

Research Article

# TAET: Two-Stage Adversarial Equalization Training on Long-Tailed Distributions

Yu-Hang Wang<sup>1</sup>, Junkang Guo<sup>1</sup>, Aolei Liu<sup>1</sup>, Kaihao Wang<sup>1</sup>, Zaitong Wu<sup>1</sup>, Zhenyu Liu<sup>1</sup>, Wenfei Yin<sup>1</sup>, Jian Liu<sup>1</sup>

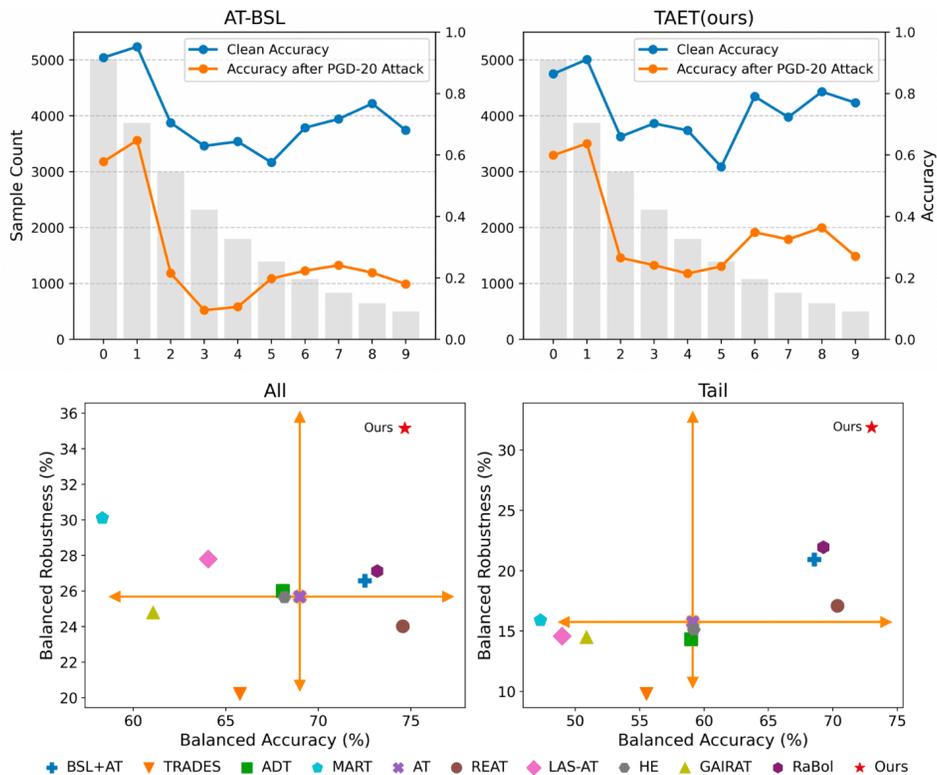
1. School of Computer Science and Information Engineering, Hefei University of Technology, China

Adversarial robustness remains a significant challenge in deploying deep neural networks for real-world applications. While adversarial training is widely acknowledged as a promising defense strategy, most existing studies primarily focus on balanced datasets, neglecting the fact that real-world data often exhibit a long-tailed distribution, which introduces substantial challenges to robustness. In this paper, we provide an in-depth analysis of adversarial training in the context of long-tailed distributions and identify the limitations of the current state-of-the-art method, AT-BSL, in achieving robust performance under such conditions. To address these challenges, we propose a novel training framework, TAET, which incorporates an initial stabilization phase followed by a stratified, equalization adversarial training phase. Furthermore, prior work on long-tailed robustness has largely overlooked a crucial evaluation metric—Balanced Accuracy. To fill this gap, we introduce the concept of Balanced Robustness, a comprehensive metric that measures robustness specifically under long-tailed distributions. Extensive experiments demonstrate that our method outperforms existing advanced defenses, yielding significant improvements in both memory and computational efficiency. We believe this work represents a substantial step forward in tackling robustness challenges in real-world applications. Our paper code can be found at <https://github.com/BuhuiOK/TAET-Two-Stage-Adversarial-Equalization-Training-on-Long-Tailed-Distributions>

Corresponding author: Jian Liu, [jianliu@mail.hfut.edu.cn](mailto:jianliu@mail.hfut.edu.cn)

# 1. Introduction

In recent years, deep learning has achieved groundbreaking advancements in computer vision<sup>[1]</sup>. However, deep neural networks remain highly susceptible to adversarial attacks<sup>[2][3]</sup>, a challenge that continues to raise concerns in both academia and industry. These attacks introduce subtle perturbations to input data, leading to erroneous predictions and revealing the vulnerabilities of neural networks to malicious disturbances, which pose significant security risks to the deployment of modern computer vision models in real-world applications<sup>[4]</sup>. To address this issue, researchers have developed a variety of techniques<sup>[5][6][7][8]</sup> aimed at enhancing adversarial robustness. Among these, Adversarial Training (AT)<sup>[9]</sup> is widely regarded as one of the most effective strategies for improving model resilience. The central concept of AT involves integrating adversarial examples into the training process, thereby fortifying the model's ability to recognize and counteract malicious perturbations and enhancing its generalization capability in practical scenarios, ultimately improving both reliability and security.



**Figure 1.** *Top:* The distribution of accuracy and adversarial robustness across different classes in a long-tail distributions, with gray bars representing the sample counts of each class. AT-BSL (left) exhibits poorer performance on the tail classes (7, 8, 9) and Class 3. *Bottom:* The evaluation includes both balanced accuracy and robustness, comparing long-tail recognition methods, adversarial training, state-of-the-art defenses, and our proposed approach. The results demonstrate that our method outperforms the others in both balanced accuracy and performance on weak classes.

Adversarial training has shown strong performance on balanced datasets (e.g., MNIST, CIFAR-10<sup>[10]</sup>, ImageNet<sup>[11]</sup>). However, real-world data often follow long-tailed distributions<sup>[12][13][14]</sup>, where certain classes are overrepresented while others are underrepresented. This imbalance complicates tasks and undermines model robustness, necessitating adversarial training strategies tailored for such distributions. Long-tail learning addresses imbalanced datasets where a few head classes dominate, and many tail classes are underrepresented, creating a skewed frequency distribution<sup>[15][16][17]</sup>. The Imbalance Ratio (IR)<sup>[18]</sup> quantifies this imbalance, with higher values indicating greater scarcity of tail class samples, exacerbating learning challenges. While long-tail learning aims to address these issues,

problems like model bias toward head classes and insufficient tail class data persist<sup>[12]</sup>. Existing robustness methods struggle with suboptimal performance on tail classes<sup>[14][19][20]</sup>, as shown in Fig. 1. In adversarial training for long-tailed distributions, accuracy improves but adversarial robustness does not scale proportionally, leading to robustness overfitting. This phenomenon results in models becoming more vulnerable to adversarial attacks despite higher accuracy. To tackle this, we propose Hierarchical Adversarial Robust Learning (HARL), a framework designed to optimize performance on underrepresented classes and improve adversarial robustness in long-tailed distributions.

Our evaluation of robustness in long-tailed distributions reveals significant limitations in current assessment metrics<sup>[14][19][20]</sup>. Although balanced accuracy<sup>[12][13][21][22]</sup> is commonly used for long-tail recognition, traditional accuracy remains the dominant metric for evaluating long-tail robustness. To address this gap, we introduce a novel metric—**Balanced Robustness**—which, alongside balanced accuracy, provides a more comprehensive framework for evaluating robustness in long-tailed distributions. To mitigate robustness overfitting<sup>[23]</sup> while optimizing training efficiency, we explore the potential of using cross-entropy loss during the early stages of adversarial training. We hypothesize that early-stage cross-entropy loss can improve both natural accuracy and adversarial robustness. Based on this insight, we propose a two-stage adversarial balanced training approach that optimizes both balanced robustness and accuracy, while minimizing memory usage and computational overhead, thereby improving practical applicability.

Our contributions are summarized as follows:

- We analyze adversarial training under long-tailed distributions and identify a key limitation: while training accuracy improves, adversarial robustness lags behind. Our root-cause analysis informs targeted enhancements to address this gap in real-world robustness.
- We introduce Hierarchical Adversarial Robust Learning (HARL), a method designed to enhance robustness for underrepresented classes in long-tail distributions, achieving competitive natural accuracy and outperforming mainstream methods across multiple datasets.
- We propose a new evaluation metric, Balanced Robustness, which, alongside balanced accuracy, offers a more effective measure of robustness in long-tail scenarios.
- We present a two-stage adversarial training framework that incorporates cross-entropy loss during the early stages to improve both natural accuracy and adversarial robustness. Comparative results

demonstrate that our approach optimizes memory efficiency and computational time, surpassing alternative methods.

## 2. Related work

### 2.1. Long-Tailed Recognition

Long-tail distribution is a prevalent form of data imbalance in training datasets, characterized by a small number of "head" classes with abundant samples, contrasted with a majority of "tail" classes that have relatively few samples<sup>[12]</sup>. Models trained on such imbalanced distributions typically show high confidence in predicting head classes, leading to "overconfidence" that severely hampers generalization to tail classes. This issue becomes particularly pronounced when evaluating rare categories in real-world applications, as illustrated in Fig. 5. Effectively addressing this challenge remains a significant research hurdle<sup>[24]</sup>.

To tackle the difficulties of long-tail recognition, various strategies have been proposed, including resampling<sup>[25]</sup>, cost-sensitive learning<sup>[26]</sup>, decoupled training<sup>[27]</sup>, and classifier design<sup>[28]</sup>. These methods aim to mitigate data imbalance by altering data distributions, adjusting loss weights, and enhancing feature learning. Recent innovations, such as Class-Conditional Sharpness-Aware Minimization (CC-SAM)<sup>[29]</sup> and Feature Cluster Compression (FCC)<sup>[30]</sup>, offer novel solutions to improve model generalization and recognition accuracy by refining feature learning and classification processes. However, despite advancements in these areas, few studies have specifically addressed improving adversarial robustness in long-tail distributions<sup>[14][19][20]</sup>, highlighting a crucial gap that warrants further investigation.

### 2.2. Adversarial Robustness

To mitigate adversarial vulnerability in deep learning models, a variety of defense strategies have been proposed, including adversarial training<sup>[9]</sup>, defensive distillation<sup>[31]</sup>, ensemble methods<sup>[32]</sup>, and data augmentation<sup>[33]</sup>, with adversarial training being one of the most robust and widely adopted approaches. Techniques such as supervised learning<sup>[34]</sup>, feature denoising<sup>[35]</sup>, and statistical filtering<sup>[36]</sup> further enhance model robustness by improving generalization and reducing the impact of adversarial perturbations. Due to the high computational cost associated with adversarial training, recent advancements have focused on improving its efficiency<sup>[37]</sup>, making it more viable for deployment in

real-world applications. The core principle of adversarial training involves generating the most challenging adversarial samples internally and optimizing the model's parameters based on these samples, which inherently enhances its robustness<sup>[9]</sup>. The objective of adversarial training can be formalized as follows:

$$\min_{\theta} \hat{\mathcal{R}}_{x'} = \frac{1}{|\mathcal{S}|} \sum_{(x,y) \in \mathcal{S}} \left( \max_{\delta \in \mathcal{B}(x)} \mathcal{L}(f_{\theta}(x + \delta), y) \right). \quad (1)$$

$\hat{\mathcal{R}}_{x'}$  represents the adversarial risk for model  $f_{\theta}$  on dataset  $\mathcal{S}$ . Here,  $|\mathcal{S}|$  denotes the number of samples in  $\mathcal{S}$ , with each sample consisting of an input-output pair  $(x, y)$ , where  $x$  is the unperturbed input and  $y$  is the ground truth label. The perturbation  $\delta$  is optimized within the constraint set  $\mathcal{B}(x)$ , defining permissible modifications to  $x$ . The formula computes the loss  $\mathcal{L}(f_{\theta}(x + \delta), y)$  for each perturbed sample, seeking the perturbation  $\delta$  that maximizes this loss.

Building upon the foundation of AT, subsequent works developed advanced adversarial training techniques such as TRADES<sup>[5]</sup>, ADT<sup>[38]</sup>, MART<sup>[8]</sup>, HE<sup>[6]</sup>, GAIRAT<sup>[7]</sup>, and LAS-AT<sup>[39]</sup>.

### 2.3. Robustness under Long-Tailed Distribution

Although long-tail recognition and adversarial robustness have garnered significant research interest, the challenge of adversarial robustness under long-tail distributions remains relatively underexplored. While real-world datasets often exhibit long-tail characteristics, only a limited number of studies have systematically addressed this issue<sup>[14][19][20]</sup>. Wu et al.<sup>[14]</sup> were among the first to investigate the impact of data imbalance on adversarial robustness, introducing the RoBal framework, which integrates a cosine classifier with a two-stage rebalancing strategy to enhance both natural and robust accuracy.

Building upon the RoBal framework, Yue et al.<sup>[19]</sup> proposed Adversarial Training with Balanced Softmax Loss (AT-BSL), a streamlined method that achieves performance comparable to RoBal while significantly reducing both training time and memory consumption. This study provides an in-depth analysis of AT-BSL, particularly focusing on the design of the Balanced Softmax Loss (BSL) in Sec. 3, offering novel insights and strategies to improve adversarial robustness under long-tail distributions.

## 3. Analysis of AT-BSL

In this section, we present Adversarial Training with Balanced Softmax Loss (AT-BSL)<sup>[19]</sup> and clarify its principles, along with defining the notation used.

### 3.1. Preliminaries

**Balanced Softmax Loss** adjusts class logits to enhance performance on long-tail datasets. It modifies the logit  $z_i$  for each class, adding a term  $b_y = \tau_b \log n_y$  for the target class  $y$  to control class influence<sup>[14]</sup>. The formula is as follows:

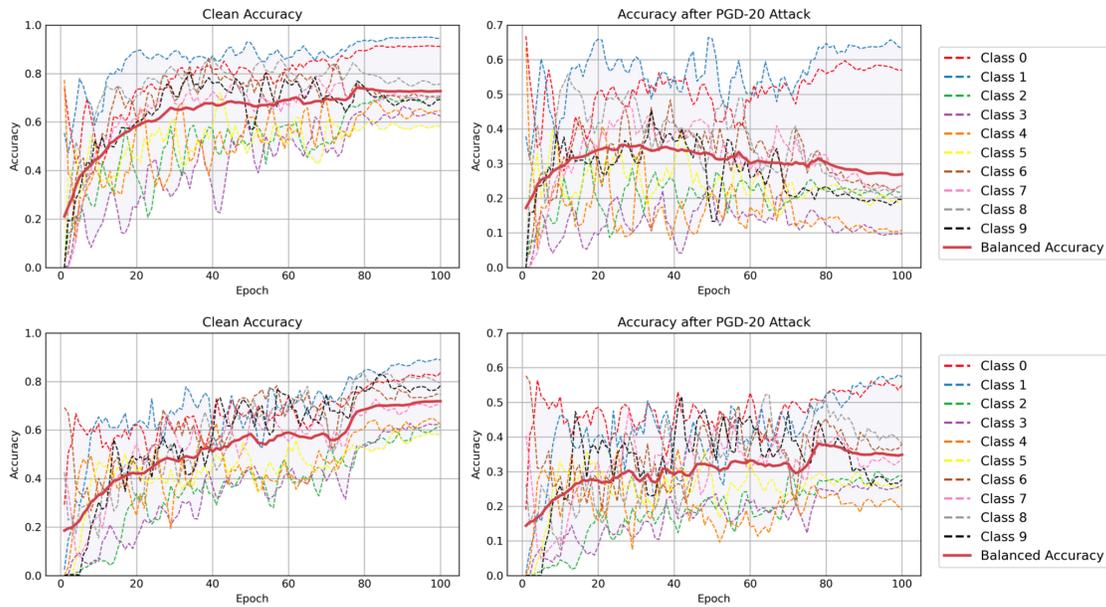
$$\begin{aligned} L_0(g(f(x)), y) &= -\log\left(\frac{e^{z_y+b_y}}{\sum_i e^{z_i+b_i}}\right) \\ &= -\log\left(\frac{n_y^{\tau_b} \cdot e^{z_y}}{\sum_i n_i^{\tau_b} \cdot e^{z_i}}\right). \end{aligned} \quad (2)$$

Here,  $g(\cdot)$  denotes the linear classifier,  $f(x)$  is the feature representation of input  $x$ ,  $y$  is the true class label,  $z_y$  the logit for class  $y$ ,  $b_y$  the bias for class  $y$ , and  $\tau_b$  a hyperparameter adjusting the logit range to control the softmax output distribution.

**Adversarial Training.** We previously outlined the objective of adversarial training<sup>[9]</sup>. Here, we detail the steps to generate adversarial samples in the AT-BSL<sup>[19]</sup> framework, as shown in the following formula:

$$x'^{(t+1)} = \text{Proj}_{\mathcal{B}_\infty} \left( x'^{(t)} + \alpha \cdot \text{sign} \left( \nabla_{x'^{(t)}} \mathcal{L}_{CE}(f(x'^{(t)}), y) \right) \right) \quad (3)$$

Here,  $x'^{(t)}$  represents the adversarial sample at step  $t$ , generated to mislead the model.  $x$  is the natural sample, and  $\mathcal{B}_\infty(x, \epsilon)$  defines the  $\ell_\infty$  norm ball around  $x$  with radius  $\epsilon$ , limiting adversarial perturbations. Proj is the projection,  $\mathcal{L}_{CE}$  the cross-entropy loss, and  $\alpha$  the step size, controlling each perturbation's magnitude.  $\nabla$  is the loss gradient with respect to the current adversarial sample at step  $t$ , and sign applies the gradient's sign to ensure updates align with gradient ascent.



**Figure 2.** *Top:* Accuracy progression during training with AT-BSL (left) and under PGD-20 attack (right).  
*Bottom:* Accuracy progression during training with our method (left) and under PGD-20 attack (right).

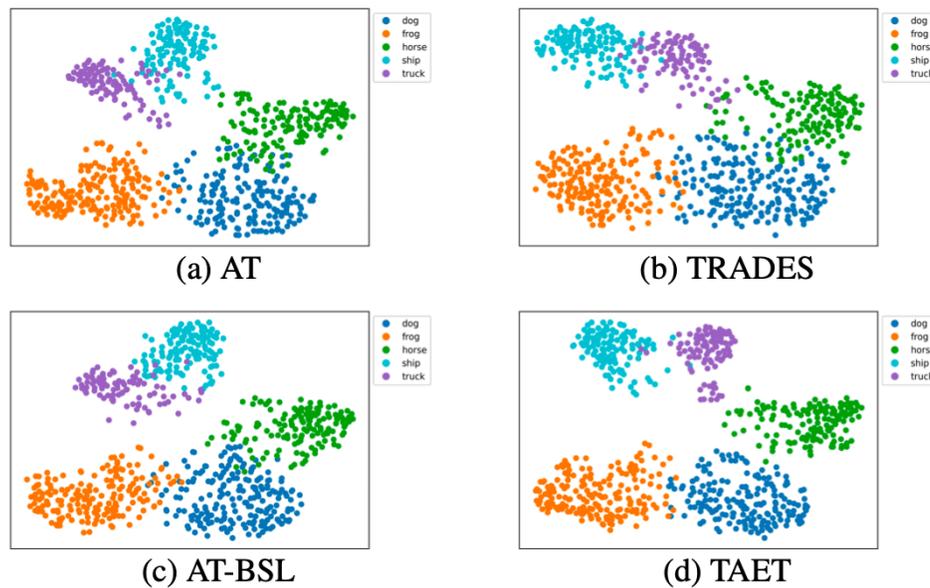
### 3.2. Limitations of AT-BSL

We conducted a thorough evaluation of AT-BSL<sup>[19]</sup>, with a particular focus on the variations in both accuracy and adversarial robustness across different classes throughout training. The key findings are as follows:

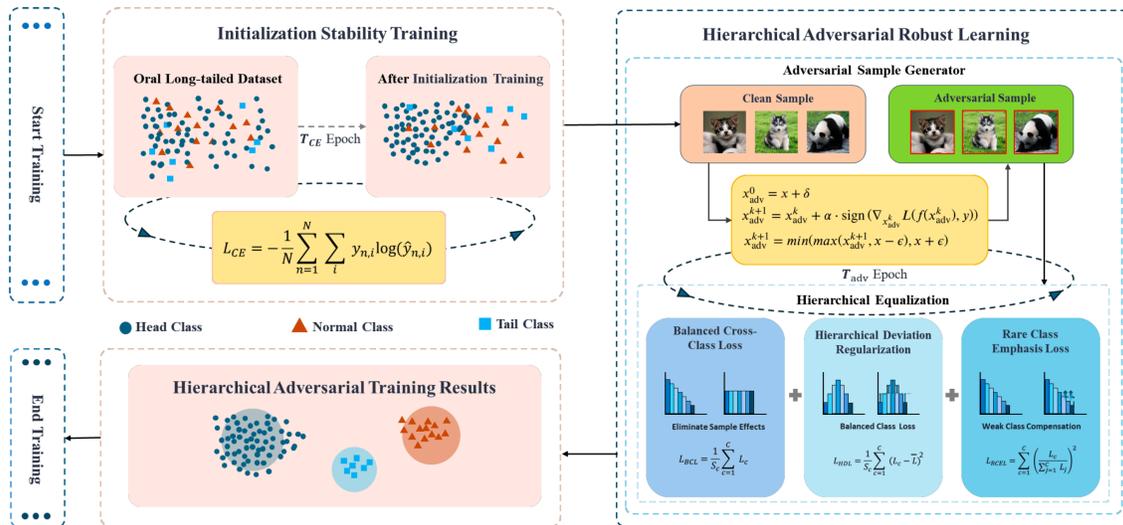
**The BSL method has substantial limitations in improving performance for underrepresented classes.** This method aims to enhance the performance of underrepresented classes by adjusting the output logits according to the class distribution, specifically by suppressing the logits of dominant classes in order to elevate those of minority classes. However, this straightforward approach fails to adequately improve the robustness of the weaker classes (as shown in Fig. 2). In contrast, our TAET method introduces a hierarchical equalization module that dynamically identifies underrepresented classes and facilitates targeted performance improvement. The experimental results indicate that our method significantly outperforms existing approaches in terms of adversarial robustness for underrepresented classes, as evidenced in Tab. 2.

**BSL is prone to robust overfitting.** Due to its fixed loss function, the model converges rapidly during training, as demonstrated in Figs. 2 and 3. The robustness of BSL peaks around the 25th epoch, after

which it gradually declines, particularly as the natural accuracy increases. This decline becomes more pronounced during the later training epochs (75-100), potentially causing instability when adjusting the learning rate. Although BSL improves accuracy, its adversarial robustness does not increase accordingly, and in some cases, it even decreases. As a result, it fails to achieve an optimal balance between accuracy and robustness. In contrast, our proposed method significantly improves both model robustness and accuracy, effectively reducing performance discrepancies across different classes. To address robust overfitting, we introduced a hybrid training strategy, which not only enhances model accuracy and robustness but also ensures stability when handling diverse inputs.



**Figure 3.** t-SNE visualization of latent space logits extracted from (a) AT, (b) TRADES, (c) AT-BSL, and (d) our proposed TAET method on tail classes (last five classes)



**Figure 4.** The TAET framework includes an Initial Stabilization Module (upper left) and the HARL module (right). The Initial Stabilization Module, based on cross-entropy loss, aims to stabilize accuracy in early training and transfers the trained model to HARL. Our HARL module consists of three components: BCL, HDL, and RCEL. A multi-step generation process creates perturbations (upper right), which are processed by normalization components (lower right).

## 4. Methodology

In Sec. 3, we conduct a thorough analysis of existing long-tail robustness solutions and identify two major limitations. First, current methods lack an effective mechanism for accurately identifying and compensating for the tail classes. To address this, we propose a novel adversarial equalization module, termed the Hierarchical Adversarial Robust Learning (HARL) framework, which consists of three key components, as detailed in Sec.4.1. Second, existing methods are prone to overfitting during training. To mitigate this issue, we introduce a two-stage adversarial equalization training approach, the steps of which are elaborated in Sec. 4.2.

### 4.1. Hierarchical Adversarial Robust Learning

Our Hierarchical Adversarial Robust Learning (HARL) framework includes an adversarial sample generation module and a hierarchical equalization component. The hierarchical equalization loss integrates Balanced Cross-Class Loss (BCL), Hierarchical Deviation Loss (HDL), and Rare Class Emphasis Loss (RCEL), working together with the adversarial sample generation module to balance accuracy and

robustness. This design addresses long-tail dataset challenges, boosting minority class performance and reducing overfitting, enhancing robustness in practical applications. Unlike BSL, which uses sample count for tail class identification, HARL identifies weak classes based on mean loss during training, strengthening both tail and weak classes. Sample-count methods often miss weak classes (e.g., Class 3 in Fig. 1), where large sample counts don't guarantee good performance, a limitation our method largely overcomes.

**Adversarial Sample Generation:** We employ a multi-step adversarial sample generation process, which begins with random perturbations and iteratively computes the gradient of the sample with respect to the loss function. This gradient guides the direction for generating subsequent adversarial samples. A primary challenge addressed in this study is the generation of balanced adversarial samples in long-tail distributions. The transformation of adversarial samples is facilitated through hierarchical equalization loss. The learning objective for this module is as follows:

$$\begin{aligned} & \arg \min_{\theta} \mathbb{E}_{(x_0, y) \sim D_{\text{lt}}} [\mathcal{L}_{\text{HEL}}(f_{\theta}(x'), y)] \\ \text{s.t.} \quad & \|x' - x_0\|_{\infty} \leq \epsilon, \end{aligned} \quad (4)$$

Here,  $\mathcal{L}_{\text{HEL}}$  represents the Hierarchical Equalization Loss<sup>[40][41]</sup>, adjusting model logits for adversarial samples<sup>[9]</sup>.  $D_{\text{lt}}$  is the long-tail dataset, and  $\epsilon$  denotes the allowable perturbation budget. The objective is to keep perturbations near the input sample  $x_0$  while significantly altering the model  $F$ 's predictions. This aligns with prior adversarial training methods, but our approach uniquely integrates long-tail recognition with adversarial training. As shown in Fig. 4, we first generate adversarial samples, then apply hierarchical equalization to enhance weak classes. With a fixed perturbation budget, our goal is to minimize adversarial loss for long-tail recognition, with the loss defined as:

$$\mathcal{L}_{\text{HEL}} = \alpha \cdot \mathcal{L}_{\text{BCL}} + \beta \cdot \mathcal{L}_{\text{HDL}} + \gamma \cdot \mathcal{L}_{\text{RCEL}} \quad (5)$$

Here,  $\alpha$ ,  $\beta$ , and  $\gamma$  are hyperparameters for loss reweighting, tuned through validation. Each loss component is detailed below.

**Balanced Cross-Class Loss (BCL)** ensures balanced losses across classes. By reweighting, it prevents excessive focus on head classes in long-tail datasets, enabling more equal contribution to total loss. BCL achieves this by averaging class losses, enhancing classifier performance across all classes under long-tail distribution, as shown below:

$$\mathcal{L}_{\text{BCL}} = \frac{1}{S_c} \sum_{c=1}^C \mathcal{L}_c \quad (6)$$

Here,  $S$  is the total number of classes, and  $L_c$  represents the loss for class  $c$ . Averaging across all classes helps the model maintain balanced attention in a long-tail setting, avoiding bias toward head classes.

**Hierarchical Deviation Loss (HDL)** adjusts the discrepancies in inter-class loss to mitigate extreme imbalances. It quantifies the deviation of each class's loss from the mean and applies a quadratic penalty to reduce this gap, thereby improving robustness against imbalanced data distributions. This hierarchical framework allows the model to better accommodate the complexities inherent in different class distributions, as illustrated below:

$$\mathcal{L}_{\text{HDL}} = \frac{1}{S_c} \sum_{c=1}^C (\mathcal{L}_c - \bar{\mathcal{L}})^2 \quad (7)$$

**Rare Class Emphasis Loss (RCEL)** addresses the challenge of long-tail distributions by focusing on rare classes, which are typically more difficult to classify. By normalizing class-wise losses and assigning higher weights to rare classes, RCEL amplifies their loss contributions. This approach encourages the model to prioritize learning from these underrepresented classes, thereby improving overall classification performance, as demonstrated below:

$$\mathcal{L}_{\text{RCEL}} = \sum_{c=1}^C \left( \frac{\mathcal{L}_c}{\sum_{j=1}^C \mathcal{L}_j} \right)^2 \quad (8)$$

## 4.2. Two-Stage Adversarial Equalization Training

In Fig. 4, we present our long-tail robustness framework, which initially employs a cross-entropy loss function. Our experiments show that cross-entropy loss<sup>[42]</sup> facilitates rapid convergence on unperturbed samples, ensuring early-stage accuracy. It provides stable gradient signals, enhancing robustness against adversarial attacks and mitigating robust overfitting. For a detailed experimental analysis, please refer to Sec. 6.

$$\mathcal{L}_{\text{CE}} = -\frac{1}{N} \sum_{n=1}^N \sum_i y_{(n,i)} \log(\hat{y}_{(n,i)}) \quad (9)$$

In the later stages of training, we incorporate adversarial training along with the hierarchical equalization component to improve model performance. To highlight the contribution of hierarchical equalization within the TAET framework, we conduct an ablation study in Sec. 6.4, evaluating the impact of each component individually. By selectively removing components, we examine how different configurations affect both accuracy and robustness, thereby identifying the optimal configuration. Fig. 6 illustrates the training procedure and component interactions.

## 5. Revisiting Evaluation Metrics for the Long-Tail Robustness

### Problem

**Balanced Accuracy:** In long-tail robustness research, common evaluation metrics include clean accuracy and post-attack robustness. However, Balanced Accuracy (BA), an important metric for long-tail learning, is often overlooked<sup>[12][13][22][15]</sup>. BA balances performance across classes by considering both true positive and true negative rates, mitigating class bias. In imbalanced datasets, traditional accuracy metrics may neglect minority classes due to the majority class's dominance. In contrast, BA ensures fair evaluation by capturing minority class performance, offering a comprehensive assessment of overall model performance. The mathematical definition of BA<sup>[43]</sup> is as follows:

$$\text{Balanced\_Accuracy} = \frac{1}{S_C} \sum_{i=1}^C A_i^{x_0} \quad (10)$$

Where  $A_i = \frac{TP_i}{TP_i + FN_i}$ ,  $TP_i$  represents the number of true positives for class  $i$ , and  $FN_i$  represents the number of false negatives for class  $i$ . The following confusion matrix provides additional detail.

**Balanced Robustness:** Balanced accuracy reflects model performance across classes<sup>[43][44]</sup>, mitigating bias from class imbalance, which is crucial for long-tail recognition. In long-tail robustness, we aim for consistent model performance across all classes under adversarial attacks. We extend balanced accuracy to balanced robustness, quantifying the model's average defense capability against attacks across classes. Specifically, balanced robustness, under adversarial samples, is defined as follows:

$$\begin{aligned} \text{Balance\_Robustness} &= \frac{1}{S_C} \sum_{i=1}^C \mathcal{R}_i^{x'} \\ &= \frac{1}{S_C} \sum_{i=1}^C \frac{TP_i^{x'}}{TP_i^{x'} + FN_i^{x'}} \end{aligned} \quad (11)$$

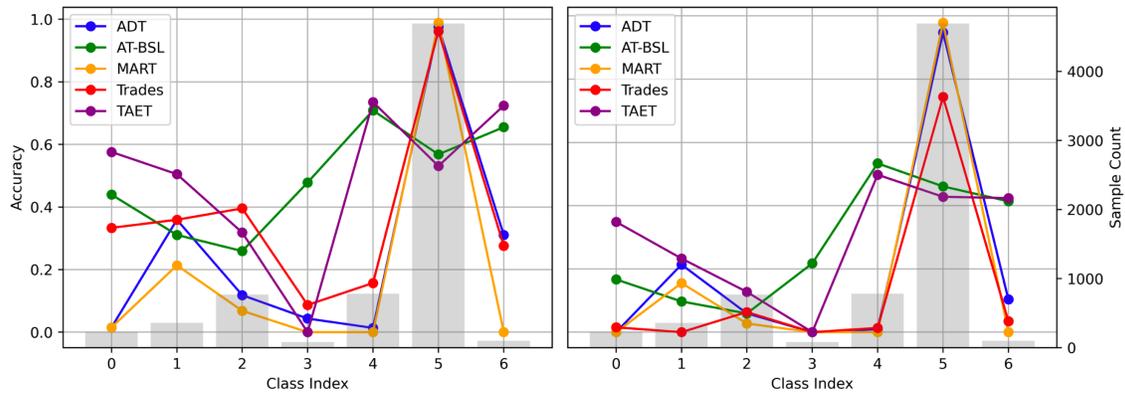
We introduce balanced robustness, a novel metric in the context of long-tail learning, to the best of our knowledge, marking its first application in this domain. We believe this metric can effectively measure the adversarial robustness of methods under long-tail distributions and will have extremely important applications in fields such as medicine. We advocate for the adoption of this metric as a standardized evaluation tool to promote consistency in the research and practical applications of long-tail robustness and to drive its further development.

Method	Clean	FGSM	PGD		CW	AA
			20	100		
AT-BSL <sup>[19]</sup>	72.74	34.13	26.86	25.62	15.67	25.26
RoBal <sup>[14]</sup>	73.18	33.14	27.12	26.98	13.44	24.13
TRADES <sup>[5]</sup>	65.77	25.73	20.23	19.59	27.42	19.63
AT <sup>[9]</sup>	69.00	32.53	25.69	24.55	15.22	24.28
MART <sup>[8]</sup>	58.33	33.67	30.10	29.36	48.12	26.04
ADT <sup>[38]</sup>	68.08	32.70	26.00	25.27	10.05	25.06
REAT <sup>[20]</sup>	74.56	31.42	24.02	22.52	11.07	22.69
LAS-AT <sup>[39]</sup>	64.04	33.04	27.80	26.74	36.77	24.71
HE <sup>[6]</sup>	68.18	32.40	25.66	24.59	12.94	24.32
GAIRAT <sup>[7]</sup>	61.07	28.90	24.79	23.81	42.62	24.59
TAET (our)	<b>74.67</b>	<b>39.59</b>	<b>35.15</b>	<b>34.29</b>	<b>57.59</b>	<b>30.57</b>

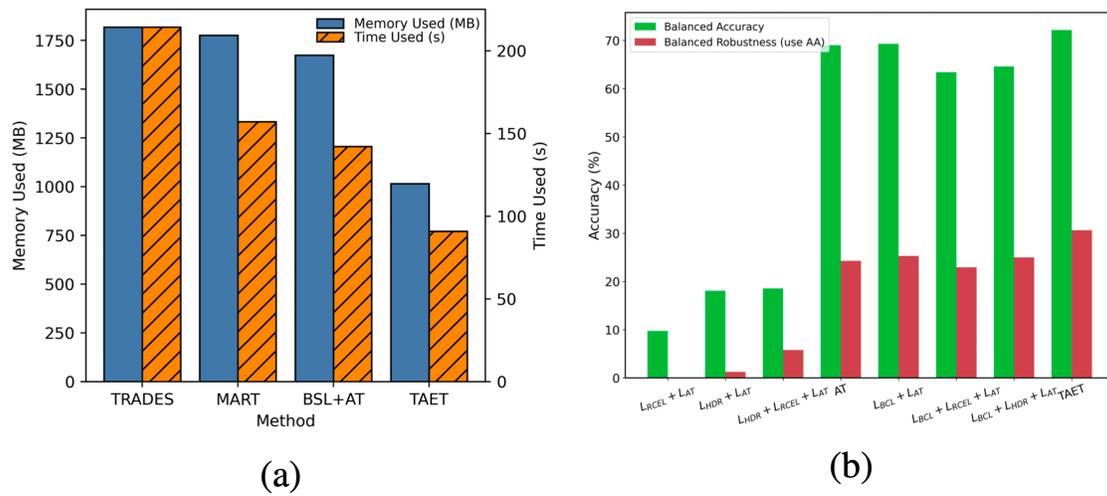
**Table 1.** Balanced accuracy and balanced robustness of various algorithms using ResNet-18 on CIFAR-10-LT with an imbalance ratio (IR) of 10. The best results are highlighted in **bold**.

Method	Clean					Attacked (PGD-20)				
	Dog	Frog	Horse	Ship	Truck	Dog	Frog	Horse	Ship	Truck
AT-BSL <sup>[19]</sup>	<u>58.63</u>	73.48	70.48	78.29	66	<u>23.02</u>	<u>26.97</u>	21.7	<u>21.7</u>	19
RoBal <sup>[14]</sup>	<b>59.13</b>	<u>74.73</u>	68.52	<u>78.96</u>	65	<b>23.74</b>	25.78	20.79	19.5	<u>20</u>
TRADES <sup>[5]</sup>	53.96	65.11	63.25	47.28	47	11.87	13.2	16.27	3.87	4
AT <sup>[9]</sup>	51.07	62.32	63.85	64.34	54	14.74	17.67	19.27	17.05	10
ADT <sup>[38]</sup>	52.15	65.58	65.66	56.59	55	16.18	16.74	18.67	10.07	10
MART <sup>[8]</sup>	37.76	52.09	58.43	47.25	41	17.62	19.06	<u>25.3</u>	8.52	9
REAT <sup>[20]</sup>	<u>58.63</u>	73.02	<b>72.89</b>	78.29	<u>69</u>	16.18	13.48	18.67	20.15	17
LAS-AT <sup>[39]</sup>	50.72	55.34	56.02	48.83	34	19.06	13.02	20.48	12.4	8
HE <sup>[6]</sup>	57.31	67.9	65.66	51.16	54	14.74	20	18.07	10.85	12
GAIRAT <sup>[7]</sup>	48.56	58.13	56.62	48.06	43	15.82	17.2	22.28	6.2	11
TAET (our)	56.11	<b>79.06</b>	<u>72.28</u>	<b>80.6</b>	<b>77</b>	<b>23.74</b>	<b>34.83</b>	<b>32.53</b>	<b>36.34</b>	<b>27</b>

**Table 2.** Performance of various methods on the last Tail classes (last five classes) under clean and attacked conditions using ResNet-18 on CIFAR-10-LT(IR=10). The best results are highlighted in **bold**, and the second-best results are underlined.



**Figure 5.** Results after training the VIT-B/16 model on the DermaMNIST for 100 epochs. The left side of the figure shows balanced accuracy of different methods in a natural setting, while the right side presents balanced robustness under a PGD-20 attack. The background section illustrates data distribution in DermaMNIST.



**Figure 6.** (a) Memory usage and time per epoch for each method; (b) Robustness and clean accuracy performance based on the effectiveness of each component in Eq. (5).

## 6. Experiment

### 6.1. Settings

Following<sup>[14]</sup>, we evaluated our approach on CIFAR-10-LT and CIFAR-100-LT, and extended experiments to the MedMNIST<sup>[45]</sup> for real-world validation. MedMNIST consists of 12 medical image datasets across various modalities (e.g., CT, X-ray, ultrasound, OCT), supporting tasks such as multi-label and binary classification. We selected DermaMNIST, a highly imbalanced dataset from HAM10000<sup>[46]</sup> with an imbalance ratio (IR) of 58.66, containing 10,015 images in 7 diagnostic categories.

The primary metrics used were **Balanced Accuracy**<sup>[43]</sup> and **Balanced Robustness**, with dataset imbalance assessed by the imbalance ratio (IR). Model robustness was evaluated under  $l_\infty$  bounded perturbations with a size of  $8/255$  using FGSM<sup>[3]</sup> and PGD<sup>[9]</sup> (20 and 100 steps, step size  $1/255$ ), as well as a 100-step CW<sup>[47]</sup> attack and AutoAttack (AA)<sup>[48]</sup>, a powerful ensemble method.

### 6.2. Main Results

The results in Tab. 1 show that on CIFAR-10-LT, our method achieves the highest balanced accuracy and robustness on ResNet-18, with a 5.31% improvement in robustness against AA<sup>[48]</sup> attacks compared to the AT-BSL<sup>[19]</sup> model. Overall, our method outperforms other methods in handling long-tail datasets.

Tab. 2 provides a detailed breakdown of clean accuracy and robustness in the Tail Class, demonstrating that our method significantly improves tail class performance. This is critical in real-world scenarios where tail class accuracy is low, such as in rare disease research. Our method also improves robustness for challenging classes like Class 3.

To validate our method on real-world long-tail datasets, Tab. 3 presents results on the DermaMNIST subset of MedMNIST<sup>[45]</sup>, showing optimal adversarial robustness and competitive balanced accuracy. Fig. 5 compares clean accuracy and adversarially perturbed accuracy across classes. Our method outperforms others on tail classes such as Class 0, 1, 2, and 4. For extremely rare classes like Class 3 and Class 6, all methods show high variability. While TAET's hierarchical equalization improves tail class performance, weight adjustments for very rare classes may still be insufficient. Increasing the hyperparameter  $\gamma$  in Eq. (5) could enhance these classes, though it might impact overall performance.

Method	Clean	FGSM	PGD		CW	AA
			20	100		
ADT	26.23	20.07	19.04	18.93	23.89	17.81
MART	18.37	17.71	16.59	16.07	17.61	15.09
TRADES	18.89	16.57	12.48	11.35	31.18	10.50
AT-BSL	<b>48.55</b>	<u>38.81</u>	<u>27.85</u>	<u>26.72</u>	<b>43.44</b>	<u>21.04</u>
TAET (our)	<u>48.42</u>	<b>40.12</b>	<b>29.26</b>	<b>28.79</b>	<u>43.11</u>	<b>21.74</b>

**Table 3.** Balanced accuracy and balanced robustness under Various Attack Methods on the MedMNIST , best results are highlighted in **bold**, and the second-best results are underlined.

### 6.3. Initial Training Epochs and Model Robustness

To address robust overfitting<sup>[23][49]</sup>, we hypothesize that limited model accuracy may constrain robustness. Thus, improving accuracy on clean samples could also enhance robustness. Previous studies used data augmentation with some success. In our approach, we apply cross-entropy loss for a set number of epochs in the initial training phase. Experiments show that using cross-entropy loss<sup>[42]</sup> for 40 epochs increased clean sample accuracy by 7.19% and balanced robustness under AA attack by 1.04%, compared to omitting cross-entropy loss. This strategy also reduces computational costs, improving both model performance and training efficiency, as shown in Figure 7.

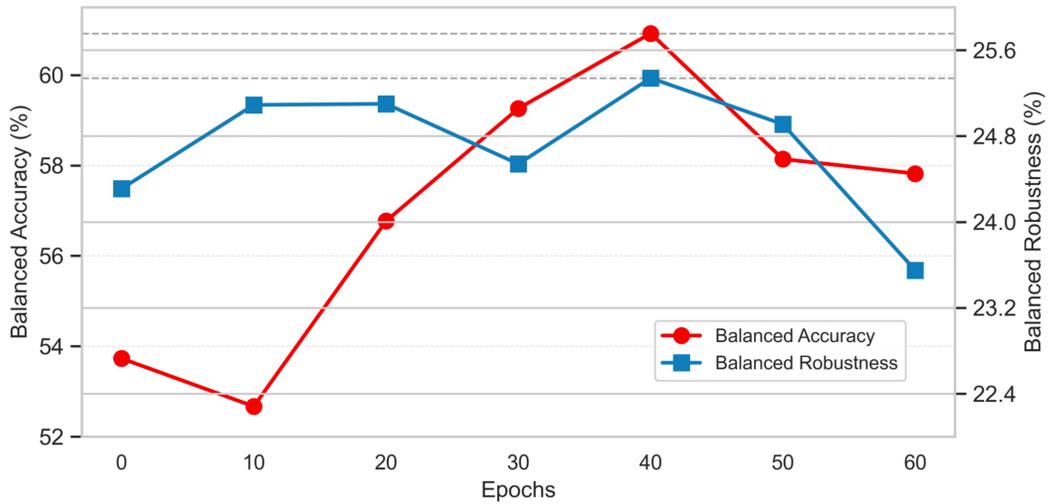


Figure 7. Results of using different epochs of cross-entropy loss in the early training stages.

#### 6.4. Ablation Study

We analyzed the proposed adversarial training on CIFAR-10-LT using ResNet-18<sup>[50]</sup> for more insights:

**Contribution of HARL.** Fig. 6 shows the impact of each component. BCL notably improves accuracy, as expected. HDL and RCEL are designed to boost weak classes and perform best in combination.

**Hyperparameter Importance Analysis.** We examined the influence of hyperparameters  $\alpha$ ,  $\beta$ , and  $\gamma$  on performance, adjusting their ratios to assess stability. Due to squared terms in the loss, value fluctuations may impact training stability. Thus, we used a 40-epoch cross-entropy training period for stable observation. Tab. 4 details results across hyperparameter settings, showing that equal weights ( $\alpha = \beta = \gamma$ ) yield the best performance. This emphasizes the importance of balanced hyperparameters in enhancing robustness and accuracy for long-tail datasets.

		$\alpha = 0.1$	$\alpha = 0.1$	$\alpha = 0.05$	$\alpha = 0.05$
		$\gamma = 0.1$	$\gamma = 0.05$	$\gamma = 0.1$	$\gamma = 0.05$
$\beta = 0.1$		<b>35.15</b>	34.7	33.14	31.24
$\beta = 0.05$		34.44	34.97	33.56	35.04

**Table 4.** Balanced robustness results of ResNet-18 on CIFAR-10-LT (IR=10) under different hyperparameter settings.

### 6.5. Results Across Long-Tail Imbalance Ratios

To evaluate the effectiveness of our method across different Imbalance Ratios (IR), we generated long-tail versions of CIFAR-10-LT with varying IRs. The results shown in Tab. 5 demonstrate that our method outperforms previous state-of-the-art techniques in terms of robustness under a variety of IR conditions. This enhanced robustness not only improves class adaptability but also boosts overall accuracy in imbalanced datasets, underscoring the practical potential of our approach for real-world applications.

IR	Method	Clean	FGSM	PGD-20	CW	AA
10	AT-BSL	72.74	34.13	26.86	15.67	25.26
	TAET	<b>74.67</b>	<b>39.59</b>	<b>35.15</b>	<b>57.59</b>	<b>30.57</b>
20	AT-BSL	66.4	28.82	22.56	15.43	21.46
	TAET	<b>66.93</b>	<b>36.82</b>	<b>33.98</b>	<b>58.41</b>	<b>27.84</b>
50	AT-BSL	58.23	26.37	20.87	16.74	19.79
	TAET	<b>59.37</b>	<b>31.69</b>	<b>29.35</b>	<b>50.15</b>	<b>25.3</b>
100	AT-BSL	<b>54.73</b>	22.51	19.17	17.34	18.29
	TAET	52.68	<b>28.68</b>	<b>26.53</b>	<b>42.3</b>	<b>22.98</b>

**Table 5.** Balanced accuracy and balanced robustness using ResNet-18 on CIFAR-10-LT under different long-tail imbalance ratios(IRs). Better results are highlighted in bold.

## 7. Conclusion

In this study, we analyze the limitations of AT-BSL in addressing long-tail robustness. Through extensive experiments, we show that the HARK strategy improves performance on both long-tail and weak classes. We propose balanced robustness as a novel metric to assess long-tail robustness, extending the concept of balanced accuracy. To mitigate robust overfitting, we introduce a two-stage adversarial equalization training (TAET) approach, which reduces overfitting while improving both adversarial robustness and clean accuracy. This method is computationally efficient and effective for practical applications. Our analysis confirms the generalizability and superiority of the proposed method. This work significantly advances adversarial training for real-world scenarios, enhancing adaptability to long-tail data. Future work will optimizing performance for long-tail distributions. Additionally, we will conduct experiments on real-world long-tail datasets to refine the model's adversarial robustness.

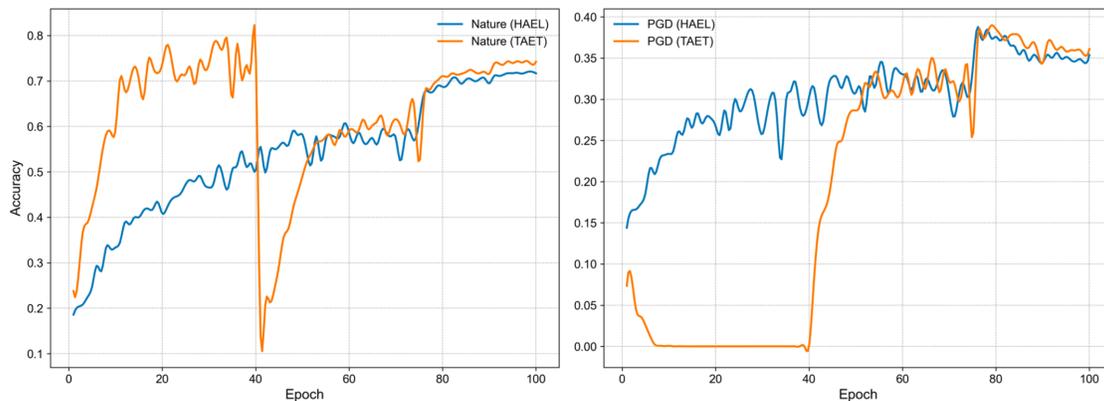
## Appendix A. Implementation Details of Experiments

### A.1. Training Details and Hyper-parameter Setting

We use ResNet-18 as the model architecture. The initial learning rate is set to 0.1, with decay factors of 10 at epochs 75 and 90, for a total of 100 epochs. We evaluate using the final epoch, and no methods employ early stopping. We use the SGD optimizer with a momentum of 0.9 and weight decay set to  $5e-4$ . In the main paper, we set the batch size to 128. For adversarial training, the maximum perturbation is set to  $8/255$ , the step size to  $2/255$ , and the number of steps to 10. To ensure fairness in the comparison, we adopt this same configuration for adversarial training across all methods. There are several hyperparameters involved, among which the most influential are  $\alpha$ ,  $\beta$ ,  $\gamma$ , and the number of epochs during the initial phase of training. For CIFAR-10-LT, CIFAR-100-LT, and Dermamnist, we set  $\alpha = \beta = \gamma = 0.1$  and use 40 epochs for training in the adversarial domain.

### A.2. Training Details and Hyper-parameter Setting

For the defense methods compared in this paper, we use their official implementations, including AT-BSL<sup>[19]</sup><sup>1</sup>, RoBAL<sup>[14]</sup><sup>2</sup>, AT<sup>[9]</sup><sup>3</sup>, TRADES<sup>[5]</sup><sup>4</sup>, MART<sup>[8]</sup><sup>5</sup>, GAIRAT<sup>[7]</sup><sup>6</sup>, LAS-AT<sup>[39]</sup><sup>7</sup>, and REAT<sup>[20]</sup><sup>8</sup>. For the attacks used in the evaluation, we implement them based on the official code and the original papers, including FGSM<sup>[3]</sup>, PGD<sup>[9]</sup>, CW<sup>[47]</sup>, and AutoAttack<sup>[48]</sup>.



**Figure 8.** Variations in clean accuracy (left) and adversarial robustness (right) during training with the application or absence of two-stage training.

## Appendix B. Additional Experiments

### *B.1. Necessity and Benefits of the TAET Two-Stage Training*

The comparison between the two-stage training method (TAET) and the single-stage approach (HARL) is illustrated in the figure. TAET employs a two-stage training strategy, where the first 40 epochs are dedicated to optimizing cross-entropy (CE), followed by adversarial loss (HARL) optimization in subsequent epochs. This methodology enables a dynamic equilibrium between natural accuracy and adversarial robustness as shown in Figure 8.

The left graph compares TAET and HARL in terms of natural accuracy. During the initial 40 epochs of CE training, TAET primarily concentrates on improving the classification accuracy of natural samples. This phase enables rapid convergence and leads to a significant performance improvement over HARL, thus establishing a strong baseline for natural performance. After completing the 40th epoch, TAET transitions to HARL optimization. Although a transient fluctuation in natural accuracy occurs during this transition, the model quickly adapts to the new training objective and ultimately achieves a natural accuracy of approximately 0.8, clearly surpassing HARL. This outcome demonstrates that TAET's two-stage training approach provides sufficient room for optimizing adversarial robustness while maintaining high natural accuracy, all while offering a substantial reduction in resource requirements.

The right graph presents the comparison of adversarial robustness under PGD attacks. During the first 40 epochs of CE training, TAET predominantly focuses on optimizing natural accuracy, with minimal attention to adversarial robustness, resulting in nearly zero PGD robustness. However, following the switch to HARL optimization at epoch 40, TAET's PGD robustness improves at an accelerated rate, eventually surpassing HARL in the later stages of training and stabilizing around 0.36. This finding underscores that TAET's strategy effectively prioritizes natural accuracy optimization before shifting focus to adversarial robustness, achieving a dynamic balance between the two objectives. Additionally, by combining CE and HARL in a complementary manner, TAET not only significantly enhances the model's ability to generalize on natural samples but also demonstrates substantial potential for improving adversarial robustness, while significantly reducing the computational burden compared to HARL. These advantages collectively result in nearly halving the time and memory consumption, as shown in Figure 6.

Method	Clean	FGSM	PGD	CW	AA
AT-BSL	77.74	43.56	36.53	22.77	34.99
TRADES	75.66	39.17	32.71	40.76	31.83
AT	77.96	44.98	37.7	24.41	36.09
MART	69.8	44.45	39.66	34.31	38.24
ADT	76.98	46.18	39.42	18.78	38.1
REAT	<b>79.72</b>	41.25	33.39	17.5	31.85
LAS-AT	76.02	44.16	38.27	51.73	36.15
HE	77.61	<b>46.18</b>	38.73	22.3	37.26
GAIRAT	71.49	41.77	36.87	54.6	33.79
TAET(ours)	76.31	45.1	<b>40.8</b>	<b>62.21</b>	<b>39.24</b>

**Table 6.** The standard accuracy performance of various methods ON CIFAR10-LT

### B.2. Conventional Evaluation Metrics on CIFAR10-LT

The experimental results on CIFAR-10-LT using accuracy and robustness metrics are shown in Tab. 6. From these results, it is evident that our method achieves high robustness while maintaining strong performance in terms of conventional accuracy, demonstrating the best performance across multiple adversarial attack scenarios, including PGD, CW, and AA. Additionally, we observe that the MART method exhibits robust performance under the long-tail distribution, but its accuracy on clean samples is relatively lower. This suggests that a promising avenue for further improvement could involve combining the MART method with long-tail recognition approaches to enhance robustness in long-tail settings while maintaining higher accuracy on natural samples.

### B.3. Different IRs

We present the results on CIFAR-10-LT and CIFAR-100-LT datasets under different Imbalance Ratios (IR) in Tabs. 7, 8. The results show that our method outperforms the previous state-of-the-art method, AT-BSL, across all evaluation metrics, further confirming the effectiveness of our approach. Notably, in the

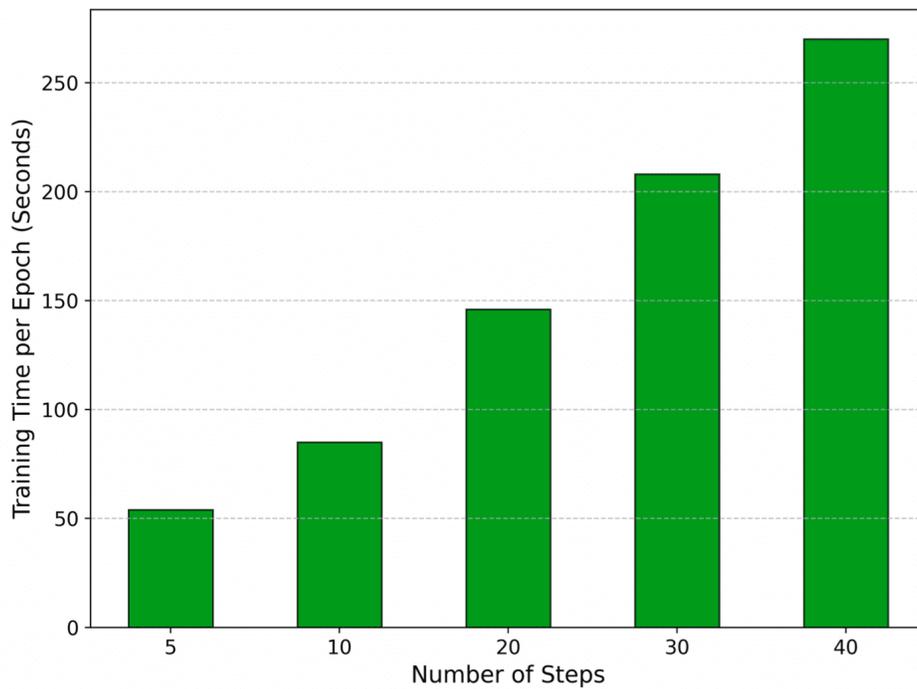
standard accuracy evaluation on the CIFAR-10 dataset, our method slightly underperforms AT-BSL<sup>[19]</sup> on clean samples. This phenomenon can be attributed to the fact that our method places more emphasis on optimizing the performance of the tail classes, which have relatively fewer samples in the overall dataset. As a result, while the balanced accuracy might be higher, the standard accuracy could be lower. Nevertheless, even under the standard accuracy evaluation, our method still demonstrates superior adversarial robustness compared to previous methods, which further highlights the superiority of our approach.

IR	Method	Clean	FGSM	PGD	CW	AA
10	AT-BSL	<b>77.74</b>	43.56	36.53	22.77	34.99
	TAET	76.31	<b>44.37</b>	<b>38.17</b>	<b>53.61</b>	<b>37.1</b>
20	AT-BSL	<b>78.3</b>	45.4	37.96	27.99	36.4
	TAET	77.57	<b>47.66</b>	<b>41.69</b>	<b>55.51</b>	<b>39.9</b>
50	AT-BSL	<b>80.68</b>	<b>52.74</b>	46.05	38.93	44.19
	TAET	77.53	51.2	<b>46.98</b>	<b>39.2</b>	<b>44.62</b>
100	AT-BSL	<b>81.92</b>	<b>54.52</b>	49.03	43.66	47.78
	TAET	74.09	51.98	<b>49.95</b>	<b>66.71</b>	<b>45.88</b>

**Table 7.** The standard accuracy performance under different Imbalance Ratios (IR) on CIRAR10-LT

IR	Method	Clean	FGSM	PGD	CW	AA
10	AT-BSL	48.41	24.77	20.81	18.8	18.71
	TAET	55.69	30.2	24.46	29.69	22.25
20	AT-BSL	48.27	24.98	20.23	20.05	18.12
	TAET	57.31	30.07	23.19	32.79	20.56
50	AT-BSL	49.46	26.19	22.99	24.43	20.35
	TAET	60.71	30.17	24.05	35.7	21.59
100	AT-BSL	48.48	28.13	24.34	28.2	21.58
	TAET	59.56	32.91	26.52	40.22	23.17

**Table 8.** The standard accuracy performance under Different Imbalance Ratios (IR) on CIFAR100-LT



**Figure 9.** Time Cost Under Different Numbers of Steps

#### B.4. Different PGD Steps

Through an analysis of the model's accuracy on clean samples and its performance under various adversarial attacks across different training steps in Tabs. 9, along with an evaluation of the associated time cost, we observed that the number of steps significantly impacts both adversarial robustness and computational efficiency. Specifically, clean accuracy remained stable as the number of steps increased, consistently ranging from 73% to 79%, indicating that the model retains strong robustness in the absence of attacks. Under adversarial attacks, model accuracy improved with more steps, particularly under FGSM and CW attacks, peaking at 30 steps. However, the associated training time increased substantially, rising from 54 seconds at 5 steps to 208 seconds at 30 steps, and further to 270 seconds at 40 steps, which imposes significant computational overhead. Notably, using 10 steps provided a reasonable balance, achieving competitive adversarial robustness and clean accuracy while requiring only 85 seconds per epoch, which is significantly lower than the time cost for 30 or 40 steps. These results suggest that selecting 10 steps offers a practical compromise between adversarial robustness, clean accuracy, and computational efficiency.

Steps	Clean	FGSM	PGD	CW	AA
5	78.86	31.21	22.7	30.52	19.77
10	74.25	40.64	35.11	60.22	31.17
20	73.27	40.21	36.62	60.25	31.36
30	73.24	41.82	37.55	61.14	33.17
40	73.89	41.00	36.43	59.90	32.20

**Table 9.** The Impact of Adversarial Training Steps on Adversarial Robustness

#### B.5. Different PGD Size

In Tabs. 10 we observed that step size adjustments significantly influence the model's adversarial robustness while having a relatively minor effect on clean sample performance. Specifically, clean accuracy remained stable across all step sizes, ranging between 73% and 74%, which highlights the model's robustness on natural samples. In contrast, adversarial robustness exhibited varying trends with

step size changes. Larger step sizes, such as  $3/255$ , notably improved robustness against stronger attacks like PGD and AA, with AA accuracy increasing from 31.17% to 37.74%. Meanwhile, robustness against FGSM attacks remained relatively stable, whereas performance under CW attacks showed some fluctuation as the step size increased. These findings suggest that increasing the step size can enhance robustness in certain adversarial scenarios but may also introduce vulnerabilities in others (e.g., CW attacks). Therefore, while appropriately increasing the step size can improve adversarial robustness without compromising clean accuracy, the optimal step size should be carefully selected based on task-specific requirements and the adversarial attack types encountered.

Size	Clean	FGSM	PGD	CW	AA
1/255	74.25	40.64	35.11	60.22	31.17
2/255	73.13	39.88	35.60	58.26	31.05
3/255	73.29	40.58	36.20	58.56	37.74

**Table 10.** The Impact of Adversarial Training Step Size on Adversarial Robustness

### *B.6. Different model*

In Tab. 11, we present a comparison between our method and the previous state-of-the-art method (AT-BSL) across different models. The results indicate that our method not only surpasses the original approach in terms of balanced accuracy but also in adversarial robustness, highlighting the superiority of our method. Furthermore, it can be observed that when training on the long-tailed dataset with ViT-b/16 (without using a pre-trained model), the AT-BSL method has completely failed, which demonstrates the effectiveness of our two-stage training approach in stabilizing model training.

Model	Method	Clean	FGSM	PGD	AA
ResNet50	AT-BSL	74.34	35.41	29.01	27.14
	TAET	<b>74.48</b>	<b>40.17</b>	<b>34.70</b>	<b>31.27</b>
ViT-B/16	AT-BSL	10.00	10.00	10.00	10.00
	TAET	<b>22.65</b>	<b>18.45</b>	<b>18.56</b>	<b>14.72</b>

**Table 11.** Comparison results with AT-BSL under different models on CIFAR-10-LT(IR = 10), with better results highlighted in bold.

## Appendix C. Time Efficiency Analysis

### C.1. Total Training Time Comparison

The total time consumption for TAET (Two-Stage Adversarial Training) and standard AT (Adversarial Training) can be expressed as:

$$\begin{aligned} T_{\text{TAET}} &= N_{\text{CE}} \cdot (F + B) + N_{\text{AT}} \cdot [F + B + \kappa \cdot (F + B_{\text{adv}})] \\ T_{\text{AT}} &= N_{\text{total}} \cdot [F + B + \kappa \cdot (F + B_{\text{adv}})] \end{aligned} \quad (12)$$

where:

- $N_{\text{CE}} = 40$ : Number of epochs for the Cross-Entropy (CE) stage.
- $N_{\text{AT}} = 60$ : Number of epochs for the Adversarial Training (AT) stage.
- $N_{\text{total}} = 100$ : Total number of training epochs.

### C.2. Acceleration Ratio Derivation

The time acceleration factor is defined as:

$$\eta = \frac{T_{\text{TAET}}}{T_{\text{AT}}} = \frac{N_{\text{CE}} \cdot \rho + N_{\text{AT}} \cdot (1 + \kappa \cdot \gamma)}{N_{\text{total}} \cdot (1 + \kappa \cdot \gamma)} \quad (13)$$

where:

- $\rho = \frac{F+B}{F+B+\kappa \cdot (F+B_{\text{adv}})} \approx 0.139$ : Relative efficiency during the CE stage
- $\gamma = \frac{F+B_{\text{adv}}}{F+B} \approx 0.619$ : Adversarial training time expansion coefficient

Substituting the parameters, we get:

$$\eta \approx \frac{40 \times 0.139 + 60 \times (1 + 10 \times 0.619)}{100 \times (1 + 10 \times 0.619)} \approx 0.608 \quad (15)$$

This indicates that TAET reduces the theoretical computational workload by approximately 39.2%.

## Appendix D. Space Efficiency Analysis

### D.1. Memory Requirement Model

The peak memory requirement during training is given by:

$$M = M_{\text{model}} + M_{\text{data}} + M_{\text{grad}} + \delta \cdot M_{\delta} \quad (15)$$

where:

- $M_{\delta} = b \cdot c \cdot h \cdot w \cdot d \approx 1629.42\text{MB}$ : Memory for adversarial components
- $b = 128$ : Batch size.
- $c = 3$ : Number of channels.
- $h = w = 32$ : Height and width of the input data.
- $d = 4$  bytes: Data type size.
- $\delta \in \{0, 1\}$ : Adversarial training activation flag.

### D.2. Memory Requirement Comparison

The memory difference between TAET and AT can be quantified as:

$$\Delta M = M_{\text{TAET}} - M_{\text{AT}} = \begin{cases} -1629.42\text{MB} & \text{CE stage} \\ 0 & \text{AT stage} \end{cases} \quad (16)$$

### D.3. Effective Memory Savings

The memory saving efficiency is defined as:

$$\xi = \frac{T_{\text{CE}}}{T_{\text{total}}} \cdot M_{\delta} = \frac{40}{100} \times 1629.42 = 651.77 \text{ MB} \quad (17)$$

This indicates that TAET reduces the memory occupation by 1629.42 MB during 40% of the training time.

## Appendix E. Algorithmic Framework for Two-Stage Adversarial Training

The following pseudocode outlines the implementation of Two-Stage Adversarial Equalization Training (TAET). This algorithm effectively balances natural accuracy and adversarial robustness through staged optimization of cross-entropy and hierarchical adversarial loss. By initially focusing on natural accuracy and later shifting to robustness objectives, TAET ensures a dynamic equilibrium between performance and defense. Its efficient design also reduces computational overhead compared to single-stage approaches, making it a practical solution for adversarial training.

---

### Algorithm 1: Two-Stage Adversarial Equalization Training

---

**Require:** Classifier  $F$ , clean samples  $x_0$  and corresponding labels  $y$ , number of classes  $K$ , optimizer, step size  $\eta$ , perturbation size  $\epsilon$ , number of PGD attack steps  $n$ , total training epochs  $T$ , epochs for standard cross-entropy training  $T_{ce}$ , loss balance weights  $\alpha, \beta, \gamma$

---

```

1:  for  $t = 1$  to  $T_{ce}$  do
2:    Forward pass  $x_0$  through  $F$ 
3:    Compute cross-entropy loss:
4:     $\mathcal{L}_{ce} = \text{CrossEntropyLoss}(F(x_0), y)$ 
5:    Backpropagate and update  $F$  to minimize  $\mathcal{L}_{ce}$ 
6:  End for
7:  for  $t = T_{ce} + 1$  to  $T$  do
      Initialize adversarial sample:
8:     $x' = x_0 + 0.001 \cdot \mathcal{N}(0, 1)$ 
9:    for  $j = 1$  to  $n$  do
      Compute adversarial loss on current adversarial sample:
10:    $\mathcal{L}_{ce} = \text{CrossEntropyLoss}(F(x'), y)$ 
      Update adversarial sample via gradient ascent:
       $\nabla x' = \text{grad}(\mathcal{L}_{ce}, x')$ 
11:    $x' = \text{Clip}(x' + \eta \cdot \text{sign}(\nabla x'), x_0 - \epsilon, x_0 + \epsilon)$ 
12:   End for
      Generate final adversarial example:
13:    $x' = \text{Clip}(x', 0.0, 1.0)$ 
14:   Forward pass adversarial sample  $x'$  through  $F$ 
      Compute Hierarchical Equalization Loss (HEL):
15:    $\mathcal{L}_{HEL} = \alpha \cdot \mathcal{L}_{bal} + \beta \cdot \mathcal{L}_{hie} + \gamma \cdot \mathcal{L}_{rare}$ 
16:   Backpropagate and update  $F$  to minimize  $\mathcal{L}_{HEL}$ 
17: End for

```

---

## Appendix F. Adversarial Attacks

The **Fast Gradient Sign Method** generates adversarial examples by perturbing the input data in the direction of the gradient of the loss function with respect to the input. The perturbation is controlled by a small step size  $\epsilon$ . The formula for FGSM is given by:

$$x' = x + \epsilon \cdot \text{sign}(\nabla_x J(\theta, x, y)) \quad (18)$$

where  $x$  is the original input image,  $\epsilon$  is the perturbation magnitude, and  $\nabla_x J(\theta, x, y)$  is the gradient of the loss function  $J(\theta, x, y)$  with respect to the input  $x$ , where  $\theta$  represents the model parameters and  $y$  is the true label. The adversarial example  $x'$  is created by adding the perturbation  $\epsilon$  to the input.

**Projected Gradient Descent (PGD)** attack is an iterative version of FGSM, where perturbations are applied multiple times in small steps. After each step, the perturbed image is projected back into a feasible region to ensure that the perturbations stay within a specified bound. The update rule for PGD is given by:

$$x^{t+1} = \Pi_{x+\mathcal{B}}(x^t + \alpha \cdot \text{sign}(\nabla_x J(\theta, x^t, y))) \quad (19)$$

where  $\alpha$  is the step size,  $\mathcal{B}$  is the set of allowed perturbations (typically the  $L_\infty$  ball), and  $\Pi$  denotes the projection operator that ensures the adversarial example stays within the allowed perturbation space.

**Carlini-Wagner (CW)** attack is an optimization-based method that seeks the minimal perturbation necessary to misclassify an input. It achieves this by minimizing a loss function that balances the perturbation size and the model's confidence in the incorrect class. The objective function for the CW attack is:

$$\begin{aligned} \min_{\delta} \quad & \|\delta\|_2 + c \cdot f(x + \delta) \\ \text{s.t.} \quad & x + \delta \in [0, 1]^n \end{aligned} \quad (20)$$

Here,  $\delta$  is the perturbation added to the original input  $x$ ,  $\|\delta\|_2$  denotes the  $L_2$  norm of the perturbation, and  $c$  is a constant that balances the importance of the perturbation size against the attack success. The function  $f(x + \delta)$  is defined as:

$$f(x + \delta) = \max_{i \neq t} (Z(x + \delta)_i) - Z(x + \delta)_t, -\kappa \quad (21)$$

In this function,  $Z(x + \delta)_i$  represents the logits (pre-softmax outputs) of the model for class  $i$ ,  $t$  is the target class in a targeted attack, and  $\kappa$  (kappa) is a confidence parameter that controls the margin by which the adversarial example's logit for the target class must exceed that of any other class. The constraint  $x + \delta \in [0, 1]^n$  ensures that the perturbed input remains within valid data bounds. The goal is

to find the smallest perturbation  $\delta$  that causes the model to misclassify the input as the target class  $t$  with high confidence.

AutoAttack (AA) is a state-of-the-art adversarial attack method that combines multiple attack strategies to generate robust adversarial examples. AA includes attack methods such as APGD, APGDT, FAB, and Square. Each attack method in AutoAttack is designed to explore different adversarial strategies, making it a versatile and powerful tool for generating adversarial examples across a wide range of models.

## Acknowledgments

This work was supported in part by the Hefei Municipal Natural Science Foundation under Grant(No.HZR2403), Natural Science Research Project of Colleges and Universities in Anhui Province (No.2022AH051889, No.2308085QF227), the Fundamental Research Funds for the Central Universities and the Fundamental Research Funds for the Central Universities of China.

## Footnotes

<sup>1</sup> <https://github.com/NISPLab/AT-BSL>

<sup>2</sup> <https://github.com/wutong16/AdversarialLong-Tail>

<sup>3</sup> <https://github.com/MadryLab/cifar10challenge>

<sup>4</sup> <https://github.com/yaodongyu/TRADES>

<sup>5</sup> <https://github.com/NISPLab/AT-BSL>

<sup>6</sup> <https://github.com/zjfheart/Geometry-aware-Instance-reweightedAdversarial-Training>

<sup>7</sup> <https://github.com/jiaxiaojunqiang/las-at>

<sup>8</sup> <https://github.com/GuanlinLee/REAT>

## References

- <sup>1</sup> [Ittannavar SS, Khot BP, Mahadevaswamy M, Havaladar RH. "Deep Learning for Computer Vision: Recent Breakthroughs and Emerging Trends." In: 2023 ICCES, 2023.](#)
- <sup>2</sup> [Szegedy C \(2013\). "Intriguing properties of neural networks". arXiv preprint arXiv:1312.6199.](#)

3. <sup>a, b, c</sup>Goodfellow IJ, Shlens J, Szegedy C (2014). "Explaining and harnessing adversarial examples". arXiv pre print arXiv:1412.6572.
4. <sup>^</sup>Long T, Gao Q, Xu L, Zhou Z (2022). "A survey on adversarial attacks in computer vision: Taxonomy, visualization and future directions". *Computers & Security*. 2022.
5. <sup>a, b, c, d, e</sup>Zhang H, Yu Y, Jiao J, Xing E, El Ghaoui L, Jordan M (2019). "Theoretically principled trade-off between robustness and accuracy". In: *ICML*.
6. <sup>a, b, c, d</sup>Pang T, Yang X, Dong Y, Xu K, Zhu J, Su H (2020). "Boosting adversarial training with hypersphere embedding". *NeurIPS*. 2020.
7. <sup>a, b, c, d, e</sup>Zhang J, Zhu J, Niu G, Han B, Sugiyama M, Kankanhalli M. Geometry-aware instance-reweighted adversarial training. In: *ICLR*; 2021.
8. <sup>a, b, c, d, e</sup>Wang Y, Zou D, Yi J, Bailey J, Ma X, Gu Q (2019). "Improving adversarial robustness requires revisiting misclassified examples". In: *ICLR*.
9. <sup>a, b, c, d, e, f, g, h, i, j</sup>Madry A, Makelov A, Schmidt L, Tsipras D, Vladu A (2018). "Towards Deep Learning Models Resistant to Adversarial Attacks". In: *ICLR*, 2018.
10. <sup>^</sup>Krizhevsky A, Hinton G, et al. (2009). "Learning multiple layers of features from tiny images". Toronto, ON, Canada.
11. <sup>^</sup>Deng J, Dong W, Socher R, Li LJ, Li K, Fei-Fei L (2009). "Imagenet: A large-scale hierarchical image database". In: *CVPR*.
12. <sup>a, b, c, d, e</sup>Zhang Y, Kang B, Hooi B, Yan S, Feng J (2023). "Deep long-tailed learning: A survey". *IEEE TPAMI*. 2023.
13. <sup>a, b, c</sup>Kang B, Xie S, Rohrbach M, Yan Z, Gordo A, Feng J, Kalantidis Y (2020). "Decoupling Representation and Classifier for Long-Tailed Recognition". In: *ICLR*.
14. <sup>a, b, c, d, e, f, g, h, i, j, k</sup>Wu T, Liu Z, Huang Q, Wang Y, Lin D (2021). "Adversarial robustness under long-tailed distribution". In: *CVPR*, 2021.
15. <sup>a, b</sup>Liu Z, Miao Z, Zhan X, Wang J, Gong B, Yu SX. "Large-scale long-tailed recognition in an open world." In: *CVPR*; 2019.
16. <sup>^</sup>Alshammari S, Wang YX, Ramanan D, Kong S (2022). "Long-tailed recognition via weight balancing". In: *CVPR*.
17. <sup>^</sup>Tan J, Wang C, Li B, Li Q, Ouyang W, Yin C, Yan J (2020). "Equalization loss for long-tailed object recognition". In: *CVPR*.

18. <sup>△</sup>Megahed FM, Chen YJ, Megahed A, Ong Y, Altman N, Krzywinski M (2021). "The class imbalance problem". *Nat Methods*. 2021.
19. <sup>a, b, c, d, e, f, g, h, i, j, k, l, m</sup>Yue X, Mou N, Wang Q, Zhao L (2024). "Revisiting Adversarial Training under Long-Tailed Distributions". In: *CVPR, 2024*.
20. <sup>a, b, c, d, e, f, g</sup>Li G, Xu G, Zhang T (2023). "Adversarial training over long-tailed distribution". *arXiv preprint arXiv:2307.10205*. Available from: <https://arxiv.org/abs/2307.10205>.
21. <sup>△</sup>Yang L, Jiang H, Song Q, Guo J (2022). "A survey on long-tailed visual recognition". *IJCV*.
22. <sup>a, b</sup>Qi T, Xie H, Li P, Ge J, Zhang Y (2023). "Balanced classification: A unified framework for long-tailed object detection". *IEEE TMM*. 2023.
23. <sup>a, b</sup>Rice L, Wong E, Kolter Z (2020). "Overfitting in adversarially robust deep learning". In: *ICML*.
24. <sup>△</sup>Chen W, Yang K, Yu Z, Shi Y, Chen CL (2024). "A survey on imbalanced learning: latest research, applications and future directions". *Artificial Intelligence Review*. 2024.
25. <sup>△</sup>Chang N, Yu Z, Wang YX, Anandkumar A, Fidler S, Alvarez JM. "Image-level or object-level? a tale of two resampling strategies for long-tailed detection." In: *ICML, 2021*.
26. <sup>△</sup>Baloch BK, Kumar S, Haresh S, Rehman A, Syed T (2019). "Focused anchors loss: Cost-sensitive learning of discriminative features for imbalanced classification". In: *Asian Conference on Machine Learning*.
27. <sup>△</sup>Cong C, Xuan S, Liu S, Zhang S, Pagnucco M, Song Y. "Decoupled optimisation for long-tailed visual recognition." In: *AAAI, 2024*.
28. <sup>△</sup>Wei XS, Xu SL, Chen H, Xiao L, Peng Y (2022). "Prototype-based classifier learning for long-tailed visual recognition". *Science China Information Sciences*. 2022.
29. <sup>△</sup>Zhou Z, Li L, Zhao P, Heng PA, Gong W (2023). "Class-conditional sharpness-aware minimization for deep long-tailed recognition". In: *CVPR*.
30. <sup>△</sup>Li J, Meng Z, Shi D, Song R, Diao X, Wang J, Xu H (2023). "Fcc: Feature clusters compression for long-tailed visual recognition". In: *CVPR*.
31. <sup>△</sup>Papernot N, McDaniel P (2017). "Extending defensive distillation". *arXiv preprint arXiv:1705.05264*. Available from: <https://arxiv.org/abs/1705.05264>.
32. <sup>△</sup>Pang T, Xu K, Du C, Chen N, Zhu J (2019). "Improving adversarial robustness via promoting ensemble diversity". In: *ICML*.
33. <sup>△</sup>Rebuffi SA, Gowal S, Calian DA, Stimberg F, Wiles O, Mann TA (2021). "Data augmentation can improve robustness". *NeurIPS*. 2021.

34. <sup>△</sup>Hendrycks D, Mazeika M, Kadavath S, Song D (2019). "Using self-supervised learning can improve model robustness and uncertainty". *NeurIPS*.
35. <sup>△</sup>Xie C, Wu Y, van der Maaten L, Yuille AL, He K (2019). "Feature denoising for improving adversarial robustness". In: *CVPR*.
36. <sup>△</sup>Ben-Eliezer O, Yogev E (2020). "The adversarial robustness of sampling". *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*.
37. <sup>△</sup>Chen EC, Lee CR (2024). "Data filtering for efficient adversarial training". *PR*.
38. <sup>△</sup><sub>a, b, c</sub>Dong Y, Deng Z, Pang T, Zhu J, Su H (2020). "Adversarial distributional training for robust deep learning". *NeurIPS*.
39. <sup>△</sup><sub>a, b, c</sub>Jia X, Zhang Y, Wu B, Ma K, Wang J, Cao X (2022). "LAS-AT: adversarial training with learnable attack strategy". In: *CVPR*.
40. <sup>△</sup>Zhao Y, Chen S, Liu S, Hu Z, Xia J (2024). "Hierarchical equalization loss for long-tailed instance segmentation". *IEEE TMM*. 2024. Published by IEEE.
41. <sup>△</sup>Ma JW, Liang M, Chen L, Tian S, Chen SL, Qin J, Yin XC (2023). "Sample weighting with hierarchical equalization loss for dense object detection". *IEEE TMM*. Published by IEEE.
42. <sup>△</sup><sub>a, b</sub>Mao A, Mohri M, Zhong Y (2023). "Cross-entropy loss functions: Theoretical analysis and applications". In: *ICML*.
43. <sup>△</sup><sub>a, b, c</sub>Brodersen KH, Ong CS, Stephan KE, Buhmann JM (2010). "The balanced accuracy and its posterior distribution". In: *ICPR*. pp. 3121–3124.
44. <sup>△</sup>Kozerawski J, Fragoso V, Karianakis N, Mittal G, Turk M, Chen M (2020). "BlT: Balancing long-tailed datasets with adversarially-perturbed images". In: *ACCV*.
45. <sup>△</sup><sub>a, b</sub>Yang J, Shi R, Wei D, Liu Z, Zhao L, Ke S (2023). "Medmnist v2—a large-scale lightweight benchmark for 2d and 3d biomedical image classification". *arXiv preprint arXiv:2301.08658*. [arXiv:2301.08658](https://arxiv.org/abs/2301.08658).
46. <sup>△</sup>Tschandl P, Rosendahl C, Kittler H (2018). "The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions". *Scientific Data*. 2018.
47. <sup>△</sup><sub>a, b</sub>Carlini N, Wagner D (2017). "Towards evaluating the robustness of neural networks". In: *IEEE SP*.
48. <sup>△</sup><sub>a, b, c</sub>Croce F, Hein M (2020). "Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks". In: *ICML*.
49. <sup>△</sup>Yu C, Han B, Shen L, Yu J, Gong C, Gong M, Liu T (2022). "Understanding robust overfitting of adversarial training and beyond". In: *ICML*.

50. He K, Zhang X, Ren S, Sun J (2016). "Deep residual learning for image recognition". In: CVPR.

## Declarations

**Funding:** No specific funding was received for this work.

**Potential competing interests:** No potential competing interests to declare.