS is also a two-stage method, we think it’s fairer to use GML together with their proposed label aware smoothing loss.

Methods	G-Mean	H-Mean	Lowest Recall
CE	21.03*	2.05*	0.00
CE + GML	36.59 (15.56 $\uparrow$ )	31.26 (29.21 $\uparrow$ )	6.00 (6.00 $\uparrow$ )
BSCE	35.16	26.24	3.00
BSCE + GML	36.52 (1.36 $\uparrow$ )	30.88 (4.64 $\uparrow$ )	7.00 (4.00 $\uparrow$ )
MiSLAS	40.16	30.93	5.00
MiSLAS + GML	40.90 (0.74 $\uparrow$ )	36.49 (5.56 $\uparrow$ )	11.00 (6.00 $\uparrow$ )
PaCo	45.29	36.42	5.00
PaCo + GML	45.47 (0.18 $\uparrow$ )	39.20 (2.78 $\uparrow$ )	9.00 (4.00 $\uparrow$ )

Table 6. Results on the CIFAR100-LT dataset with imbalance ratio 100. Our proposed method is applicable to various methods. “G-Mean” is short for “Geometric Mean” and “H-Mean” is short for “Harmonic Mean”.

$t_{old}$	$t_{new}$	Geometric Mean	Harmonic Mean	Lowest Recall
1	1	45.70	<b>41.02</b>	<b>15.00</b>
1	2	45.58	38.86	8.00
1	3	44.68	36.09	5.00
2	1	44.39	40.22	<b>15.00</b>
2	2	45.61	41.01	<b>15.00</b>
2	3	<b>46.06</b>	40.40	11.00
3	1	43.73	39.63	<b>15.00</b>

Table 7. Ablation on the effects of temperature variables when we apply our ensemble method upon PaCo [4].

the value of the temperature variables shown in Eq. (12) and Eq. (13) in the ensemble stage when we apply our method to PaCo [4]. The results are shown in Tab. 7. In our experiments on CIFAR100-LT, we set  $t_{old} = 1$  and  $t_{new} = 1$ . And as we can see from the table, the value of  $t_{old}$  does not affect the result much while  $t_{new}$  is somewhat important, especially for improving the lowest recall—higher  $t_{new}$  will hurt it. This is natural because higher  $t$  results in a smoother distribution and less information is preserved. Since the newly trained classifier is crucial for improving the lowest recall, we need to preserve more information from it.

Generally speaking, since these two hyper-parameters

Methods	G-Mean	H-Mean	Lowest Recall
CE + GML	<b>36.59</b>	<b>31.26</b>	<b>6.00</b>
(w/o re-weighting)	25.45	5.11	0.00
(w/o re-weighting, re-sampling)	32.32	23.05	4.00

Table 8. Ablation on the re-weighting design in GML.

Methods	G-Mean	H-Mean	Lowest Recall
CE + GML	36.59	31.26	6.00
PaCo + GML	<b>45.47</b>	<b>39.20</b>	<b>9.00</b>
GML (train from scratch)	36.63	30.86	<b>9.00</b>

Table 9. Training from scratch using GML.

are introduced for calibration, for different methods, we may need to tune them respectively because the calibrations of different methods are known to vary a lot in long-tailed recognition problems [30]. But usually this will not cause a headache, because the ensemble stage does not involve any training and thus can be tuned very quickly.

#### 4.5.2 Re-weighting in GML

As shown in Eq. (3), re-weighting is applied in our proposed GML, here we ablate this design. The results are shown in Tab. 8. As we can see, the performance will be relatively poor without re-weighting. And just like many two-stage decoupling methods [14], re-sampling during fine-tuning can help, but is inferior to re-weighting.

#### 4.5.3 Training from Scratch using GML

Since our main purpose is to propose a simple plug-in method that is widely applicable to various existing methods, we use GML only to re-train the classifier of an existing model. However, GML can also be used to train a model from scratch. The results are shown in Tab. 9. As we can see, the result is comparable to the result when GML




Figure 3. Bar plot of per-class recall on the imbalanced CIFAR100 (with imbalance ratio 100) before and after the fine-tuning. No ensemble is performed. Our method focuses on improving the recognition of classes that have low recall values, making sure no category is left behind.

is attached to CE, but inferior to the result when GML is applied to SOTA methods like PaCo, just as expected.

#### 4.5.4 Visualization of the Per-Class Recall

In Fig. 1, we visualize the per-class recall of model trained on the imbalanced dataset, and find they vary a lot from class to class. Since our major motivation is to make sure all classes are equally treated so that no category will be left behind, here we visualize the change of per-class recall after applying our method. We perform visualizing using two datasets, CIFAR10-LT and CIFAR100-LT, both have imbalance ratio 100. The visualization results on CIFAR10-LT and CIFAR100-LT are shown in Fig. 2 and Fig. 3, respectively. Note that here we do not perform the ensemble. As we can clearly see in both figures, our fine-tuned model has a more uniform distribution of per-class recall. These two figures justify the effectiveness of our proposed method.

## 5. Limitations and Broader Impacts

Although our method is applicable to a wide range of existing methods, we do find that when applied to some methods, the improvement is less obvious compared to others. For example, our framework does not cope well with methods like BSCE [22] and DiVE [12]. Interestingly, although DiVE is based on knowledge distillation, BSCE is used in DiVE as well to train both the teacher and student model so we attribute the failure to the fact that these method all re-weight the loss function. Two-stage decoupling methods like cRT [14] are also known to not perform well when the first stage training involves re-weighting or re-sampling. The author argues it is because re-sampling and re-weighting have negative effects on feature learning and we believe that argument applies to our method as well.

Another thing worth noting here is that we do not report the experimental results on iNaturalist2018 [5] because we find it generally not suitable in our settings. iNaturalist2018 is another widely used benchmark dataset in long-tailed recognition. It is large-scale, consisting of 437.5K training images from 8142 classes. But for the test split, it only has 24K images, so each class only has 3 images. Since our major motivation is to make sure no category is left behind, it is generally not meaningful to look at the per-class recall in iNaturalist2018 because the variance of the result can be very large due to the very small number of test images. However, in real-world applications, it is generally rare to have such small number of test images.

For broader impacts, we believe our research have positive impacts on reducing the bias of model training on an imbalanced dataset. Machine learning models are known to be biased if there is lack of training images from some categories. We believe our work will be helpful to achieve better fairness by making sure all categories are equally treated.

## 6. Conclusions

In this paper, we focused on improving some of the worst-performing classes in long-tailed recognition. We found that when trained with an imbalanced dataset, the per-class recall of the model varies a lot from class to class. Previous convention reports the arithmetic mean of per-class recall, but we argued such an evaluation scheme can be problematic. First, arithmetic mean only incurs little penalty to small numbers, making it too coarse to reflect whether there are categories that are left behind (i.e., with close to zero per-class accuracy). Second, it is not necessarily true that classes in the “Few” category perform worse than those in “Many” or “Medium”, so focusing on improving “Few” accuracy alone is not enough. For these reasons, we argued that we should pay more attention to the harmonic mean of per-class recall and we therefore proposed a novel method that aims at improving the lowest recall and the harmonic mean of recalls. Our method is a simple plug-in that is applicable to a wide range of existing methods. Specifically, our method consists of three stages, in the first stage, we use any existing method to train a model from scratch. Then we re-train the classifier using our proposed novel GML loss function. Finally, we propose a simple ensemble trick that can be used to combine the predictions from the two classifiers with nearly no extra cost. We validated the effectiveness of our method on three widely used benchmark datasets, and witnessed consistent improvements on the harmonic mean of recalls and lowest recall value, while the overall accuracy still remains high. By visualizing the distribution of per-class recall values of the fine-tuned model, we found our model indeed achieved a more balanced distribution.



## References

- [1] Shaden Alshammari, Yu-Xiong Wang, Deva Ramanan, and Shu Kong. Long-tailed recognition via weight balancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6897–6907, June 2022. [3](#)
- [2] 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 32*, pages 1567–1578, December 2019. [2](#), [6](#)
- [3] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002. [2](#)
- [4] Jiequan Cui, Zhisheng Zhong, Shu Liu, Bei Yu, and Jiaya Jia. Parametric contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 715–724, October 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [5] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4109–4118, June 2018. [8](#)
- [6] 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, June 2009. [1](#), [5](#)
- [7] Chris Drummond and Robert C. Holte. C4.5, class imbalance, and cost sensitivity: Why under-sampling beats over-sampling. In *Workshop on Learning from Imbalanced Data Sets II*, 2003. [2](#)
- [8] 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*, pages 1321–1330, August 2017. [4](#)
- [9] Hui Han, Wen-Yuan Wang, and Bing-Huan Mao. Borderline-SMOTE: A new over-sampling method in imbalanced data sets learning. In *Advances in Intelligent Computing*, pages 878–887, August 2005. [2](#)
- [10] Haibo He and Eduardo A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, 2009. [2](#)
- [11] 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, June 2016. [5](#)
- [12] Yin-Yin He, Jianxin Wu, and Xiu-Shen Wei. Distilling virtual examples for long-tailed recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 235–244, October 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)
- [13] Youngkyu Hong, Seungju Han, Kwanghee Choi, Seokjun Seo, Beomsu Kim, and Buru Chang. Disentangling label distribution for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6626–6636, June 2021. [6](#)
- [14] 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*, April 2020. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [15] 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, June 2020. [2](#)
- [16] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>. [1](#), [5](#)
- [17] Tianhong Li, Peng Cao, Yuan Yuan, Lijie Fan, Yuzhe Yang, Rogerio S. Feris, Piotr Indyk, and Dina Katabi. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6928, June 2022. [3](#)
- [18] Tianhao Li, Limin Wang, and Gangshan Wu. Self supervision to distillation for long-tailed visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 630–639, October 2021. [2](#), [3](#)
- [19] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, June 2019. [2](#), [5](#)
- [20] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *International Conference on Learning Representations*, May 2021. [3](#)
- [21] Seulki Park, Youngkyu Hong, Byeongho Heo, Sangdoon Yun, and Jin Young Choi. The majority can help the minority: Context-rich minority oversampling for long-tailed classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6887–6896, June 2022. [2](#)
- [22] Jiawei Ren, Cunjun Yu, Shunan Sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. In *Advances in Neural Information Processing Systems 33*, pages 4175–4186, December 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [23] Dvir Samuel and Gal Chechik. Distributional robustness loss for long-tail learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9495–9504, October 2021. [2](#)
- [24] Li Shen, Zhouchen Lin, and Qingming Huang. Relay back-propagation for effective learning of deep convolutional neural networks. In *Proceedings of the European Conference on Computer Vision*, pages 467–482, October 2016. [2](#)
- [25] Peng Wang, Kai Han, Xiu-Shen Wei, Lei Zhang, and Lei Wang. Contrastive learning based hybrid networks for long-tailed image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 943–952, June 2021. [3](#)

- [26] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella Yu. Long-tailed recognition by routing diverse distribution-aware experts. In *International Conference on Learning Representations*, May 2021. [1](#), [2](#), [4](#), [5](#), [6](#)
- [27] Saining Xie, Ross Girshick, Piotr Dollar, 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, July 2017. [5](#)
- [28] Yuzhe Yang and Zhi Xu. Rethinking the value of labels for improving class-imbalanced learning. In *Advances in Neural Information Processing Systems 33*, pages 19290–19301, December 2020. [3](#)
- [29] Songyang Zhang, Zeming Li, Shipeng Yan, Xuming He, and Jian Sun. Distribution alignment: A unified framework for long-tail visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2361–2370, June 2021. [3](#)
- [30] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16489–16498, June 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [31] 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, June 2020. [2](#), [4](#), [5](#)
- [32] 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, 2018. [5](#)
- [33] Jianggang Zhu, Zheng Wang, Jingjing Chen, Yi-Ping Phoebe Chen, and Yu-Gang Jiang. Balanced contrastive learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6908–6917, June 2022. [3](#)