

Correlative and Discriminative Label Grouping for Multi-Label Visual Prompt Tuning

Lei-Lei Ma^{1*}, Shuo Xu^{1*}, Ming-Kun Xie^{2†}, Lei Wang³, Dengdi Sun¹, Haifeng Zhao^{1†}

¹Anhui Provincial Key Laboratory of Multimodal Cognitive Computation,
School of Computer Science and Technology, Anhui University, Hefei, China

²Nanjing University of Aeronautics and Astronautics, Nanjing, China

³Nanjing University of Science and Technology, Nanjing, China

Abstract


Modeling label correlations has always played a pivotal role in multi-label image classification (MLC), attracting significant attention from researchers. However, recent studies have overemphasized co-occurrence relationships among labels, which can lead to overfitting risk on this overemphasis, resulting in suboptimal models. To tackle this problem, we advocate for balancing correlative and discriminative relationships among labels to mitigate the risk of overfitting and enhance model performance. To this end, we propose the Multi-Label Visual Prompt Tuning framework, a novel and parameter-efficient method that groups classes into multiple class subsets according to label co-occurrence and mutual exclusivity relationships, and then models them respectively to balance the two relationships. In this work, since each group contains multiple classes, multiple prompt tokens are adopted within Vision Transformer (ViT) to capture the correlation or discriminative label relationship within each group, and effectively learn correlation or discriminative representations for class subsets. On the other hand, each group contains multiple group-aware visual representations that may correspond to multiple classes, and the mixture of experts (MoE) model can cleverly assign them from the group-aware to the label-aware, adaptively obtaining label-aware representation, which is more conducive to classification. Experiments on multiple benchmark datasets show that our proposed approach achieves competitive results and outperforms SOTA methods on multiple pre-trained models.

1. Introduction

Multi-label image classification (MLC) aims to assign multiple labels to a single instance [56]. As a fundamental



(a) predicted image



(b) predicted probability

Figure 1. This illustrates the negative impact of co-occurrence. Due to the high co-occurrence probability $p(\text{dog}|\text{frisbee})$, e.g., 0.23, the model incorrectly infers the presence of *dog* when the label *frisbee* appears.

machine learning algorithm [31], MLC reflects real-world scenarios and is an essential component in many applications, e.g., face recognition [3], and scene understanding [57], weakly-supervised semantic segmentation [2]. Compared with single-label settings, the MLC presents significant challenges due to the exponential expansion of the output space, which usually requires large-scale datasets [45].

In real-world scenarios, co-occurrence between labels is prevalent, and numerous studies [47, 55] have leveraged label relationships to address the challenges resulting from complex label output spaces. In the context of deep learning, pioneering works [41, 44] employ recurrent neural networks (RNNs) to capture limited correlations between labels, while later studies [7, 46] utilize graph convolutional networks (GCNs) to capture global relationships in graphs or higher-order relationships in hypergraphs. More recently, for better representations, most researchers employ attention architectures to decouple an image into multiple label-aware features, as seen in methods like SSGRL [5], Q2L [30], etc.

Despite the great advantages these methods showcase, two significant drawbacks require mitigation. ❶ The pre-trained model in these methods needs to be fine-tuned, leading to increased computational and resource demands. Also, for large-scale pre-trained models (e.g., DINOv2 [35]), full fine-tuning can negatively affect performance on down-

*Equal contribution.

†Correspondence to: Haifeng Zhao (senith@163.com), Ming-Kun Xie (mkxie@nuaa.edu.cn)

stream tasks due to object tasks and data distributions [25].
② More importantly, while modeling correlations between labels can enhance model performance, overemphasizing these correlations will lead to the risks of overfitting and may misguide the model to make erroneous inferences, resulting in sub-optimal overall performance [28]. Figure 1 verifies this in an empirical example. Recent attempts to address this issue, such as BoostMLC [50] and PAT [49], still involve considerable computational costs (*e.g.*, multiple teacher models and more augmented images).

Recently, visual prompt tuning (VPT) [21] has emerged as a parameter-efficient fine-tuning (PEFT) [11, 14] method for adapting large-scale pre-trained vision transformer models to downstream tasks [17, 43, 53]. Naturally, the success of VPT has inspired us to explore using PEFT to address the first issue. However, existing VPT-like methods primarily downstream tasks in single-label prediction [16, 36] or dense prediction tasks [19, 22], which may not suit MLC where the correlation between labels is not considered during the learning representation phase.

Regarding the second drawback, the key is to mitigate the adverse effects of overemphasizing co-occurrence relationships while preserving their advantages. A Naïve idea is to model the correlative (CO) and discriminative (DC) relationships between labels separately and then balance them. In terms of CO, a reasonable strategy is to highlight class pairs with high co-occurrence probabilities within a co-occurrence graph [5, 7], thereby modeling the CO between these classes. In contrast, DC tends to emphasize the characteristics of the object class to highlight distinguishability. Therefore, class pairs with low co-occurrence probability may have higher DC relationships (discriminability). Existing studies [49, 50] tend to be one of them or require more complex models and computation. Therefore, to effectively balance these two relationships, we can naturally group the classes into multiple sets according to their co-occurrence probabilities. All classes are placed together in the CO groups via high probabilities. Conversely, for DC groups, it is advocated to divide all classes using low co-occurrence probabilities.

Drawing on the above investigations and analysis, we propose a multi-label visual prompt tuning (ML-VPT) framework that utilizes grouping classes to balance the two relationships simply and effectively. *Firstly*, we take advantage of VPT and introduce an equal number of visual prompt tokens to each group, which facilitates learning visual representations for each class subset, referred to as group-aware representations. Unlike mainstream methods [5, 30], our approach offers two great advantages: It emphasizes directly modeling label relationships based on multiple prompt tokens for each group within the visual encoder. Our method maintains a balanced focus, not overemphasizing the CO relationship between labels while considering DC. *Secondly*, each class subset comprises multiple classes, and each group

incorporates multiple visual prompt tokens, resulting in each group containing multiple group-aware representations. To more effectively highlight the distinctions between classes, we employ a MoE model to selectively identify suitable visual representations for each class within the relevant group-aware representations. MoE can simplify the mapping from group-aware to label-aware complexity and enable adaptive mapping tailored to specific images.

Overall, our contribution is two-fold: 1) We propose a novel framework for MLC that groups classes first and subsequently models them within ViT. The proposed method balances the CO and DC relationships to overcome the risk of co-occurrence overfitting and improve model performance. 2) The group-aware MoEs are implemented to map group-aware representations to label-aware ones. In a dynamic way, multiple appropriate experts with gating networks are chosen to construct label-aware representations.

2. Related Work

Multi-Label Image Classification. MLC aims to predict multiple labels associated with an instance simultaneously [48]. The mainstream MLC methods can be roughly summarized into two types: modeling label correlations and learning label-aware representations. The former typically employs GCNs or RNNs to capture dependencies between labels [5, 7, 41, 44], while the latter often utilizes attention mechanisms to learn category-specific representations, which focus on effectively processing contextual information in images [30, 33, 34, 38, 51]. Very recently, limited studies have begun to investigate overfitting and causal interventions in the context of MLC. For example, CDD [28] and IDA [29] eliminate the contextual debiasing. BoostMLC [50] and PAT [49] attempt to eliminate the risk of overfitting resulting from overemphasizing co-occurrence, but involve complex calculation processes and computational overhead.

Visual Prompt Tuning. To strike an optimal balance between computational cost and performance, Jia *et al.* pioneer VPT [21], which effectively adapts transformer models pre-trained on large-scale datasets for various downstream tasks. Subsequently, research initiatives analogous to VPT emerged [17, 43, 53], with many researchers employing VPT across various computer vision applications, such as classification [36], test-time adaptation [16], segmentation [22], and continual learning [15]. However, these studies are unsuitable for MLC, as they fail to account for label correlations, whereas our approach explicitly considers them.

Mixture of Experts. The MoE [20] dynamically selects specialized expert networks to process input data, enhancing overall model performance. MoE has been successfully applied across various fields, *e.g.*, reinforcement learning [58], domain generalization [24], multimodal large language models [4], Re-identification [27]. Recently, HQS [52] treats each label as task-specialized and employs the MoE model

for learning in MLC tasks. However, unlike HQS, which assigns experts based on label-aware representation, our approach assigns experts based on group-aware representation, which can balance correlative (CO) and discriminative (DC) relationships to eliminate the risk of overfitting.

3. Preliminaries

3.1. Notations

Let $\mathbf{x} \in \mathcal{X}$ denote the instance, and $\mathbf{y} \in \mathcal{Y}$ denote the corresponding label, where $\mathcal{X} = \mathbb{R}^d$ is the instance space, and $\mathcal{Y} = \{0, 1\}^K$ is the label space containing the K classes. For a given instance \mathbf{x} , $y^k = 1$ denotes k -th label is a relevant label associated with the instance \mathbf{x} , and vice versa. In MLC tasks, given a labeled dataset with N samples $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$. Our goal is to train a model, denoted as $f(\mathbf{x}; \theta) : \mathcal{X} \rightarrow \mathcal{Y}$, using a training dataset \mathcal{D} . The aim is to accurately predict all potential labels \mathbf{y} for unseen instances \mathbf{x} , with θ representing the learnable model parameters. Accordingly, the MLC expected risk function:

$$\mathcal{R}(f) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p(\mathbf{x}, \mathbf{y})} [\mathcal{L}(f(\mathbf{x}), \mathbf{y})], \quad (1)$$

where $\mathcal{L} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$ indicates a multi-label loss function that turns MLC task into multiple more straightforward binary classification problems.

3.2. Visual Prompt Tuning Revisit

Visual Prompt Tuning is mainly implemented in ViTs [12], which usually consists of a patch embedding layer and L stacked transformer blocks (Block). Given an image \mathbf{x} that is divided into a set of non-overlapping patches, the patch embedding layer embedded each patch with the position information \mathbf{E}_{pos} into a D -dimensional space, as follows:

$$\mathbf{E}_0 = \text{PatchEmbed}(\mathbf{x}) + \mathbf{E}_{\text{pos}}. \quad (2)$$

Next, image patch embeddings $\mathbf{E}_i = \{\mathbf{e}_i^j \in \mathbb{R}^D | j \in \mathbb{N}, 1 \leq j \leq N_e\}$ along with a class token cls_i is input into the $(i+1)$ -th transformer block:

$$[\text{cls}_i, \mathbf{E}_i] = \text{Block}_i([\text{cls}_{i-1}, \mathbf{E}_{i-1}]), \quad (3)$$

where $[\cdot, \cdot]$ indicates stacking and concatenation on the sequence length dimension.

For VPT, Jia *et al.* [21] advocate introducing a set of N learnable prompt tokens as input to i -th transformer block (Block_i), denoted as $\mathbf{P}_{i-1} = \{\mathbf{p}_{i-1,k} \in \mathbb{R}^D | k \in \mathbb{N}, 1 \leq k \leq N_p\}$. The whole VPT is formulated as follows:

$$[\text{cls}_i, -, \mathbf{E}_i] = \text{Block}_i([\text{cls}_{i-1}, \mathbf{P}_{i-1}, \mathbf{E}_{i-1}]). \quad (4)$$

3.3. How to Group Labels?

To group the labels, we take an approach [50], which performs spectral clustering on the co-occurrence graph \mathcal{G}^+

and dis-occurrence graph \mathcal{G}^- . Specifically, we define a co-occurrence matrix $\mathbf{S} \in \mathbb{R}^{K \times K}$ and an affinity matrix,

$$\mathbf{M} = \begin{cases} \mathbf{M}^+ = (\sqrt{\tau} \mathbf{S} + \sqrt{\tau} \mathbf{S}^\top) / 2, & \mathcal{G} = \mathcal{G}^+ \\ \mathbf{M}^- = \mathbf{I} - (\sqrt{\tau} \mathbf{S} + \sqrt{\tau} \mathbf{S}^\top) / 2, & \mathcal{G} = \mathcal{G}^- \end{cases} \quad (5)$$

where τ is a hyper-parameter. By applying spectral clustering to \mathbf{M} , we divide the classes \mathcal{C} into several subsets within both CO and DC, designated as $\{\mathcal{C}_t^+\}_{t=1}^{N_c}$ and $\{\mathcal{C}_t^-\}_{t=1}^{N_c}$, respectively, where N_c is the number of subsets in each set. Additionally, the number of classes $|\mathcal{C}_t^+|/|\mathcal{C}_t^-|$ within these subsets is expected to vary. More can be found in Appendix.

4. Method

In this section, we describe our novel framework, named ML-VPT, with group-aware visual prompt tuning (in §4.2) and group-aware MoE (in §4.3) for MLC, depicting the overall overview in Fig. 2. In the end, we describe in §4.4 how our CO and DC jointly optimize and infer.

4.1. Overview

To address the co-occurrence overfitting problem as well as high resource demands in training, we propose a novel multi-label learning method with VPT, and the key is to balance CO and DC, effectively reducing the overfitting risk resulting from overemphasized co-occurrence. To this end, first, we divide all classes into N_c CO groups and N_c DC groups as class subsets, according to the co-occurrence graph.

Secondly, to model CO and DC between labels, we add two sets of prompt tokens inside ViT, called CO and DC tokens $\mathbf{P}^+/\mathbf{P}^-$, by leveraging the capabilities of advanced VPT. In each subset, these tokens are evenly distributed across several groups, with each group associated with a class subset. Within the ViT, extended prompt tokens learn CO and DC relationships among labels, aiming to mitigate the adverse effects of CO relationships. Upon completing those, the ViT generates multiple group-aware representations $\mathbf{Z}^+/\mathbf{Z}^-$ for class subsets in CO and DC groups:

$$\mathbf{Z}^+ \cup \mathbf{Z}^- = \text{ViT}_\Phi([\mathbf{P}^+, \mathbf{P}^-], \mathbf{x}). \quad (6)$$

These representations are shared among various classes within the same grouping strategy, which complicates the representation from group-aware to label-aware.

Subsequently, to enhance the learning of label-aware representations $\mathbf{C}^+/\mathbf{C}^-$ from group-aware representations for each class, we use MoEs to adaptively map these representations from group-aware to label-aware:

$$\mathbf{C}^+ = \text{MoE} \left[\biguplus \text{Gate}(\mathbf{Z}^+) \right], \quad \mathbf{C}^- = \text{MoE} \left[\biguplus \text{Gate}(\mathbf{Z}^-) \right], \quad (7)$$

where \biguplus indicates the weighted sum, MoE refers to the mixture-of-experts model, and Gate is the gating network.

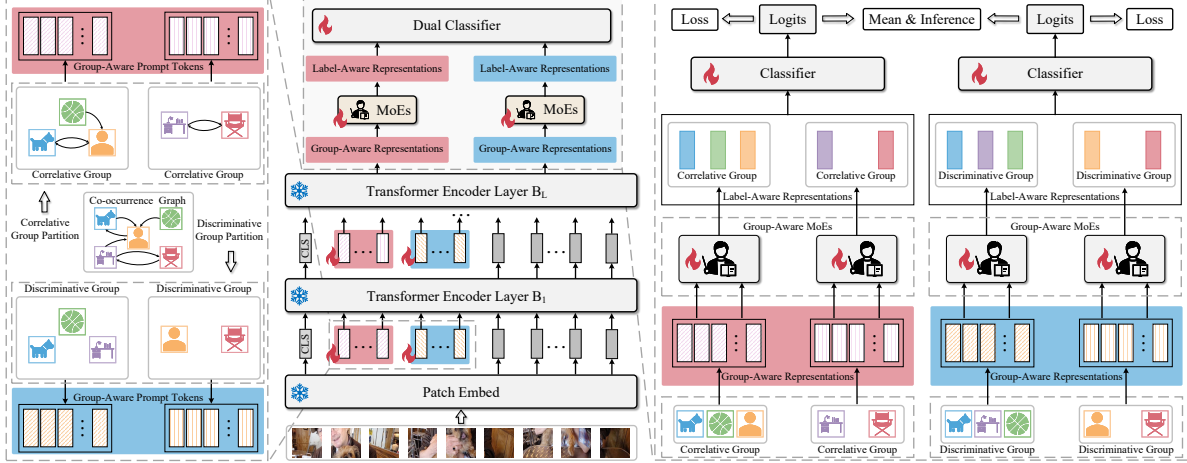


Figure 2. The overview of our framework. We apply clustering strategies to the co-occurrence graph to group labels into co-occurrence groups (CO) and discriminative groups (DC). Then, we introduce two types of prompt tokens into ViT, corresponding to CO and DC, respectively. ViT subsequently generates two types of group-aware representations. To adaptively map these representations to label-aware ones, we design Group-Aware MoEs. Finally, two distinct classifiers are employed to classify the label-aware representations.

Finally, we employ two sets of classifiers with non-shared parameters, each making predictions in their respective, and then combine the results using a weighted sum.

4.2. Group-Aware Visual Prompt Tuning

Learning the relationship between labels has long been a key focus in MLC [7, 30]. However, overemphasizing these correlations often results in inferring objects primarily from co-occurring ones, which is unsuitable for all real-world scenarios [50]. On the other hand, learning discriminative representations for each object independently can lead to incorrect inferences when context or co-occurring objects are absent [46]. To address this issue, we decompose label relationships into two aspects: ❶ In the CO, learning shared information for each group is utilized to model the relationships between co-occurring labels. ❷ Conversely, in the DC, learning unique information for each group is employed to define the relationships between labels that do not co-occur.

Specifically, we define two types of prompt tokens, referred to as CO group prompts $\{\mathbf{p}_t^+\}_{t=1}^{N_c}$ and DC group prompts $\{\mathbf{p}_t^-\}_{t=1}^{N_c}$. On the one hand, from the perspective of grouping class, each prompt token $\mathbf{p}_t^+/\mathbf{p}_t^-$ uniquely maps to a specific class subset $\mathcal{C}_t^+/\mathcal{C}_t^-$. On the other hand, in terms of representation learning, the prompt token $\mathbf{p}_t^+/\mathbf{p}_t^-$ is associated with all classes, contained by $\mathcal{C}_t^+/\mathcal{C}_t^-$, thus forming a one-to-many relationship. For example, if \mathcal{C}_t^+ includes the classes $\{\text{bicycle}; \text{car}\}$, then \mathbf{p}_t^+ is exclusively associated with these classes, which can indicate vision representations of these classes. Similarly, \mathbf{p}_t^- corresponds to \mathcal{C}_t^- . In the context of prompt tuning, we introduce a set of CO group prompts $\mathbf{P}_{i-1}^+ = \{\mathbf{p}_{i-1,t}^+ | t \in \mathbb{N}, 1 \leq t \leq N_c\}$ and a set of DC group prompts $\mathbf{P}_{i-1}^- = \{\mathbf{p}_{i-1,t}^- | t \in \mathbb{N}, 1 \leq t \leq N_c\}$ into each transformer block (Block_i), which is

designed to facilitate the learning of visual representations pertinent to specific class subsets \mathcal{C}_t^+ and \mathcal{C}_t^- . It is expressed in the following mathematical form:

$$[\text{cls}_i, -, -, \mathbf{E}_i] = \text{Block}_i([\text{cls}_{i-1}, \mathbf{P}_{i-1}^+, \mathbf{P}_{i-1}^-, \mathbf{E}_{i-1}]). \quad (8)$$

Leveraging the powerful long-range modeling advantages of the transformer, each $\mathbf{p}_{i,t}^+/\mathbf{p}_{i,t}^-$ can aggregate semantic information about the class subset \mathcal{C}_t^+ from the image embedding \mathbf{E}_i . In CO, $\mathbf{p}_{i,t}^+$ focuses on learning shared information within the class subset \mathcal{C}_t^+ . Conversely, in DC, due to the classes that do not co-occur in subsets \mathcal{C}_t^- , the DC group prompts $\mathbf{p}_{i,t}^-$ seek to capture the distinct information unique to each class within the subset. However, due to MSA, prompt tokens $\mathbf{p}_{i,t}$ may transfer information to each other. To mitigate this effect, we introduce a predefined mask to prevent information exchange between prompt tokens across groups, making them more discriminative.

4.3. Group-Aware Mixture-of-Experts

As in §4.2, each class subset was associated with one prompt token. However, this simplification does not accurately represent the real-world scenarios. For example, in the CO, the classes $\{\text{traffic light}; \text{stop sign}; \text{car}\}$ exhibit a strong co-occurrence relationship and are classified into the same group. However, when only one prompt token is used, it becomes challenging to capture the co-occurrence relationships between labels within the same group. In contrast, for the DC, the classes $\{\text{skateboard}; \text{skis}; \text{refrigerator}\}$, which are divided into the same group due to their solid discrimination relationships, tend to show significant visual differences [54]

¹The index t , in class subsets \mathcal{C}_t^+ and \mathcal{C}_t^- , may not be the same. However, for simplicity, we will uniformly denote it as t in our analysis.

in the images. When only one prompt token is assigned to each group, the model tends to share the group’s common semantic information, which in turn impedes its ability to capture the unique representations for each class.

To this end, we expand the number of prompt tokens for each group. Since each class possesses unique characteristics, it is essential to select a group-aware representation that is appropriate for each class. To achieve this, we frame the selection of group-aware representations as a multi-task learning problem, using the MoEs to play its advantages in multi-task learning [24]. And each subtask is associated with multiple group-aware representations, which are then mapped to label-aware representations via the MoE model.

More Prompt Tokens for Each Group. Specifically, we extend each prompt token $\mathbf{p}_{i-1,t}$ from 1 to N_m , defined as follows: $\mathbf{p}_{i-1,t} = \{\mathbf{p}_{i-1,t}^e \in \mathbb{R}^D | e \in \mathbb{N}, 1 \leq e \leq N_m\}$. Moreover, $\mathbf{p}_{i-1,t}^{e,+}$ indicates the correlative group prompt tokens, while $\mathbf{p}_{i-1,t}^{e,-}$ is the opposite. To reduce additional parameters, we advocate introducing MoE after the last transformer block. Then, the final transformer block’s output can be expressed as: $[\text{cls}_o, \mathbf{Z}^+, \mathbf{Z}^-, \mathbf{E}_o]$, where \mathbf{Z}^+ corresponds to the output of the CO and \mathbf{Z}^- is the opposite. More specifically, $\mathbf{z}_t^{e,+}$ is related to the e -th prompt token in the t -th discriminative group. And, we refer to \mathbf{z}_t^e as the e -th representation corresponding to the class subset \mathcal{C}_t .

Mixture-of-Experts. Subsequently, we introduce correlative and discriminative group experts \mathcal{E} , and each expert $E_t^e(\cdot)$ is implemented as an FFN with a ReLU activation and residual connection:

$$\mathcal{E}^+ = \{E_t^{e,+} | e, t \in \mathbb{N}, 1 \leq e \leq N_m, 1 \leq t \leq N_c\}, \quad (9)$$

$$\mathcal{E}^- = \{E_t^{e,-} | e, t \in \mathbb{N}, 1 \leq e \leq N_m, 1 \leq t \leq N_c\}, \quad (10)$$

$$E_t^e(\mathbf{z}_t^e) = \mathbf{z}_t^e + \sigma(\mathbf{z}_t^e \mathbf{W}_{t,\text{dn}}^e + \mathbf{b}_{t,\text{dn}}^e) \mathbf{W}_{t,\text{up}}^e + \mathbf{b}_{t,\text{up}}^e, \quad (11)$$

where \mathbf{W}_t^e and \mathbf{b}_t^e are the weights and biases of the experts, and $\sigma(\cdot)$ is a activation. To make lightweightness, the hidden dimension is kept small, e.g., 5. The two groups of experts don’t share parameters.

Gating Networks. We employ two distinct strategies for expert selection: ① Each subset within the grouping class utilizes a routing network to learn varying weights for multiple group-aware representations \mathbf{z}_t^e , followed by weighted merging by multiple experts. ② The second strategy considers each class independently; here, different classes within the corresponding class subset group employ different dynamic routing networks to determine diverse weights for multiple visual representations \mathbf{z}_t^e , subsequently merging these using multiple experts. We opt for the second strategy, as different classes necessitate distinct information, rendering identical shared visual representations insufficient.

For each class k , we define a gating network $g_k^+(\cdot)/g_k^-(\cdot)$ in CO and DC groups, respectively, implemented as a simple

fully connected network:

$$\mathbf{w}_k^+ = \text{softmax}(\mathbf{z}_t^{e,+} \mathbf{W}_k^+ + \mathbf{b}_k^+) \in \mathbb{R}^{N_m}, k \in \mathcal{C}_t^+, \quad (12)$$

$$\mathbf{w}_k^- = \text{softmax}(\mathbf{z}_t^{e,-} \mathbf{W}_k^- + \mathbf{b}_k^-) \in \mathbb{R}^{N_m}, k \in \mathcal{C}_t^-, \quad (13)$$

where $\mathbf{w}_k^+ = \{w_k^{e,+}\}_{e=1}^{N_e}$ and $\mathbf{w}_k^- = \{w_k^{e,-}\}_{e=1}^{N_e}$ denote the weights associated with class k in class subset \mathcal{C}_t^+ and \mathcal{C}_t^- .

Adaptive Label-Aware Representation. Adaptively aggregate group-aware representations into label-aware representations. The group-aware MoE is formalized as follows:

$$\mathbf{c}_k^+ = \sum_{e=1}^{N_e} w_k^{e,+} E_t^{e,+}(\mathbf{z}_t^{e,+}), \quad (14)$$

$$\mathbf{c}_k^- = \sum_{e=1}^{N_e} w_k^{e,-} E_t^{e,-}(\mathbf{z}_t^{e,-}). \quad (15)$$

where \mathbf{c}_k^+ and \mathbf{c}_k^- indicate the label-aware representations for class k , obtained from CO and DC groups, respectively.

4.4. Optimization and Inference

In this work, we adapt dual classification heads to mitigate the negative impacts arising from the accumulated prediction errors from a single classification head [49]. The logits for both correlative and discriminative groups can be predicted using two linear classifiers with a sigmoid as follows:

$$\hat{y}_k^+ = \sigma(\mathbf{W}_k^T \mathbf{c}_k^+ + b_k^+), \quad \hat{y}_k^- = \sigma(\mathbf{W}_k^T \mathbf{c}_k^- + b_k^-). \quad (16)$$

Here, \hat{y}_k^+ and \hat{y}_k^- represent the predictions for class k in the two types of , while $f^+(\mathbf{x}) = \hat{y}^+$ and $f^-(\mathbf{x}) = \hat{y}^-$ denote the predictions for all classes about instance \mathbf{x} , respectively.

Referring to Eq. (1), the overall expected risk associated with our proposed method can be reformulated as follows:

$$\mathcal{R}(f) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p(\mathbf{x}, \mathbf{y})} [\mathcal{L}(f^+(\mathbf{x}), \mathbf{y}) + \mathcal{L}(f^-(\mathbf{x}), \mathbf{y})], \quad (17)$$

where $\mathcal{L}(\cdot, \cdot)$ is ASL loss function [1]. Given an instance \mathbf{x} , the model’s final prediction is $f(\mathbf{x}) = \hat{y} = 0.5 \cdot (\hat{y}^+ + \hat{y}^-)$.

5. Experiments

5.1. Experimental Settings

Dataset and Evaluation Metric. In this work, we evaluate the effectiveness of the proposed method on four benchmark datasets, including Pascal VOC 2007 (VOC07) [13], MS-COCO 2014 (COCO) [26], NUS-WIDE (NUS) [8], and Visual Genome (VG256) [23]. For a fair comparison, we utilize widely adopted metrics: Mean Average Precision (mAP) across all classes. Additionally, consistent with previous works [7, 33], we also showcase overall precision (OP), recall (OR), F1-measure (OF1), as well as per-category precision (CP), recall (CR), and F1-measure (CF1) for detailed comparisons. It is important to emphasize that mAP, OF1, and CF1 are the most important metrics among these.

Table 1. Comparison of our method with SOTA models on COCO at 224×224 and 448×448 resolution. All metrics are in %. Since mAP, CF1, and OF1 are among the most important evaluation metrics, they are highlighted in dark gray in our method.

Method	Backbone	Pre-trained Data	Resolution: 224×224							Resolution: 448×448						
			mAP	CP	CR	CF1	OP	OR	OF1	mAP	CP	CR	CF1	OP	OR	OF1
VPT	ViT-B	ImageNet 1K	78.0	71.7	72.3	72.0	74.6	75.4	75.0	82.6	76.1	76.7	76.4	78.2	79.1	78.7
GateVPT			75.6	69.7	70.3	70.0	72.9	73.7	73.3	80.4	74.2	74.7	74.4	76.6	77.5	77.0
E2VPT			77.3	71.6	72.2	71.9	74.5	75.4	75.0	81.7	75.3	75.9	75.6	77.7	78.5	78.1
Ours			79.6	73.1	73.6	73.3	75.7	76.5	76.1	83.6	76.9	77.5	77.2	79.1	79.9	79.5
VPT	MAE	ImageNet 1K	71.0	66.0	66.6	66.3	71.1	72.0	71.6	72.2	66.9	67.6	67.3	72.3	73.1	72.7
GateVPT			66.5	62.4	62.9	62.6	68.2	69.0	68.6	69.1	64.5	65.1	64.8	70.1	70.9	70.5
E2VPT			69.8	65.0	65.6	65.3	70.4	71.2	70.8	73.0	67.8	68.4	68.1	72.8	73.6	73.2
Ours			75.2	69.3	69.9	69.6	74.0	74.9	74.4	78.8	72.8	73.4	73.1	76.8	77.6	77.2
VPT	MoCo v3	ImageNet 1K	73.0	67.6	68.1	67.9	71.8	72.6	72.2	75.9	70.2	70.8	70.5	74.1	74.9	74.5
GateVPT			70.9	65.7	66.3	66.0	70.3	71.2	70.8	74.9	69.4	70.0	69.7	73.4	74.2	73.8
E2VPT			72.7	67.3	67.9	67.6	71.7	72.5	72.1	76.0	70.3	70.9	70.6	74.2	75.0	74.6
Ours			75.1	69.1	69.7	69.4	73.2	74.0	73.6	77.3	71.4	72.0	71.7	75.2	76.0	75.6
VPT	ViT-B-21k	ImageNet 21K	81.0	74.5	75.1	74.8	77.3	78.1	77.7	85.0	78.4	79.0	78.7	80.5	81.4	81.0
GateVPT			80.8	74.2	74.8	74.5	76.9	77.8	77.3	84.5	77.8	78.4	78.1	80.1	80.9	80.5
E2VPT			81.9	75.4	76.0	75.7	77.9	78.8	78.3	85.2	78.6	79.2	78.9	80.7	81.6	81.1
Ours			83.0	76.2	76.7	76.4	78.6	79.5	79.0	86.4	79.7	80.3	80.0	81.6	82.5	82.0
VPT	DINOv2/B	LVD-142M [35]	86.1	79.5	80.1	79.8	81.9	82.8	82.4	89.7	83.2	83.8	83.5	85.2	86.1	85.7
GateVPT			85.6	79.0	79.6	79.3	81.6	82.5	82.0	89.1	82.5	83.1	82.8	84.5	85.5	85.0
E2VPT			86.3	79.8	80.4	80.1	82.2	83.1	82.6	89.6	83.1	83.7	83.4	85.2	86.1	85.6
Ours			87.5	80.8	81.4	81.1	83.0	83.9	83.4	90.6	84.2	84.8	84.5	86.0	86.9	86.4
VPT	DINOv2/S	LVD-142M [35]	80.1	73.8	74.4	74.1	77.1	77.9	77.5	84.5	78.2	78.8	78.5	81.1	82.0	81.5
GateVPT			79.4	73.2	73.8	73.5	76.5	77.3	76.9	83.3	77.0	77.7	77.3	80.0	80.9	80.4
E2VPT			80.6	74.2	74.8	74.5	77.5	78.4	77.9	84.3	78.1	78.8	78.5	81.0	81.9	81.4
Ours			83.4	76.6	77.1	76.8	79.6	80.5	80.0	87.4	80.8	81.4	81.1	83.2	84.1	83.7

Implementation Details. Following previous MLC methods [30, 33, 37], we adopt a similar experimental setup. The AdamW [32] optimizer with the one-cycle policy lr_schedule [40] is applied to train the model with maximal learning rate of 0.0005. All models are trained for 40 epochs with the early stopping. The batch size is set to 64. For data augmentation, we apply RandAugment [9] and Cutout [10]. To make the model more robust, we apply an exponential moving average to the model parameters θ , using a decay rate of 0.9997.

Pre-trained Models. To verify the robustness of our proposed method, we utilized a series of pre-trained models as backbone networks, including ViT [12], MAE [18], MoCo v3 [6], and DINOv2 [35]. For ViT, we employ ViT-B and ViT-B-21k, supervised training on ImageNet 1k [39], and ImageNet 21k [39] respectively. MAE and MoCo v3 are trained on ImageNet 1k for autoregressive and contrastive self-supervised training, respectively. For DINO v2, we utilize DINOv2/B and DINOv2/S, the former is the base model and the latter is the small variant, trained on a larger dataset for self-supervised learning.

Comparing Methods. To evaluate the performance of our model, we carried out a comprehensive set of experiments involving the following methods: VPT [21], GateVPT [53], and E2VPT [17]. These methods are primarily designed for single-label tasks, and the cls token output of the model is input into a linear classifier and optimized using the ASL [1] to adapt it to multi-label tasks. In our proposed methods, by default, the number of groups in the CO and DC groups is 5, and the number of experts in each group is 3. In the work, all comparing methods are setup and compared fairly.

5.2. Compared to State-of-the-Art (SOTA) Results

Performance on MS-COCO 2014. In Tab. 1, we report the comparison results between our method and the state-of-the-

art method on the COCO. This table demonstrates that, under the same setting (*i.e.*, the same pre-trained backbone and resolution), our method generally surpasses SOTA methods in terms of mean mAP, CF1, and OF1 at both 224×224 and 448×448 resolutions. Specifically, at the 224×224 resolution, our method achieves improvements over the SOTA methods as follows: the mAP increases by 1.2%-8.7%, the CF1 by 0.7%-7.0%, and the OF1 by 0.7%-5.8%. Similarly, at the 448×448 resolution, improvements are noted in the mAP from 0.4% to 6.6%, the CF1 from 0.8% to 5.8%, and the OF1 score from 0.7% to 4.5%, also compared with the SOTA method. These results demonstrate the effectiveness of our method. Moreover, it should be noted that GateVPT and E2VPT, which are modified from VPT, do not consistently outperform VPT. For instance, at a 224×224 resolution using ViT-B, MAE, and MoCo v3, VPT demonstrates superior performance compared to these two modifications. A potential explanation could be that GateVPT and E2VPT are tailored for single-label tasks and do not account for label correlations during the feature extraction phase. The best mAP is 90.6% achieved by our method, using DINOv2/B at a resolution of 448×448 , which can be attributed to the superior capabilities of DINOv2/B and higher resolution.

Performance on Pascal VOC 2007. For VOC07, we present the results of both our method and the comparison method in Tab. 2. For clarity and brevity, unlike previous methods [7, 30, 42], we report the same evaluation metric as those used in the COCO and do not provide the Average Precision (AP) for each class. Table 2 illustrates that in the same settings, our method generally outperforms other approaches in mAP, OF1, and CF1 at both 224×224 and 448×448 resolutions. Specifically, at the 224×224 resolution, our method shows improvements of 0.7%-8.9% in mAP, 0.7%-8.6% in

Table 2. Comparison of our method with SOTA models on VOC07 at 224×224 and 448×448 resolution. All metrics are in %.

Method	Backbone	Resolution: 224			Resolution: 448		
		mAP	CF1	OF1	mAP	CF1	OF1
VPT	ViT-B	90.4	83.4	85.7	93.4	87.4	89.1
GateVPT		87.4	80.9	82.9	89.2	83.1	85.2
E2VPT		85.8	79.4	83.7	92.4	86.5	88.3
Ours		92.9	86.4	88.1	94.3	88.6	90.1
VPT	MAE	80.3	74.1	79.2	82.8	75.9	81.6
GateVPT		85.1	75.4	80.5	86.1	79.2	84.2
E2VPT		83.3	76.7	81.5	84.6	78.2	83.2
Ours		89.2	82.7	86.4	90.4	83.9	87.4
VPT	MoCo v3	86.2	79.7	83.6	88.4	82.3	85.7
GateVPT		83.5	77.1	81.2	87.7	81.1	84.7
E2VPT		85.9	79.2	83.1	88.8	82.4	85.6
Ours		91.1	84.2	86.8	92.7	86.3	88.6
VPT	ViT-B-21k	94.2	88.0	89.6	94.4	88.7	90.4
GateVPT		94.3	88.5	89.7	95.2	89.6	91.1
E2VPT		94.1	87.9	89.4	95.1	89.3	90.7
Ours		95.0	89.2	90.3	95.6	90.0	91.3
VPT	DINOv2/B	95.6	90.2	91.6	96.1	91.0	92.3
GateVPT		95.5	89.7	91.2	96.2	91.1	92.5
E2VPT		95.4	89.9	91.4	96.1	91.1	92.4
Ours		96.4	91.0	92.3	97.0	92.1	93.1
VPT	DINOv2/S	92.8	86.4	88.5	93.8	88.1	90.1
GateVPT		92.1	85.5	87.7	93.3	87.3	89.3
E2VPT		89.6	83.1	86.3	91.2	84.8	88.0
Ours		94.3	88.3	90.3	95.7	90.2	92.0

Table 3. Comparison of our method with SOTA models on NUS and VG256 at 224×224 resolution. All metrics are in %.

Method	Backbone	NUS			VG256		
		mAP	CF1	OF1	mAP	CF1	OF1
VPT	ViT-B	65.2	62.8	73.8	42.2	44.0	57.5
GateVPT		62.8	61.5	72.9	41.1	43.0	56.5
E2VPT		63.9	62.2	73.7	42.2	44.2	57.6
Ours		65.7	63.3	74.1	44.2	45.8	58.8
VPT	MAE	60.3	59.3	73.0	37.9	40.3	55.2
GateVPT		58.2	57.6	71.9	35.3	38.2	53.2
E2VPT		60.6	59.4	73.0	37.6	40.2	55.0
Ours		61.9	60.7	73.5	40.5	42.7	57.3
VPT	MoCo v3	62.7	61.3	73.8	40.4	42.6	56.5
GateVPT		61.8	60.5	72.9	39.6	41.8	55.8
E2VPT		62.8	61.2	73.5	40.4	42.5	56.5
Ours		63.0	61.1	73.6	41.0	43.1	57.1
VPT	ViT-B-21k	67.5	64.8	74.7	46.0	47.4	60.2
GateVPT		66.9	64.4	74.3	45.4	46.8	59.5
E2VPT		67.7	64.9	74.7	46.0	47.4	60.1
Ours		68.2	65.0	75.0	47.6	48.6	61.2
VPT	DINOv2/B	68.1	65.1	75.0	49.3	50.3	62.9
GateVPT		67.2	64.6	74.5	48.6	49.7	62.4
E2VPT		67.9	65.0	75.0	49.2	50.3	62.9
Ours		68.7	65.5	75.2	50.9	51.7	64.1
VPT	DINOv2/S	65.3	63.2	74.3	45.2	46.7	60.2
GateVPT		64.5	62.6	73.7	44.3	45.9	59.3
E2VPT		65.4	63.1	74.3	45.1	46.7	60.1
Ours		66.9	64.3	74.8	47.7	48.7	61.9

CF1, and 0.6%-7.2% in OF1 compared to SOTA methods. Similarly, at the 448×448 resolution, the enhancements are 0.4%-7.6% in mAP, 0.4%-8.0% in CF1, and 0.2%-5.8% in OF1 compared to SOTA methods. These results robustly confirm the effectiveness of the proposed method.

Performance on NUS-WIDE and Visual Genome. In Tab. 3, we show the experimental results for NUS and VG256. For the VG256 setup, we follow previous work [49]. Our method achieves the best performance and outperforms SOTA methods in mAP, CF1, and OF1, under the same pre-training model and resolution setting.

5.3. Diagnostic Experiments

Impact of Different Components. To investigate the impact of different components of our method, including Group-Level Visual Prompt Tuning (GVPT) and Group-Level Mixture-of-Experts (GMoE), we conducted experi-




Figure 3. The analysis of different components. VPT stands for the vanilla model, GVPT&GMoE combines both GVPT and GMoE. Note that each model uses 30 additional prompt tokens.




Figure 4. The analysis of different grouping strategies. VPT is vanilla mode; VPT-CO indicates the VPT uses the CO groups. VPT-DC applies DC groups in VPT. GVPT incorporates both. Note that each model uses 10 additional prompt tokens and excludes MoE.

ments on two benchmark datasets: COCO and VOC07. The Fig. 3 illustrates the performance enhancements achieved by GVPT on the COCO, with improvements of 1.41% in mAP, 1.22% in CF1, and 1.1% in OF1, respectively. Moreover, incorporating GMoE into GVPT further improves performance by 0.99%, 0.72%, and 0.49%, respectively, compared to GVPT alone. On the VOC07, GVPT led to improvements of 1.77%, 1.97%, and 1.55% in mAP, CF1, and OF1, respectively. However, GMoE has limited improvement in model performance, e.g., 0.48% in mAP. These results reveal that the grouping classes in VOC07 have a relatively weak ability to learn label correlations within each group, as each group contains relatively few classes, and the distinctions between multiple tokens within each group may be minimal.

Grouping Strategies. To study the impact of the grouping strategy, we use four settings: vanilla VPT, using only the CO group, using only the DC group, and using both. As can be seen from Fig. 4, on COCO and VOC07, using the CO group and the DC group, respectively, has a certain improvement. When using both modes, the model is improved even more. We analyze that emphasizing correlative and discriminative relationships among labels has a certain gain on the MLC. While emphasizing the co-occurrence relationship, cleverly using the discriminative relationship will further compete for the model’s performance.

Number of Groups. In this study, our main idea is to group classes within the co-occurrence graph. We then aim to establish both a CO and DC relationship. Hence, additional research is required to explore the impact of group quantity




Figure 5. The performance curve varies with the increase in the number of groups. More can be found in Appendix.




Figure 6. The performance curve varies with the increase in the number of experts, in the mAP, More can be found in Appendix.

on model performance. As illustrated in Fig. 5, we present comprehensive ablation studies on the COCO and VOC07. For the COCO, model performance improves with the increased number of groups when the same pre-trained model is used. Notably, when pre-training with MoCo v3, optimal performance is achieved with 5 groups. For VOC07, optimal results are achieved when the number of groups is set to either 5 or 8. The potential reason is the dataset’s relatively small number of 20 classes, which likely impedes the effective learning of label correlations in the CO group. Additionally, the average number of labels per image in VOC07 is only 1.5, leading to weaker correlations between labels in many images compared to those in COCO.

Number of Experts. The number of experts corresponds to the number of group-level representatives within each class group, which reflects the capacity to model label relationships within each group. Accordingly, we study the influence of the number of experts on model performance. As shown in Fig. 6, on the COCO, model performance marginally improves as the number of experts increases. However, the performance on VOC07 shows slight fluctuations. This phenomenon may be attributed to the small number of classes in each group, which does not require more prompt tokens to learn the relationship between labels. Therefore, adding more experts yields limited benefits and could potentially detract from model performance.

Table 4. Computation cost and parameters number on COCO.

Method	# Total param.	# Learnable param.	FLOPs	mAP
Linear Probing	86.64 M	0.06 M	21.96 G	77.8%
Fine Tuning	86.64 M	86.64 M	21.96 G	84.2%
VPT	86.91 M	0.33 M	24.51 G	86.1%
E2VPT	87.01 M	0.43 M	24.55 G	86.3%
GateVPT	86.67 M	0.08 M	24.51 G	85.6%
GVPT (Ours)	86.98 M	0.39 M	24.51 G	87.2%
ML-VPT (Ours)	87.60 M	1.02 M	24.51 G	87.5%

Computation Cost and Parameters. To verify the first

Image	Predictions	Image	Predictions
	$p(\text{car} \text{traffic light}) = 0.62$ VPT: traffic light <input checked="" type="checkbox"/> car <input checked="" type="checkbox"/> ML-VPT: traffic light <input checked="" type="checkbox"/>		$p(\text{sink} \text{toilet}) = 0.51$ VPT: toilet <input checked="" type="checkbox"/> sink <input checked="" type="checkbox"/> ML-VPT: toilet <input checked="" type="checkbox"/>
	$p(\text{car} \text{stop sign}) = 0.45$ VPT: stop sign <input checked="" type="checkbox"/> car <input checked="" type="checkbox"/> ML-VPT: stop sign <input checked="" type="checkbox"/>		$p(\text{person} \text{cow}) = 0.32$ VPT: cow <input checked="" type="checkbox"/> person <input checked="" type="checkbox"/> ML-VPT: cow <input checked="" type="checkbox"/>

Figure 7. Visualization of model predictions on COCO. Each example includes an input image, the co-occurrence probability, and the predictions of the two models.

issue mentioned in §1, we use DINOv2/B as a pre-trained model in several methods to compare the computational overhead and parameters. As shown in Tab. 4, FineTuning not only requires more learnable parameters but also reduces model performance. Our method achieves the highest result with limited computational resources and parameters.

5.4. Case Study

To verify that our method can alleviate the co-occurrence overfitting problem, we select label pairs with high co-occurrence probabilities to study. As shown in Fig. 7, ML-VPT is compared with VPT on COCO. Under the influence of high co-occurrence probabilities, VPT is prone to over-inference, resulting in model misprediction. For example, in the second example, when toilet appears, VPT infers that there is a sink through co-occurrence classes and other contextual information (the sink has been removed). However, our method can correctly predict all labels. Similarly, in outdoor scenes, traffic light and stop sign are both misclassified as car, while in animal-related classes, the presence of a cow leads to the false inference of a person. These examples empirically demonstrate that ML-VPT can mitigate co-occurrence overfitting.

6. Conclusion

In MLC, mainstream methods tend to learn co-occurrence relationships between labels. However, an overemphasis on these relationships can lead to model overfitting. To address this issue, we introduce a novel Multi-Label Visual Prompt Tuning method designed to balance both correlative and discriminative relationships among labels. Our method offers two advantages: i) We employ the concept of grouping classes to independently model the correlative and discriminative relationships between labels, thus reducing the risk of overfitting. ii) While the grouping idea allows multiple classes to share the same group information, we further enhance our approach by integrating a group-aware MoE model that dynamically maps optimal multiple group-aware representations to a label-aware representation for each class.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China (No.62076005, No.62472004, No.61876002), the Provincial Quality Project of Education in the New Era in 2023 (Postgraduate Education No.2023lhpyfsjd009) and the University Synergy Innovation Program of Anhui Province, China (GXXT-2021-002, GXXT-2022-029, GXXT-2022-031). We thank the reviewers for their comments and suggestions. We also thank the High-performance Computing Platform of Anhui University for providing computational resources for this project.

References

- [1] Emanuel Ben-Baruch, Tal Ridnik, Itamar Friedman, Avi Ben-Cohen, Nadav Zamir, Asaf Noy, and Lihi Zelnik-Manor. Multi-label classification with partial annotations using class-aware selective loss. In *CVPR*, pages 4764–4772, 2022. 5, 6
- [2] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation via sub-category exploration. In *CVPR*, pages 8991–9000, 2020. 1
- [3] Shikai Chen, Jianfeng Wang, Yuedong Chen, Zhongchao Shi, Xin Geng, and Yong Rui. Label distribution learning on auxiliary label space graphs for facial expression recognition. In *CVPR*, pages 13984–13993, 2020. 1
- [4] Shaoxiang Chen, Zequn Jie, and Lin Ma. Llava-mole: Sparse mixture of lora experts for mitigating data conflicts in instruction finetuning mllms. *arXiv preprint arXiv:2401.16160*, 2024. 2
- [5] Tianshui Chen, Muxin Xu, Xiaolu Hui, Hefeng Wu, and Liang Lin. Learning semantic-specific graph representation for multi-label image recognition. In *ICCV*, pages 522–531, 2019. 1, 2
- [6] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021. 6
- [7] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo. Multi-label image recognition with graph convolutional networks. In *CVPR*, pages 5177–5186, 2019. 1, 2, 4, 5, 6
- [8] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yan-Tao Zheng. Nus-wide: A real-world web image database from national university of singapore. In *CIVR*, Santorini, Greece., 2009. 5
- [9] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPRW*, pages 702–703, 2020. 6
- [10] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 6
- [11] Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3):220–235, 2023. 2
- [12] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 6
- [13] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 111:98–136, 2015. 5
- [14] Zihao Fu, Haoran Yang, Anthony Man-Cho So, Wai Lam, Lidong Bing, and Nigel Collier. On the effectiveness of parameter-efficient fine-tuning. In *Proceedings of the AAAI conference on artificial intelligence*, pages 12799–12807, 2023. 2
- [15] Qiankun Gao, Chen Zhao, Yifan Sun, Teng Xi, Gang Zhang, Bernard Ghanem, and Jian Zhang. A unified continual learning framework with general parameter-efficient tuning. In *ICCV*, pages 11483–11493, 2023. 2
- [16] Yunhe Gao, Xingjian Shi, Yi Zhu, Hao Wang, Zhiqiang Tang, Xiong Zhou, Mu Li, and Dimitris N Metaxas. Visual prompt tuning for test-time domain adaptation. *arXiv preprint arXiv:2210.04831*, 2022. 2
- [17] Cheng Han, Qifan Wang, Yiming Cui, Zhiwen Cao, Wenguan Wang, Siyuan Qi, and Dongfang Liu. E²vpt: An effective and efficient approach for visual prompt tuning. In *ICCV*, pages 17445–17456, 2023. 2, 6
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 6
- [19] Lingyi Hong, Shilin Yan, Renrui Zhang, Wanyun Li, Xinyu Zhou, Pinxue Guo, Kaixun Jiang, Yiting Chen, Jinglun Li, Zhaoyu Chen, et al. Onetracker: Unifying visual object tracking with foundation models and efficient tuning. In *CVPR*, pages 19079–19091, 2024. 2
- [20] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. 2
- [21] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, pages 709–727, 2022. 2, 3, 6
- [22] Beomyoung Kim, Joonsang Yu, and Sung Ju Hwang. Eclipse: Efficient continual learning in panoptic segmentation with visual prompt tuning. In *CVPR*, pages 3346–3356, 2024. 2
- [23] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123:32–73, 2017. 5
- [24] Bo Li, Yifei Shen, Jingkang Yang, Yezhen Wang, Jiawei Ren, Tong Che, Jun Zhang, and Ziwei Liu. Sparse mixture-of-experts are domain generalizable learners. In *ICML*, 2023. 2, 5
- [25] Shikun Li, Xiaobo Xia, Hansong Zhang, Yibing Zhan, Shiming Ge, and Tongliang Liu. Estimating noise transition matrix with label correlations for noisy multi-label learning. *NeurIPS*, 35:24184–24198, 2022. 2
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and Larry Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5

- [27] Fangyi Liu, Mang Ye, and Bo Du. Learning a generalizable re-identification model from unlabelled data with domain-agnostic expert. *Visual Intelligence*, 2(1):28, 2024. 2
- [28] Ruyang Liu, Hao Liu, Ge Li, Haodi Hou, Tinghao Yu, and Tao Yang. Contextual debiasing for visual recognition with causal mechanisms. In *CVPR*, pages 12755–12765, 2022. 2
- [29] Ruyang Liu, Jingjia Huang, Thomas H Li, and Ge Li. Causality compensated attention for contextual biased visual recognition. In *ICLR*, 2023. 2
- [30] Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu. Query2label: A simple transformer way to multi-label classification. *arXiv preprint arXiv:2107.10834*, 2021. 1, 2, 4, 6
- [31] Weiwei Liu, Haobo Wang, Xiaobo Shen, and Ivor W Tsang. The emerging trends of multi-label learning. *IEEE TPAMI*, 44(11):7955–7974, 2021. 1
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. 6
- [33] Leilei Ma, Dengdi Sun, Lei Wang, Haifeng Zhao, and Bin Luo. Semantic-aware dual contrastive learning for multi-label image classification. In *ECAL*, pages 1656–1663, 2023. 2, 5, 6
- [34] Leilei Ma, Hongxing Xie, Lei Wang, Yanping Fu, Dengdi Sun, and Haifeng Zhao. Text-region matching for multi-label image recognition with missing labels. In *ACM MM*, 2024. 2
- [35] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1, 6
- [36] Sunggho Park and Hyeran Byun. Fair-vpt: Fair visual prompt tuning for image classification. In *CVPR*, pages 12268–12278, 2024. 2
- [37] Tal Ridnik, Emanuel Ben-Baruch, Nadav Zamir, Asaf Noy, Itamar Friedman, Matan Protter, and Lihi Zelnik-Manor. Asymmetric loss for multi-label classification. In *CVPR*, pages 82–91, 2021. 6
- [38] Tal Ridnik, Gilad Sharir, Avi Ben-Cohen, Emanuel Ben-Baruch, and Asaf Noy. MI-decoder: Scalable and versatile classification head. In *WACV*, pages 32–41, 2023. 2
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115:211–252, 2015. 6
- [40] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. *arXiv preprint arXiv:1708.07120*, 2017. 6
- [41] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. Cnn-rnn: A unified framework for multi-label image classification. In *CVPR*, pages 2285–2294, 2016. 1, 2
- [42] Lei Wang, Yibing Zhan, Leilei Ma, Dapeng Tao, Liang Ding, and Chen Gong. Splicemix: A cross-scale and semantic blending augmentation strategy for multi-label image classification. *IEEE TMM*, 2025. 6
- [43] Yuzhu Wang, Lechao Cheng, Chaowei Fang, Dingwen Zhang, Manni Duan, and Meng Wang. Revisiting the power of prompt for visual tuning. In *ICML*, pages 50233–50247, 2024. 2
- [44] Zhouxia Wang, Tianshui Chen, Guanbin Li, Ruijia Xu, and Liang Lin. Multi-label image recognition by recurrently discovering attentional regions. In *ICCV*, pages 464–472, 2017. 1, 2
- [45] Yunchao Wei, Wei Xia, Min Lin, Junshi Huang, Bingbing Ni, Jian Dong, Yao Zhao, and Shuicheng Yan. Hcp: A flexible cnn framework for multi-label image classification. *IEEE TPAMI*, 38(9):1901–1907, 2015. 1
- [46] Xiangping Wu, Qingcai Chen, Wei Li, Yulun Xiao, and Baotian Hu. Adahgnn: Adaptive hypergraph neural networks for multi-label image classification. In *ACM MM*, pages 284–293, 2020. 1, 4
- [47] Ming-Kun Xie and Sheng-Jun Huang. Partial multi-label learning with noisy label identification. *IEEE TPAMI*, 44(7):3676–3687, 2021. 1
- [48] Ming-Kun Xie, Jiahao Xiao, Hao-Zhe Liu, Gang Niu, Masashi Sugiyama, and Sheng-Jun Huang. Class-distribution-aware pseudo-labeling for semi-supervised multi-label learning. *NeurIPS*, 36:25731–25747, 2023. 2
- [49] Ming-Kun Xie, Jia-Hao Xiao, Pei Peng, Gang Niu, Masashi Sugiyama, and Sheng-Jun Huang. Counterfactual reasoning for multi-label image classification via patching-based training. In *ICML*, 2024. 2, 5, 7
- [50] Jiazhi Xu, Sheng Huang, Fengtao Zhou, Luwen Huangfu, Daniel Zeng, and Bo Liu. Boosting multi-label image classification with complementary parallel self-distillation. In *IJCAI*, 2022. 2, 3, 4
- [51] Jin Ye, Junjun He, Xiaojiang Peng, Wenhao Wu, and Yu Qiao. Attention-driven dynamic graph convolutional network for multi-label image recognition. In *ECCV*, pages 649–665. Springer, 2020. 2
- [52] Zihao Yin, Chen Gan, Kelei He, Yang Gao, and Junfeng Zhang. Hybrid sharing for multi-label image classification. In *ICLR*, 2024. 2
- [53] Seungryong Yoo, Eunji Kim, Dahuin Jung, Jungbeom Lee, and Sungroh Yoon. Improving visual prompt tuning for self-supervised vision transformers. In *ICML*, pages 40075–40092, 2023. 2, 6
- [54] Wang Zhan and Min-Ling Zhang. Inductive semi-supervised multi-label learning with co-training. In *KDD*, pages 1305–1314, 2017. 4
- [55] Min-Ling Zhang and Jun-Peng Fang. Partial multi-label learning via credible label elicitation. *IEEE TPAMI*, 43(10):3587–3599, 2020. 1
- [56] Min-Ling Zhang and Zhi-Hua Zhou. A review on multi-label learning algorithms. *IEEE TKDE*, 26(8):1819–1837, 2013. 1
- [57] Yuqi Zhang, Xiucheng Li, Hao Xie, Weijun Zhuang, Shihui Guo, and Zhijun Li. Multi-label action anticipation for real-world videos with scene understanding. *IEEE TIP*, 2024. 1
- [58] Zhuobin Zheng, Chun Yuan, Xinrui Zhu, Zhihui Lin, Yangyang Cheng, Cheng Shi, and Jiahui Ye. Self-supervised mixture-of-experts by uncertainty estimation. In *AAAI*, pages 5933–5940, 2019. 2