

Neural collapse inspired attraction–repulsion-balanced loss for imbalanced learning

Liang Xie^a, Yibo Yang^b, Deng Cai^{a,*}, Xiaofei He^{a,c}

^aState Key Lab of CAD & CG, Zhejiang University, Hangzhou, China

^bJD Explore Academy, Beijing, China

^cFabu Inc., Hangzhou, China

ARTICLE INFO

Article history:

Received 20 April 2022

Revised 19 October 2022

Accepted 8 January 2023

Available online 11 January 2023

Communicated by Zidong Wang

Keywords:

Long-tailed learning

Neural collapse

Machine Learning

Image Classification

ABSTRACT

Class imbalance distribution widely exists in real-world engineering. However, the mainstream optimization algorithms that seek to minimize error will trap the deep learning model in sub-optimums when facing extreme class imbalance. It seriously harms the classification precision, especially in the minor classes. The essential reason is that the gradients of the classifier weights are imbalanced among the components from different classes. In this paper, we propose Attraction–Repulsion–Balanced Loss (ARB-Loss) to balance the different components of the gradients. We perform experiments on large-scale classification and segmentation datasets, and our ARB-Loss can achieve state-of-the-art performance via only one-stage training instead of 2-stage learning like nowadays SOTA works.

© 2023 Elsevier B.V. All rights reserved.

1. Introduction

Class imbalance exists in almost all real-world scenes, and we call it long-tailed distribution when the imbalance is very extreme. Due to that standard gradient descent optimization algorithm seeking the solution to make the loss as low as possible, the learned algorithm tends to perform better on the major classes while ignoring the learning on those minor classes. It will trap the model in the sub-optimum.

Many works have recently emerged to mitigate the performance drop under the long-tailed distribution. Some works [1,2] re-sample the data to mitigate the class imbalance. However, over-sampling might cause the model to overfit these classes, while under-sampling might bring the risks of removing the important samples. Another line of works [3–8] designs specific loss function to weight different classes according to the class cardinality. However, for the large-scale dataset, the re-weighting cost functions make deep models difficult to optimize. Additionally, these methods rely on hyper-parameters, which must be adjusted according to the specific data. Recently, many methods use two-stage training to decouple the representation learning, and the class-balanced training of the classifier [7,9–11]. They first use instance-balanced sampling to learn discriminative features, and then they fix the backbone network and train the classifier via

class-balanced sampling. The common belief behind these 2-stage methods is that the instance-balanced sampling learns better and more general representations, which has been demonstrated via better experimental results. However, the 2-stage methods fail to explain the essential reason for the poor performance in long-tailed data distribution. Moreover, these methods involve many tricks to obtain better results, which brings difficulties and complexities for training and poor interpretability. Some work attempts to solve the issue by balancing the gradients. [8] filters out the gradients from major classes via an 0-1 mask when back-propagating the gradients for samples of minor classes. However, besides the heavy reliance on many hyper-parameters, they only filter out gradients that are too large and still do not essentially solve the imbalance of the remaining gradients. Therefore, the key to solving the performance drop under the long-tailed distribution is to balance each gradient component from different classes.

A recent study [12] observes the *Neural Collapse* phenomenon. When the model converges to optimum in an ideal state, *i.e.*, a balanced data distribution, four phenomenons can be observed:

1. The intra-class features will collapse to their class means.
2. The vectors of the class means will collapse to an ETF (Equiangular Tight Frame) geometric structure, *i.e.*, the class mean vectors will have consistent length and equal-sized angles between any given pair, where the pairwise-distanced are maximized.

* Corresponding author.

3. The class means and the classifier weights will dually converge to each other so that the dot product of the features and the classifier weights will obtain the maximum for matched classes while the minimum for unmatched classes.
4. The classifier decides the class according to the Euclidean distances among the feature vector and the classifier weights.

However, under an imbalanced data distribution, the law of *Neural Collapse* is no longer satisfied, and another unexpected phenomenon, *Minority Collapse* [13], occurs, i.e., some weight vectors are close and even merged. The collapsed weights can no longer effectively distinguish features of different classes, which explains the deteriorated performance of classification on imbalanced data. We illustrate these two observations in 1.

Inspired by the *Neural Collapse* phenomenon, in this paper, we propose Attraction–Repulsion–Balanced Loss, abbreviated as ARB-Loss, to balance the gradients from different classes. We first analyze that the essential reason of *Minority Collapse* is the imbalance between the *attractive* and *repulsive* gradient components from different classes. The classifier weights of minority classes receive a much stronger *repulsion* gradients from majority-class features compared to the *attractive* gradients from intra-class features. Therefore, to mitigate the *Minority Collapse* dilemma, ARB-Loss adds coefficients on the denominator of the softmax function to balance the gradients from different classes. Compared to the state-of-the-art methods for imbalanced learning problems [9,11], our ARB-Loss can achieve comparable or even better performance without bells and whistles.

The contributions of this paper can be summarized as follows:

1. We analyze that the cause of *Minority Collapse* is the imbalance between the *attractive* and *repulsive* gradient components from different classes. Based on the analyses, we propose ARB-Loss to balance the gradients from different classes to mitigate the *Minority Collapse* dilemma under the imbalanced distribution.
2. We analyze the properties of ARB-Loss. Both the theoretical and empirical results show that compared with the Cross-Entropy loss, ARB-Loss can achieve more balanced gradients and lead to a model closer to the neural collapse state.
3. We perform experiments on large-scale long-tailed image classification and segmentation tasks. The imperial results suggest that our method can achieve comparable or even better performance without bells and whistles.

In the following sections, we first review the related works of long-tailed learning and the studies of the *Neural Collapse* phenomenon in Section 2. Then, in Section 3, we rethink the learning mechanism under the traditional cross-entropy loss from the perspective of the gradients and analyze the cause of the unexpected *Minority Collapse*. In Section 4, we elaborate on our ARB-Loss and provide relevant analytical results. In Section 5, we perform experiments to show the effectiveness of our method. Finally, in Section 6, we make a conclusion for this paper and a prospective for future work.

2. Related works

2.1. Long-tailed learning

Re-sampling. Early works [1,14,15,2] re-sample the data (over-sample the samples of minor classes or under-sample the samples of major classes or both) to balance the data distribution. Over-sampling repeats the sample of minor classes, which might make the model overfit these classes and harm the generality. Under-sampling removes some samples of major classes, which might remove the key data for representation learning and bring the performance drop.

Re-weighting. Another idea is to assign different weights for different classes, even instances. [3,4,6,8] re-weight the loss according to the class cardinalities. [16] introduce the prior probabilities from a Bayesian view and balance the exponential logits via the class cardinalities. Although [16] shares a similar formulation of the loss function with ours, we derive from a completely new perspective and offer in-depth analyses. Focal loss [17,18] re-weight each instances according to their *hard* level, i.e., making the model to pay attention to the wrong-recognized samples.

Two-stage Learning. An observation is that the representations learned under the instance-balancing sampling are more general. So many works try to split learning into two stages. [7] first trains the model in a normal way in stage 1 and then uses deferred resampling to fine-tune with class-balanced resampling or uses deferred re-weighting to re-weight different classes in stage 2. [10] spatially disentangles the regular feature learning and re-balancing learning via two parallel branches.[9,11] decouple the representation learning and classification. They concluded that the instance-balanced sampling gives more general representations. They first used instance-balanced sampling to learn the rep-




Fig. 1. (a) *Neural Collapse*: when trained on the balanced data, the classifier weights and features are dually collapsed into an ETF geometric structure (Note that the figure depicts a 3D ETF structure, where the arrows are equal-length like the molecular structure of methane.); (b) *Minority Collapse*: when the cardinalities of some classes (the orange and purple classes) are much lower than others (the green and blue classes), the weights corresponding to these rare classes collapsed into similar direction. (The arrow represents the classifier weights, the light-color spot represents the head of the feature vector the dark-color square represents the feature mean; each color represents a class.) (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

representations at stage-1. Then, they freeze the feature and retrain the classifier (cRT) via techniques such as label-aware smooth (LAS) and learnable weight scaling (LWS).

2.2. Neural collapse

A recent study [12] discovers the classifier will converge to the state called *Neural Collapse* when trained on the balanced data. Especially, *Neural Collapse* is a phenomenon that the mean vectors of the intra-class features and the classifier weight vectors dually converge to an ETF (Equiangular Tight Frame) geometric structure, and the decision is based on the Euclidian distance among the feature vectors and the classifier weights. In the 3.1, we give a detailed description for the *Neural Collapse* phenomenon. Most of existing theoretical methods [19–24] only analyse the reason of *Neural Collapse* phenomenon when being trained under a balanced data distribution. However, for most real-world data, the class distribution is usually imbalanced. In the imbalanced cases, an unexpected phenomenon, *Minority Collapse* [13], occurs. The classifier weights of minority classes will converge to similar directions. These theoretical works give an explainable view of the mechanism of the classifier under balanced or imbalanced distribution. However, there is no design inspired by *Neural Collapse* for the long-tailed learning. Some recent work [25] tries to avoid the gradient imbalance via fixing the classifier weights. However, it is not guaranteed that a fixed length of classifier weight does not harm representation learning. In this paper, we reveal that the *Minority Collapse* phenomenon trained under imbalanced data distribution is essentially caused by the imbalanced gradient components. Inspired by this idea, we propose ARB-Loss from the *Neural-Collapse* view to balance the *attractive* and different *repulsive* components of the gradients for the classifier weights, which are much more interpretable compared to previous methods.

3. Rethinking the cross entropy loss

Notation & Problem Setup. For a c -class classification task, we denote $\mathcal{X} = \{\mathbf{x}_i, y_i\}, i \in \{1, \dots, n\}$ as the observed dataset, where n is the number of samples, \mathbf{x}_i and $y_i \in \{1, \dots, c\}$ are the i -th sample and its label respectively. We also use $\mathbf{y}_i \in \mathbb{R}^c$ as the corresponding vectorized label, e.g. the one-hot or smoothed one-hot label vector. Suppose that class i has n_i samples in the dataset \mathcal{X} , we use $\pi(i)$ to represent the set of sample indices that belongs to the class i , i.e., $\pi(i) = \{j : y_j = i\}$.

A vanilla L -layer deep neural network can be represented as:

$$\Phi_{\Theta}(\mathbf{x}) = \mathbf{W}_L \sigma(\mathbf{W}_{L-1} \sigma(\dots \sigma(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1) + \dots) + \mathbf{b}_{L-1}), \quad (1)$$

where \mathbf{W}_L is the final classifier weight. We omit its bias for simplicity because bias can be incorporated into the weights by adding a dimension filled with 1 on the features. The optimization target is:

$$\min_{\Theta} \sum_{i=1}^c \sum_{j \in \pi(i)} \mathcal{L}(\Phi_{\Theta}(\mathbf{x}_j), \mathbf{y}_j) + \frac{\lambda}{2} \|\Theta\|^2, \quad (2)$$

where $\Theta = \{\mathbf{W}_i, \mathbf{b}_i\}_{i=1}^L$ are the learnable parameters, \mathcal{L} is the loss function, and λ is the weight decay parameter.

We use $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_c]^T \in \mathbb{R}^{d \times c}$ to represent the classifier weights and $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_n] \in \mathbb{R}^{d \times n}$ to note the features of the last layer, i.e., $\mathbf{H} = \{\mathbf{h}_i = \sigma(\mathbf{W}_{L-1} \sigma(\dots \sigma(\mathbf{W}_1 \mathbf{x}_i + \mathbf{b}_1) + \dots) + \mathbf{b}_{L-1}) \mid i = 1, 2, \dots, n\}$. Because the norms of the weights $\{\mathbf{W}_i\}_{i=1}^L$ and bias $\{\mathbf{b}_i\}_{i=1}^{L-1}$ have supremums and the popular activation functions also have finite supremums (e.g. Sigmoid, tanh etc.) or keep linear mapping in the non-restraint part (e.g. ReLU), the norm of \mathbf{H} also have a supremum:

$$\frac{1}{n} \|\mathbf{H}\|_F^2 = \frac{1}{n} \sum_{i=1}^c \sum_{j \in \pi(i)} \|\mathbf{h}_j\|_2^2 \leq E_H, \quad (3)$$

which also means that, for a certain class i , the features also have supremum:

$$\frac{1}{n_i} \sum_{j \in \pi(i)} \|\mathbf{h}_j\|_2^2 \leq E_{H,i}, \quad i = \{1, 2, \dots, c\}, \quad (4)$$

where $E_H \in \mathbb{R}, E_{H,i} \in \mathbb{R}, i \in \{1, \dots, n\}$ are constants.

3.1. Neural collapse

[12] observes that the feature and weight vectors of the classifier will dually converge to a special geometric structure under the balanced data distribution, and this phenomenon is called *Neural Collapse*. *Neural Collapse* have four manifestations:

1. The features will converge to their class mean.
2. Under a balanced data distribution, the class mean vectors will collapse to an ETF (Equiangular Tight Frame) geometric structure, which can be represented as:

$$\mathbf{W}^* = \sqrt{\frac{c}{c-1}} \mathbf{P} \left(\mathbf{I}_c - \frac{1}{c} \mathbf{1}_c \mathbf{1}_c^T \right), \quad (5)$$

where $\mathbf{P} \in \mathbb{R}^{d \times c}$, ($d \geq c$) is a partial orthogonal matrix such that $\mathbf{P}^T \mathbf{P} = \mathbf{I}_c$, \mathbf{I}_c is the $c \times c$ identical matrix and $\mathbf{1}_c \in \mathbb{R}^{c \times 1}$ is a vector filled with 1. For the classifier weights which has collapsed to the ETF structure, they meet:

$$\mathbf{w}_i^T \mathbf{w}_j^* = \begin{cases} 1, & i = j \\ -\frac{1}{c-1}, & i \neq j \end{cases}, \quad (6)$$

where \mathbf{w}_i^* is the i -the column of \mathbf{W}^* , i.e., the weights corresponding to class i .

3. The classifier weights will dually converge to the mean vectors of the corresponding class, i.e., the classifier weight vectors will collapse to the consistent ETF geometric structure.
4. The decision of the classifier is based on the Euclidean distances of the feature vectors and each the classifier weight, i.e., deciding the class whose mean vector is nearest to the current features.

When the classifier weights and feature vectors obey these phenomenons of *Neural Collapse*, the Fisher's discriminant ratio [26] is maximized, i.e., the optimal geometric structures for classification problem.

However, for long-tailed learning, *Neural Collapse* no longer emerges. Instead, *Minority Collapse* occurs [13], which describes the phenomenon that the classifier weight vectors of minority classes converge in a similar direction, i.e.,

$$\lim_{\substack{n_m \rightarrow \infty \\ \frac{n_m}{n} \rightarrow \infty}} \mathbf{w}_i - \mathbf{w}_j = \mathbf{0}_d \quad (7)$$

where n_m is the cardinality of one of the majority classes. Due to the collapse of the minority-class weight vectors, the classifier fails to discriminate these classes correctly, which explains the fundamental reason for the poor performance under the long-tailed data distribution. In the following, we analyze the gradients derived from cross-entropy loss and reveal that the cause of *Minority Collapse* is the imbalance among gradient components from different classes.

3.2. The gradients of cross-entropy loss

The cross-entropy loss can be defined as:

$$\mathcal{L}_{ce} = - \sum_{j=1}^b \sum_{i=1}^c \mathbf{y}_{j,i} \sigma(\mathbf{z}_{j,i}) = - \sum_{j=1}^b \sum_{i=1}^c \mathbf{y}_{j,i} \log \left(\frac{\exp(\mathbf{w}_i^T \mathbf{h}_j)}{\sum_{k=1}^c \exp(\mathbf{w}_k^T \mathbf{h}_j)} \right), \quad (8)$$

where b represents the batch size, σ is the softmax function, $\mathbf{z}_{j,i} = \mathbf{w}_i^T \mathbf{h}_j$ is the logit of the j -th sample for the class i . When the label is a one-hot vector, we can derive the gradient for the classifier weights as the following:

$$\nabla_{\mathbf{w}_i} = \underbrace{\sum_{j \in \pi(i)} -(1 - p_{j,i}) \mathbf{h}_j}_{\nabla_a^{(i)}} + \underbrace{\sum_{k=1, k \neq i}^c \sum_{j \in \pi(k)} p_{j,i} \mathbf{h}_j}_{\nabla_r^{(i)}}, \quad (9)$$

where $p_{j,i} = \sigma(\mathbf{z}_{j,i})$ is the probability that sample j belongs to class i , $\pi(i)$ represents the set of sample indices that belongs to the class i . We decouple the gradients $\nabla_{\mathbf{w}_i}$ as two parts: $\nabla_a^{(i)}$ and $\nabla_r^{(i)}$. When the classifier weight \mathbf{w}_i updates, it steps along the negative gradient direction. As illustrated in 2, $\nabla_a^{(i)}$ has the inverse direction of \mathbf{h}_j and push \mathbf{w}_i and \mathbf{h}_j together, which acts like an *attraction force* from \mathbf{h}_j ; $\nabla_r^{(i)}$ has the same direction of \mathbf{h}_j and push \mathbf{w}_i and \mathbf{h}_j away, which acts like a *repulsion force* from \mathbf{h}_j .

Gradients-Imbalance. [12] observes the *Neural Collapse* phenomenon that the classifier weights and the last-layer features collapse to a c -simplex geometric structure, i.e., an equiangular tight frame (ETF), when the training data are balanced. It is reasonable and intuitive. However, for the class-imbalanced training set, *Minority Collapse* phenomenon occurs [13], i.e., the ETF vertices of

the minor classes collapse together. *Minority Collapse* limits the performance of deep learning models, especially in the minor classes. 9 reveals that the gradients are composed of the *attraction* and *repulsion* parts. To see the gradient imbalance more clearly, we perform further simplification as the following:

$$\begin{aligned} \nabla_{\mathbf{w}_i} &= \sum_{j \in \pi(i)} -q_{j,i} \mathbf{h}_j + \sum_{k=1, k \neq i}^c \sum_{j \in \pi(k)} q_{j,i} \mathbf{h}_j \\ &= \underbrace{-n_i \bar{\mathbf{h}}^{(i)}}_{\nabla_a^{(i)}} + \underbrace{\sum_{k=1, k \neq i}^c n_k \bar{\mathbf{h}}^{(k)}}_{\nabla_r^{(i)}}, \end{aligned} \quad (10)$$

where n_i is the number of samples belonging to class i and

$$\bar{\mathbf{h}}^{(i)} = \frac{1}{n_i} \sum_{j \in \pi(i)} \mathbf{h}_j \quad (11)$$

is a weighted feature mean of class i and

$$q_{j,i} = \begin{cases} 1 - p_{j,i}, & j \in \pi(i) \\ p_{j,i}, & j \notin \pi(i) \end{cases} \quad (12)$$

10 shows that the coefficients of the gradients from samples of different classes depend on the class cardinalities. When there are huge disparities among the cardinality of major classes and minor classes, the gradient norms are imbalanced, and the classifier weights are most affected by the major classes. It causes the weights corresponding to minor classes to collapse due to the *strong repulsion* effect from the features of major classes. For minority classes, the *repulsion forces* that they receive have almost the same direction. So, after training, their classifier weight vectors will be pushed in a similar direction, even merging, which explains the reason for *Minority Collapse*. We illustrate this issue in 3.

4. Methods

4.1. Balancing the gradients

Through the above analysis, the key to solving the performance drop when facing imbalanced class distribution is to balance the gradients from different class features. Our main idea to solve the gradients imbalance is to add balancing coefficients on each component of $\nabla_{\mathbf{w}_i}$. We hope $\nabla_{\mathbf{w}_i}$ has the following format:

$$\nabla_{\mathbf{w}_i} = -C(1 - p_i) \bar{\mathbf{h}}_i + \underbrace{Cp_1 \bar{\mathbf{h}}_1 + \dots + Cp_l \bar{\mathbf{h}}_{c-1}}_{\nabla_r}, \quad (13)$$

where C is a positive constant.

4.2. ARB-Loss

To make the gradients $\nabla_{\mathbf{w}_i}$, $i \in \{1, \dots, c\}$ satisfy the format of 13, we design a novel classification loss function. We name it




Fig. 2. The attraction and repulsion components of the gradients. For class i , the gradients from other classes (denoted as $\nabla_{b,j}^{(i)}$, $j \neq i$) act like *repulsion force* to push the weights \mathbf{w}_i to their inverse directions, while the gradients from class i (denoted as $\nabla_a^{(i)}$) acts like a *attraction force*. The imbalance among different components of $\nabla_r^{(i)}$ will cause the direction of \mathbf{w}_i to incline to the inverse directions of major-class features. The imbalance between $\nabla_a^{(i)}$ and $\nabla_r^{(i)}$ will cause the norm of \mathbf{w}_i to incline towards the direction of *attraction* or *repulsion*. (The thick solid arrows represent the classifier weights, the thin solid arrows represent the features, the thin dashed arrows represent the gradients, and the thin dashed arrows with the solid head represent the sum of gradients from each class. Each color represents a class, and the thin red dashed arrow represents the sum of all gradients.).




Fig. 3. The cause of Minority Collapse is the imbalance of gradients. In this figure, we draw 4-class classification, where class 1, 2 (blue, green) are major classes and 3, 4 (orange, purple) are minor classes. We describe the updating process of the classifier weights of minor classes, i.e., $\mathbf{w}_3, \mathbf{w}_4$. For those minor classes, their gradients are mainly decided by the inverse direction of features of major classes. It seems that the major components of gradients have *strong repulsion* to push the weights corresponding to minor classes to collapse together. (The color and arrow type have the same meaning as 2.).

Attraction–Repulsion–Balanced Loss, abbreviated as ARB-Loss. ARB-Loss is formulated as follows:

$$\mathcal{L}_{arb} = - \sum_{j=1}^b \sum_{i=1}^c \mathbf{y}_{j,i} \log \left(\frac{\exp(\mathbf{w}_i^T \mathbf{h}_j)}{\sum_{k=1}^c \frac{n_k}{n_i} \exp(\mathbf{w}_k^T \mathbf{h}_j)} \right), \quad (14)$$

where n_i is the number of samples belonging to class i . Note that n_i can represent the class cardinality in the whole training dataset or the mini-batch. The difference between ARB-Loss and traditional cross-entropy loss is that we add coefficients on the denominator of the softmax function.

After simple derivation, we can obtain the gradients for the classifier weights under our ARB-Loss.

$$\begin{aligned} \tilde{\nabla} \mathbf{w}_i &= \sum_{j \in \pi(i)} - (1 - \tilde{\sigma}_i(\mathbf{w}_i^T \mathbf{h}_j)) \mathbf{h}^{(j)} \\ &\quad + \sum_{k=1, k \neq i}^c \sum_{j \notin \pi(k)} \frac{n_i}{n_k} \cdot \tilde{\sigma}_k(\mathbf{w}_k^T \mathbf{h}_j) \mathbf{h}^{(j)} \\ &= \sum_{j \in \pi(i)} - \tilde{q}_{j,i}^{(i)} \mathbf{h}^{(j)} + \sum_{k=1, k \neq i}^c \sum_{j \notin \pi(k)} \frac{n_i}{n_k} \tilde{q}_{j,i}^{(k)} \mathbf{h}^{(j)} \\ &= \underbrace{-n_i \tilde{\mathbf{h}}_{a(i)}^{(i)}}_{\tilde{\nabla}_a^{(i)}} + \underbrace{\sum_{k=1, k \neq i}^c n_i \tilde{\mathbf{h}}_{r(i)}^{(k)}}_{\tilde{\nabla}_r^{(i)}}, \end{aligned} \quad (15)$$

where

$$\tilde{q}_{j,i}^{(\alpha)} = \begin{cases} 1 - \tilde{p}_{j,i}^{(\alpha)}, & j \in \pi(i) \\ \tilde{p}_{j,i}^{(\alpha)}, & j \notin \pi(i) \end{cases}, \quad (16)$$

where $\tilde{p}_{j,i}^{(\alpha)} = \tilde{\sigma}_\alpha(\mathbf{w}_i^T \mathbf{h}_j) = \frac{\exp(\mathbf{w}_i^T \mathbf{h}_j)}{\sum_{k=1, k \neq i}^c \frac{n_k}{n_i} \exp(\mathbf{w}_k^T \mathbf{h}_j)}$ is the corresponding softmax function in our case. Note that the subscript α of $\tilde{\sigma}_\alpha$ identifies the ground-truth class of the sample j . From 15, we can see that the attraction component $\tilde{\nabla}_a^{(i)}$ and each parts of the repulsion component $\tilde{\nabla}_r^{(i)}$ have the same coefficients, i.e., attraction and repulsion terms are balanced.

4.3. Analytical results

In this section, we further explore the properties of our ARB-Loss to show its superiority over the Cross-Entropy loss when training on imbalanced data.

We introduce two propositions to demonstrate that ARB-Loss can efficiently reduce the imbalance of the gradients from different classes.

Proposition 4.1. For gradient $\nabla_{\mathbf{w}_i}$ in cross entropy and $\tilde{\nabla}_{\mathbf{w}_i}$ in ARB-Loss, the proportion of the supremums of different repulsion terms' norm, $\frac{\sup \|\nabla_{b,k}\|}{\sup \|\nabla_{b,l}\|}$ and $\frac{\sup \|\tilde{\nabla}_{b,k}\|}{\sup \|\tilde{\nabla}_{b,l}\|}$ (where $k \neq l, k \neq i, l \neq i$), obey:

$$\begin{aligned} \text{CELoss} : \frac{\sup \|\nabla_{b,k}\|}{\sup \|\nabla_{b,l}\|} &\leq \frac{n_k}{n_l} \cdot C_{kl} \cdot \sqrt{\frac{E_{H,k}}{E_{H,l}}}, \\ \text{ARB-Loss} : \frac{\sup \|\tilde{\nabla}_{b,k}\|}{\sup \|\tilde{\nabla}_{b,l}\|} &\leq \tilde{C}_{kl} \cdot \sqrt{\frac{E_{H,k}}{E_{H,l}}}, \end{aligned} \quad (17)$$

where $C_{kl}, \tilde{C}_{kl} \in \mathbb{R}$ are constants.

Proposition 4.2. For gradient $\nabla_{\mathbf{w}_i}$ in cross entropy and $\tilde{\nabla}_{\mathbf{w}_i}$ in ARB-Loss, the proportions of the supremums of the attraction and repulsion terms' norm, $\frac{\sup \|\nabla_a^{(i)}\|}{\sup \|\nabla_{b,k}\|}$ and $\frac{\sup \|\tilde{\nabla}_a^{(i)}\|}{\sup \|\tilde{\nabla}_{b,k}\|}$ (where $k \neq i$), obey:

$$\text{CELoss} : \frac{\sup \|\nabla_a^{(i)}\|}{\sup \|\nabla_{b,k}\|} \leq \frac{n_i}{n_k} \cdot C_{ik} \cdot \sqrt{\frac{E_{H,i}}{E_{H,k}}}, \quad (18)$$

$$\text{ARB-Loss} : \frac{\sup \|\tilde{\nabla}_a^{(i)}\|}{\sup \|\tilde{\nabla}_{b,k}\|} \leq \tilde{C}_{ik} \cdot \sqrt{\frac{E_{H,i}}{E_{H,k}}},$$

where $C_{ik}, \tilde{C}_{ik} \in \mathbb{R}$ are constants.

4.1 shows that ARB-Loss can reduce the imbalance of the supremum of different repulsion gradient terms. The balanced repulsion gradients can remedy the issue that the classifier weights of minority classes are pushed to the inverse of the majority class weights and collapse together. 4.2 shows that ARB-Loss can reduce the imbalance between the attraction and arbitrary repulsion terms. It is important if the current class i is a minority class. If they are imbalanced in this case, the attraction force is much less than the repulsion force from majority classes and the attraction is not enough to fight against the strong repulsion that pushes the minority class features together.

Besides balancing the attraction and repulsion components of the gradients, compared to previous loss functions designed specifically for long-tailed learning, our ARB-Loss has the following advantages:

1. ARB-Loss can achieve comparable or even better results via only one-stage training, compared to previous two-stage methods [7,9,11];
2. ARB-Loss has no hyperparameters to adjust according to different data sources like [17,6–8].

5. Experiments

In this section, we perform experiments on classification and segmentation tasks to analyze our ARB-Loss. We put the description of datasets used in this section and the implementation details in the B and C.

5.1. Empirical results

5.1.1. The geometric structure of the classifier weights

We propose three balance metrics to measure how balanced the classifier weights are, i.e., how similar the classifier weights are with the ideal ETF geometric structure described by the Neural Collapse phenomenon.

$$\begin{aligned} \mathcal{B}_{\mathcal{D}}^2(\mathbf{W}) &= \frac{1}{(c-1)^2} \sum_{i=1}^c \sum_{j=1, j \neq i}^c \left(\mathbf{w}_i^T \mathbf{w}_j - \frac{1}{(c-1)^2} \sum_{i=1}^c \sum_{j=1, j \neq i}^c \mathbf{w}_i^T \mathbf{w}_j \right)^2 \\ \mathcal{B}_{\mathcal{D}'}^2(\mathbf{W}) &= \frac{1}{(c-1)^2} \sum_{i=1}^c \sum_{j=1, j \neq i}^c \left(\langle \mathbf{w}_i, \mathbf{w}_j \rangle - \frac{1}{(c-1)^2} \sum_{i=1}^c \sum_{j=1, j \neq i}^c \langle \mathbf{w}_i, \mathbf{w}_j \rangle \right)^2 \\ \mathcal{B}_{\mathcal{D}''}^2(\mathbf{W}) &= \frac{1}{c} \sum_{i=1}^c \left(\|\mathbf{w}_i\|_2 - \frac{1}{c} \sum_{i=1}^c \|\mathbf{w}_i\|_2 \right)^2 \end{aligned} \quad (19)$$

where $\langle \mathbf{w}_i, \mathbf{w}_j \rangle = \frac{\mathbf{w}_i^T \mathbf{w}_j}{\|\mathbf{w}_i\|_2 \cdot \|\mathbf{w}_j\|_2}$ is the cosine of the angle between \mathbf{w}_i and \mathbf{w}_j . $\mathcal{B}_{\mathcal{D}}^2(\mathbf{W})$ measures the balance degree of the angles of the weight vector pairs, $\mathcal{B}_{\mathcal{D}'}^2(\mathbf{W})$ measures the variance of the norm of the weight vectors, and $\mathcal{B}_{\mathcal{D}''}^2(\mathbf{W})$ measures the balance degree of the dot-product of the weight vector pairs. In the ideal case, i.e., the ETF geometric structure, the angles of all weight vector pairs, and the norms of all weight vectors are equivalent.

For better understanding, we visualize the similarity of the classifier weights in 4. In the figure, each sub-figures represents a matrix $\mathbf{M} \in \mathbb{R}^{c \times c}$ whose i, j -th entry is:

$$\mathbf{M}_{ij} = \begin{cases} \frac{\mathbf{w}_i^\top \mathbf{w}_j}{c}, & i \neq j \\ \sum_{j=1, j \neq i}^c \mathbf{w}_i^\top \mathbf{w}_j, & \\ 0, & i = j \end{cases}, \quad (20)$$

Because the diagonal of matrix \mathbf{M} is meaningless, we set them to zero. Below each sub-figure, we also indicate the values of the three balance metrics mentioned above. In the ideal state, we expect the variance of each row (except for the diagonal) of \mathbf{M} to be zeros:

$$\mathbf{M}_{ij} = \frac{1}{c-1}, \quad i, j \in \{1, \dots, c\}, i \neq j, \quad (21)$$

The ideal values of the balance metrics are:

$$\mathcal{B}_{\mathcal{L}}^2(\mathbf{W}) = \mathcal{B}_{\mathcal{D}}^2(\mathbf{W}) = \mathcal{B}_{\mathcal{E}}^2(\mathbf{W}) = 0 \quad (22)$$

In the figure, we can see that the color variance of the left part (light and dark colors) is larger than the right part (all dark colors), which demonstrates that the classifier weights trained by ARB-Loss are more dispersed than the ones trained by cross-entropy loss. ARB-Loss can remarkably reduce the variance of each row of \mathbf{M} , which means that our ARB-Loss can mitigate the *Minority Collapse* dilemma.

5.1.2. Different components of gradients

We show a visualization example in 5. Each point on the original curves (the *thin* line in the figure) represents the mean lengths of the gradients in one epoch, *i.e.*,

$$\mathbf{g}_k = \frac{1}{N/B} \sum_{n=0}^{N/B} \|\nabla_{\mathbf{w}_k}^{(9)}\| = \frac{1}{N/B} \sum_{n=0}^{N/B} \left\| \sum_{j \in \pi(k)} \frac{\partial \mathcal{L}^{(j)}}{\partial \mathbf{w}_k} \right\|, \quad (23)$$

where N is the total number of samples, B is the batch size, $\frac{\partial \mathcal{L}^{(j)}}{\partial \mathbf{w}_k}$ is the gradients of the loss on j -th sample with respect to \mathbf{w}_k . From the visualization, we can see that using cross-entropy loss on imbalanced data will introduce huge magnitude differences among different gradient components, and using our proposed ARB-Loss can balance each gradient component to make them have a similar magnitude.

5.2. Results on image classification task

Experiment results are shown in 1 and 2. For previous methods, we directly copy the results from the original papers, except for the results noted with †.

CIFAR-10(-LT) & CIFAR-100(-LT). 1 shows the results on CIFAR-10/100-LT. When performing experiments on CIFAR-10/100-LT, the class cardinalities used in ARB-Loss, are collect in a batch, *i.e.*, the $n_i, i \in \{1, \dots, c\}$ in 14 are dynamic for each different batches. Therefore, when the dataset is balanced, ARB-Loss *approximately* degrades as the cross-entropy loss. The reason why it is said *approximately* is that the samples are not balanced in a batch, although the whole training set is balanced. Therefore, we compare ARB-Loss with the plain model (using traditional cross-entropy) in the *Normal* column. From the results, although the improvements achieved by our ARB-Loss are slight, it is also significant, considering that our method only needs one-stage training while recent state-of-the-art long-tailed works need 2-stage training. Moreover, the results on *Normal* case suggest that our ARB-Loss is approximate to the traditional cross-entropy loss, although ARB-Loss is dynamic for every mini-batch.

Large-scale Image Classification Datasets. We list the results on the ImageNet-LT, Places-LT and iNaturalist2018 in 2. On the large-scale datasets, we found that using the global class cardinalities (*i.e.*, the class cardinalities in the whole training set) can obtain better performance than using dynamic class cardinalities for each batch. Therefore, for the ARB-Loss experiments in 2, we all use the global class cardinalities. The results suggest that our ARB-Loss can achieve comparable results with current state-of-the-art methods, *e.g.* MiSLAS [11]. It is worth noting that the results of ARB-Loss are obtained from only one-stage training, and the nowadays SOTA long-tail methods such as. BBN [10], MiSLAS [11], need training for 2 stages.

5.3. Results on segmentation task

We perform image segmentation experiments on two baselines: U-Net [33] and Deeplabv3+ [34]. The results are shown in 3.

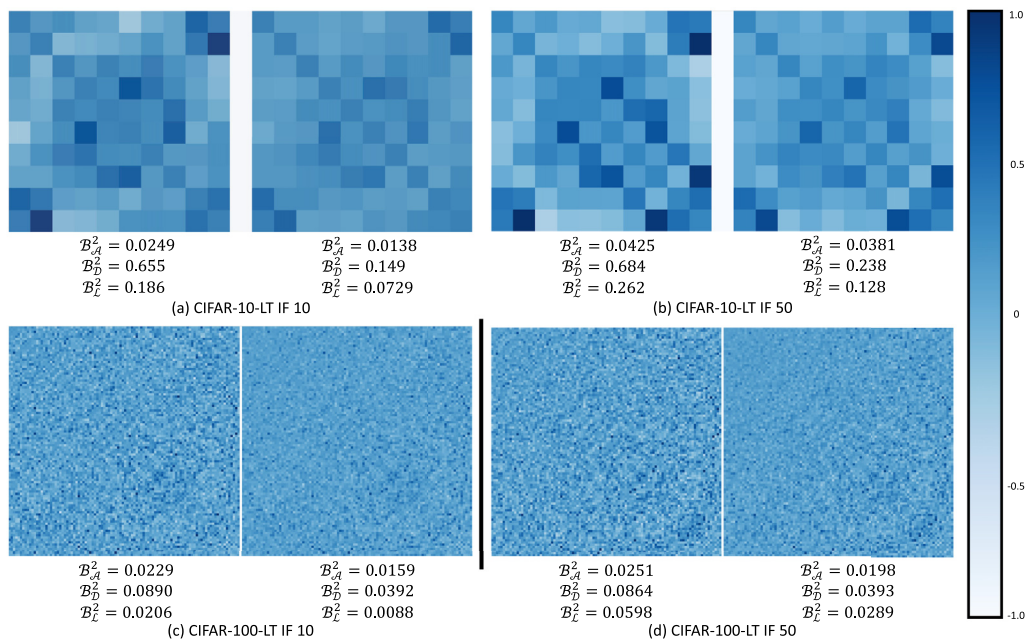


Fig. 4. The geometric structure of the classifier weights. In the (a), (b), (c), and (d) 4 sub-figures, the left part represents the classifier trained with cross-entropy loss, and the right part represents the classifier trained with ARB-Loss. The visualization details and meanings are elaborated in 5.1. The color variance of the left part is larger than the right part, which demonstrates that the classifier weights trained by ARB-Loss are more dispersed than the ones trained by cross-entropy loss.




Fig. 5. Visualization example of the different components of the gradients. In the figure, we take the training on CIFAR-10-LT dataset as an example. We plot the curve of magnitude (i.e., the vector length) of the different components $\nabla_{b,i}^{(9)}$, $i \neq 9$ of the gradients ∇_{w_0} , where w_9 the classifier weights for class 9 whose cardinality is the least. The thin and thick lines represent the original and smoothed curves, respectively. We show the results from the 40-th epoch to ignore the instability in the initial training.

Table 1

Top-1 accuracy (%) on CIFAR-10(-LT) and CIFAR-100(-LT). The experiments are performed on ResNet-32 [27]. In the table head, 100, 50, 10 represents IF = 100, 50, 10. Normal represents the original dataset without undersampling. † represents our reimplement results.

Methods	stage	CIFAR-10(-LT)				CIFAR-100(-LT)			
		100	50	10	Normal	100	50	10	Normal
Plain Model [27]	1	70.4	74.8	86.4	92.8 [†]	38.4	43.9	55.8	68.7 [†]
mixup [28]	1	73.1	77.8	87.1	92.0 [†]	39.6	45.0	58.2	68.7 [†]
LDAM + DRW [7]	1	74.9	-	86.7	-	40.3	-	57.3	-
Remix + DRW [29]	2	79.8	-	89.1	-	46.8	-	61.3	-
BBN [10]	2	79.9	82.2	88.4	-	42.6	47.1	59.2	-
MiSLAS [11]	2	82.1	85.7	90.0	-	47.0	52.3	63.2	-
ARB-Loss	1	83.3	85.7	90.2	92.6	47.2	52.6	62.1	68.6

Table 2

Top-1 accuracy (%) on ImageNet-LT, Places-LT and iNaturalist2018. The experiments of ImageNet-LT and iNaturalist2018 are performed on ResNet-50, and the ones of Places-LT are performed on ResNet-152. † represents our reimplement results.

Methods	Stage	ImageNet-LT				Places-LT				iNaturalist2018			
		Many	Median	Few	Overall	Many	Median	Few	Overall	Many	Median	Few	Overall
Plain Model [27]	1	65.5 [†]	40.3 [†]	13.1 [†]	45.6 [†]	-	-	-	-	72.8 [†]	63.7 [†]	58.4 [†]	62.3 [†]
CB-Focal [6]	1	-	-	-	-	-	-	-	-	-	-	-	61.1
RangeLoss [30]	1	-	-	-	-	41.1	35.4	23.2	35.1	-	-	-	-
Focal + DRW [17]	1	-	-	-	47.9	-	-	-	-	-	-	-	-
CE + DRW [7]	1	-	-	-	48.5	-	-	-	-	-	-	-	-
LDAM + DRW [7]	1	-	-	-	48.8	-	-	-	-	-	-	-	68.0
OLTR [31]	1	43.2	35.1	18.5	35.6	44.7	37.0	25.3	35.9	-	-	-	-
OTLR + LFME [32]	1	47.0	37.9	19.2	38.8	39.3	39.6	24.2	36.2	-	-	-	-
Remix + DRW [29]	2	-	-	-	-	-	-	-	-	-	-	-	70.5
BBN [10]	2	-	-	-	-	-	-	-	-	-	-	-	69.6
cRT + mixup [11]	2	63.9	49.1	30.2	51.7	44.1	38.5	27.1	38.3	74.2	71.1	68.2	70.2
LWS + mixup [11]	2	62.9	49.8	31.6	52.0	41.7	41.3	33.1	39.7	72.8	71.6	69.8	70.9
MiSLAS [11]	2	61.7	51.3	35.8	52.7	39.6	43.3	36.1	40.4	73.2	72.4	70.4	71.6
ARB-Loss	1	60.2	51.8	38.3	52.8	41.9	41.5	32.1	39.7	71.9	72.1	71.7	71.7

Our ARB-Loss can significantly improve pixel accuracy. The reason why the improvements on *mIoU* metric are slight might be that the denominator of IoU involved FP (*false positive*), compared to the

Table 3

The Segmentation Results on CityScapes dataset. *mAcc* is the mean of per-class pixel accuracy and *mIoU* is the mean of per-class IoU (Intersection over Union). † represents our reimplement results.

Models	U-Net		DeepLabv3+	
	mIoU	mAcc	mIoU	mAcc
baseline	67.6 [†]	74.6 [†]	79.3 [†]	86.3 [†]
ARB-Loss	67.9	79.7	79.5	89.0

denominator of pixel accuracy. In 4, we show the per-class results. We can see that the classes whose IoU gets improved while the pixel accuracy is still a little lower are almost the minor classes. If the FP item is relatively too large for the cardinality of a minor class, the IoU value will drop. Meanwhile, if these FP belongs to some major classes, it will have little effect on the value of pixel accuracy.

6. Conclusion & future works

In this paper, we analyze that the reason for the performance drop under long-tailed distributions is the imbalances of the gradients from different classes. Then, we propose Attraction-Repul

Table 4
The detailed results on CityScapes [39] dataset based on tow backbone networks: U-Net [33] and Deeplabv3+ [34].

Backbone	Loss	Metric	road	sidewalk	building	wall	fence	pole	light	sign	vegetation	terrain	sky	person	rider	car	truck	bus	train	motor	bicycle	mean
U-Net	ARB- Loss	IoU	97.01	78.92	92.83	39.73	50.86	61.05	70.68	75.34	90.98	55.96	92.83	78.12	55.11	95.46	38.71	59.17	33.86	49.98	73.33	67.89
		Acc	97.86	88.57	93.37	48.4	60.18	82.74	94.05	91.16	94.41	98.96	72.41	98.96	93.6	65.77	96.75	48.21	67.5	60.94	69.24	89.93
	Cross-Entropy	IoU	96.82	77.25	90.05	31.86	45.77	63.33	73.74	78.43	91.89	57.94	94.38	79.51	48.82	92.97	33.63	62.8	34.65	53.16	76.63	67.56
		Acc	98.61	86.47	96.33	35.54	52.81	73.4	83.91	85.31	96.35	68.11	97.97	92.0	55.94	97.44	36.51	70.04	40.91	61.39	87.48	74.55
deeplabv3+	ARB- Loss	IoU	97.99	89.97	91.93	58.15	64.64	67.51	73.46	78.18	91.91	63.1	93.16	82.36	72.08	96.62	75.61	86.64	83.97	64.86	78.24	79.49
		Acc	98.49	93.11	94.62	62.26	72.19	87.95	95.64	94.71	94.89	80.4	99.3	95.38	80.06	97.87	85.26	92.52	90.28	82.82	92.33	88.95
	Cross-Entropy	IoU	98.31	86.12	93.1	54.84	61.31	69.62	75.11	81.82	92.9	65.51	95.2	84.08	63.63	95.87	77.36	88.21	75.41	68.68	80.17	79.33
		Acc	99.08	93.07	96.82	59.6	69.51	81.66	86.41	89.41	96.83	74.68	98.37	92.93	74.01	98.12	84.8	96.04	78.08	79.29	90.54	86.28

sion-Balanced Loss (ARB-Loss) to balance the different gradient components. Moreover, we analyze the properties of our ARB-Loss to theoretically demonstrate the rationality of our design. We also give some visualization examples to analyze the geometric structure of the classifier weights and the magnitude proportions of the gradient components. Moreover, we perform experiments on large-scale classification and segmentation datasets to demonstrate the effectiveness of ARB-Loss. ARB-Loss can achieve state-of-the-art performance via only one-stage training.

For the future, the following topics are worth studying:

- The geometric structure described by *Neural Collapse* also has great inspiration for the bottlenecks of other machine-learning tasks, *i.e.*, few-shot or online learning *etc.* *How to design algorithms inspired by Neural Collapse for the other tasks is a topic worth future research.*
- *Besides the classifier, the optimization of the features extracted by the backbone networks is also affected by the imbalanced distribution. How to solve this issue is still an open question.*

In the future, we will continue the relevant research topics along the line of *Neural Collapse* and explore the underlying principles behind it.

Data availability

Data will be made available on request.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported in part by The National Key Research and Development Program of China (Grant Nos: 2018AAA0101400), in part by The National Nature Science Foundation of China (Grant Nos: 62036009, U1909203, 61936006, 62133013), in part by Innovation Capability Support Program of Shaanxi (Program No. 2021TD-05).

Appendix A. Proofs

A.1. Proof of 4.1

Proof 1. For k -th($k \neq i$) repulsion terms,

$$\begin{aligned}
 \|\nabla_{b,k}\| &= \left\| \sum_{j \in \pi(k)} p_{j,i} \mathbf{h}_j \right\| \\
 &\leq \sum_{j \in \pi(k)} p_{j,i} \|\mathbf{h}_j\| \\
 &\leq \sqrt{\sum_{j \in \pi(k)} p_{j,i}^2} \sqrt{\sum_{j \in \pi(k)} \|\mathbf{h}_j\|^2} \\
 &= B_k \sqrt{n_k \|\mathbf{h}_{(k)}\|^2} \leq B_k \sqrt{n_k E_{H,k}}
 \end{aligned}
 \tag{24}$$

where

$$B_k = \sqrt{\sum_{j \in \pi(k)} p_{j,i}^2}
 \tag{25}$$

and

$$\frac{\sum_{j \in \pi(i)} \|\mathbf{h}_j\|^2}{\|\mathbf{h}_{(i)}\|^2} = \frac{\sum_{j \in \pi(i)} \|\mathbf{h}_j\|^2}{n_i} \leq E_{H,i}. \quad (26)$$

Note that we use the subscript with a bracket to represent the mean of the squared norm of some certain class features. The brackets are used to be distinguished from the feature indices.

For $j \in \pi(k)$, $k \neq i$, there is a non-zero minimum for $p_{j,i}$ which is denoted as $p_{\min}^{(k,i)}$ and a maximum for $p_{j,i}$ which is less than 1 and is denoted as $p_{\max}^{(k,i)}$. Then,

$$\sqrt{n_k} p_{\min}^{(k,i)} \leq B_k \leq \sqrt{n_k} p_{\max}^{(k,i)} \quad (27)$$

Therefore,

$$\frac{\sup \|\nabla_{b,k}\|}{\sup \|\nabla_{b,l}\|} = \frac{B_k \sqrt{n_k E_{H,k}}}{B_l \sqrt{n_l E_{H,l}}} \leq \frac{n_k}{n_l} \cdot C_{kl} \cdot \sqrt{\frac{E_{H,k}}{E_{H,l}}} \quad (28)$$

where $C_{kl} = p_{\max}^{(k,i)} / p_{\min}^{(k,i)}$ is a constant.

Similarly, for the ARB-Loss,

$$\begin{aligned} \|\tilde{\nabla}_{b,k}\| &= \left\| \frac{n_i}{n_k} \sum_{j \in \pi(k)} \tilde{p}_{j,i} \mathbf{h}_j \right\| \\ &\leq \frac{n_i}{n_k} \sum_{j \in \pi(k)} \tilde{p}_{j,i} \|\mathbf{h}_j\| \\ &\leq \frac{n_i}{n_k} \sqrt{\sum_{j \in \pi(k)} \tilde{p}_{j,i}^2} \sqrt{\sum_{j \in \pi(k)} \|\mathbf{h}_j\|^2} \\ &= \frac{n_i}{n_k} \cdot \tilde{B}_k \sqrt{n_k \|\mathbf{h}_{(k)}\|^2} \leq \frac{n_i}{n_k} \cdot \tilde{B}_k \sqrt{n_k E_{H,k}} \end{aligned} \quad (29)$$

where,

$$\tilde{p}_{\min}^{(k,i)} \leq \tilde{B}_k = \sqrt{\sum_{j \in \pi(k)} \tilde{p}_{j,i}^2} \leq \tilde{p}_{\max}^{(k,i)} \quad (30)$$

where $\tilde{p}_{\max}^{(k,i)}$ and $\tilde{p}_{\min}^{(k,i)}$ are defined similarly with $p_{\max}^{(k,i)}$ and $p_{\min}^{(k,i)}$, i.e., $\tilde{p}_{\max}^{(k,i)} = \max_{j \in \pi(i)} \tilde{p}_{j,i}$, $\tilde{p}_{\min}^{(k,i)} = \min_{j \in \pi(k)} \tilde{p}_{j,i}$. Then, we have,

$$\frac{\sup \|\tilde{\nabla}_{b,k}\|}{\sup \|\tilde{\nabla}_{b,l}\|} = \frac{\frac{n_i}{n_k} \cdot \tilde{B}_k \sqrt{n_k E_{H,k}}}{\frac{n_i}{n_l} \cdot \tilde{B}_l \sqrt{n_l E_{H,l}}} \leq \tilde{C}_{kl} \cdot \sqrt{\frac{E_{H,k}}{E_{H,l}}} \quad (31)$$

where $\tilde{C}_{kl} = \tilde{p}_{\max}^{(k,i)} / \tilde{p}_{\min}^{(k,i)}$ is a constant.

A.2. Proof of 4.2

Proof 2.

$$\begin{aligned} \|\nabla_a^{(i)}\| &= \left\| \sum_{j \in \pi(i)} (1 - p_{j,i}) \mathbf{h}_j \right\| \\ &\leq \sum_{j \in \pi(i)} (1 - p_{j,i}) \|\mathbf{h}_j\| \\ &\leq \sqrt{\sum_{j \in \pi(i)} (1 - p_{j,i})^2} \sqrt{\sum_{j \in \pi(i)} \|\mathbf{h}_j\|^2} \\ &= A_k \sqrt{n_i \|\mathbf{h}_{(i)}\|^2} \leq A_k \sqrt{n_i E_{H,i}}, \end{aligned} \quad (32)$$

where

$$A_k = \sqrt{\sum_{j \in \pi(i)} (1 - p_{j,i})^2}. \quad (33)$$

Similar to the proof of 4.2,

$$\|\nabla_{b,k}^{(i)}\| = \left\| \sum p_{j,ij \in \pi(k)} \mathbf{h}_j \right\| \leq B_k \sqrt{n_k \|\mathbf{h}_{(k)}\|^2} \leq B_k \sqrt{n_k E_{H,k}}, \quad (34)$$

For $j \in \pi(i)$, the maximum of $p_{j,i}$ is less than 1 and we denote this maximum as $p_{\max}^{(i)}$. For $j \in \pi(k)$, $k \neq i$, there is a non-zero minimum for $p_{j,i}$ and we denote this minimum as $p_{\min}^{(k,i)}$. Then we have,

$$A_k \geq \sqrt{n_i} (1 - p_{\max}^{(i)}), \quad B_k \leq \sqrt{n_k} p_{\min}^{(k,i)}. \quad (35)$$

Therefore, the proportion of the supremum of the attraction and k -th($k \neq i$) repulsion terms' norm obey:

$$\frac{\sup \|\nabla_a^{(i)}\|}{\sup \|\nabla_{b,k}^{(i)}\|} = \frac{A_k \sqrt{n_i \|\mathbf{h}_{(i)}\|^2}}{B_k \sqrt{n_k \|\mathbf{h}_{(k)}\|^2}} \leq \frac{n_i}{n_k} \cdot C_k \cdot \sqrt{\frac{E_{H,i}}{E_{H,k}}}, \quad (36)$$

where $C_k = (1 - p_{\max}^{(i)}) / p_{\min}^{(k,i)}$ is a constant.

Similarly, for the ARB-Loss,

$$\begin{aligned} \|\tilde{\nabla}_a^{(i)}\| &= \left\| \sum_{j \in \pi(i)} (1 - \tilde{p}_{j,i}) \mathbf{h}_j \right\| \\ &\leq \sum_{j \in \pi(i)} (1 - \tilde{p}_{j,i}) \|\mathbf{h}_j\| \\ &\leq \sqrt{\sum_{j \in \pi(i)} (1 - \tilde{p}_{j,i})^2} \sqrt{\sum_{j \in \pi(i)} \|\mathbf{h}_j\|^2} \\ &= \tilde{A}_k \sqrt{n_i \|\mathbf{h}_{(i)}\|^2} \leq \tilde{A}_k \sqrt{n_i E_{H,i}} \\ \|\tilde{\nabla}_{b,k}\| &= \left\| \frac{n_i}{n_k} \sum_{j \in \pi(k)} \tilde{p}_{j,i} \mathbf{h}_j \right\| \\ &\leq \frac{n_i}{n_k} \sum_{j \in \pi(k)} \tilde{p}_{j,i} \|\mathbf{h}_j\| \\ &\leq \frac{n_i}{n_k} \sqrt{\sum_{j \in \pi(k)} \tilde{p}_{j,i}^2} \sqrt{\sum_{j \in \pi(k)} \|\mathbf{h}_j\|^2} \\ &= \frac{n_i}{n_k} \cdot \tilde{B}_k \sqrt{n_k \|\mathbf{h}_{(k)}\|^2} \leq \frac{n_i}{n_k} \cdot \tilde{B}_k \sqrt{n_k E_{H,k}} \end{aligned} \quad (37)$$

where,

$$\begin{aligned} \tilde{A}_k &= \sqrt{\sum_{j \in \pi(i)} (1 - \tilde{p}_{j,i})^2} \geq \sqrt{n_i} (1 - \tilde{p}_{\max}^{(i)}) \\ \tilde{B}_k &= \sqrt{\sum_{j \in \pi(k)} \tilde{p}_{j,i}^2} \leq \sqrt{n_k} \tilde{p}_{\max}^{(k,i)} \end{aligned} \quad (38)$$

where $\tilde{p}_{\max}^{(i)}$ and $\tilde{p}_{\min}^{(k,i)}$ are defined similarly with $p_{\max}^{(i)}$ and $p_{\min}^{(k,i)}$, i.e., $\tilde{p}_{\max}^{(i)} = \max_{j \in \pi(i)} \tilde{p}_{j,i}$, $\tilde{p}_{\min}^{(k,i)} = \min_{j \in \pi(k)} \tilde{p}_{j,i}$. Therefore,

$$\frac{\sup \|\tilde{\nabla}_a^{(i)}\|}{\sup \|\tilde{\nabla}_{b,k}\|} = \frac{\tilde{A}_k \sqrt{n_i \|\mathbf{h}_{(i)}\|^2}}{\frac{n_i}{n_k} \cdot \tilde{B}_k \sqrt{n_k \|\mathbf{h}_{(k)}\|^2}} \leq \tilde{C}_k \cdot \sqrt{\frac{E_{H,i}}{E_{H,k}}}, \quad (39)$$

where $\tilde{C}_k = (\tilde{p}_{\max}^{(i)}) / \tilde{p}_{\min}^{(k,i)}$ is a constant.

Appendix B. Datasets

B.1. Classification datasets

We perform experiments on CIFAR-10/100 [35] and three large-scale image classification datasets: ImageNet [36], iNaturalist2018 [37] and Places [38].

CIFAR-10/100 are both composed of a training set of 50000 images and a validation set of 10000 images with 10 and 100 classes, respectively. The long-tail versions of CIFAR-10/100 are generated by the same rule as [7], i.e., reducing the sample number per class according to an exponential curve controlled by an imbalance factor $IF = \frac{n_{\max}}{n_{\min}}$, where n_{\max} and n_{\min} are the sample numbers of the most and the least frequent classes.

ImageNet-LT and **Places-LT** were first proposed by [31]. ImageNet is composed of 1,281,167 training images, 50,000 validation images and 100,000 test images and spans 1000 classes. Places contains 10,624,928 images from 434 classes. **ImageNet-LT** has 115.8K images from 1000 classes whose cardinalities ranging from 5 to 1,280. **Places-LT** contains 184.5K images from 365 classes with class cardinality ranging from 5 to 4,980.

iNaturalist2018 is a dataset that originally suffers from an extreme class imbalance, which consists of 437.5K images from 8,142 classes. Besides, we also face the fine-grained classification in iNaturalist2018 dataset.

B.2. Segmentation datasets

Essentially, the segmentation is a pixel-wise classification. And it usually suffers from extremely imbalanced distributions, e.g. the sky and ground usually occupy much more pixels than foreground objects on the image captured by the car-equipped camera.

CityScapes [39] contains 5000 annotated images of the street scenes of 50 different cities and 19 semantic labels are used for evaluation. The training, validation, and test sets contain 2975, 500, and 1525 images, respectively.

Appendix C. Implementation details

We use the PyTorch [40] framework to perform all the experiments. For all experiments on the image classification task, we use an SGD optimizer with a momentum of 0.9, and an initial learning rate of 0.1. For experiments on CIFAR-10/100(-LT), we train the model on a single GPU, and the learning rate decays as its 0.1 at the 160-th and 180-th epochs. For ImageNet-LT, Places-LT, and iNaturalist2018 datasets, we train the models on 8 GPUs and use cosine learning rate schedule [41] gradually decaying from 0.2 to 0. For all experiments, we use mixup [28] data augmentation, and the batch size is 128 per GPU. For experiments on the image segmentation task, we use an SGD optimizer with an initial learning rate of 0.01, a momentum of 0.9, and a weight decay of 0.0005. And we update the learning rate by the poly policy with a power of 0.9.

Appendix D. Detailed results of image segmentation task

From the per-class results in the above table, we can see that the classes whose IoU gets improved while the pixel accuracy is still a little lower are almost the minor classes. The reason might be the differences in the definition of these two metrics. For a certain class i , IoU is defined as:

$$\text{IoU}_i = \frac{\text{TP}_i}{\text{TP}_i + \text{FP}_i + \text{FN}_i} \quad (40)$$

And pixel accuracy is defined as:

$$\text{Acc}_i = \frac{\text{TP}_i}{\text{TP}_i + \text{FN}_i} \quad (41)$$

If the FP_i item is relatively too large for the cardinality of a minor class, i.e., FP_i is too large compared to $\text{TP}_i + \text{FN}_i$, the IoU value will drop a lot. Meanwhile, if these FP_i pixels belongs to some major classes, it will have little effect on the pixel accuracy of these major classes. So, it will cause that although pixel accuracy improves a lot, the improvements in the metric of IoU are slight.

References

[1] N.V. Chawla, K.W. Bowyer, L.O. Hall, W.P. Kegelmeyer, Smote: synthetic minority over-sampling technique, *J. Artif. Intell. Res.* 16 (2002) 321–357.

[2] M. Buda, A. Maki, M.A. Mazurowski, A systematic study of the class imbalance problem in convolutional neural networks, *Neural Networks* 106 (2018) 249–259.

[3] C. Huang, Y. Li, C.C. Loy, X. Tang, Learning deep representation for imbalanced classification, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5375–5384.

[4] S. Rota Bulò, G. Neuhof, P. Kotschieder, Loss max-pooling for semantic image segmentation, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2126–2135.

[5] S.H. Khan, M. Hayat, M. Bennamoun, F.A. Sohel, R. Togneri, Cost-sensitive learning of deep feature representations from imbalanced data, *IEEE Trans. Neural Networks Learn. Syst.* 29 (8) (2017) 3573–3587.

[6] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, S. Belongie, Class-balanced loss based on effective number of samples, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 9268–9277.

[7] K. Cao, C. Wei, A. Gaidon, N. Aréchiga, T. Ma, Learning imbalanced datasets with label-distribution-aware margin loss, *Adv. Neural Inform. Process. Syst.* 32 (2019) 1567–1578.

[8] J. Tan, C. Wang, B. Li, Q. Li, W. Ouyang, C. Yin, J. Yan, Equalization loss for long-tailed object recognition, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11662–11671.

[9] B. Kang, S. Xie, M. Rohrbach, Z. Yan, A. Gordo, J. Feng, Y. Kalantidis, Decoupling representation and classifier for long-tailed recognition, in: *International Conference on Learning Representations*, 2019.

[10] B. Zhou, Q. Cui, X.-S. Wei, Z.-M. 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*, 2020, pp. 9719–9728.

[11] Z. Zhong, J. Cui, S. Liu, J. Jia, Improving calibration for long-tailed recognition, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16489–16498.

[12] V. Pappayan, X. Han, D.L. Donoho, Prevalence of neural collapse during the terminal phase of deep learning training, *Proc. Nat. Acad. Sci.* 117 (40) (2020) 24652–24663.

[13] C. Fang, H. He, Q. Long, W.J. Su, Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training, *Proceedings of the National Academy of Sciences* 118 (43).

[14] C. Drummond, Class imbalance and cost sensitivity: Why undersampling beats oversampling, in: *ICML-KDD 2003 Workshop: Learning from Imbalanced Datasets*, Vol. 3, 2003.

[15] H. Han, W.-Y. Wang, B.-H. Mao, Borderline-smote: a new over-sampling method in imbalanced data sets learning, in: *International conference on intelligent computing*, Springer, 2005, pp. 878–887.

[16] J. Ren, C. Yu, X. Ma, H. Zhao, S. Yi, et al., Balanced meta-softmax for long-tailed visual recognition, *Adv. Neural Inform. Process. Syst.* 33 (2020) 4175–4186.

[17] T.-Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, in: Focal loss for dense object detection, in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.

[18] J. Mukhoti, V. Kulharia, A. Sanyal, S. Golodetz, P. Torr, P. Dokania, Calibrating deep neural networks using focal loss, *Advances in Neural Information Processing Systems* 33.

[19] D.G. Mixon, H. Parshall, J. Pi, Neural collapse with unconstrained features, *arXiv preprint arXiv:2011.11619*.

[20] E. Weinan, S. Wojtowytsch, On the emergence of tetrahedral symmetry in the final and penultimate layers of neural network classifiers, *arXiv preprint arXiv:2012.05420*.

[21] J. Lu, S. Steinerberger, Neural collapse with cross-entropy loss, *arXiv preprint arXiv:2012.08465*.

[22] Z. Zhu, T. DING, J. Zhou, X. Li, C. You, J. Sulam, Q. Qu, A geometric analysis of neural collapse with unconstrained features, in: *Advances in Neural Information Processing Systems*, 2021.

[23] X. Han, V. Pappayan, D.L. Donoho, in: *ICLR (Ed.), Neural collapse under mse loss: Proximity to and dynamics on the central path*, 2022.

[24] W. Ji, Y. Lu, Y. Zhang, Z. Deng, W.J. Su, An unconstrained layer-peeled perspective on neural collapse, *ICLR* (2022).

[25] Y. Yang, L. Xie, S. Chen, X. Li, Z. Lin, D. Tao, Do we really need a learnable classifier at the end of deep neural network?, *arXiv preprint arXiv:2203.09081*.

[26] R.A. Fisher, The use of multiple measurements in taxonomic problems, *Annals Eugenics* 7 (2) (1936) 179–188.

[27] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[28] H. Zhang, M. Cisse, Y.N. Dauphin, D. Lopez-Paz, mixup: Beyond empirical risk minimization, in: *International Conference on Learning Representations*, 2018.

[29] H.-P. Chou, S.-C. Chang, J.-Y. Pan, W. Wei, D.-C. Juan, Remix: Rebalanced mixup, in: *European Conference on Computer Vision*, Springer, 2020, pp. 95–110.

[30] X. Zhang, Z. Fang, Y. Wen, Z. Li, Y. Qiao, Range loss for deep face recognition with long-tailed training data, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5409–5418.

[31] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, S.X. Yu, Large-scale long-tailed recognition in an open world, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2537–2546.

[32] L. Xiang, G. Ding, J. Han, Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification, *European Conference on Computer Vision*, Springer (2020) 247–263.

- [33] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer, 2015, pp. 234–241.
- [34] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, H. Adam, Encoder-decoder with atrous separable convolution for semantic image segmentation, in: Proceedings of the European conference on computer vision (ECCV), 2018, pp. 801–818.
- [35] A. Krizhevsky, G. Hinton, et al., Learning multiple layers of features from tiny images.
- [36] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in, IEEE conference on computer vision and pattern recognition, IEEE 2009 (2009) 248–255.
- [37] G. Van Horn, O. Mac Aodha, Y. Song, Y. Cui, C. Sun, A. Shepard, H. Adam, P. Perona, S. Belongie, The inaturalist species classification and detection dataset, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 8769–8778.
- [38] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, A. Torralba, Places: A 10 million image database for scene recognition, IEEE Trans. Pattern Anal. Mach. Intell. 40 (6) (2017) 1452–1464.
- [39] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, B. Schiele, The cityscapes dataset for semantic urban scene understanding, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 3213–3223.
- [40] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimshe, L. Antiga, et al., Pytorch: An imperative style, high-performance deep learning library, Adv. Neural Inform. Process. Syst. 32 (2019) 8026–8037.
- [41] I. Loshchilov, F. Hutter, Sgdr: Stochastic gradient descent with warm restarts, Learning 10 3.




Yibo Yang is a research scientist in JD Explore Academy. He obtained his Ph.D. degree from Peking University in 2021. His research interests lie in machine learning and computer vision. He has published more than 20 publications in top-tier journals and conferences. He was awarded the 2021 Wuwenjun-Outstanding Ph.D. Dissertation Award in CAAI. He was also awarded the first prize in BMTT 2021, the Innovative Award in MS COCO 2019, the Runner-up Award in MS COCO 2018.



Deng Cai is a Professor in the State Key Lab of CAD&CG, College of Computer Science at Zhejiang University, China. He received the Ph.D. degree in computer science from University of Illinois at Urbana Champaign in 2009. His research interests include machine learning, data mining and information retrieval.



Liang Xie is currently pursuing his PhD degree in the College of Computer Science and Technology at Zhejiang University and his research interests include computer vision and machine learning.



Xiaofei He is currently a full professor in the College of Computer Science at Zhejiang University, China. He received the BS degree from Zhejiang University, China, in 2000 and the PhD degree from University of Chicago in 2005, both in computer science. He joined Zhejiang University in 2007.