

LMPT: Prompt Tuning with Class-Specific Embedding Loss for Long-Tailed Multi-Label Visual Recognition

Peng Xia¹, Di Xu², Ming Hu¹, Lie Ju¹, Zongyuan Ge¹

¹Monash University, ²Imperial College London

richard.peng.xia@gmail.com, zongyuan.ge@monash.edu

Abstract

Long-tailed multi-label visual recognition (LTML) task is a highly challenging task due to the label co-occurrence and imbalanced data distribution. In this work, we propose a unified framework for LTML, namely prompt tuning with class-specific embedding loss (LMPT), capturing the semantic feature interactions between categories by combining text and image modality data and improving the performance synchronously on both head and tail classes. Specifically, LMPT introduces the embedding loss function with class-aware soft margin and re-weighting to learn class-specific contexts with the benefit of textual descriptions (captions), which could help establish semantic relationships between classes, especially between the head and tail classes. Furthermore, taking into account the class imbalance, the distribution-balanced loss is adopted as the classification loss function to further improve the performance on the tail classes without compromising head classes. Extensive experiments are conducted on VOC-LT and COCO-LT datasets, which demonstrates that our method significantly surpasses the previous state-of-the-art methods and zero-shot CLIP in LTML. Our codes are fully public at <https://github.com/richard-peng-xia/LMPT>.

1 Introduction

Long-tailed multi-label visual recognition (LTML) (Wu et al., 2020; Guo and Wang, 2021) is a common and practical task owing to the highly imbalanced data distribution (Zhang et al., 2021b) and diverse objects of real-world images (Wang et al., 2017; Ju et al., 2023). Compared with long-tailed recognition and multi-label recognition tasks, LTML is more complex and challenging, because it requires capturing multiple categories and the label co-occurrence in individual images (Chen et al., 2019a), which needs to compensate for the negative impacts caused by the long-tailed distribution (i.e., *low performance on the tail classes*).




Figure 1: The class distribution is long-tailed and the VLM compares image embeddings* to text embeddings• of the class, which means the closer the distance between the embeddings of different modalities, the higher the probability that the category of the text embeddings matches the image. (a) Person and horse in the image belong to the head classes and the tail classes respectively. (b) Zero-Shot CLIP. (c) Existing Prompt Tuning *w/o* CSE loss. (d) LMPT (Ours) *w/* CSE loss.

Several approaches have been proposed to address the LTML problem from different perspectives, such as re-sampling (Buda et al., 2018; Dong et al., 2017; Guo and Wang, 2021), re-weighting (Cao et al., 2019; Wu et al., 2020) and modeling more powerful structures (Chen et al., 2019a; Wang et al., 2016, 2017). Despite their great contributions, these works neglect to take into account two crucial aspects. *First of all*, the importance of semantic feature interaction between classes to capture label co-occurrence. However, these methods are limited to balancing the distribution of categories from the perspective of samples, without considering the feature correlation between different classes. *Second*, synchronous improvements in head-to-tail category performance, while some of these works improve the performance of tail classes at the expense of the head classes.

Recently, graphic models have been introduced

to model the semantic label correlation in a few works (Chen et al., 2019a; Wang et al., 2016), whereas these works are complex and are modeling label dependencies mainly based on the image modality without additional semantic information from other modal data. Vision-language models (VLMs) (Radford et al., 2021; Jia et al., 2021; Tian et al., 2022; Huang et al., 2022; Xia et al., 2024a) demonstrate the huge potential of text modality on semantic context feature for downstream visual tasks, especially for the prompt tuning methods (Schick and Schütze, 2021; Shin et al., 2020; Yao et al., 2021; Xia et al., 2023), which provide an efficient way to transfer pre-trained VLMs to downstream tasks by learning the task-specific prompts rather than finetuning the entire model. Nonetheless, the existing prompt tuning methods (Zhou et al., 2022b,a; Sun et al., 2022) for visual recognition simply minimize prediction errors using the classification loss (*e.g.*, cross-entropy loss) with respect to the learnable prompts, which may lead to learning general embeddings or inaccurate class-related embeddings. For instance, when presented with an image (Fig.1a) that contains both a head class [person] and a tail class [horse], the zero-shot method (Fig.1b) relies solely on the rich knowledge of the pre-trained VLMs to assess the similarity between the image and the word embeddings of the class names, while the existing prompt tuning method (Fig.1c) further learns more generalized prompt tokens to improve model performance. However, these methods do not consider the inter-class relationships, particularly between head and tail classes, which is a critical factor for LTML. This underscores the need for approaches that incorporate such relationships to improve performance in such scenarios.

Therefore, to address these issues, we present the class-specific embedding loss for prompt tuning on long-tailed multi-label visual recognition, called LMPT. The abundance of image-caption data facilitates prompt learning that encompasses more nuanced and specific textual descriptions, as well as the semantic inter-dependencies between categories (Fig.1d) that share information, such as similar features or common descriptions. This attribute is particularly critical in the identification of both head and tail classes. More specifically, we propose the class-specific embedding loss to enhance the inclusivity of class-related embeddings within prompts. By gradually approaching

the embeddings of the corresponding caption, our proposed approach enables prompt tokens to effectively judge the association between different classes with the aid of textual modality. Aiming for class imbalance and consistency improvements between head classes and tail classes, we integrate class-aware soft margin and re-weighting into the class-specific embedding loss, which serves to assign larger margins and more weights to tail classes. Notably, for images containing both head and tail classes, our approach outperforms visual models and current prompt tuning methods. Moreover, we adopt the distribution-balanced loss (Wu et al., 2020) as the classification loss. To sum up, the main contributions of this work include:

- We propose the LMPT framework to adapt pre-trained VLMs to tackle long-tailed multi-label visual recognition, where captions are easily accessible from public image-caption datasets or generated by powerful image-caption models (Wang et al., 2022).
- We present a novel class-specific embedding loss with class-aware soft margin and re-weighting to learn more fine-grained and class-related embeddings that build semantic relationships across head and tail classes with shared semantic information. Such design can benefit performance in tail classes and hard-to-recognize classes with the help of text modality.
- We verify the effectiveness of the proposed method by achieving new state-of-the-art (SOTA) results on two datasets, which outperform previous SOTA (Guo and Wang, 2021) by 9/6% and zero-shot CLIP by 6/2% on VOC-LT / COCO-LT.

2 Related Work

2.1 Long-Tailed Visual Recognition

Real-world training data usually exhibits long-tailed distribution (Zhang et al., 2021b), which presents a challenge for traditional methods due to the imbalanced class distribution. To address this problem, several approaches (Cui et al., 2022; Menon et al., 2020; Ouyang et al., 2016; Samuel and Chechik, 2021; Xia et al., 2024b) have been proposed from different aspects. One common method is to directly re-sample the training data to balance the class distribution (Drummond et al., 2003; Buda et al., 2018; Dong et al., 2017), by adjusting the sampling rate of head classes and tail

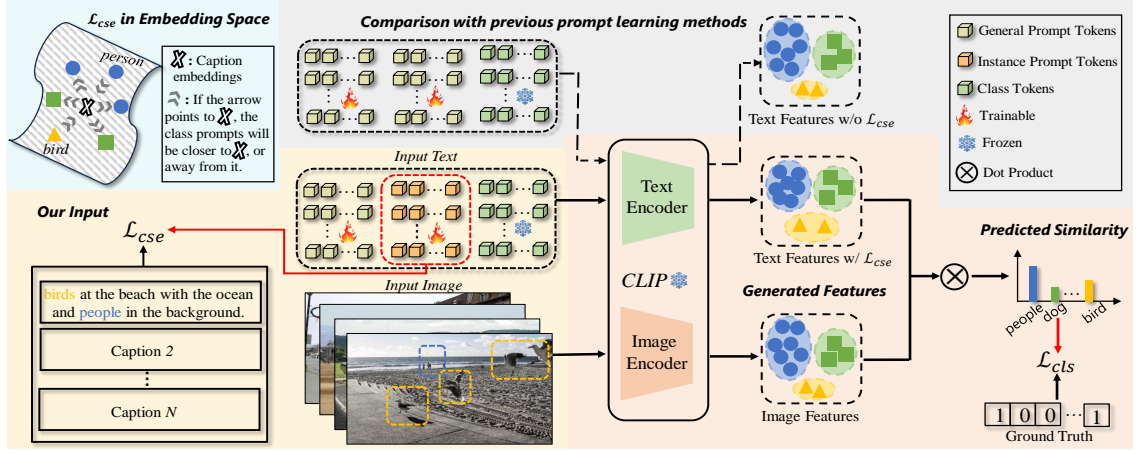


Figure 2: Overview of the architecture of our proposed method. The color blocks are defined as shown in Fig. 1.

classes, yet it might lead to the overfitting of tail classes. A better solution is to design re-weighted loss functions (Khan et al., 2017; Huang et al., 2016; Cao et al., 2019) that assign more weight to tail classes or ignore negative gradients (Tan et al., 2020) for tail classes. In addition, researchers also propose to use techniques such as transfer learning (Liu et al., 2019; Zhu and Yang, 2020) and self-supervised learning (Kang et al., 2020; Zhang et al., 2021a) to alleviate the class imbalance problem. Recently, some studies (Ma et al., 2021; Tian et al., 2022) also explore the possibility of text modality by refining visual-language representations on the long-tailed recognition tasks.

2.2 Multi-Label Visual Recognition

For multi-label visual recognition, some early methods include treating it as multiple binary image classifications (Tsoumakas and Katakis, 2007; Zhang and Zhou, 2013) or finding k-nearest neighbors (Zhang and Zhou, 2007). To locate regions of interest, some researchers (Wang et al., 2016, 2017) proposed to introduce recurrent neural networks (e.g., RNN, LSTM) to learn a joint image-label embedding. In addition, Chen et al. (Chen et al., 2019a) proposed to model the label correlations by constructing a graph based on the label co-occurrence and Ye et al. (Ye et al., 2020) updated static graph to dynamic graph convolutional network (GCN) for robust representation. Wu et al. (Wu et al., 2020) proposed a distribution-balanced loss and Guo et al. (Guo and Wang, 2021) adopted collaborative training on the uniform and re-balanced samplings to alleviate the class imbalanced problem. There is also a popular trend to align between visual and textual features (Xu et al.,

2022; Liu et al., 2021; Huang et al., 2022; Ridnik et al., 2023) for multi-label recognition.

2.3 Prompt Tuning for Vision-Language Models

Prompt tuning (Schick and Schütze, 2021; Shin et al., 2020; Yao et al., 2021) is a parameter-efficient technique used to utilize the representation ability of pre-trained vision-language models to achieve better performance instead of fine-tuning the whole model on downstream tasks. Meanwhile, large-scale vision-language models (e.g., CLIP (Radford et al., 2021), ALIGN (Jia et al., 2021)) have demonstrated impressive power to learn visual and textual features. CoOp (Zhou et al., 2022b) learns soft prompts via minimizing the classification loss and CoCoOp (Zhou et al., 2022a) further formulates the prompts in an image-conditional way to improve its generalization to unseen classes. DualCoOp (Sun et al., 2022) firstly adapts CLIP to multi-label image recognition by learning pairs of positive and negative prompts for each class, then TaI-DPT (Guo et al., 2023) extracts both coarse-grained and fine-grained embedding by treating texts as images in prompt tuning. Different from the above work, LMPT focuses on exploring the transfer ability to address long-tailed multi-label visual recognition.

3 Methodology

In this section, we present our proposed prompting tuning method, *i.e.*, LMPT, for adapting pre-trained vision-language models for long-tailed multi-label visual recognition.

3.1 Preliminaries

Consider \mathcal{D} as the dataset we use, N as the number of the dataset, C as the number of classes, and L as the fixed length of contexts for optimization. Then $(x^k, y^k, t^k) \in \mathcal{D}_{train}$, $k \in \{1, \dots, N\}$, where x^k is an input single image, $y^k = [y_1^k, \dots, y_C^k] \in \{0, 1\}^C$ is the multi-label ground-truth and $t^k = [t_1^k, \dots, t_L^k]$ is the corresponding text embedding of text description (caption). But during the test phase, only $(x^k, y^k) \in \mathcal{D}_{test}$. Let $n_i = \sum_{k=1}^N y_i^k$ denote the number of training examples that contain class i . Please note that labels for computing the class-specific embedding loss need to be processed into $\tilde{y}^k = [\tilde{y}_1^k, \dots, \tilde{y}_C^k] = [2 * y_1^k - 1, \dots, 2 * y_C^k - 1] \in \{-1, 1\}^C$, where $\{-1, 1\}$ indicates negative and positive.

3.2 Approach Overview

In order to make effective use of the linguistic modality in the long-tailed multi-label visual recognition task, we propose a novel framework (*i.e.*, LMPT), as depicted in Fig. 2. Text encoder from the pre-trained CLIP is used to encode the prompts and text descriptions (captions) of images. Only the parameters in the prompts are optimized, while the text encoder and image encoder are both kept frozen. We introduce two sorts of trainable prompts to obtain class embedding, which are jointly optimized by the classification loss \mathcal{L}_{cls} and class-specific embedding loss \mathcal{L}_{cse} . Details of the aforementioned loss functions will be introduced in the later sections.

3.3 Prompt Tuning

Formally, the vision-language model consists of an image encoder $f(\cdot)$ and a text encoder $g(\cdot)$. Following (Zhou et al., 2022a), a prompt is defined as:

$$o_i|_1^M = [V]_1 [V]_2 \dots [V]_m \dots [V]_M [\text{CLASS}], \quad (1)$$

where $i \in \{1, \dots, C\}$, $m \in \{1, \dots, M\}$, the [CLASS] token is replaced by the specific class name (*e.g.*, “cat,” “dog”, “car”), each $[V]_m$ is a learnable word embedding with the same dimension as normal word embeddings in the vocabulary (*i.e.*, 512 for CLIP), and M is a hyper-parameter specifying the number of context tokens. The prediction probability (classification output) z is then




Figure 3: The class margins (dotted lines) are enforced for generated samples by updating the decision boundary with respect to class margins.

computed as:

$$p(y = i | x) = \frac{\exp(\cos(\mathbf{g}(o_i), \mathbf{f}(x)) / \tau)}{\sum_{j=1}^C \exp(\cos(\mathbf{g}(o_j), \mathbf{f}(x)) / \tau)}, \quad (2)$$

where τ is a temperature parameter learned by CLIP and $\cos(\cdot, \cdot)$ represents cosine similarity.

3.4 Class-Specific Embedding Loss

We introduce the class-specific embedding (CSE) loss to optimize the trainable fine-grained instance prompts by learning from text embeddings of captions. It tries to minimize the cosine distance of matching patches and to increase the cosine distance of non-matching patches above the margin. Embedding loss is then computed as

$$\ell_{ebd} = \begin{cases} \Delta_i^k, & \text{if } \tilde{y}_i^k = 1, \\ \max(0, \mu - \Delta_i^k), & \text{if } \tilde{y}_i^k = -1, \end{cases} \quad (3)$$

$$\Delta_i^k = 1 - \cos(t_i^k, o_i|_m^M),$$

where μ is the margin factor. Intuitively the embedding loss penalizes positive (*i.e.*, prompts of matching classes) pairs that have large distances and negative (*i.e.*, prompts of non-matching classes) pairs that have small distance (less than μ).

LDAM (Cao et al., 2019) has inspired the development of a decision boundary that is both robust and generalizable, capable of accurately classifying features that vary within a certain range. However, when applied to long-tailed datasets characterized by a significant class imbalance, models tend to exhibit greater sensitivity to more frequent classes. As a result, the performance of these models in less frequent classes is often poor.

To address this issue, CSE loss employs the class-aware soft margin strategy to encourage the model to have the optimal trade-off between per-class margins by stimulating the minority classes to have larger margins, which can be viewed as

regularization (Wei et al., 2018). More specifically, as illustrated in Fig. 3, blue samples (head classes) are classified incorrectly, and the model update gradient is shown with pointed arrows. Green samples (medium classes) are classified correctly outside of the margin and the gradient is shown. Intuitively, the embedding loss does not give special consideration to the minority categories, but with the help of class-aware soft margin, the trade-off of μ_1 (in Fig. 3) can be optimized by shifting the decision boundary to encourage the tail classes to have larger margins. So yellow samples (tail classes) are classified correctly outside of the original margin but within the enlarged margin, and the embedding loss has no gradient for these samples. Following the trade-off between the class margins, we adopt a class-aware margin for multiple classes of the form

$$\tilde{\mu}_i \propto n_i^{-1/4} = \frac{\eta}{n_i^{1/4}}. \quad (4)$$

Here η is a hyper-parameter to be tuned. Therefore, when $y_i^k = -1$, the loss can be computed as $\max\{0, \tilde{\mu}_i - \Delta_i^k\}$.

Meanwhile, our loss can be combined with a re-weighting strategy to be more efficient when it comes to long-tailed distribution data. We then define the reference weight based on the empirical class frequencies $\{n_1, \dots, n_C\}$ on the training set:

$$w_i = \frac{(1/n_i)^\gamma}{\sum_{i=1}^C (1/n_i)^\gamma}, \quad (5)$$

where γ is a scale hyper-parameter to provide more flexibility. Hence, the re-weighted class-specific embedding loss is defined as:

$$\ell_{cse} = \begin{cases} w_i \Delta_i^k, & \text{if } \tilde{y}_i^k = 1, \\ \max\{0, w_i (\tilde{\mu}_i - \Delta_i^k)\}, & \text{if } \tilde{y}_i^k = -1, \end{cases} \quad (6)$$

$$\mathcal{L}_{cse} = \frac{\sum_{k=1}^N \ell_{cse}}{N}. \quad (7)$$

The overall process of class-specific embedding loss is outlined in Algorithm 1.

3.5 Multi-Label Classification Loss

Our method can be easily combined with the existing multi-label classification loss functions (Ridnik et al., 2021; Lin et al., 2017; Cui et al., 2019; Wu et al., 2020), regardless of whether they are designed for long-tailed distributions or not. By

Algorithm 1: Class-Specific Embedding Loss

Input: Text embeddings of textual descriptions (captions) t , labels \tilde{y} , prompt o
Output: Class-Specific Embedding Loss \mathcal{L}_{cse}

```

1 for  $k = 1, 2, \dots, N$  do
2    $\ell_{cse} = 0$ ;
3   for  $i = 1, 2, \dots, C$  do
4     Calculate class-aware soft margin  $\tilde{\mu}_i$  by Eq. 4;
5     Calculate weight  $w_i$  by Eq. 5;
6     Calculate  $\Delta_i^k = 1 - \cos(t_i^k, o_i^k)$ ;
7     if  $\tilde{y}_i^k = 1$  then
8        $\ell_{cse} = w_i \Delta_i^k$ ;
9     else
10       $\ell_{cse} = \text{ReLU}(w_i (\tilde{\mu}_i - \Delta_i^k))$ ;
11 Calculate  $\mathcal{L}_{cse}$  by Eq. 7.
```

blending the classification loss functions with our proposed CSE loss, our method facilitates prompt learning of more refined class descriptions and semantic relationships between categories, particularly between head and tail classes.

In this study, we introduce the distribution-balanced loss (Wu et al., 2020) as the classification loss function, which can be formulated as:

$$r = \alpha + \sigma \left(\beta \times \left(\frac{\frac{1}{n_i}}{\sum_{i=1}^C \frac{1}{n_i}} - \theta \right) \right), \quad (8)$$

$$v_i = -\kappa \times -\log \left(\frac{1}{n_i/N} - 1 \right), \quad (9)$$

$$\ell_{cls} = \begin{cases} -r (1 - q_i^k)^\gamma \log(q_i^k), & \text{if } y_i^k = 1, \\ -\frac{r}{\zeta} (q_i^k)^\gamma \log(1 - q_i^k), & \text{if } y_i^k = -1, \end{cases} \quad (10)$$

where $q_i^k = \sigma(z_i^k - v_i)$ is for positive instances, $q_i^k = \sigma(\zeta(z_i^k - v_i))$ is for negative ones and $\alpha, \beta, \theta, \kappa, \zeta$ are hyperparameters. Then $\mathcal{L}_{cls} = \sum_{k=1}^N \ell_{cls} / N$.

Hence, the overall training loss can be written as:

$$\mathcal{L} = \lambda \mathcal{L}_{cls} + (1 - \lambda) \mathcal{L}_{cse}, \quad (11)$$

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

4 Experiment

4.1 Benchmark Setting

Following (Wu et al., 2020; Guo and Wang, 2021), we conduct experiments on two datasets for long-tailed multi-label visual recognition: VOC-LT and COCO-LT (Wu et al., 2020). They are artificially

Datasets	VOC-LT				COCO-LT			
Methods	total	head	medium	tail	total	head	medium	tail
RN-50								
ERM	70.86	68.91	80.20	65.31	41.27	48.48	49.06	24.25
RW	74.70	67.58	82.81	73.96	42.27	48.62	45.80	32.02
Focal Loss (Lin et al., 2017) ICCV'17	73.88	69.41	81.43	71.56	49.46	49.80	54.77	42.14
RS (Shen et al., 2016) ECCV'16	75.38	70.95	82.94	73.05	46.97	47.58	50.55	41.70
ML-GCN (Chen et al., 2019b) CVPR'19	68.92	70.14	76.41	62.39	44.24	44.04	48.36	38.96
OLTR (Liu et al., 2019) CVPR'19	71.02	70.31	79.80	64.95	45.83	47.45	50.63	38.05
LDAM (Cao et al., 2019) NeurIPS'19	70.73	68.73	80.38	69.09	40.53	48.77	48.38	22.92
CB Focal (Cui et al., 2019) CVPR'19	75.24	70.30	83.53	72.74	49.06	47.91	53.01	44.85
BBN (Zhou et al., 2020) CVPR'20	73.37	71.31	81.76	68.62	50.00	49.79	53.99	44.91
DB Focal (Wu et al., 2020) ECCV'20	78.94	73.22	84.18	79.30	53.55	51.13	57.05	51.06
LTML (Guo and Wang, 2021) CVPR'21	81.44	75.68	85.53	82.69	56.90	54.13	60.59	54.47
<hr/>								
CLIP (Radford et al., 2021) ICML'21	84.30	63.60	88.03	97.03	56.19	35.73	60.52	68.45
CoOp (Zhou et al., 2022b) IJCV'22	81.34	65.10	81.54	93.37	54.94	38.06	56.67	67.51
CoCoOp (Zhou et al., 2022a) CVPR'22	78.63	64.33	80.51	87.94	46.02	36.02	50.57	48.82
DualCoOp (Sun et al., 2022) NeurIPS'22	81.03	66.45	80.53	92.33	53.11	40.48	55.20	62.11
TaI-DPT (Guo et al., 2023) CVPR'23	83.75	66.27	85.17	94.57	56.23	40.52	58.40	66.09
LMPT (ours)	85.44	66.62	88.11	97.86	58.97	41.87	61.60	69.60
ViT-B/16								
CLIP (Radford et al., 2021) ICML'21	85.77	66.52	88.93	97.83	60.17	38.52	65.06	72.28
CoOp (Zhou et al., 2022b) IJCV'22	86.02	67.71	88.79	97.67	60.68	41.97	63.18	73.85
CoCoOp (Zhou et al., 2022a) CVPR'22	84.47	64.58	87.82	96.88	61.49	39.81	64.63	76.42
LMPT (ours)	87.88	72.10	89.26	98.49	66.19	44.89	69.80	79.08

Table 1: mAP performance of the proposed method and comparison methods. Above the dotted line is the performance of image-only models and below is that of vision-language models.

sampled from two multi-label recognition benchmarks, PascalVOC (Everingham et al., 2015) and MS-COCO (Lin et al., 2014), respectively.

4.2 Experimental Settings

Metrics. As in (Liu et al., 2019), the classes are split into three groups by the number of their training examples: head classes each contain over 100 samples, medium classes each have between 20 and 100 samples, and tail classes with under 20 samples each. We use mean average precision (mAP) to evaluate the performance of long-tailed multi-label visual recognition for all the classes.

Implementation Details. We adopt CLIP ResNet-50 (He et al., 2016) or ViT-B/16 (Dosovitskiy et al., 2020) as the visual encoder and use the corresponding CLIP Transformer as the text encoder. During training, the parameters of both the two encoders are kept frozen, and only learnable prompts are optimized. SGD optimizer is adopted to learn prompt tokens, and the training epochs are set to 30. The learning rates for COCO-LT and VOC-LT are empirically initialized with $1e-4$, $5e-4$, and decay by the cosine annealing rule during training. For loss functions, η in Eq. 4, γ in Eq. 5 and λ in Eq. 11 are

set as 1.0, 1.0 and 0.5, respectively. Other hyper-parameters in DB loss are set as the same as (Wu et al., 2020).

4.3 Long-Tailed Multi-Label Visual Recognition

To evaluate the effectiveness of the proposed method, firstly we compare it with previous methods of image-only models on the two long-tailed multi-label datasets. The compared methods include Empirical Risk Minimization (ERM), a smooth version of Re-Weighting (RW) using the inverse proportion to the square root of class frequency, Re-Sampling (RS) (Shen et al., 2016), Focal Loss (Lin et al., 2017), ML-GCN (Chen et al., 2019b), OLTR (Liu et al., 2019), LDAM (Cao et al., 2019), Class-Balanced (CB) Focal (Cui et al., 2019), BBN (Zhou et al., 2020), Distribution-Balanced (DB) Focal (Wu et al., 2020) and LTML (Guo and Wang, 2021). The mAP performance of different methods is shown in Table 1. The prior best performance is achieved by LTML – mAP of 81.44% over all classes on VOC-LT and 56.90% over all classes on COCO-LT.

Furthermore, we compare zero-shot and prompt

Datasets	VOC-LT			
Methods	total	head	medium	tail
BCE	82.18	64.90	83.17	94.30
MLS	84.30	64.31	84.82	97.47
Focal Loss	85.37	66.17	87.70	97.52
CB Loss	85.25	65.37	87.71	97.20
R-BCE-Focal	84.56	66.01	86.61	97.67
ASL	86.40	69.12	88.79	98.07
DB Focal	87.88	72.10	89.26	98.49
Datasets	COCO-LT			
Methods	total	head	medium	tail
BCE	58.04	41.79	58.86	73.90
MLS	61.26	41.71	64.11	74.58
Focal Loss	54.40	37.60	59.36	62.33
CB Loss	56.45	34.61	58.77	74.52
R-BCE-Focal	60.13	38.11	64.87	72.79
ASL	64.89	43.18	68.22	78.43
DB Focal	66.19	44.89	69.80	79.08

Table 2: mAP performance of the proposed method with different multi-label loss functions.

learning methods based on CLIP on the two benchmarks. The mAP performance of these methods is shown in Table 1 as well. For a fair comparison, we initialize the prompt as the default hand-crafted one “a photo of a” for all the methods. The results show that when using ViT-B/16 as the backbone, even the overall mAP performance of zero-shot CLIP reaches 85.77% and 60.17%, which outperforms previous SOTA LTML by **4.33** points (85.77% vs.81.44%) and **3.27** points (60.17% vs.56.90%) on the two datasets, respectively. Therefore, it is meaningful to explore how to use prompt tuning based on CLIP effectively for better performance. From the perspective of prompt tuning methods, when using ResNet-50 as the backbone, the performance of our method on VOC-LT is more promising, which is **4.1** points, **6.81** points, **4.41** points and **1.69** points better than CoOp, CoCoOp, DualCoOp and TaI-DPT, which are popular prompt learning methods for single-label and multi-label recognition. The performance on COCO-LT is similar to that on VOC-LT, which is **4.03** points, **12.95** points, and **5.86** points better than CoOp, CoCoOp, and DualCoOp. When replacing the backbone with ViT-B/16, the overall mAP performance of our method can further boost up to 87.88% and 66.19% on VOC-LT and COCO-LT, which is the current new **state-of-the-art** of the two datasets.

4.4 Ablation Analysis

Components Analysis. To further analyze which component makes our methods performant for LTML, we conduct a set of ablation studies and report the results in Table 3. We first conduct exper-

iments with CLIP and the mAP performances are 85.77% on VOC-LT, 60.17% on COCO-LT, which surprisingly outperforms the prior SOTA LTML. It indicates that pre-trained VLMs demonstrate a robust capability for visual recognition, providing a solid foundation for our approach. However, the mAP performance of the tail classes outperforms the head classes by nearly 30 points on both VOC-LT and COCO-LT. Then CoOp is benefited from soft prompts and the mAP performance is improved to 86.02% on VOC-LT and 60.68% on COCO-LT, with **0.25%** and **0.51%** increments. Besides, we design the class-specific embedding loss with class-aware soft margin and re-weighting to learn more fine-grained and class-related prompts that build semantic relationships across different classes, especially for the tail classes by encouraging those classes to have larger margins and weights. The mAP performances of head, medium, and tail classes after adding the embedding loss are all significantly improved and the overall mAP surpasses CoOp by **1.26%** and **4.66%** on VOC-LT and COCO-LT, which demonstrates our embedding loss can help prompts learn fine-grained classes descriptions and semantic relationships across the classes. Finally, the integration of CASM and RW strategy further improves the mAP performance slightly, mainly for the tail performance by **0.65%** and **1.12%** on VOC-LT and COCO-LT.

Multi-Label Classification Loss Functions. We compare a number of multi-label classification loss functions, including Binary Cross-Entropy Loss (BCE), Multi-Label Soft Margin Loss (MSL), Focal Loss, CB Loss, R-BCE-Focal, Asymmetric Loss (ASL) and DB Focal. As illustrated in Table 2, DB Focal loss that takes the co-occurrence of labels and the dominance of negative labels into account works significantly better than other multi-label classification loss for the LTML task.

Effectiveness of Text Supervision. We further compare our method with fine-tuning CLIP’s image encoder when using ResNet-50 as the backbone to explore whether the significant effect of our approach is due to text supervision or simply because the CLIP’s image encoder is so powerful. In order to prevent interference with the trained CLIP’s image encoder during the fine-tuning phase, we only fine-tune a fully connected layer added at the end of the image encoder. The results are shown in Fig. 4. Obviously, fine-tuning the image encoder shows promising results, but still largely

Soft Prompt	Embedding Loss	Class-Aware Soft Margin	Re-weighting	VOC-LT					COCO-LT				
				total	head	medium	tail	avg.Δ	total	head	medium	tail	avg.Δ
				85.77	66.52	88.93	97.83		60.17	38.52	65.06	72.28	
✓				86.02	67.71	88.79	97.67	+0.29	60.68	41.97	63.18	73.85	+0.91
✓	✓			87.28	71.07	89.01	97.84	+0.51	65.34	44.27	69.39	77.96	+5.23
✓	✓	✓		87.62	72.01	89.26	98.13	+1.99	65.81	44.90	69.71	78.76	+5.79
✓	✓	✓	✓	87.88	72.10	89.26	98.49	+2.17	66.19	44.89	69.80	79.08	+5.98

Table 3: Ablation analysis on different components of the our method. “avg.Δ” average performance improvement.

underperforms LMPT, which suggests that the gradients that went through the text encoder provide more useful information.




Figure 4: mAP performance of different methods w/o text supervision on two datasets. (a) VOC-LT. (b) COCO-LT.

4.5 Case Analysis

To better understand how our method deals with long-tailed multi-label data, we performed qualitative experiments with ResNet, CLIP, and ours on COCO-LT and VOC-LT. Fig. 5 shows several cases where the model justifies its abilities for the prediction. For example, in the third column, ResNet only recognizes [person] (belongs to head classes) and fails to classify the image to [train] (belongs to tail classes), which is a pervasive challenge encountered by image-only models. The emergence of CLIP is a great remedy for this issue, owing to its huge training data and effective text supervision. Nevertheless, simple hand-crafted templates as prompts still cannot accurately identify categories as they cannot describe the characteristics of each category. Understanding the inter-class relationships, particularly among head and tail categories, presents a formidable challenge in multi-label visual recognition, which is essential for achieving optimal performance in this domain. With the aid of our approach, utilizing prompts that learn from a large corpus of image-caption data, it has become feasible to discern the semantic relationships between categories and accurately predict the relevant categories of simple objects, even in challenging scenarios such as identifying [stop sign] from images. Therefore, our proposed method demonstrates significant advantages in effectively address-

		VOC-LT		COCO-LT	
Input Image	GT	Our Prediction	CLIP's Prediction	ResNet's Prediction	
	person car dog horse	horse: 92.43 person: 4.22 car: 1.35 dog: 0.53 cow: 0.37	horse: 87.99 car: 3.52 bus: 2.49 bicycle: 2.10 person: 1.47	horse: 99.99 dog: 1e-4 aeroplane: 7e-7 bird: 3e-9 diningtable: 8e-11	
	sheep bird	sheep: 99.90 bird: 0.06 dog: 0.03 cow: 0.02 horse: 0.01	sheep: 99.90 cow: 0.06 horse: 0.03 dog: 0.01 person: 0.01	cow: 99.99 chair: 2e-3 bottle: 4e-4 bird: 3e-6 aeroplane: 2e-6	
	person stop sign train	train: 90.17 stop sign: 2.81 person: 1.56 bus: 0.06 book: 0.02	train: 99.36 remote: 0.18 bus: 0.13 car: 0.06 laptop: 0.04	traffic light: 99.90 orange: 0.11 bus: 7e-4 car: 2e-4 person: 4e-5	
	person kite chair	person: 84.77 kite: 1.85 bottle: 0.31 bird: 0.23	person: 45.70 chair: 36.16 bird: 9.14 bottle: 2.24 bicycle: 1.22	donut: 98.24 tv: 1.79 wine glass: 0.03 traffic light: 0.01 person: 7e-4	

Figure 5: Example decisions from our model, CLIP, and ResNet.

ing the intricate relationship among multiple labels and the long-tailed problem with the aid of text supervision.

5 Conclusion

In this work, we propose a new view of prompt tuning for long-tailed multi-label visual recognition by learning class-specific contexts from the alignment of prompts and textual description (caption), which complements more fine-grained features and builds semantic relationships across head and tail classes. Considering the class imbalance, a novel class-specific embedding loss with the class-aware soft margin and re-weighting strategy is introduced to promote increased generalization among the tail classes. Furthermore, we integrate a distribution-balanced loss as the classification loss function in consideration of its empirical efficacy compared to alternative loss functions. Our method exhibits significant improvement over the previous state-of-the-art (SOTA) and zero-shot CLIP on VOC-LT and COCO-LT. Additionally, We hope our approach will inspire future work in this field.

Acknowledgement

We thank Monash-Airdoc Research donation Phase 2, MRFF NCRI GA89126, and MRFF AI-Eye NCRI GA89126 for their grant.

References

- Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. 2018. A systematic study of the class imbalance problem in convolutional neural networks. Neural networks, 106:249–259.
- Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. Advances in neural information processing systems, 32.
- Tianshui Chen, Muxin Xu, Xiaolu Hui, Hefeng Wu, and Liang Lin. 2019a. Learning semantic-specific graph representation for multi-label image recognition. In Proceedings of the IEEE/CVF international conference on computer vision, pages 522–531.
- Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. 2019b. Multi-label image recognition with graph convolutional networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 5177–5186.
- Jiequan Cui, Shu Liu, Zhuotao Tian, Zhisheng Zhong, and Jiaya Jia. 2022. Reslt: Residual learning for long-tailed recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence.
- Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. 2019. Class-balanced loss based on effective number of samples. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 9268–9277.
- Qi Dong, Shaogang Gong, and Xiatian Zhu. 2017. Class rectification hard mining for imbalanced deep learning. In Proceedings of the IEEE International Conference on Computer Vision, pages 1851–1860.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.
- Chris Drummond, Robert C Holte, et al. 2003. Class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In Workshop on learning from imbalanced datasets II, volume 11, pages 1–8.
- Mark Everingham, SM Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. 2015. The pascal visual object classes challenge: A retrospective. International journal of computer vision, 111(1):98–136.
- Hao Guo and Song Wang. 2021. 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, pages 15089–15098.
- Zixian Guo, Bowen Dong, Zhilong Ji, Jinfeng Bai, Yiwen Guo, and Wangmeng Zuo. 2023. Texts as images in prompt tuning for multi-label image recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2808–2817.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778.
- Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. 2016. Learning deep representation for imbalanced classification. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5375–5384.
- Xinyu Huang, Youcai Zhang, Ying Cheng, Weiwei Tian, Ruiwei Zhao, Rui Feng, Yuejie Zhang, Yaqian Li, Yandong Guo, and Xiaobo Zhang. 2022. Idea: Increasing text diversity via online multi-label recognition for vision-language pre-training. In Proceedings of the 30th ACM International Conference on Multimedia, pages 4573–4583.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In International Conference on Machine Learning, pages 4904–4916. PMLR.
- Lie Ju, Zhen Yu, Lin Wang, Xin Zhao, Xin Wang, Paul Bonnington, and Zongyuan Ge. 2023. Hierarchical knowledge guided learning for real-world retinal disease recognition. IEEE Transactions on Medical Imaging.
- Bingyi Kang, Yu Li, Sa Xie, Zehuan Yuan, and Jiashi Feng. 2020. Exploring balanced feature spaces for representation learning. In International Conference on Learning Representations.
- Salman H Khan, Munawar Hayat, Mohammed Benamoun, Ferdous A Sohel, and Roberto Togneri. 2017. Cost-sensitive learning of deep feature representations from imbalanced data. IEEE transactions on neural networks and learning systems, 29(8):3573–3587.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In Proceedings of the IEEE international conference on computer vision, pages 2980–2988.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco:

- Common objects in context. In European conference on computer vision, pages 740–755. Springer.
- Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu. 2021. Query2label: A simple transformer way to multi-label classification. arXiv preprint arXiv:2107.10834.
- Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. 2019. Large-scale long-tailed recognition in an open world. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2537–2546.
- Teli Ma, Shijie Geng, Mengmeng Wang, Jing Shao, Jiasen Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2021. A simple long-tailed recognition baseline via vision-language model. arXiv preprint arXiv:2111.14745.
- Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. 2020. Long-tail learning via logit adjustment. arXiv preprint arXiv:2007.07314.
- Wanli Ouyang, Xiaogang Wang, Cong Zhang, and Xiaokang Yang. 2016. Factors in finetuning deep model for object detection with long-tail distribution. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 864–873.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In International Conference on Machine Learning, pages 8748–8763. PMLR.
- Tal Ridnik, Emanuel Ben-Baruch, Nadav Zamir, Asaf Noy, Itamar Friedman, Matan Protter, and Lih Zelnik-Manor. 2021. Asymmetric loss for multi-label classification. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 82–91.
- Tal Ridnik, Gilad Sharir, Avi Ben-Cohen, Emanuel Ben-Baruch, and Asaf Noy. 2023. MI-decoder: Scalable and versatile classification head. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 32–41.
- Dvir Samuel and Gal Chechik. 2021. Distributional robustness loss for long-tail learning. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 9495–9504.
- Timo Schick and Hinrich Schütze. 2021. Exploiting cloze-questions for few-shot text classification and natural language inference. In Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, pages 255–269.
- Li Shen, Zhouchen Lin, and Qingming Huang. 2016. Relay backpropagation for effective learning of deep convolutional neural networks. In European conference on computer vision, pages 467–482. Springer.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. arXiv preprint arXiv:2010.15980.
- Ximeng Sun, Ping Hu, and Kate Saenko. 2022. Dualcoop: Fast adaptation to multi-label recognition with limited annotations. arXiv preprint arXiv:2206.09541.
- Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, and Junjie Yan. 2020. Equalization loss for long-tailed object recognition. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 11662–11671.
- Changyao Tian, Wenhai Wang, Xizhou Zhu, Jifeng Dai, and Yu Qiao. 2022. VI-ltr: Learning class-wise visual-linguistic representation for long-tailed visual recognition. In European Conference on Computer Vision, pages 73–91. Springer.
- Grigorios Tsoumakas and Ioannis Katakis. 2007. Multi-label classification: An overview. International Journal of Data Warehousing and Mining (IJDWM), 3(3):1–13.
- Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. 2016. Cnn-rnn: A unified framework for multi-label image classification. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2285–2294.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. arXiv preprint arXiv:2202.03052.
- Zhouxia Wang, Tianshui Chen, Guanbin Li, Ruijia Xu, and Liang Lin. 2017. Multi-label image recognition by recurrently discovering attentional regions. In Proceedings of the IEEE international conference on computer vision, pages 464–472.
- Colin Wei, Jason Lee, Qiang Liu, and Tengyu Ma. 2018. On the margin theory of feedforward neural networks.
- Tong Wu, Qingqiu Huang, Ziwei Liu, Yu Wang, and Dahua Lin. 2020. Distribution-balanced loss for multi-label classification in long-tailed datasets. In European Conference on Computer Vision, pages 162–178. Springer.

- Peng Xia, Ze Chen, Juanxi Tian, Yangrui Gong, Ruibo Hou, Yue Xu, Zhenbang Wu, Zhiyuan Fan, Yiyang Zhou, Kangyu Zhu, et al. 2024a. Cares: A comprehensive benchmark of trustworthiness in medical vision language models. arXiv preprint arXiv:2406.06007.
- Peng Xia, Ming Hu, Feilong Tang, Wenxue Li, Wenhao Zheng, Lie Ju, Peibo Duan, Huaxiu Yao, and Zongyuan Ge. 2024b. Generalizing to unseen domains in diabetic retinopathy with disentangled representations. arXiv preprint arXiv:2406.06384.
- Peng Xia, Xingtong Yu, Ming Hu, Lie Ju, Zhiyong Wang, Peibo Duan, and Zongyuan Ge. 2023. Hg-clip: Exploring vision-language models with graph representations for hierarchical understanding. arXiv preprint arXiv:2311.14064.
- Shichao Xu, Yikang Li, Jenhao Hsiao, Chiuman Ho, and Zhu Qi. 2022. A dual modality approach for (zero-shot) multi-label classification. arXiv preprint arXiv:2208.09562.
- Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. 2021. Cpt: Colorful prompt tuning for pre-trained vision-language models. arXiv preprint arXiv:2109.11797.
- Jin Ye, Junjun He, Xiaojiang Peng, Wenhao Wu, and Yu Qiao. 2020. Attention-driven dynamic graph convolutional network for multi-label image recognition. In European conference on computer vision, pages 649–665. Springer.
- Min-Ling Zhang and Zhi-Hua Zhou. 2007. MI-knn: A lazy learning approach to multi-label learning. Pattern recognition, 40(7):2038–2048.
- Min-Ling Zhang and Zhi-Hua Zhou. 2013. A review on multi-label learning algorithms. IEEE transactions on knowledge and data engineering, 26(8):1819–1837.
- Yifan Zhang, Bryan Hooi, Lanqing Hong, and Jiashi Feng. 2021a. Test-agnostic long-tailed recognition by test-time aggregating diverse experts with self-supervision. arXiv preprint arXiv:2107.09249.
- Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. 2021b. Deep long-tailed learning: A survey. arXiv preprint arXiv:2110.04596.
- Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. 2020. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 9719–9728.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022a. Conditional prompt learning for vision-language models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 16816–16825.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022b. Learning to prompt for vision-language models. International Journal of Computer Vision, 130(9):2337–2348.
- Linchao Zhu and Yi Yang. 2020. Inflated episodic memory with region self-attention for long-tailed visual recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4344–4353.