Robust Long-Tailed Learning via Label-Aware Bounded CVaR

Hong Zhu[*]
Huawei
zhuhong8@huawei.com

Runpeng Yu[*]
National University of Singapore
r.yu@u.nus.edu

Xing Tang
Tencent
shawntang@tencent.com

Yifei Wang
Peking University
yifei_wang@pku.edu.cn

Yuan Fang
Huawei
frank.fy@huawei.com

Yisen Wang
Peking University
yisen.wang@pku.edu.cn

Abstract

Data in the real-world classification problems are always imbalanced or long-tailed, wherein the majority classes have the most of the samples that dominate the model training. In such setting, the naive model tends to have poor performance on the minority classes. Previously, a variety of loss modifications have been proposed to address the long-tailed leaning problem, while these methods either treat the samples in the same class indiscriminatingly or lack a theoretical guarantee. In this paper, we propose two novel approaches based on CVaR (Conditional Value at Risk) to improve the performance of long-tailed learning with a solid theoretical ground. Specifically, we firstly introduce a *Label-Aware Bounded CVaR* (LAB-CVaR) loss to overcome the pessimistic result of the original CVaR, and further design the optimal weight bounds for LAB-CVaR theoretically. Based on LAB-CVaR, we additionally propose a *LAB-CVaR with logit adjustment* (LAB-CVaR-logit) loss to stabilize the optimization process, where we also offer the theoretical support. Extensive experiments on real-world datasets with long-tailed label distributions verify the superiority of our proposed methods.

1 Introduction

In real-world classification problems, there always have *long-tailed* label distributions, wherein the minority classes have much fewer samples than the majority classes [32, 18, 1, 37]. In such long-tailed setting, the majority classes dominate the model training, which results in poor performance on the minority classes [23, 10]. This is particularly undesirable when the minority classes are important, such as in the applications of cancer diagnosis and credit card fraud detection [35]. Therefore, how to address the long-tailed learning problem is critical.

The extant methods for long-tailed learning could be divided into 4 categories: 1) under/over sampling of the inputs [14, 33, 3], 2) post-hoc corrections of the outputs [7, 4, 12], 3) training procedures manipulation [25, 31, 34], and 4) loss function modifications [1, 23, 5, 35]. Besides using one of them alone, methods are further combined together to complement each other, e.g. the under/over sampling method may be easily combined with the methods in the other three categories [23], and the loss modification method may be promoted by the combination with a proper training procedure [25, 1]. In this paper, we focus on the loss modification methods, which assign weights [9, 5] or relative margins [1, 23, 30] to the training loss aiming to balance the label distribution. For the class-level loss modification methods [9, 5, 35], their notable disadvantage is that they do not discriminatingly

---

[*]Both authors contributed equally to this research.

treat the samples in the same class. On the other hand, the sample-level loss modification methods [17, 15] usually lack a theoretical guarantee.

In this paper, we propose novel methods for long-tailed learning based on CVaR (Conditional Value at Risk) which could consider the difference between the samples in the same class. As one of the most popular distributionally robust optimization (DRO) methods, CVaR aims to maximize the worst-case performance, and then can naturally improve the model performance on minority classes [36, 25]. However, in the classification task in which the model is evaluated by the zero-one loss at test time, the original CVaR and DRO are proved to result in the same model performance with ERM (empirical risk minimization) [8, 36]. Different from the solutions proposed in [8] and [36], we propose another novel approach to overcome the pessimistic result of the original CVaR in the classification task. This approach is named as *label-aware bounded CVaR* (LAB-CVaR) which means that the weight bounds are varied with labels. Furthermore, we theoretically give the optimal design of the upper and lower weight bounds for each class aiming to minimize the LAB-CVaR loss. On the other hand, LAB-CVaR is a special re-weighting method. But, the re-weighting strategies themselves may cause difficulties and instability in optimization [1, 5]. Thus, we additionally propose an approach named *LAB-CVaR with logit adjustment* (LAB-CVaR-logit) to alleviate the influence of the loss weight in optimization, and theoretically study the mechanism that how the logit adjustment technique can improve the stability in optimization. In summary, our contributions are:

- We propose the LAB-CVaR loss which can overcome the pessimistic result of the original CVaR, and theoretically give its optimal design of the upper and lower weight bounds by aiming to minimize the LAB-CVaR loss.

- We propose the LAB-CVaR-logit loss which can alleviate the difficulties and instability in optimization caused by the loss re-weighting, and study its mechanism of how to reduce the influence of the loss weight.

- Extensive experiments on publicly available real-world datasets show that our approaches outperform the state-of-the-art algorithms significantly.

2   Related work

As aforementioned, we divide the extant long-tailed learning methods into four categories, and mainly focus on the loss modification category in this paper. Thus, we briefly introduce some representative loss modification methods here.

Re-weighting. Re-weighting methods try to balance the class distribution via assigning weights to the training loss. The vanilla re-weighting method is to weight classes proportionally to the inverse sample size of each class [9]. Instead of the inverse sample size, [5] defines the effective number as the expected volume of samples, and proposes to use the inverse effective number to do re-weighting. [35] proves that there exist performance trade-offs in class-weighted methods, and proposes robust approaches based on CVaR. All the class-level re-weighting methods do not discriminatingly treat the samples in the same class. Besides those class-level re-weighting methods, there also exists some sample-level re-weighting methods, such as Focal loss [17] which down-weights the well-classified samples, and GHM [15] which uses gradients to check whether the sample is well-classified. But the sample-level re-weighting methods for long-tailed learning usually lack theoretical foundation [25]. In contrast, our proposed LAB-CVaR is a special re-weighting method in which the loss weight is variable during the adversarial training, and not only considers the differences between the samples in the same class, but also has the theoretical guarantee in the design of weight bounds.

Logit adjustment. [16, 21] propose to class-wisely shift the margin in the hinge loss. [1] points out that re-weighting methods may cause difficulties and instability in optimization, and proposes the label-distribution-aware margin loss in the softmax cross-entropy. [23] proposes a unified framework for re-weighting methods and post-hoc methods with the help of logit adjustment. Different from them, our LAB-CVaR-logit extends the logit adjustment technique to make the loss consistent with the re-weighted loss, and studies the mechanism why the logit adjustment can improve the re-weighting methods in optimization.

3 Preliminaries

3.1 Classification with long-tailed learning

In a multi-class classification problem, the input space is denoted as $\mathcal{X} \in \mathbb{R}^d$, and the label space is denoted as $\mathcal{Y} = [L] \doteq \{1, \ldots, L\}$. The n labeled training data $\{(x_i, y_i)\}_{i=1}^n$ are sampled from $\mathcal{X} \times \mathcal{Y}$ with the joint distribution $P_{x,y}$. A prediction model $f : \mathbb{R}^d \to \mathbb{R}^L$ usually are trained over $P_{x,y}$ by minimizing the *empirical risk* which is defined as

$$R^\ell(f) = \frac{1}{n} \sum_{i=1}^n \ell(f(x_i), y_i), \tag{1}$$

where $\ell(\cdot)$ is a loss function, such as the Cross-Entropy loss and Hinge loss. For classification tasks, we usually care about the misclassification error of the model, which is equal to optimize the zero-one loss:

$$R^{\ell_{0/1}}(f) = \frac{1}{n} \sum_{i=1}^n \ell_{0/1}(f(x_i), y_i) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}, \tag{2}$$

where $\mathbb{1}$ is the indicator function. In Eq. (2), the misclassification error is calculated over the overall data distribution. However, when the label distribution P_y is highly skewed, there exists the *long-tailed learning* or *class imbalance* problem. Then, the overall error in Eq. (2) is not a suitable measure of model performance in such case, because a trivial model which classifies every sample to the majority label will attain a low misclassification error [23]. To cope with this, we should separately consider the class-wise misclassification error to improve the model performance over the minority class. The class-wise misclassification error for the class j is defined as

$$R_j^{\ell_{0/1}}(f) = \frac{1}{n_j} \sum_{y_i = j} \ell_{0/1}(f(x_i), y_i), \tag{3}$$

where n_j is the sample size of the class j, and it is obvious that $n = \sum_{j=1}^L n_j$, $R^{\ell_{0/1}} = \sum_{j=1}^L \frac{n_j}{n} R_j^{\ell_{0/1}}$. Without loss of generality, we assume $n_1 \leq n_2 \leq \cdots \leq n_L$. Further more, the balanced error rate (BER), which averages each of the per-class error rates, namely $\frac{1}{L} \sum_{j=1}^L R_j^{\ell_{0/1}}(f)$, is widely used to evaluate the model performance [2, 1, 23].

3.2 DRO and αCVaR

Eq. (1) is also known as the empirical risk minimization (ERM) loss which implicitly assumes that the training and test data have the same distribution. This assumption usually does not hold in real-world scenarios in which there exist unknown distribution shifts between the training and test data, especially in the long-tailed learning problem, the label distribution is highly skewed in the training data, but is balanced even inversely skewed in the test data. The distributionally robust optimization (DRO) is a useful technique to solve the unknown distribution shift problem [27] whose effectiveness has been verified in long-tailed learning [26, 25]. DRO minimizes the worst-case distribution Q within a ball w.r.t. divergence D around the training distribution P. The DRO loss is defined as [36]

$$\text{DRO}^\ell(f) = \sup_{Q \ll P} \{\mathbb{E}_Q[\ell(f(x), y)], D(Q||P) \leq \rho\}, \tag{4}$$

where $Q \ll P$ means Q is absolute continuous to P, i.e. $P(x, y) = 0 \Rightarrow Q(x, y) = 0$, and ρ is the constraint radius of the distribution ball, which controls the degree of distribution shift.

Different divergence functions (e.g. Kullback-Leibler divergence, Wasserstein metric) derive different DRO risks. In this work, we focus on one of the most popular DRO risks, the conditional value-at-risk at level α (α-CVaR), which is derived from the Rényi divergence by $D_\epsilon(P||Q) = \frac{1}{\epsilon-1} \log \int (\frac{dP}{dQ})^\epsilon dQ$ where $\epsilon \to \infty$ and $\rho = -\log \alpha$ [36]. The α-CVaR loss is defined as

$$\alpha\text{-CVaR}^\ell(f) = \max_{\boldsymbol{w} \in \Delta_n, 0 \leq w_i \leq \frac{1}{\alpha n}} \sum_{i=1}^n w_i \ell(f(x_i), y_i), \tag{5}$$

where $\Delta_n = \{(w_1, \cdots, w_n) : \sum_{i=1}^n w_i = 1\} \in \mathbb{R}^n$ and $\alpha \in (0, 1)$. Because of the constraints of $\sum_{i=1}^n w_i = 1$ and $0 \leq w_i \leq \frac{1}{\alpha n}$, α-CVaR is always trained on the worst α fraction of the training data.

4 The Proposed Approach

4.1 Label-aware bounded CVaR

In the long-tailed classification scenario, the minority classes tend to have worse classification accuracy than the majority classes [1, 23]. Then DRO and α-CVaR which pay more attention to the worst samples are naturally expected to be helpful to improve the model performance over the minority classes. However, [8] theoretically points out a surprising phenomenon that DRO has the same minimizer as ERM in the classification scenario in which the model is evaluated by zero-one loss at test time. [8] thinks the reason for this phenomenon is that the range of the test distributions to which the DRO tries to be robust is too wide, then the group DRO [8, 24, 29] is proposed, which reduces the freedom degree of the distribution shift by forcing the samples in the same group to have the same loss weight. The drawback of group DRO is that it cannot discriminatingly treat the samples in the same group. Meanwhile, [36] also finds that α-CVaR has the same pessimistic phenomenon, and thinks the reason is that the classifiers are deterministic, then proposes the Boosted CVaR Classification framework to learn randomized classifiers to circumvent the problem. The drawback of Boosted CVaR Classification is the low efficiency in its boosting procedure. On the other hand, both the group DRO and Boosted CVaR Classification do not specially design their algorithms for long-tailed learning.

In this paper, we find that when the bound of the loss weight w_i in CVaR is varied across the labels, rather than all the w_i share the same bound $[0, \frac{1}{\alpha n}]$, the minimizers of CVaR and ERM are no longer the same. So beyond the group DRO and Boosted CVaR Classification, we propose a novel approach to overcome the pessimistic result of α-CVaR in the classification task, and the approach is named as *label-aware bounded CVaR* (LAB-CVaR) which means that the weight bounds are label-wisely designed. All the detailed proofs of the theory results in this section are shown in Appendix A.

Definition 4.1. *Define the loss of LAB-CVaR as:*

$$\textit{LAB-CVaR}^\ell(f) = \max_{\boldsymbol{w} \in \Delta_n, \frac{1}{\beta_j n} \leq w_i \leq \frac{1}{\alpha_j n}} \sum_{i=1}^n w_i \ell(f(x_i), y_i), \tag{6}$$

where $\alpha_j \in \mathbb{R}^+$ and $\beta_j \in \mathbb{R}^+$ are label-wise, and $y_i = j$.

Proposition 4.2. *For zero-one loss, the relationship holds: $\textit{LAB-CVaR}^{\ell_{0/1}}(f) = \min\{1 - \sum_{j=1}^L \frac{n_j}{\beta_j n} + \sum_{j=1}^L \frac{n_j}{\beta_j n} R_j^{\ell_{0/1}}(f), \sum_{j=1}^L \frac{n_j}{\alpha_j n} R_j^{\ell_{0/1}}(f)\}$.*

From the definition of LAB-CVaR in Eq. (6), we can see that LAB-CVaR degrades to α-CVaR when $\alpha_j = \alpha$ and $\frac{1}{\beta_j n} \to 0$. Then according to Proposition 4.2, we get the same result with [36] that α-CVaR$^{\ell_{0/1}}(f) = \min\{1, \frac{1}{\alpha} R^{\ell_{0/1}}(f)\}$, which reveals that the minimizer of ERM is also the minimizer of α-CVaR$^{\ell_{0/1}}(f)$, and means α-CVaR has no better performance than ERM in classification task. On the other hand, when the α_j and β_j are label-aware, the minimizer of LAB-CVaR$^{\ell_{0/1}}(f)$ is equal to the minimizer of one of the two class-weighted loss, $\sum_{j=1}^L \frac{n_j}{\beta_j n} R_j^{\ell_{0/1}}(f)$ and $\sum_{j=1}^L \frac{n_j}{\alpha_j n} R_j^{\ell_{0/1}}(f)$. Thus LAB-CVaR has the different minimizer with that of ERM, and could be expected to have better performance than ERM when α_j and β_j are properly designed.

Besides the label-aware bound design in LAB-CVaR which helps to overcome the pessimistic result of α-CVaR, the another import design in LAB-CVaR is the non-zero lower bound $\frac{1}{\beta_j n}$. Because of the non-zero lower bound, some samples are re-weighted by the upper bound and the others are re-weighted by the lower bound, thus LAB-CVaR takes advantage of all the samples and at the same time discriminatingly treats the samples in the same class. In the following, we theoretically design the optimal upper and lower bounds aiming to minimize the LAB-CVaR loss.

Proposition 4.3. *With probability at least $1 - \delta$, LAB-CVaR loss is bounded by*

$$\sum_{j=1}^L \frac{1}{\beta_j} \sqrt{n_j} \frac{4L * C(\mathcal{F})}{n} \lesssim \mathbb{E}_{P_{x,y}} \textit{LAB-CVaR}^{\ell_{0/1}}(f) \lesssim \sum_{j=1}^L \frac{1}{\alpha_j} \sqrt{n_j} \frac{4L * C(\mathcal{F})}{n}, \tag{7}$$

where we use $\lesssim$ to hide some constant factors and low order terms, $\mathcal{F}$ is the family of hypothesis class, and $C(\mathcal{F})$ is some proper complexity measure of $\mathcal{F}$.

4

Assume the weighted sum of α_j is constrained by the form that $\sum_{j=1}^{L} n_j^{k_1} \alpha_j = \tau_1$, where $\tau_1 \in \mathbb{R}^+$, k_1 is a variable exponent, then in order to minimize the upper bound of LAB-CVaR in Proposition 4.3, α_j should be proportional to $n_j^{1/4 - k_1/2}$, because according to Cauchy–Schwarz inequality, $\sum_{j=1}^{L} \frac{1}{\alpha_j} \sqrt{n_j} = \frac{1}{\tau_1} (\sum_{j=1}^{L} n_j^{k_1} \alpha_j)(\sum_{j=1}^{L} \frac{1}{\alpha_j} \sqrt{n_j}) \geq \frac{1}{\tau_1} (\sum_{j=1}^{L} n_j^{1/4 + k_1/2})^2$, if and only if $n_j^{k_1} \alpha_j \propto \frac{1}{\alpha_j} \sqrt{n_j}$. Similarly, assume $\sum_{j=1}^{L} n_j^{k_2} \beta_j = \tau_2$, we have $\beta_j \propto n_j^{1/4 - k_2/2}$ to minimize the lower bound of LAB-CVaR. Then the optimal α_j and β_j are

$$\alpha_j^* = \tau_1 \frac{n_j^{1/4 - k_1/2}}{\sum_{j=1}^{L} n_j^{1/4 + k_1/2}}, \quad \beta_j^* = \tau_2 \frac{n_j^{1/4 - k_2/2}}{\sum_{j=1}^{L} n_j^{1/4 + k_2/2}}, \tag{8}$$

where $k_1 < \frac{1}{2}$ and $k_2 < \frac{1}{2}$ ensure the classes with smaller sample size have larger weight bounds, and $\alpha_j^* < \beta_j^*$ ensures the upper bound is larger than the lower bound. Then the corresponding LAB-CVaR with the optimal bounds is

$$\text{LAB-CVaR}^{\ell*}(f) = \max_{\boldsymbol{w} \in \Delta_n, \frac{1}{\beta_j^* n} \leq w_i \leq \frac{1}{\alpha_j^* n}} \sum_{i=1}^{n} w_i \ell(f(x_i), y_i). \tag{9}$$

Eq. (9) is a linear programming problem which can be easily solved by the simplex algorithm [6]. Further, in our experiment section, we set $k_1 = k_2$. On one hand, this can reduce the hyper-parameter number. On the other hand, it will be easy to ensure $\alpha_j < \beta_j$ just by a simple way, setting $\tau_1 = \eta \tau_2$ with $0 < \eta < 1$.

4.2 Label-aware bounded CVaR with logit adjustment

In practice, the zero-one loss is non-convex and non-continuous, so we usually use a surrogate loss instead of the zero-one loss at training time for the optimization tractability, such as the softmax cross-entropy loss for multiclass classification. A softmax cross-entropy loss over the sample (x_i, y_i) could be written as

$$\ell_{sc}(f(x_i), y_i) = -\log\left[\frac{\exp(f_j(x_i))}{\sum_{j' \in [L]} \exp(f_{j'}(x_i))}\right] = \log[1 + \sum_{j' \neq j} \exp(f_{j'}(x_i) - f_j(x_i))], \tag{10}$$

where $y_i = j$, $f_j(x_i)$ is the logit of the class j, namely *logit* refers to the input of the softmax function.

From Eq. (9), we can see that LAB-CVaR is a special re-weighting method in which the loss weight is variable during the adversarial training. Without loss of generality, we assume $n_1 \leq n_2 \leq \cdots \leq n_L$. Then in our LAB-CVaR, the largest weight ratio between two samples is $\frac{\beta_L^*}{\alpha_1^*}$ which is very large when the labels are extremely imbalanced. Some recent work [1, 5] find that the large weight ratio can cause difficulties and instability in optimization, which is also observed in our experiments (see Figure 2). In this section, we try to reduce the influence of the loss weight in optimization by combining the logit adjustment technique [23] with our LAB-CVaR.

Proposition 4.4. *Define the weighted softmax cross-entropy loss with the logit adjustment as:*

$$R_{logit}^{\ell_{sc}}(f, \pi) = \sum_{i=1}^{n} \pi_j w_i \log[1 + \sum_{j' \neq j} \exp(f_{j'}(x_i) + \log \pi_{j'} - f_j(x_i) - \log \pi_j)]. \tag{11}$$

Then for any $\pi \in \mathbb{R}_+^L$, the loss $R_{logit}^{\ell_{sc}}(f, \pi)$ is consistent with the original weighted softmax loss $\sum_{i=1}^{n} w_i \ell_{sc}(f(x_i), y_i)$, namely these two losses have the same optimal minimizer.

Let $s_1 = (x_{i_1}, y_{i_1})$ and $s_2 = (x_{i_2}, y_{i_2})$ be two samples in the most minority class *1* and the most majority class L respectively. And let the logit $f_j(x) = W_j^T \Phi(x)$, where W_j is the classifier for class j, $\Phi(x)$ is the feature extractor. In our opinion, the large weight of s_1 makes the sample have more influence on the classifier W_t ($1 < t < L$) to push itself away from the class t, however, this may make the class t closer with s_2. From Proposition 4.4, we can see that the influence of the loss weight can be transferred to the logit term, then an intuitive solution to reduce the influence of the loss weight is to make the more minority classes have smaller π_j, which can smooth the $\pi_j w_i$.

5

The influence of the loss weight of (x_i, y_i) can be observed by the gradient over the classifier W_t. Let $\ell_{log}(f(x_i), y_i) = \pi_j w_i \log[1 + \sum_{j' \neq j} \exp(f_{j'}(x_i) + \log \pi_{j'} - f_j(x_i) - \log \pi_j)]$, then the norm of the gradient of $\ell_{log}(f(x_i), y_i)$ over the W_t with $\pi = [\pi_1, ..., \pi_n]$ is

$$g(x_i, y_i, \pi, t) = \|\frac{\partial \ell_{log}(f(x_i), y_i)}{\partial W_t}\| = \pi_j w_i \frac{\pi_t \exp(f_t(x_i))}{\sum_{j' \in [L]} \pi_{j'} \exp(f_{j'}(x_i))} \|\Phi(x_i)\|, \ y_i \neq t. \quad (12)$$

Through a simple proof, we have that, if the more minority class has the smaller π_j, this intuitive solution indeed reduces the influence of the loss weight. The simple proof is that, when $\pi_1 \leq \cdots \leq \pi_L$, $\frac{g(x_{i_1}, y_{i_1}, \pi, t)}{g(x_{i_2}, y_{i_2}, \pi, t)} = \frac{\pi_1 w_{i_1}}{\pi_L w_{i_2}} \frac{\pi_t \exp(f_t(x_{i_1}))}{\sum_{j' \in [L]} \pi_{j'} \exp(f_{j'}(x_{i_1}))} \frac{\sum_{j' \in [L]} \pi_{j'} \exp(f_{j'}(x_{i_2}))}{\pi_t \exp(f_t(x_{i_2}))} \frac{\|\Phi(x_{i_1})\|}{\|\Phi(x_{i_2})\|} \leq \frac{\pi_1 w_{i_1}}{\pi_L w_{i_2}} \frac{\pi_t \exp(f_t(x_{i_1}))}{\pi_1 \sum_{j' \in [L]} \exp(f_{j'}(x_{i_1}))} \frac{\pi_L \sum_{j' \in [L]} \exp(f_{j'}(x_{i_2}))}{\pi_t \exp(f_t(x_{i_2}))} \frac{\|\Phi(x_{i_1})\|}{\|\Phi(x_{i_2})\|} = \frac{g(x_{i_1}, y_{i_1}, 1, t)}{g(x_{i_2}, y_{i_2}, 1, t)}$, which indicates that the loss of (x_{i_1}, y_{i_1}) has relatively smaller gradient after being logit adjusted with π.

In this paper, we adopt the policy that $\pi_j = n\alpha_j^*$ which is the reciprocal of the upper weight bound of each class. Then we have our second approach named *label-aware bounded CVaR with logit adjustment* (LAB-CVaR-logit):

$$\text{LAB-CVaR-logit}^{\ell_{sc}}(f) = \sum_{i=1}^{n} n\alpha_j^* w_i \log[1 + \sum_{j' \neq j} \exp(f_{j'}(x_i) - f_j(x_i) + \log \frac{\alpha_{j'}^*}{\alpha_j^*})], \quad (13)$$

where w_i is calculated in Eq. (9), namely we firstly get w_i by LAB-CVaR, and secondly use LAB-CVaR-logit to train the model.

4.3 Comparison to existing work

LHCVaR [35]. LHCVaR is also a label-aware bounded method, our LAB-CVaR is different from it at the following aspects: 1) There is no lower weight bound design in LHCVaR. Furthermore, the upper weight bound designed in LHCVaR is intuitive and lacks theoretical guarantee, whereas we get the optimal α_j^* by the theoretical analysis. 2) In our work, α_j is just the bound of w_i in class j, and w_i is still varied across samples even in the same class. While in LHCVaR, the samples in the same class must have the same w_i, which will cause the loss of the ability to discriminatingly treat the samples in the same class.

Logit adjustment [23]. [23] studies how the logit adjustment technique can make the loss consistent with the balanced loss, whereas we extend the technique to make the loss consistent with the re-weighted loss. The logit adjustment in [23] is a special case of our LAB-CVaR-logit, because when $k_1 = k_2 = -\frac{3}{2}$ and $\tau_1 = \tau_2$, then $w_i \propto n_j^{-1}$, and LAB-CVaR degrades to the vanilla class-weighted method. Moreover, there is no study in [23] for the mechanism that the logit adjustment technique can improve the re-weighted methods.

5 Experiments

We evaluate the performance of the proposed LAB-CVaR and LAB-CVaR-logit on the imbalanced version of the real-world natural language processing dataset (imbalanced version of IMDB Review [20]) and computer vision datasets (the imbalanced version of CIFAR10, CIFAR100 [13] and Tiny ImageNet [28]).

Baselines. We compare our methods with baselines in the following 4 categories: 1) The classic Empirical Risk Minimization (ERM), which trains network by unweighted cross entropy loss. 2) Three loss function re-weighting methods: the vanilla re-weighting (Vanilla RW), the Class-Balanced Softmax Cross Entropy Loss using Effective Number (CB RW) [5], and the Focal (Focal RW) Loss [17]. The performance gaps among these re-weighting methods are attributed to the difference in their weight designs. Vanilla RW and CB RW use a constant and unified weight for each class during the training. While, Focal RW uses a loss-dependent weight for each sample (see Table 3 in Appendix B). For all of these methods, we do a rescaling of the weighted loss function to make the weight 1 on average in every minibatch. 3) Two logit adjustment methods: Label-Distribution-Aware Margin (LDAM) Loss [1] and Logit-Adjusted (LA) Loss [23]. These methods shift the logits of the samples in each class by a constant before feeding the logits to the softmax function and then the

Table 1: The error rate (ER%) of each class and the balanced error rate (BER%) on imbalanced IMDB Review. Our method LAB-CVaR-logit achieves the best performance.

	Positive Class ER%	Negative Class ER%	BER%
ERM	67.03±20.49	**3.2**±1.22	35.12±9.78
Vanilla RW	24.46±5.13	19.96±5.30	22.21±4.87
CB RW [5]	23.49±2.52	17.56±3.45	20.52±0.84
Focal RW [17]	48.49±6.71	5.41±1.35	26.95±2.68
LDAM [1]	50.31±1.01	3.83±0.87	27.09±0.33
LDAM+DRW [1]	46.41±3.24	4.44±1.19	25.42±1.06
LA [23]	26.20±9.68	18.06±1.23	22.13±4.94
α-CVaR	67.19±25.40	3.79±2.33	35.49±11.63
LHCVaR [35]	33.56±7.77	23.06±4.27	28.31±2.98
LAB-CVaR	23.96±2.13	16.92±1.98	20.44±0.92
LAB-CVaR-logit	**22.63**±0.75	17.42±0.35	**20.02**±0.38

cross entropy loss. The difference between these two methods lies in their designs of the shifting constant (see Table 3 in Appendix B). As shown in their paper, LDAM can significantly benefit from the Deferred Re-Weighting (DRW) strategy. So, we add LDAM+DRW as another baseline. 4) The original CVaR loss and one of its variants Label Heterogeneous CVaR (LHCVaR) [35] that dynamically optimizes the class weights while training with only an upper bound of the weights.

Datasets. The IMDB Review is a binary classification dataset. The Cifar10, Cifar100 and Tiny ImageNet are multiclass classification datasets. All four datasets originally have class-balanced training and validation sets. To create the imbalanced version of the IMDB Review, we keep the class-balanced validation set unchanged and randomly remove 90% of positive reviews in the training set. To create the imbalanced version of the Cifar10, Cifar100, and Tiny ImageNet, we again keep the class-balanced validation sets unchanged and follow [1, 23] to downsample the training sets according to exponential decay. Specifically, the j-th class are downsampled to $round(\lambda^{L-j} n_j)$, and λ are designed for each dataset to make the imbalance ratio [2] 100. To be consistent with the name in [23], imbalanced versions of Cifar10, Cifar100, and Tiny ImageNet are named Cifar10-LT, Cifar100-LT, and Tiny ImageNet-LT, respectively.

Experimental Settings. For the imbalanced IMDB Review, we use the two-layer bidirectional LSTM and the Adam optimizer. For Cifar10-LT and Cifar100-LT, we use the ResNet32 and the SGD optimizer. For Tiny ImageNet, we use the ResNet18 and the SGD optimizer. All results reported in this section are averaged over five independent runs and are in the format of $average \pm standard\ deviation$. More experimental details can be found in Appendix B.

5.1 Results and analysis

For the imbalanced IMDB Review, the error rate (ER%) of each class on the validation set and the balanced error rate (BER%) on the entire validation set are reported in the Table 1. For the Cifar10-LT and Cifar100-LT, the balanced error rate (BER%) over the entire validation set and the error rate of the worst class (WER%) in the validation dataset are reported in the Table 2. For the Tiny ImageNet-LT, because the WERs of all algorithms are 100.00%, besides the balanced error rate (BER%) over the entire validation set, the error rate of the worst 20 classes (WER20%) in the validation dataset are reported in the Table 2.

From Table 1 and Table 2, we can generally see that, under the evaluation of BER and WER, our proposed LAB-CVaR-logit consistently outperforms all the state-of-the-art baselines. On the other hand, our LAB-CVaR does not perform as well as LDAM+DRW and LA in multiclass classification datasets, this is because LAB-CVaR is a re-weighting method which has the optimization problem as aforementioned in Section 4.2, but LAB-CVaR still performs better or comparably to all of the re-weighting baselines. Overall, those results verify the effectiveness of our proposed methods.

---

[2]The ratio between the sample size of the last class (the most frequent class) and that of the first class (the least frequent class).

Table 2: Performance on CIFAR10-LT, CIFAR100-LT and Tiny ImageNet-LT measured by the balanced error rate (BER%) and the error rate of the worst class (WER%). Because all the methods have 100.00% WER on the Tiny ImageNet-LT, the error rate of the worst 20 classes (WER20%) is further evaluated. Our method LAB-CVaR-logit achieves the best performance.

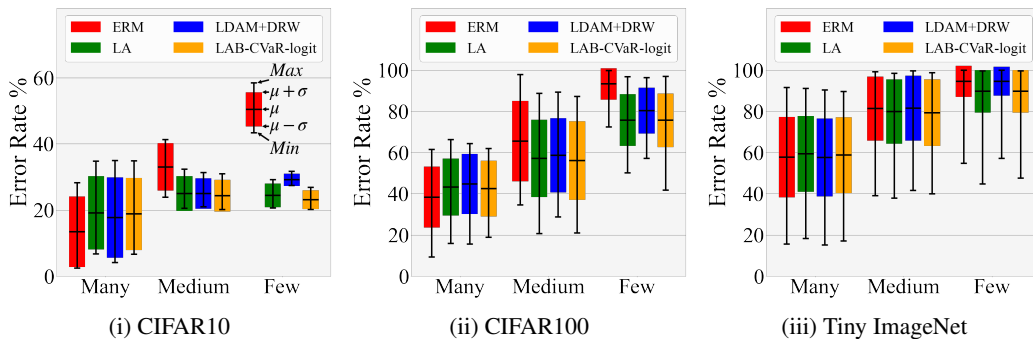|	CIFAR10-LT		CIFAR100-LT		Tiny ImageNet-LT	
	BER%	WER%	BER%	WER%	BER%	WER20%
ERM	30.45±0.70	51.78±2.67	64.42±0.45	100.00±0.00	77.11±0.38	99.80±0.15
Vanilla RW	32.59±1.36	50.42±3.88	73.12±1.62	99.00±0.00	81.17±0.34	99.10±0.08
CB RW [5]	29.07±1.26	40.86±2.33	64.35±0.80	99.00±0.00	77.01±0.27	99.90±0.05
Focal RW [17]	30.28±0.43	53.02±2.43	62.21±0.69	99.80±0.40	76.80±0.41	99.90±0.16
LDAM [1]	29.68±0.20	49.82±3.61	64.16±0.50	100.00±0.00	77.08±0.30	99.80±0.26
LDAM+DRW [1]	23.38±0.41	33.38±1.29	60.38±0.39	97.40±1.36	76.58±0.23	98.50±0.35
LA [23]	22.53±0.37	33.58±1.69	57.91±1.04	95.60±1.02	75.69±0.12	97.90±0.27
α-CVaR	30.41±0.34	53.30±2.85	65.80±0.73	100.00±0.00	76.96±0.28	100.00±0.00
LHCVaR [35]	28.06±0.55	45.14±1.55	63.32±0.44	99.80±0.40	77.15±0.20	100.00±0.00
LAB-CVaR	26.93±0.27	40.24±1.84	62.96±0.28	99.20±0.40	75.87±0.04	99.70±0.12
LAB-CVaR-logit	**21.82**±0.22	**32.98**±0.78	**57.29**±0.91	**95.40**±0.49	**75.16**±0.23	**97.52**±0.09

Figure 1: Box plots of the group-wise error rate on imbalanced CIFAR10, CIFAR100 and Tiny ImageNet. Our method LAB-CVaR-logit achieves the best performance when the sample size of the class is Medium and Few. For an explanation of the box plot structure, please refer to the box plot of the ERM method on the CIFAR10 dataset corresponding to the group Few (in subfigure 1i), where μ and σ refer to the mean and standard deviation, respectively.

On the binary classification dataset IMDB Review, the methods (e.g. ERM) overfitting on the negative class tend to have higher and unstable ER on the positive class. Our methods, though with higher negative class ERs, achieve the best positive class ERs and average BERs with small variances.

Ablation study. The results of LAB-CVaR-logit, LAB-CVaR, and α-CVaR in the Table 1 and Table 2 constitute an ablation study. The effectiveness of our label-aware bound design is proved by the phenomenon that LAB-CVaR outperforms α-CVaR, systematically. The effectiveness of logit adjustment in our method is proved by the phenomenon that LAB-CVaR-logit outperforms LAB-CVaR, systematically. Furthermore, the results of the α-CVaR in the Table 1 and Table 2 is not better than that of the ERM, which is consistent with the insight in our Proposition 4.2.

Group-wise error rate. To investigate where the performance gain of our method compared to baselines comes from. We follow the protocol used in [23, 11, 19] to group the class into three subsets according to the sample size of the class. Note that the last class has the largest sample size n_L in our experiments. Accordingly, we assign the j-th class to the group "Many" if $n_j \geq 0.2n_L$, the group "Medium" if $0.04n_L \leq n_j < 0.2n_L$, the group "Few" otherwise. As shown in Figure 1, for all three datasets, ERM has the lowest error rate in the "Many" group, and our method has the lowest error rate in the "Medium" and "Few" groups.
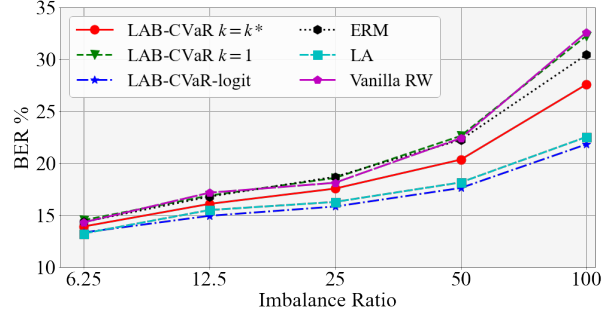
Figure 2: The balanced error rate (BER) of some re-weighting methods and logit adjustment methods on imbalanced CIFAR10 for different imbalance ratio. Base on the results, we analyze the optimization problem of re-weighting methods. Here, $k \triangleq 1/4 - k_1/2 = 1/4 - k_2/2$ is the exponent of the numerator in Eq. (8).

Performance under different imbalance ratios. To analyze how the performances of the methods change under different imbalance levels, we synthesize imbalanced CIFAR10 datasets with different imbalance ratios for further comparison. As shown in Figure 2, our LAB-CVaR-logit constantly outperforms other methods under all imbalance ratios, and the performance gap between our method and the ERM increases when the imbalance ratio increases. More importantly, the re-weighting methods (Vanilla RW and LAB-CVaR $k = 1$) are worse than ERM when the imbalance ratio is high, this phenomenon is consistent with the analysis in Section 4.2, which verifies that the re-weighting method has optimization problems in the highly imbalanced scenarios. Meanwhile, the corresponding logit adjustment methods (LA and LAB-CVaR-logit) can alleviate the optimization problems of the re-weighting methods. We can also find that the LAB-CVaR $k = k^*$ is better than ERM under all the imbalance ratios, this is because k^* is always less than 1, which could reduce the loss weight ratio. However, LAB-CVaR $k = k^*$ is still worse than the logit adjustment methods.

6 Conclusion

We propose two approaches to address the long-tailed learning problem, label-aware bounded CVaR (LAB-CVaR), and LAB-CVaR with logit adjustment (LAB-CVaR-logit). These two approaches significantly outperform the state-of-the-art baselines on a variety of imbalanced real-world datasets. Specifically, the LAB-CVaR theoretically designs the optimal weight bounds to minimize its loss bound, and LAB-CVaR-logit adopts the logit adjustments to alleviate the optimization problem caused by re-weighting. The effectiveness of our label-aware bounded CVaR also inspires the study of the sample-aware bounded CVaR to make the best of the differences among the samples in future work.

References

[1] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *arXiv preprint arXiv:1906.07413*, 2019.

[2] Philip K Chan and Salvatore J Stolfo. Learning with non-uniform class and cost distributions: Effects and a distributed multi-classifier approach. In *In Workshop Notes KDD-98 Workshop on Distributed Data Mining*. Citeseer, 1998.

[3] Nitesh V Chawla. Data mining for imbalanced datasets: An overview. *Data mining and knowledge discovery handbook*, pages 875–886, 2009.

[4] Guillem Collell, Drazen Prelec, and Kaustubh Patil. Reviving threshold-moving: a simple plug-in bagging ensemble for binary and multiclass imbalanced data. *arXiv preprint arXiv:1606.08698*, 2016.

[5] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019.

[6] George B Dantzig. Origins of the simplex method. In *A history of scientific computing*, pages 141–151. 1990.

[7] Tom Fawcett and Foster J Provost. Combining data mining and machine learning for effective user profiling. In *KDD*, volume 96, pages 8–13, 1996.

[8] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. Does distributionally robust supervised learning give robust classifiers? In *International Conference on Machine Learning*, pages 2029–2037. PMLR, 2018.

[9] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5375–5384, 2016.

[10] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019.

[11] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *ArXiv*, abs/1910.09217, 2020.

[12] Byungju Kim and Junmo Kim. Adjusting decision boundary for class imbalanced learning. *IEEE Access*, 8:81674–81685, 2020.

[13] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.

[14] Miroslav Kubat, Stan Matwin, et al. Addressing the curse of imbalanced training sets: one-sided selection. In *Icml*, volume 97, page 179. Citeseer, 1997.

[15] Buyu Li, Yu Liu, and Xiaogang Wang. Gradient harmonized single-stage detector. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8577–8584, 2019.

[16] Yaoyong Li, Hugo Zaragoza, Ralf Herbrich, John Shawe-Taylor, and Jaz Kandola. The perceptron algorithm with uneven margins. In *ICML*, volume 2, pages 379–386, 2002.

[17] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.

[18] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019.

[19] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2532–2541, 2019.

[20] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, A. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *ACL*, 2011.

[21] Hamed Masnadi-Shirazi and Nuno Vasconcelos. Risk minimization, probability elicitation, and cost-sensitive svms. In *ICML*, 2010.

[22] Shahar Mendelson. Lower bounds for the empirical minimization algorithm. *IEEE Transactions on Information Theory*, 54(8):3797–3803, 2008.

[23] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*, 2020.

[24] Yonatan Oren, Shiori Sagawa, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust language modeling. *arXiv preprint arXiv:1909.02060*, 2019.

[25] Qi Qi, Yi Xu, Rong Jin, Wotao Yin, and Tianbao Yang. Attentional biased stochastic gradient for imbalanced classification. *arXiv preprint arXiv:2012.06951*, 2020.

[26] Qi Qi, Yan Yan, Zixuan Wu, Xiaoyu Wang, and Tianbao Yang. A simple and effective framework for pairwise deep metric learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16*, pages 375–391. Springer, 2020.

[27] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.

[28] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115:211–252, 2015.

[29] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019.

[30] Dvir Samuel and Gal Chechik. Distributional robustness loss for long-tail learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9495–9504, 2021.

[31] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *Advances in Neural Information Processing Systems*, 33:1513–1524, 2020.

[32] Grant Van Horn and Pietro Perona. The devil is in the tails: Fine-grained classification in the wild. *arXiv preprint arXiv:1709.01450*, 2017.

[33] Byron C Wallace, Kevin Small, Carla E Brodley, and Thomas A Trikalinos. Class imbalance, redux. In *2011 IEEE 11th international conference on data mining*, pages 754–763. Ieee, 2011.

[34] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. *Advances in Neural Information Processing Systems*, 30, 2017.

[35] Ziyu Xu, Chen Dan, Justin Khim, and Pradeep Ravikumar. Class-weighted classification: Trade-offs and robust approaches. In *International Conference on Machine Learning*, pages 10544–10554. PMLR, 2020.

[36] Runtian Zhai, Chen Dan, Arun Suggala, J Zico Kolter, and Pradeep Ravikumar. Boosted cvar classification. *Advances in Neural Information Processing Systems*, 34, 2021.

[37] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *arXiv preprint arXiv:2110.04596*, 2021.

A Appendix: Theoretical Analyses

A.1 Proof of Proposition 4.2

Base on the definition of LAB-CVaR$^{\ell_{0/1}}$, we have

$$
\text{LAB-CVaR}^{\ell_{0/1}}(f) = \max_{\boldsymbol{w} \in \Delta_n, \frac{1}{\beta_j n} \leq w_i \leq \frac{1}{\alpha_j n}} \sum_{i=1}^{n} w_i \mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}
$$

$$
= \sum_{i=1}^{n} \frac{1}{\beta_j n} \mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\} \tag{14}
$$

$$
+ \max_{\boldsymbol{w} \in \Delta'_n, 0 \leq w_i \leq (\frac{1}{\alpha_j n} - \frac{1}{\beta_j n})} \sum_{i=1}^{n} w_i \mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\},
$$

where $\Delta'_n = \{(w_1, \cdots, w_n) : \sum_{i=1}^{n} w_i = 1 - \sum_{j=1}^{L} \frac{n_j}{\beta_j n}\} \in \mathbb{R}^n$.

Because $\max_{\boldsymbol{w} \in \Delta'_n, 0 \leq w_i \leq (\frac{1}{\alpha_j n} - \frac{1}{\beta_j n})} w_i \mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\} = (\frac{1}{\alpha_j n} - \frac{1}{\beta_j n})\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}$ whether $\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\} = (\frac{1}{\alpha_j n} - \frac{1}{\beta_j n})\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}$ is equal to 1 or 0, and note that, LAB-CVaR$^{\ell_{0/1}}(f) = 1 - \sum_{j=1}^{L} \frac{n_j}{\beta_j n}$ when $\sum_{i=1}^{n}(\frac{1}{\alpha_j n} - \frac{1}{\beta_j n})\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\} > 1 - \sum_{j=1}^{L} \frac{n_j}{\beta_j n}$, then we get

$$
\text{LAB-CVaR}^{\ell_{0/1}}(f) = \min\{1 - \sum_{j=1}^{L} \frac{n_j}{\beta_j n}, \sum_{i=1}^{n}(\frac{1}{\alpha_j n} - \frac{1}{\beta_j n})\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}
$$

$$
+ \sum_{i=1}^{n} \frac{1}{\beta_j n}\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}
$$

$$
= \min\{1 - \sum_{j=1}^{L} \frac{n_j}{\beta_j n} + \sum_{i=1}^{n} \frac{1}{\beta_j n}\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}, \tag{15}
$$

$$
\sum_{i=1}^{n} \frac{1}{\alpha_j n}\mathbb{1}\{y_i \neq \max_{y' \in [L]} f_{y'}(x_i)\}
$$

$$
= \min\{1 - \sum_{j=1}^{L} \frac{n_j}{\beta_j n} + \sum_{j=1}^{L} \frac{n_j}{\beta_j n} R_j^{\ell_{0/1}}(f), \sum_{j=1}^{L} \frac{n_j}{\alpha_j n} R_j^{\ell_{0/1}}(f)\}
$$

A.2 Proof of Proposition 4.3

In our label-aware bounded CVaR, $\frac{1}{\beta_j n} \leq w_i \leq \frac{1}{\alpha_j n}$, thus

$$
\sum_{j=1}^{L} \frac{n_j}{\beta_j n} R_j^{\ell_{0/1}}(f) \leq \text{LAB-CVaR}^{\ell_{0/1}}(f) \leq \sum_{j=1}^{L} \frac{n_j}{\alpha_j n} R_j^{\ell_{0/1}}(f). \tag{16}
$$

Inspired by the proof of the Theorem 2 of [1], we prove the loss bound for each class j separately and then union bound over all the classes. From the proof of the Theorem 1 of [35], we can see that, with probability at least $1 - \delta/L$,

$$
\frac{n_j}{n} \mathbb{E}_{P_{x,y}} R_j^{\ell_{0/1}} \leq \frac{n_j}{n} R_j^{\ell_{0/1}} + 4L \mathbb{E}_{P_{x,y}}[\frac{n_j}{n} \mathfrak{R}_j(\Pi_1(\mathcal{F}))] + \sqrt{\frac{\log \frac{L}{\delta}}{2n}}, \tag{17}
$$

where $\mathcal{F}$ is the family of hypothesis class, $\mathfrak{R}_j(\cdot)$ is the empirical Rademacher complexity of class j, $\Pi_1(\mathcal{F}) = \{x \mapsto f_j(x) : j \in [L], f \in \mathcal{F}\}$

From the proof of the Corollary 2 of [35], we have $\mathbb{E}_{P_{x,y}}[\frac{n_j}{n}\mathfrak{R}_j(\mathcal{F})] \leq C(\mathcal{F})\frac{\sqrt{n_j}}{n}$, where $C(\mathcal{F})$ is some proper complexity measure of $\mathcal{F}$, then

$$\frac{n_j}{n}\mathbb{E}_{P_{x,y}}R_j^{\ell_{0/1}} \leq \frac{n_j}{n}R_j^{\ell_{0/1}} + 4L\frac{\sqrt{n_j}}{n}C(\mathcal{F}) + \sqrt{\frac{\log\frac{L}{\delta}}{2n}}. \tag{18}$$

So $\mathbb{E}_{P_{x,y}}[\sum_{j=1}^{L}\frac{n_j}{\alpha_j n}R_j^{\ell_{0/1}}(f)] \leq \sum_{j=1}^{L}\frac{n_j}{\alpha_j n}R_j^{\ell_{0/1}}(f) + \sum_{j=1}^{L}\frac{1}{\alpha_j}\sqrt{n_j}\frac{4L*C(\mathcal{F})}{n} + \sqrt{\frac{\log\frac{L}{\delta}}{2n}}\sum_{j=1}^{L}\frac{1}{\alpha_j}$.
Note that, $\frac{n_j}{n}R_j^{\ell_{0/1}}(f)]$ is an arbitrary empirical estimation of $\frac{n_j}{n}\mathbb{E}_{P_{x,y}}R_j^{\ell_{0/1}}$, then

$$\mathbb{E}_{P_{x,y}}[\sum_{j=1}^{L}\frac{n_j}{\alpha_j n}R_j^{\ell_{0/1}}(f)] \leq \sup_{P_{x,y}}\sum_{j=1}^{L}\frac{n_j}{\alpha_j n}R_j^{\ell_{0/1}}(f) + \sum_{j=1}^{L}\frac{1}{\alpha_j}\sqrt{n_j}\frac{4L*C(\mathcal{F})}{n} + \sqrt{\frac{\log\frac{L}{\delta}}{2n}}\sum_{j=1}^{L}\frac{1}{\alpha_j}. \tag{19}$$

Remove the low order term $\sqrt{\frac{\log\frac{L}{\delta}}{2n}}\sum_{j=1}^{L}\frac{1}{\alpha_j}$ and hide the constant term $\sup_{P_{x,y}}\sum_{j=1}^{L}\frac{n_j}{\alpha_j n}R_j^{\ell_{0/1}}(f)$, we get

$$\mathbb{E}_{P_{x,y}}[\sum_{j=1}^{L}\frac{n_j}{\alpha_j n}R_j^{\ell_{0/1}}(f)] \lesssim \sum_{j=1}^{L}\frac{1}{\alpha_j}\sqrt{n_j}\frac{4L*C(\mathcal{F})}{n}, \tag{20}$$

where we use $\lesssim$ to hide some constant factors and low order terms.

From the Theorem 1.1 of [22], we can have $\frac{n_j}{n}\mathbb{E}_{P_{x,y}}R_j^{\ell_{0/1}} \geq \inf_{f\in\mathcal{F}}\frac{n_j}{n}R_j^{\ell_{0/1}} + \frac{c\sqrt{n_j}}{n}$, where c is a constant depending only on $\mathcal{F}$. Because $\frac{4L*C(\mathcal{F})\sqrt{n_j}}{n} < \frac{4L*C(\mathcal{F})}{\sqrt{n}}$, there must exist a constant $\mu \in \mathbb{R}^+$ which satisfies $\frac{c\sqrt{n_j}}{n} \geq \frac{4L*C(\mathcal{F})\sqrt{n_j}}{n} - \mu$ for all the class j. Then, $\frac{n_j}{n}\mathbb{E}_{P_{x,y}}R_j^{\ell_{0/1}} \geq \inf_{f\in\mathcal{F}}\frac{n_j}{n}R_j^{\ell_{0/1}} + \frac{4L*C(\mathcal{F})\sqrt{n_j}}{n} - \mu$. Similarly with Eq. (19) and Eq. (20), we get

$$\sum_{j=1}^{L}\frac{1}{\beta_j}\sqrt{n_j}\frac{4L*C(\mathcal{F})}{n} \lesssim \mathbb{E}_{P_{x,y}}[\sum_{j=1}^{L}\frac{n_j}{\beta_j n}R_j^{\ell_{0/1}}(f)]. \tag{21}$$

Recall Eq. (16), we get the result in Proposition 4.3.

By minimizing the bounds $\sum_{j=1}^{L}\frac{1}{\alpha_j}\sqrt{n_j}\frac{4L*C(\mathcal{F})}{n}$ and $\sum_{j=1}^{L}\frac{1}{\beta_j}\sqrt{n_j}\frac{4L*C(\mathcal{F})}{n}$, we get the optimal α^* and β^*, and in fact, the coefficient $4L*C(\mathcal{F})$ does not influence the results of α^* and β^*.

A.3 Proof of Proposition 4.4

The proof is inspired by the proof of Theorem 1 of [23].

Denote $\eta_y(x) = P(y|x)$ in the training distribution $P_{x,y}$. Because $P(y|x) \propto \exp(f_y(x))$ could be regarded as estimates of $P_{y|x}$ [23], we have $f_y^*(x) = \log\eta_y(x)$.

The class-weighted loss is

$$\mathbb{E}_{P_{x,y}}[\ell_\pi(f(x), y)] = \sum_{y\in[L]}\pi_y P(y)\mathbb{E}_{P_{x|y}}[\ell_{soft}(f(x), y)]. \tag{22}$$

Denote $\tilde{\eta}_y(x) = \tilde{P}(y|x)$ in the class-weighted training distribution $\tilde{P}_{x,y}$. From Eq. (22), we can see that $\tilde{P}(y) = \pi_y P(y)$ and $\tilde{P}(x|y) = P(x|y)$. Then we have

$$\tilde{\eta}_y(x) = \tilde{P}(y|x) = \frac{\tilde{P}(x|y)\tilde{P}(y)}{\tilde{P}(x)} = \frac{P(x|y)\pi_y P(y)}{\tilde{P}(x)} = \eta_y(x)\pi_y\frac{P(x)}{\tilde{P}(x)}. \tag{23}$$

So the optimal model of the class-weighted loss is $\tilde{f}_y^*(x) = \log\tilde{\eta}_y(x) = \log\eta_y(x)\pi_y\frac{P(x)}{\tilde{P}(x)} = f_y^* + \log\pi_y + \log\frac{P(x)}{\tilde{P}(x)}$. Because $\log\frac{P(x)}{\tilde{P}(x)}$ does not depend on y, we have $\arg\max_{y\in[L]}\tilde{f}_y^*(x) = \arg\max_{y\in[L]}(f_y^*(x) + \log\pi_y)$.

Table 3: The designs of class weight and shifting constant in our baselines.

Reweighting Methods	Weight of the j-th Class	Hyperparameter
Vanilla RW	$\frac{1}{n_j}$	
CB RW	$\frac{1-\gamma}{1-\gamma^{n_j}}$	$\gamma \in (0,1)$
	Weight of the i-th Sample	Hyperparameter
Focal RW	$(1 - \ell(f(x_i), y_i))^\gamma$	$\gamma > 0$
Logit Adjustment Methods	Shifting Constant of the j-th Class	Hyperparameter
LDAM	$\begin{cases} -\frac{C}{n_j^{\frac{1}{4}}}, & j \text{ is the true label} \\ 0, & \text{otherwise} \end{cases}$	$C > 0$
LA	$-\tau \log n_j$	$\tau > 0$

Table 4: Hyperparameters and their spaces for random search.

Algorithm	Hyperparameter	Search Space
CB RW	γ	$\{0.5, 0.7, 0.8, 0.9, 0.99, 0.999, 0.9999\}$
Focal RW	γ	$\{0.2, 0.5, 0.8, 1, 2, 5, 10, 15\}$
LDAM	max_m	$\{0.2, 0.5, 0.8, 1, 2, 5, 10, 15\}$
LDAM+DRW	max_m	$\{0.2, 0.5, 0.8, 1, 2, 5, 10, 15\}$
	DRW Epoch	int(Total Epoch $\times r$) for r in $\{0.6, 0.8\}$
α-CVaR	α	$\{0.2, 0.5, 0.8, 1, 2, 5, 10, 15\}$
LHCVaR	c	$\{0.2, 0.5, 0.8\}$
	k	$\{0.2, 0.5, 0.8, 1, 2, 5, 10, 15\}$
LAB-CVaR	τ_1	$\{1, 2, 5\}$
	η	$\{1/2, 1/3, 1/6, 1/11, 1/16\}$
	$k_1 = k_2$	$\{0.2, 0.5, 0.8, 1, 2, 5\}$
LAB-CVaR-logit	τ_1	$\{1, 2, 5\}$
	η	$\{1/2, 1/3, 1/6, 1/11, 1/16\}$
	$k_1 = k_2$	$\{0.2, 0.5, 0.8, 1, 2, 5\}$

If we add a logit $\log \psi_y$ into the class-weighted loss, the corresponding optimal model $\bar{f}_y^*(x)$ satisfies $\bar{f}_y^*(x) + \log \psi_y = \log \tilde{P}(y|x) = \tilde{f}_y^*(x)$.

Consequently, when $\psi_y = \pi_y$, $\bar{f}_y^*(x)$ is consistent with $f_y^*(x)$.

Note that, $P_{x,y}$ could be any distribution in the above proof, hence we can assume $P_{x,y}$ is the distribution after loss re-weighting by w_i, and get the Proposition 4.4.

B Appendix: Experiments

B.1 The differences in the class weight designs and the shifting constants

The class weight designs of Vanilla RW, CB RW [3] and Focal RW, and the shifting constants of the LDAM and LA are shown in the Table 3.

---

[3]The hyperparameter used in CB RW is denoted by β in [5]. To avoid notation conflicts, we denote it by γ.

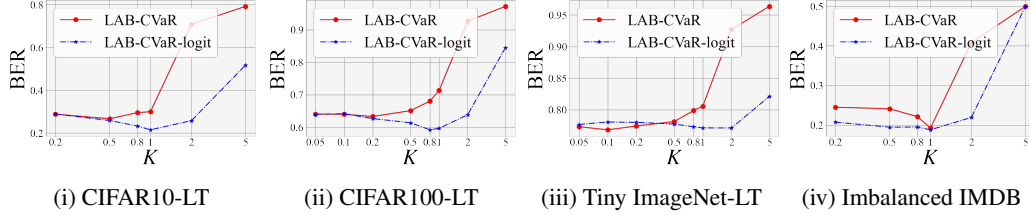(i) CIFAR10-LT   (ii) CIFAR100-LT   (iii) Tiny ImageNet-LT   (iv) Imbalanced IMDB

Figure 3: Performance of LAB-CVaR and LAB-CVaR-logit with different k on the CIFAR10-LT, Cifar100-LT, Tiny ImageNet-LT and imbalanced IMDB Review.

B.2 Implementation Details

IMDB Review. For IMDB Review, we train the two-layer bidirectional LSTM with batch size of 128 for 10 epochs. We use Adam optimizer with $\beta = 0.9$ and $\beta_1 = 0.999$ for training. The initial learning rate is 0.02 and rescaled by 0.1 at the 5-th epoch.

Cifar10 and Cifar100. For Cifar10 and Cifar100, we train the ResNet32 with a batch size of 128 for 200 epochs. We use stochastic gradient descent with momentum of 0.9, weight decay of 2×10^{-4} for training. The initial learning rate is 0.1 and rescaled by 0.01 at the 160-th and 180-th epoch. To be augmented, the training images are horizontally flipped with probability 0.5, padded 4 pixels on each side and randomly cropped to 32×32.

Tiny ImageNet. For Tiny ImageNet, we train the ResNet18 with a batch size of 128 for 120 epochs. We use stochastic gradient descent with momentum of 0.9, weight decay of 2×10^{-4} for training. The initial learning rate is 0.1 and rescaled by 0.1 at the 90-th epoch. To be augmented, the training images are horizontally flipped with probability 0.5, padded 8 pixels on each side and randomly cropped to 64×64.

The hyperparameter search space is shown in Table 4. Based on our experimental observation (see Figure 3), for LAB-CVaR and LAB-CVaR-logit, we would like to suggest choose the exponent k less than 1 and close to 1, respectively.