

LIFT+: Lightweight Fine-Tuning for Long-Tail Learning

Jiang-Xin Shi, Tong Wei, and Yu-Feng Li, *Senior Member, IEEE*

Abstract—The fine-tuning paradigm has emerged as a prominent approach for addressing long-tail learning tasks in the era of foundation models. However, the impact of fine-tuning strategies on long-tail learning performance remains unexplored. In this work, we disclose that existing paradigms exhibit a profound misuse of fine-tuning methods, leaving significant room for improvement in both efficiency and accuracy. Specifically, we reveal that heavy fine-tuning (fine-tuning a large proportion of model parameters) can lead to non-negligible performance deterioration on tail classes, whereas lightweight fine-tuning demonstrates superior effectiveness. Through comprehensive theoretical and empirical validation, we identify this phenomenon as stemming from inconsistent class conditional distributions induced by heavy fine-tuning. Building on this insight, we propose LIFT+, an innovative lightweight fine-tuning framework to optimize consistent class conditions. Furthermore, LIFT+ incorporates semantic-aware initialization, minimalist data augmentation, and test-time ensembling to enhance adaptation and generalization of foundation models. Our framework provides an efficient and accurate pipeline that facilitates fast convergence and model compactness. Extensive experiments demonstrate that LIFT+ significantly reduces both training epochs (from ~ 100 to ≤ 15) and learned parameters (less than 1%), while surpassing state-of-the-art approaches by a considerable margin. The source code is available at <https://github.com/shijxcs/LIFT-plus>.

Index Terms—Long-tail learning, foundation model, class-imbalanced learning, vision-language model.




Fig. 1: Comparison of LIFT+ with state-of-the-art methods (using CLIP as foundation model). The x-axis indicates the number of training epochs, and the y-axis represents the test accuracy. The size of each marker reflects the quantity of learned parameters. Gray labels denote methods that utilize external data. LIFT+ consistently achieves higher performance with lower costs and is even comparable with methods that leverage external data.

1 INTRODUCTION

LONG-TAIL learning tackles the challenge of learning from highly imbalanced data, where a small subset of dominant classes (head classes) possess abundant training samples, while the remaining classes (tail classes) have only limited training samples. The inherent challenge posed by long-tail data is prevalent across various scenarios, such as image classification [1], [2], [3], instance segmentation

[4], [5], [6], and object detection [7], [8], [9]. This problem has attracted considerable research attention, leading to numerous approaches proposed to enhance generalization, particularly for tail classes. Existing approaches can be categorized into three main paradigms: i) data manipulation techniques [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], ii) representation learning strategies [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], and iii) model output adjustment mechanisms [32], [33], [34], [35], [36], [37], [38], [39], [40], [41]. Despite the considerable progress achieved by these methods, a significant performance gap persists compared to models trained on rich data, indicating room for substantial improvement.

Recent studies in long-tail learning have shown that fine-tuning foundation models such as CLIP [42] and ViT [43] can substantially improve performance, eliminating the necessity of training deep neural networks from scratch. For instance, BALLAD [44], VL-LTR [45], Wang et al. [46] employ the CLIP model, while RAC [47] and LPT [48]

- Jiang-Xin Shi and Yu-Feng Li are with the National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China, and the School of Artificial Intelligence, Nanjing University, Nanjing 210023, China.
E-mail: {shijx, liyf}@lamda.nju.edu.cn.
- Tong Wei is with the School of Computer Science and Engineering, Southeast University, Nanjing 210096, China, and the Key Laboratory of Computer Network and Information Integration, Southeast University, Ministry of Education, China.
E-mail: weit@seu.edu.cn.

Manuscript received XXXX, 20XX; revised XXXX, 20XX.
(Corresponding author: Yu-Feng Li.)

adopt the ViT model pre-trained on ImageNet-21K. However, although these works introduce a new paradigm for long-tail learning, they fail to provide a systematic analysis of the fine-tuning strategies under long-tail scenarios. This oversight potentially leads to suboptimal methodologies and unnecessary computational overhead. Specifically, existing approaches typically come at the cost of i) prolonged training durations (≈ 100 epochs); ii) a two-stage procedure, which limits the deployment flexibility; iii) large external datasets ($\approx 10^6$ samples) to facilitate the training process.

To bridge this gap, we first provide a comprehensive analysis of fine-tuning strategies for long-tail learning. Our empirical evidence reveals that heavy fine-tuning (fine-tuning a large proportion of model parameters) can lead to *non-negligible performance deterioration* on tail classes. Through systematic investigation, we identify that full fine-tuning distorts the intrinsic intra-class distance distributions. Theoretically, we prove that such distortions break the consistency assumption of class-conditional distributions, consequently resulting in biased predictions. Building upon these theoretical and empirical findings, we recognize that optimizing a small proportion of pre-trained weights could simultaneously enhance model discriminability while preserving the intra-class distributions. Motivated by this insight, we propose an efficient and accurate long-tail learning framework named *Lightweight Fine-Tuning* (LIFT+). The proposed approach facilitates rapid convergence and compact models through adaptive lightweight fine-tuning.

LIFT+ is a single-stage framework that usually achieves convergence in fewer than 15 training epochs without requiring external data. Furthermore, we improve LIFT+ with semantic-aware classifier initialization, which provides a robust starting point for the model, enabling better generalization while maintaining computational efficiency. LIFT+ also introduces a minimalist data augmentation strategy to eliminate trivial augmentation approaches and further accelerate the training period. Additionally, we incorporate test-time ensembling to mitigate the inherent limitations of foundation models and boost generalization capability. Figure 1 gives a performance comparison on three typical long-tail datasets, ImageNet-LT, Places-LT, and iNaturalist 2018. Notably, LIFT+ consistently outperforms state-of-the-art methods without auxiliary data by an average margin of 2.1% in accuracy, while requiring substantially fewer learned parameters and reduced training epochs.

In conclusion, the contributions of this paper can be summarized as follows:

- We first reveal the critical limitation of heavy fine-tuning that distorts the tail-class performance;
- Through empirical and theoretical validation, we discover that optimizing a small proportion of parameters can effectively alleviate the performance degradation;
- We propose LIFT+, an efficient and accurate long-tail learning framework that leverages lightweight fine-tuning with minimal computational overhead;
- Extensive experiments demonstrate that LIFT+ consistently outperforms state-of-the-art methods with lower computational costs on multiple benchmarks (an average of 2.1% accuracy improvements, with fewer than 15 training epochs and less than 1% learned parameters).

This paper presents an extended version of our previous work, LIFT, which was initially proposed at ICML 2024 [49]. The current version offers four major advancements. First, we introduce LIFT+, a more systematic framework that addresses long-tail learning challenges through input processing, representation enhancement, and output optimization. This unified framework establishes a versatile research pipeline for future studies. Second, we propose a novel minimalist data augmentation strategy that fundamentally overcomes the limitations of conventional approaches, reducing training time by 5 epochs (25-50% of the original training cost) while maintaining superior performance. Third, we extend our analysis to incorporate more foundation models as baselines, including ImageNet-21K pre-trained models, ViT-Large, ResNet and DINO-ViT. We also conduct more comprehensive empirical studies to verify the advantages of our framework. Last but not least, we reconstruct our source code by rectifying some mistakes and further improving the computational efficiency, enabling full reproducibility using a single NVIDIA 4090 GPU with only 24GB of memory. We rerun all experiments and find that the performance is further improved in the updated pipeline.

2 RELATED WORK

Long-Tail Learning via Deep Learning. Traditional approaches typically employ convolutional neural networks such as ResNet and ResNeXt for long-tail learning [50], [51]. To address the long-tail challenge, existing methods follow three primary directions: i) data manipulation [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], which involves designing re-sampling and data augmentation strategies. Interestingly, Shi et al. [15] discover that re-sampling is sensitive to irrelevant context and may not always help enhance representation. Ahn et al. [16] find that data augmentation may adversely affect the augmented class while benefiting the non-augmented classes. ii) representation learning [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], which focuses on improving feature extraction. This direction incorporates various advanced techniques employed, such as supervised contrastive learning [20], [21], [22], [23], [30], [31], self-supervised learning [24], [28], and neural collapse [25], [26], [29]. iii) model output adjustment [32], [33], [34], [35], [36], [37], [38], [39], [40], [41], which optimizes unbiased loss functions during training or applies post-hoc calibrations. Moreover, ensembling learning methods [52], [53], [54], [55], [56] aim to combine multiple diverse experts and optimization objectives to improve both head and tail classes. Additionally, many studies adopt a two-stage training paradigm [57], [58], [59], [60], [61], in which the first stage learns representations and the second stage learns the classifier. In contrast to the aforementioned works, our work presents an end-to-end training approach that leverages the strengths of foundation models. Moreover, we propose novel techniques to enhance input processing, feature representation, and output optimization within a unified framework, providing a versatile and extensible pipeline for future research.

Long-Tail Learning via Foundation Model. Fine-tuning foundation models such as CLIP [42] and ViT [43] has attracted widespread attention [62], [63], [64], [65], [66], and

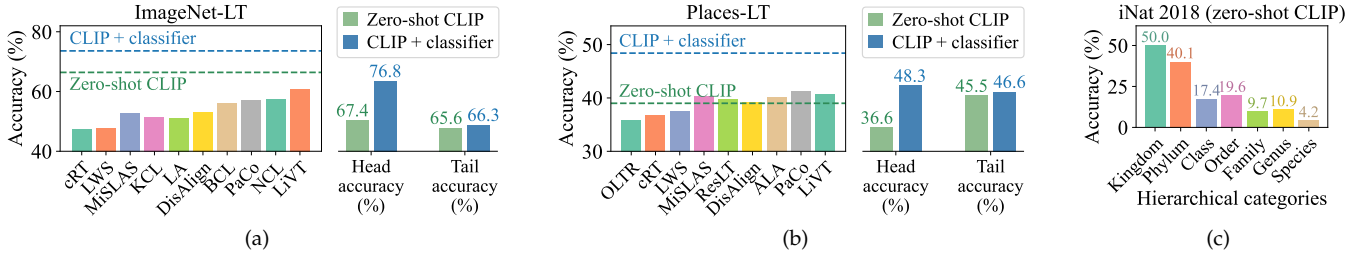


Fig. 2: (a-b) On ImageNet-LT and Places-LT, zero-shot CLIP has surpassed many prior methods. By simply introducing an additional classifier, the accuracy further increases. However, the improvements mainly come from the head classes, while the tail classes only achieve marginal enhancements. (c) On iNaturalist 2018, zero-shot CLIP encounters challenges in achieving high accuracy for fine-grained long-tail categories.

has emerged as an effective strategy to mitigate class imbalance by leveraging their powerful representation learning capabilities [44], [45], [46], [47], [48], [49], [67], [68], [69], [70], [71], [72], [73]. For instance, BALLAD [44] first employs full fine-tuning for the foundation model, then freezes the backbone and optimizes a linear adapter on re-sampled data. VL-LTR [45] first incorporates auxiliary image-text web data to fine-tune the vision-language models, then freezes the text encoder, and optimizes the image encoder along with a language-guided head. RAC [47] jointly trains an encoder with a retrieval module to leverage external datasets (*e.g.*, ImageNet-21K) for input augmentation. LPT [48] adopts a two-phase prompt learning methodology to adapt the foundation model. Wang et al. [46] propose to add a decoder after the pre-trained model to extract relevant features. Chen et al. [73] investigate the biases of foundation models and empirically demonstrate that ensembling different models can effectively improve the performance. Although various approaches have been proposed for adapting foundation models, current research lacks a systematic analysis regarding the impact of fine-tuning strategies on long-tail learning. Moreover, it is important to note that existing methods typically suffer from prolonged training durations and often rely on auxiliary external data. In contrast, our proposed approach exhibits a remarkable capability to achieve fast convergence without requiring external data. Furthermore, our method is versatile and inclusive, allowing for seamless integration with different lightweight fine-tuning methods and various backbones.

3 HOW FINE-TUNING AFFECTS LONG-TAIL

3.1 Preliminary

Unlike conventional neural networks, foundation models that employ the Transformer architecture [43], [74] are more elegantly designed while exhibiting exceptional generalization capabilities. It has demonstrated remarkable adaptability across various computer vision [43], [75] and natural language processing tasks [74], [76]. Taking the image classification task as an example, given an image x , it first partitions the image into multiple patches, then feeds the patches into an extractor ϕ to obtain the corresponding feature $\phi(x)$. For a downstream K -classification task, the model computes the predicted logits for each class as $z_k = \phi(x)w_k^\top + b_k$, where w_k and b_k denote task-specific classifier weights and

bias. Due to space constraints, the comprehensive architecture and inference processes of foundation models are presented in Section 6.

The recent vision-language foundation model, CLIP, further underscores its efficacy by demonstrating impressive zero-shot performance [42]. Given an image x and considering K candidate classes, CLIP first generates K textual prompts t_1, \dots, t_K , each representing a descriptive phrase, such as “a photo of a cat” or “a photo of a dog”. It then extracts the image feature $\phi(x)$ and the prompt features $\psi(t_1), \dots, \psi(t_K)$. To predict the label for the given image, CLIP compares the cosine similarity between the image and each class prompt:

$$y_{\text{pred}} = \arg \max_{k \in [K]} \langle \phi(x)P_I, \psi(t_k)P_T \rangle \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes cosine similarity. P_I and P_T are projection matrices that map the extracted features to a shared latent space with consistent dimensionality.

3.2 Long-Tail Learning with Foundation Model

We first evaluate the performance of CLIP on typical long-tail datasets, as illustrated in Figures 2a to 2c. The results reveal that zero-shot CLIP outperforms most conventional methods. Furthermore, by freezing the backbone and training an additional classifier, the performance can be further improved, which validates the high-quality representations learned by CLIP. However, the majority of performance gains are dominated by the head classes, while the tail classes exhibit only marginal improvements.

Moreover, while zero-shot CLIP exhibits impressive performance in general scenarios, its effectiveness diminishes on specialized long-tail datasets. A notable case is the iNaturalist 2018 dataset, which presents a fine-grained long-tail challenge with a hierarchical categorization spanning from 7 kingdoms to 8142 species. Although zero-shot CLIP achieves high accuracy in predicting coarse-grained categories (*e.g.*, “kingdom” and “phylum”), it performs poorly in classifying fine-grained species, *i.e.*, long-tail classes.

3.3 Heavy Fine-Tuning Hurts

Although the foundation model has shown commendable performance in downstream long-tail learning tasks, it has several limitations. As discussed previously, while CLIP





Fig. 3: Comparison of different fine-tuning manners. Full fine-tuning improves head-class accuracy but severely degrades tail-class performance, even when employing balanced loss and classifier initialization.



(a) Classifier fine-tuning.

(b) Full fine-tuning.

Fig. 4: Inter-class feature similarities (heatmaps) and intra-class distance distributions from tail classes (histograms) on ImageNet-LT. Classifier fine-tuning limits head-class performance due to high inter-class similarities. Full fine-tuning optimizes inter-class similarities but leads to inconsistent distribution between train and test data on tail classes.

achieves significant improvements on head classes, its performance gains on tail classes remain suboptimal. This observation naturally raises a critical question: *Is the current fine-tuning strategy insufficient?*

However, the results in Figure 3 reveal a concerning finding: full fine-tuning enhances head-class accuracy while at the expense of degrading tail-class performance. Note that we have optimized the balanced logit-adjusted loss [34] and applied a balanced classifier initialization (which will be introduced in Section 4.2), the prediction bias still persists. Most strikingly, the performance deterioration of tail classes renders the overall accuracy of full fine-tuning even worse than that of classifier fine-tuning (71.6% vs. 73.6% on ImageNet-LT, and 46.6% vs. 48.4% on Places-LT).

To uncover the reasons behind it, we identify that full fine-tuning may distort the representation of tail classes. To validate this, we quantify the model representation from two perspectives: i) *inter-class feature similarities*, which is measured by the cosine similarities between class mean features; and ii) *intra-class distance distributions*, which is characterized by the cosine similarities between individual samples and their corresponding class mean features. This approach is adopted since the class-conditional distribution cannot be directly observed. We compute these metrics and present the results in Figure 4. Notably, Figure 4b reveals that full fine-tuning enhances feature discriminability by reducing inter-class similarities to approximately zero, rendering the features of different classes nearly orthogonal. However, it also distorts intra-class distributions, leading to a distribution shift between train and test data on tail classes. Consequently, using the fine-tuned model to estimate tail-class data will inevitably result in an underestimated class-conditional probability. We provide further theoretical evidence to analyze this matter.

Proposition 3.1. *The underestimated class-conditional probability $P(\phi(\mathbf{x}) | y = j)$ leads to an underestimated loss on class j and a biased prediction towards other classes.*

Proof. Denote P_s and P_t as the probability distributions in the source (training) and target (test) domains, respectively. In long-tail learning scenarios, $P_s(y)$ follows a long-tail distribution and $P_t(y)$ follows a uniform distribution, *i.e.*, $P_t(y = k) \equiv 1/K$. Consequently, we have

$$\begin{aligned} P_s(y = j | \phi(\mathbf{x})) &= \frac{P_s(y = j | \phi(\mathbf{x}))}{\sum_{k=1}^K P_s(y = k | \phi(\mathbf{x}))} \\ &= \frac{P_t(y = j | \phi(\mathbf{x})) \cdot \frac{P_s(y = j | \phi(\mathbf{x}))}{P_t(y = j | \phi(\mathbf{x}))}}{\sum_{k=1}^K P_t(y = k | \phi(\mathbf{x})) \cdot \frac{P_s(y = k | \phi(\mathbf{x}))}{P_t(y = k | \phi(\mathbf{x}))}} \\ &= \frac{P_t(y = j | \phi(\mathbf{x})) \cdot \frac{P_s(\phi(\mathbf{x}) | y = j)}{P_t(\phi(\mathbf{x}) | y = j)} \cdot \frac{P_s(y = j)}{P_t(y = j)}}{\sum_{k=1}^K P_t(y = k | \phi(\mathbf{x})) \cdot \frac{P_s(\phi(\mathbf{x}) | y = k)}{P_t(\phi(\mathbf{x}) | y = k)} \cdot \frac{P_s(y = k)}{P_t(y = k)}} \\ &= \frac{P_t(y = j | \phi(\mathbf{x})) \cdot P_s(y = j) \cdot \zeta_{s-t}(j)}{\sum_{k=1}^K P_t(y = k | \phi(\mathbf{x})) \cdot P_s(y = k) \cdot \zeta_{s-t}(k)} \quad (2) \end{aligned}$$

$$\begin{aligned} \mathcal{L}(\phi(\mathbf{x}), y = j) &= -\log P_s(y = j | \phi(\mathbf{x})) \\ &= -\log \frac{\exp(z_j + \log P_s(y = j) + \log \zeta_{s-t}(j))}{\sum_{k=1}^K \exp(z_k + \log P_s(y = k) + \log \zeta_{s-t}(k))} \quad (3) \end{aligned}$$

where $\zeta_{s-t}(j) = \frac{P_s(\phi(\mathbf{x}) | y = j)}{P_t(\phi(\mathbf{x}) | y = j)}$ and z_j is the predicted logit on class j . For samples \mathbf{x} belonging to class j , the underestimated $P_t(\phi(\mathbf{x}) | y = j)$ leads to an underestimation of $\mathcal{L}(\phi(\mathbf{x}), y = j)$, and consequently results in a biased optimization towards other classes. \square

Remark 3.2. Theorem 3.1 reveals that the performance degradation in full fine-tuning stems from inconsistent class-conditional distributions among tail classes. Existing approaches such as Balanced-Softmax [33] and Logit-adjusted loss [34], [35], [77], [78] inherently assume a consistent class-conditional distribution between the source and target domains, *i.e.*, $\zeta_{s-t}(j) = 1$. However, as empirically demonstrated in Figure 4b, full fine-tuning violates this assumption. While an ideal solution would involve estimating the underlying class-conditional distribution, this is infeasible in practice due to the scarcity of tail-class data. Alternatively, another viable approach is to mitigate such distribution distortions. To achieve this goal, we introduce lightweight fine-tuning in Section 4.

To address generalization challenges with long-tail data, recent studies have explored two-stage training procedures [44] or incorporated external training data [45], [47]. However, these strategies typically incur substantial training overhead or require additional data sources, which hinders their practical deployment. In response to this, we introduce LIFT+, an efficient and accurate lightweight fine-tuning framework tailored for long-tail learning.




Fig. 5: Fine-tuning a small proportion of all parameters (e.g., 0.1%-2%) yields superior performance. As the proportion increases, performance deteriorates even when we search for the best learning rate.




Fig. 6: Inter-class feature similarities (heatmaps) and intra-class distributions from tail classes (histograms) on ImageNet-LT. Both arbitrary and structured lightweight fine-tuning perform well in optimizing inter-class similarities and preserving intra-class distributions.

3.4 Lightweight Fine-Tuning Helps

To mitigate intra-class distribution distortion, a straightforward approach is to constrain the number of learned parameters. Formally, for each weight matrix $\mathbf{W} \in \mathbb{R}^{d_1 \times d_2}$ in the foundation model, we optimize a specified proportion α of parameters in \mathbf{W} , while keeping the rest frozen. This approach allows only $\alpha d_1 d_2$ parameters to be optimized. In implementation, a sparse binary mask $\mathbf{M} \in \{0, 1\}^{d_1 \times d_2}$ is employed to select the optimized parameters:

$$\mathbf{X}\mathbf{W} \rightarrow \mathbf{X}(\mathbf{W} \circ \mathbf{M}) + \underbrace{\mathbf{X}(\mathbf{W} \circ (1 - \mathbf{M}))}_{\text{gradient detached}} \quad (4)$$

where $\|\mathbf{M}\|_0 = \alpha d_1 d_2$. In Figure 5, we investigate the impact of varying proportions of fine-tuned parameters. The result demonstrates the advantages of lightweight fine-tuning, since even a small proportion (e.g., 0.1%) yields significant performance improvements. In contrast, as the proportion increases, the performance faces a risk of degradation, thereby underscoring the drawbacks of heavy fine-tuning. Moreover, heavy fine-tuning is sensitive to hyperparameters, as it requires careful searching of the learning rate to achieve optimal results.

It is noteworthy that the optimized parameters are selected in an arbitrary manner, constituting what we term as *arbitrary lightweight fine-tuning*. Despite its simplicity, the performance improvement is remarkably significant. This suggests that lightweight fine-tuning is crucial to enhance long-tail learning even without specific fine-tuning strategies. To gain a deeper understanding, we visualize both inter-class feature similarities and intra-class distributions in Figure 6a. The results demonstrate that arbitrary lightweight fine-tuning achieves feature separability comparable to that of full fine-tuning, while effectively maintaining consis-

tent class conditions between training and test data. Furthermore, prediction accuracy substantiates the efficacy of arbitrary lightweight fine-tuning, as it matches the high performance of full-fine-tuning on head classes and even achieves superior performance on tail classes.

By investigating the above arbitrary strategy, we demonstrate that the key to effective fine-tuning lies in learning a small set of parameters rather than optimizing the entire model. Motivated by this insight, we further explore *structured lightweight fine-tuning*, which learns a small set of task-specific parameters in a structured manner. Some related ideas, although already used in foundation models, are not directly designed for long-tail learning. We presented more detailed introductions in Section 6 due to space constraints. Figure 6b illustrates the potential of a very simple idea of structured lightweight fine-tuning. As shown, this method optimizes inter-class similarities to almost orthogonal while preserving intra-class distributions undistorted. Notably, the performance surpasses arbitrary lightweight fine-tuning on both head and tail classes. These results indicate that the well-designed lightweight module can further enhance the feature separability as well as maintain the consistency of class-conditional distributions, and will be more accurate.

4 LIFT+: EFFICIENT AND ACCURATE LONG-TAIL LEARNING FRAMEWORK

4.1 The Proposed Fine-Tuning Method

Our analyses justify that lightweight fine-tuning can significantly alleviate performance degradation and enhance generalization. Based on this finding, we propose *Lightweight Fine-Tuning* (LIFT+), an efficient and accurate framework for long-tail learning. The LIFT+ framework is versatile and inclusive, accommodating the integration of various lightweight fine-tuning methods, including arbitrary lightweight fine-tuning proposed in Section 3.4, and various structured lightweight fine-tuning methods, such as BitFit, VPT, Adapter, LoRA, and AdaptFormer. Although these methods have shown efficacy in foundation models, their performance remains suboptimal to full fine-tuning in previous studies [79], [80], [81], [82], [83]. In contrast, our work first identifies the limitations of full fine-tuning, thereby revealing the untapped potential of these lightweight fine-tuning methods in long-tail learning. Due to space constraints, the technical details are presented in Section 6.

Crucially, lightweight fine-tuning maintains the consistency of class-conditional distributions. This property enables us to adopt an unbiased LA loss for optimization:

$$\mathcal{L}(\mathbf{x}, y = j) = -\log \frac{\exp(z_j + \log P(y = j))}{\sum_{k \in [K]} \exp(z_k + \log P(y = k))} \quad (5)$$

where $y = j$ represents the ground-truth label of \mathbf{x} and z_j denotes the predicted logit. $P(y = j)$ signifies the class prior probability, which can be estimated from the training data.

In order to have better adaptability among tasks, in addition to lightweight fine-tuning, a classifier is essential for discerning the features for different tasks. The linear classifier is commonly adopted due to its simplicity and versatility. Given a feature vector $\phi(\mathbf{x})$, the predicted logit for class j is computed as $z_j = \phi(\mathbf{x})\mathbf{w}_j^\top + b_j$. However,




Fig. 7: Left: LIFT+ with different fine-tuning methods. Right: A simple illustration of the proposed LIFT+ framework.




Fig. 8: An overview of the proposed semantic-aware initialization. The text encoder is discarded after a single forward pass to initialize the classifier.

when training with long-tail data, the classifier weight norms $\|w_k\|_2$ ($1 \leq k \leq K$) tend to exhibit an imbalanced distribution, leading to biased predictions [57], [84]. To overcome this issue and draw on the strengths of CLIP, which optimizes cosine distances in feature space, inspired by [84], we propose to optimize an enhanced cosine classifier $z_j = \sigma \cdot \langle \phi(x), w_j \rangle$, where σ is a scaling factor. This design inherently eliminates the influence of the classifier norm by employing cosine similarity, thereby effectively mitigating biased predictions.

4.2 Semantic-Aware Initialization

It has been demonstrated that an uninitialized classifier can adversely affect model fine-tuning [85]. Therefore, it is crucial to set an appropriate initial state for the classifier. A straightforward method is to apply linear probing by optimizing balanced loss. Another approach is to utilize the class mean features. However, these two approaches not only require extracting features of training data, but are also unreliable with scarce tail-class data. To overcome it, we propose to leverage the semantic knowledge from the text modality of CLIP. Specifically, we generate hand-crafted textual prompts (e.g., "a photo of a [CLASS].") and compute their text features $\psi(t_1), \dots, \psi(t_K)$, which are then utilized to initialize the classifier weights as $w_j = \psi(t_j)P_T P_I^T$. In this way, we can converse image-text matching $z_j = \langle \phi(x)P_I, \psi(t_j)P_T \rangle$ to feature-classifier matching $z_j = \langle \phi(x), w_j \rangle$, where the initial performance of the model is still close to CLIP. We call this way *semantic-aware initialization* (SAI). Unlike prior methods that fine-tune both the image encoder and the text encoder in optimization processes [44], [45], SAI relies on a single forward pass of the text encoder for each class description. After that, the text encoder is discarded. This simple approach allows us to achieve a better initial state of the classifier with small




Fig. 9: (a) Comparison of conventional and minimalist data augmentation. (b) The cropped scale scheduling function.

Algorithm 1 Minimalist Data Augmentation

Input: Image x (resolution $h \times w$), cropped scale range (λ_0, λ_1) , epoch $t \in [T]$, scheduling function $g(\cdot) : [0, 1] \rightarrow [0, 1]$.
Output: Image x' (resolution $a \times a$).
 1: $\Delta\lambda_0 = (\lambda_1 - \lambda_0)g(\frac{t-1}{T-1})$.
 2: Sample $\lambda \sim (\lambda_0 + \Delta\lambda_0, \lambda_1)$.
 3: $h' = w' = \sqrt{hw\lambda}$.
 4: Randomly crop x' ($h' \times w'$) from x .
 5: Resize x' to $a \times a$.

computational overhead. Figure 8 gives an overview of the proposed semantic-aware initialization.

4.3 Minimalist Data Augmentation

The above methods incorporate tailored strategies at both the model representation and the output levels, including a lightweight fine-tuning module, an unbiased logit-adjusted loss, an enhanced cosine classifier, and a semantic-aware initialization method. However, such a framework still uses conventional data augmentation [86], [87], [88], [89] at the input level. The conventional data augmentation strategy is proposed to train a deep model from scratch, which consists of random cropping, stretching, resizing, and flipping of a given image. The purpose is to generate more diverse samples and to cover the underlying data distribution. However, when it comes to fine-tuning foundation models, the key point is not to simulate richer data, but to adapt more efficiently and precisely to downstream data. Moreover, conventional data augmentation faces the risk of generating distorted or trivial images, hindering robust generalization, especially for tail classes where training samples are severely limited. In our learning framework, the number of training epochs is quite small (e.g., 5-15), which means that an image is augmented only a few times, and the quality of augmentation is quite crucial.

To address this dilemma, we further introduce a novel *minimalist data augmentation* (MDA) strategy. Specifically, we

Algorithm 2 Test-Time Ensembling**Input:** Image \mathbf{x} (resolution $a \times a$), expanded size e .**Output:** Predicted logits \mathbf{z} .

- 1: Resize \mathbf{x} to \mathbf{x}' sized $(a + e) \times (a + e)$.
- 2: Crop the center $a \times a$ region of \mathbf{x}' , denoted by \mathbf{x}^c .
- 3: Split \mathbf{x}^c evenly into m patches $[\mathbf{x}_1^p; \dots; \mathbf{x}_m^p]$.
- 4: Calculate the feature $\phi(\mathbf{x}^c)$ and then the logits \mathbf{z}^c .
- 5: Crop the $a \times a$ region from each corner (top left, top right, bottom left, and bottom right) of \mathbf{x}' , repeat procedure 3-4 and obtain the corresponding logits $\mathbf{z}^{tl}, \mathbf{z}^{tr}, \mathbf{z}^{bl}, \mathbf{z}^{br}$.
- 6: $\mathbf{z} = \text{Average}([\mathbf{z}^c, \mathbf{z}^{tl}, \mathbf{z}^{tr}, \mathbf{z}^{bl}, \mathbf{z}^{br}])$.

eliminate stretching and flipping operations and only crop images with consistent proportions to ensure image fidelity. In addition, we design a scheduling function $g(\cdot)$ to control the scale of cropped images from local regions to the global level based on the training process. Through this approach, we ensure that critical information is not neglected, especially at the end of training. As illustrated in Figure 9a, conventional data augmentation generates many redundant or misleading training samples, while the proposed MDA is minimalist and essential. The detailed procedure is presented in Algorithm 1. Following the implementation of multiple previous works [86], [87], [88], [89] and the official setting of PyTorch¹, we set $(\lambda_0, \lambda_1) = (0.08, 1)$. The scheduling function is set as a simple circular convex function $g(t) = 1 - \sqrt{1 - t^2}$ (plot in Figure 9b). The analysis of different scheduling functions is given in Section 5.4. Compared to conventional data augmentation, MDA is even more efficient due to its pruning of redundant steps.

4.4 Test-Time Ensembling

It is well-established that applying random perturbations to each input can lead to improved generalization during tests [90], [91], [92]. Recent works further discover that scaling the test-time computation can improve the output and can be even more effective than scaling the model parameters [93], [94]. This principle is particularly suitable for enhancing Transformer-based models, where an image is split into multiple patches, potentially resulting in the partitioning of continuous patterns into discrete patches. Such partitioning aggravates the data scarcity issues, especially for tail classes with limited data. To enhance generalization robustness, we propose to aggregate the predictions from a set of perturbed versions of the image. Formally, given a test image \mathbf{x} , its predicted logits \mathbf{z} are obtained by averaging the predictions from M perturbed versions:

$$\mathbf{z} = \log \text{P}(\mathbf{y} | \mathbf{x}) = \frac{1}{M} \sum_{i=1}^M \log \text{P}(\mathbf{y} | \alpha_i(\mathbf{x})) \quad (6)$$

Here, $\alpha_i(\mathbf{x})$ represents different perturbed versions of \mathbf{x} . We term this technique *test-time ensembling* (TTE).

In practice, we employ a simple but effective perturbation approach, *i.e.*, using different cropping regions. This approach helps mitigate the bias introduced by image patches. The detailed procedure is presented in Algorithm 2.

¹The official PyTorch implementation of image cropping: <https://pytorch.org/vision/stable/generated/torchvision.transforms.RandomResizedCrop.html>

Conventionally, an image is first resized, center-cropped, and then split into patches before being fed into the Transformer model. However, this approach inevitably leads to the partitioning of important patterns across different patches, thus impeding generalization. By applying diverse croppings, patterns that were separated in one cropping can be preserved in another. It is crucial to emphasize that the expanded size e should not be a multiple of the patch size; otherwise, the five cropped images will share a large portion of the same patches, rendering the expected diversity unattainable. For example, when the patch size is 16, we suggest setting e to 8, 24, 40 rather than 16, 32, 48. We provide an empirical analysis of different expanded sizes in Section 5.4. The proposed TTE strategy can be seamlessly integrated into LIFT+ with minimal computational overhead while consistently improving performance. For fair comparison, we will make necessary statements when equipping any method with TTE.

5 EMPIRICAL STUDY

5.1 Experimental Settings

We evaluate our approach on four widely-used long-tail datasets, including ImageNet-LT [3], Places-LT [3], iNaturalist 2018 [2], and CIFAR-100-LT [32]. ImageNet-LT has 115.8K images from 1000 classes, with class frequencies ranging from 1280 to 5 samples. Places-LT contains 62.5K images from 365 classes, with class frequencies ranging from 4980 to 5 samples. iNaturalist 2018 consists of 437.5K images distributed across 8142 species, with the number of images per species varying from 1000 to 2. CIFAR-100-LT is constructed with multiple imbalance ratio options, including 100, 50, and 10. The models are trained on long-tail datasets and subsequently assessed on their corresponding balanced test datasets. In addition to measuring overall accuracy, we follow the evaluation protocol introduced by [3] to report accuracy among three splits of classes: head classes (containing more than 100 images), medium classes (20 to 100 images), and tail classes (fewer than 20 images).

5.2 Implementation Details

For all experiments, we use the SGD optimizer with a batch size of 128, a weight decay of 5×10^{-4} , and a momentum of 0.9. For LIFT+, the learning rate is 0.02; and for full fine-tuning, we search for the optimal learning rate from $\{0.03, 0.02, 0.01, 0.005, 0.002, 0.001\}$ considering its weak stability. For ImageNet-LT, Places-LT, and CIFAR-100-LT, we train the model for only 5 epochs; and for iNaturalist 2018, we train 15 epochs since it has much more data. We set the bottleneck dimensionality $r = 2^{\lceil \log_2(\frac{K}{2L}) \rceil}$ for Adapter and AdaptFormer such that it learns even fewer parameters than the classifier (please refer to Section 6 for detailed analysis). The scaling factor σ of the cosine classifier is set to 25 (please refer to Table 7 for analysis). The expanded size e for TTE is set to 24 (please refer to Figure 15 for analysis). To ensure the reproducibility, we fix the random seed to 0 (please refer to Table 11 for analysis). For further implementation details, please refer to the source code. All experiments in this paper are conducted on a single NVIDIA 4090 GPU with only 24GB of memory.

TABLE 1: Comparison with state-of-the-arts methods on ImageNet-LT. † denotes methods using external data. Empty cells “-” indicate unreported results in their original papers.

Methods	Backbone	Learned Params.	Epochs	All Head	Med.	Tail
Training from scratch						
cRT [57]	RN-50	23.51M	90+10	47.3	58.8	44.0 26.1
LWS [57]	RN-50	23.51M	90+10	47.7	57.1	45.2 29.3
MiSLAS [58]	RN-50	23.51M	180+10	52.7	62.9	50.7 34.3
KCL [20]	RN-50	23.51M	200	51.5	61.8	49.4 30.9
LA [34]	RN-50	23.51M	90	51.1	-	- -
DisAlign [36]	RN-50	23.51M	90	52.9	61.3	52.2 31.4
RIDE [53]	RN-50	23.51M	100	55.4	66.2	52.3 36.5
rwSAM [24]	RN-50	23.51M	100	55.5	-	- -
BCL [23]	RN-50	23.51M	100	56.0	-	- -
PaCo [22]	RN-50	23.51M	400	57.0	-	- -
NCL [95]	RN-50	23.51M	400	57.4	-	- -
FeatRecon [31]	RN-50	23.51M	90	56.8	-	- -
LiVT [96]	ViT-B/16	85.80M	100	60.9	73.6	56.4 41.0
Fine-tuning foundation model (from CLIP)						
BALLAD [44]	ViT-B/16	149.62M	50+10	75.7	79.1	74.5 69.8
VL-LTR [45]†	ViT-B/16	149.62M	100	77.2	84.5	74.6 59.3
GML [97]†	ViT-B/16	149.62M	100	78.0	-	- -
Decoder [46]	ViT-B/16	21.26M	~18	73.2	-	- -
LIFT+ (Ours)	ViT-B/16	0.62M	5	77.0	79.9	76.1 71.7
w/ TTE (Ours)	ViT-B/16	0.62M	5	78.3	81.0	77.5 73.4

5.3 Comparison with State-of-the-art Methods

Results on ImageNet-LT. We report the test accuracy in Table 1. While existing approaches such as VL-LTR [45] and GML [97] rely on extensive auxiliary data to facilitate fine-tuning, our method LIFT+ achieves competitive performance with significantly lower costs. Reliance on external data not only incurs substantial computational overhead but also limits practical applicability, as such data is often unavailable in real-world scenarios. Notably, LIFT+ requires only 5 epochs of training and fine-tunes far fewer model parameters (*i.e.*, from 21.26M to 0.62M). Compared to methods that do not use auxiliary data, LIFT+ demonstrates even more pronounced advantages, surpassing the best existing method by 1.3% in accuracy. It is worth noting that results using the ImageNet-21K pre-trained foundation model are not reported, considering the class overlap between it and ImageNet-LT. RAC [47] and LPT [48] are not included, since they default to employing the ImageNet-21K pre-trained foundation model, and the results on ImageNet-LT were not reported in their original papers.

Results on Places-LT. The results in Table 2 shows that LIFT+ achieves significantly superior performance over existing methods on Places-LT. In particular, LIFT+ surpasses VL-LTR by 1.4%, despite its use of external training data. When taking a closer look at the performance on tail classes, we can see that LIFT+ achieves particularly significant improvements compared to state-of-the-art methods, *i.e.*, from 48.4 to 51.6. It is worth noting that, when applying to models pre-trained from ImageNet-21K, we use class mean features instead of semantic-aware initialization due to the absence of a corresponding text encoder. Although LPT achieves higher performance when using ImageNet-21K pre-trained model, it needs to adopt a two-stage framework with a total of 80 epochs (40 in each stage). In comparison, LIFT+

TABLE 2: Comparison with state-of-the-arts methods on Places-LT. † denotes methods using external data. Empty cells “-” indicate unreported results in their original papers.

Methods	Backbone	Learned Params.	Epochs	All Head	Med.	Tail
Training from scratch (with ImageNet pre-trained weights)						
OLTR [3]	RN-152	58.14M	30	35.9	44.7	37.0 25.3
cRT [57]	RN-152	58.14M	90+10	36.7	42.0	37.6 24.9
LWS [57]	RN-152	58.14M	90+10	37.6	40.6	39.1 28.6
MiSLAS [58]	RN-152	58.14M	90+10	40.4	39.6	43.3 36.1
ResLT [54]	RN-152	58.14M	30	39.8	39.8	43.6 31.4
DisAlign [36]	RN-152	58.14M	30	39.3	40.4	42.4 30.1
ALA [77]	RN-152	58.14M	30	40.1	43.9	40.1 32.9
PaCo [22]	RN-152	58.14M	30	41.2	36.1	47.9 35.3
LiVT [96]	ViT-B/16	85.80M	100	40.8	48.1	40.6 27.5
Fine-tuning foundation model (from CLIP)						
BALLAD [44]	ViT-B/16	149.62M	50+10	49.5	49.3	50.2 48.4
VL-LTR [45]†	ViT-B/16	149.62M	100	50.1	54.2	48.5 42.0
Decoder [46]	ViT-B/16	21.26M	~34	46.8	-	- -
LIFT+ (Ours)	ViT-B/16	0.18M	5	51.5	50.8	52.0 51.6
w/ TTE (Ours)	ViT-B/16	0.18M	5	52.1	51.5	52.8 51.7
Fine-tuning foundation model (from ImageNet-21K)						
RAC [47]†	ViT-B/16	85.80M	30	47.2	48.7	48.3 41.8
LPT [48]	ViT-B/16	1.01M	40+40	50.1	49.3	52.3 46.9
LIFT+ (Ours)	ViT-B/16	0.18M	5	48.1	48.4	48.7 46.3
w/ TTE (Ours)	ViT-B/16	0.18M	5	49.0	49.2	49.7 47.3

TABLE 3: Comparison with state-of-the-arts methods on iNaturalist 2018. † denotes using external data. Empty cells “-” indicate unreported results in their original papers.

Methods	Backbone	Learned Params.	Epochs	All Head	Med.	Tail
Training from scratch						
cRT [57]	RN-50	23.51M	90+10	65.2	69.0	66.0 63.2
LWS [57]	RN-50	23.51M	90+10	65.9	65.0	66.3 65.5
MiSLAS [58]	RN-50	23.51M	200+30	71.6	73.2	72.4 70.4
DiVE [13]	RN-50	23.51M	90	69.1	70.6	70.0 67.6
DisAlign [36]	RN-50	23.51M	90	69.5	61.6	70.8 69.9
ALA [77]	RN-50	23.51M	90	70.7	71.3	70.8 70.4
RIDE [53]	RN-50	23.51M	100	72.6	70.9	72.4 73.1
+CR [27]	RN-50	23.51M	200	73.5	71.0	73.8 74.3
+OTmix [17]	RN-50	23.51M	210	73.0	71.3	72.8 73.8
BCL [23]	RN-50	23.51M	100	71.8	-	- -
PaCo [22]	RN-50	23.51M	400	73.2	70.4	72.8 73.6
NCL [95]	RN-50	23.51M	400	74.2	72.0	74.9 73.8
GML [97]	RN-50	23.51M	400	74.5	-	- -
FeatRecon [31]	RN-50	23.51M	100	72.9	-	- -
LOS [61]	RN-50	23.51M	200	70.8	69.2	70.7 71.3
LiVT [96]	ViT-B/16	85.80M	100	76.1	78.9	76.5 74.8
Fine-tuning foundation model (from CLIP)						
VL-LTR [45]†	ViT-B/16	149.62M	100	76.8	-	- -
Decoder [46]	ViT-B/16	21.26M	~5	59.2	-	- -
LIFT+ (Ours)	ViT-B/16	4.75M	15	79.2	73.0	79.1 80.8
w/ TTE (Ours)	ViT-B/16	4.75M	15	80.5	74.4	80.4 82.3
Fine-tuning foundation model (from ImageNet-21K)						
RAC [47]†	ViT-B/16	85.80M	20	80.2	75.9	80.5 81.1
LPT [48]	ViT-B/16	1.01M	80+80	76.1	-	- 79.3
LIFT+ (Ours)	ViT-B/16	4.75M	15	81.1	74.5	81.5 82.4
w/ TTE (Ours)	ViT-B/16	4.75M	15	82.1	75.4	82.3 83.6

requires only 5 training epochs. Moreover, the number of learned parameters is substantially smaller (0.18M).

TABLE 4: Comparison with state-of-the-art methods on CIFAR-100-LT with various imbalance ratios. Empty cells “-” indicate unreported results in their original papers.

Methods	Backbone	Learned Params.	Epochs	Imbalance Ratio		
				100	50	10
Training from scratch						
LDAM [32]	RN-32	0.46M	200	42.0	46.6	58.7
BBN [10]	RN-32	0.46M	200	42.6	47.0	59.1
DiVE [13]	RN-32	0.46M	200	45.4	51.1	62.0
MiSLAS [58]	RN-32	0.46M	200+10	47.0	52.3	63.2
ResLT [54]	RN-32	0.46M	200	45.3	50.0	60.8
BS [33]	RN-32	0.46M	400	50.8	54.2	63.0
PaCo [22]	RN-32	0.46M	400	52.0	56.0	64.2
BCL [23]	RN-32	0.46M	200	51.9	56.6	64.9
FeatRecon [31]	RN-32	0.46M	200	52.5	57.0	65.3
Fine-tuning foundation model (from CLIP)						
LiVT [96]	ViT-B/16	85.80M	100	58.2	-	69.2
BALLAD [44]	ViT-B/16	149.62M	50+10	77.8	-	-
LIFT+ (Ours)	ViT-B/16	0.10M	5	81.7	83.1	84.7
w/ TTE (Ours)	ViT-B/16	0.10M	5	83.4	84.3	85.9
Fine-tuning foundation model (from ImageNet-21K)						
LPT [48]	ViT-B/16	1.01M	40+40	89.1	90.0	91.0
LIFT+ (Ours)	ViT-B/16	0.10M	5	89.3	90.1	91.6
w/ TTE (Ours)	ViT-B/16	0.10M	5	89.7	90.7	92.0

Results on iNaturalist 2018. We present the results in Table 3. In general, LIFT+ achieves state-of-the-art performance on this challenging dataset across different foundation models. It consistently outperforms a series of existing approaches, including VL-LTR, LPT, and RAC. Although Decoder [46] employs fewer training epochs, its performance trails far behind LIFT+. Compared to LPT, LIFT+ surpasses by 5.0% in accuracy and reduces the training cost from 160 epochs (80 per stage) to only 15 epochs. Although LPT uses fewer learned parameters, we can adjust the parameters of LIFT+ to reach a lower quantity (*e.g.*, reduce the bottleneck dimensionality r to 32, requiring only 0.62M parameters). In this case, LIFT+ achieves an accuracy of 80.4% (without TTE) / 81.5% (with TTE), which outperforms LPT by at least 4.3%. In fact, due to the large number of classes in iNaturalist 2018, the classifier itself already comprises 6.25M parameters. Therefore, the parameters introduced by LIFT+ do not impose too much overhead.

Results on CIFAR-100-LT. The results in Table 4 demonstrate that LIFT+ consistently outperforms LiVT, BALLAD, and various training-from-scratch approaches across all imbalance ratios. In particular, our method achieves a 3.9% accuracy improvement over BALLAD, while requiring only 5 training epochs. When adapting to foundation models pre-trained from ImageNet-21K, we employ class mean features to initialize the classifier due to the lack of a corresponding text encoder. It should be noted that the inherent class overlaps between CIFAR-100-LT and ImageNet-21K² can elevate the baseline performance to an exceptionally high level. Despite this, our method maintains its superior performance, surpassing LPT with fewer training epochs and learned parameters.

2. Classes in CIFAR-100-LT: <https://www.cs.toronto.edu/~7Ekriz/cifar.html>; Classes in ImageNet-21K: https://storage.googleapis.com/bit_models/imagenet21k_wordnet_lemmas.txt

TABLE 5: Ablation study of each component in LIFT+. The baseline involves learning a cosine classifier using LA loss.

SLF	SAI	MDA	TTE	ImageNet-LT				Places-LT			
				All	Head	Med.	Tail	All	Head	Med.	Tail
				69.3	77.1	71.0	41.2	43.9	49.4	47.7	25.0
✓				74.9	79.6	75.0	61.6	48.7	50.4	50.3	42.0
✓	✓			76.9	79.8	76.1	71.4	51.0	50.6	51.6	50.5
✓	✓	✓		77.0	79.9	76.1	71.7	51.5	50.8	52.0	51.6
✓	✓	✓	✓	78.3	81.0	77.5	73.4	52.1	51.5	52.8	51.7

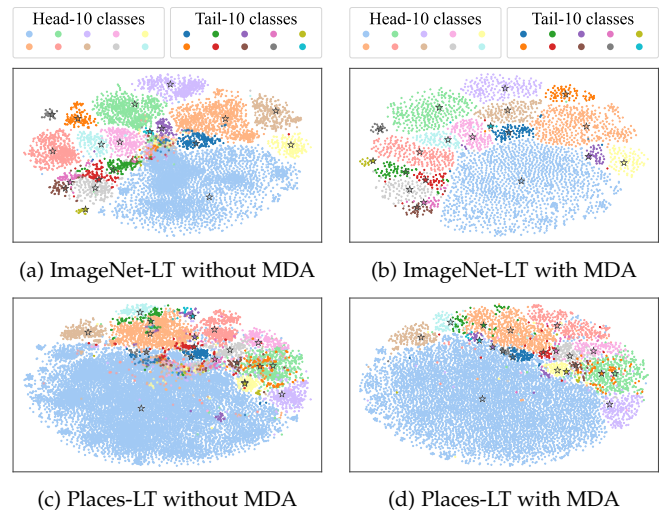


Fig. 10: t-SNE visualization of class-wise features under different augmentation methods. Stars indicate the class centers computed from test features.

5.4 Component Analysis and Ablation Study

Effect of Each Component. To assess the effectiveness of each component, we conduct a systematic ablation study on key components in LIFT+ including i) structured lightweight fine-tuning (SLF), ii) semantic-aware initialization (SAI), iii) minimalist data augmentation (MDA), and iv) test-time ensembling (TTE). The results are presented in Table 5. Specifically, i) SLF substantially improves performance across all classes; ii) SAI consistently improves generalization, particularly on tail classes; iii) MDA further boosts model performance. More crucially, MDA significantly reduces training durations, as will be discussed in Section 5.4; iv) TTE provides additional performance gains for both head and tail classes.

Impact of Minimalist Data Augmentation. We further examine the proposed minimalist data augmentation (MDA) by visualizing class-wise features on ImageNet-LT and Places-LT. For each dataset, we select the first 10 head classes and the last 10 tail classes, generate 10 augmented views for each image, and extract features using the pre-trained CLIP-ViT-B/16 model. As shown in Figure 10, the pre-trained model already provides a preliminary feature representation. However, conventional augmentation without MDA introduces noticeable feature confusion, with some tail-class samples shifted into head-class regions. Such noisy transformations may distort optimization signals and hinder convergence, especially for tail classes with limited data. In contrast, MDA produces more separable fea-

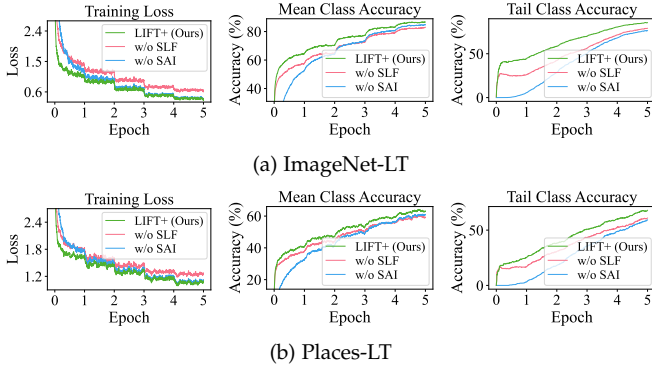


Fig. 11: Convergence curves of training loss and accuracy.

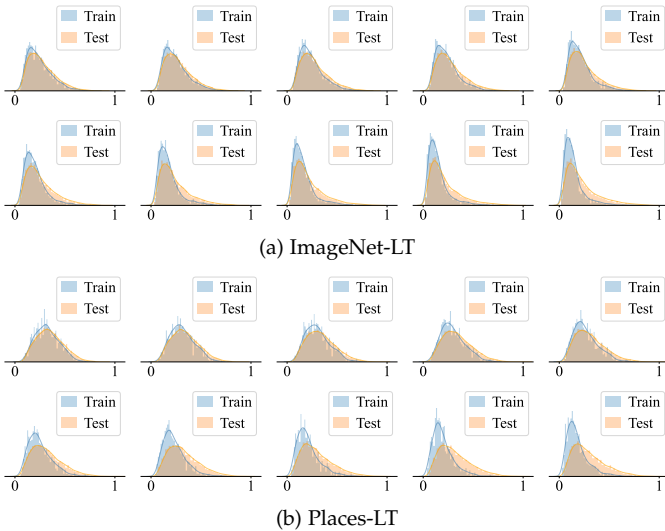


Fig. 12: Intra-class distributions of tail classes under different fine-tuning proportions (from 0.1% to 100%).

tures across classes, indicating that it better preserves class boundaries in the pre-trained feature space.

Impact on Convergence Efficacy. In Figure 11, we illustrate the convergence curve of training loss, mean class accuracy, and tail class accuracy. The results show that, without structured lightweight fine-tuning (SLF), the training loss and accuracy converge suboptimally in all cases. Without semantic-aware initialization (SAI), the training loss is slightly affected, while the class accuracy decreases by a large margin, especially on tail classes. These results underscore the significance of SLF and SAI in ensuring optimal convergence.

Boundary conditions of fine-tuning distortion. To further investigate the boundary conditions under which heavy fine-tuning induces distribution distortion, we visualize the intra-class distribution of tail classes across different fine-tuning proportions in Figure 12. For each dataset, the top row corresponds to proportions of 0.1%, 0.2%, 0.5%, 1%, and 2%, and the bottom row corresponds to 5%, 10%, 20%, 50%, and 100%. The results reveal that as the fine-tuning proportion increases, the intra-class distributions of tail classes exhibit a gradual shift between training and test sets. The boundary conditions vary across datasets but fall within the range of approximately 2% to 10%.




Fig. 13: Performance of LIFT+ with different training epochs.




Fig. 14: Weight norms of the linear classifier on three long-tail datasets. Classes are sorted by their frequency in the training dataset. On iNaturalist 2018, the weight norms are much more skewed, leading to a suboptimal performance.

Notably, the prediction accuracy reported in Figure 5 also begins to decline within the same proportional range. These observations further underscore the connection between fine-tuning proportion and intra-class distribution shift in affecting model performance.

Comparison of Different Training Epochs. To investigate the training efficiency of LIFT+, we run LIFT and LIFT+ with different epochs and report the results in Figure 13. The learning rate for each experiment is modified to $lr = (\text{base_epochs} \cdot \text{base_lr}) / \text{epochs}$. On ImageNet-LT and Places-LT, while LIFT maintains high accuracy under prolonged training durations, it struggles to achieve competitive performance within a few training epochs. In contrast, LIFT+ demonstrates superior efficiency, attaining an accuracy comparable to LIFT with only 5 epochs. On iNaturalist 2018, when training for 5 epochs, LIFT+ achieves an overall accuracy of 74.9% (without TTE) / 76.2% (with TTE), surpassing Decoder [46] by more than 15%. We default to train LIFT+ for 15 epochs on this dataset; however, further performance gains remain attainable. As illustrated in Figure 13, when training for more epochs (*e.g.*, 25 epochs), LIFT+ achieves an additional performance improvement of 1%. However, this will increase the computational overhead, so we abort this approach in LIFT+.

Comparison of Different Losses. Building upon our theoretical analysis, we propose to optimize the unbiased logit-adjusted loss (LA). To further validate our proposal, we compare the performance of LIFT+ with different losses, including cross-entropy loss (CE), focal loss [98], label-distribution-aware margin loss (LDAM) [32], class-balanced loss (CB) [99], generalized re-weighting loss (GRW) [36], and label distribution disentangling loss (LADE) [35]. The results are shown in Table 6, where the LA loss achieves the highest average overall accuracy across these three datasets. In contrast, other losses, such as LDAM, CB, and GRW, cannot achieve satisfactory performance in all cases. The LADE performs well on iNaturalist 2018 but is suboptimal on the other two datasets. Moreover, its form is far more complex compared with LA.

Comparison of Different Classifiers. We propose to use an enhanced cosine classifier in LIFT+. Furthermore, we investigate the linear classifier $z_k = \mathbf{w}_k^\top \mathbf{f} + b$, the L2-

TABLE 6: Performance of LIFT+ with different losses.

Losses	ImageNet-LT				Places-LT				iNaturalist 2018			
	All	Head	Med.	Tail	All	Head	Med.	Tail	All	Head	Med.	Tail
CE	72.1	86.0	69.0	43.3	41.8	56.3	37.6	24.8	74.7	82.8	75.8	71.2
Focal	72.2	85.3	69.3	45.3	42.3	55.8	38.2	26.7	72.4	80.9	73.7	68.4
LDAM	70.0	86.1	67.1	35.0	40.5	56.7	36.1	20.8	76.3	84.2	77.3	72.9
CB	76.8	82.2	76.3	63.5	50.0	52.0	51.3	43.0	78.6	73.5	79.1	79.3
GRW	76.9	82.2	76.3	63.6	50.0	51.9	51.5	43.0	78.7	73.7	79.2	79.4
LADE	78.0	80.6	77.1	74.0	51.9	51.7	52.1	51.5	80.9	75.5	80.7	82.4
LA	78.3	81.0	77.5	73.4	52.1	51.5	52.7	51.7	80.5	74.4	80.4	82.3

TABLE 7: Performance of LIFT+ with different classifiers.

Classifiers	ImageNet-LT				Places-LT				iNaturalist 2018				
	All	Head	Med.	Tail	All	Head	Med.	Tail	All	Head	Med.	Tail	
Linear	78.2	81.1	77.3	73.1	52.1	51.7	52.4	51.9	76.7	75.7	77.7	75.8	
L2-normalized	78.3	81.0	77.2	74.6	52.1	51.4	52.5	52.5	80.1	74.9	79.9	81.7	
Cosine	$\sigma = 15$	75.7	80.4	76.5	59.5	49.6	52.4	53.1	36.7	76.8	73.2	77.0	77.3
	$\sigma = 20$	77.7	80.7	77.4	70.2	51.6	51.7	53.3	47.7	79.9	74.7	79.6	81.7
	$\sigma = 25$	78.3	81.0	77.5	73.4	52.1	51.5	52.7	51.7	80.5	74.4	80.4	82.3
	$\sigma = 30$	78.4	81.3	77.4	73.8	51.9	51.2	52.3	52.2	80.4	75.1	80.4	81.7
	$\sigma = 35$	78.3	81.6	77.0	73.6	51.5	51.0	51.9	51.7	79.8	74.9	79.8	80.9

TABLE 8: Comparison of LIFT+ with different classifier initialization methods.

Methods	ImageNet-LT				Places-LT			
	All	Head	Med.	Tail	All	Head	Med.	Tail
Random	76.5	80.9	76.4	64.4	50.2	51.3	51.8	44.3
Linear probing	77.1	81.8	76.4	66.3	49.9	51.5	50.8	44.8
Class mean features	77.5	81.3	76.8	69.4	51.2	51.3	52.2	48.6
Semantic-aware	78.3	81.0	77.5	73.4	52.1	51.5	52.8	51.7

normalized classifier $z_j = \frac{w_j^\top}{\|w_j\|_2} \mathbf{f}$, as well as the cosine classifier with $\sigma \in \{15, 20, 25, 30, 35\}$. The results in Table 7 show that the linear classifier performs well on ImageNet-LT and Places-LT, but diminishes on the more challenging iNaturalist 2018 dataset. This can be inferred from the skewed distribution of classifier weight norms illustrated in Figure 14. By removing the impact of weight norms, the L2-normalized classifier and the cosine classifier achieve higher performance, especially on tail classes. Moreover, when adopting the cosine classifier, setting σ to 25 or 30 leads to superior performance.

Comparison of Different Initialization Strategies. In Table 8, we evaluate four types of classifier initialization strategies. The linear probing strategy yields limited improvements due to the inherent challenge posed by long-tail data. Using class mean features to initialize the classifier achieves a notable improvement, particularly on tail classes. In LIFT+, we adopt textual features because they transfer semantic relations between classes during fine-tuning. The results show that our proposed strategy is significantly superior to other methods. This experiment also demonstrates that a good starting point for parameter optimization can lead to a better solution.

Comparison of Different Scheduling Functions. In the proposed MDA module, the cropped scale range starts from (λ_0, λ_1) and gradually converges to the maximum value λ_1 by adding a variable $\Delta\lambda_0$. We compare five different scheduling functions for $\Delta\lambda_0$, including minimal, convex, linear, concave, and maximal functions. The definition and

TABLE 9: Definitions of different scheduling functions.

No. of func.	Definitions
① (Minimal)	$g(t) = 0$
② (Convex)	$g(t) = 1 - \sqrt{1 - t^2}$
③ (Linear)	$g(t) = t$
④ (Concave)	$g(t) = \sqrt{1 - (1 - t)^2}$
⑤ (Maximal)	$g(t) = 1$

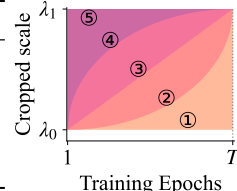


TABLE 10: Comparison of LIFT+ with different scheduling functions for MDA.

No.	ImageNet-LT				Places-LT				iNaturalist 2018			
	All	Head	Med.	Tail	All	Head	Med.	Tail	All	Head	Med.	Tail
①	78.1	80.8	77.2	73.1	51.9	51.5	52.7	50.7	80.1	73.6	79.9	82.2
②	78.3	81.0	77.5	73.4	52.1	51.5	52.8	51.7	80.5	74.4	80.4	82.3
③	78.3	81.2	77.3	73.5	52.1	51.7	52.7	51.7	80.0	74.7	79.8	81.7
④	78.3	81.3	77.4	73.5	52.0	51.3	52.7	51.6	79.9	75.1	79.7	81.4
⑤	78.3	81.4	77.3	73.1	51.8	51.5	52.2	51.4	79.8	75.8	79.7	80.9




Fig. 15: Performance with different expanded size e . When applying TTE, setting e to a multiple of the patch size 16 yields suboptimal improvements.

curves are shown in Table 9, and the comparison results are reported in Table 10. On ImageNet-LT and Places-LT, the performance difference is not significant. On iNaturalist 2018, the convex function exhibits an obvious enhancement.

Comparison of Different Expanded Sizes. To analyze the effect of TTE, we compare different expanded sizes e and report the results in Figure 15. Remarkably, TTE consistently improves performance with different expanded sizes. Note that $e = 0$ still achieves improvements on ImageNet-LT and iNaturalist 2018, because the aspect ratios of the images in

TABLE 11: Comparison of LIFT+ with different random seeds.

Random seeds	ImageNet-LT				Places-LT			
	All	Head	Med.	Tail	All	Head	Med.	Tail
0	78.3	81.0	77.5	73.4	52.1	51.5	52.8	51.7
1	78.4	81.3	77.4	73.7	52.1	51.4	52.8	51.6
2	78.2	81.0	77.3	73.8	52.1	51.3	52.9	51.7
3	78.3	81.3	77.4	73.0	52.0	51.4	52.9	51.2
4	78.4	81.2	77.5	73.7	52.0	51.3	52.8	51.4
Average	78.3	81.2	77.4	73.5	52.1	51.4	52.8	51.5

TABLE 12: Comparison of different fine-tuning methods. All baselines use SAI, MDA, and TTE for a fair comparison.

Methods	ImageNet-LT				Places-LT				
	All	Head	Med.	Tail	All	Head	Med.	Tail	
Zero-shot CLIP	67.9	69.0	67.3	67.2	39.9	37.5	39.1	45.9	
Classifier fine-tuning	74.6	77.9	73.8	67.7	49.6	49.3	50.7	47.8	
Full fine-tuning	73.2	81.3	72.8	51.9	47.5	51.9	49.4	35.0	
LIFT+ (Arbitrary)	77.7	80.9	76.8	72.5	51.8	51.7	52.4	50.8	
LIFT+ w/	BitFit	77.0	79.8	76.3	71.4	51.6	50.8	52.6	50.7
	VPT	77.6	80.1	76.8	72.9	51.9	51.4	52.7	51.2
	Adapter	78.0	81.0	77.1	72.6	51.9	51.4	52.7	51.1
	LoRA	77.9	80.4	76.9	73.9	51.6	51.4	52.1	50.9
	AdaptFormer	78.3	81.0	77.5	73.4	52.1	51.5	52.8	51.7

these two datasets are not always 1:1, which enables TTE to generate partially diverse croppings. More crucially, setting the expanded size as a multiple of the patch size 16 results in suboptimal enhancement, as the generated images will share many repeated patches, thus reducing patch diversity. Generally, setting $e = 24$ achieves significant improvements on different datasets.

Comparison of Different Random Seeds. To evaluate the robustness of LIFT+ across different runs, we conduct experiments with different fixed random seeds and report the results in Table 11. The results show that LIFT+ achieves consistent performance under different random seeds, with the variation in overall performance remaining within 0.1. Notably, the average results across these seeds are also competitive. These findings indicate the robustness of LIFT+ under different operational conditions, as the choice of random seed does not introduce significant variability. To ensure the reproducibility, we set the random seed to 0 for all experiments.

5.5 More Advantages of LIFT+

Adaptability to Varying Fine-Tuning Methods. LIFT+ is a general framework in which many lightweight fine-tuning methods can be integrated. In addition to zero-shot CLIP, classifier fine-tuning, and full fine-tuning, we test LIFT+ with arbitrary lightweight fine-tuning and five structured lightweight methods, BitFit [79], VPT [80], Adapter [81], LoRA [82], and AdaptFormer [83]. The results in Table 12 demonstrate that arbitrary lightweight fine-tuning surpasses the baseline methods by a large margin across both overall and tail-class performance. This underscores the efficacy of lightweight fine-tuning even in the absence of specific strategies. Furthermore, the integration of structured methods leads to improved performance. Specifically, LIFT+ with AdaptFormer performs best on both datasets.

TABLE 13: Results on ImageNet-LT with ViT-L/14.

Methods	Backbone	Learned Params.	Epochs	All	Head	Med.	Tail
CLIP	ViT-L/14	-	-	73.6	74.6	73.1	72.3
Decoder [46]	ViT-L/14	39.79M	~18	79.3	-	-	-
LIFT+ (Ours)	ViT-L/14	0.86M	5	82.5	84.7	81.7	78.8
w/ TTE (Ours)	ViT-L/14	0.86M	5	83.1	85.4	82.3	79.4

TABLE 14: Results on Places-LT with ViT-L/14.

Methods	Backbone	Learned Params.	Epochs	All	Head	Med.	Tail
CLIP	ViT-L/14	-	-	39.9	38.1	39.2	45.1
Decoder [46]	ViT-L/14	39.79M	~34	48.4	-	-	-
LIFT+ (Ours)	ViT-L/14	0.27M	5	53.3	52.3	53.9	53.6
w/ TTE (Ours)	ViT-L/14	0.27M	5	53.7	52.9	54.3	54.0

TABLE 15: Results on iNaturalist 2018 with ViT-L/14.

Methods	Backbone	Learned Params.	Epochs	All	Head	Med.	Tail
CLIP	ViT-L/14	-	-	5.8	11.1	5.4	4.9
Decoder [46]	ViT-L/14	39.79M	~5	72.3	65.5	73.2	73.0
LIFT+ (Ours)	ViT-L/14	6.37M	15	84.0	79.1	83.9	85.4
w/ TTE (Ours)	ViT-L/14	6.37M	15	84.9	79.3	84.7	86.5

TABLE 16: Results on iNaturalist 2018 with ViT-L/14 (336×336 pixels resolution).

Methods	Backbone	Learned Params.	Epochs	All	Head	Med.	Tail
CLIP	ViT-L/14	-	-	6.2	11.5	5.8	5.3
LIFT+ (Ours)	(336px)	6.37M	15	86.8	82.4	86.9	87.7
w/ TTE (Ours)	(336px)	6.37M	15	87.3	83.4	87.2	88.3

TABLE 17: Results on ImageNet-LT with ResNet-50.

Methods	Backbone	Learned Params.	Epochs	All	Head	Med.	Tail
CLIP	RN-50	-	-	57.6	58.7	56.9	56.8
LIFT+ w/ SSF	RN-50	0.002M	10	65.8	70.9	65.0	53.9
LIFT+ w/ BitFit	RN-50	0.026M	10	65.9	70.1	65.1	57.2

TABLE 18: Results on Places-LT with ResNet-50.

Methods	Backbone	Learned Params.	Epochs	All	Head	Med.	Tail
CLIP	RN-50	-	-	35.2	33.1	34.6	40.5
LIFT+ w/ SSF	RN-50	0.002M	5	45.8	46.6	47.4	40.9
LIFT+ w/ BitFit	RN-50	0.026M	10	46.6	47.3	47.9	42.2

Extension to Different Backbones. LIFT+ can be extended to various backbones. In addition to ViT-B/16, we also assess LIFT+ based on the larger ViT-L/14. The results in Tables 13 to 15 show that LIFT+ surpasses the state-of-the-art method Decoder [46] by a considerable margin, showcasing improvements of 3.2% on ImageNet-LT, 4.9% on Places-LT, and 11.7% on iNaturalist 2018. Furthermore, the model with higher resolution (336×336 pixels) yields better performance. These results underscore the adaptability of LIFT+ across variant backbones.

The adaptation of LIFT+ to the widely used ResNet [87] is also worth exploring. However, due to the absence

TABLE 19: Results on ImageNet-LT with DINO-ViT.

Methods	Backbone	Learned Params.	Epochs	All Head	Med. Tail
LIFT+ (Ours)	ViT-B/16	0.62M	5	73.6	77.0 72.6 67.1
w/ TTE (Ours)	ViT-B/16	0.62M	5	74.4	77.8 73.5 67.9

TABLE 20: Results on Places-LT with DINO-ViT.

Methods	Backbone	Learned Params.	Epochs	All Head	Med. Tail
LIFT+ (Ours)	ViT-B/16	0.18M	5	44.4	45.1 45.6 40.4
w/ TTE (Ours)	ViT-B/16	0.18M	5	45.1	46.1 46.1 40.7

TABLE 21: Results on iNaturalist 2018 with DINO-ViT.

Methods	Backbone	Learned Params.	Epochs	All Head	Med. Tail
LIFT+ (Ours)	ViT-B/16	4.75M	15	71.6	64.3 72.4 72.5
w/ TTE (Ours)	ViT-B/16	4.75M	15	72.7	65.2 73.5 73.7

of a dedicated lightweight fine-tuning method tailored for ResNet, it is challenging to directly integrate LIFT+ with ResNet. Despite this limitation, we investigate some intuitive strategies, including i) incorporating a scaling and shifting (SSF) [100] module after the backbone, and ii) fine-tuning solely the bias terms (BitFit) of ResNet. The results are presented in Tables 17 and 18. LIFT+ delivers remarkable performance enhancements over zero-shot CLIP. Furthermore, when compared with conventional ResNet-based approaches in Tables 1 and 2, LIFT+ achieves superior results with fewer learned parameters and training epochs.

Moreover, one may be concerned with the generalizability of LIFT+ across different foundation models. To access this, we conduct additional experiments by adapting LIFT+ to DINO-ViT. The results are presented in Tables 19 to 21. When using DINO as the foundation model, LIFT+ achieves lower performance than with CLIP, which may be attributed to the distinct pre-training paradigms and representation characteristics of these foundation models. Nonetheless, the results remain competitive against most existing methods, demonstrating the effectiveness and generalizability of LIFT+ across different foundation models.

6 TECHNICAL DETAILS OF LIGHTWEIGHT FINE-TUNING FOR LONG-TAIL LEARNING

Preliminary to the Transformer Architecture. The Transformer model architecture comprises an embedding layer followed by a series of Transformer blocks. Specifically, it first partitions an input image \mathbf{x} into m discrete patches $\{\mathbf{x}_i\}_{i=1}^m$. These patches are then embedded into sequences of d -dimensional tokens $\mathbf{E}^0 = \text{Embedding}([\mathbf{x}_1; \dots; \mathbf{x}_m]) \in \mathbb{R}^{m \times d}$. The embeddings are subsequently passed through L successive Transformer blocks $\{\Phi^l\}_{l=1}^L$ within the model:

$$\mathbf{X}^l = \Phi^l(\mathbf{X}^{l-1}) \begin{cases} \hat{\mathbf{X}}^l = \text{MSA}(\text{LN}(\mathbf{X}^{l-1})) + \mathbf{X}^{l-1} \\ \mathbf{X}^l = \text{FFN}(\text{LN}(\hat{\mathbf{X}}^l)) + \hat{\mathbf{X}}^l \end{cases} \quad (7)$$

In practice, \mathbf{X}^0 is normally set as \mathbf{E}^0 , and will prepend a learnable token $\mathbf{c}^0 \in \mathbb{R}^d$ when performing classification tasks, *i.e.*, $\mathbf{X}^0 = [\mathbf{c}^0; \mathbf{E}^0]$. The feature is extracted

TABLE 22: Components and parameter quantity of ViT.

Layers	Components	Variables	#Params.
Embedding	Projection	-	$(d_0 + 1)d$
	Class Token	\mathbf{c}^0	d
	Positional	-	$(m + 1)d$
Block-1 ($l = 1$)	LN	γ, β	$2d$
	MSA	$\{\mathbf{W}_h^Q, \mathbf{b}_h^Q\}_{h=1}^H$ $\{\mathbf{W}_h^K, \mathbf{b}_h^K\}_{h=1}^H$ $\{\mathbf{W}_h^V, \mathbf{b}_h^V\}_{h=1}^H$ $\mathbf{W}^O, \mathbf{b}^O$	$4d^2 + 4d$
	LN	γ, β	$2d$
	FFN	$\mathbf{W}_1, \mathbf{b}_1$ $\mathbf{W}_2, \mathbf{b}_2$	$8d^2 + 5d$
\vdots	\vdots	\vdots	\vdots
Block- L	\dots	\dots	$12d^2 + 13d$
Normalization	LN	γ, β	$2d$

from the same position of the last-layer sequence, which is $\phi(\mathbf{x}) = \text{LN}(\mathbf{c}^L)$. Moreover, MSA represents multi-head self-attention, FFN indicates feed-forward network, and LN denotes layer normalization [101]. The definitions are

$$\text{MSA}(\mathbf{X}) = \text{Concat}_{h=1}^H \left(\text{Softmax} \left(\frac{\mathbf{Q}_h \mathbf{K}_h^\top}{\sqrt{d/H}} \right) \mathbf{V}_h \right) \mathbf{W}^O$$

where $(\mathbf{Q}_h, \mathbf{K}_h, \mathbf{V}_h) = (\mathbf{X} \mathbf{W}_h^Q, \mathbf{X} \mathbf{W}_h^K, \mathbf{X} \mathbf{W}_h^V)$ (8)

$$\text{FFN}(\mathbf{X}) = \text{ReLU}(\mathbf{X} \mathbf{W}_1) \mathbf{W}_2 \quad (9)$$

$$\text{LN}(\mathbf{X}_i) = \frac{\mathbf{X}_i - \mathbb{E}[\mathbf{X}_i]}{\sqrt{\text{Var}[\mathbf{X}_i]}} \cdot \gamma + \beta \quad (10)$$

where H is the number of heads, $\mathbf{W}_h^Q, \mathbf{W}_h^K, \mathbf{W}_h^V \in \mathbb{R}^{d \times \frac{d}{H}}$, $\mathbf{W}^O \in \mathbb{R}^{d \times d}$, $\mathbf{W}_1 \in \mathbb{R}^{d \times 4d}$, $\mathbf{W}_2 \in \mathbb{R}^{4d \times d}$, $\gamma \in \mathbb{R}^d$ and $\beta \in \mathbb{R}^d$ are model parameters. Note that the bias terms are omitted for simplification.

The architectural components and parameter quantities of a Vision Transformer (ViT) are summarized in Table 22. The total number of parameters in ViT can be calculated as $12Ld^2 + (13L + m + d_0 + 5)d$, where L represents the number of blocks, m corresponds to the number of image patches, d_0 denotes the dimensionality of each image patch, and d indicates the dimensionality of the embedded features. For example, in ViT-B/16 with 224×224 resolution, where $L = 12$, $m = \frac{224}{16} \times \frac{224}{16} = 196$, $d_0 = 16 \times 16 \times 3 = 768$, and $d = 768$, the parameter quantity amounts to approximately 85.80M. Note that models from different sources (*e.g.*, CLIP or ImageNet-21K pre-trained ViT) may have slight differences due to implementation details, while their results are always within the standard configuration [42], [43].

LIFT+ with Structured Lightweight Fine-Tuning Methods. In this paper, we explore lightweight fine-tuning by learning only a small set of parameters for adaptation. Concretely, we propose a versatile and inclusive framework LIFT+, allowing the incorporation of a range of lightweight

TABLE 23: Parameter quantities for structured lightweight fine-tuning modules in a Transformer block.

Modules	Components	Variables	#Params.
	LN-bias	β	d
BitFit	MSA-bias	$\{b_h^Q\}_{h=1}^H$	$4d$
		$\{b_h^K\}_{h=1}^H$	
		$\{b_h^V\}_{h=1}^H$	
		b^O	
	LN-bias	β	d
	FFN-bias	b_1^l b_2^l	$5d$
VPT	Prompts	P^l	pd
Adapter	LN	γ, β	$2d$
	Projection	$W_{\text{down}}, b_{\text{down}}$ $W_{\text{up}}, b_{\text{up}}$	$(2r + 1)d + r$
LoRA	Projection	W_{down} (for W_Q)	$4rd$
		W_{up} (for W_Q)	
		W_{down} (for W_V)	
		W_{up} (for W_V)	
	LN	γ, β	$2d$
AdaptFormer	Projection	$W_{\text{down}}, b_{\text{down}}$ $W_{\text{up}}, b_{\text{up}}$	$(2r + 1)d + r$
	Scaling	s	1

modules, including but not limited to:

- *Bias-terms Fine-tuning (BitFit)* [79], which aims to fine-tune only the bias terms of the model. Formally, given a projection $\mathbf{X}\mathbf{W} + \mathbf{b}$, it freezes \mathbf{W} and optimizes \mathbf{b} .
- *Visual Prompt Tuning (VPT)* [80], which prepends learnable prompts $P^l \in \mathbb{R}^{p \times d}$ at the entrance of each layer to extend $\mathbf{X}^l = [c^l; E^l]$ to $[c^l; P^l; E^l]$.
- *Adapter* [81], which attaches a bottleneck module to FFN layer to reconstruct $\text{FFN}(\cdot)$ to $\text{Adapter}(\text{FFN}(\cdot))$. The definition is $\text{Adapter}(\mathbf{X}) = \text{ReLU}(\text{LN}(\mathbf{X})\mathbf{W}_{\text{down}})\mathbf{W}_{\text{up}}$, where $\mathbf{W}_{\text{down}} \in \mathbb{R}^{d \times r}$ and $\mathbf{W}_{\text{up}} \in \mathbb{R}^{r \times d}$ ($r \ll d$).
- *Low-Rank Adapter (LoRA)* [82], which optimizes low-rank matrices \mathbf{W}_{down} and \mathbf{W}_{up} to reparameterize a matrix \mathbf{W} to $\mathbf{W} + \mathbf{W}_{\text{down}}\mathbf{W}_{\text{up}}$. In practice, it is often applied to the weights in the MSA module, such as W_Q and W_V .
- *AdaptFormer* [83], which extends the sequential Adapter to a parallel one. Formally, it computes $s \cdot \text{Adapter}(\hat{\mathbf{X}}^l)$ and adds it to \mathbf{X}^l in Equation (7), where s is a scaling parameter that can be either specified or learnable³.

We summarize the parameter quantities of the above lightweight fine-tuning modules in Table 23. For all modules, the quantities are at the polynomial level of d . Note that the prompt length p for VPT and the bottleneck dimensionality r for Adapter, LoRA, and AdaptFormer are much smaller than d . Compared to a standard Transformer block,

3. The usage of learnable scaling parameter is not proposed in the main paper but is implemented in their source code: <https://github.com/ShoufaChen/AdaptFormer/blob/main/models/adaptformer.py>

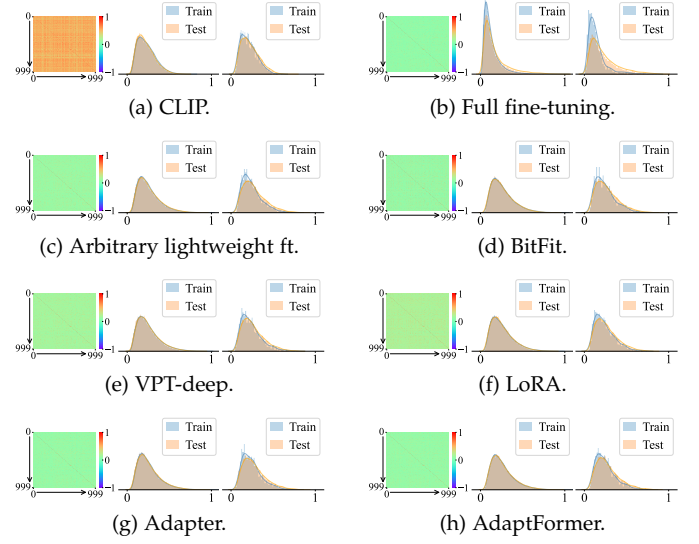


Fig. 16: Visualization of the inter-class feature similarities (the heatmaps) and intra-class distance distributions from head classes (the left histograms) and tail classes (the right histograms) on ImageNet-LT.

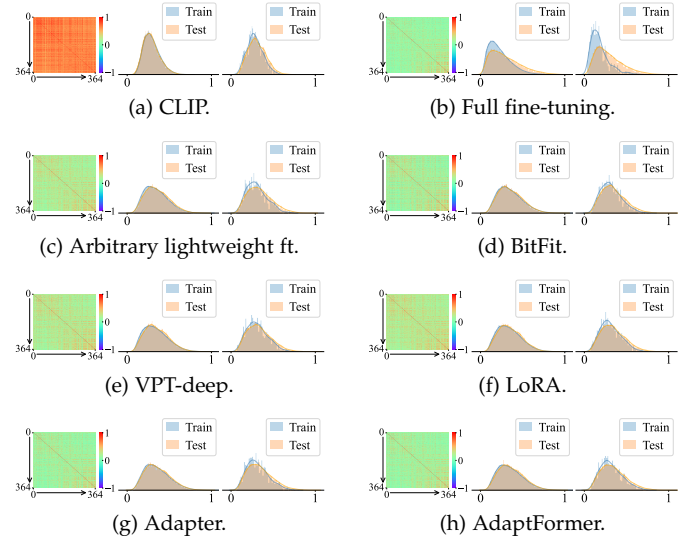


Fig. 17: Visualization of the inter-class feature similarities (the heatmaps) and intra-class distance distributions from head classes (the left histograms) and tail classes (the right histograms) on Places-LT.

these lightweight fine-tuning modules are significantly low-complexity ($\mathcal{O}(d)$ vs. $\mathcal{O}(d^2)$).

In addition to the lightweight fine-tuning module, a task-specific classifier is required, which contains approximately Kd parameters (omitting bias terms), where K represents the number of classes. In LIFT+, we set the bottleneck dimensionality $r = 2^{\lfloor \log_2(\frac{K}{2L}) \rfloor} \leq \frac{K}{2L}$ for the AdaptFormer module, so that the total parameter quantity is $L \cdot 2rd \leq Kd$ (ignoring constant terms). As a result, it learns even fewer parameters than the classifier.

Learned Representations of Different Fine-Tuning




Fig. 18: Comparative analysis of learned parameters by adjusting the bottleneck dimensionality r . In the yellow region, the incorporated module contains fewer learned parameters than the classifier. The blue region is just the opposite.




Fig. 19: The learned scaling parameters of the AdaptFormer modules in different layers. LIFT+ with AdaptFormer adaptively learns a distinct scaling parameter for each layer.

Methods. In Figures 16 and 17, we visualize the inter-class feature separabilities and intra-class distance distributions based on the representation learned by (a) original CLIP, (b) full fine-tuning, (c) arbitrary lightweight fine-tuning, and (d-i) structured lightweight fine-tuning methods. Compared to the original CLIP, all of these lightweight fine-tuning methods yield more discriminative representations, achieving feature separability comparable to full fine-tuning. More importantly, in contrast to full fine-tuning, lightweight methods preserve undistorted intra-class distributions. For both head and tail classes, the features of training and test data exhibit nearly overlapping distributions. This property contributes to their stable performance improvements, particularly on tail classes, as demonstrated in Table 12.

Impact of the Quantity of Learned Parameters. In LIFT+, the number of learned parameters can be flexibly adjusted. We investigate the impact of parameter quantity by adjusting bottleneck dimensionality r and present the results in Figure 18. The results reveal that the performance is robust to the adjustment of dimensionalities. Generally, when the lightweight parameters approximate the classifier parameters, the model achieves adequate improvements without incurring substantial computational overhead.

Advantage of LIFT+ with AdaptFormer. In each layer, the output of the AdaptFormer module can be multiplied by a learnable scaling parameter s before being added to the corresponding block. Therefore, we can compare the values of s to analyze the effects of the module for different layers. The comparison results are presented in Figure 19. It is inspiring that LIFT+ with AdaptFormer adaptively learns suitable scaling parameters for different layers. For example, the values of the last layers tend to be larger, which indicates that the adaptations of the last several layers contribute more significantly to downstream classification tasks.

Lightweight Fine-Tuning vs. Partial Fine-Tuning. Partial fine-tuning [75] offers an intuitive approach to reduce the number of learned parameters. Specifically, it fine-tunes only the last k layers while keeping the remaining frozen. In Figure 20, we compare partial fine-tuning and lightweight fine-tuning (LIFT+) on ImageNet-LT, Places-LT, and iNatu-




Fig. 20: Partially fine-tuning the last k layers. Similar to full fine-tuning, we search the learning rate from $\{0.03, 0.02, 0.01, 0.005, 0.002, 0.001\}$. For LIFT+, the learning rate is fixed to 0.02. The optimal learning rate and fine-tuned layers need to be elaborately selected for best performance. In contrast, LIFT+ consistently demonstrates superior performance.

ralist 2018. Similar to full fine-tuning, partial fine-tuning is also sensitive to the learning rate. For shallow fine-tuned layers (e.g., $k = 0, 1, 2$), higher learning rates produce better performance. When fine-tuning more layers (e.g., $k = 9, 12$), high learning rates adversely lead to severe accuracy degradation. Moreover, identifying the optimal number of fine-tuned layers presents another non-trivial challenge, even if using the ideal learning rate, as the optimal k varies dramatically across datasets: 1 for ImageNet-LT and Places-LT, and 6 for iNaturalist 2018. Remarkably, LIFT+ maintains strong performance across all datasets without requiring hyperparameter adjustments.

7 CONCLUSION

This paper investigates long-tail learning with the foundation model. We first identify the critical limitation of heavy fine-tuning in distorting tail-class performance and discover that lightweight fine-tuning can effectively mitigate this issue. Based on this insight, we propose LIFT+, a versatile and inclusive framework tailored for long-tail learning. The proposed framework facilitates efficient and accurate fine-tuning, achieving state-of-the-art performance in fewer than 15 training epochs without relying on any external data, while consistently outperforming numerous baseline methods across a range of long-tail datasets, including ImageNet-LT, Places-LT, iNaturalist 2018, and CIFAR-100-LT. We emphasize the ease of training and hope that our approach serves as an inspiration for further advancements in the field of long-tail learning.

Limitation: LIFT+ adopts semantic-aware initialization by harnessing semantic knowledge of CLIP. However, it faces challenges when applied to vision-only foundation models. In such cases, we opt for class mean features as a practical alternative. Fortunately, the results on DINO and ImageNet-21K pre-trained ViT indicate that class mean features serve as a viable substitute. As delineated in Table 8, this approach yields notable improvements compared to other initialization methods; however, its efficacy on tail classes remains inferior to that of semantic-aware strategies. It remains an intriguing challenge regarding how to exploit available information for classifier initialization in visual-only foundation models. We identify this as a promising direction for future research.

Broader Impact: This paper aims to advance the field of long-tail learning with foundation models. Our empirical

validations are conducted on common long-tail datasets, including ImageNet-LT, Places-LT, iNaturalist 2018, and CIFAR-100-LT. However, in more critical domains such as object detection and instance segmentation, and high-stakes applications such as autonomous driving, medical image diagnosis, and industrial anomaly detection, the presence of long-tail distributions may lead to severely biased and unreliable predictions. Although fine-tuning foundation models offer strong adaptability, its direct application to such scenarios may result in significant performance degradation if not carefully designed. By highlighting these risks, this paper aims to inspire further research into robust and reliable long-tail learning algorithms in real-world scenarios.

ACKNOWLEDGMENTS

This research was supported by the Jiangsu Science Foundation (BK20243012, BG2024036), the National Natural Science Foundation of China (625B2084, 62576162, 62576174), and the Fundamental Research Funds for the Central Universities (022114380023).

REFERENCES

- [1] Y.-X. Wang, D. Ramanan, and M. Hebert, "Learning to model the tail," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 7029–7039.
- [2] G. Van Horn, O. Mac Aodha, Y. Song, Y. Cui, C. Sun, A. Shepard, H. Adam, P. Perona, and S. Belongie, "The inaturalist species classification and detection dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8769–8778.
- [3] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, and S. X. Yu, "Large-scale long-tailed recognition in an open world," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2537–2546.
- [4] A. Gupta, P. Dollar, and R. Girshick, "LVIS: A dataset for large vocabulary instance segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5356–5364.
- [5] T. Wang, Y. Li, B. Kang, J. Li, J. Liew, S. Tang, S. Hoi, and J. Feng, "The devil is in classification: A simple framework for long-tail instance segmentation," in *Proceedings of the 16th European Conference on Computer Vision*, 2020, pp. 728–744.
- [6] J. Wang, W. Zhang, Y. Zang, Y. Cao, J. Pang, T. Gong, K. Chen, Z. Liu, C. C. Loy, and D. Lin, "Seesaw loss for long-tailed instance segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9695–9704.
- [7] Y. Li, T. Wang, B. Kang, S. Tang, C. Wang, J. Li, and J. Feng, "Overcoming classifier imbalance for long-tail object detection with balanced group softmax," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10991–11000.
- [8] T. Wang, Y. Zhu, C. Zhao, W. Zeng, J. Wang, and M. Tang, "Adaptive class suppression loss for long-tail object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3103–3112.
- [9] B. Li, Y. Yao, J. Tan, G. Zhang, F. Yu, J. Lu, and Y. Luo, "Equalized focal loss for dense long-tailed object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6990–6999.
- [10] B. Zhou, Q. Cui, X.-S. Wei, and Z.-M. Chen, "BBN: Bilateral-branch network with cumulative learning for long-tailed visual recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9719–9728.
- [11] H.-P. Chou, S.-C. Chang, J.-Y. Pan, W. Wei, and D.-C. Juan, "Remix: Rebalanced mixup," in *Proceedings of the 16th European Conference on Computer Vision*, 2020, pp. 95–110.
- [12] Y. Yang and Z. Xu, "Rethinking the value of labels for improving class-imbalanced learning," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 19290–19301.
- [13] Y.-Y. He, J. Wu, and X.-S. Wei, "Distilling virtual examples for long-tailed recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 235–244.
- [14] S. Park, Y. Hong, B. Heo, S. Yun, and J. Y. Choi, "The majority can help the minority: Context-rich minority oversampling for long-tailed classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6887–6896.
- [15] J.-X. Shi, T. Wei, Y. Xiang, and Y.-F. Li, "How re-sampling helps for long-tail learning?" in *Advances in Neural Information Processing Systems*, vol. 36, 2023, pp. 75669–75687.
- [16] S. Ahn, J. Ko, and S.-Y. Yun, "CUDA: Curriculum of data augmentation for long-tailed recognition," in *International Conference on Learning Representations*, 2023.
- [17] J. Gao, H. Zhao, Z. Li, and D. Guo, "Enhancing minority classes by mixing: An adaptive optimal transport approach for long-tailed classification," in *Advances in Neural Information Processing Systems*, vol. 36, 2023, pp. 60329–60348.
- [18] Z. Jian, Y. Chen, Y. Wang, J. Yao, M. Wang, and Q. Wu, "Supervised exploratory learning for long-tailed visual recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 1870–1880.
- [19] Z. Li and Y. Jia, "Conmix: Contrastive mixup at representation level for long-tailed deep clustering," in *The Thirteenth International Conference on Learning Representations*, 2025.
- [20] B. Kang, Y. Li, S. Xie, Z. Yuan, and J. Feng, "Exploring balanced feature spaces for representation learning," in *International Conference on Learning Representations*, 2021.
- [21] P. Wang, K. Han, X.-S. Wei, L. Zhang, and L. Wang, "Contrastive learning based hybrid networks for long-tailed image classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 943–952.
- [22] J. Cui, Z. Zhong, S. Liu, B. Yu, and J. Jia, "Parametric contrastive learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 715–724.
- [23] J. Zhu, Z. Wang, J. Chen, Y.-P. P. Chen, and Y.-G. Jiang, "Balanced contrastive learning for long-tailed visual recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6908–6917.
- [24] H. Liu, J. Z. HaoChen, A. Gaidon, and T. Ma, "Self-supervised learning is more robust to dataset imbalance," in *International Conference on Learning Representations*, 2022.
- [25] Y. Yang, S. Chen, X. Li, L. Xie, Z. Lin, and D. Tao, "Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network?" in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 37991–38002.
- [26] G. Peifeng, Q. Xu, P. Wen, Z. Yang, H. Shao, and Q. Huang, "Feature directions matter: Long-tailed learning via rotated balanced representation," in *Proceedings of the 40th International Conference on Machine Learning*, 2023, pp. 27542–27563.
- [27] Y. Ma, L. Jiao, F. Liu, S. Yang, X. Liu, and L. Li, "Curvature-balanced feature manifold learning for long-tailed classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 15824–15835.
- [28] A. Kukleva, M. Böhle, B. Schiele, H. Kuehne, and C. Rupprecht, "Temperature schedules for self-supervised contrastive methods on long-tail data," in *International Conference on Learning Representations*, 2023.
- [29] J. Gao, H. Zhao, D. D. Guo, and H. Zha, "Distribution alignment optimization through neural collapse for long-tailed classification," in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 14969–14987.
- [30] Y. Liu, L. Yang, and Y. Wang, "Long-tailed classification with multi-granularity semantics," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 4285–4294.
- [31] L. Yi, J. Yao, W. Lyu, H. Ling, R. Douady, and C. Chen, "Geometry of long-tailed representation learning: Rebalancing features for skewed distributions," in *The Thirteenth International Conference on Learning Representations*, 2025.
- [32] K. Cao, C. Wei, A. Gaidon, N. Arechiga, and T. Ma, "Learning imbalanced datasets with label-distribution-aware margin loss," in *Advances in Neural Information Processing Systems*, vol. 32, 2019, pp. 1565–1576.
- [33] J. Ren, C. Yu, S. Sheng, X. Ma, H. Zhao, S. Yi, and H. Li, "Balanced meta-softmax for long-tailed visual recognition," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 4175–4186.

- [34] A. K. Menon, S. Jayasumana, A. S. Rawat, H. Jain, A. Veit, and S. Kumar, "Long-tail learning via logit adjustment," in *International Conference on Learning Representations*, 2021.
- [35] Y. Hong, S. Han, K. Choi, S. Seo, B. Kim, and B. Chang, "Disentangling label distribution for long-tailed visual recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6626–6636.
- [36] S. Zhang, Z. Li, S. Yan, X. He, and J. Sun, "Distribution alignment: A unified framework for long-tail visual recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 2361–2370.
- [37] D. Samuel and G. Chechik, "Distributional robustness loss for long-tail learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9495–9504.
- [38] T. Wei, H. Wang, W.-W. Tu, and Y.-F. Li, "Robust model selection for positive and unlabeled learning with constraints," *Science China Information Sciences*, vol. 65, no. 11, p. 212101, 2022.
- [39] B. Han, "Wrapped cauchy distributed angular softmax for long-tailed visual recognition," in *Proceedings of the 40th International Conference on Machine Learning*, 2023, pp. 12 368–12 388.
- [40] Z. Zhao, H. Wen, P. Wang, S. Wang, Z. Wang, Q. Zhang, and Y. Wang, "Balancing model efficiency and performance: Adaptive pruner for long-tailed data," in *Proceedings of the 42nd International Conference on Machine Learning*, 2025, pp. 77 723–77 739.
- [41] S. Li, Q. Xu, Z. Yang, Z. Wang, L. Zhang, X. Cao, and Q. Huang, "Focal-sam: Focal sharpness-aware minimization for long-tailed classification," in *Proceedings of the 42nd International Conference on Machine Learning*, 2025, pp. 36 624–36 651.
- [42] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *Proceedings of the 38th International Conference on Machine Learning*, 2021, pp. 8748–8763.
- [43] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *International Conference on Learning Representations*, 2021.
- [44] T. Ma, S. Geng, M. Wang, J. Shao, J. Lu, H. Li, P. Gao, and Y. Qiao, "A simple long-tailed recognition baseline via vision-language model," *arXiv preprint arXiv:2111.14745*, 2021.
- [45] C. Tian, W. Wang, X. Zhu, J. Dai, and Y. Qiao, "VL-LTR: Learning class-wise visual-linguistic representation for long-tailed visual recognition," in *Proceedings of the 17th European Conference on Computer Vision*, 2022, pp. 73–91.
- [46] Y. Wang, Z. Yu, J. Wang, Q. Heng, H. Chen, W. Ye, R. Xie, X. Xie, and S. Zhang, "Exploring vision-language models for imbalanced learning," *International Journal of Computer Vision*, vol. 132, pp. 224–237, 2024.
- [47] A. Long, W. Yin, T. Ajanthan, V. Nguyen, P. Purkait, R. Garg, A. Blair, C. Shen, and A. van den Hengel, "Retrieval augmented classification for long-tail visual recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6959–6969.
- [48] B. Dong, P. Zhou, S. Yan, and W. Zuo, "LPT: Long-tailed prompt tuning for image classification," in *International Conference on Learning Representations*, 2023.
- [49] J.-X. Shi, T. Wei, Z. Zhou, J.-J. Shao, X.-Y. Han, and Y.-F. Li, "Long-tail learning with foundation model: Heavy fine-tuning hurts," in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 45 014–45 039.
- [50] L. Yang, H. Jiang, Q. Song, and J. Guo, "A survey on long-tailed visual recognition," *International Journal of Computer Vision*, vol. 130, no. 7, pp. 1837–1872, 2022.
- [51] Y. Zhang, B. Kang, B. Hooi, S. Yan, and J. Feng, "Deep long-tailed learning: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 9, pp. 10 795–10 816, 2023.
- [52] L. Xiang, G. Ding, and J. Han, "Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification," in *Proceedings of the 16th European Conference on Computer Vision*, 2020, pp. 247–263.
- [53] X. Wang, L. Lian, Z. Miao, Z. Liu, and S. Yu, "Long-tailed recognition by routing diverse distribution-aware experts," in *International Conference on Learning Representations*, 2021.
- [54] J. Cui, S. Liu, Z. Tian, Z. Zhong, and J. Jia, "ResLT: Residual learning for long-tailed recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 3695–3706, 2022.
- [55] Y. Zhang, B. Hooi, L. Hong, and J. Feng, "Self-supervised aggregation of diverse experts for test-agnostic long-tailed recognition," in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 34 077–34 090.
- [56] J.-X. Shi, T. Wei, and Y.-F. Li, "Residual diverse ensemble for long-tailed multi-label text classification," *SCIENCE CHINA Information Sciences*, 2024.
- [57] B. Kang, S. Xie, M. Rohrbach, Z. Yan, A. Gordo, J. Feng, and Y. Kalantidis, "Decoupling representation and classifier for long-tailed recognition," in *International Conference on Learning Representations*, 2020.
- [58] Z. Zhong, J. Cui, S. Liu, and J. Jia, "Improving calibration for long-tailed recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16 489–16 498.
- [59] T. Wei and K. Gan, "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*, 2023, pp. 3469–3478.
- [60] G. Nam, S. Jang, and J. Lee, "Decoupled training for long-tailed classification with stochastic representations," in *International Conference on Learning Representations*, 2023.
- [61] S. Sun, H. Lu, J. Li, Y. Xie, T. Li, X. Yang, L. Zhang, and J. Yan, "Rethinking classifier re-training in long-tailed recognition: Label over-smooth can balance," in *The Thirteenth International Conference on Learning Representations*, 2025.
- [62] A. P. Steiner, A. Kolesnikov, X. Zhai, R. Wightman, J. Uszkoreit, and L. Beyer, "How to train your vit? data, augmentation, and regularization in vision transformers," *Transactions on Machine Learning Research*, 2022.
- [63] K. Zhou, J. Yang, C. C. Loy, and Z. Liu, "Learning to prompt for vision-language models," *International Journal of Computer Vision*, vol. 130, no. 9, pp. 2337–2348, 2022.
- [64] K. Zhou, J. Yang, C. Loy, and Z. Liu, "Conditional prompt learning for vision-language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 816–16 825.
- [65] B. X. Yu, J. Chang, H. Wang, L. Liu, S. Wang, Z. Wang, J. Lin, L. Xie, H. Li, Z. Lin, Q. Tian, and C. W. Chen, "Visual tuning," *ACM Computing Surveys*, vol. 56, no. 12, 2024.
- [66] Z. Zhou, M. Yang, J.-X. Shi, L.-Z. Guo, and Y.-F. Li, "DeCoOp: Robust prompt tuning with out-of-distribution detection," in *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [67] A. Iscen, A. Fathi, and C. Schmid, "Improving image recognition by retrieving from web-scale image-text data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 19 295–19 304.
- [68] P. Xia, D. Xu, L. Ju, M. Hu, J. Chen, and Z. Ge, "LMPT: Prompt tuning with class-specific embedding loss for long-tailed multi-label visual recognition," *arXiv preprint arXiv:2305.04536*, 2023.
- [69] X. He, S. Fu, X. Ding, Y. Cao, and H. Wang, "Uniformly distributed category prototype-guided vision-language framework for long-tail recognition," in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 5027–5037.
- [70] Y. Song, M. Li, and B. Wang, "Long-tailed visual recognition via improved cross-window self-attention and trivialaugmt," *IEEE Access*, vol. 11, pp. 49 601–49 610, 2023.
- [71] B. Li, Y. Yao, J. Tan, R. Gong, J. Lu, and Y. Luo, "Rectify representation bias in vision-language models for long-tailed recognition," *Neural Networks*, vol. 172, p. 106134, 2024.
- [72] S. Zhao, X. Wen, J. Liu, C. Ma, C. Yuan, and X. Qi, "Learning from neighbors: Category extrapolation for long-tail learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025, pp. 30 483–30 492.
- [73] J. Chen, B. Qin, J. Li, H. Chen, and B. Su, "Rethinking the bias of foundation model under long-tailed distribution," in *Proceedings of the 42nd International Conference on Machine Learning*, 2025, pp. 8910–8930.
- [74] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 5998–6008.
- [75] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 000–16 009.
- [76] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [77] Y. Zhao, W. Chen, X. Tan, K. Huang, and J. Zhu, “Adaptive logit adjustment loss for long-tailed visual recognition,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 3, 2022, pp. 3472–3480.
- [78] M. Li, Y.-m. Cheung, and Y. Lu, “Long-tailed visual recognition via gaussian clouded logit adjustment,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6929–6938.
- [79] E. B. Zaken, Y. Goldberg, and S. Ravfogel, “BitFit: Simple parameter-efficient fine-tuning for transformer-based masked language-models,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2022, pp. 1–9.
- [80] M. Jia, L. Tang, B.-C. Chen, C. Cardie, S. Belongie, B. Hariharan, and S.-N. Lim, “Visual prompt tuning,” in *Proceedings of the 17th European Conference on Computer Vision*, 2022, pp. 709–727.
- [81] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. De Larousilhe, A. Gesmundo, M. Attariyan, and S. Gelly, “Parameter-efficient transfer learning for NLP,” in *Proceedings of the 36th International Conference on Machine Learning*, 2019, pp. 2790–2799.
- [82] E. J. Hu, Yelong shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “LoRA: Low-rank adaptation of large language models,” in *International Conference on Learning Representations*, 2022.
- [83] S. Chen, C. GE, Z. Tong, J. Wang, Y. Song, J. Wang, and P. Luo, “AdaptFormer: Adapting vision transformers for scalable visual recognition,” in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 16 664–16 678.
- [84] T. Wei, W.-W. Tu, Y.-F. Li, and G.-P. Yang, “Towards robust prediction on tail labels,” in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 1812–1820.
- [85] A. Kumar, A. Raghunathan, R. M. Jones, T. Ma, and P. Liang, “Fine-tuning can distort pretrained features and underperform out-of-distribution,” in *International Conference on Learning Representations*, 2022.
- [86] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *International Conference on Learning Representations*, 2015.
- [87] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [88] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le, “AutoAugment: Learning augmentation strategies from data,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 113–123.
- [89] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, “RandAugment: Practical automated data augmentation with a reduced search space,” in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 18 613–18 624.
- [90] Y. Sun, X. Wang, Z. Liu, J. Miller, A. Efros, and M. Hardt, “Test-time training with self-supervision for generalization under distribution shifts,” in *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 9229–9248.
- [91] D. Wang, E. Shelhamer, S. Liu, B. Olshausen, and T. Darrell, “Tent: Fully test-time adaptation by entropy minimization,” in *International Conference on Learning Representations*, 2021.
- [92] Z. Zhou, L.-Z. Guo, L.-H. Jia, D. Zhang, and Y.-F. Li, “ODS: Test-time adaptation in the presence of open-world data shift,” in *Proceedings of the 40th International Conference on Machine Learning*, 2023, pp. 42 574–42 588.
- [93] C. Snell, J. Lee, K. Xu, and A. Kumar, “Scaling llm test-time compute optimally can be more effective than scaling model parameters,” *arXiv preprint arXiv:2408.03314*, 2024.
- [94] Z. Chen, W. Wang, Y. Cao, Y. Liu, Z. Gao, E. Cui, J. Zhu, S. Ye, H. Tian, Z. Liu *et al.*, “Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling,” *arXiv preprint arXiv:2412.05271*, 2024.
- [95] J. Li, Z. Tan, J. Wan, Z. Lei, and G. Guo, “Nested collaborative learning for long-tailed visual recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 6949–6958.

- [96] Z. Xu, R. Liu, S. Yang, Z. Chai, and C. Yuan, “Learning imbalanced data with vision transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 15 793–15 803.
- [97] M.-K. Suh and S.-W. Seo, “Long-tailed recognition by mutual information maximization between latent features and ground-truth labels,” in *Proceedings of the 40th International Conference on Machine Learning*, vol. 202, 2023, pp. 32 770–32 782.
- [98] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 2980–2988.
- [99] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, “Class-balanced loss based on effective number of samples,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 9268–9277.
- [100] D. Lian, D. Zhou, J. Feng, and X. Wang, “Scaling & shifting your features: A new baseline for efficient model tuning,” in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 109–123.
- [101] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” *arXiv preprint arXiv:1607.06450*, 2016.



Jiang-Xin Shi received the BSc degree in 2020. He is currently working toward the PhD degree at the School of Artificial Intelligence, the National Key Laboratory for Novel Software Technology at Nanjing University, China. His research interests focus on long-tail learning. He served as the Program Committee Member for top-tier conferences, e.g., ICML/NeurIPS/ICLR/KDD/AAAI, and the Senior Program Committee Member for IJCAI 2025.



Tong Wei received his Ph.D. in Computer Science from Nanjing University, China, in 2021. He is currently an associate professor at the School of Computer Science and Engineering at Southeast University. His research focuses on trustworthy machine learning, and he has published over 30 papers in top-tier conferences and journals. Dr. Wei has also served as an Area Chair for premier conferences, including ICML'25, NeurIPS'25, and IJCAI'25.



Yu-Feng Li (Senior Member, IEEE) received the BSc and PhD degrees in computer science from Nanjing University, China in 2006 and 2013, respectively. He is currently a full professor at the School of Artificial Intelligence, the National Key Laboratory for Novel Software Technology at Nanjing University. He is a member of the LAMDA group, led by Prof. Zhi-Hua Zhou (IEEE/ACM/AAAI Fellow). His research interests focus on robust and reliable machine learning. He published over 90 academic papers in top-tier journals and conferences in the field, with around 7000 citations. He serves as journal associate/action editor for Artificial Intelligence, Machine Learning, Neural Networks, etc. He served as program co-chair for IEEE Big Comp 2020/CCML 2021, and area chairs for top-tier conferences, e.g., ICML/NeurIPS/ICLR/IJCAI. The research work has been selected for the IJCAI 2021 Early-Career Spotlight Talk. He won the PAKDD Early-Career Research Award 2024. He is the associated program co-chair for IJCAI 2025.