

SSE-SAM: Balancing Head and Tail Classes Gradually Through Stage-Wise SAM

Xingyu Lyu^{1,2}, Qianqian Xu^{1*}, Zhiyong Yang², Shaojie Lyu³, Qingming Huang^{2, 1, 4*}

¹Key Lab. of Intelligent Information Processing, Institute of Computing Tech., CAS

²School of Computer Science and Tech., University of Chinese Academy of Sciences

³Tencent Corporate

⁴BDKM, University of Chinese Academy of Sciences

lvxingyu20@mails.ucas.ac.cn, xuqianqian@ict.ac.cn, yangzhiyong21@ucas.ac.cn, shaojielv@tencent.com, qmhuang@ucas.ac.cn

Abstract

Real-world datasets often exhibit a long-tailed distribution, where vast majority of classes known as tail classes have only few samples. Traditional methods tend to overfit on these tail classes. Recently, a new approach called Imbalanced SAM (ImbSAM) is proposed to leverage the generalization benefits of Sharpness-Aware Minimization (SAM) for long-tailed distributions. The main strategy is to merely enhance the smoothness of the loss function for tail classes. However, we argue that improving generalization in long-tail scenarios requires a careful balance between head and tail classes. We show that neither SAM nor ImbSAM alone can fully achieve this balance. For SAM, we prove that although it enhances the model’s generalization ability by escaping saddle point in the overall loss landscape, it does not effectively address this for tail-class losses. Conversely, while ImbSAM is more effective at avoiding saddle points in tail classes, the head classes are trained insufficiently, resulting in significant performance drops. Based on these insights, we propose Stage-wise Saddle Escaping SAM (SSE-SAM), which uses complementary strengths of ImbSAM and SAM in a phased approach. Initially, SSE-SAM follows the majority sample to avoid saddle points of the head-class loss. During the later phase, it focuses on tail-classes to help them escape saddle points. Our experiments confirm that SSE-SAM has better ability in escaping saddles both on head and tail classes, and shows performance improvements.

Code — <https://github.com/Zemdalk/SSE-SAM>

Introduction

Over the past few decades, deep neural networks have made significant progress in various fields, largely due to the use of well-curated datasets (Deng et al. 2009; Krizhevsky, Hinton et al. 2009). However, real-world data often exhibit a long-tail distribution (Ouyang et al. 2016; Zhang et al. 2021; Wei and Gan 2023; Li et al. 2024), where a few dominant classes (head classes) hold the majority of samples and most other classes (tail classes) have very few. Modern over-parameterized neural networks tend to overfit on head classes and underperform on tail classes due to inadequate training for the latter (Buda, Maki, and Mazurowski 2018; Van Horn and Perona 2017).

*Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Several methods have been developed to address this imbalance. Most of these approaches use re-weighting techniques, such as LDAM (Cao et al. 2019), logit-adjustment (Menon et al. 2021; Wang et al. 2023), and Vector Scaling (VS) Loss (Kini et al. 2021), to adjust class weights and focus on tail classes in the loss function. This adjustment significantly enhances model performance in datasets with long-tail distributions. However, these re-weighting techniques can still cause overfitting in tail classes (Hawkins 2004).

To address overfitting, many strategies involving smoothness regularization have been proposed (Kang et al. 2019; Alshammari et al. 2022). Sharpness-Aware Minimization (SAM) (Foret et al. 2021), known for its strong theoretical basis and excellent results, has garnered considerable attention. According to Rangwani et al. (2022), SAM effectively helps models escape saddle points, heading for a flat local minimum. Zhou et al. (2023) introduced Imbalanced SAM (ImbSAM) for long-tail learning, which bypasses the SAM optimization for head classes to prevent their dominance in the optimization process, showing notable performance improvements in long-tail scenarios. However, empirical studies indicate that ImbSAM significantly reduces the generalization ability of head classes compared to SAM. Thus, enhancing long-tail generalization requires a careful balance between head and tail classes.

To find a better solution, we first examine the properties of SAM and ImbSAM. We show that SAM theoretically has a greater component on the maximum negative curvature direction than Stochastic Gradient Descent (SGD), suggesting a stronger saddle escape capability. However, in long-tail distributions, head classes mainly influence the direction and magnitude of perturbations in SAM, often neglecting tail classes. In contrast, empirical studies show that ImbSAM reduces the misleading effects of head classes, helping tail classes escape saddles more effectively. However, it results in insufficient training of head classes.

From the analysis, SAM and ImbSAM offer complementary benefits, but it’s challenging to escape saddle points in both head and tail classes simultaneously. To tackle this, we introduce a new approach called Stage-wise Saddle Escaping SAM (SSE-SAM), designed to gradually overcome saddle points in both head and tail classes. SSE-SAM divides the training into two phases. Initially, it employs the small loss assumption to help head classes escape saddle points effec-




Figure 1: **Illustration of proposed SSE-SAM.** SSE-SAM splits the training process into two stages. The first stage helps head classes escape saddle points, while the second stage assists tail classes. The graphs on the right show an initial improvement in head class performance during the first stage, followed by a significant enhancement in tail class performance during the second stage. For detailed experimental results, see Sec..

tively. Later, it shifts attention to tail classes to facilitate their escape from saddle points. Our experiments demonstrate that SSE-SAM is more effective at escaping saddles in both head and tail classes.

Our contributions in this paper are summarized as follows:

- We theoretically show that SAM, compared with SGD, has a greater component along the maximum negative direction of the loss landscape. This finding demonstrates the more efficient saddle escaping ability of SAM and lays the foundation for the following analysis.
- We analyze the complementary characteristics of SAM and ImbSAM in saddle escaping ability. We observe the distinct effects of aiding escape from saddle points by SAM and ImbSAM, both theoretically and empirically.
- We propose a Stage-wise Saddle Escaping SAM (SSE-SAM) algorithm. SSE-SAM splits the training process into two stages of saddle escaping to achieve balance in head and tail classes. Our method achieves remarkable performance across multiple experiments.

Related Work

In this section, we briefly discuss related work here. More detailed discussion can be found in Appendix. In the end, we explain the notation conventions in this paper.

Long-tail Learning

Class re-balancing techniques are mainstream methods in long-tail learning (Zhang et al. 2023), including re-sampling (Chawla et al. 2002; Buda, Maki, and Mazurowski 2018), class-sensitive learning (Lin et al. 2017; Cao et al. 2019; Kini et al. 2021; Du et al. 2024), and logit adjustment (Menon et al. 2021; Wang et al. 2023). These strategies aim to provide a more balanced training environment, enhancing the representation and performance of tail classes.

Smoothness Regularization Methods

Smoothness regularization algorithms promote convergence to flatter minima, which can significantly improve generalization (Keskar et al. 2016; Jiang et al. 2019). A notable method

within this domain is Sharpness-Aware Minimization (SAM) (Foret et al. 2021), which has demonstrated notable success in guiding models toward smoother convergence and thus, better generalization. There are abundant theoretical (Woodworth et al. 2020; Andriushchenko and Flammarion 2022; Wen, Ma, and Li 2022; Wen, Li, and Ma 2023) and practical (Du et al. 2021; Kwon et al. 2021; Mi et al. 2022; Zhou et al. 2023) researches surrounding SAM, underscoring its broad applicability.

Escaping Saddle Points

Saddle points are regions in the loss landscape where the function’s gradient is zero but the Hessian matrix is indefinite (Lee 2018). Models converging to saddle points typically exhibit poor generalization (Dauphin et al. 2014). There are extensive theoretical studies on saddle point problem (Daneshmand et al. 2018; Jin et al. 2021; Rangwani et al. 2022; Hsieh et al. 2023), and many algorithms have been developed to efficiently escape from saddle points (Palaniappan and Bach 2016; Jin et al. 2017; Staib et al. 2019; Criscitiello and Boumal 2019; Zhang and Li 2021; Huang et al. 2022).

Preliminary

As we have mentioned, our method is a two-stage saddle escaping method based on SAM and ImbSAM algorithm. In this section, we first introduce the related concepts and algorithms.

Sharpness-Aware Minimization

Sharpness-Aware Minimization (SAM) (Foret et al. 2021) aims to regularize sharpness of loss landscape by finding the highest loss value in the neighborhood of current weight w controlled by ρ , then minimize such value:

$$\min_w \max_{\|\epsilon\| \leq \rho} \mathcal{L}(w + \epsilon). \quad (1)$$

The solution of this minimax problem is approximated by first order Taylor expansion, which gives $\epsilon \approx \rho \frac{\nabla \mathcal{L}(w)}{\|\nabla \mathcal{L}(w)\|}$. The SAM loss function is thus defined as follows:

$$\mathcal{L}^{\text{SAM}}(w) = \mathcal{L}(w + \epsilon) = \mathcal{L}\left(w + \rho \frac{\nabla \mathcal{L}(w)}{\|\nabla \mathcal{L}(w)\|}\right). \quad (2)$$

Imbalanced SAM

Imbalanced SAM (ImbSAM) is a variant of SAM tailored for long-tail distributions (Zhou et al. 2023). It corrects the bias towards dominant classes found in SAM by allowing data from less represented classes to contribute more to optimization during training.

The loss function of SAM can be decomposed as

$$\mathcal{L}^{\text{SAM}}(\mathbf{w}) = \mathcal{L}(\mathbf{w} + \epsilon) = \mathcal{L}_{\text{head}}(\mathbf{w} + \epsilon) + \mathcal{L}_{\text{tail}}(\mathbf{w} + \epsilon). \quad (3)$$

ImbSAM removes the perturbation term in $\mathcal{L}_{\text{head}}$ to focus the sharpness minimization on tail classes. Consequently, the loss function of ImbSAM is structured as follows:

$$\begin{aligned} \mathcal{L}^{\text{ImbSAM}}(\mathbf{w}) &= \mathcal{L}_{\text{head}}(\mathbf{w}) + \mathcal{L}_{\text{tail}}(\mathbf{w} + \epsilon_{\text{tail}}) \\ &= \mathcal{L}_{\text{head}}(\mathbf{w}) + \mathcal{L}_{\text{tail}}\left(\mathbf{w} + \rho \frac{\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})}{\|\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})\|}\right). \end{aligned} \quad (4)$$

Geometric Measurement of Escaping Saddle Points

We use two methods to measure the ability of models in escaping saddles.

From theoretical aspect, to analyze the saddle escaping process, we have to analyze the direction of maximum negative curvature¹ around saddles. Here we briefly summarize the differential geometry theories showing that the direction of the eigenvector corresponding to the smallest eigenvalue of Hessian around saddle point is the direction of maximum negative curvature.

Differential geometry suggests that the principal direction of maximum negative curvature in the loss landscape corresponds to the eigenvector associated with the most negative eigenvalue of the Weingarten matrix (Lee 2018) \mathbf{W} :

$$\mathbf{W} = \left(\mathbf{I} - \frac{\nabla \mathcal{L} \cdot \nabla \mathcal{L}^\top}{\alpha^2} \right) \frac{\nabla^2 \mathcal{L}}{\alpha}, \quad (5)$$

where $\alpha := \sqrt{1 + \|\nabla \mathcal{L}\|^2}$. Proof of such property can be found in Appendix.

Denote by \mathcal{H} the Hessian of \mathcal{L} at \mathbf{w}_0 , $\nabla^2 \mathcal{L}(\mathbf{w}_0)$. Also, denote λ_{\min} , \mathbf{v}_{\min} as the most negative eigenvalue and its corresponding eigenvector of \mathcal{H} . Since the gradient $\nabla \mathcal{L}(\mathbf{w}_0)$ vanishes near saddle points, \mathcal{H} effectively approximates \mathbf{W} . In this sense, \mathcal{H} is a good approximation of \mathbf{W} . Therefore \mathbf{v}_{\min} can effectively represent the direction of maximum negative curvature of loss landscape.

From empirical aspect, we focus on how well our model escapes from saddles. Following previous works (Li et al. 2018; Rangwani et al. 2022), we use $\lambda_{\text{ratio}} = |\lambda_{\max}/\lambda_{\min}|$ to measure the non-convexity of the loss landscape, where high λ_{ratio} indicates convergence to points with less negative curvature, therefore away from saddle points.

Roadmap

This paper is structured as follows. The next two sections provide a detailed analysis of the strengths and weaknesses

¹Here "maximum" means the absolute value reaches maximum among all negative curvatures, *i.e.*, the smallest curvature which is negative. This expression is used throughout this paper.

of SAM and ImbSAM in escaping saddle points. This examination reveals their complementary dynamics, leading to the development of a stagewise approach, presented in the following section. Experimental evidence to validate the efficacy of our proposed SSE-SAM algorithm is then presented in the subsequent section. The final section of the paper offers concluding remarks. All proofs are deferred to Appendix.

Analysis of SAM in Escaping Saddle Points

Advantage: SAM Escapes from Saddles More Efficiently Compared with SGD

As mentioned in previous discussions, the theoretical power of SAM comes partially from its ability to escape from saddle points and heading to a flat local/global minimum. Inspired by this, we first delve deeper into the theoretical properties of SAM's saddle point escaping ability.

Building upon the foundation laid in the preceding section on geometric measurement of escaping saddle points, we evaluate this capability by examining how close the learned weight is to the subspace associated with the most negative curvature. Moreover, we only need to examine it around the saddle point, where the Weingarten matrix roughly equals to the Hessian, *i.e.*, $\mathbf{W} \approx \nabla^2 \mathcal{L}$. The most negative curvature subspace is roughly the subspace spanned by the eigenvector associated with the minimum eigenvalue of the Hessian λ_{\min} ². In this sense, a greater component on this subspace implies a more effective direction for descent around the saddle. Therefore, we compare the update directions of SAM and ImbSAM by examining the inner product between these directions and the eigenvector. The result is shown in the following theorem.

Theorem 1. *Let $\mathcal{W} := \{\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_t, \tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_t\}$, where \mathbf{w}_0 is the starting weight, \mathbf{w}_t denotes the weight obtained after optimizing for t steps using SGD, $\tilde{\mathbf{w}}_t$ denotes the weight obtained after optimizing for t steps using SAM. All weights in \mathcal{W} are around saddle point. Let $L := \max_{\mathbf{w} \in \mathcal{W}} \|\nabla \mathcal{L}(\mathbf{w})\| > 0$. Given any $\mu \geq 1$, when $\lambda_{\min} \leq -\frac{L}{\eta\rho} \left(\frac{1+\mu}{2}\eta + \sqrt{\left(\frac{1+\mu}{2}\right)^2 \eta^2 + \frac{(2+\mu)\eta\rho}{L}} \right)$, under second order Taylor approximation, we have*

$$\begin{aligned} |\langle \mathbf{v}_{\min}, \tilde{\mathbf{w}}_t - \mathbf{w}_0 \rangle| &\geq \mu |\langle \mathbf{v}_{\min}, \mathbf{w}_t - \mathbf{w}_0 \rangle| \\ &+ ((\mu^2 + \mu)t - 2\mu^2 - \mu) \frac{|\langle \mathbf{v}_{\min}, \nabla \mathcal{L}(\mathbf{w}_0) \rangle|}{|\lambda_{\min}|}. \end{aligned} \quad (6)$$

For all $\mu \geq 1$, when $t \geq 2$, we have that $(\mu^2 + \mu)t - 2\mu^2 - \mu \geq \mu > 0$. Therefore, SAM leads to a component on \mathbf{v}_{\min} at least μ times greater than SGD around the saddle point, showing a better potential to find a descent path toward a smooth and flat local minimum (SAM naturally punishes sharp local minimums through its objective function). Above all, SAM facilitates the saddle point escaping ability of the optimization algorithm and leads to better generalization by sharpness regularization.

² λ_{\min} is mostly negative as shown in the experiments.

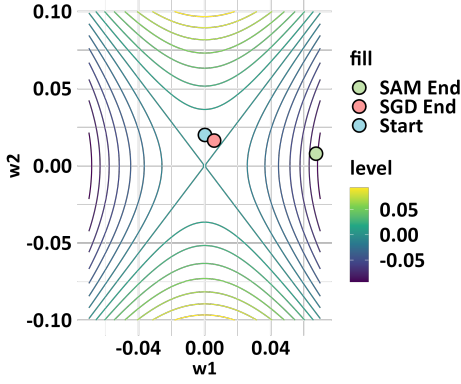


Figure 2: Illustration of Thm.1.

Fig.2 intuitively illustrates Thm.1. Starting from a position close to a saddle point, SAM moves substantially more towards the direction of \mathbf{v}_{\min} than SGD, showing better saddle escaping capability.

Remark 2. Rangwani et al. (2022) also establish a saddle point escaping theory for SAM. However, their result relies on the Correlated Negative Curvature (CNC) assumption (Daneshmand et al. 2018), which assumes the existence of $\gamma > 0$ such that $\mathbb{E}[\langle \mathbf{v}_{\min}, \nabla L(\mathbf{w}) \rangle^2] \geq \gamma$. If CNC holds, they show that SAM will increase the second moment $\mathbb{E}[\langle \mathbf{v}_{\min}, \nabla L(\mathbf{w}) \rangle^2]$. However, this assumption is not typically met near saddle points, where $\nabla L(\mathbf{w})$ approaches zero. Moreover, the CNC assumption is based on a baseline called CNC-PGD algorithm, which does not align with the settings in standard SGD (Daneshmand et al. 2018). Compared with this existing work, our result is much more general, since it not only get rid of the CNC assumption but also captures the dynamic trajectory after t steps. As t increases, the different behavior between SAM and SGD becomes increasingly apparent.

Limitation: Suboptimal Optimization for Tail Classes in Escaping Saddles

The aforementioned analysis is based on the overall loss landscape. When it comes to long-tail distribution, such theories are no longer validated. Given that head classes vastly outnumber tail classes in such datasets, the shape of the loss surface is predominantly influenced by head classes, effectively overshadowing the characteristics of tail classes which remain underexplored. Zhou et al. (2023) empirically show that while optimization process of SAM significantly benefits head classes, it leads to suboptimal optimization direction of tail classes due to the overwhelming existence of head classes.

Delving deeper into this problem, we provide a formal theoretical analysis of such suboptimality.

Decomposing loss function of SAM as $\mathcal{L}^{\text{SAM}}(\mathbf{w}) = \mathcal{L}(\mathbf{w} + \epsilon) = \mathcal{L}_{\text{head}}(\mathbf{w} + \epsilon) + \mathcal{L}_{\text{tail}}(\mathbf{w} + \epsilon)$, we have

$$\epsilon = \underbrace{\rho \cdot \frac{\nabla \mathcal{L}_{\text{head}}(\mathbf{w})}{\|\nabla \mathcal{L}(\mathbf{w})\|}}_{\epsilon_{\text{head}}} + \underbrace{\rho \cdot \frac{\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})}{\|\nabla \mathcal{L}(\mathbf{w})\|}}_{\epsilon_{\text{tail}}}. \quad (7)$$

Because $|S_{\text{head}}| \gg |S_{\text{tail}}|$, we have

$$\mathcal{L}_{\text{head}} = \sum_{(\mathbf{x}_i, y_i) \in S_{\text{head}}} \ell_i(\mathbf{w}) \gg \mathcal{L}_{\text{tail}} = \sum_{(\mathbf{x}_i, y_i) \in S_{\text{tail}}} \ell_i(\mathbf{w}), \quad (8)$$

where ℓ_i is the loss for single sample (\mathbf{x}_i, y_i) . In fact, we have the following propositions:

Proposition 3. Assume that the samples are drawn from i.i.d distribution, let $m = |S_{\text{head}}|, n = |S_{\text{tail}}|$ then only if $\frac{\xi_{\text{head}}}{\xi_{\text{tail}}} = \Theta(\frac{n^2}{m^2})$, we have

$$\frac{\mathbb{E}_{\text{head}}[\|\nabla \mathcal{L}(\mathbf{w})\|^2]}{\mathbb{E}_{\text{tail}}[\|\nabla \mathcal{L}(\mathbf{w})\|^2]} = \Theta(1). \quad (9)$$

where

$$\xi_{\text{head}} = \|\boldsymbol{\mu}^{(1)}\|^2 + \sum_{j=1}^d \sigma_j^{(1)2}, \quad \xi_{\text{tail}} = \|\boldsymbol{\mu}^{(2)}\|^2 + \sum_{j=1}^d \sigma_j^{(2)2},$$

$\boldsymbol{\mu}^{(1)}, \boldsymbol{\mu}^{(2)}$ are the mean of instance gradient $\nabla \ell_i(\mathbf{w})$ for head/tail class, $\sigma_j^{(1)2}, \sigma_j^{(2)2}$ are the variance of the j -th dimension of the instance gradient $\nabla \ell_i(\mathbf{w})$ for head/tail class.

For long-tail distribution, we have $m^2 \gg n^2$, since it is almost impossible to observe comparable gradients for head and tail distribution. In this sense, we have $\epsilon_{\text{head}} \gg \epsilon_{\text{tail}}$ the optimization will overlook the update for the tail-class. Furthermore, the following proposition shows the closeness between ϵ_{head} and ϵ in terms of their angles.

Proposition 4. Denote by θ_{head} the angle between ϵ_{head} and ϵ , θ_{tail} the angle between ϵ_{tail} and ϵ , ψ the angle between ϵ_{head} and ϵ_{tail} . When $\|\epsilon_{\text{head}}\| \gg \|\epsilon_{\text{tail}}\|$, we have

$$\begin{aligned} \cos \theta_{\text{head}} &\approx 1, \\ \cos \theta_{\text{tail}} &\approx \cos \psi. \end{aligned} \quad (10)$$

Therefore, both in terms of direction and magnitude, ϵ_{head} is close to ϵ . **The smoothness regularization direction for tail classes is dominated by head classes, leading to suboptimal regularization trajectory for tail classes.** Consequently, this results in less effective generalization enhancements for tail classes compared to head classes in SAM.

Analysis of ImbSAM in Escaping Saddles

Advantage: Efficiently Escaping Saddle Points on Tail Classes

Previous section shows that SAM fails to assist tail classes in escaping saddle points. To address this issue, ImbSAM propose a simple and effective solution, *i.e.*, removing the perturbation term ϵ_{head} in SAM, resulting in a new objective $\mathcal{L}^{\text{ImbSAM}} = \mathcal{L}_{\text{head}}(\mathbf{w}) + \mathcal{L}_{\text{tail}}\left(\mathbf{w} + \rho \frac{\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})}{\|\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})\|}\right)$.

Previous studies have already shown the effectiveness of ImbSAM in long-tail distribution (Zhou et al. 2023). The benefit of only punishing tail-class sharpness seems straightforward, because it directly eliminate the influence of the head class. Our empirical study further confirms this inference. As is shown in Fig.4, ImbSAM has larger λ_{ratio} on tail classes compared with SGD and SAM, rendering it better at escaping saddles on tail classes and improving generalization ability.

Limitation: Insufficient Optimization of Head Classes

Although ImbSAM achieves better performance on tail classes, its effectiveness in head classes still needs validation. To this end, we record the head class performance in Tab.1, showing that ImbSAM underperforms SAM in these classes. This shortfall stems from inadequate optimization of head classes. Due to the elimination of ϵ_{head} in head classes, ImbSAM actually has worse saddle escaping performance on head classes compared with SAM. Empirical Hessian analysis offers more general and intuitive illustration of the degenerate saddle escaping ability of ImbSAM. See Fig.4, where a higher λ_{ratio} indicates better saddle escaping ability. ImbSAM works well on tail classes (Fig.4c), but has unsatisfactory results on head classes (Fig.4b) and on a global scale (Fig.4a). Finally, based on our theoretical analysis in Thm.1, we offer a counter-example that ImbSAM shows degenerate saddle point escaping capabilities.

Proposition 5. Consider a special case where $\mathcal{H}_{\text{tail}} = r\mathcal{H}$, $\mathcal{H}_{\text{head}} = (1-r)\mathcal{H}$ with $r = \frac{|\mathcal{S}_{\text{tail}}|}{|\mathcal{S}|} \approx 0$. Let \mathbf{w}'_t denote the weight obtained after optimizing for t steps by ImbSAM. Other variables are defined the same as in Thm.1. For any $\mu \geq 1$, we have that:

(1) when $|\langle \mathbf{v}_{\min}, \mathbf{w}'_t - \mathbf{w}_0 \rangle| < \mu |\langle \mathbf{v}_{\min}, \mathbf{w}_t - \mathbf{w}_0 \rangle|$, as long as $\lambda_{\min} \leq -\frac{L}{\eta\rho} \left(\frac{1+\mu}{2}\eta + \sqrt{\left(\frac{1+\mu}{2}\right)^2 \eta^2 + \frac{(2+\mu)\eta\rho}{L}} \right)$, it holds that $|\langle \mathbf{v}_{\min}, \tilde{\mathbf{w}}_t - \mathbf{w}_0 \rangle| \geq \mu |\langle \mathbf{v}_{\min}, \mathbf{w}_t - \mathbf{w}_0 \rangle|$;

(2) when $|\langle \mathbf{v}_{\min}, \tilde{\mathbf{w}}_t - \mathbf{w}_0 \rangle| < \mu |\langle \mathbf{v}_{\min}, \mathbf{w}_t - \mathbf{w}_0 \rangle|$, it also holds that $|\langle \mathbf{v}_{\min}, \mathbf{w}'_t - \mathbf{w}_0 \rangle| < \mu |\langle \mathbf{v}_{\min}, \mathbf{w}_t - \mathbf{w}_0 \rangle|$.

Prop.5(1) states that, even when ImbSAM has worse saddle escaping ability, i.e., $|\langle \mathbf{v}_{\min}, \mathbf{w}'_t - \mathbf{w}_0 \rangle| < \mu |\langle \mathbf{v}_{\min}, \mathbf{w}_t - \mathbf{w}_0 \rangle|$, under mild conditions, SAM still can have better saddle escaping ability. However, according to Prop.5(2), when SAM has worse saddle escaping ability, ImbSAM cannot have better saddle escaping ability as well. Collectively, these observations indicate that ImbSAM has a reduced ability to escape from saddles compared to SAM.

Stagewise Saddle Escaping SAM (SSE-SAM) Saddle Point Escaping and Generalization

Before introducing the proposed algorithm, we first analyze the connection between saddle point escaping and generalization. We have the following theorem.

Theorem 6. Suppose loss function \mathcal{L} is upper bounded by M . For any $\rho > 0$ and any distribution \mathcal{D} , with probability at least $1 - \delta$ over the choice of the training set $S \sim \mathcal{D}$, there exists a constant $0 \leq c \leq 1$ such that

$$\begin{aligned} \mathcal{L}_{\mathcal{D}}(\mathbf{w}) &\leq \mathcal{L}_S(\mathbf{w}) + \rho^2 \sqrt{\frac{d}{4\pi}} \max_{\|\epsilon\| \leq \rho} \lambda_{\max}(\nabla^2 L(\mathbf{w} + c\epsilon)) \\ &+ \frac{M}{\sqrt{n}} + \left[\left(\frac{1}{4} d \log \left(1 + \frac{\|\mathbf{w}\|^2 (\sqrt{d} + \sqrt{\log n})^2}{d\rho^2} \right) \right) + \frac{1}{4} \right. \\ &\left. + \log \frac{n}{\delta} + 2 \log(6n + 3d) \right] / (n-1)^{\frac{1}{2}}, \end{aligned} \quad (11)$$

where d is the number of parameters and n is the size of training set S .

Note that the upper bound can be written as

$$\underbrace{\mathcal{L}_S(\mathbf{w})}_{(A)} + \underbrace{\rho^2 \sqrt{\frac{d}{4\pi}} \max_{\|\epsilon\| \leq \rho} \lambda_{\max}(\nabla^2 \mathcal{L}(\mathbf{w} + c\epsilon))}_{(B)} + r \left(\frac{\|\mathbf{w}\|^2}{\rho^2}, \tilde{O}\left(\frac{1}{\sqrt{n}}\right) \right). \quad (12)$$

Here r is an increasing function on $\|\mathbf{w}\|^2/\rho^2$ and is of order $\tilde{O}(1/\sqrt{n})$. The success of SAM can be viewed as a two-stage process. By escaping from saddle point along the most negative eigenvector \mathbf{v}_{\min} , (A) is consistently minimized while (B) remains stable. After escaping from saddles, SAM seeks for smooth local/global minima. (Wen, Li, and Ma 2023) proves that SAM minimizes λ_{\max} , which corresponds to term (B). In all, SAM escapes from saddle point and finds smooth local minima, minimizing both (A) and (B), which greatly improves generalization. This result bridges the gap between escaping from saddle points and improving model generalization, and dynamically analyzes the behavior of SAM algorithm, which lays foundation for the following algorithm.

Stagewise Saddle Escaping SAM (SSE-SAM) Algorithm

Previous sections analyze the complementary characteristics of SAM and ImbSAM. To sum up, SAM excels at saddle points escaping in the overall loss landscape, but fails in helping tail classes to do so. Conversely, while ImbSAM does focus on helping tail classes escape saddles, the elimination of head classes perturbation also hurts the generalization ability of head classes.

Therefore, an important question arises: *can we combine the advantages of both while avoiding their drawbacks?* Integrating these approaches seems challenging due to the apparent contradictions in saddle escaping abilities between SAM and ImbSAM. Nonetheless, recognizing that it is difficult to achieve it under a single objective, we propose a two-stage strategy tailored to exploit the strengths of each method while compensating for their weaknesses:

(1) **In the first stage, we apply SAM with a focus on head classes to ensure a rapid and effective escape from saddle points.** According to small loss assumption (Xia et al. 2022), the **head-classes**, often bear way smaller loss value, can naturally follow the trajectory to **escape the saddle points**. Meanwhile, increasing ρ lowers the upper bound for λ_{\min} in Thm.1, facilitating an easier saddle escape. Hence, we construct ρ and ϵ separately for head- and tail-classes. By increasing ρ_{tail} , it is easier for tail-class to escape saddle points. Overall, the SSE-SAM objective for the first stage is:

$$\mathcal{L}_1^{\text{SSE-SAM}}(\mathbf{w}) = \mathcal{L}_{\text{head}}(\mathbf{w} + \epsilon_{\text{head}}) + \mathcal{L}_{\text{tail}}(\mathbf{w} + \epsilon_{\text{tail}}) \quad (13)$$

where $\epsilon_{\text{head}} = \rho_{\text{head}} \frac{\nabla \mathcal{L}_{\text{head}}(\mathbf{w})}{\|\nabla \mathcal{L}_{\text{head}}(\mathbf{w})\|}$, $\epsilon_{\text{tail}} = \rho_{\text{tail}} \frac{\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})}{\|\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})\|}$.

(2) **In the second stage, we resort to the ImbSAM algorithm to further assist tail classes to escape saddles.** The

SSE-SAM objective for the second stage is:

$$\mathcal{L}_2^{\text{SSE-SAM}}(\mathbf{w}) = \mathcal{L}_{\text{head}}(\mathbf{w}) + \mathcal{L}_{\text{tail}}(\mathbf{w} + \epsilon_{\text{tail}}) \quad (14)$$

where $\epsilon_{\text{tail}} = \rho_{\text{head}} \frac{\nabla \mathcal{L}_{\text{tail}}(\mathbf{w})}{\|\nabla \mathcal{L}(\mathbf{w})\|}$.

(3) **We introduce a hyperparameter γ (where $0 < \gamma < 1$) to mark the transition point between the two stages.** Specifically, if the total training duration is T epochs, the model will switch from the first stage objective to the second stage objective after γT epochs.

For further details and pseudo-code, please refer to Appendix. Note that when $\gamma = 1$, SSE-SAM reverts to the SAM algorithm; when $\gamma = 0$, SSE-SAM corresponds to the ImbSAM algorithm. This demonstrates that γ serves as a pivotal factor in balancing the synergistic benefits of both SAM and ImbSAM.

Experiments

In this section, we first present our experimental results, demonstrating the remarkable performance of our method. Next, we provide a Hessian analysis to compare saddle escaping abilities, which aligns with our theoretical analysis. More additional experiment details are presented in Appendix.

Results

Our experiment results on CIFAR-100-LT and CIFAR-10-LT are presented in Tab.1. All hyperparameters are set according to the model with the highest *overall accuracy*: $\rho = 0.05$ for SAM, $\rho = 0.10$ for ImbSAM, and $\rho_{\text{head}} = 0.05, \rho_{\text{tail}} = 0.10, \gamma = 0.70$ for SSE-SAM. SSE-SAM achieves the highest overall accuracy among all methods and baselines. We also run our model on large scale dataset ImageNet, where we set $\rho = 0.05$ for SAM, $\rho = 0.15$ for ImbSAM and $\rho_{\text{head}} = 0.05, \rho_{\text{tail}} = 0.15, \gamma = 0.80$ for SSE-SAM. The results can be found in Tab.2.

To further demonstrate the improved trade-off effect between head and tail classes, we trained two groups of models. Each group consists of three models: SAM (ext.)³, ImbSAM, and our SSE-SAM. SSE-SAM and SAM share ρ_{head} , and all three models share ρ_{tail} . The results are presented in Tab.3. In each group, the evaluation results for Many class accuracy generally follow the pattern “**SAM > SSE-SAM > ImbSAM**”, indicating that SSE-SAM indeed reduces the insufficient training of head classes to some extent, aligning with our expectations. Meanwhile, Medium and Few class accuracy typically adhere to the pattern “**SSE-SAM > ImbSAM > SAM**”, suggesting that SSE-SAM more effectively escapes saddle points in tail classes. Additionally, it is notable that SSE-SAM generally outperforms ImbSAM in tail classes, exceeding our expectations.

Finally, we delve into the dynamic training process of SAM, ImbSAM, and our SSE-SAM. The results are presented in Fig.3. The data indicates that after the transitional point, the accuracy for medium and few classes in SSE-SAM exceeds that of ImbSAM, while the accuracy for many classes remains between that observed with SAM and ImbSAM.

³We use SAM (ext.) to denote models trained with the extended version of SAM, *i.e.*, the first stage objective in SSE-SAM.

	CIFAR-100-LT				CIFAR-10-LT			
	Many	Med.	Few	Avg	Many	Med.	Few	Avg
CE	74.3	43.1	11.6	44.6	92.3	78.6	54.8	77.0
+SAM	75.3	45.6	12.9	46.2	93.2	79.8	60.4	79.4
+ImbSAM	72.3	50.4	16.6	47.9	92.2	79.2	65.8	80.4
+SSE-SAM	71.7	54.5	17.1	49.3	93.0	80.5	72.8	83.2
LDAM	74.3	48.4	17.7	48.3	90.8	76.0	61.5	77.6
+SAM	76.3	49.6	19.4	49.9	91.2	77.4	62.6	78.5
+ImbSAM	72.0	55.4	23.3	51.6	89.1	78.9	71.7	80.8
+SSE-SAM	72.3	54.7	21.5	50.9	91.1	81.8	69.7	81.9
LA	75.0	45.7	16.2	47.1	91.8	80.0	65.4	80.4
+SAM	75.2	49.9	21.2	50.2	92.3	80.8	73.3	83.1
+ImbSAM	68.7	54.1	23.5	50.0	91.1	80.5	74.0	82.8
+SSE-SAM	68.8	58.7	28.6	53.2	92.5	81.1	79.2	85.1
VS	73.5	46.7	19.3	47.8	92.0	80.1	65.1	80.4
+SAM	73.0	53.9	24.6	51.8	89.9	81.2	81.9	84.9
+ImbSAM	68.9	55.2	25.4	51.0	90.9	80.0	76.6	83.3
+SSE-SAM	69.7	57.2	30.5	53.5	91.6	82.6	79.0	85.1

Table 1: **Comparison of overall and split accuracy (%) on CIFAR-100-LT and CIFAR-10-LT.** ‘Med.’ denotes Medium classes, ‘Avg’ denotes overall accuracy. We set IF = 100.

	Many	Med.	Few	Avg.
CB	36.9	32.7	16.8	33.2
τ -norm	59.1	46.9	30.7	49.4
cRT	62.5	47.4	29.5	50.3
LWS	61.8	48.6	33.5	51.2
DisAlign	61.3	52.2	31.4	52.9
DRO-LT	64.0	49.8	33.1	53.5
CE	54.8	43.8	25.7	45.1
CE+SAM	54.9	44.8	28.4	46.0
CE+ImbSAM	54.9	44.6	29.2	46.1
CE+SSE-SAM	56.0	46.0	31.2	47.5

Table 2: Comparison of overall and split accuracy (%) on ImageNet-LT.

Overall, the all-class accuracy of SSE-SAM experiences further improvement upon entering the second stage of training. This enhancement solidifies the superior performance of SSE-SAM relative to the other methods.

Hessian Analysis

We utilize λ_{ratio} , as described in Preliminary section, to evaluate convexity. We analyze the λ_{ratio} for all classes, head classes, and tail classes across models trained using SGD, SAM ($\rho = 0.05$), ImbSAM ($\rho = 0.10$), SAM (ext.) ($\rho_{\text{head}} = 0.05, \rho_{\text{tail}} = 0.10$), and our SSE-SAM ($\rho_{\text{head}} = 0.05, \rho_{\text{tail}} = 0.10, \gamma = 0.7$). The outcomes are illustrated in Fig.4. Fig.4a distinctly demonstrates that SSE-SAM shows superior overall performance compared to the other methods. Specifically, for head classes, we analyze λ_{ratio} for class 0, where Fig.4b shows that SSE-SAM outperforms all other methods. In the case of class 80, which represents tail classes,

Models	ρ_{head}	ρ_{tail}	IF=100				IF=50			
			Many	Med.	Few	Avg	Many	Med.	Few	Avg
SGD	0.00	0.00	74.3	43.1	11.6	44.6	73.6	41.4	15.4	49.9
SAM	0.05	0.05	75.3	45.6	12.9	46.2	74.7	43.9	19.1	52.1
ImbSAM	0.00	0.05	73.0	<u>50.1</u>	<u>15.2</u>	<u>47.6</u>	71.4	<u>48.5</u>	23.2	<u>53.3</u>
SSE-SAM	0.05	0.05	<u>73.1</u>	51.3	15.8	48.3	<u>73.4</u>	49.7	<u>22.4</u>	54.5
SAM (ext.)	0.05	0.10	75.2	48.6	14.9	47.8	74.4	46.4	21.3	53.4
ImbSAM	0.00	0.10	<u>72.3</u>	50.4	<u>16.6</u>	<u>47.9</u>	70.9	<u>50.2</u>	<u>24.1</u>	<u>54.0</u>
SSE-SAM	0.05	0.10	71.7	54.5	17.1	49.3	<u>72.2</u>	51.9	26.1	55.6

Table 3: Comparison of overall and split accuracy (%) of different IFs and ρ 's. We set $\gamma = 0.70$ for SSE-SAM models. We use CE as loss function and train all models on CIFAR-100-LT.

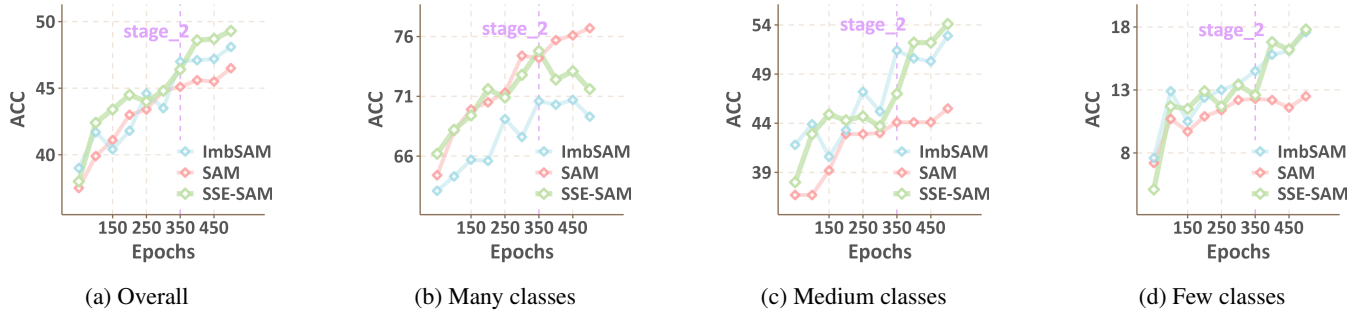


Figure 3: Test accuracy on CIFAR-100-LT across different classes.

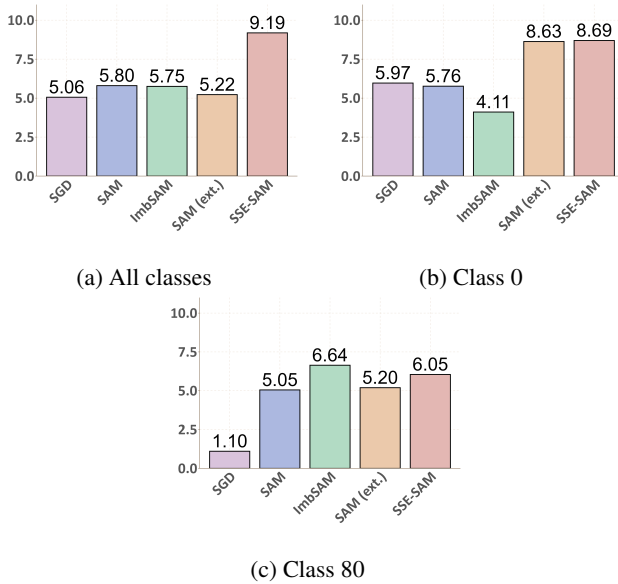


Figure 4: Comparative Analysis of λ_{ratio} for Different Models. These results validate SSE-SAM's effective balance between head and tail classes compared with SAM and ImbSAM.

Fig.4c reveals that SSE-SAM's capacity to escape saddle points lies between ImbSAM and SAM (ext.), all sharing the

same ρ_{tail} . Concerning ρ_{head} , our analysis indicates that the predominance of head classes results in suboptimal escape from saddles in tail classes. Here, the ρ_{head} of ImbSAM is 0, indicating an absence of suppressive effect from head classes; the ρ_{head} of SAM (ext.) is 0.05, displaying the most significant suppressive effect; the ρ_{head} of SSE-SAM transitions from 0.05 to 0, reflecting a suppressive impact that lies between the other two methods. These findings corroborate our theoretical expectations.

Conclusion

In this work, we theoretically show that SAM has better ability in escaping saddle points compared with ImbSAM. Under long-tail distribution, we analyze that SAM is good at escaping saddles for head classes but fails to efficiently escape saddles for tail classes. ImbSAM, on the contrary, focuses on escaping saddles for tail classes while suffers from insufficient training of head classes. Leveraging this complementary effect, we propose a two-stage saddle escaping algorithm SSE-SAM. At the first stage, SSE-SAM uses SAM to help head classes escape saddles. At the next stage, SSE-SAM sets $\rho_{\text{head}} = 0$ to focus on helping tail classes in escaping saddles. Experiment results show that SSE-SAM escapes saddles efficiently both on head and tail classes, achieving outstanding performance among all compared baselines.

Acknowledgements

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, U23B2051, 61931008, 62122075, 62206264, and 92370102, in part by Youth Innovation Promotion Association CAS, in part by the Strategic Priority Research Program of the Chinese Academy of Sciences, Grant No. XDB0680000, in part by the Innovation Funding of ICT, CAS under Grant No.E000000, in part by the Tencent Marketing Solution Rhino-Bird Focused Research Program.

References

- Alshammari, S.; Wang, Y.-X.; Ramanan, D.; and Kong, S. 2022. Long-tailed recognition via weight balancing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6897–6907.
- Andriushchenko, M.; and Flammarion, N. 2022. Towards understanding sharpness-aware minimization. In *International Conference on Machine Learning*, 639–668.
- Buda, M.; Maki, A.; and Mazurowski, M. A. 2018. A systematic study of the class imbalance problem in convolutional neural networks. *Neural networks*, 106: 249–259.
- Cao, K.; Wei, C.; Gaidon, A.; Arechiga, N.; and Ma, T. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 1567–1578.
- Chawla, N. V.; Bowyer, K. W.; Hall, L. O.; and Kegelmeyer, W. P. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16: 321–357.
- Criscitello, C.; and Boumal, N. 2019. Efficiently escaping saddle points on manifolds. In *Advances in Neural Information Processing Systems*, 5985–5995.
- Daneshmand, H.; Kohler, J.; Lucchi, A.; and Hofmann, T. 2018. Escaping saddles with stochastic gradients. In *International Conference on Machine Learning*, 1155–1164.
- Dauphin, Y. N.; Pascanu, R.; Gulcehre, C.; Cho, K.; Ganguli, S.; and Bengio, Y. 2014. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *Advances in Neural Information Processing Systems*, 2933–2941.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Du, C.; Wang, Y.; Song, S.; and Huang, G. 2024. Probabilistic Contrastive Learning for Long-Tailed Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Du, J.; Yan, H.; Feng, J.; Zhou, J. T.; Zhen, L.; Goh, R. S. M.; and Tan, V. Y. 2021. Efficient sharpness-aware minimization for improved training of neural networks. In *International Conference on Learning Representations*.
- Foret, P.; Kleiner, A.; Mobahi, H.; and Neyshabur, B. 2021. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*.
- Hawkins, D. M. 2004. The problem of overfitting. *Journal of chemical information and computer sciences*, 44(1): 1–12.
- Hsieh, Y.-P.; Karimi Jaghargh, M. R.; Krause, A.; and Mertikopoulos, P. 2023. Riemannian stochastic optimization methods avoid strict saddle points. In *Advances in Neural Information Processing Systems*.
- Huang, M.; Chen, X.; Ji, K.; Ma, S.; and Lai, L. 2022. Efficiently escaping saddle points in bilevel optimization. *arXiv preprint arXiv:2202.03684*.
- Jiang, Y.; Neyshabur, B.; Mobahi, H.; Krishnan, D.; and Bengio, S. 2019. Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*.
- Jin, C.; Ge, R.; Netrapalli, P.; Kakade, S. M.; and Jordan, M. I. 2017. How to escape saddle points efficiently. In *International Conference on Machine Learning*, 1724–1732.
- Jin, C.; Netrapalli, P.; Ge, R.; Kakade, S. M.; and Jordan, M. I. 2021. On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points. *Journal of the ACM (JACM)*, 68(2): 1–29.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2019. Decoupling Representation and Classifier for Long-Tailed Recognition. In *International Conference on Learning Representations*.
- Keskar, N. S.; Mudigere, D.; Nocedal, J.; Smelyanskiy, M.; and Tang, P. T. P. 2016. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. In *International Conference on Learning Representations*.
- Kini, G. R.; Paraskevas, O.; Oymak, S.; and Thrampoulidis, C. 2021. Label-imbalanced and group-sensitive classification under overparameterization. In *Advances in Neural Information Processing Systems*, 18970–18983.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Kwon, J.; Kim, J.; Park, H.; and Choi, I. K. 2021. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*, 5905–5914.
- Lee, J. M. 2018. *Introduction to Riemannian manifolds*, volume 2. Springer.
- Li, H.; Xu, Z.; Taylor, G.; Studer, C.; and Goldstein, T. 2018. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, 6391–6401.
- Li, W.; Lyu, F.; Shang, F.; Wan, L.; and Feng, W. 2024. Long-Tailed Learning as Multi-Objective Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3190–3198.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 2980–2988.
- Menon, A. K.; Jayasumana, S.; Rawat, A. S.; Jain, H.; Veit, A.; and Kumar, S. 2021. Long-tail learning via logit adjustment. In *International Conference on Learning Representations*.

- Mi, P.; Shen, L.; Ren, T.; Zhou, Y.; Sun, X.; Ji, R.; and Tao, D. 2022. Make sharpness-aware minimization stronger: A sparsified perturbation approach. In *Advances in Neural Information Processing Systems*, 30950–30962.
- Ouyang, W.; Wang, X.; Zhang, C.; and Yang, X. 2016. Factors in finetuning deep model for object detection with long-tail distribution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 864–873.
- Palaniappan, B.; and Bach, F. 2016. Stochastic variance reduction methods for saddle-point problems. In *Advances in Neural Information Processing Systems*, 1408–1416.
- Rangwani, H.; Aithal, S. K.; Mishra, M.; et al. 2022. Escaping saddle points for effective generalization on class-imbalanced data. In *Advances in Neural Information Processing Systems*, 22791–22805.
- Staib, M.; Reddi, S.; Kale, S.; Kumar, S.; and Sra, S. 2019. Escaping saddle points with adaptive gradient methods. In *International Conference on Machine Learning*, 5956–5965.
- Van Horn, G.; and Perona, P. 2017. The devil is in the tails: Fine-grained classification in the wild. *arXiv preprint arXiv:1709.01450*.
- Wang, Z.; Xu, Q.; Yang, Z.; He, Y.; Cao, X.; and Huang, Q. 2023. A Unified Generalization Analysis of Re-Weighting and Logit-Adjustment for Imbalanced Learning. In *Advances in Neural Information Processing Systems*, 48417–48430.
- Wei, T.; and Gan, K. 2023. Towards realistic long-tailed semi-supervised learning: Consistency is all you need. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3469–3478.
- Wen, K.; Li, Z.; and Ma, T. 2023. Sharpness minimization algorithms do not only minimize sharpness to achieve better generalization. In *Advances in Neural Information Processing Systems*, 1024–1035.
- Wen, K.; Ma, T.; and Li, Z. 2022. How does sharpness-aware minimization minimize sharpness? In *International Conference on Learning Representations*.
- Woodworth, B.; Gunasekar, S.; Lee, J. D.; Moroshko, E.; Savarese, P.; Golan, I.; Soudry, D.; and Srebro, N. 2020. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, 3635–3673.
- Xia, X.; Liu, T.; Han, B.; Gong, M.; Yu, J.; Niu, G.; and Sugiyama, M. 2022. Sample Selection with Uncertainty of Losses for Learning with Noisy Labels. In *International Conference on Learning Representations*.
- Zhang, C.; and Li, T. 2021. Escape saddle points by a simple gradient-descent based algorithm. In *Advances in Neural Information Processing Systems*, 8545–8556.
- Zhang, S.; Li, Z.; Yan, S.; He, X.; and Sun, J. 2021. Distribution alignment: A unified framework for long-tail visual recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2361–2370.
- Zhang, Y.; Kang, B.; Hooi, B.; Yan, S.; and Feng, J. 2023. Deep long-tailed learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9): 10795–10816.
- Zhou, Y.; Qu, Y.; Xu, X.; and Shen, H. 2023. ImbSAM: A Closer Look at Sharpness-Aware Minimization in Class-Imbalanced Recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, 11345–11355.