

# Open-Vocabulary Multi-Label Classification via Multi-Modal Knowledge Transfer

Sunan He<sup>1,2,3,\*†</sup>, Taian Guo<sup>3,\*</sup>, Tao Dai<sup>1,‡</sup>, Ruizhi Qiao<sup>3</sup>, Xiujun Shu<sup>3</sup>, Bo Ren<sup>3</sup>, Shu-Tao Xia<sup>2,4</sup>

<sup>1</sup> College of Computer Science and Software Engineering, Shenzhen University

<sup>2</sup> Tsinghua Shenzhen International Graduate School, Tsinghua University

<sup>3</sup> YouTu Lab, Tencent

<sup>4</sup> Research Center of Artificial Intelligence, Peng Cheng Laboratory

hsn20@mails.tsinghua.edu.cn, {taianguo, ruizhiqiao, xiujunshu, timren}@tencent.com

daitao.edu@gmail.com, xiaast@sz.tsinghua.edu.cn

## Abstract

Real-world recognition system often encounters the challenge of unseen labels. To identify such unseen labels, multi-label zero-shot learning (ML-ZSL) focuses on transferring knowledge by a pre-trained textual label embedding (e.g., GloVe). However, such methods only exploit *single-modal* knowledge from a language model, while ignoring the rich semantic information inherent in image-text pairs. Instead, recently developed open-vocabulary (OV) based methods succeed in exploiting such information of image-text pairs in object detection, and achieve impressive performance. Inspired by the success of OV-based methods, we propose a novel open-vocabulary framework, named multi-modal knowledge transfer (MKT), for multi-label classification. Specifically, our method exploits *multi-modal* knowledge of image-text pairs based on a vision and language pre-training (VLP) model. To facilitate transferring the image-text matching ability of VLP model, knowledge distillation is employed to guarantee the consistency of image and label embeddings, along with prompt tuning to further update the label embeddings. To further enable the recognition of multiple objects, a simple but effective two-stream module is developed to capture both local and global features. Extensive experimental results show that our method significantly outperforms state-of-the-art methods on public benchmark datasets.

## Introduction

Multi-label recognition, which aims to recognize all the relevant labels in an image, is a fundamental task in computer vision applications, such as scene understanding, surveillance systems and self-driving cars. In real-world applications, multi-label recognition systems should learn tens of thousands of labels, locate them in images, and even deal with many unseen labels. To date, classic multi-label classification methods trained and tested with seen labels are far from fulfilling the requirements for real applications, where plenty of unseen labels exist.

\*These authors contributed equally.

†This work was done during an internship at Tencent.

‡Corresponding author: Tao Dai

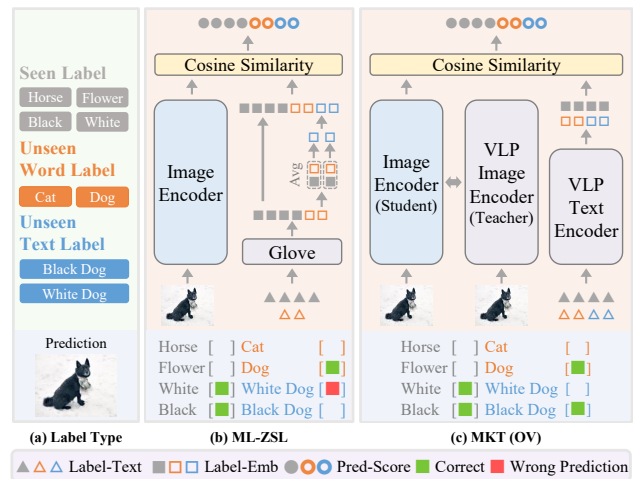


Figure 1: The overall framework of the classic multi-label zero-shot learning (ML-ZSL), and our multi-modal knowledge transfer (MKT) method. (b) ML-ZSL only exploits single-modal knowledge of language-based models (e.g., GloVe), and may fail to recognize unseen text labels (e.g., ‘Black Dog’). (c) Instead, our MKT succeeds in predicting it by jointly exploring multi-modal knowledge of vision and language pre-training (VLP) models. (Best viewed in color.)

To identify the unseen labels in an image, many multi-label zero-shot learning (ML-ZSL) methods (Huynh and Elhamifar 2020; Gupta et al. 2021; Ben-Cohen et al. 2021; Narayan et al. 2021) have been recently developed by transferring knowledge between seen and unseen labels. However, most existing methods (Zhang, Gong, and Shah 2016; Huynh and Elhamifar 2020; Gupta et al. 2021; Ben-Cohen et al. 2021; Narayan et al. 2021) contain two main issues. **First**, these methods solely exploit *single-modal* knowledge by a pre-trained textual label embeddings like GloVe (Pennington, Socher, and Manning 2014) (as shown in Figure 1 (b)), while ignoring the visual semantic image-text pair information. **Second**, although such textual label embeddings (e.g., GloVe) handle word labels (e.g., label of ‘cat’) well, they cannot be easily extended to text labels (e.g., label of

‘black cat’), thus hindering the flexibility of the models. As shown in Figure 1, ML-ZSL fails to recognize the unseen text label of ‘black dog’, while our MKT succeeds in predicting this label by jointly exploring *multi-modal* knowledge of vision and language models.

To explore such multi-modal knowledge, recently developed open-vocabulary (OV) methods (Gu et al. 2022; Huynh et al. 2022; Ghiasi et al. 2022; Du et al. 2022; Ma et al. 2022) have been proposed based on vision and language pre-training (VLP) models. Such OV-based methods trained on billions of image-text pairs contain powerful image-text matching ability, and have achieved remarkable performance in computer vision tasks like object detection. However, how to extend such OV-based methods to multi-label classification, including unseen text labels, is less explored.

Motivated by the above observations, we propose a novel open-vocabulary framework, named multi-modal knowledge transfer (MKT), for multi-label classification. Unlike the previous ML-ZSL methods that exploit only language-based information, our MKT utilizes *multi-modal* knowledge from image-text pairs from a vision and language pre-training (VLP) model. As shown in Figure 1(c), our MKT mainly consists of an image encoder to extract image features, and a VLP image/text encoder to extract image/label embeddings. Specifically, to facilitate transferring the image-text matching ability of VLP models, knowledge distillation and prompt tuning are introduced to guarantee the consistency of image and label embeddings. In practice, knowledge distillation makes image embeddings align better with its relevant label embeddings, while prompt tuning adapts the label embeddings to better support classification task. Besides, to further improve the ability of feature expressions, we propose a simple but effective two-stream feature extraction module to capture both local and global features to extract more discriminative features. In this way, our MKT framework can capture the rich semantic information inherent in image-text pairs of VLP models.

The main contributions can be summarized as follows:

1. We propose an open-vocabulary based multi-modal knowledge transfer (MKT) framework for multi-label classification, which exploits the semantic multi-modal information in image-text pairs based on VLP models. To the best of our knowledge, this is the first work to explore open-vocabulary multi-label classification task.
2. Our MKT framework mainly consists of an image encoder to extract image features, and a VLP image/text encoder to extract image/label embeddings. To guarantee the consistency of image and label embeddings, a knowledge distillation strategy is incorporated into our MKT framework, along with prompt tuning to update the label embeddings iteratively. Besides, to further improve the ability of feature expressions of our method, we propose a two-stream feature extraction module by jointly capturing local and global features.
3. Extensive results show that our MKT method significantly outperforms the previous ML-ZSL methods and establishes a new state of the art for open-vocabulary multi-label classification on two large-scale benchmark

datasets, namely NUS-WIDE and Open Images.

## Related Works

### Multi-Label Zero-Shot Learning

The goal of standard multi-label classification task is to predict a set of labels in an image. A vanilla approach is to train a binary classifier for each label present in the training dataset without considering the dependence among the labels (Tsoumakas and Katakis 2007; Read et al. 2011). To capture the label correlation, structure learning (Gong et al. 2014; Wang et al. 2016; Zhu et al. 2017; Wang et al. 2017) and graph methods (Li et al. 2016; Lee et al. 2018; Chen et al. 2019) are introduced in this task. Recently, vision transformer based methods have received much attention due to the powerful ability of capturing the global dependency (Lanchantin et al. 2021; Cheng et al. 2022). Although these methods have achieved promising results in multi-label classification, they cannot handle unseen labels, thus limiting their real applications.

To identify the unseen labels, zero-shot learning (ZSL) usually utilizes semantic information like attributes or word embeddings (Mikolov et al. 2013; Xian, Schiele, and Akata 2017). In particular, Lampert *et al.* (Lampert, Nickisch, and Harmeling 2009) proposed two attribute-based paradigms with direct attribute prediction (DAP) and indirect attribute prediction (IAP). The former aims to learn multiple attribute classifiers (Lampert, Nickisch, and Harmeling 2014), while the latter uses seen class proportions for prediction (Zhang and Saligrama 2015). While they can recognize to single unseen label, they cannot handle multi-label problem.

As an extension of ZSL, multi-label zero-shot learning (ML-ZSL) is developed to identify multiple seen and unseen labels in an image. The keys to this task are the alignment of image embeddings with its relevant label embeddings and the relation between seen and unseen label embeddings. To this end, Fast0Tag (Zhang, Gong, and Shah 2016) and ZS-SDL (Ben-Cohen et al. 2021) aim to find principal directions of an image along which the relevant labels rank higher. LESA (Huynh and Elhamifar 2020) and BiAM (Narayan et al. 2021) introduce attention module to capture both local and global features for better recognition of multiple objects. On the other hand, GAN-MLZSL (Gupta et al. 2021) introduces generative adversarial networks (GANs) to tackle the problem of multi-label feature synthesis from corresponding multi-label class embedding.

However, most existing ML-ZSL works exploit only single-modal knowledge via a language model (*e.g.*, GloVe). Due to the lack of visual information, these language-based models cannot capture visual consistency among labels, thus limiting the generalization ability. By contrast, we attempt to explore multi-modal knowledge from VLP models to leverage the consistency of image and label embeddings and can handle multiple word and text unseen labels.

### Open-Vocabulary Classification

With recent great development in vision and language pre-training model, open-vocabulary classification emerges as an alternative way to predict arbitrary labels. Large-scale

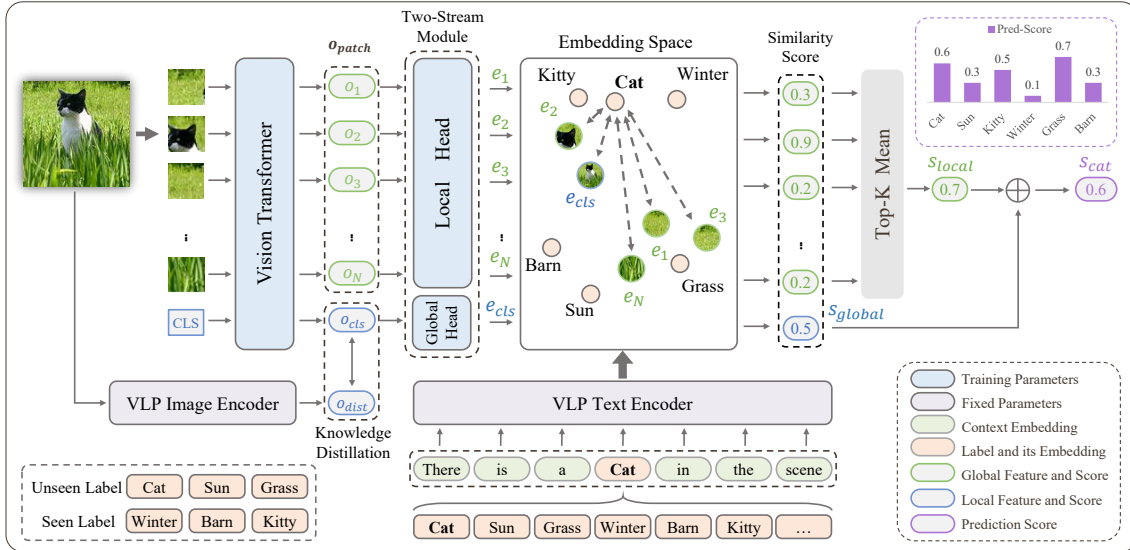


Figure 2: The overall framework of our multi-modal knowledge transfer (MKT) model for open-vocabulary multi-label classification. Our MKT mainly consists of a vision and language pre-training (VLP) model and a vision transformer model. The VLP model aims to extract multi-modal knowledge of input image-text pairs, while vision transformer is used to extract semantic features of input images. Moreover, knowledge distillation is used to guarantee the consistency of image and its relevant label embeddings, along with prompt tuning to further update the label embeddings. (Best viewed in color.)

pre-trained models first become prevalent in natural language processing (NLP), such as BERT (Devlin et al. 2018) and GPT2 (Radford et al. 2019). Based on large-scale language corpus (Raffel et al. 2020) and multiple task-agnostic pre-training objectives (Devlin et al. 2018), these pre-trained models achieve promising results in downstream tasks. Recently, Vision and Language Pre-training (VLP) models (Lu et al. 2019; Chen et al. 2020; Li et al. 2020; Li et al. 2020; Kim, Son, and Kim 2021) have received much attention in multi-modal tasks. For example, with billions of image-text pairs as training samples, CLIP (Radford et al. 2021) and ALIGN (Jia et al. 2021) have achieved impressive performance in image-text matching task. By transferring this matching ability to the classification task, we can achieve arbitrary text label prediction. Specifically, for any concept, we can generate its label embedding through the text encoder of VLP model and calculate its similarity to image embedding for classification. Due to the large scale training corpus, we can excavate label embedding of an unbounded vocabulary and achieve open-vocabulary (OV) classification.

Some works have explored the OV classification in object detection (Zareian et al. 2021; Gu et al. 2022; Du et al. 2022; Ma et al. 2022; Zang et al. 2022) and image segmentation (Huynh et al. 2022; Ghiasi et al. 2022). They usually replace the classification head with label embeddings and achieve impressive performance in arbitrary text concept recognition. Moreover, to boost the classification ability, knowledge distillation (Hinton et al. 2015) and prompt tuning (Li and Liang 2021) are introduced to facilitate transferring the image-text matching ability (Zhou et al. 2022).

However, most existing OV works focus on single label classification task. Multi-label classification is more practi-

cal and challenging because the models need to recognize multiple objects and cannot be trained with contrastive loss directly. In this work, we first explore the multi-label open-vocabulary classification task and propose a novel multi-modal knowledge transfer (MKT) framework by jointly exploiting multi-modal knowledge of the image-text pairs based on vision and language pre-training models.

## Multi-modal Knowledge Transfer

### Preliminary

Similar to the ML-ZSL problem, suppose we have two disjoint label sets  $Y^S$  and  $Y^U$ , where  $Y^S$  denotes seen labels present in the training set and  $Y^U$  denotes unseen labels without training images. Let  $(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N)$  be  $N$  training sample, where  $\mathbf{x}_i$  denotes the  $i$ -th training samples and  $\mathbf{y}_i \in Y^S$  denotes the labels present in the image. In the standard zero-shot learning (ZSL) task, the goal is to learn a classifier  $f_{ZSL} : X \rightarrow Y^U$  to identify the relevant unseen labels for a given image. Note that in a more challenging and realistic setup of generalized zero-shot learning (GZSL) task, the classifier needs to identify both seen and unseen labels present in the test image, i.e.,  $f_{GZSL} : X \rightarrow Y^U \cup Y^S$ .

### The Overall Framework

As illustrated in Figure 2, we show the overall architecture of our multi-modal knowledge transfer (MKT) method, which mainly consists of a vision transformer and a vision and language pre-training (VLP) model. Specifically, We utilize the vision transformer (Dosovitskiy et al. 2021) as our backbone network to extract semantic features from input images. Due to its powerful ability in learning visual representations, we

choose CLIP (Radford et al. 2021) as our VLP model to extract semantic multi-modal knowledge from both VLP image and text encoders. Concretely, the label embedding is first generated based on the VLP text encoder, followed by further updates through prompt tuning. Moreover, knowledge distillation is introduced to facilitate the alignment between image embeddings and its relevant labels.

### Vision Transformer with Two-Stream Module

Denote an input image as  $\mathbf{x} \in \mathbb{R}^{C \times H \times W}$ , where  $H \times W$  is the size of the image and  $C$  is the number of channels. Following (Dosovitskiy et al. 2021), we reshape it into a sequence of flattened 2D patches  $\mathbf{x}_{patch} \in \mathbb{R}^{N \times (P^2 \cdot C)}$ , where  $P$  denotes the size of each patch and the total number of patches is  $N = HW/P^2$ . Followed by a trainable linear projection,  $\mathbf{x}_{patch}$  is mapped into  $\bar{\mathbf{x}}_{patch} \in \mathbb{R}^{N \times D}$ , where  $D$  is input embedding dimension. Then the processing of the  $k$ -th block in vision transformer is formulated as

$$\begin{aligned} \mathbf{x}_0 &= [\mathbf{E}_{cls}, \bar{\mathbf{x}}_{patch}] + \mathbf{E}_{pos}, \\ \mathbf{y}_k &= \mathbf{x}_{k-1} + \text{MSA}(\text{NORM}(\mathbf{x}_{k-1})), \\ \mathbf{x}_k &= \mathbf{y}_k + \text{MLP}(\text{NORM}(\mathbf{y}_k)), \end{aligned} \quad (1)$$

where  $\mathbf{E}_{cls}$  is the class token embedding and  $\mathbf{E}_{pos}$  is the position embedding.  $[\cdot, \cdot]$  means concatenation.  $\text{MLP}(\cdot)$ ,  $\text{NORM}(\cdot)$ ,  $\text{MSA}(\cdot)$  denote multilayer perceptron, norm layer and multi-head self-attention, respectively.

Denote the output of vision transformer as  $\mathbf{x}_L = [\mathbf{o}_{cls}, \mathbf{o}_{patch}]$ , where  $\mathbf{o}_{cls}$  and  $\mathbf{o}_{patch}$  correspond to the output of class and patch tokens, respectively.  $\mathbf{o}_{cls}$  represents the global feature and  $\mathbf{o}_{patch}$  denotes the local features.

To identify multiple labels in an image, we propose a simple two-stream module consisting of local head  $\Theta_L(\cdot)$  and global head  $\Theta_G(\cdot)$ , mapping local and global features into embedding space respectively,

$$\mathbf{e}_{cls} = \Theta_G(\mathbf{o}_{cls}), \mathbf{e}_{patch} = \Theta_L(\mathbf{o}_{patch}), \quad (2)$$

where  $\mathbf{e}_{patch} = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_N]$  and  $\mathbf{e}_{cls}$  are local and global feature embeddings respectively.

Then, final prediction score is formulated as

$$s_i = \langle \mathbf{z}_i, \mathbf{e}_{cls} \rangle + \text{TopK}([\langle \mathbf{z}_i, \mathbf{e}_1 \rangle, \langle \mathbf{z}_i, \mathbf{e}_2 \rangle, \dots, \langle \mathbf{z}_i, \mathbf{e}_N \rangle]), \quad (3)$$

where  $\mathbf{z}_i \in \mathbb{R}^{1 \times D_e}$  is a label embedding and  $\text{TopK}(\cdot)$  is the  $top-k$  mean pooling.  $\langle \cdot, \cdot \rangle$  denotes inner product.

The ranking loss  $\mathcal{L}_{rank}$  on prediction scores are used to train the network:

$$\mathcal{L}_{rank} \triangleq \sum_i \sum_{p \in \mathbf{y}_i, m \notin \mathbf{y}_i} \max(1 + s_i^p - s_i^m, 0), \quad (4)$$

where  $\mathbf{y}_i \in Y^S$  is the target labels of an image  $i$ .  $s_i^m$  and  $s_i^p$  denote the scores of negative and positive labels.

### Knowledge Distillation for Alignment

As a key point to generalize to unseen labels, the alignment of an image embedding with its associated seen label embeddings plays a critical role in open-vocabulary classification. We take CLIP (Radford et al. 2021) as our VLP model, consisting of an image encoder and a text encoder. Considering

that the pre-training task of CLIP is to match the paired image and text, the image embedding generated by the CLIP image encoder should be similar to its relevant label embeddings generated by the CLIP text encoder. Thus, we introduce knowledge distillation to facilitate the alignment between the embeddings of an image and its relevant labels.

Denote the teacher model (*i.e.*, CLIP image encoder) as  $\Phi_I^{CLIP}(\cdot)$ , then the process of distillation is formulated as

$$\mathcal{L}_{dist} \triangleq \|\Phi_I^{CLIP}(\mathbf{x}) - \mathbf{o}_{cls}\|_1 = \|\mathbf{o}_{dist} - \mathbf{o}_{cls}\|_1, \quad (5)$$

where  $\mathbf{x}$  is an image input,  $\mathbf{o}_{cls}$  is the global features generated by the student model (*i.e.*, our vision backbone), and  $\mathbf{o}_{dist}$  denotes the output of CLIP image encoder. The reason for distillation on the global features instead of the local is twofold. First, both  $\mathbf{o}_{cls}$  and the output of CLIP image encoder are corresponded to the CLS token. Moreover, the local features  $\mathbf{o}_{patch}$  corresponding to different input patches are expected to be discriminative instead of identical in order to facilitate the recognition of multiple objects.

### Prompt Tuning for Label Embedding

Following (Radford et al. 2021), we first design a manual prompt template as ‘‘There is a  $\{label\}$  in the scene’’. We fill up the blank in this template with label name and treat the whole sentence as the input of CLIP text encoder. The output of CLIP text encoder is utilized as the label embedding. Due to the different training objectives, we argue that the label embeddings generated by pre-trained CLIP text encoder are not optimal for multi-label classification. Thus, we propose to further fine-tune the label embedding. However, it is very hard to fine-tune the entire text encoder due to the mode collapse problem caused by insufficient training samples. Motivated by CoOp (Zhou et al. 2022), we introduce prompt tuning for the adaptation of label embedding. During the tuning process, all parameters except for the context embedding of the prompt template, which illustrated as the dotted box in Figure 2, are fixed. We show that compared with the hand-crafted prompt, continuous search in embedding space based on CLIP text encoder facilitates the learning of optimal context embedding for our task.

### Loss Functions

We divide the training process of our method into two stages. In the first stage, label embedding is generated by the pre-trained CLIP text encoder, and the vision encoder is trained with the objectives of ranking loss and distillation loss,

$$\mathcal{L}_{stage1} = \mathcal{L}_{rank} + \lambda \mathcal{L}_{dist}, \quad (6)$$

where  $\lambda$  is the weight factor of knowledge distillation.

In the second stage, we only finetune the context embedding with the objective of ranking loss,

$$\mathcal{L}_{stage2} = \mathcal{L}_{rank}. \quad (7)$$

## Experiments

### Experiments Setup

**Datasets:** In the NUS-WIDE dataset, there are 81 human verified labels, in addition to 925 labels based on Flickr

Method	Setting	Task	NUS-WIDE			Open-Images			
			F1(K=3)	F1(K=5)	mAP	F1(K=10)	F1(K=20)	mAP	WmAP
LESA (M=10)		ZSL	31.6	28.7	19.4	1.4	1.0	41.7	-
		GZSL	14.4	16.8	5.6	17.4	14.3	45.4	-
ZS-SDL	ZS	ZSL	30.5	27.8	25.9	10.7	8.3	62.9	-
		GZSL	18.5	21.0	12.1	37.8	32.9	75.3	-
BiAM*		ZSL	32.7	29.8	25.9	7.0	5.5	65.6	72.9
		GZSL	15.4	18.2	9.4	14.8	9.7	<b>81.7</b>	85.0
CLIP-FT		ZSL	23.5	21.7	30.5	19.1	11.1	66.2	88.2
		GZSL	20.3	23.2	16.8	40.2	35.4	77.5	85.9
MKT	OV	ZSL	<b>34.1</b>	<b>31.1</b>	<b>37.6</b>	<b>19.7</b>	<b>11.4</b>	<b>68.1</b>	<b>89.2</b>
		GZSL	<b>22.0</b>	<b>25.4</b>	<b>18.3</b>	<b>40.5</b>	<b>35.4</b>	81.4	<b>89.8</b>

Table 1: State-of-the-art comparison for ZSL and GZSL tasks on the NUS-WIDE and Open Images datasets. The results are reported in terms of mAP, as well as precision (P), recall (R), and F1 score at  $K \in \{3, 5\}$  for NUS-WIDE and  $K \in \{10, 20\}$  for Open Images. ‘\*’ means that the results are reproduced based on official pre-trained models. Bold indicates the best score.

user tags. Similar to LESA (Huynh and Elhamifar 2020), we treat 925 labels as seen labels and the other 81 labels as unseen labels. Following official train/test split, we utilize 161,789 images for training and 107,859 images for testing. The **Open Images** (v4) dataset is more challenging because it consists of 9M training images and 125,456 testing images. Similar to LESA, we treat 7,186 labels with more than 100 images in training set as seen and the most frequent 400 test labels that are not present in training data as unseen.

**Metrics:** Following LESA, we use mean Average Precision (mAP) and F1 score at  $top-K$  predictions to evaluate our method. The mAP reflects the ranking accuracy of each label across all images and the F1 score reflects the label ranking accuracy of each image.

**Implementation Details:** We use the ImageNet-1K pre-trained ViT-B/16 as our vision backbone. As for the two-stream module, the local head consists of two linear layers, and the global head is a linear projection layer. To generate label embedding and conduct knowledge distillation on vision encoder, we select the pre-trained CLIP with ViT-B/16 image encoder as our VLP model. Patch projection of ViT-B/16 yields  $14 \times 14 = 196$  patches for an image with a resolution of  $224 \times 224$ . The  $k$  for  $top-k$  pooling is set to 18, and the weight of knowledge distillation  $\lambda$  is set to 1. In the first stage, we use AdamW optimizer with base learning rate of 0.001 and weight decay of 0.005. We adjust base learning rate of the AdamW optimizer to 0.00003 during the second stage for fine-tuning the context embedding. On NUS-WIDE, we train the model for 20 epochs with the mini-batch of 128 and 10 epochs with the mini-batch of 16 in the first and second stage, respectively. Considering the large scale of Open Images, the model is trained for 4 epochs and 2 epochs in each stage with the same batch size as above.

### State-of-the-art Comparison

In this experiment, we compare our model with traditional ML-ZSL methods. Also, we fine-tune the pre-trained CLIP on base categories with ranking loss and denote it as CLIP-

FT. As a new OV-ML baseline, CLIP-FT surpasses most existing ML-ZSL methods on mAP. The experimental results on zero-shot learning(ZSL) and generalized zero-shot learning(GZSL) tasks are shown in Table 1. The mAP and F1 scores at  $top-K$  ( $K \in \{3, 5\}$  for NUS-WIDE and  $K \in \{10, 20\}$  for Open Images) are reported.

On NUS-WIDE, the recently proposed BiAM (Narayan et al. 2021), which utilizes a bi-level attention module to enrich the features, acquires the best results in ZSL task with mAP score of 25.9%. MKT surpasses BiAM with an absolute gain of 11.7% mAP and improves the F1 score by absolute gains of 1.4% and 1.3% at  $K=3$  and  $K=5$ , respectively. In GZSL task, the approach of ZS-SDL (Ben-Cohen et al. 2021) achieves the best scores with 12.1% mAP. MKT improves the mAP by an absolute gain of 6.5% and reaches state of the art in terms of F1 score with 22.0% at  $K=3$  and 25.4% at  $K=5$ . Compared with CLIP-FT, MKT shows significant improvement on both ZSL and GZSL task.

On Open Images, following BiAM, we also calculate mAP weighted on different sample numbers(denoted as WmAP). ZS-SDL reaches the state of the art before in terms of F1 score in both ZSL and GZSL tasks. MKT achieves consistent improvement over it with absolute gains of 9.0%/2.7% and 3.1%/2.5% at  $K=10$  and  $K=20$  on ZSL/GZSL task. In comparison with previous best results on mAP/WmAP metric, MKT outperforms BiAM by 2.5%/16.3% on ZSL and have a comparable performance on GZSL task. MKT also surpasses CLIP-FT on both tasks.

### Ablation Studies

**Effects of knowledge distillation and prompt tuning:** To study the impacts of knowledge distillation and prompt tuning, we conduct experiments with different training schemes and illustrate the results in Table 2. We take the first row as the baseline for the following comparisons, which is trained without knowledge distillation and prompt tuning. It shows that the introduction of knowledge distillation improves the performance on both ZSL and GZSL tasks. We conjecture

Distill	Prompt	Task	mAP	F1 (K = 3)	F1 (K = 5)
✗	✗	ZSL	32.4	29.4	26.5
		GZSL	16.8	21.0	24.0
✓	✗	ZSL	37.3	32.5	29.5
		GZSL	18.2	21.7	24.9
✗	✓	ZSL	32.5	29.5	26.4
		GZSL	16.8	21.1	24.1
✓	✓	ZSL	<b>37.6</b>	<b>34.1</b>	<b>31.1</b>
		GZSL	<b>18.3</b>	<b>22.0</b>	<b>25.4</b>

Table 2: Impact of knowledge distillation and prompt tuning.

Embedding	Task	mAP	F1 (K = 3)	F1 (K = 5)
GloVe	ZSL	27.1	22.8	21.4
	GZSL	16.1	20.6	23.4
CLIP	ZSL	<b>32.4</b>	<b>29.4</b>	<b>26.5</b>
	GZSL	<b>16.8</b>	<b>21.0</b>	<b>24.0</b>

Table 3: Impact of label embedding. For a fair comparison, we only change label embedding and train both models without knowledge distillation or prompt tuning.

that knowledge distillation not only facilitates the image embedding to align better with VLP model based label embedding but also suppresses the overfitting of the model to seen labels. Moreover, we observe that prompt tuning can further improve performance. It can be attributed to the reason that the prompt-tuned context embedding tends to pay more attention to the visual information that benefits image classification. Compared with the baseline in the first row, MKT shows significant improvement with the combination of knowledge distillation and prompt tuning.

**Comparison of label embedding:** Because prediction results are based on the similarity between image and label embeddings, label embedding has a significant impact on model performance. Table 3 shows the results of baseline model with VLP model based and GloVe based label embeddings. Compared with the model based on GloVe embedding, the VLP embedding based model achieves superior performance on both ZSL and GZSL task. We speculate that language models like GloVe or Bert cannot capture visual consistency among similar labels because of the lack of visual information during the training process, thus limiting the generalization ability to unseen labels. To validate our assumption, we conduct a label retrieval experiment. We select 62 common labels in NUS-WIDE and divide them into 14 major categories based on their visual and semantic similarity. Both language models (*i.e.*, GloVe and Bert) and VLP models (*i.e.*, CLIP and its prompt-tuned version) are utilized to generate label embeddings. All embeddings are normalized, and cosine similarity is used to retrieve the most similar embeddings. Figure 3 illustrates the retrieval results with the overall Top-3 accuracy and examples of retrieved labels. Notice that compared with language model, VLP model can capture both semantic and visual consistency between la-




Figure 3: Results of label retrieval. Overall Top-3 accuracy and examples of retrieved labels are reported. Retrieved labels belonging to the same major category with the query label are considered to be correct (in green).




Figure 4: Distribution of global and local predictions.

els. For instance, “girls” contains similar visual information with its retrieved labels “man”, “kid” and “person”. We argue that label embedding with both visual and semantic consistency facilitates the generalization to unseen labels.

**Effect of the two-stream module:** To demonstrate the effectiveness of our proposed two-stream module, we conduct ablation studies of both local and global heads. Table 4 shows the results in terms of mAP and F1 score on NUS-WIDE. Notice that the global head only model performs well on mAP while the local head only model achieves better F1 score in ZSL task. We speculate that this is due to the fact that the global representation is more general while the local representation is more discriminative. As illustrated in Figure 4, the local head tends to predict higher scores than the global head. While the more discriminative feature allows relevant labels to stand out, it also makes the model more sensitive to noise, leading to wrong predictions. On the other hand, compared to F1 score, mAP is more susceptible to the wrong predictions with high scores. Therefore, the local head only model acquires better F1 score and inferior mAP. With the combination of local and global heads, the two-stream module can acquire more discriminative predictions with resistance to noise, leading to higher performance.

**Varying the hyper-parameters:** Here, we explore the effect of knowledge distillation and variation of  $k$  value in the local head. Knowledge distillation aims to transfer zero-shot classification ability. We are more concerned about its performance on unseen labels. Figure 5a illustrates the results of ZSL task with respect to distillation weight  $\lambda$ . Notice that when  $\lambda$  is smaller than 1, the performance of our approach improves because knowledge distillation facilitates the alignment of image and label embeddings. However,



Local	Global	Task	mAP	F1 (K = 3)	F1 (K = 5)
✓	✗	ZSL	29.1	<b>29.9</b>	<b>27.2</b>
		GZSL	<u>15.7</u>	<u>20.8</u>	<u>23.8</u>
✗	✓	ZSL	<u>30.3</u>	23.3	21.4
		GZSL	15.5	19.4	22.1
✓	✓	ZSL	<b>32.4</b>	<u>29.4</u>	<u>26.5</u>
		GZSL	<b>16.8</b>	<b>21.0</b>	<b>24.0</b>

Table 4: Effectiveness of the two-stream module. Bold indicates the best, and underline indicates the second best.




Figure 5: Impact of hyper-parameters. The results of ZSL task with respect to distillation weight  $\lambda$  and GZSL task with respect to  $k$  for  $top-k$  in local head are presented.

there is a drop in performance when  $\lambda$  is larger than 2. We argue that too large  $\lambda$  may impair the learning of classification objective  $\mathcal{L}_{rank}$ . The two-stream module is designed to improve the recognition of multiple labels, so we focus more on GZSL tasks. Figure 5b illustrates the results of GZSL when altering  $k$  value in the local head. As  $k$  increases, F1 score reaches the highest when  $k=18$ . We argue that when  $k$  is too small, the local head output is sensitive to noise. On the other hand, if  $k$  is too large, the output will be less discriminative. For example, if  $k$  is set as the total patch number,  $top-k$  pooling will be equal to global average pooling. In contrast to F1 score, mAP tends to increase while  $k$  value increases. When  $k$  is small, the local head output tends to be discriminative but sensitive to noise, resulting in a lower mAP value. As  $k$  increases, the output becomes moderate and more resistant to noise, leading to a higher mAP value.

### Qualitative Assessment

In this section, we visualize both predictions and attention maps on several samples. Figure 6 presents predictions of CLIP, BiAM and our approach on ZSL and GZSL tasks respectively. Compared with CLIP, our approach produces more diverse predictions because the two-stream module captures discriminative features. Compared to BiAM, our model with VLP based label embedding can identify semantic and visual similarity among labels. For example, in the last sample of Figure 6, label “plane”, “airplane” and “aircraft” are synonymous and should have similar scores. Figure 7 illustrates the comparison of attention maps between BiAM and ours. The results show that our method can capture relevant regions more precisely. For instance, in the first column, BiAM pays attention to large irrelevant areas while our method exactly focuses on the boat region.

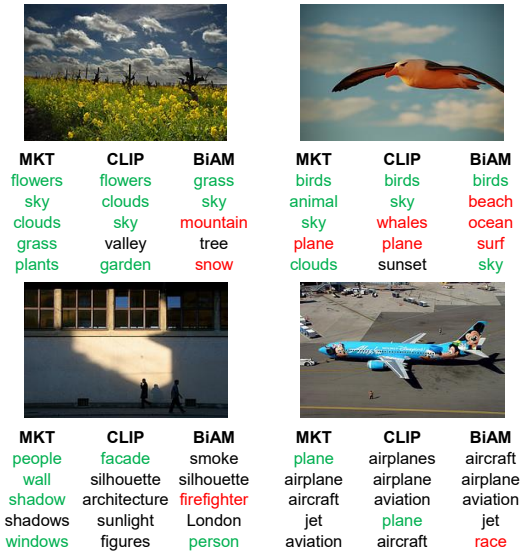


Figure 6: Comparison of predictions. The top row shows the prediction in ZSL task, and the bottom is the prediction in GZSL task. True positive predictions are shown in green and the red font denotes apparently incorrect predictions.

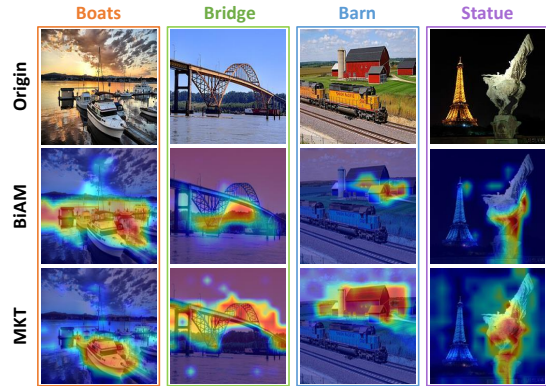


Figure 7: Comparison of Grad-CAM visualization.

### Conclusion

In this work, we propose an open-vocabulary based multi-modal knowledge transfer (MKT) framework for multi-label classification, which jointly exploits semantic multi-modal information in image-text pairs based VLP models. To facilitate transferring the image-text matching ability of VLP model to classification, knowledge distillation and prompt tuning are introduced. Additionally a two-stream module is proposed to capture both local and global features, leading to significant performance gains in multi-label tasks. Extensive results demonstrate that our model surpasses previous ML-ZSL methods and establishes a new state of the art for open-vocabulary multi-label classification on NUS-WIDE and Open Images datasets. This is the first work in open-vocabulary multi-label classification and it is expected to encourage future works to explore multi-modal knowledge applications in classification.

## Acknowledgements

This work is supported in part by the National Natural Science Foundation of China under Grant 62171248, the Natural Science Foundation of Guangdong Province 2021A1515011807, Shenzhen Science and Technology Program (Grant No.RCYX20200714114523079, JCYJ20220818101012025) and the PCNL KEY project (PCL2021A07), and Shenzhen Science and Technology Innovation Commission (Research Center for Computer Network (Shenzhen) Ministry of Education).

## References

- Ben-Cohen, A.; Zamir, N.; Ben-Baruch, E.; Friedman, I.; and Zelnik-Manor, L. 2021. Semantic Diversity Learning for Zero-Shot Multi-Label Classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 640–650.
- Chen, Y.-C.; Li, L.; Yu, L.; Kholy, A. E.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. J. 2020. UNITER: Learning UNiversal Image-TEXT Representations. In *European Conference on Computer Vision (ECCV 2020)*.
- Chen, Z.-M.; Wei, X.-S.; Wang, P.; and Guo, Y. 2019. Multi-Label Image Recognition With Graph Convolutional Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Cheng, X.; Lin, H.; Wu, X.; Shen, D.; Yang, F.; Liu, H.; and Shi, N. 2022. Mltr: Multi-label classification with transformer. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. N. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR 2021: The Ninth International Conference on Learning Representations*.
- Du, Y.; Wei, F.; Zhang, Z.; Shi, M.; Gao, Y.; and Li, G. 2022. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14084–14093.
- Ghiasi, G.; Gu, X.; Cui, Y.; and Lin, T.-Y. 2022. Scaling open-vocabulary image segmentation with image-level labels. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*, 540–557. Springer.
- Gong, Y.; Jia, Y.; Leung, T.; Toshev, A.; and Ioffe, S. 2014. Deep Convolutional Ranking for Multilabel Image Annotation. In *ICLR 2014 : International Conference on Learning Representations (ICLR) 2014*.
- Gu, X.; Lin, T.; Kuo, W.; and Cui, Y. 2022. Open-vocabulary Object Detection via Vision and Language Knowledge Distillation. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Gupta, A.; Narayan, S.; Khan, S.; Khan, F. S.; Shao, L.; and van de Weijer, J. 2021. Generative Multi-Label Zero-Shot Learning. *arXiv preprint arXiv:2101.11606*.
- Hinton, G.; Vinyals, O.; Dean, J.; et al. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7).
- Huynh, D.; and Elhamifar, E. 2020. A Shared Multi-Attention Framework for Multi-Label Zero-Shot Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Huynh, D.; Kuen, J.; Lin, Z.; Gu, J.; and Elhamifar, E. 2022. Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7020–7031.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 4904–4916. PMLR.
- Kim, W.; Son, B.; and Kim, I. 2021. ViLT: Vision-and-Language Transformer Without Convolution or Region Supervision. In *ICML 2021: 38th International Conference on Machine Learning*, 5583–5594.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2009. Learning to detect unseen object classes by between-class attribute transfer. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 951–958.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2014. Attribute-Based Classification for Zero-Shot Visual Object Categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3): 453–465.
- Lanchantin, J.; Wang, T.; Ordonez, V.; and Qi, Y. 2021. General Multi-label Image Classification with Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16478–16488.
- Lee, C.-W.; Fang, W.; Yeh, C.-K.; and Wang, Y.-C. F. 2018. Multi-Label Zero-Shot Learning With Structured Knowledge Graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Li, G.; Duan, N.; Fang, Y.; Gong, M.; and Jiang, D. 2020. Unicoder-VL: A Universal Encoder for Vision and Language by Cross-Modal Pre-Training. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07): 11336–11344.
- Li, Q.; Qiao, M.; Bian, W.; and Tao, D. 2016. Conditional Graphical Lasso for Multi-Label Image Classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.



- Li, X.; Yin, X.; Li, C.; Zhang, P.; Hu, X.; Zhang, L.; Wang, L.; Hu, H.; Dong, L.; Wei, F.; Choi, Y.; and Gao, J. 2020. Oscar: Object-Semantics Aligned Pre-training for Vision-Language Tasks. In *European Conference on Computer Vision*, 121–137.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, 4582–4597. Association for Computational Linguistics.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.
- Ma, Z.; Luo, G.; Gao, J.; Li, L.; Chen, Y.; Wang, S.; Zhang, C.; and Hu, W. 2022. Open-vocabulary one-stage detection with hierarchical visual-language knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14074–14083.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems 26*, volume 26, 3111–3119.
- Narayan, S.; Gupta, A.; Khan, S.; Khan, F. S.; Shao, L.; and Shah, M. 2021. Discriminative Region-Based Multi-Label Zero-Shot Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 8731–8740.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1532–1543.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML 2021: 38th International Conference on Machine Learning*, 8748–8763.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21(140): 1–67.
- Read, J.; Pfahringer, B.; Holmes, G.; and Frank, E. 2011. Classifier chains for multi-label classification. *Machine learning*, 85(3): 333–359.
- Tsoumakas, G.; and Katakis, I. 2007. Multi-label classification: An overview. *International Journal of Data Warehousing and Mining (IJDWM)*, 3(3): 1–13.
- Wang, J.; Yang, Y.; Mao, J.; Huang, Z.; Huang, C.; and Xu, W. 2016. CNN-RNN: A Unified Framework for Multi-Label Image Classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, Z.; Chen, T.; Li, G.; Xu, R.; and Lin, L. 2017. Multi-Label Image Recognition by Recurrently Discovering Attentional Regions. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Xian, Y.; Schiele, B.; and Akata, Z. 2017. Zero-Shot Learning - the Good, the Bad and the Ugly. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Zang, Y.; Li, W.; Zhou, K.; Huang, C.; and Loy, C. C. 2022. Open-Vocabulary DETR with Conditional Matching. In Avidan, S.; Brostow, G. J.; Cissé, M.; Farinella, G. M.; and Hassner, T., eds., *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part IX*, volume 13669 of *Lecture Notes in Computer Science*, 106–122. Springer.
- Zareian, A.; Rosa, K. D.; Hu, D. H.; and Chang, S.-F. 2021. Open-Vocabulary Object Detection Using Captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14393–14402.
- Zhang, Y.; Gong, B.; and Shah, M. 2016. Fast Zero-Shot Image Tagging. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5985–5994.
- Zhang, Z.; and Saligrama, V. 2015. Zero-Shot Learning via Semantic Similarity Embedding. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.
- Zhu, F.; Li, H.; Ouyang, W.; Yu, N.; and Wang, X. 2017. Learning Spatial Regularization With Image-Level Supervisions for Multi-Label Image Classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.