

# EXPLORING WEIGHT BALANCING ON LONG-TAILED RECOGNITION PROBLEM

Naoya Hasegawa & Issei Sato

The University of Tokyo

{hasegawa-naoya410, sato}@g.ecc.u-tokyo.ac.jp

## ABSTRACT

Recognition problems in long-tailed data, in which the sample size per class is heavily skewed, have gained importance because the distribution of the sample size per class in a dataset is generally exponential unless the sample size is intentionally adjusted. Various methods have been devised to address these problems. Recently, weight balancing, which combines well-known classical regularization techniques with two-stage training, has been proposed. Despite its simplicity, it is known for its high performance compared with existing methods devised in various ways. However, there is a lack of understanding as to why this method is effective for long-tailed data. In this study, we analyze weight balancing by focusing on neural collapse and the cone effect at each training stage and found that it can be decomposed into an increase in Fisher’s discriminant ratio of the feature extractor caused by weight decay and cross entropy loss and implicit logit adjustment caused by weight decay and class-balanced loss. Our analysis enables the training method to be further simplified by reducing the number of training stages to one while increasing accuracy. Code is available at <https://github.com/HN410/Exploring-Weight-Balancing-on-Long-Tailed-Recognition-Problem>.

## 1 INTRODUCTION

Datasets with an equal number of samples per class, such as MNIST (Lecun et al., 1998) and CIFAR100 (Krizhevsky, 2009), are often used, when we evaluate classification models and training methods in machine learning. However, it is empirically known that the size distribution in the real world often shows a type of exponential distribution called Pareto distribution (Reed, 2001), and the same is true for the number of per-class samples in classification problems (Li et al., 2017; Spain & Perona, 2007). Such distributions are called long-tailed data due to the shape of the distribution since some classes (head classes) are often sampled and many others (tail classes) are not sampled very often. Long-tailed recognition (LTR) is used to attempt to improve the accuracy of classification models on uniform distribution when training data shows such a distribution. There is a problem in LTR that the head classes have large sample size; thus, the output is biased toward them. This reduces the overall and tail class accuracy because tail classes make up the majority (Zhang et al., 2021).

Various methods have been developed for LTR, such as class-balanced loss (CB) (Cui et al., 2019), augmenting samples of tail classes (Wang et al., 2021), two-stage learning (Kang et al., 2020), and enhancing feature extractors (Liu et al., 2023; Yang et al., 2022); see Appendix A.2 for more related research. Alshammari et al. (2022) proposed a simple method, called *weight balancing* (WB), that empirically outperforms previous complex state-of-the-art methods. WB simply combines two classic techniques, weight decay (WD) (Hanson & Pratt, 1989) and MaxNorm (Hinton et al., 2012), with two-stage learning. WD and MaxNorm is known to prevent overlearning (Hinton, 1989); however, it is not known why WB significantly improves the LTR performance.

**Contribution** In this work, we analyze the effectiveness of WB in LTR focusing on neural collapse (NC) (Papayan et al., 2020) and the cone effect (Liang et al., 2022). We first decompose WB into five components: WD, MaxNorm, cross entropy (CE), CB, and two-stage learning. We then show that each component has the following useful properties.

- 1st stage: WD and CE increase the Fisher’s discriminant ratio (FDR) (Fisher, 1936) of features.
  - Degrade the inter-class cosine similarities (Theorem 1 in Sec. 4.2).
  - Decrease the scaling parameters of batch normalization (BN) (Sec. 4.3). This has a positive effect on feature training.
  - Facilitate improvement of FDR as features pass through layers (Sec. 4.4).
- 1st stage: WD increases the norms of features for tail classes (Sec. 4.5).
- 2nd stage: WD and CB perform implicit logit adjustment (LA) by making the norm of classifier’s weights higher for tail classes. MaxNorm facilitates this effect. This stage does not work well for datasets with a small class number (Theorem 2 in Sec. 5.1).

The above analysis is useful in the following points: 1. it provides a guideline for the design of learning in LTR; 2. our theorem offers an insight into how to suppress the cone effect (Liang et al., 2022) that negatively affects deep learning classification models; 3. WB can be further simplified by removing the second stage. Specifically, we only need to learn in LTR by WD, feature regularization (FR), and an equiangular tight frame (ETF) classifier for the linear layer to extract features that are more linearly separable and adjusting the norm of classifier’s weights by LA after the training.

## 2 RELATED WORK

Papayan et al. (2020) investigated the set of phenomena that occurs when the terminal phase of training, which they termed NC. NC can be briefly described as follows. Feature vectors converge to their class means (NC1); the class means of feature vectors converge to a simplex ETF (Strohmer & Heath, 2003) (NC2); the class means of feature vectors and the corresponding weights of the linear classifier converge in the same direction (NC3); and when making predictions, models converge to predict the class, the mean feature vector of which is closest to the feature in Euclidean distance (NC4).

Papayan et al. (2020) claimed that NC leads to increased generalization accuracy and robustness to adversarial samples. The number of studies have been conducted on the conditions under which NC occurs (Han et al., 2022; Ji et al., 2022; Lu & Steinerberger, 2021; Papayan et al., 2020). Rangamani & Banburski-Fahey (2022) theoretically and empirically showed that the occurrence of NC needs WD.

There have also been studied on NC when the model is trained with imbalanced data. Fang et al. (2021) theoretically and experimentally proved “minority collapse” in which features corresponding to classes with small numbers of samples tend to converge in the same direction, even if they belong to different classes. Yang et al. (2022) attempted to solve this problem by fixing the weights of the linear layers to an ETF. Thrampoulidis et al. (2022) tried to generalize NC to imbalanced data by demonstrating that the features converge towards simplex-encoded-labels interpolation (SELI), an extension of ETF in the terminal phase of training. They also demonstrated that minority collapse does not occur under the appropriate regularization including WD.

WD is a regularization method often used for deep neural networks (DNNs) with BN (Ioffe & Szegedy, 2015), which is scale invariant. Zhang et al. (2019) revealed that WD increases the effective learning rate which facilitates regularization in such models. Training of networks containing BN layers with WD is often studied in terms of training dynamics (Li & Arora, 2020; Lobacheva et al., 2021; Wan et al., 2021). Summers & Dinneen (2020) and Kim et al. (2022) studied whether to apply WD to the scaling and shifting parameters of the BN for such networks. WD is also known to have many other positive effects on DNNs, e.g., smoothing loss landscape (Li et al., 2018; Lyu et al., 2022), causing NC (Rangamani & Banburski-Fahey, 2022), and making filters sparser (Mehta et al., 2019) and low-ranked (Galanti et al., 2023). WD is also effective for imbalanced data (Alshammari et al., 2022; Thrampoulidis et al., 2022).

Liang et al. (2022) found a phenomenon they termed “cone effect” observed in neural networks. The phenomenon is a tendency in which features from DNNs with activation functions are prone to have high cosine similarities. It is widely observed regardless of the modality of data, the structure of models, and whether the models are trained. It also indicates that features from different classes tend to exhibit high cosine similarity. Thus, it has a negative impact on DNNs for classification since features with lower inter-class cosine similarity are more linearly separated. We also reconfirmed such phenomena in Sec. 4.2. Our analysis provides guidelines to prevent the effect.

### 3 PRELIMINARIES

This section defines the notations used in the paper and presents the strategies for LTR. See Appendix A.1 for the table of these notations. Suppose a multiclass classification with  $C$  classes by samples  $\mathcal{X} \subset \mathbb{R}^p$  and labels  $\mathcal{Y} \equiv \{1, 2, \dots, C\}$ . The training dataset  $\mathcal{D} \equiv \{(\mathbf{x}_i, y_i) | \mathbf{x}_i \in \mathcal{X}, y_i \in \mathcal{Y}\}_{i=1}^N$  consists of  $\mathcal{D}_k \equiv \{(\mathbf{x}_i, y_i) \in \mathcal{D} | y_i = k\}$ . We define  $N$  as  $\|\mathcal{D}\|$  and  $N_k$  as  $\|\mathcal{D}_k\|$ , the number of samples. Without loss of generality, assume the classes are sorted in descending order by the number of samples. In other words,  $\forall k \in \{1, 2, \dots, C-1\}$ ,  $N_k \geq N_{k+1}$  holds. Imbalance factor  $\rho = \frac{N_1}{N_C} = \frac{\max_k N_k}{\min_k N_k}$  indicates the extent to which the training dataset is imbalanced and  $\rho \gg 1$  holds in LTR. Therefore, the number of samples in each class  $N_k$  satisfies  $N_k = N_1 \rho^{-\frac{k-1}{C-1}}$ . Define  $\bar{N}$  as  $C \left( \sum_{k=1}^C \left( \frac{1}{N_k} \right) \right)^{-1}$ , the harmonic mean of the number of samples per class. In LTR, the test dataset  $\mathcal{D}'$  used for accuracy evaluation is class-balanced, i.e., each class has the same number of samples.

Consider a network  $\mathbf{f}(\cdot; \Theta) : \mathcal{X} \rightarrow \mathbb{R}^C$  parameterized by  $\Theta = \{\theta_l\}$ . It outputs a logit  $\mathbf{z}_i = \mathbf{f}(\mathbf{x}_i; \Theta)$ . The network is further divided into a feature extractor  $\mathbf{g}(\cdot; \Theta_g) : \mathcal{X} \rightarrow \mathbb{R}^d$  and a classifier  $\mathbf{h}(\cdot; \Theta_h) : \mathbb{R}^d \rightarrow \mathbb{R}^C$  with  $d$  as the number of dimensions for features. This means  $\mathbf{f}(\mathbf{x}_i; \Theta) = \mathbf{h}(\mathbf{g}(\mathbf{x}_i; \Theta_g); \Theta_h)$ . We often abbreviate  $\mathbf{g}(\mathbf{x}_i; \Theta_g)$  to  $\mathbf{g}(\mathbf{x}_i)$ . Let  $\boldsymbol{\mu}_k \equiv \frac{1}{N_k} \sum_{(\mathbf{x}_i, k) \in \mathcal{D}_k} \mathbf{g}(\mathbf{x}_i)$  be the inner-class mean of the features for class  $k$  and  $\boldsymbol{\mu} \equiv \frac{1}{C} \sum_{k=1}^C \boldsymbol{\mu}_k$  be the mean of the inner-class mean of the features. We used a linear layer for the classifier. In other words,  $\mathbf{h}(\mathbf{v}; \Theta_h) = \mathbf{W}^\top \mathbf{v}$  holds for a feature  $\mathbf{v} \in \mathbb{R}^d$  with  $\mathbf{W} = \mathbb{R}^{d \times C}$  as a weight matrix. Denote the  $k$ th column vector of  $\mathbf{W}$  by  $\mathbf{w}_k$ ; thus,  $\Theta_h = \{\mathbf{w}_k\}$ . The loss function is denoted by  $\ell(\mathbf{z}_i, y_i) : \mathbb{R}^C \times \mathbb{R}^C \rightarrow \mathbb{R}$ . Let  $\ell_{\text{CE}}$  and  $\ell_{\text{CB}}$  be the loss function of CE and CB, respectively. By rewriting  $F(\Theta; \mathcal{D}) \equiv \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{f}(\mathbf{x}_i; \Theta), y_i)$ , parameters  $\Theta$  are optimized as follows in the absence of regularization:

$$\Theta^* = \arg \min_{\Theta} F(\Theta; \mathcal{D}). \quad (1)$$

We evaluated the method by accuracy on test dataset and FDR. FDR is the ratio of the inter-class variance to the inner-class variance and indicates the ease of linear separation of the features. For example, the FDR of training features is  $\text{Tr}(S_W^{-1} S_B)$  where  $S_B = \sum_{k=1}^C N_k (\boldsymbol{\mu}_k - \boldsymbol{\mu})(\boldsymbol{\mu}_k - \boldsymbol{\mu})^\top$  and  $S_W = \sum_{k=1}^C \sum_{\mathbf{x}_i \in \mathcal{D}_k} (\mathbf{x}_i - \boldsymbol{\mu}_k)(\mathbf{x}_i - \boldsymbol{\mu}_k)^\top$ . The FDR of test features is similar.

**Regularization Methods** We implemented WD as L2 regularization with a hyperparameter  $\lambda$  because we used stochastic gradient descent (SGD) for the optimizer. Thus, optimization is written as  $\Theta^* = \arg \min_{\Theta} F(\Theta; \mathcal{D}) + \frac{\lambda}{2} \sum_{k \in \mathcal{Y}} \|\mathbf{w}_k\|_2^2$ . MaxNorm is a regularization that places a restriction on the upper bound of the norm of weights. As with Alshammari et al. (2022), for the convenience of hyperparameter tuning, the weights to be constrained are only the those belonging to the classifier  $\Theta_h \subset \Theta$ . Thus, we compute  $\Theta^* = \arg \min_{\Theta} F(\Theta; \mathcal{D})$ , s.t.  $\forall \mathbf{w}_k \in \Theta_h$ ,  $\|\mathbf{w}_k\|_2^2 \leq \eta_k^2$ , where  $\eta_k$  is a hyperparameter. Since constrained optimization is difficult to solve in neural networks, we implemented it using projected gradient descent. In our implementation, this projects the weights that do not satisfy the constraint to the range where the constraint is satisfied by updating  $\mathbf{w}_k$  to  $\min\left(1, \frac{\eta_k}{\|\mathbf{w}_k\|_2}\right) \mathbf{w}_k$ .

**Weight Balancing** WB adopts two-stage training (Kang et al., 2020). In the first stage, the parameters of the entire model  $\Theta$  are optimized with CE using WD. In the second stage,  $\Theta_g$  is fixed and only  $\Theta_h$  is optimized with CB using WD and MaxNorm. For simplicity, let  $\beta$  of CB be 1 in this paper. Thus,  $\ell_{\text{CB}}(\mathbf{z}_i, y_i)$  is  $-\frac{\bar{N}}{N_{y_i}} \log\left(\frac{\exp((\mathbf{z}_i)_{y_i})}{\sum_{j=1}^C \exp((\mathbf{z}_i)_j)}\right)$ .<sup>1</sup> Therefore, in the second stage of WB,  $F_{\text{WB}}(\mathbf{W}; \mathcal{D}) \equiv \frac{1}{N} \sum_{i=1}^N \ell_{\text{CB}}(\mathbf{z}_i, y_i) + \frac{\lambda}{2} \sum_{k \in \mathcal{Y}} \|\mathbf{w}_k\|_2^2$  is optimized for  $\mathbf{W}$ . Note that we do not consider MaxNorm in this paper, as we reveal that it only changes the initial values and do not impose any intrinsic constraints on the optimization presented in Appendix. A.5.

**Logit Adjustment** Multiplicative LA (Kim & Kim, 2020) changes the norm of the per-class weights of the linear layer. The weights of the linear layer corresponding to class  $k$ ,  $\mathbf{w}'_k$  are adjusted

<sup>1</sup>In the paper of Cui et al. (2019),  $\bar{N}$  is set to be 1, but in their implementation,  $\bar{N}$  is equal to the harmonic mean. We adopt the latter since our experiments were based on their implementation.

Table 1: FDRs for each dataset of models trained with each method. A higher FDR indicates that features are more easily linearly separable. For all datasets, the method with WD or CE produces a higher FDR and using both results in the highest FDR.

Method	CIFAR10-LT		CIFAR100-LT		mini-ImageNet-LT	
	Train	Test	Train	Test	Train	Test
CE w/o WD	$8.17 \times 10^1$	$2.17 \times 10^1$	$1.28 \times 10^2$	$4.16 \times 10^1$	$7.56 \times 10^1$	$4.28 \times 10^1$
CB w/o WD	$4.43 \times 10^1$	$1.50 \times 10^1$	$8.17 \times 10^1$	$2.42 \times 10^1$	$4.68 \times 10^1$	$2.93 \times 10^1$
CE w/ WD	<b><math>2.60 \times 10^3</math></b>	<b><math>3.89 \times 10^1</math></b>	<b><math>2.87 \times 10^4</math></b>	<b><math>1.07 \times 10^2</math></b>	<b><math>6.58 \times 10^2</math></b>	<b><math>1.01 \times 10^2</math></b>
CB w/ WD	$3.39 \times 10^2$	$3.04 \times 10^1$	$2.12 \times 10^4$	$6.74 \times 10^1$	$4.84 \times 10^2$	$6.84 \times 10^1$
WD w/o BN	$2.19 \times 10^2$	$3.42 \times 10^1$	$5.77 \times 10^2$	$7.95 \times 10^1$	$1.61 \times 10^2$	$6.73 \times 10^1$
WD fixed BN	<b><math>1.63 \times 10^3</math></b>	<b><math>4.16 \times 10^1</math></b>	<b><math>2.04 \times 10^4</math></b>	<b><math>1.05 \times 10^2</math></b>	<b><math>3.94 \times 10^2</math></b>	<b><math>1.04 \times 10^2</math></b>

to  $\frac{1}{\mathbb{P}(Y=k)^\gamma \|\mathbf{w}_k\|_2}$ , where  $\gamma > 0$  is a hyperparameter. Kim & Kim (2020) trained using projected gradient descent so that the weights of the linear layer always satisfy  $\forall k, \|\mathbf{w}_k\|_2 = 1$ , but in our study, for comparison, we trained with regular gradient descent, normalized the norm post-hoc, and then adjusted the norm. Additive LA (Menon et al., 2020) is a method to adjust the output to minimize the average per-class error by adding a different constant for each class to the logit at prediction. That is, the logit  $\mathbf{z}_i$  at prediction is adjusted to  $\mathbf{z}'_i$  as  $(\mathbf{z}'_i)_k = (\mathbf{z}_i)_k - \tau \log \mathbb{P}(Y = k)$ , where  $\tau$  is a hyperparameter. LA generally increases the probability of classifying inputs into tail classes.

**ETF Classifier** NC indicates that the matrix of classifier weights in deep learning models converges to an ETF when trained with CE (Papayan et al., 2020). The idea of ETF classifiers is to train only the feature extractor by fixing the linear layer to an ETF from the initial step. ETF classifiers’ weights  $\mathbf{W} \in \mathbb{R}^{d \times C}$  satisfy  $\mathbf{W} = \sqrt{E_W \frac{C}{C-1}} \mathbf{U} (\mathbf{I}_C - \frac{1}{C} \mathbf{1}_C \mathbf{1}_C^\top)$ , where  $\mathbf{U} \in \mathbb{R}^{d \times C}$  is a matrix such that  $\mathbf{U}^\top \mathbf{U}$  is an identity matrix,  $\mathbf{I}_C \in \mathbb{R}^{C \times C}$  is an identity matrix, and  $\mathbf{1}_C$  is a  $C$ -dimensional vector, all elements in which are 1. Also,  $E_W$  is a hyperparameter, which was set to 1 in our experiments.

## 4 ROLE OF FIRST TRAINING STAGE IN WB

First, we analyze the effect of WD and CE in the first training stage. In Sec. 4.1, we present the datasets and models used in the analysis. Then, we examine the cosine similarity and FDR of the features from the models trained with each training method for the first stage of training in Secs. 4.2, 4.3, and 4.4. In Sec. 4.5, we identify the properties of training features trained with WD, which is the key to the success of WB. The results of the experiments that is not included in the main text are presented in Appendix A.5.

### 4.1 SETTINGS

We used CIFAR10, CIFAR100 (Krizhevsky, 2009), mini-ImageNet (Vinyals et al., 2016), and ImageNet (Deng et al., 2009) as the datasets and followed Cui et al. (2019), Vigneswaran et al. (2021), and Liu et al. (2019) to create long-tailed datasets, CIFAR10-LT, CIFAR100-LT, mini-ImageNet-LT, and ImageNet-LT. We also experimented with a tabular dataset, Helena (Guyon et al., 2019), to verify the effectiveness of our method for other than image dataset. Each class is classified into one of three groups, *Many*, *Medium*, and *Few*, depending on the number of training samples  $N_k$ .

We used ResNeXt50 (Xie et al., 2017) for ImageNet-LT and ResNet34 (He et al., 2016) for the other datasets. In WB, parameters were trained in two stages (Kang et al., 2020); in the first stage, we trained the weights of the entire model  $\Theta$ , while in the second stage, we trained the weights of the classifiers  $\Theta_h$  with the weights of the feature extractors  $\Theta_g$  fixed.

To investigate the properties of neural network models, we also experimented with models using a multilayer perceptron (MLP) and a residual block (ResBlock) (He et al., 2016). The number of blocks is added after the model name to distinguish them (e.g., MLP3). We trained and evaluated these models using MNIST (Lecun et al., 1998).

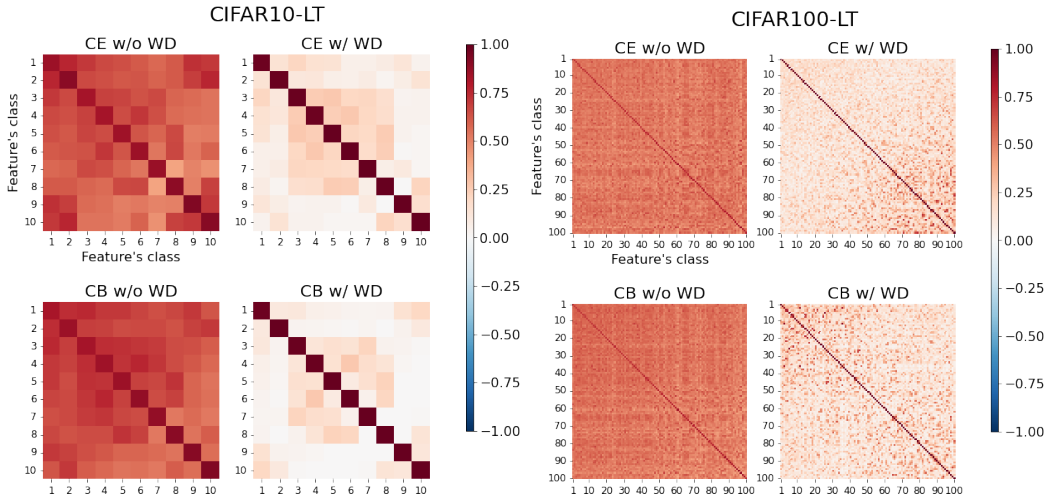


Figure 1: Heatmaps showing average cosine similarities of training features between two classes. WD maintains high cosine similarity between the same classes and reduces the cosine similarity between the different classes.

For training losses, we used CE unless otherwise stated. See Appendix A.4 for detailed settings.

#### 4.2 WD AND CE DEGRADE INTER-CLASS COSINE SIMILARITIES

We first examined the role of WD in the first stage of training. We trained the model in two ways: without WD (w/o WD) and with WD (w/ WD). In addition, we compared two types of loss: one using CE and the other using CB, for a total of four training methods. We investigated the FDR, the inter-class and inner-class mean cosine similarity of the features, and the norm of the mean features per class. Due to space limitations, some of results of the mini-ImageNet-LT excluding FDR are presented in Appendix A.5.

Table 1 lists the FDRs of the models trained with each method, and Figure 1 shows the cosine similarities of the training features per class. It can be seen that the combination of WD and CE achieves the highest FDR. This fact is consistent with the conclusion that WD is necessary for the NC in balanced datasets, as shown by Rangamani & Banburski-Fahey (2022) and suggests that this may also hold for long-tailed data.

When WD is not applied, the cosine similarity is generally higher even between features of different classes. This is natural because neural networks have the cone effect found in Liang et al. (2022). The cone effect is a phenomenon in which features from a DNN tend to have high cosine similarities to each other even if they belong to different classes. However, the methods with WD result in lower cosine similarities between features of different classes and higher cosine similarities between features of the same classes, leading to a higher FDR. This shows that WD prevents the cone effect, as supported by the following theorem.

**Theorem 1.** For all  $(\mathbf{x}_i, y_i), (\mathbf{x}_j, y_j) \in \mathcal{D}$  s.t.  $y_i \neq y_j$ , if  $\mathbf{W}$  is an ETF and there exists  $\epsilon$  and  $L$  s.t.  $\left\| \frac{\partial \ell_{\text{CE}}}{\partial \mathbf{g}(\mathbf{x}_i)} \right\|_2, \left\| \frac{\partial \ell_{\text{CE}}}{\partial \mathbf{g}(\mathbf{x}_j)} \right\|_2 \leq \epsilon < \frac{1}{C}$  and  $\|\mathbf{g}(\mathbf{x}_i)\|_2, \|\mathbf{g}(\mathbf{x}_j)\|_2 \leq L \leq 2\sqrt{2} \log(C-1)$ , the following holds:

$$\cos(\mathbf{g}(\mathbf{x}_i), \mathbf{g}(\mathbf{x}_j)) \leq 2\delta\sqrt{1-\delta^2}, \quad (2)$$

where  $\delta \equiv \frac{1}{L} \frac{C-1}{C} \log\left(\frac{(C-1)(1-\epsilon)}{\epsilon}\right) \in \left(\frac{1}{\sqrt{2}}, 1\right]$  and  $\cos(\cdot, \cdot)$  means cosine similarity of the two vectors.

This theorem states that the cone effect is suppressed when the following two conditions hold: 1. the weight matrix of the linear layer is an ETF; 2. the norms of the features are sufficiently small. Sufficient training causes NC and makes the weight matrix of the linear layer converge to an ETF

Table 2: Means and standard deviations of all scaling parameters in BN layers of models trained with each method for each dataset. WD greatly reduces the average scaling parameter.

Method	CIFAR10-LT	CIFAR100-LT	mini-ImageNet-LT
CE w/o WD	$1.00 \pm 0.0218$	$1.00 \pm 0.0289$	$0.998 \pm 0.0498$
CE w/ WD	$0.0395 \pm 0.0273$	$0.0649 \pm 0.0548$	$0.0863 \pm 0.0707$

Table 3: Means and standard deviations of all shifting parameters in BN layers of models trained with each method for each dataset.

Method	CIFAR10-LT	CIFAR100-LT	mini-ImageNet-LT
CE w/o WD	$-0.0236 \pm 0.0255$	$-0.0444 \pm 0.0348$	$-0.0904 \pm 0.0698$
CE w/ WD	$-0.00310 \pm 0.0244$	$-0.00667 \pm 0.0309$	$-0.0186 \pm 0.0448$

(Papayan et al., 2020). ETF Classifiers are also effective: they are methods to fix the weight matrix of the linear layer to an ETF from the beginning to satisfy this condition. As for the second condition, WD contributes to satisfy the condition. In fact, as shown in Figure 2, WD also indirectly degrades the norm of the features by regularizing the weights causing low inter-class cosine similarity. This also suggests that explicit FR leads to better feature extractor training. We consider that in Sec. 5.2.

#### 4.3 WD AND CE DECREASE SCALING PARAMETERS OF BN

Next, we split the effect of WD on the model into that on the convolution layers and that on the BN layers. To examine the impact on convolution, we analyzed the FDR of the features from the models trained with WD applied only to the convolution layers (WD w/o BN). The bottom halves of Table 1 shows that the FDRs of the methods with restricted WD application. These results indicate that when WD is applied only to the convolution layers, the FDR improves sufficiently, which means that enabling WD for the convolution layers is essential for improving accuracy. One reason for this may be an increase in the effective learning rate (Zhang et al., 2019).

**Small scaling parameters of BN facilitates feature learning** We also examined the effect on BN layers and found that WD reduces the mean of the BN’s scaling parameters; see Tables 2 and 3. Note that the standard deviation of the shifting parameters remains almost unchanged. We investigated how the FDR changes when models are trained with the scaling parameters of BN fixed to one common small value (WD fixed BN). In this method, we applied WD only to the convolution layers and fixed the shifting parameters of BN to 0. We selected the optimal value for the scaling parameters from  $\{0.01, 0.02, \dots, 0.20\}$  using the validation dataset. Table 1 shows that a small value of the scaling parameters of BN is crucial for boosting FDR; even setting the same value for all scaling parameters in the entire model also works. In this case, the scaling parameters of BN does not directly affect FDR, as it only multiplies the features uniformly by a constant. This suggests that the improvement in FDR caused by applying WD to BN layers is mainly because smaller scaling parameters have a positive effect on the learning dynamics and improve FDR. Although Kim et al. (2022) attribute this to the increase in the effective learning rate of the scaling parameters, they have a positive effect on FDR even when the scaling parameters are not trained, suggesting that there are other significant effects. For example, smaller scaling parameters reduce the norm of the feature, promoting the effects discussed in Sec. 4.2. We also examined the effect of high relative variance in shifting parameters of BN in Appendix A.5.

#### 4.4 WD AND CE FACILITATE IMPROVEMENT OF FDR AS FEATURES PASS THROUGH LAYERS

Table 4 compares the FDR of two sets of features; ones from each model trained with each method (before) and ones output from the model, passed through a randomly initialized linear layer, then ReLU-applied (after). This experiment revealed that the features obtained from the models trained with WD have a bias to improve their FDR regardless of the class imbalance. The detailed experimental set-up is described in Appendix A.5. While training without WD does not increase the FDR in

Table 4: FDRs of features from each model trained with each method and features passed through randomly initialized linear layer and ReLU after model. C10, C100, and mIm are the abbreviations for CIFAR10, CIFAR100, and mini-ImageNet respectively. Red text indicates higher values than before, blue text indicates lower values. The FDR of features obtained from models trained with WD improves by passing the features through a random initialized layer and ReLU.

Model	Dataset	CE		WD w/o BN		WD	
		before	after	before	after	before	after
ResNet34	C100	$7.51 \times 10^1$	$7.11 \times 10^1$	$2.29 \times 10^2$	$2.34 \times 10^2$	$3.98 \times 10^2$	$4.10 \times 10^2$
	C100-LT	$4.16 \times 10^1$	$4.02 \times 10^1$	$7.95 \times 10^1$	$7.89 \times 10^1$	$1.07 \times 10^2$	$1.13 \times 10^2$
	C10	$4.77 \times 10^1$	$5.56 \times 10^1$	$7.96 \times 10^1$	$1.02 \times 10^2$	$1.81 \times 10^2$	$2.35 \times 10^2$
	C10-LT	$2.17 \times 10^1$	$2.36 \times 10^1$	$3.42 \times 10^1$	$3.94 \times 10^1$	$3.89 \times 10^1$	$4.30 \times 10^1$
	mIm	$7.34 \times 10^1$	$6.66 \times 10^1$	$1.81 \times 10^2$	$1.86 \times 10^2$	$3.63 \times 10^2$	$3.93 \times 10^2$
	mIm-LT	$4.28 \times 10^1$	$3.97 \times 10^1$	$6.73 \times 10^1$	$6.30 \times 10^1$	$1.01 \times 10^2$	$1.08 \times 10^2$
MLP3		$1.58 \times 10^2$	$1.54 \times 10^2$	$2.48 \times 10^2$	$3.53 \times 10^2$	$2.36 \times 10^2$	$7.96 \times 10^2$
MLP4		$1.98 \times 10^2$	$1.96 \times 10^2$	$3.60 \times 10^2$	$4.97 \times 10^2$	$4.12 \times 10^2$	$1.43 \times 10^3$
MLP5	MNIST	$2.44 \times 10^2$	$2.37 \times 10^2$	$4.91 \times 10^2$	$6.76 \times 10^2$	$6.10 \times 10^2$	$2.47 \times 10^3$
ResBlock1		$1.39 \times 10^2$	$1.36 \times 10^2$	$2.20 \times 10^2$	$2.69 \times 10^2$	$2.44 \times 10^2$	$5.81 \times 10^2$
ResBlock2		$1.66 \times 10^2$	$1.59 \times 10^2$	$2.97 \times 10^2$	$3.65 \times 10^2$	$4.44 \times 10^2$	$9.55 \times 10^2$




Figure 2: Norm of mean per-class training features produced from models trained with each method. Features learned with methods with WD all demonstrate that the norms of the *Many* classes’ features tend to be smaller than those of the *Few* classes.

most cases, the features from models with WD-applied convolution improve the FDR over the before FDR in most cases. Moreover, the after-FDR increases significantly when WD is applied to the whole model. Note that this phenomenon is valid when the features go through only one randomized layer. Experiments have also shown that passing features through multiple layers of randomly initialized linear layers and ReLUs does not necessarily improve FDR; see Appendix A.5. These results mean that only increasing randomized layers does not enhance FDR. However, features trained with WD are less likely to be degraded even when going through one randomized layer. This suggests that WD and CE facilitate the training of models.

#### 4.5 WD INCREASES NORMS OF FEATURES FOR TAIL CLASSES

We also investigated the properties appearing in the norm of the training features from each model investigated in Sec. 4.2; Sec. 5.1 reveals that this property is crucial to the effectiveness of WB. Figure 2 shows the norm of the per-class mean training features. Note that this phenomenon does not occur with test features. The method without WD shows almost no relationship between the number of samples and the norm of the features for each class. However, when WD is applied, the norms of the *Many*-class features drop significantly, to between one-half and one-fifth that of *Few* classes. While this phenomenon has been observed in step-imbalanced data (Thrampoulidis et al., 2022), we found that it also occurs in long-tailed data.

## 5 ROLE OF SECOND TRAINING STAGE IN WB

In this section, we theoretically analyze how the second stage of training operates in Sec. 5.1, which leads to further simplification and empirical performance improvement of WB in Sec. 5.2.

### 5.1 WD AND CB PERFORM IMPLICIT LOGIT ADJUSTMENT

We found that WB encourages the model to output features with a high FDR in the first stage of training. How does the second stage of training improve accuracy? Alshammari et al. (2022) found that WB trains models so that the norm of the linear layer becomes more significant for the tail classes in the second stage of training but does not explicitly train in this way. We found the following theorem that shows the second stage of WB is equivalent to multiplicative LA under certain assumptions. In the following theorem, we assume  $\mu = \mathbf{0}$ , which is valid when NC occurs, i.e., when per-class mean features follow ETF or SELI (Papayan et al., 2020; Thrampoulidis et al., 2022). See Appendix. A.3 for the theorem and proof when this does not hold.

**Theorem 2.** *Assume  $\mu = \mathbf{0}$ . For any  $k \in \mathcal{Y}$ , if there exists  $\mathbf{w}_k^*$  s.t.  $\left. \frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \right|_{\mathbf{w}_k = \mathbf{w}_k^*} = \mathbf{0}$  and  $\|\mathbf{w}_k^*\|_2 = O\left(\frac{1}{\lambda \rho C}\right)$ , we have*

$$\forall k \in \mathcal{Y}, \left\| \mathbf{w}_k^* - \frac{\bar{N}}{\lambda N} \boldsymbol{\mu}_k \right\|_2 = O\left(\frac{1}{\lambda^2 \rho^2 C^2}\right), \quad (3)$$

where  $O$  is a Bachmann-Landau O-notation. For example, the notation in the form  $O\left(\frac{1}{x}\right)$  indicates that the term is at most a constant multiple of  $\frac{1}{x}$  if  $x$  is sufficiently large. Note that  $\frac{\bar{N}}{\lambda N}$  is independent of  $k$ . This theorem states that if the number of classes or the imbalance factor is sufficiently large, and if NC has occurred in the first stage, there exists linear layer weights that are constant multiples of the corresponding features and sufficiently close to the stationary point in the optimization. In other words, if the previous conditions are satisfied, the norm of the linear layer is proportional to the norm of the corresponding per-class mean feature. We show that the norm of training features are larger for *Few* classes in Sec. 4.5. Thus, the linear layer is trained to be larger for the *Few* classes. Besides, the features and the corresponding linear layers' weights are aligned when NC occurs, so the orientation of the linear layers' weights hardly changes during this training. This is the same operation as when multiplicative LA makes adjustments to long-tailed data, increasing the norm of the linear layer for the tail class (Kim & Kim, 2020). Indeed, if there exists some  $c_0$  and  $\gamma_0$ , and for all  $k$ ,  $\|\boldsymbol{\mu}_k\|_2 = c_0 \mathbb{P}(Y = k)^{\gamma_0}$  holds, then the second-stage training is equivalent to multiplicative LA. In this theorem, we do not consider MaxNorm because we found that MaxNorm makes the norm of the weights in the linear layer closer to the same value before training, only facilitating convergence. For more details, see Appendix A.5.

Note also that this theorem does not guarantee that the second stage of WB is an operation equivalent to multiplicative LA when the number of classes is small. In fact, even though the norm of the training features is larger for *Few* classes (Figure 2), WB fails to make the norm of the weights greater for *Few* classes in CIFAR10-LT: see Appendix A.5. This suggests that replacing the second stage of training with LA would be more generic and we confirm its validity in the following section.

### 5.2 IS TWO-STAGE LEARNING REALLY NECESSARY?

We experimentally demonstrated the role of each training stage of WB as indicated by Theorems 1 and 2. For this purpose, we develop a method with a combination of essential operations and verify its high accuracy and FDR. In the first stage of training, FR and an ETF classifier are used to satisfy the conditions of Theorem 1. We fix the linear layer and train the feature extractor with CE, WD, and FR; then adjust the norm of the linear layer with multiplicative LA. To show that these are effective operations implicitly carried out with WB, we compare the combination with WB. For the ablation study, we experimented by comparing WB with the following methods.

- Training with CE or CB only (CE, CB).
- Training with WD and CE (WD).



Table 5: FDRs and accuracy of models trained with each method for CIFAR100-LT. In addition to WD, training with FR and an ETF classifier further improves the FDR. However, only using these does not greatly improve accuracy; LA increases the accuracy of *Medium* and *Few* classes and enhances the average accuracy significantly as a result. In all cases, multiplicative LA improves the accuracy more than additive LA.

Method	LA	FDR		Accuracy (%)			
		Train	Test	<i>Many</i>	<i>Medium</i>	<i>Few</i>	Average
CE	N/A	$1.28 \times 10^2_{\pm 0.9 \times 10^1}$	$4.16 \times 10^1_{\pm 1.0 \times 10^0}$	$64.6 \pm 0.9$	$36.9 \pm 0.8$	$11.9 \pm 0.7$	$38.5 \pm 0.6$
CB	N/A	$8.17 \times 10^1_{\pm 8.0 \times 10^0}$	$2.42 \times 10^1_{\pm 0.6 \times 10^0}$	$47.6 \pm 1.0$	$23.2 \pm 0.6$	$5.61 \pm 0.4$	$26.1 \pm 0.4$
WD	N/A	$2.87 \times 10^4_{\pm 6.0 \times 10^3}$	$1.07 \times 10^2_{\pm 0.2 \times 10^1}$	$75.9 \pm 0.5$	$45.3 \pm 0.6$	$13.9 \pm 0.7$	$46.0 \pm 0.4$
	Add	$2.87 \times 10^4_{\pm 6.0 \times 10^3}$	$1.07 \times 10^2_{\pm 0.2 \times 10^1}$	$70.7 \pm 0.9$	$45.7 \pm 0.9$	<b><math>30.9 \pm 0.9</math></b>	$49.6 \pm 0.4$
	Mult	$2.87 \times 10^4_{\pm 6.0 \times 10^3}$	$1.07 \times 10^2_{\pm 0.2 \times 10^1}$	$72.6 \pm 0.7$	$48.5 \pm 0.8$	$29.5 \pm 0.9$	$50.8 \pm 0.4$
WB	N/A	$2.94 \times 10^4_{\pm 6.0 \times 10^3}$	$1.07 \times 10^2_{\pm 0.2 \times 10^1}$	$73.8 \pm 0.7$	$50.2 \pm 0.6$	$25.6 \pm 1.0$	$50.6 \pm 0.2$
	N/A	$3.33 \times 10^4_{\pm 2.0 \times 10^3}$	$1.13 \times 10^2_{\pm 0.1 \times 10^1}$	$76.3 \pm 0.3$	$46.0 \pm 0.4$	$15.5 \pm 0.6$	$46.9 \pm 0.2$
WD&ETF	Add	$3.33 \times 10^4_{\pm 2.0 \times 10^3}$	$1.13 \times 10^2_{\pm 0.1 \times 10^1}$	$73.8 \pm 0.4$	$48.9 \pm 0.3$	$25.8 \pm 0.4$	$50.2 \pm 0.2$
	Mult	$3.33 \times 10^4_{\pm 2.0 \times 10^3}$	$1.13 \times 10^2_{\pm 0.1 \times 10^1}$	$70.4 \pm 0.7$	$51.4 \pm 0.3$	<b><math>31.7 \pm 0.6</math></b>	$51.7 \pm 0.3$
WD&FR &ETF	N/A	<b><math>8.81 \times 10^4_{\pm 2.0 \times 10^3}</math></b>	<b><math>1.22 \times 10^2_{\pm 0.1 \times 10^1}</math></b>	<b><math>77.9 \pm 0.3</math></b>	$46.8 \pm 1.0$	$15.3 \pm 0.3$	$47.6 \pm 0.5$
	Add	<b><math>8.85 \times 10^4_{\pm 2.0 \times 10^3}</math></b>	<b><math>1.22 \times 10^2_{\pm 0.1 \times 10^1}</math></b>	$75.1 \pm 0.3$	$49.3 \pm 1.0$	$26.2 \pm 0.8$	$50.9 \pm 0.3$
	Mult	<b><math>8.85 \times 10^4_{\pm 2.0 \times 10^3}</math></b>	<b><math>1.22 \times 10^2_{\pm 0.1 \times 10^1}</math></b>	$74.2 \pm 0.3$	<b><math>52.9 \pm 0.9</math></b>	$29.9 \pm 0.8$	<b><math>53.0 \pm 0.3</math></b>

- Using an ETF classifier as the linear layer and training only the feature extractor with WD and CE (WD&ETF).
- Using an ETF classifier as the linear layer and training only the feature extractor with WD, FR, and CE (WD&FR&ETF).

As for the bottom three methods, we also compared with additive LA (Menon et al., 2020), with multiplicative LA (Kim & Kim, 2020), and without LA. To ensure that the model correctly classifies both samples of the head classes and tail classes, we consider the average of the per-class accuracy of all classes and each group, i.e., *Many*, *Medium*, and *Few*.

Table 5 lists the FDRs and accuracy of each method for CIFAR100-LT. We show the results on the other datasets including tabular data (CIFAR10-LT, mini-ImageNet-LT, ImageNet-LT, and Helena) and the other model (ResNeXt50) in Appendix A.5. First, training with the ETF classifier in addition to WD increases the FDR for both training and test features. Using FR further improves the FDR. However, it does not boost the average accuracy much. In contrast, LA enhances the classification accuracy of *Medium* and *Few* classes in all cases; thus, it also increase the final average accuracy significantly. It was observed that multiplicative LA improved accuracy to a greater extent than additive LA and that WD&FR&ETF&Multiplicative LA outperformed WB on average accuracy despite our method requiring only one-stage training.

## 6 CONCLUSION

We theoretically and experimentally demonstrated how WB improves accuracy in LTR by investigating each components of WB. In the first stage of training, WD and CE cause following three effects, decrease in inter-class cosine similarity, reduction in the scaling parameters of BN, and improvement of features in the ease of increasing FDR, which enhance the FDR of the features. In the second stage, WD, CB, and the norm of features that increased in the tail classes work as multiplicative LA, by improving the classification accuracy of the tail classes. Our analysis also reveals a training method that achieves higher accuracy than WB by improving each stage based on the objectives. This method is simple, thus we recommend trying it first. A limitation is that our experiments were limited to MLP, ResNet, and ResNeXt and the usefulness in other models such as ViT (Dosovitskiy et al., 2021) is unknown. Future research will include experiments and development of useful methods for such models.

## ACKNOWLEDGEMENTS

We would like to express my sincere gratitude to the members of Issei Sato Laboratory for their invaluable discussions. We are also grateful to the reviewers for their diligent critique, which has significantly improved the quality of this paper. This work was supported by JSPS KAKENHI Grant Number 20H04239 Japan.

## REFERENCES

- Shaden Alshammari, Yu-Xiong Wang, Deva Ramanan, and Shu Kong. Long- Tailed Recognition via Weight Balancing. *c2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6887–6897, March 2022. ISSN 10636919. doi: 10.1109/cvpr52688.2022.00677. URL <https://arxiv.org/abs/2203.14197>. arXiv: 2203.14197 ISBN: 9781665469463.
- Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning Imbalanced Datasets with Label-Distribution-Aware Margin Loss. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/621461af90cadfdaf0e8d4cc25129f91-Abstract.html>.
- Peng Chu, Xiao Bian, Shaopeng Liu, and Haibin Ling. Feature Space Augmentation for Long-Tailed Data. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (eds.), *Computer Vision – ECCV 2020*, Lecture Notes in Computer Science, pp. 694–710, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58526-6. doi: 10.1007/978-3-030-58526-6\_41.
- Yin Cui, Menglin Jia, Tsung Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019-June:9260–9269, January 2019. ISSN 10636919. doi: 10.1109/CVPR.2019.00949. URL <http://arxiv.org/abs/1901.05555>. arXiv: 1901.05555 ISBN: 9781728132938.
- 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*, pp. 248–255. IEEE, June 2009. ISBN 978-1-4244-3992-8. doi: 10.1109/CVPR.2009.5206848.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, January 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Cong Fang, Hangfeng He, Qi Long, and Weijie J. Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences of the United States of America*, 118(43), January 2021. ISSN 10916490. doi: 10.1073/pnas.2103091118. URL <https://arxiv.org/abs/2101.12699>. arXiv: 2101.12699.
- R. A. Fisher. The Use of Multiple Measurements in Taxonomic Problems. *Annals of Eugenics*, 7(2): 179–188, 1936. doi: 10.1111/j.1469-1809.1936.tb02137.x.
- Tomer Galanti, Zachary S. Siegel, Aparna Gupte, and Tomaso Poggio. SGD and Weight Decay Provably Induce a Low-Rank Bias in Neural Networks, January 2023. URL <http://arxiv.org/abs/2206.05794>. arXiv:2206.05794 [cs, stat].
- Aditya Sharad Golatkar, Alessandro Achille, and Stefano Soatto. Time Matters in Regularizing Deep Networks: Weight Decay and Data Augmentation Affect Early Learning Dynamics, Matter Little Near Convergence. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL [https://proceedings.neurips.cc/paper\\_files/paper/2019/hash/87784eca6b0dealdff92478fb786b401-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2019/hash/87784eca6b0dealdff92478fb786b401-Abstract.html).

- Isabelle Guyon, Lisheng Sun-Hosoya, Marc Boullé, Hugo Jair Escalante, Sergio Escalera, Zhengying Liu, Damir Jajetic, Bisakha Ray, Mehreen Saeed, Michèle Sebag, Alexander Statnikov, Wei-Wei Tu, and Evelyne Viegas. Analysis of the AutoML Challenge Series 2015–2018. In Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren (eds.), *Automated Machine Learning: Methods, Systems, Challenges*, The Springer Series on Challenges in Machine Learning, pp. 177–219. Springer International Publishing, Cham, 2019. ISBN 978-3-030-05318-5. doi: 10.1007/978-3-030-05318-5\_10. URL [https://doi.org/10.1007/978-3-030-05318-5\\_10](https://doi.org/10.1007/978-3-030-05318-5_10).
- R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality Reduction by Learning an Invariant Mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pp. 1735–1742, June 2006. doi: 10.1109/CVPR.2006.100. ISSN: 1063-6919.
- X. Y. Han, Vardan Papyan, and David L. Donoho. Neural Collapse Under MSE Loss: Proximity to and Dynamics on the Central Path. In *International Conference on Learning Representations*, January 2022. URL [https://openreview.net/forum?id=w1UbvWH\\_R3](https://openreview.net/forum?id=w1UbvWH_R3).
- Stephen José Hanson and Lorien Pratt. Comparing Biases for Minimal Network Construction with Back-Propagation. *Advances in neural information processing systems 1*, 1:177–185, 1989. URL <http://portal.acm.org/citation.cfm?id=89851.89872>. Publisher: Morgan-Kaufmann ISBN: 1-558-60015-9.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2016-Decem, pp. 770–778, December 2016. ISBN 978-1-4673-8850-4. doi: 10.1109/CVPR.2016.90. URL <http://arxiv.org/abs/1512.03385>. arXiv: 1512.03385 ISSN: 10636919.
- Geoffrey E. Hinton. Connectionist learning procedures. *Artificial Intelligence*, 40(1):185–234, September 1989. ISSN 0004-3702. doi: 10.1016/0004-3702(89)90049-0. URL <https://www.sciencedirect.com/science/article/pii/0004370289900490>.
- Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. July 2012. doi: 10.48550/arxiv.1207.0580. URL <http://arxiv.org/abs/1207.0580>. arXiv: 1207.0580.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pp. 448–456. pmlr, 2015.
- Wenlong Ji, Yiping Lu, Yiliang Zhang, Zhun Deng, and Weijie J. Su. An Unconstrained Layer-Peeled Perspective on Neural Collapse. In *International Conference on Learning Representations*, January 2022. URL <https://openreview.net/forum?id=WZ3yjh8coDg>.
- Arlind Kadra, Marius Lindauer, Frank Hutter, and Josif Grabocka. Well-tuned Simple Nets Excel on Tabular Datasets. In *Advances in Neural Information Processing Systems*, November 2021. URL <https://openreview.net/forum?id=d3k38LTDCyO>.
- 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*, 2020. URL <https://openreview.net/forum?id=rlgRTCvFvB>.
- Bingyi Kang, Yu Li, Sa Xie, Zehuan Yuan, and Jiashi Feng. Exploring Balanced Feature Spaces for Representation Learning. In *International Conference on Learning Representations*, March 2023. URL <https://openreview.net/forum?id=OqtLIabPTit>.
- Bum Jun Kim, Hyeon Choi, Hyeonah Jang, Dong Gu Lee, Wonseok Jeong, and Sang Woo Kim. Guidelines for the Regularization of Gammas in Batch Normalization for Deep Residual Networks, May 2022. URL <http://arxiv.org/abs/2205.07260>. arXiv:2205.07260 [cs].
- Byungju Kim and Junmo Kim. Adjusting decision boundary for class imbalanced learning. *IEEE Access*, 8:81674–81685, December 2020. ISSN 21693536. doi: 10.1109/ACCESS.2020.2991231. URL <https://arxiv.org/abs/1912.01857>. arXiv: 1912.01857.

- Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. *Science Department, University of Toronto, Tech.*, pp. 1–60, 2009. ISSN 1098-6596. doi: 10.1.1.222.9220. arXiv: 1011.1669v3 ISBN: 9788578110796.
- Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, January 1998. ISSN 1558-2256. doi: 10.1109/5.726791. Conference Name: Proceedings of the IEEE.
- Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the Loss Landscape of Neural Nets. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL [https://proceedings.neurips.cc/paper\\_files/paper/2018/hash/a41b3bb3e6b050b6c9067c67f663b915-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2018/hash/a41b3bb3e6b050b6c9067c67f663b915-Abstract.html).
- 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 (CVPR)*, pp. 6918–6928, 2022. URL [https://openaccess.thecvf.com/content/CVPR2022/html/Li\\_Targeted\\_Supervised\\_Contrastive\\_Learning\\_for\\_Long-Tailed\\_Recognition\\_CVPR\\_2022\\_paper.html](https://openaccess.thecvf.com/content/CVPR2022/html/Li_Targeted_Supervised_Contrastive_Learning_for_Long-Tailed_Recognition_CVPR_2022_paper.html).
- Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. WebVision Database: Visual Learning and Understanding from Web Data, August 2017. URL <http://arxiv.org/abs/1708.02862>. arXiv:1708.02862 [cs].
- Zhiyuan Li and Sanjeev Arora. An Exponential Learning Rate Schedule for Deep Learning. In *International Conference on Learning Representations*, March 2020. URL <https://openreview.net/forum?id=rJg8TeSFDH>.
- Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Zou. Mind the Gap: Understanding the Modality Gap in Multi-modal Contrastive Representation Learning. In *Advances in Neural Information Processing Systems*, 2022. URL <http://arxiv.org/abs/2203.02053>. arXiv: 2203.02053.
- Jialun Liu, Yifan Sun, Chuchu Han, Zhaopeng Dou, and Wenhui Li. Deep representation learning on long-tailed data: A learnable embedding augmentation perspective. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 2967–2976, February 2020. doi: 10.1109/CVPR42600.2020.00304. URL <https://arxiv.org/abs/2002.10826>. arXiv: 2002.10826 ISSN: 10636919.
- Xuantong Liu, Jianfeng Zhang, Tianyang Hu, He Cao, Yuan Yao, and Lujia Pan. Inducing Neural Collapse in Deep Long-tailed Learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 11534–11544. PMLR, 2023.
- 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 (CVPR)*, April 2019. doi: 10.48550/arxiv.1904.05160. URL <https://arxiv.org/abs/1904.05160>. arXiv: 1904.05160.
- Ekaterina Lobacheva, Maxim Kodryan, Nadezhda Chirkova, Andrey Malinin, and Dmitry P. Vetrov. On the periodic behavior of neural network training with batch normalization and weight decay. *Advances in Neural Information Processing Systems*, 34:21545–21556, 2021.
- Alexander Long, Wei Yin, Thalaisyasingam Ajanthan, Vu Nguyen, Pulak Purkait, Ravi Garg, Alan Blair, Chunhua Shen, and Anton van den Hengel. Retrieval Augmented Classification for Long-Tail Visual Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6959–6969, 2022. URL [https://openaccess.thecvf.com/content/CVPR2022/html/Long\\_Retrieval\\_Augmented\\_Classification\\_for\\_Long-Tail\\_Visual\\_Recognition\\_CVPR\\_2022\\_paper.html](https://openaccess.thecvf.com/content/CVPR2022/html/Long_Retrieval_Augmented_Classification_for_Long-Tail_Visual_Recognition_CVPR_2022_paper.html).
- Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, August 2017. doi: 10.48550/arxiv.1608.03983. URL <https://arxiv.org/abs/1608.03983>. arXiv: 1608.03983.

- Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. September 2018. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Jianfeng Lu and Stefan Steinerberger. Neural Collapse with Cross-Entropy Loss, January 2021. URL <http://arxiv.org/abs/2012.08465>. arXiv:2012.08465 [cs, math].
- Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the Generalization Benefit of Normalization Layers: Sharpness Reduction. In *Advances in Neural Information Processing Systems*, October 2022. URL <https://openreview.net/forum?id=xp5VOBxTxZ>.
- Teli Ma, Shijie Geng, Mengmeng Wang, Jing Shao, Jiasen Lu, Hongsheng Li, Peng Gao, and Yu Qiao. A Simple Long-Tailed Recognition Baseline via Vision-Language Model, November 2021. URL <https://arxiv.org/abs/2111.14745v1>.
- Yanbiao Ma, Licheng Jiao, Fang Liu, Yuxin Li, Shuyuan Yang, and Xu Liu. Delving into Semantic Scale Imbalance. In *The Eleventh International Conference on Learning Representations*, September 2022. URL <https://openreview.net/forum?id=07tc5kKRIO>.
- Dushyant Mehta, Kwang In Kim, and Christian Theobalt. On implicit filter level sparsity in convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 520–528, 2019.
- Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. *International Conference on Learning Representations*, July 2020. doi: 10.48550/arxiv.2007.07314. URL <http://arxiv.org/abs/2007.07314>. arXiv: 2007.07314.
- Vardan Papyan, X. Y. Han, and David L. Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences of the United States of America*, 117(40):24652–24663, August 2020. ISSN 10916490. doi: 10.1073/pnas.2015509117. URL <http://arxiv.org/abs/2008.08186>. arXiv: 2008.08186.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763. PMLR, July 2021. URL <https://proceedings.mlr.press/v139/radford21a.html>. ISSN: 2640-3498.
- Akshay Rangamani and Andrzej Banburski-Fahey. Neural Collapse in Deep Homogeneous Classifiers and The Role of Weight Decay. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, volume 2022-May, pp. 4243–4247, 2022. ISBN 978-1-66540-540-9. doi: 10.1109/ICASSP43922.2022.9746778. ISSN: 15206149.
- William J. Reed. The Pareto, Zipf and other power laws. *Economics Letters*, 74(1):15–19, December 2001. ISSN 01651765. doi: 10.1016/S0165-1765(01)00524-9. Publisher: North-Holland.
- Merrielle Spain and Pietro Perona. Measuring and Predicting Importance of Objects in Our Visual World, November 2007. URL <https://resolver.caltech.edu/CaltechAUTHORS:CNS-TR-2007-002>. Num Pages: 8 Place: Pasadena, CA Publisher: California Institute of Technology.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014. ISSN 1533-7928. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Thomas Strohmer and Robert W Heath. Grassmannian frames with applications to coding and communication. *Applied and Computational Harmonic Analysis*, 14(3):257–275, 2003. ISSN 10635203. doi: 10.1016/S1063-5203(03)00023-X. URL <https://www.sciencedirect.com/science/article/pii/S106352030300023X>. arXiv: math/0301135.
- Cecilia Summers and Michael J. Dinneen. Four Things Everyone Should Know to Improve Batch Normalization. In *International Conference on Learning Representations*, March 2020. URL <https://openreview.net/forum?id=HJx8HANFDH>.

- Christos Thrampoulidis, Ganesh R. Kini, Vala Vakilian, and Tina Behnia. Imbalance Trouble: Revisiting Neural-Collapse Geometry. *Advances in Neural Information Processing Systems*, August 2022. doi: 10.48550/arxiv.2208.05512. URL <http://arxiv.org/abs/2208.05512>. arXiv: 2208.05512.
- Changyao Tian, Wenhai Wang, Xizhou Zhu, Jifeng Dai, and Yu Qiao. VL-LTR: Learning Class-wise Visual-Linguistic Representation for Long-Tailed Visual Recognition. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision – ECCV 2022*, Lecture Notes in Computer Science, pp. 73–91, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-19806-9. doi: 10.1007/978-3-031-19806-9\_5.
- Mariya Toneva, Adam Trischler, Alessandro Sordoni, Yoshua Bengio, Remi Tachet Des Combes, and Geoffrey J. Gordon. An empirical study of example forgetting during deep neural network learning. *7th International Conference on Learning Representations, ICLR 2019*, December 2019. doi: 10.48550/arxiv.1812.05159. URL <http://arxiv.org/abs/1812.05159>. arXiv: 1812.05159.
- Rahul Vigneswaran, Marc T. Law, Vineeth N. Balasubramanian, and Makarand Tapaswi. Feature generation for long-tail classification. In *Proceedings of the Twelfth Indian Conference on Computer Vision, Graphics and Image Processing, ICVGIP '21*, pp. 1–9, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 978-1-4503-7596-2. doi: 10.1145/3490035.3490300. URL <https://dl.acm.org/doi/10.1145/3490035.3490300>.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. Matching Networks for One Shot Learning. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL [https://papers.nips.cc/paper\\_files/paper/2016/hash/90e1357833654983612fb05e3ec9148c-Abstract.html](https://papers.nips.cc/paper_files/paper/2016/hash/90e1357833654983612fb05e3ec9148c-Abstract.html).
- Ruosi Wan, Zhanxing Zhu, Xiangyu Zhang, and Jian Sun. Spherical Motion Dynamics: Learning Dynamics of Normalized Neural Network using SGD and Weight Decay. In *Advances in Neural Information Processing Systems*, volume 34, pp. 6380–6391. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/326a8c055c0d04f5b06544665d8bb3ea-Abstract.html>.
- Jianfeng Wang, Thomas Lukasiewicz, Xiaolin Hu, Jianfei Cai, and Zhenghua Xu. RSG: A Simple but Effective Module for Learning Imbalanced Datasets. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3783–3792, June 2021. doi: 10.1109/CVPR46437.2021.00378. ISSN: 2575-7075.
- 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*, pp. 1492–1500, 2017. URL [https://openaccess.thecvf.com/content\\_cvpr\\_2017/html/Xie\\_Aggregated\\_Residual\\_Transformations\\_CVPR\\_2017\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2017/html/Xie_Aggregated_Residual_Transformations_CVPR_2017_paper.html).
- Yibo Yang, Shixiang Chen, Xiangtai Li, Liang Xie, Zhouchen Lin, and Dacheng Tao. Inducing Neural Collapse in Imbalanced Learning: Do We Really Need a Learnable Classifier at the End of Deep Neural Network? In *Advances in Neural Information Processing Systems*, October 2022. URL [https://openreview.net/forum?id=A6EmxI3\\_Xc](https://openreview.net/forum?id=A6EmxI3_Xc).
- Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization. *7th International Conference on Learning Representations, ICLR 2019*, October 2019. doi: 10.48550/arxiv.1810.12281. URL <http://arxiv.org/abs/1810.12281>. arXiv: 1810.12281.
- Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep Long-Tailed Learning: A Survey. October 2021. doi: 10.48550/arxiv.2110.04596. URL <http://arxiv.org/abs/2110.04596>. arXiv: 2110.04596.

## A APPENDIX

## A.1 ACRONYM AND NOTATION

Tables 6 and 7 summarizes the abbreviation and notations used in this paper, respectively.

Table 6: Table of abbreviation.

Abbreviation	Definition
BN	batch normalization (Ioffe & Szegedy, 2015)
CB	class-balanced loss (Cui et al., 2019)
CE	cross entropy
DNN	deep neural network
ETF	equiangular tight frame
FDR	Fisher’s discriminant ratio (Fisher, 1936)
FR	feature regularization
LA	logit adjustment (Kim & Kim, 2020; Menon et al., 2020)
LTR	long-tailed recognition
MLP	multilayer perceptron
NC	neural collapse (Papayan et al., 2020)
ResBlock	residual block (He et al., 2016)
SELI	simplex-encoded-labels interpolation
SGD	stochastic gradient descent
WB	weight balancing (Alshammari et al., 2022)
WD	weight decay

Table 7: Table of notations.

Variable	Definition
$C$	number of classes
$d$	number of dimensions for features
$N / N_k$	number of samples / of class $k$
$\bar{N}$	harmonic mean of number of samples per class
$\lambda, \zeta$	hyper parameter of weight decay / feature regularization
$\rho$	imbalance factor: $\frac{\max_k N_k}{\min_k N_k}$
$\mathcal{X} / \mathcal{Y}$	domain of samples / labels
$\mathbf{x} / y / \mathbf{z}$	sample / label / logit
$\mathcal{D} / \mathcal{D}_k$	dataset of all classes / class $k$
$\mathbf{f} / \mathbf{g} / \mathbf{h}$	neural network of whole model / feature extractor / linear classifier
$\mathbf{g}(\mathbf{x}) / \bar{\mathbf{g}}(\mathbf{x})$	feature of sample $\mathbf{x}$ / normalized feature of sample $\mathbf{x}$
$\boldsymbol{\mu} / \boldsymbol{\mu}_k$	mean of inner-class mean of features / inner-class mean of features for class $k$
$\Theta / \Theta_g / \Theta_h$	set of parameters of $\mathbf{f} / \mathbf{g} / \mathbf{h}$
$\theta$	parameter in $\Theta$
$\mathbf{W} / \mathbf{w}_k$	linear layers’ weight matrix / vector of class $k$
$\ell / \ell_{\text{CE}} / \ell_{\text{CB}}$	loss function / of CE / of CB
$F / F_{\text{WB}}$	objective function / in the second training stage of WB
$\cos(\cdot, \cdot)$	cosine similarity of two vectors
$\tau / \gamma$	hyperparameter of additive LA / multiplicative LA

## A.2 RELATED WORK

## A.2.1 LONG-TAILED RECOGNITION

There are three main approaches to LTR; “Class Re-balancing, Information Augmentation, and Module Improvement” (Zhang et al., 2021). “Class Re-balancing” adjusts the imbalance in the number of samples per class at various stages to prevent accuracy deterioration. It includes logit

Table 8: Comparison of existing LTR methods with our simplified method. A symbol ✓ indicates that the component is required, - indicates that the component is not required. We regard the deferred re-balancing optimization schedule (Cao et al., 2019) as two-stage training. Our method does not require complex technique like many other methods.

Components \ Methods	Kim & Kim (2020)	Menon et al. (2020)	Yang et al. (2022)	Ma et al. (2022)	Liu et al. (2023)	Alshammari et al. (2022)	Cao et al. (2019)	Kang et al. (2023)	Li et al. (2022)	Long et al. (2022)	Ma et al. (2021)	Tian et al. (2022)	Ours
Number of training stages	1	1	1	3	2	2	2	2	2	1*	2	2*	1
Devising loss functions	-	-	✓	✓	-	✓	✓	✓	✓	✓	✓	✓	-
Resampling	-	-	-	-	✓	-	-	✓	✓	-	✓	-	-
Contrastive learning	-	-	-	-	-	-	-	✓	✓	-	✓	✓	-
Training of linear layer	✓	✓	-	✓	✓	✓	✓	✓	✓	✓	<sup>2</sup>	<sup>3</sup>	-
Extra text encoder	-	-	-	-	-	-	-	-	-	✓	✓	✓	-
Extra image encoder	-	-	-	-	-	-	-	-	-	✓	-	-	-
Extra dataset	-	-	-	-	-	-	-	-	-	-	-	✓	-

adjustment (LA) (Kim & Kim, 2020; Menon et al., 2020) and balancing the loss function such as CB and dynamic semantic-scale-balanced loss (Ma et al., 2022). “Information Augmentation” prevents the accuracy from degrading by supplementing the information of the tail classes that lacks the number of samples (Chu et al., 2020; Liu et al., 2020; Wang et al., 2021). “Module Improvement” improves accuracy by increasing the performance of each module of the network individually, e.g., training feature extractors and classifiers separately (Kang et al., 2020), fixing the linear layer to the ETF classifier for the better feature extractor (Yang et al., 2022), and regularizing to occur NC (Liu et al., 2023). In recent years, as Kang et al. (2023) found contrastive learning (Hadsell et al., 2006) is effective for imbalanced data, many methods have been proposed to use it (Kang et al., 2023; Li et al., 2022; Ma et al., 2021; Tian et al., 2022). Ma et al. (2021), Long et al. (2022), and Tian et al. (2022) also include vision-language models such as CLIP (Radford et al., 2021) and leverage text features to improve accuracy. However, these usually have problems such as slow convergence and complex models (Liu et al., 2023). WB is a combination of “Class Re-balancing” and “Module Improvement”. Table 8 compares our simplification of WB with typical existing methods, including those that have achieved SOTA. While many methods use complex innovations, we aimed to improve on a simple structure.

**Two-stage learning** Two-stage learning (Kang et al., 2020) is a method for improving accuracy by dividing LTR training into two stages: feature extractor training and classifier training. It is used in numerous methods (Cao et al., 2019; Ma et al., 2021; 2022; Li et al., 2022; Tian et al., 2022; Liu et al., 2023; Kang et al., 2023), including WB, because of its simple but significant improvement in accuracy. Note that since our work analyzes WB, the formulation of two-stage learning is based on WB. For example, the classifier weights are initialized randomly at the start of the second stage of training in Kang et al. (2020) but not in WB.

<sup>2</sup>This method does not require linear layers.

<sup>3</sup>This method requires the LGR head to be trained (Tian et al., 2022).



## A.3 PROOF

## A.3.1 PROOF OF THEOREM 1

*Proof.* Define  $p_k(\mathbf{g}(\mathbf{x}_i))$  as  $\frac{\exp(\mathbf{w}_k^\top \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\mathbf{w}_l^\top \mathbf{g}(\mathbf{x}_i))}$ . For  $\mathbf{x} \in \{\mathbf{x}_i, \mathbf{x}_j\}$ , we have

$$\begin{aligned} \left\| \frac{\partial \ell_{CE}}{\partial \mathbf{g}(\mathbf{x})} \right\|_2 &= \left\| \frac{\partial \ell_{CE}}{\partial \mathbf{g}(\mathbf{x})} \right\|_2 \|\mathbf{w}_y\|_2 \geq \left| \left\langle \frac{\partial \ell_{CE}}{\partial \mathbf{g}(\mathbf{x})}, \mathbf{w}_y \right\rangle \right| \\ &= \left| -(1 - p_y(\mathbf{g}(\mathbf{x}))) \mathbf{w}_y^\top \mathbf{w}_y + \sum_{l \neq y} p_l(\mathbf{g}(\mathbf{x})) \mathbf{w}_l^\top \mathbf{w}_y \right| \\ &= \left| -(1 - p_y(\mathbf{g}(\mathbf{x}))) - \frac{1}{C-1} \sum_{l \neq y} p_l(\mathbf{g}(\mathbf{x})) \right| = \frac{C}{C-1} (1 - p_y(\mathbf{g}(\mathbf{x}))) \\ &\geq (1 - p_y(\mathbf{g}(\mathbf{x}))). \end{aligned} \quad (4)$$

Therefore,

$$\begin{aligned} (1 - p_y(\mathbf{g}(\mathbf{x}))) &\leq \left\| \frac{\partial \ell_{CE}}{\partial \mathbf{g}(\mathbf{x})} \right\|_2 \leq \epsilon \\ \Rightarrow p_y(\mathbf{g}(\mathbf{x})) &\geq 1 - \epsilon. \end{aligned} \quad (5)$$

For  $p_y(\mathbf{g}(\mathbf{x}))$ , using Jensen's inequality, we get

$$\begin{aligned} p_y(\mathbf{g}(\mathbf{x})) &= \frac{\exp(\mathbf{w}_y^\top \mathbf{g}(\mathbf{x}))}{\sum_{l=1}^C \exp(\mathbf{w}_l^\top \mathbf{g}(\mathbf{x}))} = \frac{\exp(\mathbf{w}_y^\top \mathbf{g}(\mathbf{x}))}{\exp(\mathbf{w}_y^\top \mathbf{g}(\mathbf{x})) + (C-1) \sum_{l \neq y} \frac{1}{C-1} \exp(\mathbf{w}_l^\top \mathbf{g}(\mathbf{x}))} \\ &\leq \frac{\exp(\mathbf{w}_y^\top \mathbf{g}(\mathbf{x}))}{\exp(\mathbf{w}_y^\top \mathbf{g}(\mathbf{x})) + (C-1) \exp\left(\left(\frac{1}{C-1} \sum_{l \neq y} \mathbf{w}_l\right)^\top \mathbf{g}(\mathbf{x})\right)} \\ &= \frac{\exp(\mathbf{w}_y^\top \mathbf{g}(\mathbf{x}))}{\exp(\mathbf{w}_y^\top \mathbf{g}(\mathbf{x})) + (C-1) \exp\left(-\frac{1}{C-1} \mathbf{w}_y^\top \mathbf{g}(\mathbf{x})\right)} \\ &= \frac{1}{1 + (C-1) \exp\left(-\frac{C}{C-1} \mathbf{w}_y^\top \mathbf{g}(\mathbf{x})\right)}. \end{aligned} \quad (6)$$

Define  $\bar{\mathbf{g}}(\mathbf{x}_i)$  as  $\frac{\mathbf{g}(\mathbf{x}_i)}{\|\mathbf{g}(\mathbf{x}_i)\|_2}$ . By (5) and (6), we have

$$\begin{aligned} (C-1) \exp\left(-\frac{C}{C-1} \mathbf{w}_y^\top \mathbf{g}(\mathbf{x})\right) &\leq \frac{\epsilon}{1-\epsilon} \\ \Rightarrow \mathbf{w}_y^\top \bar{\mathbf{g}}(\mathbf{x}) &\geq \frac{1}{\|\mathbf{g}(\mathbf{x})\|_2} \frac{C-1}{C} \log\left(\frac{(C-1)(1-\epsilon)}{\epsilon}\right) \\ \Rightarrow \mathbf{w}_y^\top \bar{\mathbf{g}}(\mathbf{x}) &\geq \frac{1}{L} \frac{C-1}{C} \log\left(\frac{(C-1)(1-\epsilon)}{\epsilon}\right) \equiv \delta. \end{aligned} \quad (7)$$

Note that  $\epsilon \leq \frac{1}{C}$ ,  $L \leq 2\sqrt{2} \log(C-1) < 2\sqrt{2} \log(C-1) \frac{C}{C-1}$ , which means  $\delta > \frac{1}{\sqrt{2}}$ . Also,  $\delta$  is a lower bound of the inner product of the two unit vectors. Therefore, we have

$$\delta \in \left(\frac{1}{\sqrt{2}}, 1\right]. \quad (8)$$

Let  $\angle(\mathbf{s}, \mathbf{t}) \equiv \arccos(\cos(\mathbf{s}, \mathbf{t}))$ , which represents the angle between vectors  $\mathbf{s}$  and  $\mathbf{t}$ . For example, we have  $\angle(\mathbf{w}_{y_i}, \mathbf{w}_{y_j}) = \arccos\left(-\frac{1}{C-1}\right)$ . Using (8), the following holds.

$$\begin{aligned} \angle(\bar{\mathbf{g}}(\mathbf{x}_i), \bar{\mathbf{g}}(\mathbf{x}_j)) &\geq \angle(\mathbf{w}_{y_i}, \mathbf{w}_{y_j}) - \angle(\mathbf{w}_{y_i}, \bar{\mathbf{g}}(\mathbf{x}_i)) - \angle(\mathbf{w}_{y_j}, \bar{\mathbf{g}}(\mathbf{x}_j)) \\ &\geq \angle(\mathbf{w}_{y_i}, \mathbf{w}_{y_j}) - 2 \max_{\mathbf{x} \in \{\mathbf{x}_i, \mathbf{x}_j\}} (\angle(\mathbf{w}_y, \bar{\mathbf{g}}(\mathbf{x}))) \\ &\geq 0. \end{aligned} \quad (9)$$

Rewrite  $\angle(\mathbf{w}_{y_i}, \mathbf{w}_{y_j})$  and  $\max_{\mathbf{x} \in \{\mathbf{x}_i, \mathbf{x}_j\}} (\angle(\mathbf{w}_y, \bar{\mathbf{g}}(\mathbf{x})))$  as  $\alpha$  and  $\beta$  respectively. Since the cosine is monotonically decreasing in  $[0, \pi]$ , we get

$$\begin{aligned} \bar{\mathbf{g}}(\mathbf{x}_i)^\top \bar{\mathbf{g}}(\mathbf{x}_j) &\leq \cos(\alpha - 2\beta) \\ &= \cos \alpha (2 \cos^2 \beta - 1) + 2 \sin \alpha \sin \beta \cos \beta \\ &\leq \cos \alpha (2 \cos^2 \beta - 1) + 2\sqrt{1 - \cos^2 \beta} \cos \beta \end{aligned} \quad (10)$$

Note that  $x\sqrt{1-x^2}$  is monotonically decreasing and  $2x^2 \geq 1$  holds in  $\left[\frac{1}{\sqrt{2}}, 1\right]$ . Therefore, from (7), (8), and (10), the following holds.

$$\cos(\mathbf{g}(\mathbf{x}_i), \mathbf{g}(\mathbf{x}_j)) \leq 2\delta\sqrt{1-\delta^2}. \quad (11)$$

□

### A.3.2 PROOF OF THEOREM 2

Define  $\hat{\mathbf{w}}_k$  be  $\frac{\bar{N}}{\lambda \bar{N}} \boldsymbol{\mu}_k$ . We prove a more strict theorem. You can easily derive Theorem 2 from the following theorem.

**Theorem 3.** For any  $k \in \mathcal{Y}$ , if there exists  $\mathbf{w}_k^*$  s.t.  $\frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \Big|_{\mathbf{w}_k = \mathbf{w}_k^*} = \mathbf{0}$  and  $\|\mathbf{w}_k^*\|_2 = O\left(\frac{1}{\lambda \rho C}\right)$ , we have

$$\forall k \in \mathcal{Y}, \|\mathbf{w}_k^* - \hat{\mathbf{w}}_k\|_2 \leq \frac{\bar{N}}{\lambda N} \|\boldsymbol{\mu}\|_2 + O\left(\frac{1}{\lambda^2 \rho^2 C^2}\right). \quad (12)$$

Before proving the theorem, we first present the following lemmas.

**Lemma 1.** For the dataset of Imbalance rate  $\rho$ , the following holds.

$$\frac{\bar{N}}{N} = O\left(\frac{1}{\rho C}\right). \quad (13)$$

*Proof.* Since  $N_k = n\rho^{-\frac{k-1}{C-1}}$  holds, the harmonic mean and the number of all samples are

$$\bar{N} = \frac{C}{\sum_{k=1}^C \frac{1}{N_k}} = \frac{nC}{\sum_{k=1}^C \rho^{\frac{k-1}{C-1}}}, \quad (14)$$

$$N = \sum_{k=1}^C N_k = n \sum_{k=1}^C \rho^{-\frac{k-1}{C-1}} = \frac{n \sum_{k=1}^C \rho^{\frac{k-1}{C-1}}}{\rho}. \quad (15)$$

Therefore,

$$\frac{\bar{N}}{N} = \frac{\rho C}{\left(\sum_{k=1}^C \rho^{\frac{k-1}{C-1}}\right)^2} = \frac{\rho C \left(\rho^{\frac{1}{C-1}} - 1\right)^2}{\left(\rho^{\frac{C}{C-1}} - 1\right)^2}. \quad (16)$$

From (16),  $\frac{\bar{N}}{N} = O\left(\frac{1}{\rho}\right)$  ( $\rho \rightarrow \infty$ ) is obvious. In addition, the following holds.

$$\lim_{C \rightarrow \infty} \frac{\bar{N}C}{N} = \lim_{C \rightarrow \infty} \frac{\rho C^2}{\left(\sum_{k=1}^C \rho^{\frac{k-1}{C-1}}\right)^2} \leq \lim_{C \rightarrow \infty} \frac{\rho C^2}{C^2} = \rho. \quad (17)$$

Thus,  $\frac{\bar{N}}{N} = O\left(\frac{1}{C}\right)$  ( $C \rightarrow \infty$ ) is also satisfied.  $\square$

**Lemma 2.** When  $\mathbf{w}_k = \frac{\bar{N}}{\lambda N} \boldsymbol{\mu}_k$  is satisfied, we have

$$\forall k \in \mathcal{Y}, \frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \leq \frac{\bar{N}}{N} \boldsymbol{\mu} + O\left(\frac{1}{\lambda \rho^2 C^2}\right). \quad (18)$$

*Proof of Lemma 2.* For any  $k \in \mathcal{Y}$ , it holds that

$$\frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} = \frac{\bar{N}}{N} \sum_{j=1}^C \left( \frac{1}{N_j} \sum_{(\mathbf{x}_i, j) \in \mathcal{D}_j} \frac{\exp(\mathbf{w}_k^\top \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\mathbf{w}_l^\top \mathbf{g}(\mathbf{x}_i))} \mathbf{g}(\mathbf{x}_i) \right) - \frac{\bar{N}}{N} \boldsymbol{\mu}_k + \lambda \mathbf{w}_k. \quad (19)$$

In addition, using Lemma 1, we derive the following:

$$\begin{aligned} \frac{\exp(\mathbf{w}_k^\top \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\mathbf{w}_l^\top \mathbf{g}(\mathbf{x}_i))} &= \frac{1}{C} + \left( \frac{\exp(\mathbf{w}_k^\top \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\mathbf{w}_l^\top \mathbf{g}(\mathbf{x}_i))} - \frac{1}{C} \right) \\ &\leq \frac{1}{C} + \left( \frac{1 + O(\mathbf{w}_k^\top \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C (1 + O(\mathbf{w}_l^\top \mathbf{g}(\mathbf{x}_i)))} - \frac{1}{C} \right) \\ &= \frac{1}{C} + \frac{1}{C} \left( \left(1 + O\left(\frac{1}{\lambda \rho C}\right)\right) \left(1 - O\left(\frac{1}{\lambda \rho C^2}\right)\right) - 1 \right) \\ &= \frac{1}{C} + O\left(\frac{1}{\lambda \rho C^2}\right). \end{aligned} \quad (20)$$

Using this, the assumption and (19), we get

$$\frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \leq \frac{\bar{N}}{N} \boldsymbol{\mu} + O\left(\frac{1}{\lambda \rho^2 C^2}\right). \quad (21)$$

$\square$

*Proof of Theorem 3.*

$$\begin{aligned} \frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \Big|_{\mathbf{w}_k = \hat{\mathbf{w}}_k} &= \frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \Big|_{\mathbf{w}_k = \hat{\mathbf{w}}_k} - \frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \Big|_{\mathbf{w}_k = \mathbf{w}_k^*} \\ &= \lambda (\hat{\mathbf{w}}_k - \mathbf{w}_k^*) \\ &\quad + \frac{\bar{N}}{N} \sum_{j=1}^C \left( \frac{1}{N_j} \sum_{(\mathbf{x}_i, j) \in \mathcal{D}_j} \left( \frac{\exp(\hat{\mathbf{w}}_k^\top \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\hat{\mathbf{w}}_l^\top \mathbf{g}(\mathbf{x}_i))} - \frac{\exp(\mathbf{w}_k^{*\top} \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\mathbf{w}_l^{*\top} \mathbf{g}(\mathbf{x}_i))} \right) \mathbf{g}(\mathbf{x}_i) \right). \end{aligned} \quad (22)$$

Here, similar to (20), the following can be derived;

$$\frac{\exp(\hat{\mathbf{w}}_k^\top \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\hat{\mathbf{w}}_l^\top \mathbf{g}(\mathbf{x}_i))} - \frac{\exp(\mathbf{w}_k^{*\top} \mathbf{g}(\mathbf{x}_i))}{\sum_{l=1}^C \exp(\mathbf{w}_l^{*\top} \mathbf{g}(\mathbf{x}_i))} = \frac{1}{C} - \frac{1}{C} + O\left(\frac{1}{\lambda \rho C^2}\right) = O\left(\frac{1}{\lambda \rho C^2}\right). \quad (23)$$

Using (22) and (23),

$$\begin{aligned}\lambda(\hat{\mathbf{w}}_k - \mathbf{w}_k^*) &= \left. \frac{\partial F_{\text{WB}}}{\partial \mathbf{w}_k} \right|_{\mathbf{w}_k = \hat{\mathbf{w}}_k} + O\left(\frac{1}{\lambda \rho^2 C^2}\right) \\ &\leq \frac{\bar{N}}{N} \boldsymbol{\mu} + O\left(\frac{1}{\lambda \rho^2 C^2}\right).\end{aligned}\quad (24)$$

Therefore,

$$\|\mathbf{w}_k^* - \hat{\mathbf{w}}_k\|_2 \leq \frac{\bar{N}}{\lambda N} \|\boldsymbol{\mu}\|_2 + O\left(\frac{1}{\lambda^2 \rho^2 C^2}\right).\quad (25)$$

□

#### A.4 SETTINGS

Our implementation is based on Alshammari et al. (2022) and Vigneswaran et al. (2021). For comparison, many of the experimental settings also follow Alshammari et al. (2022).

##### A.4.1 DATASETS

We created validation datasets from the portions of the training datasets because CIFAR10 and CIFAR100 have only training and test data. As with Liu et al. (2019), only 20 samples per class were taken from the training dataset to compose the validation dataset, and the training dataset was composed of the rest of the data. We set  $N_1$  to 4980 for CIFAR10 and 480 for CIFAR100. In our experiments, we set imbalance factor  $\rho$  to 100. We call the class  $k$  *Many* if the number of training samples satisfies  $1000 < N_k \leq 4980$  (resp.  $100 < N_k \leq 480$ ), *Medium* if the number of training samples fullfills  $200 \leq N_k \leq 1000$  (resp.  $20 \leq N_k \leq 100$ ), and *Few* otherwise for CIFAR10-LT (resp. CIFAR100-LT). For mini-ImageNet-LT and ImageNet-LT, the same applies as for CIFAR100-LT.

##### A.4.2 MODELS

As for MLPs, one module block consists of three layers: a linear layer outputting 1024 dimensional features, a BN layer, and a ReLU layer. The layers are stacked in sequence and the blocks composed of them are as well. As for ResBlocks, each block has the same structure as He et al. (2016), except that the linear layer outputs 1024 dimensional features. These blocks are combined into a sequential, and at the bottom of it, we further concatenate an MLP block. In other words, the input first passes through the linear layer, BN, and ReLU before flowing into the residual blocks. In both cases, a classifier consisting of a linear layer and a softmax activation layer is on the top and does not count as one block. For example, in MLP3, the features pass through the three blocks and the classifier in sequence.

##### A.4.3 EVALUATION METRICS

Unless otherwise noted, we used the following values for hyperparameters for the ResNet. The optimizer was SGD with momentum = 0.9 and cosine learning rate scheduler (Loshchilov & Hutter, 2017) to gradually decrease the learning rate from 0.01 to 0. The batch size was 64, and the number of epochs was 320 for the first stage and 10 for the second stage. As loss functions, we used naive CE and CB. As for CB, we used class-balanced CE with  $\beta = 0.9999$ . We set  $\lambda$  of WD to 0.005 in the first stage and 0.1 in the second stage. For mini-ImageNet-LT, we set  $\lambda$  for the first stage to 0.003. We set  $\zeta$  of FR to 0.01. We calculated the MaxNorm’s threshold  $\eta$  as in Alshammari et al. (2022). We searched the optimal  $\tau$  and  $\gamma$  for the LA by cross-validation using the validation data. We chose the optimal value of  $\tau$  from  $\{1.00, 1.05, \dots, 2.00\}$  and  $\gamma$  from  $\{0.00, 0.05, \dots, 1.00\}$ .

For ImageNet-LT, We searched hyperparameters using the validation dataset for the ones not published in Alshammari et al. (2022) and used the following values. We reduced the learning rate gradually from 0.05 to 0. The number of epochs was 200 for the first stage. We set  $\lambda$  of WD to 0.00024 in the first stage and 0.00003 in the second stage. We set  $\zeta$  of FR to 0.0001. The other hyperparameters were the same as in CIFAR100-LT.

Table 9: FDRs and accuracy of models trained with an ETF classifier for CIFAR100-LT.

Method	FDR		Accuracy (%)			
	Train	Test	<i>Many</i>	<i>Medium</i>	<i>Few</i>	Average
WD&ETF	$3.33 \times 10^4 \pm 2.3 \times 10^3$	$1.13 \times 10^2 \pm 0.1 \times 10^1$	$76.3 \pm 0.3$	$46.0 \pm 0.4$	$15.5 \pm 0.6$	$46.9 \pm 0.2$
WD&ETF&DR	$8.25 \times 10^5 \pm 7.0 \times 10^4$	$9.46 \times 10^1 \pm 0.8 \times 10^0$	$72.9 \pm 0.2$	$43.0 \pm 0.6$	$20.0 \pm 0.9$	$46.0 \pm 0.4$

We trained MLPs and ResBlocks with  $\lambda$  set to 0.01 and the number of epochs set to 150. Other parameters were the same as above. FDRs and accuracy reported in our experiments were obtained by averaging the results from five training runs with different random seeds. We conducted experiments on an NVIDIA A100.

#### A.4.4 WHY WE DO NOT USE DR LOSS

This section explains why we use ETF classifiers but not the dot-regression loss (DR) proposed in Yang et al. (2022). Table 9 compares FDRs and accuracy when we use WD and an ETF Classifier for the linear layer with two options, using CE or DR for the loss. This table shows that CE is superior in the test FDR and average accuracy. Theorem 1 also claims that training by CE and WD prevents the cone effect. It does not apply to DR, which is a squared loss function. Thus, it is difficult to prove whether DR is equally effective in preventing the cone effect.

### A.5 EXPERIMENTS

#### A.5.1 EXPERIMENTS OF SEC.4.2

The right half of Figure 4 show the result of Sec.4.2 for mini-ImageNet-LT. We also examined the forget score (Toneva et al., 2019) of each method in Sec.4.2. Phenomena similar to the ones we showed in Sec.4.2 can be observed in the forgetting score of features trained with CB: Figure 3 shows that the forgetting score is higher on average with CB than with CE. In the absence of WD, this phenomenon is particularly evident for the *Many* classes, possibly because CB gives each image in *Many* classes less weight during training, which inhibits learning. These may be some of the reasons for the poor accuracy of CB compared with CE in training feature extractors reported by Kang et al. (2020).

#### A.5.2 EXPERIMENTS OF SEC. 4.3

**High relative variance in shifting parameters of BN temporarily degrades FDR** In Sec.4.3, we investigated the effect of the small scaling parameters of BN. What then is the impact of higher relative variance in the shifting parameters of BN? Our experiments have shown that they temporarily worsen the FDR in Figure 6. In the experiments, we trained MLPs with MNIST using three different methods: without WD (Naive), with WD (WD all), and with WD except for any shifting parameters of BN (WD w/o shift). We retrieved the output features of untrained and these trained models for each layer, applied ReLU to them, and examined their FDRs. We refer to these as ReLUFDRs. Figure 6 shows the increase and decrease in the ReLUFDRs of the intermediate outputs from the MLPs trained with each method. This figure indicates that the ReLUFDRs of the models trained without WD monotonically and gradually increases as the features pass through the layers. However, this is not the case for the models trained with WD. In particular, ReLUFDRs decrease drastically when we train the shifting parameters and apply the scaling and shifting parameters of the BN in the blocks close to the last layer. In this case, the ReLUFDRs increase when the features pass through the linear layer more than the ReLUFDRs decrease when the features pass through the BN later. These results suggest that applying WD to the scaling and shifting parameters of BN have a positive effect on training of the linear layer.

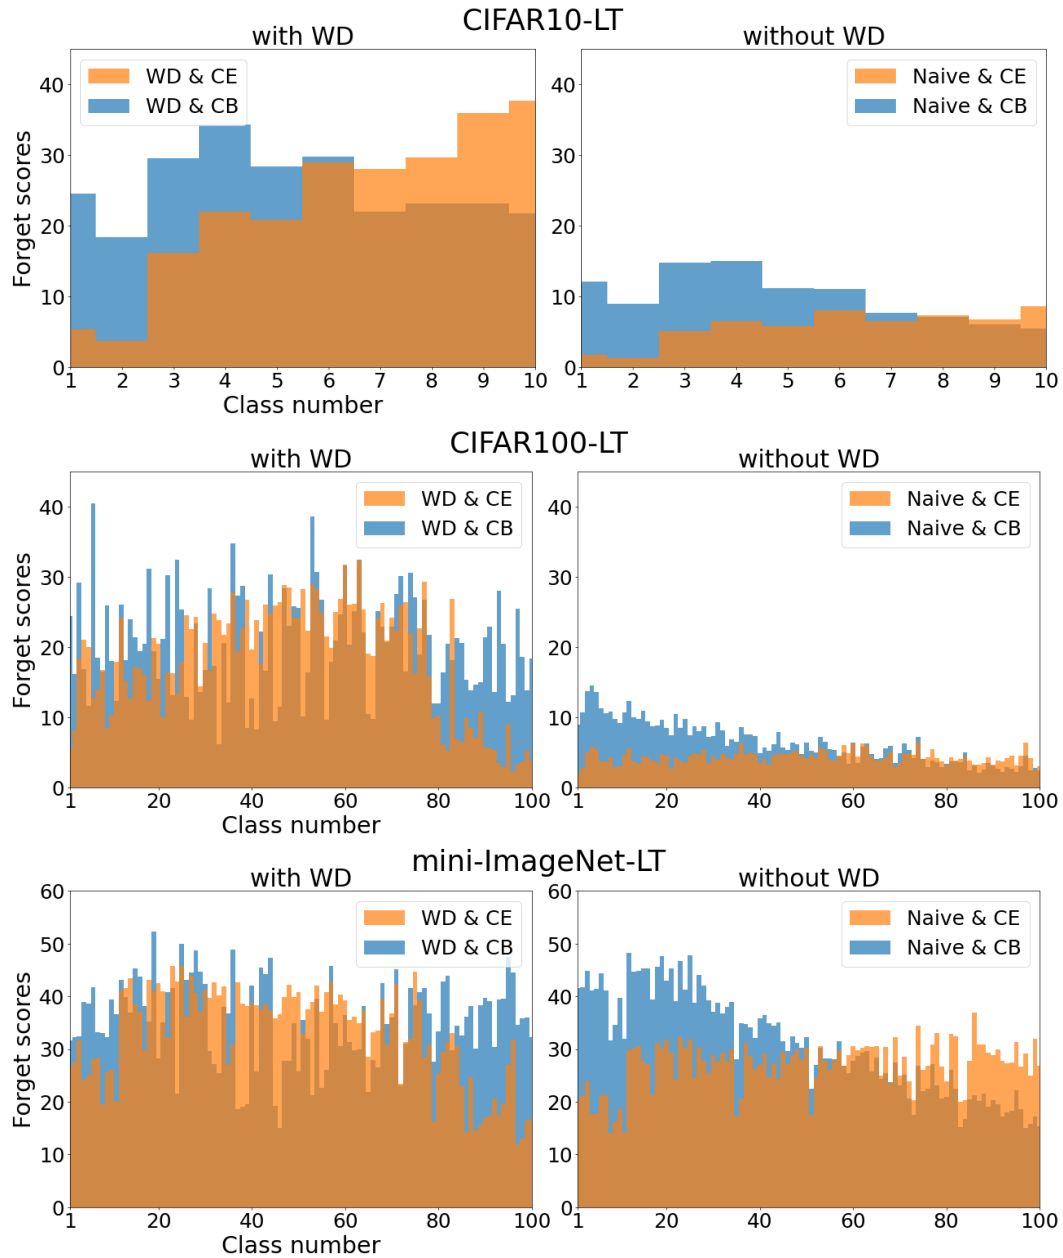


Figure 3: Average forgetting scores per class when models are trained with each method. These indicate higher forgetting scores when the models are trained with CB; this is particularly noticeable in the *Many* classes without WD.

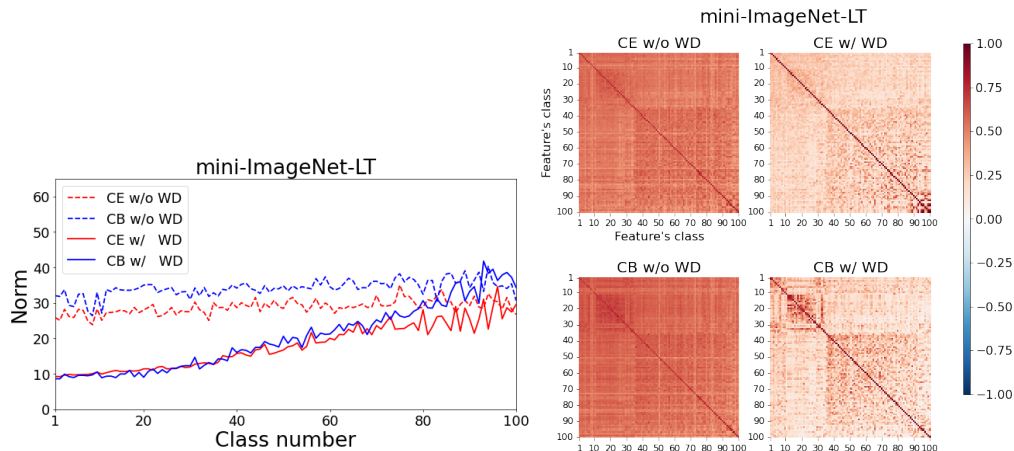


Figure 4: Results for mini-ImageNet-LT. (Left) Norm of mean per-class training features produced from models trained with each method. (Right) Heatmaps showing average cosine similarities of training features between two classes.




Figure 5: Norm ratio of mean per-class training features produced from models trained with each dataset and each method. Note that the vertical axis shows the ratio of the norm of the weights for each class with the one for the class of which sample size is the largest. Models trained with both methods have almost identical linear layer norms.




Figure 6: ReLU FDR of intermediate outputs of each MLP trained with each method. Training with WD decreases ReLU FDR significantly when features pass through the scaling and shifting layer in the BN of the final layer.

Table 10: Change in FDR as learned features pass through randomized linear layers and ReLUs multiple times. Red text indicates an increase compared to the FDR of the feature when the number of layers passed is one less, while blue text indicates a decrease.

Method	Number of times through randomized linear layers			
	0	1	2	3
CE	$2.44 \times 10^2$	$2.37 \times 10^2$	$2.17 \times 10^2$	$1.96 \times 10^2$
WD w/o BN	$4.91 \times 10^2$	$6.76 \times 10^2$	$6.56 \times 10^2$	$6.22 \times 10^2$
WD	$6.10 \times 10^2$	$2.47 \times 10^3$	$2.34 \times 10^3$	$1.60 \times 10^3$

### A.5.3 EXPERIMENTS OF SEC. 4.4

This subsection describes the experimental setup of Sec. 4.4 and additional experiments. We trained models for the balanced datasets using the same parameters settings as in Sec. 4.2. The FDRs of "before" are measured from the features obtained by models trained in the same way as the experiments in Sec. 4.2. These features are passed through an extra linear layer, and FDRs of them are measured as the "after". We initialized the weight ( $\in \mathbb{R}^{d \times d}$ ) and bias ( $\in \mathbb{R}^{d \times 1}$ ) of the extra linear layer with random values following a uniform distribution in  $[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}]$  and fixed them.

Table 10 shows the relationship between the FDRs and numbers of times features from MLP5 are applied by randomly initialized linear layers and ReLUs. The FDRs increase the first time, but gradually decrease the second and subsequent times.

### A.5.4 EXPERIMENTS OF SEC. 5.1

First, we experimentally show MaxNorm is not necessary for the second stage of WB. In the second training stage in WB, we observed that normalizing the norm of the weights to one before training (WD + WD & Norm & CB) instead of applying MaxNorm every epoch gives almost identical results as shown in Figure 5. This phenomenon is consistent with the fact that the regularization loses its effect beyond a certain epoch (Golotkar et al., 2019).

Figure 5 also shows that WB has difficulty in doing implicit LA when the number of classes is small. While the weights of the linear layer are larger for  $F_{ew}$  classes in the datasets with a sufficiently large number of classes, they remain lower for  $F_{ew}$  classes in CIFAR10-LT. These results are consistent with the conclusions drawn from Theorem 2.



Table 11: FDRs and accuracy of models trained with restricted WD and ETF Classifier for CIFAR100-LT.

Method	LA	FDR		Accuracy (%)			
		Train	Test	Many	Medium	Few	Average
WD	N/A	$1.94 \times 10^2$	$5.89 \times 10^1$	<b>74.6</b> $\pm 0.5$	44.4 $\pm 0.7$	13.4 $\pm 0.6$	45.1 $\pm 0.5$
w/o BN	Add	$1.94 \times 10^2$	$5.89 \times 10^1$	69.3 $\pm 0.8$	49.4 $\pm 0.5$	<b>30.4</b> $\pm 0.9$	<b>50.3</b> $\pm 0.6$
& ETF	Mult	$1.94 \times 10^2$	$5.89 \times 10^1$	67.6 $\pm 0.9$	<b>51.0</b> $\pm 0.5$	<b>31.5</b> $\pm 1.3$	<b>50.6</b> $\pm 0.7$
WD	N/A	<b><math>5.47 \times 10^4</math></b>	<b><math>1.05 \times 10^2</math></b>	<b>75.2</b> $\pm 0.6$	45.4 $\pm 0.9$	16.2 $\pm 0.4$	46.5 $\pm 0.1$
fixed BN	Add	<b><math>5.42 \times 10^4</math></b>	<b><math>1.05 \times 10^2</math></b>	71.9 $\pm 0.6$	44.8 $\pm 0.6$	26.5 $\pm 0.5$	48.4 $\pm 0.3$
&ETF	Mult	<b><math>5.42 \times 10^4</math></b>	<b><math>1.05 \times 10^2</math></b>	72.2 $\pm 0.8$	45.7 $\pm 0.5$	26.9 $\pm 0.4$	48.9 $\pm 0.3$

Table 12: FDRs and accuracy of models trained with each method for CIFAR10-LT.

Method	LA	FDR		Accuracy (%)			
		Train	Test	Many	Medium	Few	Average
CE	N/A	$8.17 \times 10^1$	$2.17 \times 10^1$	87.7 $\pm 0.3$	67.2 $\pm 3.1$	47.4 $\pm 2.2$	69.4 $\pm 1.0$
CB	N/A	$4.43 \times 10^1$	$1.50 \times 10^1$	81.7 $\pm 1.3$	60.6 $\pm 2.0$	42.0 $\pm 4.7$	63.5 $\pm 1.3$
WD	N/A	$2.97 \times 10^3$	$4.21 \times 10^1$	<b>89.1</b> $\pm 1.8$	76.6 $\pm 1.7$	61.5 $\pm 5.0$	77.1 $\pm 1.0$
	Add	$3.40 \times 10^3$	$4.21 \times 10^1$	81.6 $\pm 3.8$	<b>79.5</b> $\pm 1.2$	<b>79.7</b> $\pm 3.1$	<b>80.4</b> $\pm 0.7$
	Mult	$3.40 \times 10^3$	$4.21 \times 10^1$	86.2 $\pm 2.9$	<b>80.3</b> $\pm 1.0$	<b>74.0</b> $\pm 3.9$	<b>80.8</b> $\pm 0.5$
WB	N/A	$1.82 \times 10^3$	$3.86 \times 10^1$	<b>87.9</b> $\pm 2.4$	77.6 $\pm 1.5$	67.9 $\pm 3.2$	78.8 $\pm 0.6$
WD&ETF	N/A	$3.90 \times 10^4$	<b><math>4.38 \times 10^1</math></b>	<b>89.7</b> $\pm 1.9$	<b>73.9</b> $\pm 5.6$	59.1 $\pm 1.8$	75.8 $\pm 1.5$
	Add	$3.59 \times 10^4$	<b><math>4.38 \times 10^1</math></b>	<b>87.2</b> $\pm 3.4$	<b>77.1</b> $\pm 5.1$	72.5 $\pm 2.6$	<b>79.8</b> $\pm 1.1$
	Mult	$3.59 \times 10^4$	<b><math>4.38 \times 10^1</math></b>	<b>86.9</b> $\pm 4.2$	<b>79.0</b> $\pm 4.4$	73.2 $\pm 2.9$	<b>80.4</b> $\pm 0.8$
WD&FR &ETF	N/A	<b><math>5.56 \times 10^5</math></b>	<b><math>4.78 \times 10^1</math></b>	<b>90.9</b> $\pm 0.6$	76.7 $\pm 1.4$	56.2 $\pm 2.0$	76.2 $\pm 0.4$
	Add	<b><math>7.58 \times 10^5</math></b>	<b><math>4.78 \times 10^1</math></b>	88.8 $\pm 1.1$	<b>78.7</b> $\pm 1.2$	70.3 $\pm 2.1$	<b>80.2</b> $\pm 0.5$
	Mult	<b><math>7.58 \times 10^5</math></b>	<b><math>4.78 \times 10^1</math></b>	<b>89.9</b> $\pm 0.9$	<b>80.8</b> $\pm 1.3$	66.3 $\pm 2.1$	<b>80.1</b> $\pm 0.4$
WD	N/A	$9.17 \times 10^1$	$2.95 \times 10^1$	88.3 $\pm 1.4$	75.6 $\pm 3.6$	59.1 $\pm 1.8$	75.7 $\pm 1.6$
w/o BN	Add	$9.26 \times 10^1$	$2.95 \times 10^1$	85.9 $\pm 2.5$	<b>79.0</b> $\pm 3.2$	74.3 $\pm 1.5$	<b>80.3</b> $\pm 1.5$
& ETF	Mult	$9.26 \times 10^1$	$2.95 \times 10^1$	85.5 $\pm 2.5$	<b>79.0</b> $\pm 3.2$	<b>75.8</b> $\pm 1.0$	<b>80.7</b> $\pm 1.8$
WD	N/A	$7.02 \times 10^3$	$4.29 \times 10^1$	<b>89.5</b> $\pm 1.9$	74.3 $\pm 3.2$	58.9 $\pm 5.8$	75.7 $\pm 1.3$
fixed BN	Add	$6.60 \times 10^3$	$4.29 \times 10^1$	79.3 $\pm 9.2$	<b>76.9</b> $\pm 3.6$	<b>78.7</b> $\pm 5.4$	<b>78.4</b> $\pm 3.2$
&ETF	Mult	$6.60 \times 10^3$	$4.29 \times 10^1$	<b>84.5</b> $\pm 6.9$	<b>78.5</b> $\pm 4.1$	<b>77.1</b> $\pm 5.6$	<b>80.5</b> $\pm 2.6$

#### A.5.5 EXPERIMENTS OF SEC.5.2

We show the FDRs and accuracy of methods with restricted WD for CIFAR100-LT in Table 11. Table 12, 13, and Table 14 show the results for CIFAR10-LT, mini-ImageNet-LT, and ImageNet-LT respectively. We set  $\zeta$  for CIFAR10-LT, mini-ImageNet-LT, and ImageNet-LT to 0.02, 0.001, and 0.0001 respectively.

We also experimented with the ResNeXt50 (Xie et al., 2017) as with Alshammari et al. (2022). We used the same hyperparameters as the ResNet. Table 15 presents the results for CIFAR100-LT.

Table 13: FDRs and accuracy of models trained with each method for mini-ImageNet-LT.

Method	LA	FDR		Accuracy (%)			
		Train	Test	Many	Medium	Few	Average
CE	N/A	$7.56 \times 10^1$ $\pm 1.6 \times 10^0$	$4.28 \times 10^1$ $\pm 0.3 \times 10^0$	$72.1 \pm 0.5$	$34.1 \pm 0.6$	$18.3 \pm 0.5$	$42.7 \pm 0.2$
CB	N/A	$4.68 \times 10^1$ $\pm 0.7 \times 10^0$	$2.93 \times 10^1$ $\pm 0.2 \times 10^0$	$59.4 \pm 0.5$	$25.0 \pm 0.4$	$15.7 \pm 0.4$	$34.3 \pm 0.2$
WD	N/A	$6.58 \times 10^1$ $\pm 2.8 \times 10^1$	$1.01 \times 10^2$ $\pm 0.1 \times 10^1$	<b><math>81.7 \pm 0.4</math></b>	<b><math>43.3 \pm 0.5</math></b>	$20.3 \pm 0.7$	$49.8 \pm 0.3$
	Add	$6.51 \times 10^1$ $\pm 3.3 \times 10^1$	$1.01 \times 10^2$ $\pm 0.1 \times 10^1$	$80.5 \pm 0.3$	$42.1 \pm 0.6$	$39.3 \pm 0.2$	<b><math>54.7 \pm 0.3</math></b>
WB	Mult	$6.51 \times 10^1$ $\pm 3.3 \times 10^1$	$1.01 \times 10^2$ $\pm 0.1 \times 10^1$	$79.9 \pm 0.2$	$41.6 \pm 0.8$	$40.3 \pm 0.3$	<b><math>54.6 \pm 0.4</math></b>
	N/A	$6.52 \times 10^1$ $\pm 3.0 \times 10^1$	$1.00 \times 10^2$ $\pm 0.1 \times 10^1$	$80.1 \pm 0.2$	<b><math>43.9 \pm 0.7</math></b>	$38.1 \pm 0.5$	<b><math>54.8 \pm 0.3</math></b>
WD&ETF	N/A	$8.20 \times 10^1$ $\pm 6.5 \times 10^1$	$1.09 \times 10^2$ $\pm 0.2 \times 10^1$	<b><math>81.7 \pm 0.7</math></b>	<b><math>42.9 \pm 0.4</math></b>	$21.6 \pm 1.2$	$50.1 \pm 0.2$
	Add	$8.27 \times 10^1$ $\pm 8.5 \times 10^1$	$1.09 \times 10^2$ $\pm 0.2 \times 10^1$	$79.6 \pm 0.8$	<b><math>43.0 \pm 0.6</math></b>	$37.5 \pm 0.7$	$54.1 \pm 0.2$
	Mult	$8.27 \times 10^1$ $\pm 8.5 \times 10^1$	$1.09 \times 10^2$ $\pm 0.2 \times 10^1$	$78.8 \pm 1.1$	<b><math>42.7 \pm 1.0</math></b>	$39.4 \pm 1.0$	<b><math>54.3 \pm 0.3</math></b>
WD&FR &ETF	N/A	<b><math>1.32 \times 10^3</math></b> $\pm 1.2 \times 10^2$	$1.20 \times 10^2$ $\pm 0.2 \times 10^1$	<b><math>81.7 \pm 0.6</math></b>	<b><math>43.1 \pm 0.7</math></b>	$20.8 \pm 0.7$	$49.9 \pm 0.5$
	Add	<b><math>1.31 \times 10^3</math></b> $\pm 1.1 \times 10^2$	$1.20 \times 10^2$ $\pm 0.2 \times 10^1$	$79.6 \pm 0.5$	<b><math>43.5 \pm 0.6</math></b>	$34.9 \pm 0.9$	$53.6 \pm 0.6$
	Mult	<b><math>1.31 \times 10^3</math></b> $\pm 1.1 \times 10^2$	$1.20 \times 10^2$ $\pm 0.2 \times 10^1$	$78.8 \pm 0.3$	<b><math>43.5 \pm 0.8</math></b>	$37.9 \pm 0.8$	<b><math>54.2 \pm 0.5</math></b>
WD w/o BN & ETF	N/A	$1.21 \times 10^2$ $\pm 0.1 \times 10^1$	$5.85 \times 10^1$ $\pm 0.3 \times 10^0$	$78.6 \pm 0.4$	$41.7 \pm 0.6$	$21.7 \pm 0.7$	$48.6 \pm 0.3$
	Add	$1.21 \times 10^2$ $\pm 0.1 \times 10^1$	$5.85 \times 10^1$ $\pm 0.3 \times 10^0$	$75.1 \pm 0.2$	$40.3 \pm 0.7$	<b><math>41.7 \pm 0.6</math></b>	$52.9 \pm 0.3$
	Mult	$1.21 \times 10^2$ $\pm 0.1 \times 10^1$	$5.85 \times 10^1$ $\pm 0.3 \times 10^0$	$76.2 \pm 0.3$	$41.0 \pm 0.5$	$39.6 \pm 0.6$	$52.9 \pm 0.3$
WD fixed BN &ETF	N/A	$8.55 \times 10^1$ $\pm 2.0 \times 10^1$	<b><math>1.25 \times 10^2</math></b> $\pm 0.1 \times 10^1$	$80.6 \pm 0.2$	<b><math>42.8 \pm 0.5</math></b>	$20.3 \pm 0.5$	$49.3 \pm 0.1$
	Add	$8.56 \times 10^1$ $\pm 2.1 \times 10^1$	<b><math>1.25 \times 10^2</math></b> $\pm 0.1 \times 10^1$	$78.7 \pm 0.3$	$41.6 \pm 0.4$	$36.0 \pm 0.7$	$52.9 \pm 0.3$
	Mult	$8.56 \times 10^1$ $\pm 2.1 \times 10^1$	<b><math>1.25 \times 10^2</math></b> $\pm 0.1 \times 10^1$	$80.5 \pm 0.2$	<b><math>43.6 \pm 0.3</math></b>	$25.6 \pm 0.5$	$51.1 \pm 0.1$

Table 14: FDRs and accuracy of models trained with each method for ImageNet-LT.

Method	LA	FDR		Accuracy (%)			
		Train	Test	Many	Medium	Few	Average
CE	N/A	$1.35 \times 10^2$ $\pm 0.2 \times 10^1$	$1.58 \times 10^2$ $\pm 0.2 \times 10^1$	$54.7 \pm 0.2$	$30.1 \pm 0.4$	$12.1 \pm 0.2$	$37.1 \pm 0.3$
CB	N/A	$1.00 \times 10^2$ $\pm 0.8 \times 10^1$	$1.34 \times 10^2$ $\pm 0.9 \times 10^1$	$47.9 \pm 2.4$	$24.1 \pm 2.0$	$8.19 \pm 1.1$	$31.1 \pm 2.0$
WD	N/A	$3.79 \times 10^1$ $\pm 0.8 \times 10^1$	$2.99 \times 10^1$ $\pm 0.3 \times 10^1$	$67.4 \pm 0.5$	$42.0 \pm 0.4$	$15.2 \pm 0.3$	$48.1 \pm 0.4$
	Add	$3.77 \times 10^1$ $\pm 0.8 \times 10^1$	$2.99 \times 10^1$ $\pm 0.3 \times 10^1$	$62.9 \pm 0.4$	$48.8 \pm 0.5$	$35.2 \pm 0.3$	$52.4 \pm 0.4$
WB	Mult	$3.77 \times 10^1$ $\pm 0.8 \times 10^1$	$2.99 \times 10^1$ $\pm 0.3 \times 10^1$	$62.8 \pm 0.4$	<b><math>49.3 \pm 0.6</math></b>	$34.4 \pm 0.4$	<b><math>52.5 \pm 0.4</math></b>
	N/A	$3.78 \times 10^1$ $\pm 0.8 \times 10^1$	$2.99 \times 10^1$ $\pm 0.3 \times 10^1$	$62.4 \pm 0.5$	<b><math>50.0 \pm 0.5</math></b>	$31.1 \pm 0.5$	$52.2 \pm 0.4$
WD&ETF	N/A	$6.98 \times 10^1$ $\pm 1.1 \times 10^1$	$4.56 \times 10^1$ $\pm 0.7 \times 10^1$	<b><math>68.4 \pm 0.4</math></b>	$43.6 \pm 0.4$	$17.7 \pm 0.7$	$49.6 \pm 0.4$
	Add	$6.95 \times 10^1$ $\pm 1.0 \times 10^1$	$4.56 \times 10^1$ $\pm 0.7 \times 10^1$	$64.0 \pm 0.3$	<b><math>49.6 \pm 0.3</math></b>	<b><math>35.5 \pm 0.4</math></b>	<b><math>53.2 \pm 0.3</math></b>
	Mult	$6.95 \times 10^1$ $\pm 1.0 \times 10^1$	$4.56 \times 10^1$ $\pm 0.7 \times 10^1$	$63.4 \pm 0.4$	<b><math>50.0 \pm 0.4</math></b>	$35.1 \pm 0.5$	<b><math>53.2 \pm 0.3</math></b>
WD&FR &ETF	N/A	<b><math>1.03 \times 10^3</math></b> $\pm 0.4 \times 10^2$	<b><math>5.46 \times 10^2</math></b> $\pm 0.6 \times 10^1$	<b><math>68.1 \pm 0.2</math></b>	$43.0 \pm 0.2$	$17.0 \pm 0.5$	$49.1 \pm 0.2$
	Add	<b><math>1.03 \times 10^3</math></b> $\pm 0.4 \times 10^2$	<b><math>5.46 \times 10^2</math></b> $\pm 0.6 \times 10^1$	$63.3 \pm 0.3$	$48.7 \pm 0.1$	<b><math>36.4 \pm 0.7</math></b>	$52.6 \pm 0.2$
	Mult	<b><math>1.03 \times 10^3</math></b> $\pm 0.4 \times 10^2$	<b><math>5.46 \times 10^2</math></b> $\pm 0.6 \times 10^1$	$62.4 \pm 0.4$	$49.3 \pm 0.2$	<b><math>36.0 \pm 0.5</math></b>	$52.6 \pm 0.2$

Table 15: FDRs and accuracy of ResNeXt50 trained with each method for CIFAR100-LT.

Method	LA	FDR		Accuracy (%)			
		Train	Test	Many	Medium	Few	Average
CE	N/A	$2.04 \times 10^2$ $\pm 0.6 \times 10^1$	$7.32 \times 10^1$ $\pm 1.1 \times 10^0$	$59.9 \pm 0.4$	$32.8 \pm 0.7$	$9.51 \pm 0.3$	$34.8 \pm 0.4$
CB	N/A	$1.37 \times 10^2$ $\pm 1.6 \times 10^1$	$5.19 \times 10^1$ $\pm 1.1 \times 10^0$	$44.8 \pm 1.4$	$20.7 \pm 0.9$	$4.62 \pm 0.4$	$23.9 \pm 0.9$
WD	N/A	$8.36 \times 10^4$ $\pm 7.4 \times 10^3$	$2.06 \times 10^2$ $\pm 0.3 \times 10^1$	<b><math>77.9 \pm 0.2</math></b>	$48.3 \pm 0.5$	$14.9 \pm 0.8$	$48.0 \pm 0.3$
	Add	$8.41 \times 10^4$ $\pm 7.9 \times 10^3$	$2.06 \times 10^2$ $\pm 0.3 \times 10^1$	$73.4 \pm 0.5$	$48.6 \pm 0.8$	<b><math>32.9 \pm 0.8</math></b>	$52.2 \pm 0.2$
WB	Mult	$8.41 \times 10^4$ $\pm 7.9 \times 10^3$	$2.06 \times 10^2$ $\pm 0.3 \times 10^1$	$73.7 \pm 0.6$	$49.2 \pm 1.1$	<b><math>33.8 \pm 1.5</math></b>	$52.8 \pm 0.2$
	N/A	$3.19 \times 10^5$ $\pm 2.0 \times 10^4$	$2.06 \times 10^2$ $\pm 0.2 \times 10^1$	<b><math>77.5 \pm 0.2</math></b>	$51.2 \pm 0.9$	$21.1 \pm 0.6$	$50.8 \pm 0.3$
WD&ETF	N/A	$1.67 \times 10^5$ $\pm 2.4 \times 10^4$	$2.02 \times 10^2$ $\pm 0.3 \times 10^1$	<b><math>77.7 \pm 0.6</math></b>	$48.8 \pm 0.9$	$17.4 \pm 0.5$	$48.9 \pm 0.4$
	Add	$1.67 \times 10^5$ $\pm 2.4 \times 10^4$	$2.02 \times 10^2$ $\pm 0.3 \times 10^1$	$72.7 \pm 0.9$	$50.3 \pm 1.0$	$31.1 \pm 0.5$	$52.0 \pm 0.6$
WD&FR &ETF	Mult	$1.67 \times 10^5$ $\pm 2.4 \times 10^4$	$2.02 \times 10^2$ $\pm 0.3 \times 10^1$	$74.6 \pm 0.7$	<b><math>54.0 \pm 1.0</math></b>	$30.9 \pm 0.9$	<b><math>53.8 \pm 0.4</math></b>
	N/A	<b><math>3.19 \times 10^5</math></b> $\pm 1.8 \times 10^4$	$2.06 \times 10^2$ $\pm 0.2 \times 10^1$	<b><math>77.4 \pm 0.3</math></b>	$49.8 \pm 0.5$	$18.2 \pm 0.4$	$49.4 \pm 0.3$
WD&FR &ETF	Add	<b><math>3.17 \times 10^5</math></b> $\pm 1.8 \times 10^4$	$2.06 \times 10^2$ $\pm 0.2 \times 10^1$	$76.3 \pm 0.3$	$51.9 \pm 0.4$	$26.2 \pm 0.4$	$52.2 \pm 0.2$
	Mult	<b><math>3.17 \times 10^5</math></b> $\pm 1.8 \times 10^4$	$2.06 \times 10^2$ $\pm 0.2 \times 10^1$	$72.9 \pm 0.5$	<b><math>53.8 \pm 1.1</math></b>	<b><math>33.4 \pm 0.6</math></b>	<b><math>54.0 \pm 0.3</math></b>

#### A.5.6 EXPERIMENTS ON TABULAR DATA

We considered analyzing tabular data as an experiment for data with completely different characteristics and features from image data. We used Helena, a dataset for classification with 100 classes. Since the data is not divided for validation and test, we randomly extracted 20 samples per class without duplicates. The distribution of the training data is similar to that of long-tailed data with  $\rho \simeq 40$ . We call the class  $k$  *Many* if the number of training samples satisfies  $500 < N_k$ , *Medium* if the number of training samples fullfills  $200 \leq N_k \leq 500$ , and *Few* otherwise for this dataset.

Kadra et al. (2021) show that even an MLP with carefully tuned regularization outperforms state-of-the-art models in the classification of tabular data. Following them, we used a 9-layer MLP with 512 dimensions per layer for the feature extractor. The model is trained by AdamW (Loshchilov & Hutter, 2018) for 400 epochs in the first stage and 10 epochs in the second stage. Thus, WD is not implemented in L2 regularization in this dataset, but is built into the optimizer. By a hyperparameter search with the validation data, we set the hyperparameters as follows. We used dropout (Srivastava et al., 2014) and set the hyperparameter to 0.15. The initial learning rate was 0.001 for the first stage and 0.0004 for the second stage. We set  $\lambda$  of WD to 0.15 for the first stage and 0.0003 for the second stage. We set  $\zeta$  of FR to 0.001.

Table 16 presents the results. As in the experiment with image data, the proposed method outperforms the existing methods in both test FDR and average accuracy. Note that we used AdamW for the optimizer instead of SGD and that the proposed method also succeeds in this case. This result indicates that the proposed method works for optimizers other than SGD.

#### A.6 BROADER IMPACTS

Our research provides theoretical and experimental evidence for the effectiveness of an existing ad-hoc method. We also show that the original method can be simplified based on this theory. The theory is valid for general DNNs and does not concern any social issues such as privacy. On the contrary, it reduces the number of training stages to one while maintaining a higher level of accuracy, thus reducing the computational cost and the negative impact on the environment.

Table 16: FDRs and accuracy of MLP trained with each method for Helena.

Method	LA	FDR		Accuracy (%)			
		Train	Test	Many	Medium	Few	Average
CE	N/A	$5.40 \times 10^1$ $\pm 5.7 \times 10^0$	$6.48 \times 10^1$ $\pm 1.1 \times 10^0$	$34.7 \pm 0.8$	$20.8 \pm 1.1$	$17.2 \pm 0.5$	$24.3 \pm 0.5$
CB	N/A	$4.31 \times 10^1$ $\pm 7.1 \times 10^0$	$6.53 \times 10^1$ $\pm 0.7 \times 10^0$	$26.4 \pm 0.8$	$27.5 \pm 1.0$	<b><math>25.1 \pm 1.8</math></b>	$26.3 \pm 0.8$
WD	N/A	<b><math>7.61 \times 10^1</math></b> $\pm 1.0 \times 10^0$	$6.53 \times 10^1$ $\pm 0.6 \times 10^0$	<b><math>36.3 \pm 0.6</math></b>	$20.8 \pm 0.4$	$16.6 \pm 1.1$	$24.7 \pm 0.6$
	Add	<b><math>7.61 \times 10^1</math></b> $\pm 1.0 \times 10^0$	$6.53 \times 10^1$ $\pm 0.6 \times 10^0$	$33.4 \pm 0.4$	$26.3 \pm 0.6$	$24.5 \pm 0.8$	$28.1 \pm 0.2$
	Mult	<b><math>7.61 \times 10^1</math></b> $\pm 1.0 \times 10^0$	$6.53 \times 10^1$ $\pm 0.6 \times 10^0$	$31.6 \pm 0.3$	$27.1 \pm 0.5$	$21.9 \pm 1.0$	$26.9 \pm 0.3$
WB	N/A	<b><math>7.63 \times 10^1</math></b> $\pm 0.8 \times 10^0$	$6.55 \times 10^1$ $\pm 0.7 \times 10^0$	$35.0 \pm 0.5$	$26.0 \pm 0.8$	$22.7 \pm 1.4$	$28.0 \pm 0.5$
WD&ETF	N/A	<b><math>7.73 \times 10^1</math></b> $\pm 7.2 \times 10^0$	$6.86 \times 10^1$ $\pm 0.3 \times 10^0$	<b><math>36.2 \pm 0.6</math></b>	$22.6 \pm 1.1$	$17.7 \pm 0.3$	$25.6 \pm 0.4$
	Add	<b><math>7.73 \times 10^1</math></b> $\pm 7.2 \times 10^0$	$6.86 \times 10^1$ $\pm 0.3 \times 10^0$	$32.6 \pm 0.9$	$28.8 \pm 0.6$	$25.5 \pm 0.9$	<b><math>29.0 \pm 0.4</math></b>
	Mult	<b><math>7.73 \times 10^1</math></b> $\pm 7.2 \times 10^0$	$6.86 \times 10^1$ $\pm 0.3 \times 10^0$	$32.1 \pm 1.0$	<b><math>30.1 \pm 0.6</math></b>	$25.4 \pm 0.9$	<b><math>29.2 \pm 0.3</math></b>
WD&FR &ETF	N/A	<b><math>8.49 \times 10^1</math></b> $\pm 14.9 \times 10^0$	<b><math>7.04 \times 10^1</math></b> $\pm 0.7 \times 10^0$	<b><math>36.1 \pm 1.3</math></b>	$21.6 \pm 0.6$	$17.4 \pm 1.0$	$25.2 \pm 0.3$
	Add	<b><math>8.49 \times 10^1</math></b> $\pm 14.9 \times 10^0$	<b><math>7.04 \times 10^1</math></b> $\pm 0.7 \times 10^0$	$31.3 \pm 1.0$	$28.7 \pm 0.6$	<b><math>27.5 \pm 0.6</math></b>	<b><math>29.2 \pm 0.4</math></b>
	Mult	<b><math>8.49 \times 10^1</math></b> $\pm 14.9 \times 10^0$	<b><math>7.04 \times 10^1</math></b> $\pm 0.7 \times 10^0$	$31.9 \pm 1.1$	<b><math>29.6 \pm 0.9</math></b>	$25.6 \pm 0.5$	<b><math>29.1 \pm 0.5</math></b>