
A Unified Generalization Analysis of Re-Weighting and Logit-Adjustment for Imbalanced Learning

Zitai Wang^{1,2} Qianqian Xu^{3*} Zhiyong Yang⁴
Yuan He⁵ Xiaochun Cao⁶ Qingming Huang^{4,3,7*}

¹ SKLOIS, Institute of Information Engineering, CAS

² School of Cyber Security, University of Chinese Academy of Sciences

³ Key Lab. of Intelligent Information Processing, Institute of Computing Tech., CAS

⁴ School of Computer Science and Tech., University of Chinese Academy of Sciences

⁵ Alibaba Group

⁶ School of Cyber Science and Tech., Shenzhen Campus of Sun Yat-sen University

⁷ BDKM, University of Chinese Academy of Sciences

wangzitai@iie.ac.cn xuqianqian@ict.ac.cn

yangzhiyong21@ucas.ac.cn heyuan.hy@alibaba-inc.com

caoxiaochun@mail.sysu.edu.cn qmhuang@ucas.ac.cn

Abstract

Real-world datasets are typically imbalanced in the sense that only a few classes have numerous samples, while many classes are associated with only a few samples. As a result, a naïve ERM learning process will be biased towards the majority classes, making it difficult to generalize to the minority classes. To address this issue, one simple but effective approach is to modify the loss function to emphasize the learning on minority classes, such as re-weighting the losses or adjusting the logits via class-dependent terms. However, existing generalization analysis of such losses is still coarse-grained and fragmented, failing to explain some empirical results. To bridge this gap, we propose a novel technique named data-dependent contraction to capture how these modified losses handle different classes. On top of this technique, a fine-grained generalization bound is established for imbalanced learning, which helps reveal the mystery of re-weighting and logit-adjustment in a unified manner. Furthermore, a principled learning algorithm is developed based on the theoretical insights. Finally, the empirical results on benchmark datasets not only validate the theoretical results but also demonstrate the effectiveness of the proposed method.

1 Introduction

In recent years, machine learning has achieved great success with the help of well-collected datasets, where the number of samples is artificially balanced among classes [1, 2]. However, the real-world datasets are generally imbalanced in the sense that only a few classes have numerous samples (*i.e.*, the majority ones), while the others are associated with only a few samples (*i.e.*, the minority ones) [3–5]. Owing to this issue, a naïve Empirical Risk Minimization (ERM) learning process will be biased towards the majority classes, and the generalization on the minority ones becomes challenging. Hence, the imbalanced learning problem has attracted increasing attention in recent years [6–9].

One simple yet effective approach for imbalanced learning is to modify the naïve loss function, such that the learning process can pay more attention to the minority classes (Please refer to Appendix A for more orthogonal approaches). In this direction, existing approaches generally fall into two categories: re-weighting [10, 11] and logit-adjustment [12–16]. The former category assigns larger weights to

*Corresponding authors.

the losses of the minority classes. Although intuitive, this approach might lead to difficulties and instability in optimization [11, 12, 17]. To tackle this issue, Cao et al. [12] propose an effective scheme named Deferred Re-Weighting (DRW), where the re-weighting approach is applied only during the terminal phase of training. The latter category adjusts the logits by class-dependent terms. For example, the Label Distribution Aware Margin (LDAM) loss enforces larger margins for minority classes to achieve strong regularization [12]. The Logit-Adjustment (LA) loss [13] and the Class-Dependent Temperatures (CDT) loss [14] utilize additive and multiplicative terms to adjust the logits, respectively. Most recently, Kini et al. [16] combine the two types of terms and proposes a unified loss named Vector-Scaling (VS) for imbalanced learning.

Although existing loss-modification methods have achieved promising performance, the theoretical insights are still fragmented and coarse-grained. To be specific, Cao et al. [12] and Ren et al. [18] utilize the classic margin theory to explain the necessity of the additive terms in the LDAM loss. However, the theory fails to explain the significant improvement induced by the DRW scheme. Menon et al. [13] analyzes the Fisher consistency property [19] of the additive terms in the LA loss, while providing no further generalization analysis. Kini et al. [16] provides a generalization analysis of the VS loss, but the results can only explain the role of the multiplicative terms under the assumption that a linear model is trained on linearly separable data. Besides, we find that the VS loss is rather incompatible with the DRW scheme, which is also out of the scope of existing theory. Hence, a gap still exists between the theory and the practice of the loss-modification approaches.

To bridge this gap, this paper provides a systematical and fine-grained analysis of loss-modification approaches. After revisiting prior arts, we find that the only property of the loss function utilized in existing proofs is the classic Lipschitz continuity [19, 20]. However, this property is global in nature such that the whole analysis provides no insight into how the losses handle different classes. Inspired by this observation, we extend the classic Lipschitz continuity with a local technique. In this way, the local Lipschitz constants on different classes exactly correspond to the class-dependent terms of the modified loss functions. And a fine-grained generalization bound is established by a novel technique named data-dependent contraction. By applying this bound to the VS loss, the mystery of re-weighting and logit-adjustment is finally uncovered. Last but not least, a principled learning algorithm is proposed based on our theoretical insights.

To sum up, the main contributions of this paper are listed as follows:

- **New technique.** We extend the classic Lipschitz continuity and propose a novel technique named data-dependent contraction to obtain a fine-grained generalization bound for imbalanced learning.
- **Theoretical insights.** Based on the fine-grained bound, a systematical analysis succeeds in explaining the role of re-weighting and logit-adjustment in a unified manner, as well as some empirical results that are out of the scope of existing theories.
- **Principled Algorithm.** A principled algorithm is proposed based on the insights, where the re-weighting term is aligned with the generalization bound, and the multiplicative logit-adjustment term is removed during the DRW phase to avoid the incompatibility between terms.
- **Empirical Validation.** The empirical results on multiple benchmark datasets not only validate the theoretical results, but also demonstrate the superiority of the proposed method.

2 Preliminary

We first introduce the basic notations and the imbalanced learning problem in Sec.2.1. Then, we briefly review existing generalization analysis for imbalanced learning in Sec.2.2.

2.1 Notations and Problem Definition

We assume that the samples are drawn *i.i.d.* from a product space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, where \mathcal{X} is the input space and $\mathcal{Y} = \{1, \dots, C\}$ is the label space. Let $\mathcal{S} = \{(\mathbf{x}^{(n)}, y^{(n)})\}_{n=1}^N$ be the imbalanced training set sampled from the imbalanced distribution \mathcal{D} defined on \mathcal{Z} , $\mathcal{S}_y = \{\mathbf{x} \mid (\mathbf{x}, y) \in \mathcal{S}\}$ be the set of samples from the class y , $N_y := |\mathcal{S}_y|$ denote the size of \mathcal{S}_y , and $\pi_y := N_y/N$. Without loss of generality, we assume that $N_1 \geq N_2 \geq \dots \geq N_C$.

Let \mathcal{D}_{bal} be the balanced distribution defined on \mathcal{Z} . Specifically, a class y is first uniformly sampled from \mathcal{Y} , and then the input \mathbf{x} is sampled from the class-conditional distribution $\mathcal{D}_y := \mathbb{P}[\mathbf{x} \mid y]$.

Then, our task is to learn a score function $f : \mathcal{X} \rightarrow \mathbb{R}^C$ to minimize the risk defined on the balanced distribution:

$$\mathcal{R}_{\text{bal}}(f) := \frac{1}{C} \sum_{y=1}^C \mathcal{R}_y(f) = \frac{1}{C} \sum_{y=1}^C \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_y} [M(f(\mathbf{x}), y)], \quad (1)$$

where \mathcal{R}_y is the risk defined on the class y , and $M : \mathbb{R}^C \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is the measure that evaluates the model performance at $z \in \mathcal{Z}$. For example, one of the most popular choices is to check whether the top-1 prediction is right: $M(f(\mathbf{x}), y) = \mathbf{1}[y \notin \arg \max_{y' \in \mathcal{Y}} f(\mathbf{x})_{y'}]$, where $\mathbf{1}[\cdot]$ is the indicator function. Since M is generally non-differential and thus hard to optimize, one has to select a differential surrogate loss $L : \mathbb{R}^C \times \mathcal{Y} \rightarrow \mathbb{R}_+$, which induces the following surrogate risk:

$$\mathcal{R}_{\text{bal}}^L(f) := \frac{1}{C} \sum_{y=1}^C \mathcal{R}_y^L(f) = \frac{1}{C} \sum_{y=1}^C \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_y} [L(f(\mathbf{x}), y)]. \quad (2)$$

Let $\mathcal{G} := \{L \circ f : f \in \mathcal{F}\}$ denote the hypothesis set. Next, we consider a family of loss functions named Vector-Scaling (VS) [16]:

$$L_{\text{VS}}(f(\mathbf{x}), y) = -\alpha_y \log \left(\frac{e^{\beta_y f(\mathbf{x})_y + \Delta_y}}{\sum_{y'} e^{\beta_{y'} f(\mathbf{x})_{y'} + \Delta_{y'}}} \right). \quad (3)$$

The advantage behind this loss family is two-fold. On one hand, the VS loss generalizes popular re-weighting and logit-adjustment methods. For example, when $\alpha_y = 1, \beta_y = 1, \Delta_y = 0$, it becomes the traditional CE loss [19]. When $\beta_y = 1, \Delta_y = 0$, re-weighting terms $\alpha_y = \pi_y^{-1}$ and $\alpha_y = (1-p)/(1-p^{N_y}), p \in (0, 1)$ recover the classic balanced loss [10] and Class-Balanced (CB) loss [11], respectively. $\alpha_y = 1, \beta_y = 1, \Delta_y = \tau \log \pi_y, \tau > 0$ yield the LA loss [13]. When $\alpha_1 = 1, \beta_y = (N_y/N_1)^\gamma, \Delta_y = 0, \gamma > 0$, we can deduce the CDT loss [14]. On the other hand, an ideal surrogate loss should be Fisher consistent such that minimizing $\mathcal{R}_{\text{bal}}^L(f)$ not only can put more emphasis on minority classes, but also helps bound $\mathcal{R}_{\text{bal}}(f)$ [19, 21]. Fortunately, prior arts [13] have shown that a subset of the VS loss family satisfies such a property:

$$L_{\text{Fisher}}(f(\mathbf{x}), y) := \frac{\delta_y}{\pi_y} \log \left[1 + \sum_{y' \neq y} \frac{\delta_{y'}}{\delta_y} e^{f(\mathbf{x})_{y'} - f(\mathbf{x})_y} \right], \quad (4)$$

where δ_y is an arbitrary positive constant.

2.2 Existing Generalization Analysis for Imbalanced Learning

In balanced learning, we can directly minimize the empirical balanced risk defined on the balanced datasets \mathcal{S}_{bal} sampled from \mathcal{D}_{bal} :

$$\widehat{\mathcal{R}}_{\text{bal}}^L(f) := \frac{1}{N} \sum_{(\mathbf{x}, y) \in \mathcal{S}_{\text{bal}}} L(f(\mathbf{x}), y). \quad (5)$$

Then, the generalization guarantee is available by traditional concentration techniques [19]. However, in imbalanced learning, we can only minimize the empirical risk on the imbalanced dataset \mathcal{S} :

$$\widehat{\mathcal{R}}^L(f) := \frac{1}{N} \sum_{(\mathbf{x}, y) \in \mathcal{S}} L(f(\mathbf{x}), y). \quad (6)$$

To handle this issue, Cao et al. [12] and Ren et al. [18] aggregate the class-wise generalization bound directly with a union bound over class-wise results [19]:

Proposition 1 (Union bound for Imbalanced Learning [12]). *Given the function set \mathcal{F} and a loss function $L : \mathbb{R}^C \times \mathcal{Y} \rightarrow [0, M]$, then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the training set \mathcal{S} , the following generalization bound holds for all $g \in \mathcal{G}$:*

$$\mathcal{R}_{\text{bal}}^L(f) = \frac{1}{C} \sum_{y=1}^C \mathcal{R}_y^L(f) \lesssim \frac{1}{C} \sum_{y=1}^C \left(\widehat{\mathcal{R}}_y^L(f) + \widehat{\mathcal{C}}_{\mathcal{S}_y}(\mathcal{G}) + 3M \sqrt{\frac{\log 2C/\delta}{2N_y}} \right), \quad (7)$$

where $\widehat{\mathcal{R}}_y^L(f)$ is the empirical risk on \mathcal{S}_y ; $\widehat{\mathcal{C}}_{\mathcal{S}}(\mathcal{G}) := \mathbb{E}_{\xi} [\sup_{g \in \mathcal{G}} \frac{1}{N} \sum_{n=1}^N \xi^{(n)} g(\mathbf{z}^{(n)})]$ denotes the empirical complexity of the function set \mathcal{G} , and $\xi := (\xi^{(1)}, \xi^{(2)}, \dots, \xi^{(N)})$ are sampled from independent distributions such as the uniform distribution with $\{1, -1\}$; \lesssim denotes the asymptotic notation that omits undominated terms, that is, $f(t) \lesssim g(t) \iff \exists$ a constant $C > 0, f(t) \leq C \cdot g(t)$.

To further bound the complexity term $\hat{\mathcal{C}}_{\mathcal{S}_y}(\mathcal{G})$, Cao et al. [12] assume that the loss function L satisfies the Lipschitz continuity and applies the traditional contraction lemma [20]:

Definition 1 (Lipschitz Continuity). *Let $\|\cdot\|$ denote the 2-norm. Then, we say the loss function $L(f, y)$ is Lipschitz continuous with constant μ if for any $f, f' \in \mathcal{F}$, $\mathbf{x} \in \mathcal{S}$,*

$$|L(f, y) - L(f', y)| \leq \mu \cdot \|f(\mathbf{x}) - f'(\mathbf{x})\|. \quad (8)$$

Lemma 1 (Contraction Lemma). *Assume that the loss function $L(f, \mathbf{x})$ is Lipschitz continuous with a constant μ . Then, the following inequality holds:*

$$\hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{G}) \leq \mu \cdot \hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{F}). \quad (9)$$

Finally, the standard margin-based generalization bound [22] is directly applied to obtain the upper bound of $\hat{\mathcal{C}}_{\mathcal{S}_y}(\mathcal{F})$. However, this union bound has the following limitations:

- Theoretically, this generalization bound is coarse-grained and not sharp enough. To be specific, the differences among different loss functions lie in the choice of $\alpha_y, \beta_y, \Delta_y$. However, the Lipschitz continuity, which is the only property of L utilized in the proof, is global in nature and thus obscures these differences. Although the margin theory can provide some theoretical insights into the role of Δ_y , the roles of α_y, β_y are still a mystery. Besides, since

$$\text{Bound}(\mathcal{R}_{\text{bal}}^L(f)) = \frac{1}{C} \text{Bound}\left(\sum_y \mathcal{R}_y^L(f)\right) \leq \frac{1}{C} \sum_y \text{Bound}(\mathcal{R}_y^L(f)), \quad (10)$$

a sharper bound might be available if we can bound $\mathcal{R}_{\text{bal}}^L(f)$ directly.

- Empirically, although the induced LDAM loss outperforms the CE loss, the improvement is not so significant. Fortunately, when combining the Deferred Re-Weighting (DRW) technique [12], where $\alpha_y = (1-p)/(1-p^{N_y})$, $p \in (0, 1)$ [11] during the terminal phase of training, the improvement becomes much more impressive. However, Eq.(7) fails to explain this phenomenon.

Recently, Kini et al. [16] provide a generalization analysis for the VS loss. However, the results, which only hold for linear models with linearly separable data, can only explain the roles of β_y . For the role of Δ_y , they resort to analyzing the gradient of the VS loss and provide a coarse-grained analysis.

To sum up, existing generalization analysis for imbalanced learning is coarse-grained and fragmented. Next, we aim to build a more fine-grained and systematical generalization bound that can unify the roles of both re-weighting and logit-adjustment.

3 Fine-Grained Generalization Analysis for Imbalanced Learning

In Sec.3.1, we first establish a sharp generalization bound based on a novel technique named data-dependent contraction. Then, in Sec.3.2, we apply this generalization bound to the VS loss to provide a series of theoretical insights. Finally, in Sec.3.3, a principled algorithm is proposed based on the theoretical insights.

3.1 Generalization Bound Induced By Data-Dependent Contraction

Different from Eq.(7), we hope to build a direct bound between $\mathcal{R}_{\text{bal}}^L(f)$ and $\hat{\mathcal{R}}^L(f)$. To this end, our analysis is based on the following lemma, whose proof can be found in Appendix B:

Lemma 2. *Given the function set \mathcal{F} and a loss function $L : \mathbb{R}^C \times \mathcal{Y} \rightarrow [0, M]$, then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the training set \mathcal{S} , the following generalization bound holds for all $g \in \mathcal{G}$:*

$$\mathcal{R}_{\text{bal}}^L(f) \lesssim \Phi(L, \delta) + \frac{1}{C\pi_C} \cdot \hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{G}), \quad (11)$$

where $\Phi(L, \delta) := \frac{1}{C\pi_C} [\hat{\mathcal{R}}^L(f) + 3M\sqrt{\frac{\log 2/\delta}{2N}}]$ contains the empirical risk on \mathcal{S} and the δ term.

Remark 1. *Recall that $\pi_C := N_C/N$, $N_1 \geq N_2 \geq \dots \geq N_C$. Hence, this lemma reveals how the model performance depends on the imbalance degree of the data.*

As shown in Sec.2.2, the fine-grained analysis is unavailable due to the global nature of the classic Lipschitz continuous property. In view of this, we extend this traditional definition with a local technique [23]:

Definition 2 (Local Lipschitz Continuity). *We say the loss function $L(f, y)$ is local Lipschitz continuous with constants $\{\mu_y\}_{y=1}^C$ if for any $f, f' \in \mathcal{F}$, $y \in \mathcal{Y}$, $\mathbf{x} \in \mathcal{S}_y$,*

$$|L(f, y) - L(f', y)| \leq \mu_y \cdot \|f(\mathbf{x}) - f'(\mathbf{x})\|. \quad (12)$$

Then, the following data-dependent contraction inequality helps us obtain a sharper bound, whose proof is given in Appendix C.

Assumption 1. *Next, we assume that $\hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{F}) \sim \mathcal{O}(1/\sqrt{N})$. Note that this result holds for kernel-based models with traditional techniques [19] and neural networks with the latest techniques [24, 25]. And the prior arts also adopt this assumption [12].*

Lemma 3 (Data-Dependent Contraction). *Assume that the loss function $L(f, \mathbf{x})$ is local Lipschitz continuous with constants $\{\mu_y\}_{y=1}^C$. Then, the following inequality holds under Asm.1:*

$$\hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{G}) \lesssim \hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{F}) \sum_{y=1}^C \mu_y \sqrt{\pi_y}, \quad (13)$$

Combining Lem.2 and Lem.3, we have the following theorem:

Theorem 1 (Data-Dependent Bound for Imbalanced Learning). *Given the function set \mathcal{F} and a loss function $L : \mathbb{R}^C \times \mathcal{Y} \rightarrow [0, M]$, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the training set \mathcal{S} , the following generalization bound holds for all $f \in \mathcal{F}$:*

$$\mathcal{R}_{bal}^L(f) \lesssim \Phi(L, \delta) + \frac{\hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{F})}{C\pi_C} \sum_{y=1}^C \mu_y \sqrt{\pi_y}. \quad (14)$$

At the first glance, Eq.(14) seems a little loose since $\sum_{y=1}^C \sqrt{\pi_y} > 1$. In fact, this intuition holds when local Lipschitz continuity degenerates to Def.1. However, when μ_y is decreasing w.r.t. π_y , a shaper bound might be available. To build an intuitive understanding, we present the following proposition, whose proof can be found in Appendix D.

Proposition 2. *Assume that $\mu_y \propto N_y^{-\kappa}$, $\kappa > 0$. Then, when $\kappa > 1$, the data-dependent bound presented in Thm.1 is sharper than the union bound defined in Prop.1.*

3.2 Application to the VS Loss

Next, we apply Thm.1 to the VS loss to reveal the role of both re-weighting and logit-adjustment. To this end, it is necessary to analyze the local Lipschitz property of the VS loss, whose proof is presented in Appendix E.

Lemma 4. *Assume that the score function is bounded. Then, the VS loss is local Lipschitz continuous with constants $\{\mu_y\}_{y=1}^C$, where*

$$\mu_y = \alpha_y \tilde{\beta}_y [1 - \text{softmax}(\beta_y B_y(f) + \Delta_y)], \quad (15)$$

$\tilde{\beta}_y := \sqrt{\beta_y^2 + \left(\sum_{y' \neq y} \beta_{y'}\right)^2}$; *softmax*(\cdot) denotes the softmax function; $B_y(f)$ denotes the minimal prediction on the ground-truth class y , i.e., $B_y(f) := \min_{\mathbf{x} \in \mathcal{S}_y} f(\mathbf{x})_y$.

Remark 2. $B_y(f)$ is closely related to the minimal margin defined by $\text{margin}_y^\downarrow := \min_{\mathbf{x} \in \mathcal{S}_y} (f(\mathbf{x})_y - \max_{j \neq y} f(\mathbf{x})_j)$. It is not difficult to check that $B_y(f) - \text{margin}_y^\downarrow \leq \max_{\mathbf{x} \in \mathcal{S}_y, j \neq y} f(\mathbf{x})_j$. Hence, as we improve the model performance on class y , the RHS of the above inequality, i.e., the gap between $B_y(f)$ and $\text{margin}_y^\downarrow$ will decrease, and both the minimal margin and $B_y(f)$ will increase.

Then, combining Thm.1 and Lem.4, we have the following proposition, which reveals how the existing loss-oriented methods improve generalization performance by exploiting the data priors.

Proposition 3 (Data-Dependent Bound for the VS Loss). *Given the function set \mathcal{F} and the VS loss L_{VS} , for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the training set \mathcal{S} , the following generalization bound holds for all $f \in \mathcal{F}$:*

$$\mathcal{R}_{bal}^L(f) \lesssim \Phi(L_{VS}, \delta) + \frac{\hat{\mathcal{C}}_{\mathcal{S}}(\mathcal{F})}{C\pi_C} \sum_{y=1}^C \alpha_y \tilde{\beta}_y \sqrt{\pi_y} [1 - \text{softmax}(\beta_y B_y(f) + \Delta_y)]. \quad (16)$$

From Eq.(16), we have the following insights, whose empirical validation can be found in Sec.4.2.

(In1) Why re-weighting and logit-adjustment are necessary? Due to the term $\sqrt{\pi_y}$ and $B_y(f)$, the generalization bound is also imbalanced among classes. Both re-weighting and logit-adjustment can obtain a sharper generalization bound by assigning different weights to the classes with different $\sqrt{\pi_y}$ and $B_y(f)$. In this process, α_y mainly rebalances the generalization performance among classes, *i.e.*, $\sqrt{\pi_y}$, while β_y and Δ_y focus on adjusting the imbalance of the terms $B_y(f)$ among classes.

(In2) Why the deferred scheme is necessary? As pointed out in [11, 17], weighting up the minority classes will cause difficulties and instability in optimization, especially when the distribution is extremely imbalanced. To fix this issue, Cao et al. [12] develops a deferred scheme, where $\alpha_y = 1$ and $(1-p)/(1-p^{N_y})$, $p \in (0, 1)$ during the initial and terminal phase of training, respectively. Although this scheme shows significant improvement, there is still a lack of theoretical explanation.

Fortunately, Prop.3 can give us some inspiration. Specifically, although a weighted loss can boost the optimization on the minority classes, it is harmful to the further improvement on the majority classes, as shown in Fig.3. Hence, the majority/minority classes will have relatively small/large $B_y(f)$ respectively, and the generalization bound becomes even looser. By contrast, in the DRW scheme, we have $\alpha_y = 1$ during the initial phase of training. Such a warm-up phase will encourage the model to focus on the majority classes and induce a small $B_y(f)$ for both majority and minority classes after weighting up the minority classes. On top of this, the generalization bound can become sharper, which explains the effectiveness of the deferred scheme.

(In3) How does our result explain the design of existing losses? On one hand, for re-weighting losses, α_y should decrease as π_y increases, which is consistent with the balanced loss with $\alpha_y = \pi_y^{-1}$ [10] and $\alpha_y = (1-p)/(1-p^{N_y})$, $p \in (0, 1)$ [11]. On the other hand, from the insight **(In2)**, we know that when $\alpha_y = 1$, $B_y(f)$ will be increasing *w.r.t.* π_y . Hence, for logit-adjustment losses, both β_y and Δ_y should increase as π_y increases. This insight is consistent with the LDAM loss ($\Delta_y \propto -N_y^{-1/4}$) [12], the logit-adjusted loss ($\Delta_y = \tau \log \pi_y$) [13], and the CDT loss ($\beta_y = (N_y/N_1)^\gamma$) [14].

(In4) Are re-weighting and logit-adjustment fully compatible? **(a)** Unfortunately, the answer is negative. To be specific, the re-weighting term α_y is decreasing *w.r.t.* π_y , whereas the multiplicative logit-adjustment term β_y is increasing *w.r.t.* π_y . As a result, $\tilde{\beta}_y$ will weaken the effect of α_y . **(b)** Fortunately, α_y is compatible with the additive logit-adjustment term Δ_y since both terms can induce a sharper generalization bound.

3.3 Principled Learning Algorithm induced by the Theoretical Insights

In this part, we present a principled learning algorithm induced by the theoretical insights in Sec.3.2. First, according to **(In1)-(In3)**, it is crucial to comprehensively utilize re-weighting, logit-adjustment, and the DRW scheme, as they all contribute to improving the generalization bound. Second, according to **(In4)**, we propose a Truncated Logit-Adjustment (TLA) scheme to avoid the conflict between α_y and β_y . In this scheme, β_y still increases *w.r.t.* π_y during the initial phase of training but is truncated to 1 during the terminal phase of training. Third, we set $\alpha_y \propto \pi_y^{-\nu}$, $\nu > 0$ to align α_y with $\sqrt{\pi_y}$, which we name Aligned DRW (ADRW). Note that such a re-weighting scheme also follows the Fisher consistency property presented in [13]. Finally, the overall algorithm is summarized in Alg.1, where the logit-adjustment methods mentioned in Sec.2.1 are all reasonable options for β_y , Δ_y in the line 5 and Δ_y in the line 7.

Algorithm 1: Principled Learning Algorithm induced by the Theoretical Insights

Require: Training set $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^N$ and a model f parameterized by Θ .

- 1: Initialize the model parameters Θ randomly.
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: $\mathcal{B} \leftarrow \text{SampleMiniBatch}(\mathcal{S}, m)$ ▷ A mini-batch of m samples
 - 4: **if** $t < T_0$ **then**
 - 5: Set $\alpha = 1, \beta_y, \Delta_y$ ▷ Adjust logits during the initial phase
 - 6: **else**
 - 7: Set $\alpha_y \propto \pi_y^{-\nu}, \beta_y = 1, \Delta_y, \nu > 0$ ▷ TLA and ADRW
 - 8: **end if**
 - 9: $L(f, \mathcal{B}) \leftarrow \frac{1}{m} \sum_{(x,y) \in \mathcal{B}} L_{\text{vs}}(f(x), y)$ ▷ Calculate the loss
 - 10: $\Theta \leftarrow \Theta - \eta \nabla_{\Theta} L(f, \mathcal{B})$ ▷ One SGD step
 - 11: Optional: anneal the learning rate η . ▷ Required when $t = T_0$
 - 12: **end for**
-




Figure 1: The balanced accuracy of the CE loss and the LDAM loss *w.r.t.* $\alpha_y \propto \pi_y^{-\nu}$ on the CIFAR datasets, where the imbalance ratio $\rho = 100$. Both re-weighting and logit-adjustment boost the model performance, which is consistent with the theoretical insight **(In1)** and **(In4-b)**.

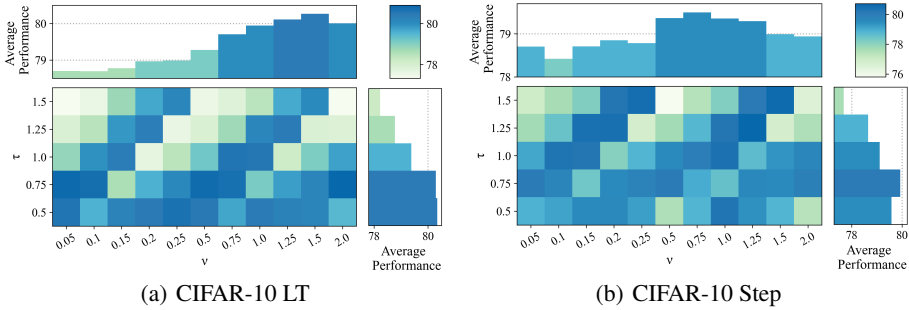


Figure 2: Sensitivity analysis of VS+ADRW *w.r.t.* $\alpha_y \propto \pi_y^{-\nu}$ and $\Delta_y = \tau \log \pi_y$ on the CIFAR-10 dataset, where the imbalance ratio $\rho = 100$. Both re-weighting and logit-adjustment boost the model performance, which is consistent with the theoretical insights **(In1)** and **(In4-b)**.

4 Experiments

4.1 Experiment Protocols

Here, we briefly introduce the experiment protocols, and more details can be found in Appendix F.

Datasets. We conduct the experiments on four popular benchmark datasets for imbalanced learning. **(a) CIFAR-10 and CIFAR-100:** Following the protocol in [26, 11, 12], we consider two types of imbalance: long-tailed imbalance (LT) and step imbalance (Step). For both imbalance types, we

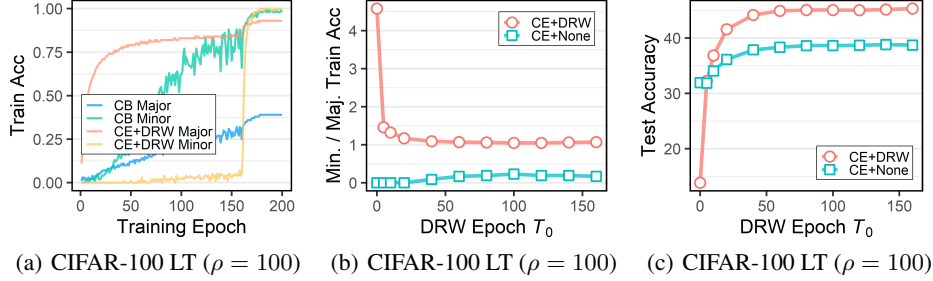


Figure 3: (a) Training accuracy of CE+DRW ($T_0 = 160$) and the CB loss *w.r.t.* training epoch. (b) $\widehat{Acc}_{\min}/\widehat{Acc}_{\max}$ *w.r.t.* the DRW epoch T_0 , where \widehat{Acc}_{\min} and \widehat{Acc}_{\max} denote the training accuracy of the best model on the minority/majority classes, respectively. (c) The test accuracy of the best model *w.r.t.* the DRW epoch T_0 . We can find that the DRW scheme balances the training accuracy between the majority classes and the minority classes and thus improves the model performance on the test set, which is consistent with the theoretical insight (In2).




Figure 4: The balanced accuracy of the VS loss *w.r.t.* $\beta_y = (N_y/N_1)\gamma$ on the CIFAR datasets, where the imbalance ratio $\rho = 10$. We can find that VS+DRW performs inferior to VS+None, especially when γ is large, which is consistent with the theoretical insight (In4-a).

report the balanced accuracy averaged over 5 random seeds with an imbalance ratio $\rho := N_1/N_C \in \{10, 100\}$. **(b) ImageNet-LT and iNaturalist:** We use the long-tailed version of the ImageNet dataset² [2] proposed by [27], and iNaturalist³ [5] is a real-world long-tailed dataset.

Baselines and Competitors. For the CIFAR datasets, we aim to validate the theoretical results and the performance gain induced by the proposed method. Hence, we select the following baselines: the CE loss (CE) [19], the LDAM loss (LDAM) [12], LDAM+DRW [12], and the VS loss (VS) [16] that generalizes the LA loss [13] and the CDT loss [14]. We tune all the hyperparameters according to the suggestions in the original papers. For the ImageNet-LT and iNaturalist datasets, we select state-of-the-art methods, listed in Tab.2, as the competitors to validate the effectiveness of the method.

Implementation Details. We implement three instances of the proposed learning algorithm: the CE loss equipped with the ADRW scheme (**CE+ADRW**), the LDAM loss equipped with the ADRW scheme (**LDAM+ADRW**), and the VS loss equipped with the TLA and the ADRW scheme (**VS+TLA+ADRW**). We tune the hyperparameter ν , and the other hyperparameters follow those used in the baselines. In addition, we incorporate the Sharpness-Aware Minimization (SAM) technique [28] to facilitate the optimization of the minority classes, allowing them to escape saddle points and converge to flat minima [29].

4.2 Theory Validation

In this part, we aim to validate our theoretical insights presented in Sec.3.2 on the CIFAR datasets. Some more empirical results can be found in Appendix G.

¹<https://www.cs.toronto.edu/~kriz/cifar.html>. Licensed MIT.

²<https://image-net.org/index.php>. Licensed MIT.

³https://github.com/visipedia/inat_comp/tree/master/2017. Licensed MIT.

Table 1: The balanced accuracy averaged over 5 random seeds on the CIFAR datasets. The best and the runner-up method for each protocol are marked with **red** and **blue**, respectively. The best baseline model is marked with underline.

Dataset	CIFAR-10				CIFAR-100			
Imbalance Type	LT		Step		LT		Step	
Imbalance Ratio	100	10	100	10	100	10	100	10
w/o SAM								
CE	71.5 \pm 0.4	87.0 \pm 0.2	64.8 \pm 0.9	85.1 \pm 0.3	38.3 \pm 0.4	56.7 \pm 0.4	38.6 \pm 0.2	54.4 \pm 0.3
LDAM	73.8 \pm 0.4	86.4 \pm 0.4	65.8 \pm 0.6	85.0 \pm 0.3	39.9 \pm 0.7	55.7 \pm 0.5	39.2 \pm 0.0	50.5 \pm 0.2
VS	78.8 \pm 0.2	<u>88.7\pm0.1</u>	76.1 \pm 0.7	<u>88.3\pm0.1</u>	41.8 \pm 0.7	<u>58.4\pm0.2</u>	<u>46.2\pm0.3</u>	<u>59.9\pm0.2</u>
CE+DRW	75.8 \pm 0.3	87.9 \pm 0.3	72.2 \pm 0.8	88.0 \pm 0.3	40.8 \pm 0.6	58.1 \pm 0.3	45.4 \pm 0.4	59.1 \pm 0.3
LDAM+DRW	77.7 \pm 0.4	87.5 \pm 0.2	77.8 \pm 0.5	87.8 \pm 0.3	<u>42.7\pm0.5</u>	57.5 \pm 0.3	45.3 \pm 0.6	56.9 \pm 0.2
VS+DRW	<u>80.1\pm0.1</u>	<u>88.6\pm0.1</u>	<u>78.2\pm0.2</u>	88.1 \pm 0.1	41.3 \pm 0.4	57.6 \pm 0.3	44.0 \pm 0.3	58.0 \pm 0.3
CE+ADRW	78.6 \pm 0.5	88.2 \pm 0.3	75.5 \pm 0.6	88.5 \pm 0.2	41.8 \pm 0.6	58.3 \pm 0.4	46.5 \pm 0.3	59.2 \pm 0.3
LDAM+ADRW	79.1 \pm 0.2	87.6 \pm 0.2	78.5 \pm 0.4	88.1 \pm 0.2	43.0\pm0.2	58.0 \pm 0.1	45.8 \pm 0.2	57.6 \pm 0.3
VS+TLA+DRW	80.8\pm0.2	88.8\pm0.1	80.0\pm0.1	89.2\pm0.1	43.0 \pm 0.4	58.9\pm0.1	46.8\pm0.1	60.0\pm0.3
VS+TLA+ADRW	81.1\pm0.2	89.0\pm0.2	80.9\pm0.2	89.3\pm0.1	43.4\pm0.6	59.2\pm0.2	47.8\pm0.1	60.5\pm0.3
w/ SAM								
CE+DRW	80.5 \pm 0.2	89.8 \pm 0.2	79.5 \pm 0.3	90.2 \pm 0.2	44.7 \pm 0.6	60.7 \pm 0.4	48.5 \pm 0.3	61.7 \pm 0.2
LDAM+DRW	81.6 \pm 0.2	89.4 \pm 0.2	81.2 \pm 0.7	89.4 \pm 0.1	45.2 \pm 0.3	59.9 \pm 0.2	49.1 \pm 0.2	59.3 \pm 0.2
VS	82.6 \pm 0.2	<u>90.0\pm0.1</u>	83.2 \pm 0.4	90.5 \pm 0.1	45.9 \pm 0.3	61.0 \pm 0.3	47.4 \pm 0.3	61.6 \pm 0.3
CE+ADRW	82.6 \pm 0.2	90.1\pm0.1	82.8\pm0.9	90.2\pm0.3	44.9 \pm 0.6	61.0\pm0.4	48.9 \pm 0.2	62.1\pm0.2
LDAM+ADRW	83.0\pm0.1	89.7 \pm 0.1	82.4 \pm 0.3	90.0 \pm 0.2	46.3\pm0.4	60.3 \pm 0.3	49.3\pm0.4	60.3 \pm 0.2
VS+TLA+ADRW	83.6\pm0.2	90.3\pm0.2	83.8\pm0.1	90.8\pm0.1	46.4\pm0.6	61.9\pm0.3	49.1\pm0.2	62.3\pm0.3

Validation of (In1) and (In4-b). We report the model performance of the baselines *w.r.t.* the hyperparameters in Fig.1 and Fig.2. From these results, we can find that (1) Both **CE+ADRW** and **LDAM** perform better than **CE**. In other words, either re-weighting or logit-adjustment can boost the model performance. (2) **LDAM+ADRW** outperforms both **CE+ADRW** and **LDAM**, and for VS+ADRW, increasing the hyperparameters ν and τ appropriately can also bring performance gains. All these results validate the compatibility between re-weighting and the additive logit-adjustment.

Validation of (In2). We present a series of results in Fig.3, where the training accuracy of different classes represents the corresponding $B_f(y)$. To be specific, Fig.3(a) demonstrates the trend of training accuracy on the majority classes and the minority classes. For the CB loss, the learning process only focuses on the minority classes and hinders the performance improvement on the majority classes (**CB Major** *v.s.* **CB Minor**). Hence, an extremely imbalanced $B_y(f)$ induces a poor generalization performance. By contrast, the DRW scheme first focuses on the majority classes (**CE+DRW Major** *v.s.* **CE+DRW Minor** with the training epoch $t \leq T_0 = 160$) and then pays more attention to the minority classes during the terminal phase of training (**CE+DRW Major** *v.s.* **CE+DRW Minor** with the training epoch $t > T_0 = 160$). Benefiting from the DRW scheme, both the majority classes and the minority classes are well-trained and thus have a balanced term $B_y(f)$ (Fig.3(b)), leading to a corresponding improvement on the test accuracy (Fig.3(c)). Even if we remove the re-weighting term (**CE+None**), the imbalance degree of $B_y(f)$ is still consistent with the test performance. Note that the learning rate of CE+None is also decreased at the corresponding epoch T_0 , making the line in Figure 3(b) and 3(c) not constant.

Validation of (In4-a). In Fig.4, we present the model performance of the VS loss *w.r.t.* the hyperparameter of the multiplicative logit-adjustment γ . We can find that **VS+DRW** performs inferior to **VS**, especially when γ is large. This phenomenon validates the incompatibility between re-weighting and multiplicative logit-adjustment.

4.3 Performance Comparison

We report the empirical results on the CIFAR datasets in Tab.1, where the imbalance ratio $\rho \in \{10, 100\}$. From these results, we have the following observations: (1) The proposed learning algorithm consistently outperforms the baselines, especially when the dataset is more imbalanced.

Table 2: The balanced accuracy on the ImageNet-LT and iNaturalist datasets.

Method	One stage	ImageNet-LT				iNaturalist			
		Many	Med.	Few	All	Many	Med.	Few	All
OLTR [27]	×	43.2	35.1	18.5	35.6	59.0	64.1	64.9	63.9
LFMR [30]	×	47.1	35.0	17.5	37.2	-	-	-	-
BBN [31]	×	-	-	-	-	49.4	70.8	65.3	66.3
cRT [32]	×	61.8	46.2	27.3	49.6	69.0	66.0	63.2	65.2
τ -norm [32]	×	59.1	46.9	30.7	49.4	65.6	65.3	65.5	65.6
DiVE [33]	×	64.1	50.4	31.5	53.1	70.6	70.0	67.6	69.1
DisAlign [34]	×	61.3	52.2	31.4	52.9	69.0	71.1	70.2	70.6
WB [35]	×	62.5	50.4	41.5	53.9	71.2	70.4	69.7	70.2
CE [32]	✓	65.9	37.5	7.7	44.4	72.2	63.0	57.2	61.7
CE+CB [11]	✓	39.6	32.7	16.8	33.2	53.4	54.8	53.2	54.0
Focal [11]	✓	36.4	29.9	16.0	30.5	-	-	-	61.1
De-confound [36]	✓	62.7	48.8	31.6	51.8	-	-	-	-
DRO-LT [37]	✓	64.0	49.8	33.1	53.5	-	-	-	69.7
SAM [29]	✓	62.0	52.1	34.8	53.1	64.1	70.5	71.2	70.1
Ours	✓	62.9	52.6	37.1	54.1	64.7	70.7	72.1	70.7

Such performance gains validate the effectiveness of the proposed methods (2) Both re-weighting and logit-adjustment can improve the model performance, which is consistent with the theoretical insight **(In1)** and **(In4-b)**. (3) When $\rho = 10$ or on the CIFAR-100 dataset, VS+DRW performs inferior to VS. Fortunately, when equipped with the proposed TLA scheme, VS+TLA+DRW outperforms both VS and VS+DRW. These results again validate our theoretical insight **(In4-a)**. (4) When $\rho = 10$, CE+ADRW outperforms LDAM+ADRW, and similar counter-intuitive phenomena are also observed in [29]. We conjecture that in this case, re-weighting is enough to rebalance the generalization bound, and the additional LDAM loss might induce other issues such as inconsistency.

We present the overall balanced accuracy on the ImageNet-LT and iNaturalist datasets in Tab.2, where SAM and Ours denotes LDAM+DRW+SAM and VS+TLA+ADRW+SAM, respectively. These results demonstrate that the proposed learning algorithm outperforms the competitors, especially the one-stage ones, which again confirms the effectiveness of the proposed learning algorithm.

5 Conclusion and Future Work

In this work, with the proposed local Lipschitz property and the data-dependent contraction technique, we present a unified generalization analysis of the loss-modification approaches for imbalanced learning. Benefiting from this fine-grained analysis, we not only reveal the role of both re-weighting and logit-adjustment approaches but also explain some empirical phenomena that are out of the scope of existing theories. Moreover, a principled learning algorithm is proposed based on the theoretical insights. Finally, extensive experimental results on benchmark datasets validate our theoretical analysis and the effectiveness of the proposed method.

Theoretically, one important future work is to provide a systematical Fisher consistency analysis for the VS loss, providing more insights to design re-weighting and logit-adjustment terms. Methodologically, it might be a promising direction to design an adaptive scheme that can automatically determine the hyperparameters of the learning algorithm.

Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grant 2018AAA0102000, in part by National Natural Science Foundation of China: 62236008, U21B2038, U2001202, 61931008, 62122075, 61976202 and 62206264, in part by the Fundamental Research Funds for the Central Universities, in part by Youth Innovation Promotion Association CAS, in part by the Strategic Priority Research Program of Chinese Academy of Sciences (Grant No. XDB28000000) and in part by the Innovation Funding of ICT, CAS under Grant No. E000000.

References

- [1] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. pages 1–60, 2009.
- [2] 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. *Int. J. Comput. Vis.*, 115:211–252, 2015.
- [3] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88:303–338, 2010.
- [4] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *European Conference on Computer Vision*, pages 740–755, 2014.
- [5] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alexander Shepard, Hartwig Adam, Pietro Perona, and Serge J. Belongie. The inaturalist species classification and detection dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8769–8778, 2018.
- [6] Haibo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE Trans. Knowl. Data Eng.*, 21:1263–1284, 2009.
- [7] Grant Van Horn and Pietro Perona. The devil is in the tails: Fine-grained classification in the wild. *CoRR*, abs/1709.01450, 2017.
- [8] Kemal Oksuz, Baris Can Cam, Sinan Kalkan, and Emre Akbas. Imbalance problems in object detection: A review. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43:3388–3415, 2021.
- [9] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *CoRR*, abs/2110.04596, 2021.
- [10] Katharina Morik, Peter Brockhausen, and Thorsten Joachims. Combining statistical learning with a knowledge-based approach - A case study in intensive care monitoring. In *International Conference on Machine Learning*, pages 268–277, 1999.
- [11] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge J. Belongie. Class-balanced loss based on effective number of samples. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019.
- [12] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Aréchiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *Annual Conference on Neural Information Processing Systems*, pages 1565–1576, 2019.
- [13] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. In *International Conference on Learning Representations*, 2021.
- [14] Han-Jia Ye, Hong-You Chen, De-Chuan Zhan, and Wei-Lun Chao. Identifying and compensating for feature deviation in imbalanced deep learning. *CoRR*, abs/2001.01385, 2020.
- [15] Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, and Junjie Yan. Equalization loss for long-tailed object recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11659–11668, 2020.
- [16] Ganesh Ramachandra Kini, Orestis Paraskevas, Samet Oymak, and Christos Thrampoulidis. Label-imbalanced and group-sensitive classification under overparameterization. In *Annual Conference on Neural Information Processing Systems*, pages 18970–18983, 2021.
- [17] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5375–5384, 2016.

- [18] Jiawei Ren, Cunjun Yu, Shunan Sheng, Xiao Ma, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Balanced meta-softmax for long-tailed visual recognition. In *Annual Conference on Neural Information Processing Systems*, 2020.
- [19] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. The MIT Press, 2012.
- [20] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23. Springer Science & Business Media, 1991.
- [21] Aditya Krishna Menon, Harikrishna Narasimhan, Shivani Agarwal, and Sanjay Chawla. On the statistical consistency of algorithms for binary classification under class imbalance. In *International Conference on Machine Learning*, pages 603–611, 2013.
- [22] Sham M. Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Annual Conference on Neural Information Processing Systems*, pages 793–800, 2008.
- [23] Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Localized rademacher complexities. In *Annual Conference on Computational Learning Theory*, pages 44–58, 2002.
- [24] Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. In *Conference On Learning Theory*, pages 297–299, 2018.
- [25] Philip M. Long and Hanie Sedghi. Generalization bounds for deep convolutional neural networks. In *International Conference on Learning Representations*, 2020.
- [26] Mateusz Buda, Atsuto Maki, and Maciej A. Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- [27] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019.
- [28] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021.
- [29] Harsh Rangwani, Sumukh K. Aithal, Mayank Mishra, and R. Venkatesh Babu. Escaping saddle points for effective generalization on class-imbalanced data. In *Advances in Neural Information Processing Systems*, pages 22791–22805, 2022.
- [30] Liuyu Xiang, Guiguang Ding, and Jungong Han. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *European Conference on Computer Vision*, pages 247–263, 2020.
- [31] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. BBN: bilateral-branch network with cumulative learning for long-tailed visual recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9716–9725, 2020.
- [32] 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.
- [33] Yin-Yin He, Jianxin Wu, and Xiu-Shen Wei. Distilling virtual examples for long-tailed recognition. In *IEEE/CVF International Conference on Computer Vision*, pages 235–244, 2021.
- [34] Songyang Zhang, Zeming Li, Shipeng Yan, Xuming He, and Jian Sun. Distribution alignment: A unified framework for long-tail visual recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2361–2370, 2021.
- [35] Shaden Alshammari, Yu-Xiong Wang, Deva Ramanan, and Shu Kong. Long-tailed recognition via weight balancing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6887–6897, 2022.

- [36] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *Annual Conference on Neural Information Processing Systems*, 2020.
- [37] Dvir Samuel and Gal Chechik. Distributional robustness loss for long-tail learning. In *IEEE/CVF International Conference on Computer Vision*, pages 9475–9484, 2021.
- [38] Lu Yang, He Jiang, Qing Song, and Jun Guo. A survey on long-tailed visual recognition. *Int. J. Comput. Vis.*, 130:1837–1872, 2022.
- [39] Miroslav Kubat and Stan Matwin. Addressing the curse of imbalanced training sets: One-sided selection. In *International Conference on Machine Learning*, pages 179–186, 1997.
- [40] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16: 321–357, 2002.
- [41] Inderjeet Mani and I Zhang. knn approach to unbalanced data distributions: a case study involving information extraction. In *Workshop on Learning from Imbalanced Datasets, International Conference on Machine Learning*, pages 1–7, 2003.
- [42] Chris Drummond, Robert C Holte, et al. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on Learning from Imbalanced Datasets II, International Conference on Machine Learning*, pages 1–8, 2003.
- [43] Justin M. Johnson and Taghi M. Khoshgoftaar. Survey on deep learning with class imbalance. *J. Big Data*, 6:27, 2019.
- [44] Guillem Collell, Drazen Prelec, and Kaustubh R. Patil. Reviving threshold-moving: a simple plug-in bagging ensemble for binary and multiclass imbalanced data. *CoRR*, abs/1606.08698, 2016.
- [45] Jiarui Cai, Yizhou Wang, and Jenq-Neng Hwang. ACE: ally complementary experts for solving long-tailed recognition in one-shot. In *IEEE/CVF International Conference on Computer Vision*, pages 112–121, 2021.
- [46] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella X. Yu. Long-tailed recognition by routing diverse distribution-aware experts. In *International Conference on Learning Representations*, 2021.
- [47] Jiequan Cui, Zhisheng Zhong, Shu Liu, Bei Yu, and Jiaya Jia. Parametric contrastive learning. In *IEEE/CVF International Conference on Computer Vision*, pages 695–704, 2021.
- [48] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 16489–16498, 2021.
- [49] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [50] Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*, pages 1139–1147, 2013.
- [51] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.

- [52] Charles R. Harris, K. Jarrod Millman, Stéfan van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with numpy. *Nat.*, 585:357–362, 2020.
- [53] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake VanderPlas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.*, 12:2825–2830, 2011.
- [54] Sébastien Marcel and Yann Rodriguez. Torchvision the machine-vision package of torch. In *Proceedings of the 18th International Conference on Multimedia 2010, Firenze, Italy, October 25-29, 2010*, pages 1485–1488, 2010.