
Rethinking the Bias of Foundation Model under Long-tailed Distribution

Jiahao Chen^{*123} Bin Qin^{*45} Jiangmeng Li^{*45} Hao Chen⁶ Bing Su¹²³

Abstract

Long-tailed learning has garnered increasing attention due to its practical significance. Among the various approaches, the fine-tuning paradigm has gained considerable interest with the advent of foundation models. However, most existing methods primarily focus on leveraging knowledge from these models, overlooking the inherent biases introduced by the imbalanced training data they rely on. In this paper, we examine how such imbalances from pre-training affect long-tailed downstream tasks. Specifically, we find the imbalance biases inherited in foundation models on downstream tasks as parameter imbalance and data imbalance. During fine-tuning, we observe that parameter imbalance plays a more critical role, while data imbalance can be mitigated using existing re-balancing strategies. Moreover, we find that parameter imbalance cannot be effectively addressed by current re-balancing techniques, such as adjusting the logits, during training, unlike data imbalance. To tackle both imbalances simultaneously, we build our method on causal learning and view the incomplete semantic factor as the confounder, which brings spurious correlations between input samples and labels. To resolve the negative effects of this, we propose a novel backdoor adjustment method that learns the true causal effect between input samples and labels, rather than merely fitting the correlations in the data. Notably, we achieve an average performance increase of about 1.67% on each dataset.

1. Introduction

Real-world data often follows a long-tailed distribution, where the majority of instances are concentrated in head classes, leaving only a small number of instances for each tail class. This scarcity of samples makes generalization on them challenging, and naive learning on such data tends to introduce an undesirable bias toward dominant labels. Recently, with the advent of foundation models, downstream performance can be significantly improved by fine-tuning these models on labeled data, often yielding superior results compared to training from scratch, while also minimizing training costs (Wang et al., 2022b;a). Consequently, there has been considerable interest in exploring how to fine-tune foundation models and leverage their strong generalization capabilities to enhance learning at downstream tasks.

Recent works such as LIFT (Shi et al., 2024), LPT (Dong et al., 2022), and VL-LTR (Tian et al., 2022) demonstrate that properly fine-tuning foundation models like CLIP (Radford et al., 2021) can significantly enhance long-tail learning performance. VL-LTR improves visual recognition, particularly for tail classes, by collecting class descriptions from the internet and jointly learning both visual and text representations. LIFT, on the other hand, reveals that heavy fine-tuning hurts and Parameter-Efficient Fine-Tuning (Chen et al., 2022a; Jia et al., 2022) (PEFT) based methods can preserve as much information from the foundation model as possible, which is important for the downstream imbalanced learning. However, these methods tend to *overemphasize the use of foundation models to solve downstream data imbalance while overlooking their inherent biases*, as shown in Fig. 1. Large-scale datasets used to train foundation models, such as LAION (Schuhmann et al., 2021), also follow a long-tailed distribution, which can negatively impact downstream tasks (Zhu et al., 2024; Wen et al., 2024). Therefore, fine-tuned models are influenced by dual long-tailed distributions (upstream and downstream), and only considering data imbalance is insufficient.

In this paper, we explore *how the imbalance of foundation models impacts downstream imbalanced tasks in PEFT-based methods*. Since the pre-training data is often inaccessible, its influence is primarily reflected in the pre-trained weights, or parameters, which we refer to as **parameter imbalance**. In contrast, downstream data is accessible and

^{*}Equal contribution ¹Gaoling School of Artificial Intelligence, Renmin University of China ²Beijing Key Laboratory of Research on Large Models and Intelligent Governance ³Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE ⁴Institute of Software Chinese Academy of Sciences ⁵University of Chinese Academy of Sciences ⁶Electrical and Computer Engineering, Carnegie Mellon University. Correspondence to: Bing Su <subingats@gmail.com>.




Figure 1. Previous methods focus on using the downstream data to fine-tune the foundation model while ignoring that the pre-training data has a potential influence of bias (dashed line).

directly affects the downstream task, which we define as **data imbalance**. Through fine-tuning, we find that both types of imbalance influence downstream tasks, but parameter imbalance plays a more significant role, as shown in Fig. 3 and Fig. 2. In addition, we locate that samples belonging to the tail classes grouped by the data and parameter imbalance at the same time are influenced extremely. Due to the inaccessibility of pre-training data, we estimate the label prior and extend the current Generalized Logit Adjustment (GLA) (Zhu et al., 2024) into the training phase and named the method GLA-Train. Unlike Logit Adjustment (Menon et al., 2020)(LA), which effectively addresses data imbalance, we find that GLA-Train cannot be used to alleviate parameter imbalance, as shown in Tab. 3. After measuring the quality of feature representation via K-Nearest-Neighbor (KNN) (Cover & Hart, 1967) accuracy, we surprisingly find that LA slightly enhances the feature representation of tail classes with the improvement mostly coming from the classifier, as shown in Tab. 4. When the imbalance is integrated into the parameter, it is difficult to alleviate it by solely adjusting the logit. Therefore, we conclude that parameter imbalance is fundamentally distinct from data imbalance, presenting a critical challenge that cannot be effectively addressed through re-balancing methods.

To tackle both parameter and data imbalances simultaneously, we build a causal structure graph and find that the incomplete semantic factor is the confounder, encouraging the model to learn spurious correlations between input samples and labels, thus restricting its generalization ability. For instance, if the class “dog” belongs to the tail classes due to parameter imbalance, the foundation model may lack sufficient semantic information, causing it to capture only partial features, such as the dog’s head, to represent the class. When this model is adapted to a downstream task where “dog” is also a tail class due to data imbalance, fine-tuning struggles to learn the complete set of relevant semantic features, further limiting its generalization. We denote the dog’s head as a typical example of an incomplete semantic factor. To inhibit the confounding effect, we adopt the backdoor criterion in causal inference to realize our backdoor adjustment. After applying our method, we achieve a more balanced performance over all the classes across different

downstream datasets.

Our contributions are summarized as follows:

- We address a practical challenge where both the data used to train the foundation model and the data for downstream tasks exhibit imbalance, referred to as parameter imbalance and data imbalance, respectively. During fine-tuning, we observe that parameter imbalance has a more significant impact than data imbalance, and existing re-balancing methods are ineffective at addressing it. Furthermore, attempts to mitigate parameter imbalance provide limited improvement to the quality of representations.
- We find that the incomplete semantic factor encourages the model to learn spurious correlations between input samples and labels, which restricts its generalization ability. To solve that, we construct a causal graph and propose a backdoor adjustment method to eliminate the confounder negative impact.
- We achieve at least 1.5%, 1.5%, 2.0% performance gains on ImageNet-LT (Deng et al., 2009), Places365-LT (Liu et al., 2019), and iNaturalist2018 (Van Horn et al., 2018) compared with state-of-the-art methods.

2. Preliminary

Notation. We aim to solve a C -way classification problem with instances $x \in \mathcal{X}$ and labels $y \in \mathcal{Y} = [C] = \{1, \dots, C\}$, where \mathcal{X} and \mathcal{Y} denote the input space and output space. For pre-training, we denote \mathcal{D}_P as the distribution of the training set and we cannot access it at the fine-tuning phase. For downstream tasks, we denote \mathcal{D}_S and \mathcal{D}_T as the distribution for training and test, respectively. In the context of long-tailed learning, the number of samples varies across classes, i.e., $\mathbb{P}_S(Y = 1) \neq \mathbb{P}_S(Y = 2) \neq \dots \neq \mathbb{P}_S(Y = C)$, where $\mathbb{P}_S(Y)$ denotes the class prior of \mathcal{D}_S . In contrast, the test set T is sampled from the distribution \mathcal{D}_T , where each class c has an equal probability, i.e., $\mathbb{P}_T(Y = c) = 1/C$. Our target is to learn a hypothesis $f : \mathcal{X} \rightarrow \mathcal{Y}$ to estimate $\mathbb{P}_S(Y = y | x)$ from the training set and generalize it to the test set.

Logit adjustment (LA). Under the label shift assumption, we have $\mathbb{P}_S(x | Y = y) = \mathbb{P}_T(x | Y = y)$ but $\mathbb{P}_S(Y = y | x) \neq \mathbb{P}_T(Y = y | x)$ for each class y . LA (Menon et al., 2020) bridges the gap between the posterior of the imbalanced training set and the balanced test set. As shown in Eq. 1, we can get the posterior of the test set from the

Table 1. The Zero-Shot (ZS) performance of different type of CLIPs on different benchmarks.

	ImageNet-LT			Places365-LT			iNaturalist2018		
	D-Many	D-Medium	D-Few	D-Many	D-Medium	D-Few	D-Many	D-Medium	D-Few
CLIP	67.87	66.27	66.03	36.59	38.01	45.53	3.60	4.38	4.10
OpenCLIP	67.15	65.02	64.31	41.40	38.06	40.28	2.45	2.80	2.38
MetaCLIP	70.98	68.70	68.57	38.79	37.93	40.44	5.07	6.37	6.00

training set by introducing a scaling factor.

$$\begin{aligned}
 \mathbb{P}_T(Y = y | \mathbf{x}) &= \frac{\mathbb{P}_T(\mathbf{x} | Y = y)\mathbb{P}_T(Y = y)}{\mathbb{P}_T(\mathbf{x})} \\
 &\propto \mathbb{P}_S(\mathbf{x} | Y = y)\mathbb{P}_T(Y = y) \quad (1) \\
 &\propto \frac{\mathbb{P}_T(Y = y)}{\mathbb{P}_S(Y = y)}\mathbb{P}_S(Y = y | \mathbf{x})
 \end{aligned}$$

There are two types of LA, either applied post-hoc to a trained model or enforced in the loss during training, and the latter can achieve better performance. When integrating it to the criterion, we get the final LA loss, as shown in Eq. 2, where $f_y(\mathbf{x} | \boldsymbol{\theta}, \phi) \propto \mathbb{P}(Y = y | \mathbf{x})$, $\boldsymbol{\theta}$, and ϕ denote the output logit of class y , the parameter of the foundation model, and the additional fine-tuned parameter, respectively. If not specified, we freeze the $\boldsymbol{\theta}$ and only optimize the ϕ .

$$\begin{aligned}
 L_{LA}(f(\mathbf{x} | \boldsymbol{\theta}, \phi), y) &= \log \left[1 \right. \\
 &\quad \left. + \sum_{y' \neq y} \left(\frac{\mathbb{P}_S(Y = y')}{\mathbb{P}_S(Y = y)} \right) \cdot \frac{e^{f_{y'}(\mathbf{x} | \boldsymbol{\theta}, \phi)}}{e^{f_y(\mathbf{x} | \boldsymbol{\theta}, \phi)}} \right] \quad (2)
 \end{aligned}$$

Generalized logit adjustment (GLA). Since the label prior $\mathbb{P}_P(Y = y)$ of the pre-training data is inaccessible and we cannot use Eq. 1 to estimate the posterior (from $\mathbb{P}_S(\mathbf{x} | Y = y)$, $\mathbb{P}_S(Y = y)$, and $\mathbb{P}_S(Y = y | \mathbf{x})$ to $\mathbb{P}_P(\mathbf{x} | Y = y)$, $\mathbb{P}_P(Y = y)$, and $\mathbb{P}_P(Y = y | \mathbf{x})$). GLA (Zhu et al., 2024) estimates the prior following the Eq. 3 on the validation set, where L_{CE} denotes the Cross-Entropy (CE) loss. After getting the estimated $\hat{\mathbb{P}}_P(Y)$, we can adjust the logit following Eq. 1.

$$\begin{aligned}
 \hat{\mathbb{P}}_P(Y) &= \min_{\mathbf{q}} \max_{\lambda \geq 0, v} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_T} L_{CE}(f(\mathbf{x} | \boldsymbol{\theta}) - \log \mathbf{q}, y) \\
 &\quad - \sum_i \lambda_i \mathbf{q}_i + v(1 - \sum_{i \in [C]} \mathbf{q}_i) \quad (3)
 \end{aligned}$$

GLA assumes that the Zero-Shot (ZS) and Fine-Tuned (FT) models have diverse predictions (Zhu et al., 2024). In this way, we can achieve unbiased predictions by simply ensembling the output logit of two models, as shown in Eq. 4, where $f(\mathbf{x} | \boldsymbol{\theta})$ and $f(\mathbf{x} | \boldsymbol{\theta}, \phi)$ denote the output logit of ZS and FT, respectively. Since the adjustments of the two models are individual, we split Eq. 4 into GLA-ZS and GLA-FT, eliminating the bias of the foundation model and

Table 2. The Average Accuracy Gap between different foundation models on Places365-LT with ZS, CE, and LA.

$f - g$	ZS	CE	LA
		Adaptformer	
CLIP–OpenCLIP	10.38	3.19	3.49
CLIP–MetaCLIP	11.56	3.22	3.32
OpenCLIP–MetaCLIP	12.03	3.03	3.40
		VPT	
CLIP–OpenCLIP	10.38	3.6	3.94
CLIP–MetaCLIP	11.56	3.42	3.87
OpenCLIP–MetaCLIP	12.03	3.10	3.53

downstream data, respectively.

$$\begin{aligned}
 f'_y(\mathbf{x}) &= \underbrace{f_y(\mathbf{x} | \boldsymbol{\theta}) - \log \hat{\mathbb{P}}_P(Y = y)}_{GLA-ZS} \\
 &\quad + \underbrace{f_y(\mathbf{x} | \boldsymbol{\theta}, \phi) - \log \mathbb{P}_S(Y = y)}_{GLA-FT} \quad (4)
 \end{aligned}$$

3. The influence of parameter imbalance

3.1. Parameter imbalance is more important

We consider the problem of the pre-training data and the downstream training data being both imbalanced. For the parameter imbalance, we provide a formal theoretical definition in Definition 3.1. Furthermore, we offer a more detailed explanation of the relationship between parameter imbalance and data imbalance in Appendix Section A.3. Since the $\mathbb{P}_P(Y)$ is not accessible, we use the estimated prior $\hat{\mathbb{P}}_P(Y)$ to substitute $\mathbb{P}_P(Y)$ for the following analysis.

Definition 3.1. Let $\mathbb{P}_P(Y)$ denote the label of the pre-training data. We have $\mathbb{P}_P(Y = 1) \neq \mathbb{P}_P(Y = 2) \neq \dots \neq \mathbb{P}_P(Y = C)$. We use the imbalanced factor (IF) $IF = \max_{c \in [C]} \mathbb{P}_P(Y) / \min_{c \in [C]} \mathbb{P}_P(Y)$ to measure the degree of parameter imbalance.

Most previous works (Dong et al., 2022; Shi et al., 2024; Tian et al., 2022) focus on eliminating the data bias to improve the performance but ignore the parameter imbalance. In the following, we will explore the impact of the superposition of two imbalances.

Experiment setup. We conduct experiments with ViT-B/16 (Dosovitskiy et al., 2020) on the ImageNet-LT, the




Figure 2. The data and parameter imbalance on (a) ImageNet-LT and (b) Places365-LT. The class indices of the left picture are sorted relying on the data imbalance while the right picture relies on the parameter imbalance. Curves are smoothed for better visualization. Parameter imbalance occupies a more vital role.

P-Many P-Medium P-Few	(a) CE			P-Many P-Medium P-Few	(b) LA		
	D-Many	D-Medium	D-Few		D-Many	D-Medium	D-Few
	75.05	53.06	34.31		72.51	66.94	63.34
	58.10	35.93	21.13		49.95	50.88	47.35
	40.72	25.36	10.56		35.78	39.87	31.12

Figure 3. The performance of different groups with (a) CE and (b) LA on the Places365-LT dataset. The three rows, from top to bottom, represent the performance of P-Many, P-Medium, and P-Few, respectively. The three columns, from left to right, represent the performance of D-Many, D-Medium, and D-Few, respectively.

Places365-LT, and iNaturalist2018 datasets. To better explore the influence of pre-trained data, we select CLIP (Radford et al., 2021), OpenCLIP (Cherti et al., 2023), and MetaCLIP (Xu et al., 2023a) as foundation models, which are pre-trained on WIT, LAION (Schuhmann et al., 2021), and MetaData, respectively. Previous works (Shi et al., 2024; Dong et al., 2022) find that we can achieve promising performance by fine-tuning only a small proportion of parameters. Therefore, we select two typical PEFT techniques, Adaptformer (Chen et al., 2022a) and Visual Prompt Tuning (VPT) (Jia et al., 2022), as the basic methods. The learning rate, number of epochs, and parameter initialization strategies follows (Shi et al., 2024). Following OLTR (Liu et al., 2019), we split the classes into three groups named “D-Many”, “D-Medium”, and “D-Few” relying on the number of samples. Similarly, for parameter imbalance, we split the classes into three groups named “P-Many”, “P-Medium”, and “P-Few” relying on $\mathbb{P}_P(Y)$. More details are in the Appendix Sec. A.

Different foundation models serve different parameter imbalance. We report the ZS performance on different long-tailed datasets, as shown in Tab. 1. Since the imbalanced downstream dataset does not influence the foundation models, it is natural that we cannot observe the data imbal-

ance, i.e., “D-Many” and “D-Few” should have a similar performance. However, on the Places365-LT, “D-Few” is better than “D-Many” for the CLIP model and other foundation models are nearly similar. Since they all share the architecture, the main differences come from the imbalance pre-training data, i.e., parameter imbalance. To better measure the difference, we define Average Accuracy Gap $\Delta_{avg}^{f,g}$, where f, g and $Acc(f, c)$ denote two arbitrary hypotheses and the c -th class-wise accuracy of model f , respectively.

$$\Delta_{avg}^{f,g} = \frac{1}{C} \sum_{c=1}^C |Acc(f, c) - Acc(g, c)| \quad (5)$$

As shown in Tab. 2, different foundations have merely 10% differences on the same downstream task. When we adapt them to the downstream, the gap is decreased but still exists.

The parameter imbalance occupies. When we adapt the foundation model to the downstream task, the fine-tuned model is influenced by parameter and data imbalance. For analysis, we conduct experiments with different criteria, CE and LA, on the CLIP model. As shown in Fig. 2, CE suffers from the data imbalance heavily although we adapt from the foundation model, which is not influenced by it. When changing the criterion from CE to LA, the data imbalance is relieved effectively and we achieve a more balanced performance. As for parameter imbalance, although the fine-tuning-based method can alleviate it, the bias still exists. We primarily give two explanations for this phenomenon. One is that current criteria do not consider the parameter imbalance explicitly. Another is that the PEFT method fixes the parameters of the foundation model, while retaining strong generalization capabilities, the parameter imbalance is also preserved. As for data imbalance, only a small proportion of parameters is influenced by it, and the re-balancing technique, i.e. LA, can easily eliminate it. Therefore, parameter imbalance plays a more vital role than data imbalance.

Tail-Tail hurts. Parameter imbalance is a critical factor that impedes further improvement. To analyze this, we consider two types of imbalances in conjunction and divide the validation sets into nine groups. For instance, some

classes can be categorized as “D-Many” due to data imbalance, while simultaneously classified as “P-Few” because of parameter imbalance. As shown in Fig. 3, we observe that classes falling into both “D-Few” and “P-Few” classes exhibit the poorest performance. Moreover, when we shift the criterion from CE to LA, we find that the performance of “P-Few” classes within the “D-Many” group deteriorates. This indicates that previous claims suggesting LA can alleviate parameter imbalance do so at the expense of “P-Few” class performance, failing to address the parameter imbalance.

3.2. Analysis of adjustment for parameter imbalance

Addressing parameter imbalance is a crucial task. Drawing inspiration from methods used to tackle data imbalance, a natural approach to mitigating this issue is to adjust the output logits, similar to the Logit Adjustment (LA) technique. The existing Generalized Logit Adjustment (GLA) method addresses parameter imbalance by modifying the logits according to the estimated label priors. Thus, in this section, we explore whether parameter imbalance can be effectively eliminated through the simple adjustment of output logits.

GLA fails during training. Firstly, We follow GLA to estimate the prior of the pre-training data and then adjust the output logit following Eq. 4. As shown in Tab. 3, we achieve a more balanced performance and verify its effectiveness. However, GLA achieves unbiased prediction by modeling the parameter imbalance and data imbalance respectively, which is more complicated. Referring to LA, which models the adjustment into the criterion, a natural improvement is extending GLA into training. This approach mitigates parameter bias by introducing additional unbiased parameters. We denote this method as GLA-Train and constitute the optimization target as shown in Eq. 2. Since the parameter imbalance does not influence the classifier, we only use Eq. 6 to learn the unbiased representation and introduce an additional stage for classifier re-training (Kang et al., 2019).

$$L_{GLA}(f(\mathbf{x} | \boldsymbol{\theta}, \phi), y) = \log \left[1 + \sum_{y' \neq y} \left(\frac{\mathbb{P}_S(Y = y') \hat{\mathbb{P}}_P(Y = y')}{\mathbb{P}_S(Y = y) \hat{\mathbb{P}}_P(Y = y)} \right) \cdot \frac{e^{f_{y'}(\mathbf{x} | \boldsymbol{\theta}, \phi)}}{e^{f_y(\mathbf{x} | \boldsymbol{\theta}, \phi)}} \right] \quad (6)$$

The results are presented in Tab.3. In comparison with GLA, GLA-Train offers minimal additional benefit in addressing both parameter and data imbalance, achieving a performance similar to that of LA, as shown in Fig.3. (More experimental results are in the Appendix Sec. I.2.) Since Eq. 6 explicitly models both data and parameter imbalance, it is important to note that only the data imbalance is effectively corrected while parameter imbalance is ignored. This indicates that parameter imbalance is fundamentally different from data imbalance and cannot be resolved through

Table 3. Ther performance of GLA, GLA-ZS, GLA-FT, and GLA-T based on the CLIP model. We denote GLA-T as the GLA-Train.

	Type	D-Many	D-Medium	D-Few	Overall
ImageNet-LT	GLA	79.76	76.48	74.12	77.42
	GLA-ZS	70.23	68.61	68.90	69.27
	GLA-FT	79.67	76.15	73.29	77.12
	GLA-T	80.21	75.93	72.01	77.10
Places365-LT	GLA	51.22	52.00	53.34	51.98
	GLA-ZS	40.58	39.57	46.63	41.31
	GLA-FT	50.41	52.23	52.18	51.56
	GLA-T	51.36	52.30	50.91	51.48

Table 4. The KNN accuracy. We denote GLA-T as the GLA-Train.

		D-Many	D-Medium	D-Few	All
ImageNet-LT	CE	74.61	70.02	66.14	71.25
	LA	74.42	70.35	66.43	71.36
	GLA-T	74.33	70.35	66.31	71.26
Places365-LT	CE	44.53	46.01	47.92	45.85
	LA	44.14	46.23	48.75	45.93
	GLA-T	44.33	46.14	48.66	45.93

simple adjustment alone.

Feature representation analysis. To explore how the re-balance influences the feature representation, we utilize the balanced validation sets of ImageNet-LT and Places365-LT to assess the feature quality of the test set through KNN accuracy. As presented in Tab. 4, we observe that all three methods yield comparable accuracy. The re-balancing-based methods provide slight improvements for the “D-Few” and “D-Medium” classes, but this comes at the expense of performance in the “D-Many” class. This suggests that the enhanced performance for tail classes can be attributed to a more accurate classifier, which can be achieved by employing re-balancing techniques like Logit Adjustment (LA). We also conduct experiments to explore the influence of classifier in the Appendix Sec. I.3. However, addressing parameter imbalance—manifested in the parameters of foundation models—proves to be more challenging. It is difficult to mitigate the negative effects of parameter imbalance solely through the classifier. Consequently, re-balancing methods cannot effectively alleviate parameter imbalance.

4. Method

As previously analyzed, the parameter imbalance becomes dominant after adapting to the downstream task, and we cannot eliminate it solely by adjusting the logits during the training phase. Therefore, exploring an effective method to address both parameter imbalance and data imbalance simultaneously is essential. In this section, instead of focusing on re-balancing methods, we analyze the issue from a causal perspective. By identifying the latent confounder,

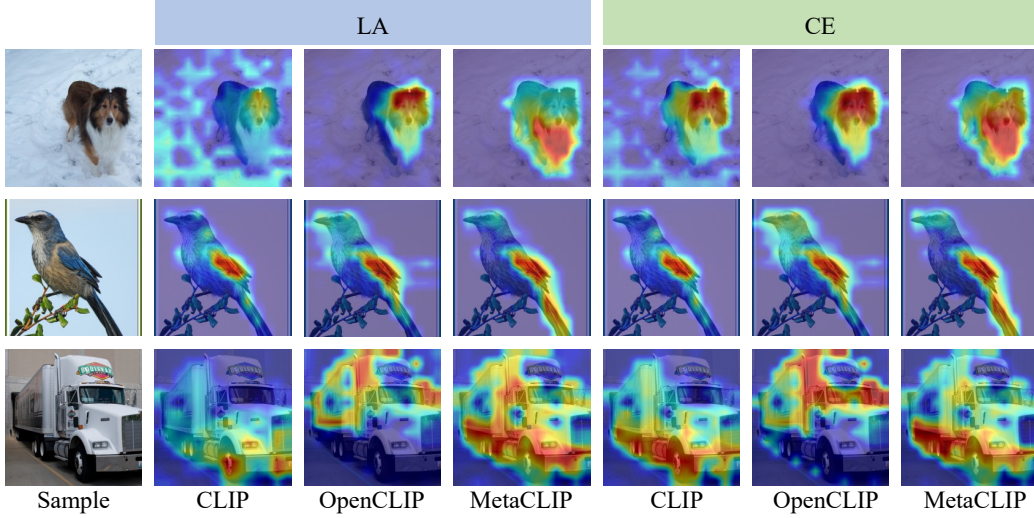


Figure 4. We randomly select three samples of different tail classes and visualize the heatmap via Grad-CAM (Selvaraju et al., 2017) with CE and LA. Different models draw attention to different areas and merging them can acquire unbiased semantic information.




Figure 5. The framework of our proposed method. (a) In the confounded setting, C influences both X and Y , leading to the backdoor path $X \leftarrow C \rightarrow Y$, thereby introducing confounding bias. (b) After intervening on X , its parent nodes C and U are severed, eliminating the unstable backdoor path and leading to a more reliable estimation of the causal effect between X and Y .

we can apply the backdoor criterion to estimate its negative impact, thereby achieving a more balanced performance across various classes.

4.1. A causal view for long-tailed learning

We model the long-tailed image classification process with a causal structure graph as shown in Fig. 5. Here, we denote X , C , U , and D as an imbalanced dataset, the incomplete semantic factor, the inaccessible semantic factor, and the parameter imbalance, respectively. B is the data balanced representation and Y is the predicted label distribution. In Fig. 5, $X \rightarrow B$ denotes we extract the data balanced representation from an imbalanced dataset. We can achieve this by applying re-balancing based method like LA. $B \rightarrow Y$ denotes the inference pipeline that we predict the label relying on the given representation. $D \rightarrow C$ presents the incom-

plete semantic factor depending on the parameter imbalance. From the perspective of data generation, $U \rightarrow X \leftarrow C$ denotes the dataset can be generated by giving the accessible semantic factor and inaccessible semantic factor, i.e., $P(X) = \sum P(X|C, U)P(C)P(U)$. $C \rightarrow Y$ indicates that the predicted label distributions follow their own preferences for incomplete semantic factors. We also give additional experimental evidence for $C \rightarrow Y$ in Sec. I.9.

Therefore, the incomplete semantic factor C in our problem setting, acting as a confounder (Pearl, 2009), can create a backdoor path $X \leftarrow C \rightarrow Y$, leading to spurious correlation between X and Y . If we ignore the influence of incomplete semantic factor and learn the posterior $\mathbb{P}_S(Y = y | \mathbf{x})$ to estimate $\mathbb{P}_T(Y = y | \mathbf{x})$, the spurious correlation will be modeled, which leads to a biased model. As shown in Fig. 4, we visualize given samples on different fine-tuned models, where the corresponding foundation models serve different parameter imbalances, i.e. incomplete semantic factor. For the specific picture, different models are drawn in different interesting areas. The relationship between the confounding path $X \leftarrow C \rightarrow Y$ is unstable. For example, in the first row of Fig. 4, OpenCLIP is more attracted by the head while MetaCLIP is more interested in the body of the dog. Therefore, if a test sample belonging to this class but the body is obscured, MetaCLIP is easier to give the wrong prediction. We analyze that each incomplete semantic factor can provide sufficient information to distinguish the given class from others on the training set. However, it also indicates that the incomplete semantic factor can prohibit the model from learning other relevant information, which limits its generalization. Therefore, eliminating the influence of incomplete semantic factor is vital for imbalanced learning.

4.2. Backdoor adjustment

A more generalized fine-tuned model should be independent of incomplete semantic factor (i.e. parameter imbalance), which inhibits confounding effects from C . Therefore, instead of estimating $\mathbb{P}(Y = y | \mathbf{x})$, we use backdoor criterion and estimate $\mathbb{P}(Y = y | do(\mathbf{x}))$, where $do()$ is exploited to cut off the connection from the C to X . Considering that each sample in X uniquely corresponds to a balanced representation in B , which indicates that the mapping between X and B is injective. Thus there only exists a certain \mathbf{b} such that $\mathbb{P}(\mathbf{b} | \mathbf{x}) = 1$ and $\mathbb{P}(\mathbf{b}' | \mathbf{x}) = 0$ for $\mathbf{b} \neq \mathbf{b}'$. Then, we propose our backdoor adjustment:

$$\begin{aligned} \mathbb{P}(Y = y | do(\mathbf{x})) &= \sum_{\mathbf{b} \in B} \sum_{c \in C} \mathbb{P}(Y = y | \mathbf{b}, c) \mathbb{P}(c) \mathbb{P}(\mathbf{b} | \mathbf{x}) \\ &= \sum_{c \in C} \mathbb{P}(Y = y | \mathbf{b}, c) \mathbb{P}(c) \end{aligned} \quad (7)$$

For simplicity, we assume $\mathbb{P}(c) = 1/M$, where M is the number of incomplete semantic factors. Intuitively, $\mathbb{P}(Y = y | do(\mathbf{x}))$ can be estimated by fusing $\mathbb{P}(Y = y | \mathbf{b}, c)$, where \mathbf{b} represents the balanced representation. For different incomplete semantic factors $\{c_1, c_2, \dots, c_M\}$, we utilize a re-balancing method, such as Logit Adjustment (LA), to compute a set of data balanced outputs $\{\mathbb{P}(Y = y | \mathbf{b}, c_1), \mathbb{P}(Y = y | \mathbf{b}, c_2), \dots, \mathbb{P}(Y = y | \mathbf{b}, c_M)\}$. These outputs can then be merged using fusion weights, following Eq. 7, to obtain an unbiased estimation. Since iterating over all possible incomplete semantic factors is impractical, we approximate these factors using models like CLIP, OpenCLIP, and MetaCLIP, thereby addressing Eq. 7 and simplifying the process of balancing both parameter and data imbalances simultaneously.

5. Experiment

We conduct experiments on the ImageNet-LT, Places365-LT, and iNaturalist2018 datasets. Most of our experimental setup follows previous descriptions in Sec. 3.1. For a fair comparison, we only compare our method with LiVT (Xu et al., 2023b), LPT (Dong et al., 2022), VL-LTR (Tian et al., 2022), RAC (Long et al., 2022), LIFT (Shi et al., 2023), Decoder (Wang et al., 2024), GML (Suh & Seo, 2023) and BALLAD (Ma et al., 2021), which are trained based on the ViT (Dosovitskiy et al., 2020). If not specified, We use the Adaptformer to fine-tune the foundation model (We also report the results based on VPT in the Appendix). More details are in the Appendix Sec. A.

5.1. Results

Places365-LT. The results are presented in Tab. 5, where our method demonstrates superior overall performance compared to other approaches. Specifically, in comparison to

VL-LTR, which leverages additional data for representation learning, our method achieves an overall performance improvement of **2.91%**. It is worth noting that while VL-LTR achieves the highest performance in the “D-Many” group, this comes at the cost of reduced performance in the “D-Medium” and “D-Few” groups. In contrast, our method delivers consistently higher and more balanced performance across all groups, indicating its effectiveness in addressing long-tailed distributions.

ImageNet-LT. The results in Tab. 6 demonstrate that our method consistently improves performance across all categories. Specifically, our approach provides gains of **2.01%**, **2.65%**, and **3.49%** for the “D-Many”, “D-Medium”, and “D-Few” classes, respectively, compared to LIFT. Notably, the largest improvements are observed in the tail classes, underscoring the significance of our backdoor adjustment technique in addressing challenges faced by these underrepresented classes.

iNaturalist2018. The results presented in Tab. 7 demonstrate that our method surpasses several baseline approaches, achieving an overall performance improvement of **1.91%** compared to RAC, a method that leverages additional data for training. This performance gain is particularly noteworthy considering the scale of the iNaturalist2018 dataset, which comprises 8,142 distinct classes. The complexity of large-scale classification tasks often introduces significant challenges due to the vast number of classes and the inherent class imbalance. Despite these obstacles, our method proves to be highly effective, outperforming competitors even without the need for extra training data.

5.2. Ablation study

The influence of incomplete semantic factor The most critical hyperparameter in our backdoor adjustment method is the number of incomplete semantic factors, denoted as M . In our experiments, we set $M = 3$, utilizing the incomplete semantic factors from CLIP, OpenCLIP, and MetaCLIP to obtain the final prediction score. To further investigate the impact of M on performance, we conducted experiments shown in Tab. 8. As the number of incomplete semantic factors increases, our backdoor adjustment demonstrates improved performance. For instance, $M = 3$ outperforms $M = 1$, with OpenCLIP showing improvements of **2.27%**, **2.58%**, and **3.29%** on the “D-Many”, “D-Medium”, and “D-Few” on the ImageNet-LT dataset, respectively. Notably, the most significant improvement was observed in the “D-Few” classes, where the tail classes benefit greatly from a more balanced prediction due to the diverse range of incomplete semantic factors used. This implies that our method is effective at improving the performance on long-tailed distributions, where the tail classes usually suffer from inad-

Table 5. Results on Places365-LT.

Method	Backbone	Extra data	Params.	Epochs	D-Many	D-Medium	D-Few	All
LiVT (Xu et al., 2023b)	ViT-B/16	✗	85.80M	100	48.1	40.6	27.5	40.8
LPT (Dong et al., 2022)	ViT-B/16	✗	1.01M	80	49.3	52.3	46.9	50.1
VL-LTR (Tian et al., 2022)	ViT-B/16	✓	149.62M	100	54.2	48.5	42.0	50.1
RAC (Long et al., 2022)	ViT-B/16	✓	85.80M	30	48.7	48.3	41.8	47.2
BALLAD (Ma et al., 2021)	ViT-B/16	✓	149.62M	60	49.3	50.2	48.4	49.5
Decoder (Wang et al., 2024)	ViT-B/16	✗	21.26	34	-	-	-	46.8
LIFT (Shi et al., 2024)	ViT-B/16	✗	0.18M	10	51.3	52.2	50.5	51.5
Ours	ViT-B/16	✗	0.54M	10	52.52	53.62	52.54	53.01

Table 6. Results on ImageNet-LT

Method	Epochs	D-Many	D-Medium	D-Few	All
LiVT	100	73.6	56.4	41.0	60.9
VL-LTR	100	84.5	74.6	59.3	77.2
BALLAD	60	79.1	74.5	69.8	75.7
Decoder	18	-	-	-	73.2
GML	100	-	-	-	78.0
LIFT	10	80.2	76.1	71.5	77.0
Ours	10	82.21	78.75	74.99	79.57

Table 7. Results on iNaturalist2018

Method	Epochs	D-Many	D-Medium	D-Few	All
LiVT	100	78.9	76.5	74.8	76.1
LPT	160	-	-	79.3	76.1
VL-LTR	100	-	-	-	76.8
RAC	20	75.9	80.5	81.1	80.2
Decoder	5	-	-	-	59.2
LIFT	20	72.4	79.0	81.1	79.1
Ours	20	82.30	81.67	82.36	82.01

equate training samples and biased parameter distributions.

Backdoor adjustment relieves the parameter imbalance.

To verify that our method effectively alleviates parameter imbalance, we report the performance results in Fig. 6, using the division principles that align with the parameter imbalance of CLIP. When compared with zero-shot (ZS) and cross-entropy (CE) base-lines in Fig. 3, our backdoor adjustment achieves an overall improvement across all groups. For instance, in the "P-Few" group, our method provides performance gains of 1.32%, 1.42%, and

P-Few P-Medium P-Many	73.21	68.58	65.12
	51.67	51.84	48.87
	37.10	41.29	32.62
	D-Many	D-Medium	D-Few

Figure 6. The performance of different groups with our method on Places365-LT. After applying our method, samples belonging to both "P-Few" and "D-Few" exhibit significant performance gains, verifying the effectiveness of our method.

Table 8. The ablation study with the number of semantic factors M . Open and Meta denote the OpenCLIP and MetaCLIP, respectively.

	CLIP	Open	Meta	ImageNet-LT		
				D-Many	D-Medium	D-Few
$M = 1$	✓	✓		80.25	76.05	71.53
				79.94	76.17	71.70
			✓	80.45	77.06	72.64
$M = 2$	✓	✓		81.65	77.99	73.91
	✓		✓	81.86	78.26	74.33
		✓	✓	81.64	78.02	74.08
$M = 3$	✓	✓	✓	82.21	78.75	74.99

1.50% over Logit Adjustment (LA) in the "D-Many", "D-Medium", and "D-Few" groups, respectively. These results demonstrate that our method effectively enhances the performance of tail classes while maintaining or improving the performance of head classes, showcasing its ability to balance performance across the board without negative trade-offs.

6. Related work

Long-tailed learning. Previous methods primarily focus on re-balancing tail classes, employing techniques such as re-weighting (Cui et al., 2019) or re-sampling (Ren et al., 2020; Guo & Wang, 2021; Kim et al., 2020). The fundamental concept behind these methods is to place greater emphasis on tail classes to alleviate the effects of imbalanced bias. For instance, in (Cui et al., 2019), different classes are assigned a weight based on the effective number of samples in the final loss function. Similarly, Logit adjustment (Menon et al., 2020) re-balances tail classes by adjusting the output logits according to the prior for each class. Recently, foundation models, which are pre-trained on vast amounts of curated data, have demonstrated their utility in producing generalizable features transferable across various tasks. Several approaches have integrated foundation models to address long-tailed learning (Dong et al., 2022; Shi et al., 2024; Tian et al., 2022; Ma et al., 2021). However, these methods typically only consider the data distribution bias in downstream tasks, overlooking the inherent imbalance within foundation models themselves (Zhu et al., 2024).

In this work, we investigate the compound effects of both pre-trained and downstream imbalances, and we propose a simple approach to mitigate these effects.

Downstream fine-tuning. Fine-tuning techniques can be broadly categorized into full fine-tuning and Parameter-Efficient Fine-Tuning (PEFT) (Jia et al., 2022; Chen et al., 2022a; Zaken et al., 2021; Houlsby et al., 2019; Hu et al., 2021). PEFT refers to methods that adapt pre-trained models to specific tasks while minimizing the number of parameters that require updating. When data samples are limited, PEFT often outperforms full fine-tuning. Visual Prompt Tuning (VPT) (Jia et al., 2022) introduces two variants, VPT-Shallow and VPT-Deep, which insert prompts into different transformer layers. In contrast, AdaptFormer (Chen et al., 2022a) introduces lightweight modules that add only a small number of parameters, yet outperform fully fine-tuned models on various benchmarks. Despite their efficiency, PEFT methods focus on introducing fewer parameters while preserving the foundational model’s information. However, this also retains the model’s inherent limitations, such as imbalances. In this work, we provide an in-depth analysis of how the biases of foundation models impact downstream tasks and propose a novel method to solve the problem.

7. Conclusion

In this paper, we analyze how foundation model bias affects downstream imbalanced tasks, formally defining parameter and data imbalance to guide our study. While fine-tuning effectively addresses data imbalance, parameter imbalance persists, limiting performance. We find that re-balancing methods offer little benefit to feature representation. To tackle this, we construct a causal structure graph, identifying incomplete semantic factors from parameter imbalance as key confounders, and propose a backdoor adjustment method to mitigate their impact. Experimental results verify the effectiveness of our method. In future work, we extend our method to other tasks, such as object detection and segmentation under long-tailed distributions.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgement

This work is supported in part by the National Natural Science Foundation of China No. 62376277, and Doubao large model fund.

This work is supported by fund for building world-class universities (disciplines) of Renmin University of China.

This work is supported by the 2023 BDA project (MIIT).

References

- Ahuja, K., Shanmugam, K., Varshney, K., and Dhurandhar, A. Invariant risk minimization games. In *International Conference on Machine Learning*, pp. 145–155. PMLR, 2020.
- Ahuja, K., Caballero, E., Zhang, D., Gagnon-Audet, J.-C., Bengio, Y., Mitliagkas, I., and Rish, I. Invariance principle meets information bottleneck for out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 34:3438–3450, 2021.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Chen, S., Ge, C., Tong, Z., Wang, J., Song, Y., Wang, J., and Luo, P. Adaptformer: Adapting vision transformers for scalable visual recognition. *Advances in Neural Information Processing Systems*, 35:16664–16678, 2022a.
- Chen, Y., Zhou, K., Bian, Y., Xie, B., Wu, B., Zhang, Y., Ma, K., Yang, H., Zhao, P., Han, B., et al. Pareto invariant risk minimization: Towards mitigating the optimization dilemma in out-of-distribution generalization. *arXiv preprint arXiv:2206.07766*, 2022b.
- Cherti, M., Beaumont, R., Wightman, R., Wortsman, M., Ilharco, G., Gordon, C., Schuhmann, C., Schmidt, L., and Jitsev, J. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2818–2829, 2023.
- Cover, T. and Hart, P. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.
- Cui, Y., Jia, M., Lin, T.-Y., Song, Y., and Belongie, S. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9268–9277, 2019.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Deng, Y., Yang, Y., Mirzasoleiman, B., and Gu, Q. Robust learning with progressive data expansion against spurious correlation. *Advances in neural information processing systems*, 36:1390–1402, 2023.

- Dong, B., Zhou, P., Yan, S., and Zuo, W. Lpt: Long-tailed prompt tuning for image classification. *arXiv preprint arXiv:2210.01033*, 2022.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Guo, H. and Wang, S. Long-tailed multi-label visual recognition by collaborative training on uniform and re-balanced samplings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15089–15098, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Houlsby, N., Giurghi, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pp. 2790–2799. PMLR, 2019.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Jia, M., Tang, L., Chen, B.-C., Cardie, C., Belongie, S., Hariharan, B., and Lim, S.-N. Visual prompt tuning. In *European Conference on Computer Vision*, pp. 709–727. Springer, 2022.
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., and Kalantidis, Y. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019.
- Kim, C. D., Jeong, J., and Kim, G. Imbalanced continual learning with partitioning reservoir sampling. In *European Conference on Computer Vision*, pp. 411–428. Springer, 2020.
- Krueger, D., Caballero, E., Jacobsen, J.-H., Zhang, A., Binias, J., Zhang, D., Le Priol, R., and Courville, A. Out-of-distribution generalization via risk extrapolation (rex). In *International conference on machine learning*, pp. 5815–5826. PMLR, 2021.
- Lin, Y., Zhu, S., Tan, L., and Cui, P. Zin: When and how to learn invariance without environment partition? *Advances in Neural Information Processing Systems*, 35: 24529–24542, 2022.
- Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., and Yu, S. X. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2537–2546, 2019.
- Long, A., Yin, W., Ajanthan, T., Nguyen, V., Purkait, P., Garg, R., Blair, A., Shen, C., and van den Hengel, A. Retrieval augmented classification for long-tail visual recognition. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6949–6959, 2022. doi: 10.1109/CVPR52688.2022.00683.
- Ma, T., Geng, S., Wang, M., Shao, J., Lu, J., Li, H., Gao, P., and Qiao, Y. A simple long-tailed recognition baseline via vision-language model. *arXiv preprint arXiv:2111.14745*, 2021.
- Menon, A. K., Jayasumana, S., Rawat, A. S., Jain, H., Veit, A., and Kumar, S. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*, 2020.
- Pearl, J. *Causality*. Cambridge university press, 2009.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Ren, J., Yu, C., Ma, X., Zhao, H., Yi, S., et al. Balanced meta-softmax for long-tailed visual recognition. *Advances in Neural Information Processing Systems*, 33: 4175–4186, 2020.
- Robey, A., Pappas, G. J., and Hassani, H. Model-based domain generalization. *Advances in Neural Information Processing Systems*, 34:20210–20229, 2021.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Rosenfeld, E., Ravikumar, P., and Risteski, A. The risks of invariant risk minimization. *arXiv preprint arXiv:2010.05761*, 2020.
- Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., and Komatsuzaki, A. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626, 2017.

- Shi, J.-X., Wei, T., Zhou, Z., Han, X.-Y., Shao, J.-J., and Li, Y.-F. Parameter-efficient long-tailed recognition. *arXiv preprint arXiv:2309.10019*, 2023.
- Shi, J.-X., Wei, T., Zhou, Z., Shao, J.-J., Han, X.-Y., and Li, Y.-F. Long-tail learning with foundation model: Heavy fine-tuning hurts. In *Forty-first International Conference on Machine Learning*, 2024.
- Suh, M.-K. and Seo, S.-W. Long-tailed recognition by mutual information maximization between latent features and ground-truth labels. In *International Conference on Machine Learning*, pp. 32770–32782. PMLR, 2023.
- Tang, K., Huang, J., and Zhang, H. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *Advances in neural information processing systems*, 33:1513–1524, 2020.
- Tian, C., Wang, W., Zhu, X., Dai, J., and Qiao, Y. Vl-ltr: Learning class-wise visual-linguistic representation for long-tailed visual recognition. In *European Conference on Computer Vision*, pp. 73–91. Springer, 2022.
- Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., and Belongie, S. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8769–8778, 2018.
- Wang, Y., Yu, Z., Wang, J., Heng, Q., Chen, H., Ye, W., Xie, R., Xie, X., and Zhang, S. Exploring vision-language models for imbalanced learning. *International Journal of Computer Vision*, 132(1):224–237, 2024.
- Wang, Z., Zhang, Z., Ebrahimi, S., Sun, R., Zhang, H., Lee, C.-Y., Ren, X., Su, G., Perot, V., Dy, J., et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *European Conference on Computer Vision*, pp. 631–648. Springer, 2022a.
- Wang, Z., Zhang, Z., Lee, C.-Y., Zhang, H., Sun, R., Ren, X., Su, G., Perot, V., Dy, J., and Pfister, T. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 139–149, 2022b.
- Wen, X., Zhao, B., Chen, Y., Pang, J., and Qi, X. Generalization beyond data imbalance: A controlled study on clip for transferable insights. *arXiv preprint arXiv:2405.21070*, 2024.
- Xu, H., Xie, S., Tan, X. E., Huang, P.-Y., Howes, R., Sharma, V., Li, S.-W., Ghosh, G., Zettlemoyer, L., and Feichtenhofer, C. Demystifying clip data. *arXiv preprint arXiv:2309.16671*, 2023a.
- Xu, Z., Liu, R., Yang, S., Chai, Z., and Yuan, C. Learning imbalanced data with vision transformers. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023b.
- Zaken, E. B., Ravfogel, S., and Goldberg, Y. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*, 2021.
- Zhu, B., Tang, K., Sun, Q., and Zhang, H. Generalized logit adjustment: Calibrating fine-tuned models by removing label bias in foundation models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Zhu, J., Wang, Z., Chen, J., Chen, Y.-P. P., and Jiang, Y.-G. Balanced contrastive learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6908–6917, 2022.

A. Dataset and experimental setup

A.1. Dataset introduction

ImageNet-LT ImageNet-LT (Liu et al., 2019) is a subset of the ImageNet (Deng et al., 2009) dataset, designed to address the challenges of long-tailed distributions. It contains a total of 118.5K images, with the number of samples per class varying significantly—from 1280 samples for the most represented class to just 5 for the least represented class. The distribution of samples per class is determined by a down-sampled Pareto distribution, emphasizing the disparity in class representation.

Places365-LT Places365-LT (Liu et al., 2019) is derived from the Places-2 dataset and consists of 62.5K images spread across 365 categories. This dataset exhibits a more pronounced imbalance compared to ImageNet-LT. In Places365-LT, the largest class contains 4980 images, while the smallest class is represented by only 5 images, resulting in an imbalance factor (IF) of 996.

iNaturalist2018 The iNaturalist2018 dataset (Van Horn et al., 2018) is focused on natural biological images and presents a significant challenge due to its heavily imbalanced distribution. It encompasses 437.5K images across 8142 classes, making it one of the largest datasets in terms of class diversity. Additionally, iNaturalist2018 poses a fine-grained classification challenge, as it requires distinguishing between similar species and categories, adding another layer of complexity to the classification task.

A.2. Experimental setup

We present the details about the hyper-parameters of our experiments on different datasets in Tab. 9, where Lr, epochs denote the initial learning rate and training epochs, respectively. We denote batch_size in Tab. 9 as the training batch size during the fine-tuning phase.

For training resources, all experiments are conducted on Intel(R) Xeon(R) Gold 5318Y CPU @ 2.10GHz with a single RTX A40 GPU. Normally, a GPU with 24GB of memory is sufficient for the reproduction.

Table 9. Hyper-parameters used in our experiments on different datasets.

Dataset	Lr	Epochs	Batch_size
Places365-LT	0.01	10	128
ImageNet-LT	0.01	10	128
iNaturalist2018	0.01	20	128

A.3. Parameter imbalance and Data imbalance

Previous methods have effectively addressed data imbalance by providing a clear definition of the issue. Following the approach of OLTR (Liu et al., 2019), we categorize the classes into three groups: “D-Many”, “D-Medium”, and “D-Few”. Specifically, classes with more than 100 training samples are classified as “D-Many”, those with 20–100 samples are categorized as “D-Medium”, and classes with fewer than 20 samples fall into the “D-Few” group. This categorization allows for a systematic examination of the impact of data imbalance across varying levels of sample availability.

To address parameter imbalance, where access to pre-training data is unavailable, we adopt the estimation method provided by GLA (Zhu et al., 2024) (Eq. 3) to estimate the label prior. Using this estimation, we divide the classes uniformly into three groups: the top 30% of classes based on the label prior are designated as “P-Many”, the next 30–60% as “P-Medium”, and the remaining 40% as “P-Few”. This approach provides a structured way to assess parameter imbalance across different subsets of classes.

By applying these criteria, any given dataset can be simultaneously split based on both parameter imbalance and data imbalance. Importantly, a single class can belong to multiple groups across the two categorizations; for instance, a class may be grouped as “P-Many” while also being classified as “D-Few”. This dual categorization facilitates a more nuanced analysis of the interplay between parameter and data imbalance.

B. Detailed discussions on incomplete semantic factors

We first provide the definition of confounders in the causal framework: When a third variable Z influences both X and Y , we say that X and Y are confounded. Such a variable Z is referred to as a confounder of X and Y (Pearl, 2009). Specifically, it satisfies the fork structure $X \leftarrow Z \rightarrow Y$. This latent common cause Z can create a spurious correlation between X and Y , making the observed statistical relationship $\mathbb{P}(Y|X)$ potentially misleading. To address this, the causal relationship $\mathbb{P}(Y|do(X))$ is used as a replacement for the correlation $\mathbb{P}(Y|X)$.

Secondly, we explain why the incomplete semantic factor C in our problem setting can create a backdoor path $X \leftarrow C \rightarrow Y$, leading to a spurious correlation. The example in the first row in Fig.4 explains why the parameter-imbalanced model D leads to an incomplete semantic factor C . For example, the semantic factor available to OpenCLIP is in the head of the object, while the semantic factor available to MetaCLIP is in the body of the object.

We emphasize that the semantic factors obtained due to model imbalance are incomplete. Therefore, for rigor, we have modified the causal diagram in Fig. 5 to include the inaccessible semantic factor U . From the perspective of data generation, C and U together constitute X , thus $X \leftarrow C$ holds. On the other hand, the path $C \rightarrow Y$ indicates that the predicted label distributions follow their own parameter imbalance. A parameter-imbalanced model, OpenCLIP, consistently predicts a “head of the object” label distribution (i.e., making more accurate predictions on test samples where the head of the object is not occluded), regardless of the test distributions. Similarly, the parameter-imbalanced model MetaCLIP exhibits a distinct “body of the object” preference across different test distributions.

The relationship between the confounding path $X \leftarrow C \rightarrow Y$ is unstable. When the training set model predicts the label with the help of the body semantics of the object, if the body of a sample of a certain class in the test example is occluded, the model will give an incorrect prediction. From this example, it can be seen that an ideal model needs to learn the causal relationship $\mathbb{P}(Y|do(X))$ between the input and the label, rather than fitting the unstable confounding path $X \leftarrow C \rightarrow Y$, in order to achieve generalization between different distributions.

C. Discussion of other causal-based methods in long-tailed learning

Research has explored the application of causal theory to long-tailed learning (Tang et al., 2020; Zhu et al., 2022). In this section, we briefly compare our method with previous approaches.

In (Tang et al., 2020), the authors analyze SGD momentum as a confounder and use backdoor adjustment to mitigate its influence. In contrast, we identify incomplete semantic factors as confounders and disentangle their impact. The key differences are as follows:

- **Motivation:** Our method investigates how foundation model biases influence imbalanced downstream tasks. In (Tang et al., 2020), the focus is on how SGD momentum affects long-tailed learning.
- **Confounder:** We treat incomplete semantic factors as the confounder, whereas (Tang et al., 2020) identifies momentum as the confounder. This distinction leads to different causal graphs.
- **Method:** We compute $P(Y|do(X))$ by leveraging diverse scenarios and methods. In (Tang et al., 2020), the imbalance is viewed as causing momentum bias, and $P(Y|x, m)$ is computed through de-confounded training to address this. In contrast, our approach focuses on incomplete semantic factors, computing $P(Y|b, c)$ by utilizing different foundation models to capture varied semantic information, ensuring the model’s inference is not constrained by incomplete semantics.

In (Zhu et al., 2022), the authors propose xERM, which aims to ensure that imbalanced learning methods perform well on both imbalanced and balanced test distributions. They consider the selected domain as the confounder, while we focus on incomplete semantic factors. The differences are as follows:

- **Motivation:** Our method examines how foundation model biases affect imbalanced downstream tasks. In (Zhu et al., 2022), the authors analyze the influence of different domains under imbalanced distributions.
- **Confounder:** We treat incomplete semantic factors as the confounder, while (Zhu et al., 2022) identifies the selected domain as the confounder. This leads to distinct causal graphs.

- **Method:** We compute $P(Y|do(X))$ using diverse approaches. In (Zhu et al., 2022), the selected domain is addressed by computing $P(Y|x, s)$ through de-confounded training. In contrast, we focus on incomplete semantic factors, computing $P(Y|b, c)$ by leveraging different foundation models to capture richer semantic information, ensuring robust inference unaffected by incomplete semantics.

D. Discussion with Invariant Risk Minimization (IRM)

Invariant Risk Minimization (IRM) aims to enhance out-of-distribution generalization by identifying and optimizing invariant features, thereby reducing bias in deep learning models (Arjovsky et al., 2019; Lin et al., 2022; Deng et al., 2023). In linear systems, IRM has strong theoretical guarantees and a clear connection to causal theory. Building upon IRM, numerous variants have emerged (Ahuja et al., 2020; Krueger et al., 2021; Robey et al., 2021; Ahuja et al., 2021), aiming to address some of IRM’s challenges, such as its failure in nonlinear tasks (Rosenfeld et al., 2020), the requirement for extensive domain information (Lin et al., 2022), and optimization difficulties in deep neural networks (Chen et al., 2022b).

While IRM aims to identify and optimize invariant features across domains to reduce bias, our method leverages a causal learning framework that identifies and mitigates biases introduced by spurious correlations, which are induced by the backdoor path. Rather than assuming the existence of invariant features, we treat the incomplete semantic factor as a confounder and apply a backdoor adjustment method to learn the true causal effect, offering a more flexible solution to both parameter and data imbalance.

E. Discussion on parameter imbalance of LLM

We also believe that current LLMs suffer from parameter imbalance. We analyze this issue from two perspectives, the influence of parameter imbalance without fine-tuning and with fine-tuning under downstream data:

Without Fine-Tuning Since LLMs are trained on diverse corpora, some domains (e.g., news articles, Wikipedia, and general web content) dominate the pre-training process. As a result, the model parameters may encode more precise representations of frequent patterns while underrepresenting rare or specialized knowledge. This imbalance becomes apparent when applying LLMs to highly domain-specific tasks, such as specialized scientific research or low-resource languages, where the model struggles due to insufficient parameter representation.

Fine-tuning When we use LLMs for customized tasks, such as role-playing or simulating a specific individual’s tone and manner of speaking, parameter imbalance can also have a significant impact. Since pre-training data is often biased toward general linguistic patterns found in publicly available texts, the model may struggle to accurately capture and reproduce highly personalized or domain-specific speech styles. For example, if an LLM is fine-tuned to simulate the conversational style of a historical figure or a niche expert, but the pre-training corpus contains limited examples of their actual speech patterns, the model may default to generic language structures instead of faithfully mimicking the target style. Even with sufficient fine-tuning data, the model’s prior knowledge—encoded in the pre-trained parameters—can cause it to retain biases from its original training distribution, making it difficult to fully adapt to the new task.

F. Overall inference pipeline

During the fine-tuning phase, our method can be decomposed into two stages. In the first stage, we apply logit adjustment to fine-tune various foundation models (CLIP, OpenCLIP, and MetaCLIP), addressing the data imbalance present in the downstream dataset. In the second stage, we input each test sample into the fine-tuned models, obtaining output logits without additional fine-tuning. As described in Eq. 7, we then ensemble these logits using the importance weight $P(c)$ to correct for parameter imbalance, ultimately producing the final prediction score.

G. The selection of foundation model

As shown in Fig. 5, the path $D \rightarrow C$ represents the incomplete semantic factor that arises due to parameter imbalance, stemming from the imbalance in the pre-training data of the foundation model. Consequently, we have chosen CLIP, OpenCLIP, and MetaCLIP because they are pre-trained on distinct datasets, resulting in varying degrees of parameter imbalance and differing incomplete semantic factors.

In causal theory, Eq. 7 should cover the value space of C as fully as possible for accurate causal effect estimation. Since C has infinite values, we approximate it finitely, with broader coverage improving estimation accuracy. In this way, model selection should take into account the path $D \rightarrow C$. If two foundation models are trained on similar pre-training datasets, only one should be selected, as choosing both would not significantly increase the coverage of the value space of C .

H. Introduction to the ViT and PEFT

H.1. Vision Transformer

The Vision Transformer (ViT) is a deep learning model designed to process image data by leveraging the architecture of transformers, originally developed for natural language processing tasks. Unlike traditional convolutional neural networks (CNNs), which rely on convolutional layers to capture spatial relationships in images, ViT treats an image as a sequence of patches. Each image is divided into fixed-size patches, and these patches are embedded and processed as tokens, similar to words in a sentence. Through self-attention mechanisms, ViT captures global dependencies between patches, allowing it to effectively model long-range relationships within images. ViT has demonstrated strong performance in a variety of computer vision tasks, often surpassing CNNs, particularly as the availability of large-scale image datasets has grown.

For a pre-trained Vision Transformer (ViT) composed of N blocks, denoted as $B = \{B_1, B_2, \dots, B_N\}$, each input image x_i is divided into m patches, resulting in patch embeddings $\mathbf{E}_i^{(0)} = \{e_{i,1}^{(0)}, e_{i,2}^{(0)}, \dots, e_{i,m}^{(0)}\}$. Together with the classification token (CLS), these patch embeddings are passed through the ViT backbone, producing an output embedding $\{\text{CLS}_i^{(N)}, e_{i,1}^{(N)}, e_{i,2}^{(N)}, \dots, e_{i,m}^{(N)}\}$.

Each block consists of a multi-head self-attention mechanism followed by a feed-forward layer, both of which incorporate layer normalization and a residual connection. In the self-attention layer, the patch embeddings and the class token are updated based on the similarity matrix, as shown in Eq. 8. Here, $\mathbf{Q}_i^{(n-1)}$, $\mathbf{K}_i^{(n-1)}$, and $\mathbf{V}_i^{(n-1)}$ are derived from $[\text{CLS}_i^{(n-1)}, \mathbf{E}_i^{(n-1)}]$ and are input to different linear transformations within block B_n , where d represents the feature dimension. The self-attention layer updates the tokens based on the values in the self-attention matrix.

$$[\widetilde{\text{CLS}}_i^{(n-1)}, \widetilde{\mathbf{E}}_i^{(n-1)}] = \text{softmax}\left(\frac{(\mathbf{Q}_i^{(n-1)})^T \mathbf{K}_i^{(n-1)}}{\sqrt{d}}\right) \mathbf{V}_i^{(n-1)} \quad (8)$$

where d is the feature dimension. The resulting tokens $[\widetilde{\text{CLS}}_i^{(n-1)}, \widetilde{\mathbf{E}}_i^{(n-1)}]$ from the multi-head self-attention layer are then passed into the feed-forward layer, producing the output $[\text{CLS}_i^{(n)}, \mathbf{E}_i^{(n)}]$ of block B_n .

H.2. Adaptformer

AdaptFormer is a lightweight adaptation framework designed to enhance the efficiency and flexibility of large pre-trained vision transformers (ViTs) in downstream tasks. By introducing a low-rank adaptation (LoRA) module within the feed-forward layers of the transformer, AdaptFormer allows for the modification of the model's capabilities without requiring a full re-training of all parameters. This approach preserves the pre-trained knowledge in the original transformer while enabling it to adapt effectively to new tasks with minimal additional parameters. As a result, AdaptFormer achieves strong performance in transfer learning tasks, providing a balance between computational efficiency and task-specific adaptability.

Let \mathbf{X} be the input to the feed-forward layer, and the original weight matrix of the layer be $\mathbf{W} \in \mathbb{R}^{d \times d}$. In AdaptFormer, an additional low-rank module is added as a modification. This is represented by two matrices, $\mathbf{A} \in \mathbb{R}^{d \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times d}$, where $r \ll d$ represents the low-rank dimension. The output of the modified feed-forward layer is given by:

$$\mathbf{Y} = \mathbf{W}\mathbf{X} + \alpha(\mathbf{A}\mathbf{B})\mathbf{X} \quad (9)$$

Here, α is a scaling factor that controls the influence of the low-rank adaptation. The term $\mathbf{A}\mathbf{B}$ represents the low-rank adaptation added to the original weight transformation, allowing the model to adapt with a minimal number of additional parameters.

H.3. Visual Prompt Tuning

Visual Prompt Tuning is an innovative approach designed to adapt pre-trained vision models for specific tasks by utilizing learnable visual prompts. Instead of fine-tuning all the parameters of a large vision model, this method introduces a small set

Table 10. The ablation study with the different number of semantic factors of M for VPT.

	CLIP	OpenCLIP	MetaCLIP	ImageNet-LT			Places365-LT		
				D-Many	D-Medium	D-Few	D-Many	D-Medium	D-Few
$M = 1$	✓	✓	✓	78.63	75.42	71.25	50.70	51.72	49.07
				78.62	75.40	70.91	51.45	51.53	50.68
				79.86	76.47	72.19	50.91	52.06	49.86
$M = 2$	✓	✓	✓	80.42	77.17	73.19	51.80	52.66	51.33
				80.75	77.96	73.54	51.55	53.20	50.82
				80.54	77.52	73.41	51.89	53.09	51.22
$M = 3$	✓	✓	✓	81.09	78.37	74.16	51.97	53.26	51.34

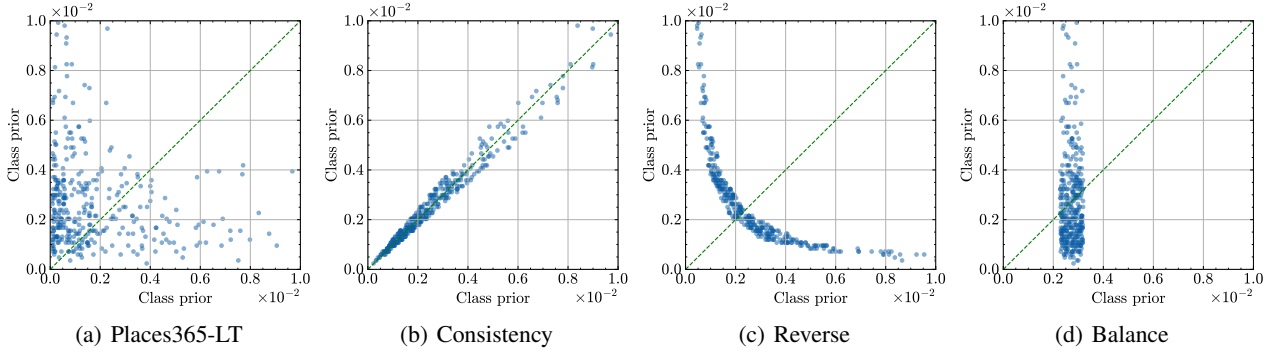


Figure 7. We constitute different parameter imbalanced datasets (a) the original Places365-LT, (b) the data imbalance and parameter imbalance are consistent, (c) the data imbalance and parameter imbalance are reversed, and (d) the downstream dataset is balanced. Each dot in the figure represent the pair $(\mathbb{P}_S(Y = y), \hat{\mathbb{P}}_P(Y = y))$.

of visual prompts that can be learned during the adaptation process. These prompts are essentially additional visual tokens that are concatenated to the input image embeddings, guiding the model to focus on relevant features for the target task. This approach significantly reduces the computational burden associated with full model fine-tuning while retaining the rich representations learned from extensive pre-training. By leveraging visual prompts, practitioners can achieve competitive performance in various computer vision applications, such as image classification and object detection, with minimal adjustments to the model architecture.

The process of Visual Prompt Tuning can be represented mathematically by introducing a set of learnable visual prompts into the embedding space of a pre-trained vision model. Let $\mathbf{P} = \{\mathbf{P}^{(1)}, \dots, \mathbf{P}^{(N)}\}$, where $\mathbf{P}^{(n)} \in \mathbb{R}^{l \times d}$ is the prompts of the n -th block, where l is the number of prompts. The visual prompt tuning process can be described by the following equation:

$$\text{CLS}_i^{(n)}, \mathbf{E}_i^{(n)} = B_n(\text{CLS}_i^{(n-1)}, \mathbf{P}^{(n-1)}, \mathbf{E}_i^{(n-1)}) \quad (10)$$

This concatenated representation is then fed into the transformer architecture of the vision model to adapt it for specific downstream tasks. The learnable visual prompts \mathbf{P} are updated during the tuning process to optimize task performance while keeping the original model parameters fixed.

I. Additional experiments

I.1. The performance for visual prompt tuning

We also conduct experiments using different PEFT-based methods, such as Visual Prompt Tuning (VPT), with the results shown in Tab. 10. A similar phenomenon is observed with Adaptorformer for fine-tuning the foundation model. For instance, when $M = 3$ outperforms $M = 1$, OpenCLIP demonstrates improvements of **2.38%**, **2.97%**, and **3.25%** on the “D-Many”, “D-Medium”, and “D-Few” groups of the ImageNet dataset, respectively. These experiments further indicate that our method generalizes well to different types of PEFT-based methods.

Table 11. Ther performance of GLA, GLA-ZS, GLA-FT, and GLA-Train based on the CLIP model.

	Type	D-Many	D-Medium	D-Few	P-Many	P-Medium	P-Few	Overall
ImageNet-LT	GLA	79.76	76.48	74.12	88.34	78.58	65.30	77.42
	GLA-ZS	70.23	68.61	68.90	83.63	70.36	53.79	69.27
	GLA-FT	79.67	76.15	73.29	88.07	77.92	65.32	77.12
	GLA-Train	80.21	75.93	72.01	88.05	77.93	65.32	77.10
Places365-LT	GLA	51.22	52.00	53.34	69.07	50.70	35.87	51.98
	GLA-ZS	40.58	39.57	46.63	61.33	39.69	22.56	41.31
	GLA-FT	50.41	52.23	52.18	67.71	49.57	37.15	51.56
	GLA-Train	51.36	52.30	50.91	67.70	49.58	37.17	51.48

Table 12. The results of decoupling training and LA on ImageNet-LT dataset.

	D-Many	D-Medium	D-Few	All
LA	80.25	76.05	71.53	77.06
Decoupling	80.30	75.90	70.4	76.80

I.2. More experimental results for GLA

In Tab.3, we present the results of GLA, GLA-ZS, GLA-FT, and GLA-Train on ImageNet-LT and Places365-LT, focusing on performance influenced by data imbalance due to space constraints. Additionally, in Tab.11, we provide results highlighting the performance under parameter imbalance.

I.3. The influence of classifiers

Decoupling training (Kang et al., 2019) has demonstrated its effectiveness under long-tailed distributions, suggesting that a generalizable representation can be learned without re-balancing. However, this conclusion is largely based on experiments using CNN architectures trained from scratch, and its influence on fine-tuning-based methods remains unclear. To explore this, we conducted experiments following a similar pipeline, where the backbone is fine-tuned using Cross Entropy (CE) while the classifier is trained with Logit Adjustment (LA). As shown in Tab. 12, decoupling training achieves a performance similar to LA. Combined with previous findings, we observe that LA primarily alleviates data imbalance by driving a more balanced classifier, though it offers limited improvement in the representation of tail classes.

I.4. The influence of data imbalance alone

In our paper, we address both data imbalance and parameter imbalance simultaneously, with analyses primarily based on commonly used datasets. To better understand the impact of parameter imbalance, we consider the following scenarios: (1) Consistency: The data imbalance and parameter imbalance are consistent, where the head classes grouped by the data imbalance are also the head classes grouped by the parameter imbalance. (2) Reverse: The data imbalance and parameter imbalance are reversed, where the head classes grouped by the data imbalance are the tail classes grouped by the parameter imbalance. (3)Balance: The downstream data is balanced. To simulate these scenarios, we construct additional imbalanced datasets by downsampling the Places365 dataset, as illustrated in Fig. 7. We evaluate performance using the Places365-LT balanced test set. To reduce the influence of randomness, we create three datasets for each scenario and report the average performance across them.

The results, shown in Tab. 13, indicate that when the downstream data is balanced, the model achieves the best performance, as this setup aligns most closely with the distribution of the downstream task. Interestingly, the Reverse scenario outperforms

Table 13. The results of Uniform, Resverse, and Balance.

Balance	Uniform	Reverse
53.63	50.80	51.2

Table 14. The imbalanced factor of parameter imbalance of different foundation models on Places365-LT.

	CLIP	OpenCLIP	MetaCLIP
IF	57.50	63.25	60.20

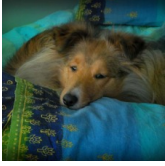
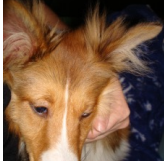
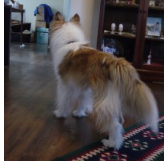
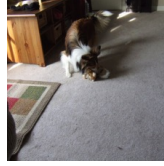
	C=0		C=1	
				
OpenCLIP	0.7177	0.6689	0.3032	0.2932
MetaCLIP	0.2634	0.2432	0.7180	0.7090

Figure 8. Prediction scores of fine-tuned OpenCLIP and MetaCLIP. OpenCLIP is more confident about the sample if the head is exposed ($C = 0$), while MetaCLIP is more sensitive to the body ($C = 1$).

the Consistency scenario. In the Reverse setting, the tail classes (in terms of parameter imbalance) can be viewed as head classes (in terms of data imbalance), suggesting that parameter imbalance can be mitigated by increasing the number of training samples for tail classes. However, in real-world applications, collecting sufficient data for tail classes is often challenging, and simply adding more data is not a practical solution to this issue.

I.5. Discussion

Our method tackles the parameter imbalance in PEFT under long-tailed distributions, providing a balanced solution that improves both model performance and fairness. As shown in Tab. 15, increasing M allows for the identification of more incomplete semantic factors, enhancing the model’s ability to address these confounding elements. However, this improvement comes with substantially higher computational costs, underscoring the importance of optimizing M to balance performance gains and efficiency, particularly in resource-constrained scenarios.

I.6. Interpreting the confounder

In this section, we conduct experiments to investigate why the incomplete semantic factor acts as a confounder in the causal graph. As illustrated in the first row of Fig. 4, the fine-tuned OpenCLIP model predominantly focuses on the head, whereas MetaCLIP primarily focuses on the body. We define $C = 0$ to represent the head and $C = 1$ to represent the body for this particular class.

To assess the influence of the incomplete semantic factor, we randomly select samples from this class and examine the prediction results of OpenCLIP and MetaCLIP. As shown in Fig. 8, when the dog’s head is occluded, OpenCLIP tends to produce lower-confidence predictions. Conversely, when the body is occluded and the head remains visible, OpenCLIP is more likely to provide higher-confidence predictions. Since both models are fine-tuned using the same principle, the observed differences in predictions can be attributed to incomplete semantic factors, thereby supporting our hypothesis.

Table 15. Computation costs.

	FLOPs
$M = 1$	239.06
$M = 2$	478.13
$M = 3$	717.19

Table 16. Differences between $\mathbb{P}(Y | X)$ and $\mathbb{P}(Y | do(X))$. Image1, Image2, Image3, and Image4 denote the image in Fig. 8 from left to right, respectively.

	Image1	Image2	Image3	Image4
$\mathbb{P}(Y do(X))$	0.5614	0.5443	0.5837	0.5742
$\mathbb{P}(Y X)$	0.5431	0.2221	0.2210	0.1031




Figure 9. The comparison between fine-tuned model and training from scratch.

To show the existence of confounding bias directly, we measure differences between $\mathbb{P}(Y | X)$ and $\mathbb{P}(Y | do(X))$. As shown in Tab. 16, there is a significant difference between $\mathbb{P}(Y | X)$ and $\mathbb{P}(Y | do(X))$, which also verifies the existence of confounding bias.

I.7. More experimentsParameter imbalance

In addition, to better support our point, we train a model from scratch on ImageNet-LT with ResNet-50. Since the parameter of ResNet-50 (He et al., 2016) is randomly initialized, the trained model is only influenced by the data imbalance. After getting the model, we calculate the accuracy of each class and sort them relying on the estimated pre-training data label prior $\mathbb{P}_P(Y)$. Since the parameter imbalance does not influence the trained model, the performance curve after sorting should not present a downward trend. As shown in Fig. 8, the result verifies our point. Therefore, in this comparison experiment, the downward trend in the right image in Fig. 2 is attributed to the imbalance of pre-training data, which can also verify the impact of parameter imbalance.

I.8. Visualization of our backdoor adjustment

As illustrated in Fig. 10, we present heatmaps generated after applying our backdoor adjustment, showcasing the effectiveness of our method in mitigating the influence of incomplete semantic factors. By addressing these confounders, our method enhances focus on the entire object rather than isolated parts.

In the first row, OpenCLIP primarily concentrates on the head, while MetaCLIP shifts its attention to the body. In contrast, our method successfully integrates both the head and body, yielding a more holistic understanding of the object. Similarly, in the second row, our method expands its focus to include the head, wings, and tail, offering a more comprehensive representation of the object. Finally, in the last row, our method effectively captures the entire structure of the truck, demonstrating its robustness and adaptability to diverse objects and scenarios.

I.9. Experimental evidence for $C \rightarrow Y$

In this section, we will verify the existence of $C \rightarrow Y$ in experiments. We constructed a dataset consisting of six classes: “airplane”, “automobile”, “bird”, “cat”, “dog”, and “truck”. For each class, we define distinct incomplete semantic factors, as shown in Tab. 17, where C denotes the incomplete semantic factor:

For example, the term “airplane head” refers to the front part of the aircraft, including the cockpit, while “airplane tail” represents the rear section, including the winglets. To construct the dataset, we used Stable Diffusion (Rombach et al., 2022) by providing prompts in the format: “a photo of a {class name}’s {head or tail}”. For instance, to generate samples for the airplane class, we used the prompts “a photo of an airplane’s head” and “a photo of an airplane’s tail”, respectively, ensuring

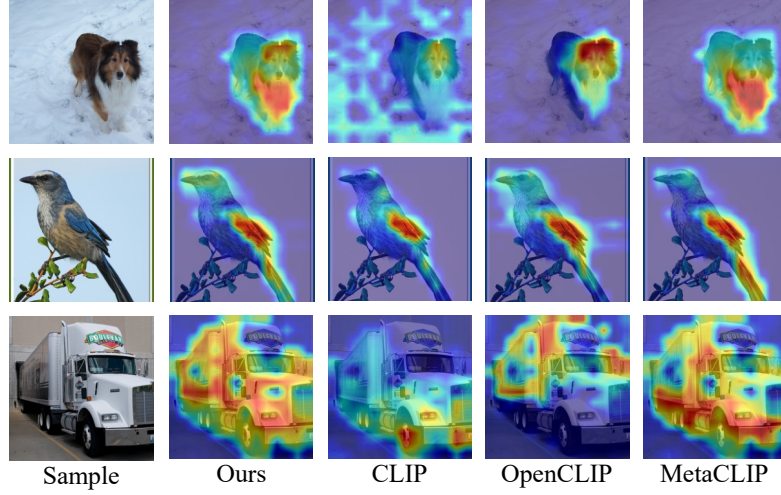


Figure 10. The comparison of our method with other fine-tuning based methods.

Table 17. The six classes and their corresponding incomplete semantic factors.

airplane	airplane head	airplane tail
automobile	automobile head	automobile tail
bird	bird head	bird body
cat	cat head	cat body
dog	dog head	dog body
truck	truck head	truck tail

sufficient samples for each semantic factor. We also visualize some samples in Fig. 11.

The training dataset is created under the following conditions: blue1) For each incomplete semantic factor, 50 samples are generated. This dataset is referred to as A_{train} . 2) For $C = 0$, 90 samples are generated, and for $C = 1$, 10 samples are generated. This dataset is referred to as B_{train} .

For evaluation, the test dataset is constructed under the following conditions: 1) For each incomplete semantic factor, 50 samples are generated. This dataset is referred to as A_{test} . 2) For $C = 0$, 90 samples are generated, and for $C = 1$, 10 samples are generated. This dataset is referred to as B_{test} .

We use ResNet-32 (He et al., 2016) as the backbone and train it from scratch on A_{train} and B_{train} , respectively. The trained models are then evaluated on A_{test} and B_{test} , respectively. As shown in Fig. 12, when the model is trained on B_{train} but evaluated on A_{test} , we find the performance drops (compared with testing on B_{test}), which indicates the confounder influences the model and cannot make right predictions. Our findings can be summarized as follows:




Figure 11. The visualization of the “dog” class: $C = 0$ represents samples generated for the dog’s head, while $C = 1$ corresponds to samples of the dog’s body.




Figure 12. The performance curve illustrates four evaluation scenarios: A_{train}, A_{test} indicates the model is trained on A_{train} and evaluated on A_{test} ; A_{train}, B_{test} represents the model trained on A_{train} and evaluated on B_{test} ; B_{train}, A_{test} corresponds to the model trained on B_{train} and evaluated on A_{test} ; and B_{train}, B_{test} denotes the model trained on B_{train} and evaluated on B_{test} .

Table 18. The experiments of randomness.

	Places365-LT	ImageNet-LT
LIFT	51.4 (0.09)	77.0(0.04)
Ours	53.03 (0.08)	79.59(0.03)

1. The model trained on A_{train} , when used for inference, can be interpreted as estimating $\mathbb{P}(Y \mid do(X))$. We observe that the model trained on A_{train} performs similarly on both A_{test} and B_{test} . This result indicates that the trained model does not learn the correlation.

2. The model trained on B_{train} , when used for inference, can be interpreted as estimating $\mathbb{P}(Y \mid X)$. We observe that the model trained on B_{train} exhibits different performance between A_{test} and B_{test} . This difference arises due to the influence of the confounder.

From these experiments, we verify that C is a confounder.

I.10. Randomness

To ensure that the observed performance improvement is attributable to the advantages of our method and not random variation, we conduct experiments with five different random seeds and report the mean and standard deviation for both LIFT and our method on Places365-LT and ImageNet-LT. As shown in Tab. 18, our method outperforms LIFT with a smaller standard deviation (reported in brackets), demonstrating that the performance gains are not due to randomness.

In addition, we also conduct a significance test (T-test) between the results of our method and the LIFT, obtaining a $p_{value1} = 4.57 \times 10^{-9}$ and $p_{value2} = 3.62 \times 10^{-10}$ for Places365-LT and ImageNet-LT, respectively. This indicates that our method is statistically significant, with a probability of making no more than a 5% error.