

# VL-LTR: Learning Class-wise Visual-Linguistic Representation for Long-Tailed Visual Recognition

Changyao Tian<sup>1\*†</sup>, Wenhai Wang<sup>3\*</sup>, Xizhou Zhu<sup>2\*</sup>, Jifeng Dai<sup>2✉</sup>, Yu Qiao<sup>3</sup>  
<sup>1</sup>Chinese University of Hong Kong   <sup>2</sup>SenseTime   <sup>3</sup>Shanghai AI Laboratory  
tcyhost@buaa.edu.cn   {wangwenhai, qiaoyu}@pjlab.org.cn  
{zhuwalter, daijifeng}@sensetime.com

**Abstract.** Recently, computer vision foundation models such as CLIP and ALIGN, have shown impressive generalization capabilities on various downstream tasks. But their abilities to deal with the long-tailed data still remain to be proved. In this work, we present a novel framework based on pre-trained visual-linguistic models for long-tailed recognition (LTR), termed VL-LTR, and conduct empirical studies on the benefits of introducing text modality for long-tailed recognition tasks. Compared to existing approaches, the proposed VL-LTR has the following merits. (1) Our method can not only learn visual representation from images but also learn corresponding linguistic representation from noisy class-level text descriptions collected from the Internet; (2) Our method can effectively use the learned visual-linguistic representation to improve the visual recognition performance, especially for classes with fewer image samples. We also conduct extensive experiments and set the new state-of-the-art performance on widely-used LTR benchmarks. Notably, our method achieves 77.2% overall accuracy on ImageNet-LT, which significantly outperforms the previous best method by over 17 points, and is close to the prevailing performance training on the full ImageNet. Code is available at <https://github.com/ChangyaoTian/VL-LTR>.

**Keywords:** Long-tailed Recognition, Vision-Language Models

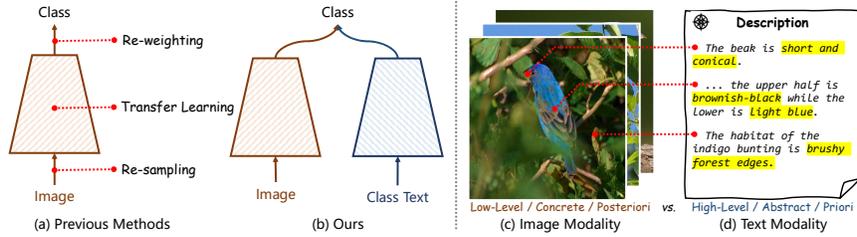
## 1 Introduction

Real-world data always presents a long-tailed distribution, where only a few head classes encompass most of the data, and most tail classes have very few samples. Such phenomenon is not conducive to the practical application of deep-learning based models. Because of this, a number of works have emerged and tried to alleviate the class imbalance problem from different aspects, such as re-sampling the training data [3,5,41], re-weighting the loss functions [22,8,28], or employing transfer learning methods [49,54,29] (see Figure 1 (a)). Despite their great contributions, most of these works still restrict themselves to only relying on the image modality for solving this problem.

---

\* Authors contributed equally. ✉ Corresponding author.

† The work is done when Changyao Tian is an intern at SenseTime Research.



**Fig. 1: Comparison of different long-tailed recognition (LTR) frameworks and different modalities.** (a) Previous LTR methods [3,5,22,8,49,54] mainly focus on the class imbalance problem on the image modality, while (b) our method addresses the LTR task by combining the advantages of image and text modalities. (c) and (d) give intuitive explanations for the correlations and differences between the image and text modalities.

As illustrated in Figure 1 (c)(d), there are some inner connections between images and text descriptions of the same class, especially when it comes to some visual concepts and attributes. However, different from the image modality that usually presents concrete low-level features (*e.g.*, shape, color, texture) of the object or scene, the text modality typically contains much high-level and abstract information. Furthermore, text descriptions are prior knowledge that can be summarized by experts, which could be useful when there are no sufficient images to learn general class-wise representation for recognition.

Although there have been some visual-linguistic approaches [15,58,36] for visual recognition, their performance is still not satisfactory, due to the gap between image and text representation and the lack of robustness to noisy text. Recently, the rise of visual-linguistic foundation models [37,42,19] has provided an effective way to learn powerful representation that can connect the image and text modalities. Motivated by this, we present a visual-linguistic framework for long-tailed recognition, termed VL-LTR, which can utilize the advantages of both visual and linguistic representation for visual recognition tasks as shown in Figure 1 (b). Our method mainly consists of two key components, which are (1) a class-wise visual-linguistic pre-training (CVLP) framework for linking images and text descriptions at the class level, and (2) a language-guided recognition (LGR) head designed to perform long-tailed recognition according to the learned visual-linguistic representation.

Overall, the proposed VL-LTR possesses the following merits. (1) Compared to the visual-linguistic pre-training [7,53,6,48,21], our method can learn visual-linguistic representation at the class level, and take the advantages of class-wise linguistic representation to improve visual recognition performance, especially in the long-tailed scenario; (2) Compared to previous visual-linguistic classifiers [15,58,36], our method can not only effectively bridge the gap between visual and linguistic representation, but also be more flexible and robust to noisy text descriptions.

To verify the effectiveness of our method, we conduct extensive experiments on three challenging long-tailed recognition (LTR) benchmarks, including ImageNet-LT [29], Places-LT [29], and iNaturalist 2018 [46]. As shown in Figure 2, using ResNet-50 [14] as backbone, our method achieves an overall accuracy of 70.1% on ImageNet-LT, which is 10.1 points higher than the previous best method PaCo [7] (ResNeXt-101 [50]). For tail classes, the medium and few-shot accuracy of our method reaches 67.0% and 50.8% respectively, which significantly outperform that of the prior arts [7,53,6,48] as well.

In summary, our main contributions are three-fold.

(1) We provide a detailed analysis on the connection and differences between image and text modalities, and point out that class descriptions can serve as a supplement to images, which is conducive to long-tailed visual recognition.

(2) We present a new visual-linguistic framework for long-tailed visual recognition (VL-LTR), which contains two tailored components, including a class-wise text-image pre-training (CVLP) to bridge the class-level images and text descriptions, and a language-guided recognition (LGR) head to perform classification based on the learned visual-linguistic representation.

(3) The proposed VL-LTR has achieved state-of-the-art performance on prevailing ImageNet-LT, Places-LT, and iNaturalist 2018 datasets. Notably, our method gets the best overall accuracy of 77.2% on ImageNet-LT, outperforming the old record by 17.2 points, and even approaching the performance training on the full ImageNet [9].

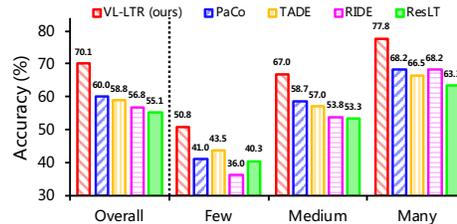


Fig. 2: **Performance comparison on ImageNet-LT [29].** Our VL-LTR (ResNet-50 [14]) significantly outperforms prior arts, including PaCo [7], TADE [53], RIDE (4 Experts) [48], and ResLT [6], which use heavier ResNeXt-50/101 [50] as backbone.

## 2 Related Work

### 2.1 Long-Tailed Visual Recognition

Class re-balanced strategy [13,22,2,11,49,18] has been comprehensively studied for long-tailed visual recognition. One type of the class re-balanced strategy is Data Resampling [11,1,3,13,5,41], which generates class-balanced data by adjusting the sampling rate of tail classes and head classes, yet they might take the risk of over-fitting on data-scarced classes. Besides that, some recent methods [5,23] augment tail class samples with head classes ones, to alleviate the over-fitting problem. Another kind of class re-balanced strategy is to design re-weighting loss functions, where tail classes would be emphasized by using large weights or margins [8,22,16,49,2], or ignoring negative gradients for tail classes [44].

In addition, researchers also address the long-tailed recognition task from the aspect of transfer learning [29,57,51,21,56]. Liu *et al.* [29] and Zhu *et al.* [57] transfer

knowledge from head classes’s features to tail classes by maintaining memory bank and modeling intra-class variance, respectively. After that, Samuel *et al.* [39] proposes a late-fusion framework for long-tail learning with class descriptors. Some decoupling methods [56,21] also can be regarded as transferring head classes frozen feature to tail classes when fine-tuning classifiers. Recently, some studies [20,47,7,40,53] also transfer the representation learned by contrastive learning or self-supervised learning for long-tailed problems.

The aforementioned methods mainly focus on addressing the class imbalance problem based on image modalities, while rarely exploring the possibility of integrating text modalities on this problem.

## 2.2 Visual-linguistic Model

In this section, we mainly discuss visual-linguistic pre-training and classification related to our work.

Visual-linguistic pre-training [32,43,4,27,52,25,12,24,33,37] have achieved great success on a number of downstream vision tasks. Zhang *et al.* [52] show the importance of visual features in visual-linguistic pre-training and obtain more strong visual representations from large object detectors. Li *et al.* [27] find that a larger transformer visual-linguistic model can learn more powerful representation from a larger visual-linguistic corpus. In addition, Huang *et al.* [17,42] proposed a visual-linguistic pre-training model by extracting patch features from the convolutional layers without the proposal computation. Recently, CLIP [37] and ALIGN [19] learns powerful visual-linguistic representation via contrastive learning on large-scale image-text pairs.

Prior to these works, there have been some visual-linguistic approaches [15,58,36] designed for tasks related to image classification. He *et al.* [15] propose a two-stream model, which directly combines visual and linguistic representation for fine-grained image classification. Mu *et al.* [36] present a few-shot visual recognition model that is regularized with text descriptions during training. Similar to He *et al.* [15], Zhuang *et al.* [58] design a multi-modal model for automatic fish classification, with a CNN encoder for images and a RNN encoder for class text. However, these methods (1) cannot effectively model the connection between images and text, leading to a considerable gap between visual and linguistic representation; and (2) require high-quality text annotations, which is usually expensive and thus limits their practical application.

## 3 Methodology

### 3.1 Overall Architecture

In order to make effective use of the linguistic modality in the visual recognition task, we propose a two-stage framework, as depicted in Figure 3. (1) The first stage is class-wise visual-linguistic pre-training (CVLP), which is used to link the images and text descriptions of the same class via contrastive learning. (2) In the second stage, a language-guided recognition (LGR) head is designed to collect the overall linguistic representation of each class to guide the image recognition. As a result, the proposed VL-LTR

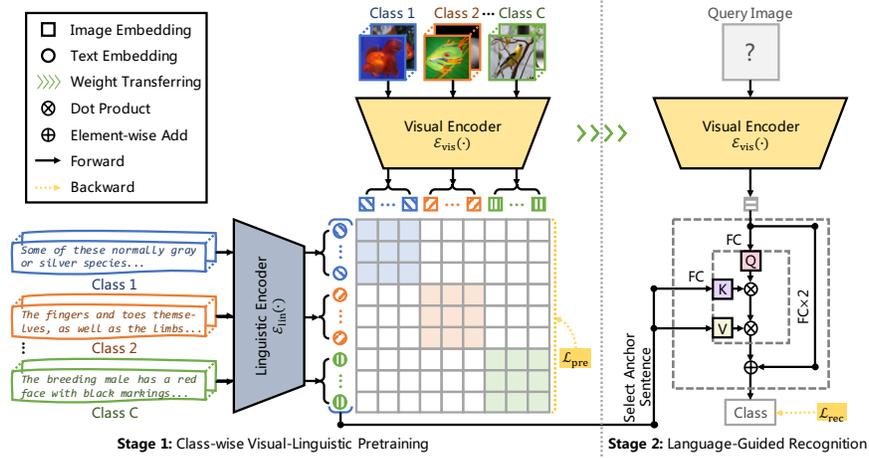


Fig. 3: **Overall architecture of VL-LTR.** The entire model has two stages. In the first stage, class-wise visual-linguistic pre-training (CVLP) takes both the images and text of each class as inputs, learning to connect the representation of the two modalities through class-wise contrastive learning. In the second stage, the language-guided recognition (LGR) head uses the learned visual-linguistic representation to perform image classification.

is able to combine the advantages of visual and linguistic representation and achieve impressive long-tailed recognition performance.

When training VL-LTR models, we first pre-train the visual and linguistic encoders by class-wise visual-linguistic contrastive learning and the pre-training loss  $\mathcal{L}_{\text{pre}}$ . During pre-training, an image and a sentence from the same class would be regarded as a positive pair, and otherwise is a negative pair. After pre-training, the weights of the linguistic encoder are frozen, and the anchor sentences of each class are then selected by filtering out the low-scored sentences in text descriptions. The visual encoder and LGR head are fine-tuned by the recognition loss  $\mathcal{L}_{\text{rec}}$ . Details of the aforementioned loss functions will be introduced in the later sections.

In the inference phase, given a query image and pre-populated text embeddings of anchor sentences, we first feed the image to the visual encoder and obtain an image embedding. Then, the image embedding passes through the LGR head and is categorized into a class according to the image embedding itself as well as the text embeddings of anchor sentences.

### 3.2 Class-wise Visual-Linguistic Pre-training

The goal of this stage is to learn the visual-linguistic representation of images and text descriptions at the class level. To this end, we design a class-wise visual-linguistic pre-training (CVLP) framework. Unlike previous works [37,19] that use instance-wise

image-text pairs for pre-training, our framework is expected to fuse the class-wise linguistic information into the visual space.

During pre-training, as shown in Figure 3, we first randomly sample a batch of images  $\mathcal{I} = \{I_i\}_{i=1}^N$ , and the corresponding text sentences  $\mathcal{T} = \{T_i\}_{i=1}^N$ , where  $N$  denotes the batch size. Then, the images  $\mathcal{I}$  and texts  $\mathcal{T}$  are fed to the visual encoder  $\mathcal{E}_{\text{vis}}(\cdot)$  and linguistic encoder  $\mathcal{E}_{\text{lin}}(\cdot)$  respectively, yielding image and text embeddings as Eqn. 1:

$$E_i^I = \mathcal{E}_{\text{vis}}(I_i), \quad E_i^T = \mathcal{E}_{\text{lin}}(T_i), \quad (1)$$

where both  $E_i^I$  and  $E_i^T$  are of  $D$  dimensions. After that, a class-wise contrastive learning (CCL) loss is used to optimize the visual and linguistic encoders. Let us denote the cosine similarity of  $E_i^I$  and  $E_j^T$  as  $S_{i,j}$ , and then the CCL loss can be formulated as:

$$\begin{aligned} \mathcal{L}_{\text{ccl}} &= \mathcal{L}_{\text{vis}} + \mathcal{L}_{\text{lin}} \\ &= -\frac{1}{|\mathcal{T}_i^+|} \sum_{T_j \in \mathcal{T}_i^+} \log \frac{\exp(S_{i,j}/\tau)}{\sum_{T_k \in \mathcal{T}} \exp(S_{i,k}/\tau)} \\ &\quad -\frac{1}{|\mathcal{I}_i^+|} \sum_{I_j \in \mathcal{I}_i^+} \log \frac{\exp(S_{j,i}/\tau)}{\sum_{I_k \in \mathcal{I}} \exp(S_{k,i}/\tau)}, \end{aligned} \quad (2)$$

where  $\mathcal{L}_{\text{vis}}$  and  $\mathcal{L}_{\text{lin}}$  denote the loss of visual and linguistic side respectively, while  $\mathcal{T}_i^+$  denotes a subset of  $\mathcal{T}$ , in which each text shares the same class with the image  $I_i$ . Correspondingly, all images in  $\mathcal{I}_i^+$  share the same class with the text  $T_i$ .  $\tau$  is a learnable parameter with an initial value of 0.07.

In addition to CCL, we also distill the knowledge from the CLIP [37] pre-trained model, to reduce the risk of over-fitting caused by limited text corpus in the pre-training stage. The distillation loss  $\mathcal{L}_{\text{dis}}$  can be written as Eqn. 3:

$$\begin{aligned} \mathcal{L}_{\text{dis}} &= -\frac{\exp(S'_{i,i}/\tau)}{\sum_{T_j \in \mathcal{T}} \exp(S'_{i,j}/\tau)} \log \frac{\exp(S_{i,i}/\tau)}{\sum_{T_k \in \mathcal{T}} \exp(S_{i,k}/\tau)} \\ &\quad -\frac{\exp(S'_{i,i}/\tau)}{\sum_{I_j \in \mathcal{I}} \exp(S'_{j,i}/\tau)} \log \frac{\exp(S_{i,i}/\tau)}{\sum_{I_k \in \mathcal{I}} \exp(S_{k,i}/\tau)}. \end{aligned} \quad (3)$$

Here,  $S'$  is the cosine similarity matrix produced by the frozen CLIP model.

Our pre-training framework has two merits as follows: (1) It is convenient to add new training samples for image or text modality in our framework, since the image and text description for a specific class is independent of each other, which greatly reduces the cost of data collection; (2) The text description of each image sample is different in each iteration, which serves as an additional regularization to prevent the model from learning some fixed trivial correlation within a certain image-text pair, and thus our framework is robust to the noisy text from the Internet.

### 3.3 Language-Guided Recognition

In this stage, we design (1) an anchor sentence selection strategy to filter out noise texts, and (2) a language-guided recognition head to effectively use visual and linguistic representation learned in the pre-training stage.

**Anchor Sentence Selection.** Most text descriptions in our corpus are crawled from the Internet, which are noisy and might degrade the recognition performance. To address this problem, we propose an anchor sentence selection (AnSS) strategy to find the most discriminative sentences for each class. Specifically, we first construct a “special” image batch  $I'$ , which contains at most 50 images (if any) of each class. Then, for each text sentence  $T_i$ , we score each sentence  $T_i$  by computing the  $\mathcal{L}_{\text{lin}}$  between the sentence and the image batch  $I'$ . Finally, we select  $M$  text sentences with the smallest  $\mathcal{L}_{\text{lin}}$  as the anchor sentences for the follow-up visual recognition.

**Language-Guided Recognition Head.** After obtaining the anchor sentences of each class, we design a language-guided recognition (LGR) head, to adjust the weights of these sentences based on the attention scores with the input image. In this way, visual and linguistic features can be flexibly and dynamically combined according to the query image.

As shown in Figure 3, given an image embedding  $E^I \in \mathbb{R}^D$ , as well as the embeddings of all classes’ anchor sentences  $E^T \in \mathbb{R}^{C \times M \times D}$ , where  $C$  is the class number, and  $M$  is the maximum number of sentences for each class. Then the LGR head can be formulated as:

$$Q = \text{Linear}(\text{LayerNorm}(E^I)), \quad (4)$$

$$K = \text{Linear}(\text{LayerNorm}(E^T)), \quad V = E^T, \quad (5)$$

$$G = \sigma\left(\frac{QK^T}{\sqrt{D}}\right)V, \quad (6)$$

$$P = P^I + P^T = \sigma(\text{MLP}(E^I)) + \sigma(\langle E^I, G \rangle / \tau). \quad (7)$$

Here,  $Q \in \mathbb{R}^D$ ,  $K \in \mathbb{R}^{C \times M \times D}$ , and  $V \in \mathbb{R}^{C \times M \times D}$  are query, key and value of the attention operation.  $G \in \mathbb{R}^{C \times D}$  is the gather of the  $M$  anchor sentence embeddings of each class.  $\sigma(\cdot)$  denotes Softmax function.  $\text{MLP}(\cdot)$  denotes two linear layers sandwich a ReLU in the middle.  $\langle E^I, G \rangle$  is the cosine similarity of  $E^I$  and  $G$ .  $P$  is the classification probability of the image  $I_q$ ,  $P^I$  and  $P^T$  are the classification probabilities based on visual and linguistic representation, respectively.

### 3.4 Loss Function

As mentioned in Section 3.1, the training process of our method has two stages, namely pre-training and fine-tuning respectively. In the pre-training stage, the visual encoder and linguistic encoder are jointly optimized by the CCL loss  $\mathcal{L}_{\text{ccl}}$  and distillation loss  $\mathcal{L}_{\text{dis}}$ . So the overall pre-training loss can be written as:

$$\mathcal{L}_{\text{pre}} = \lambda \mathcal{L}_{\text{ccl}} + (1 - \lambda) \mathcal{L}_{\text{dis}}, \quad (8)$$

where  $\lambda \in [0, 1]$  is a hyperparameter to balance  $\mathcal{L}_{\text{ccl}}$  and  $\mathcal{L}_{\text{dis}}$ .

In the fine-tuning stage, after computing the classification probabilities  $P^I$  and  $P^T$ , we simply calculate their corresponding CrossEntropy loss  $\mathcal{L}_{CE}$  with the ground truth label  $\mathbf{y}$  as Eqn. 9:

$$\mathcal{L}_{rec} = \mathcal{L}_{CE}(P^I, \mathbf{y}) + \mathcal{L}_{CE}(P^T, \mathbf{y}). \quad (9)$$

## 4 Experiments

### 4.1 Datasets

We perform extensive experiments on three challenging long-tailed visual recognition benchmarks, namely ImageNet-LT [29], Places-LT [29], and iNaturalist 2018 [46]. Among these benchmarks, ImageNet-LT is constructed from ImageNet-2012 [9] by sampling a subset following the Pareto distribution with the power value  $\alpha = 6$ , which contains 1,000 classes. The training set has 115.8K images, and the number of images per class ranges from 1,280 to 5 images. Both the validation set and the test set are balanced, containing 20K and 50K images respectively. We select the hyperparameters on the validation set and report numerical results on the test set. Similar to ImageNet-LT, Places-LT is a long-tailed version of the large-scale scene classification dataset Places [55]. It consists of 62.5K images from 365 categories with class cardinality ranging from 5 to 4,980. iNaturalist 2018 is a real-world, naturally long-tailed dataset, which is composed of 8,142 fine-grained species. The training set contains 437.5K images, and its imbalance factor is equal to 500. We use the official validation set to test our approach, which has 3 images per class.

We also collect the class-level text descriptions for the three datasets. The text descriptions mainly come from Wikipedia<sup>1</sup>, an open-source online encyclopedia that contains millions of articles for free. We first use the original class name as an initial query to get the best matching entry on Wikipedia. After cleaning and filtering out some obviously irrelevant sections such as “references” or “external links” of these entries, we split the left into sentences to form the original text candidate set for each class. Noting that some classes have relatively much fewer sentences than others, we also add 80 additional prompt sentences for each class to alleviate the data scarcity problem. These sentences, which are in the form of ‘a photo of a {label}’, are auto-generated based on the prompt templates provided in [37].

### 4.2 Evaluation Protocol

Following common practices [29,56,7], we evaluate our proposed models on the corresponding balanced validation/test set and report the overall top-1 accuracy. To diagnose the source of improvement, we also report the top-1 accuracy of the three subsets split by the number of training samples in each class, namely many-shot ( $\geq 100$  samples), medium-shot (20~100 samples), and few-shot ( $\leq 20$  samples).

<sup>1</sup> <https://en.wikipedia.org/>

### 4.3 Experiments on ImageNet-LT

**Settings.** To verify the effectiveness of our method, we conduct extensive experiments on ImageNet-LT [29]. We use ResNet-50 [14] or ViT-Base/16 [10] as the visual encoder, and a 12-layer Transformer [38] as the linguistic encoder. All models are optimized by AdamW [31] with a momentum of 0.9 and a weight decay of  $5 \times 10^{-2}$ . We use the same data augmentation as [45] (w/o distillation). In the pre-training phase, the maximum length of text tokens is set to 77 (including [SOS] and [EOS] tokens), and the pre-trained weights of CLIP [37] is loaded. The initial learning rate is set to  $5 \times 10^{-5}$  and decays following the cosine schedule [30]. During this phase, models are pre-trained for 50 epochs, with a mini-batch size of 256. In the fine-tuning phase, we select 64 sentences for each class and fine-tune models with the mini-batch size of 128 for another 50 epochs. We set the initial learning rate to  $1 \times 10^{-3}$  and still decrease it with the cosine schedule. In both stages, we adopt the input size of  $224 \times 224$  and the square-root data sampling strategy [34,35] unless specifically mentioned.

For a fair comparison, we also build a baseline that is only based on visual modality while keeping other settings exactly the same as our proposed method, except that the baseline models are directly initialized with CLIP pre-trained weights and fine-tuned for 100 epochs. In addition, we re-implement and report the performance of some representative methods as well, such as  $\tau$ -normalized, cRT, NCM, and LWS [21], which are all initialized with CLIP pre-trained weights.

**Results.** In Table 1, we can see that our VL-LTR models are superior to conventional vision-based methods with similar visual encoders (*i.e.*, backbones). For example, when using ResNet-50 (R-50) [14] as the backbone, the overall accuracy of our method reaches 70.1%, which outperforms baseline by 9.6 points (70.1% vs.60.5%), and 10.1 points better than previous best PaCo [7] (70.1% vs.60.0%).

Moreover, from the aspect of few-shot accuracy, the performance of our method is more promising, which is 16.3 points and 7.3 points better than baseline (50.8% vs.34.5%) and the second-best method [53] (50.8% vs.43.5%). When replacing the backbone with heavy ViT-Base/16 (ViT-B) [10], the overall accuracy of our method can further boost up to 77.2%, which is the current new state-of-the-art of ImageNet-LT, and near the prevailing performance (*i.e.*, 80%) training on the full ImageNet [9].

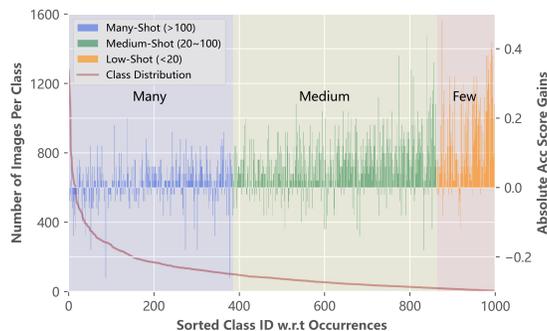


Fig. 4: Absolute accuracy score of our method over the baseline using ViT-Base/16 [10] as the backbone on ImageNet-LT [29]. Our method enjoys more performance gains on classes with fewer image samples.

Method	Backbone	Accuracy (%)			
		Overall	Many	Medium	Few
Cross Entropy [26]	ResNeXt-50	44.4	65.9	37.5	7.7
OLTR [29]	ResNeXt-50	46.3	-	-	-
SSD [26]	ResNeXt-50	56.0	66.8	53.1	35.4
RIDE (4 Experts) [48]	ResNeXt-50	56.8	68.2	53.8	36.0
TADE [53]	ResNeXt-50	58.8	66.5	57.0	43.5
smDRAGON [39]	ResNeXt-50	50.1	-	-	-
ResLT [6]	ResNeXt-101	55.1	63.3	53.3	40.3
PaCo [7]	ResNeXt-101	60.0	68.2	58.7	41.0
NCM [21]	ResNeXt-152	51.3	60.3	49.0	33.6
cRT [21]	ResNeXt-152	52.4	64.7	49.1	29.4
$\tau$ -normalized [21]	ResNeXt-152	52.8	62.2	50.1	35.8
LWS [21]	ResNeXt-152	53.3	63.5	50.4	34.2
NCM [21]	ResNet-50*	49.2	58.9	46.6	31.1
cRT [21]	ResNet-50*	50.8	63.3	47.2	27.8
$\tau$ -normalized [21]	ResNet50*	51.2	60.9	48.4	33.8
LWS [21]	ResNet-50*	51.5	62.2	48.6	31.8
Zero-Shot CLIP [37]	ResNet-50*	59.8	60.8	59.3	58.6
Baseline	ResNet-50*	60.5	74.4	56.9	34.5
VL-LTR (ours)	ResNet-50*	<b>70.1</b>	<b>77.8</b>	<b>67.0</b>	<b>50.8</b>
VL-LTR (ours)	ViT-Base*	<b>77.2</b>	<b>84.5</b>	<b>74.6</b>	<b>59.3</b>

Table 1: **Results on ImageNet-LT [29]**. Our method outperforms prior arts when using a similar backbone network. “\*” indicates the corresponding backbone is initialized with CLIP [37] weights.

In Figure 4, we visualize the class-level performance improvement, which is measured by the absolute accuracy gains of our method against the baseline, both of which use ViT-B as the visual backbone. We see that there is more gains on tail classes, which indicates that our method can help mitigate the data-scarce problem under long-tail settings by introducing class-level text descriptions.

#### 4.4 Experiments on Places-LT

**Settings.** We also investigate our method on Places-LT [29], a dataset with a different domain. The experimental setting of Places-LT is the same as Section 4.3.

**Results.** As reported in Table 2, using ResNet-50 (R-50) as backbone, our model achieves 48.0% overall accuracy, surpassing counterparts by at least 6.8 points (48.0% vs.41.2%), including state-of-the-art PaCo [7], TADE [53], and ResLT [6], while all of them use ResNet-152 [14] as backbone. The performance are also impressive for the medium- (47.2%) and few-shot (38.4%) classes. Once again, the model with ViT-Base/16 (ViT-B) [10] gives the top overall accuracy of 50.1%, which is a new state-of-the-art on this benchmark.

Method	Backbone	Accuracy (%)			
		Overall	Many	Medium	Few
OLTR [29]	ResNet-152	35.9	44.7	37.0	25.3
ResLT [6]	ResNet-152	39.8	39.8	43.6	31.4
TADE [53]	ResNet-152	40.9	40.4	43.2	36.8
PaCo [7]	ResNet-152	41.2	36.1	47.9	35.3
NCM [21]	ResNet-152	36.4	40.4	37.1	27.3
cRT [21]	ResNet-152	36.7	42.0	37.6	24.9
$\tau$ -normalized [21]	ResNet-152	37.9	37.8	40.7	31.8
LWS [21]	ResNet-152	37.6	40.6	39.1	28.6
smDRAGON [39]	ResNet-50	38.1	-	-	-
NCM [21]	ResNet-50*	30.8	37.1	30.6	19.9
cRT [21]	ResNet-50*	30.5	38.5	29.7	17.6
$\tau$ -normalized [21]	ResNet-50*	31.0	34.5	31.4	23.6
LWS [21]	ResNet-50*	31.3	36.0	32.1	20.7
Zero-Shot CLIP [37]	ResNet-50*	38.0	37.5	37.5	40.1
Baseline	ResNet-50*	39.7	50.8	38.6	22.7
VL-LTR (ours)	ResNet-50*	<b>48.0</b>	<b>51.9</b>	<b>47.2</b>	<b>38.4</b>
VL-LTR (ours)	ViT-Base*	<b>50.1</b>	<b>54.2</b>	<b>48.5</b>	<b>42.0</b>

Table 2: **Results on Places-LT [29]**. “\*” indicates the corresponding backbone is initialized with CLIP [37] weights.

#### 4.5 iNaturalist 2018

**Settings.** We further test our VL-LTR on iNaturalist 2018, a long-tailed fine-grained benchmark. Following the common practice [45], we adopt a long training schedule. To be specific, our models are pre-trained for 100 epochs, and fine-tuned for 360 epochs. The initial learning rate of the pre-training and fine-tuning phase is set to  $5 \times 10^{-4}$  and  $2 \times 10^{-5}$ , respectively. Correspondingly, the baseline has the same fine-tuning epochs and initial learning rate as the proposed method. All other experimental settings are the same as Section 4.3.

**Results.** Table 3 shows the top-1 accuracy on iNaturalist 2018 of different methods. We see that when using ResNet-50 (R-50) [14] as the backbone, our models can achieve a 74.6% overall accuracy, surpassing previous methods with the same backbone by at least 1.4 points. Besides that, when equipped with a strong backbone ViT-Base/16 (ViT-B) [10], our model can have an overall accuracy of 76.8%, which outperforms the state-of-the-art PaCo (ResNet-152) by 1.6 points (76.8% vs.75.2%). Moreover, our model can also benefit from a larger image input size (*i.e.*,  $384 \times 384$ ), and achieve 81.0% top-1 accuracy, which is 1.5 points higher than DeiT-B/16-384 [45] (81.0% vs.79.5%).

#### 4.6 Ablation Study

**Settings.** In order to provide a deep analysis of our proposed method, we also conduct ablation studies on the ImageNet-LT dataset. In these experiments, we use ResNet-50 as the default backbone. All other settings remain the same as Section 4.3 unless specifically mentioned.

Method	Backbone	Accuracy (%)
CB-Focal [2]	ResNet-50	61.1
LDAM+DRW [2]	ResNet-50	68.0
BBN [56]	ResNet-50	69.6
SSD [26]	ResNet-50	71.5
RIDE (4 experts) [48]	ResNet-50	72.6
smDRAGON [39]	ResNet-50	69.1
ResLT [6]	ResNet-50	72.3
TADE [53]	ResNet-50	72.9
PaCo [7]	ResNet-50	73.2
NCM [21]	ResNet-50	63.1
cRT [21]	ResNet-50	67.6
$\tau$ -normalized [21]	ResNet-50	69.3
LWS [21]	ResNet-50	69.5
NCM [21]	ResNet-50*	65.3
cRT [21]	ResNet-50*	69.9
$\tau$ -normalized [21]	ResNet-50*	71.2
LWS [21]	ResNet-50*	71.0
Zero-Shot CLIP [37]	ResNet-50*	3.4
Baseline	ResNet-50*	72.6
VL-LTR (ours)	ResNet-50*	<b>74.6</b>
PaCo [7]	ResNet-152	75.2
DeiT-B/16 [45]	-	73.2
DeiT-B/16-384 [45]	-	79.5
VL-LTR (ours)	ViT-Base*	<b>76.8</b>
VL-LTR-384 (ours)	ViT-Base*	<b>81.0</b>

Table 3: **Results on iNaturalist 2018** [46]. “\*” indicates the corresponding backbone is initialized with CLIP weights. “\*-384” means the input size of  $384 \times 384$ .

**Class-wise Visual-Linguistic Pre-training.** To examine the effectiveness of our class-wise visual-linguistic pre-training (CVLP) framework, we remove it by directly performing the fine-tuning process on the pre-trained weights of CLIP [37]. As reported in the #1 and #2 of Table 4, the model with CVLP outperforms the one without CVLP by 7.3 points on the overall accuracy. Such gap might be attributed to the inconsistency between image and text representation, which can be alleviated by our CVLP.

To verify this, we visualize some concepts by retrieving images with the greatest cosine similarity. As shown in Figure 6, both CLIP and our method can learn common visual concepts, such as the “blue” color, but CLIP [37] fails to capture rare concepts, such as “stick” shape and “spot” texture. More examples are provided in the supplementary material.

**CLIP Pre-trained Weights.** To analyze the influence of CLIP pre-trained weights, we train our method with randomly initialized weights. Comparing the #1 and #3 of Table 4, we can see that initializing with CLIP pre-trained weights benefits our VL-LTR. We also plot the curves of training and validation loss in the fine-tuning stage in Figure 5, where CLIP pre-trained weights (see red curves) can help alleviate the over-fitting problem. This phenomenon is caused by the limited text corpus for pre-training. There

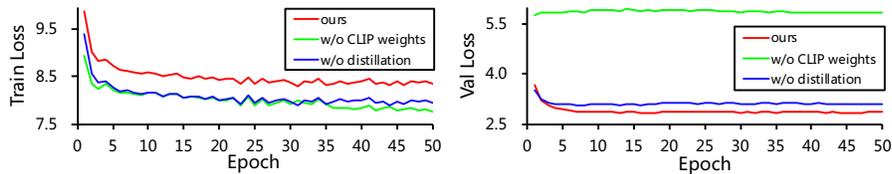


Fig. 5: Train and validation loss curves of VL-LTR (ResNet-50) on ImageNet-LT [29] under different settings. Both without CLIP [37] weights (w/o CLIP) and without distillation lead to a certain degree of overfitting.

#	CLIP Weights	Pre-training		Fine-tuning		Accuracy (%)
		w/o $\mathcal{L}_{dis}$	w/ $\mathcal{L}_{dis}$	Head	SS	
1	✓	-	✓	LGR	AnSS	<b>70.1</b>
2	✓	-	-	LGR	AnSS	62.8
3	-	✓	-	LGR	AnSS	46.8
4	✓	✓	-	LGR	AnSS	66.2
5	✓	-	✓	FC	-	62.1
6	✓	-	✓	KNN	-	63.9
7	✓	-	✓	LGR	Cut Off	69.7

Table 4: Ablation studies on ImageNet-LT [29]. “Head” denotes the recognition head used in the fine-tuning stage, and “SS” denotes the sentence selection strategy.

are only 1000 class descriptions (about 127K sentences) for ImageNet-LT, it is easy to overfit an image to a specific set of sentences without a pre-trained linguistic encoder.

**Distillation Loss.** Similar to the role of pre-trained weights, the distillation loss  $\mathcal{L}_{dis}$  is also used to reduce the risk of over-fitting in the pre-training phase. Comparing the red and blue curves in Figure 5, the over-fitting problem is alleviated in the model with  $\mathcal{L}_{dis}$ . From the #1 and #4 of Table 4, we also see that the model with  $\mathcal{L}_{dis}$  performs better than the one without  $\mathcal{L}_{dis}$  (70.1% vs.66.2%).

**Linguistic-Guided Recognition.** We verify the effectiveness of linguistic-guided recognition (LGR) by comparing it with other recognition heads, including FC (vision-based), and KNN (vision-language-based). As reported in #1, #5, and #6 of Table 4, the proposed LGR performs better than FC and KNN by 8.0% and 6.2% points in overall accuracy respectively. It is notable that, as a vision-language-based recognition head, KNN also works better than FC. These results indicate the effectiveness of LGR and the power of visual-linguistic representation.

**Anchor Sentence Selection.** We study the effectiveness of anchor sentence selection (AnSS) by replacing it with “Cut Off” strategy, where we simply select the first  $M$  sentences from text descriptions as the anchor sentences for visual recognition. As shown in Table 4, the model with AnSS (see the #1 of Table 4) outperforms the model with “Cut Off” on the overall accuracy, which proves the effectiveness of AnSS to filter out some noisy sentences. Note that, AnSS is a training-free module, which can bring considerable improvements in noisy scenes.

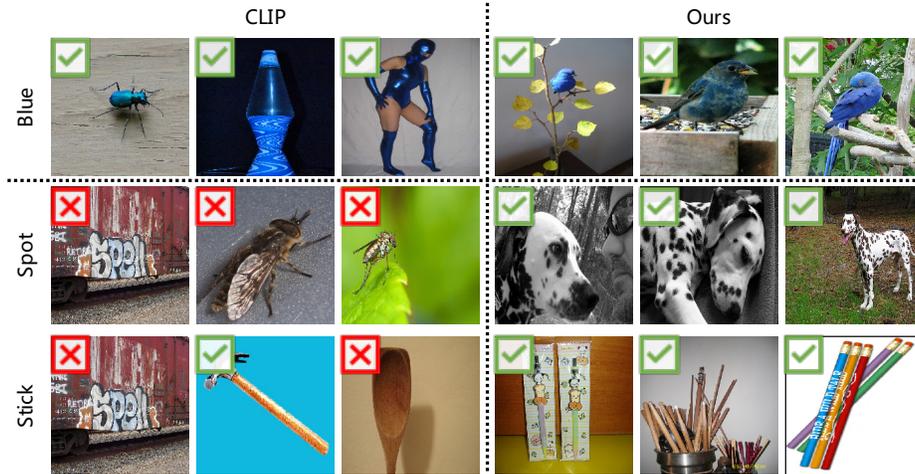


Fig. 6: **Concept visualization**, where “freq” and “rare” mean concepts appear frequently and rarely, respectively. Our method can effectively learn common visual concepts, and even the rare concepts where CLIP [37] makes mistakes, such as “spot” texture and “stick” shape.

#### 4.7 Limitations

Although the proposed VL-LTR achieves good performance on multiple long-tailed recognition benchmarks, it still has some flaws. First, due to the limited text corpus, our method currently relies on existing pre-trained foundation models to learn high-quality linguistic representation. Second, like most LTR works [26,7,21], our VL-LTR is a two-stage method as well, which does not support end-to-end training. But we believe these problems could be well addressed in the future with the enrichment of text data and the development of visual-linguistic model.

## 5 Conclusions

In this work, we introduce VL-LTR, a new visual-linguistic framework for long-tailed recognition. We develop a class-level visual-linguistic pre-training (CVLP) to connect images and text descriptions at class level, and a language-guided recognition (LGR) head to make effective use of visual-linguistic representation for visual recognition. Extensive experiments on various long-tailed recognition benchmarks verify that our method works better than well-designed vision-based methods. We hope this work could provide a strong baseline for vision-language-based long-tailed visual recognition.

## References

1. Buda, M., Maki, A., Mazurowski, M.A.: A systematic study of the class imbalance problem in convolutional neural networks. *IEEE Trans. Neural Netw.* **106**, 249–259 (2018) [3](#)
2. Cao, K., Wei, C., Gaidon, A., Aréchiga, N., Ma, T.: Learning imbalanced datasets with label-distribution-aware margin loss. In: *Proc. Advances in Neural Inf. Process. Syst.* (2019) [3](#), [12](#)
3. Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research* **16**, 321–357 (2002) [1](#), [2](#), [3](#)
4. Chen, Y.C., Li, L., Yu, L., El Kholy, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740* (2019) [4](#)
5. Chu, P., Bian, X., Liu, S., Ling, H.: Feature space augmentation for long-tailed data. In: *Proc. Eur. Conf. Comp. Vis.* (2020) [1](#), [2](#), [3](#)
6. Cui, J., Liu, S., Tian, Z., Zhong, Z., Jia, J.: Reslt: Residual learning for long-tailed recognition. *arXiv preprint arXiv:2101.10633* (2021) [2](#), [3](#), [10](#), [11](#), [12](#)
7. Cui, J., Zhong, Z., Liu, S., Yu, B., Jia, J.: Parametric contrastive learning. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2021) [2](#), [3](#), [4](#), [8](#), [9](#), [10](#), [11](#), [12](#), [14](#)
8. Cui, Y., Jia, M., Lin, T.Y., Song, Y., Belongie, S.: Class-balanced loss based on effective number of samples. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2019) [1](#), [2](#), [3](#)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2009) [3](#), [8](#), [9](#)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *Proc. Int. Conf. Learn. Representations* (2021) [9](#), [10](#), [11](#)
11. Drummond, C., Holte, R.C., et al.: C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In: *Workshop on learning from imbalanced datasets II*. vol. 11, pp. 1–8 (2003) [3](#)
12. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. In: *Proc. Advances in Neural Inf. Process. Syst.* (2020) [4](#)
13. Han, H., Wang, W.Y., Mao, B.H.: Borderline-smote: a new over-sampling method in imbalanced data sets learning. In: *Proc. Int. Conf. Intell. Comput.* (2005) [3](#)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2016) [3](#), [9](#), [10](#), [11](#)
15. He, X., Peng, Y.: Fine-grained image classification via combining vision and language. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2017) [2](#), [4](#)
16. Huang, C., Li, Y., Loy, C.C., Tang, X.: Learning deep representation for imbalanced classification. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2016) [3](#)
17. Huang, Z., Zeng, Z., Liu, B., Fu, D., Fu, J.: Pixel-bert: Aligning image pixels with text by deep multi-modal transformers. *arXiv preprint arXiv:2004.00849* (2020) [4](#)
18. Jamal, M.A., Brown, M., Yang, M.H., Wang, L., Gong, B.: Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2020) [3](#)
19. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: *Proc. Int. Conf. Mach. Learn.* (2021) [2](#), [4](#), [5](#)
20. Kang, B., Li, Y., Xie, S., Yuan, Z., Feng, J.: Exploring balanced feature spaces for representation learning. In: *Proc. Int. Conf. Learn. Representations* (2020) [4](#)

21. Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. In: Proc. Int. Conf. Learn. Representations (2019) [2](#), [3](#), [4](#), [9](#), [10](#), [11](#), [12](#), [14](#)
22. Khan, S.H., Hayat, M., Bennamoun, M., Sohel, F.A., Togneri, R.: Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE Trans. Neural Netw. & Learn. Syst.* **29**(8), 3573–3587 (2017) [1](#), [2](#), [3](#)
23. Kim, J., Jeong, J., Shin, J.: M2m: Imbalanced classification via major-to-minor translation. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2020) [3](#)
24. Li, L., Chen, Y., Cheng, Y., Gan, Z., Yu, L., Liu, J.: HERO: hierarchical encoder for video+language omni-representation pre-training. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020 (2020) [4](#)
25. Li, L., Gan, Z., Liu, J.: A closer look at the robustness of vision-and-language pre-trained models. arXiv preprint arXiv:2012.08673 (2020) [4](#)
26. Li, T., Wang, L., Wu, G.: Self supervision to distillation for long-tailed visual recognition. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2021) [10](#), [12](#), [14](#)
27. Li, X., Yin, X., Li, C., Zhang, P., Hu, X., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., et al.: Oscar: Object-semantics aligned pre-training for vision-language tasks. In: Proc. Eur. Conf. Comp. Vis. (2020) [4](#)
28. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proc. IEEE Int. Conf. Comp. Vis. (2017) [1](#)
29. Liu, Z., Miao, Z., Zhan, X., Wang, J., Gong, B., Yu, S.X.: Large-scale long-tailed recognition in an open world. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2019) [1](#), [3](#), [8](#), [9](#), [10](#), [11](#), [13](#)
30. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. In: Proc. Int. Conf. Learn. Representations (2017) [9](#)
31. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: Proc. Int. Conf. Learn. Representations (2019) [9](#)
32. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: Proc. Advances in Neural Inf. Process. Syst. (2019) [4](#)
33. Lu, J., Goswami, V., Rohrbach, M., Parikh, D., Lee, S.: 12-in-1: Multi-task vision and language representation learning. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2020) [4](#)
34. Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., Van Der Maaten, L.: Exploring the limits of weakly supervised pretraining. In: Proc. Eur. Conf. Comp. Vis. (2018) [9](#)
35. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Proc. Advances in Neural Inf. Process. Syst. (2013) [9](#)
36. Mu, J., Liang, P., Goodman, N.D.: Shaping visual representations with language for few-shot classification. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020 (2020) [2](#), [4](#)
37. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J.: Learning transferable visual models from natural language supervision. arXiv preprint arXiv:2103.00020 (2021) [2](#), [4](#), [5](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
38. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019) [9](#)
39. Samuel, D., Atzmon, Y., Chechik, G.: From generalized zero-shot learning to long-tail with class descriptors. In: IEEE Winter. Conf. Applications of Comp. Vis. (2021) [4](#), [10](#), [11](#), [12](#)
40. Samuel, D., Chechik, G.: Distributional robustness loss for long-tail learning. arXiv preprint arXiv:2104.03066 (2021) [4](#)

41. Shen, L., Lin, Z., Huang, Q.: Relay backpropagation for effective learning of deep convolutional neural networks. In: Proc. Eur. Conf. Comp. Vis. (2016) [1](#), [3](#)
42. Su, W., Zhu, X., Cao, Y., Li, B., Lu, L., Wei, F., Dai, J.: VI-bert: Pre-training of generic visual-linguistic representations. In: Proc. Int. Conf. Learn. Representations (2019) [2](#), [4](#)
43. Tan, H., Bansal, M.: LXMERT: learning cross-modality encoder representations from transformers. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019 (2019) [4](#)
44. Tan, J., Wang, C., Li, B., Li, Q., Ouyang, W., Yin, C., Yan, J.: Equalization loss for long-tailed object recognition. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2020) [3](#)
45. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: Proc. Int. Conf. Mach. Learn. (2021) [9](#), [11](#), [12](#)
46. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2018) [3](#), [8](#), [12](#)
47. Wang, P., Han, K., Wei, X.S., Zhang, L., Wang, L.: Contrastive learning based hybrid networks for long-tailed image classification. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2021) [4](#)
48. Wang, X., Lian, L., Miao, Z., Liu, Z., Yu, S.X.: Long-tailed recognition by routing diverse distribution-aware experts. In: Proc. Int. Conf. Learn. Representations (2020) [2](#), [3](#), [10](#), [12](#)
49. Wang, Y.X., Ramanan, D., Hebert, M.: Learning to model the tail. In: Proc. Advances in Neural Inf. Process. Syst. (2017) [1](#), [2](#), [3](#)
50. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2017) [3](#)
51. Yin, X., Yu, X., Sohn, K., Liu, X., Chandraker, M.: Feature transfer learning for face recognition with under-represented data. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2019) [3](#)
52. Zhang, P., Li, X., Hu, X., Yang, J., Zhang, L., Wang, L., Choi, Y., Gao, J.: Vinvl: Revisiting visual representations in vision-language models. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2021) [4](#)
53. Zhang, Y., Hooi, B., Hong, L., Feng, J.: Test-agnostic long-tailed recognition by test-time aggregating diverse experts with self-supervision. arXiv preprint arXiv:2107.09249 (2021) [2](#), [3](#), [4](#), [9](#), [10](#), [11](#), [12](#)
54. Zhong, Y., Deng, W., Wang, M., Hu, J., Peng, J., Tao, X., Huang, Y.: Unequal-training for deep face recognition with long-tailed noisy data. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2019) [1](#), [2](#)
55. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. IEEE Trans. Pattern Anal. Mach. Intell. **40**(6), 1452–1464 (2017) [8](#)
56. Zhou, B., Cui, Q., Wei, X.S., Chen, Z.M.: Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2020) [3](#), [4](#), [8](#), [12](#)
57. Zhu, L., Yang, Y.: Inflated episodic memory with region self-attention for long-tailed visual recognition. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2020) [3](#)
58. Zhuang, P., Wang, Y., Qiao, Y.: Wildfish++: A comprehensive fish benchmark for multimedia research. IEEE Trans. Multimedia **23**, 3603–3617 (2020) [2](#), [4](#)