

Regularized Mask Tuning: Uncovering Hidden Knowledge in Pre-trained Vision-Language Models

Kecheng Zheng^{†,1,2} Wei Wu^{†,4} Ruili Feng⁴ Kai Zhu⁴ Jiawei Liu⁴
Deli Zhao³ Zheng-Jun Zha⁴ Wei Chen¹ Yujun Shen²

¹State Key Lab of CAD&CG, Zhejiang University ²Ant Group ³Alibaba Group ⁴USTC

Abstract

Prompt tuning and adapter tuning have shown great potential in transferring pre-trained vision-language models (VLMs) to various downstream tasks. In this work, we design a new type of tuning method, termed as **regularized mask tuning**, which masks the network parameters through a learnable selection. Inspired by neural pathways, we argue that the knowledge required by a downstream task already exists in the pre-trained weights but just gets concealed in the upstream pre-training stage. To bring the useful knowledge back into light, we first identify a set of parameters that are important to a given downstream task, then attach a binary mask to each parameter, and finally optimize these masks on the downstream data with the parameters frozen. When updating the mask, we introduce a novel gradient dropout strategy to regularize the parameter selection, in order to prevent the model from forgetting old knowledge and overfitting the downstream data. Experimental results on 11 datasets demonstrate the consistent superiority of our method over previous alternatives. It is noteworthy that we manage to deliver **18.73%** performance improvement compared to the zero-shot CLIP via masking an average of only **2.56%** parameters. Furthermore, our method is synergistic with most existing parameter-efficient tuning methods and can boost the performance on top of them. Project page can be found [here](#).

1. Introduction

The advent of large-scale pre-trained vision-language models (VLMs) [30] has ushered in a new era of incorporating language features to supervise the image encoder for a wide range of downstream visual tasks, such as few-shot learning [47] and open-world detection [8]. Thanks to the multimodal architecture and millions of text-image pairs from the web, VLMs exhibit exceptional zero-shot transferability in downstream tasks. To further enhance

[†] indicates equal contribution.

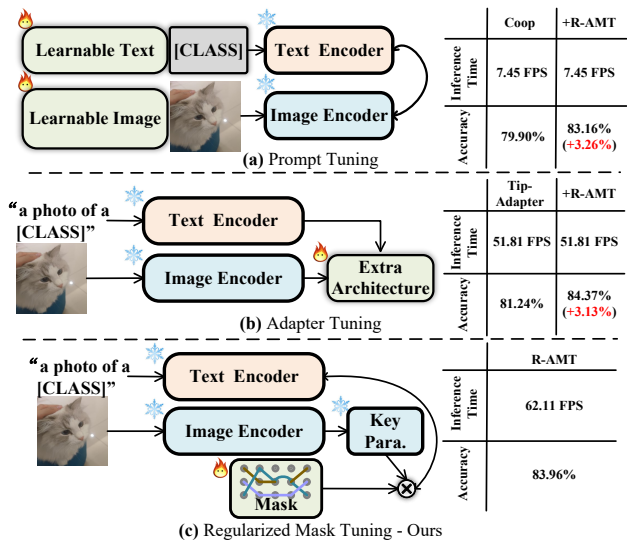


Figure 1. **Concept diagrams** of (a) prompt tuning [47, 39], (b) adapter tuning [12, 43], and (c) our *regularized mask tuning*. The tables on the right of (a)(b) demonstrate the inference time and accuracy of the existing tuning method before and after combining with our regularized mask tuning method (R-AMT). The R-AMT significantly boosts their performance without introducing additional inference time. “Key Para.” refers to the identified key parameters (e.g., multi-head self-attention). **Flames** and **snowflakes** refer to learnable and frozen parameters, respectively.

the transferability of VLMs, researchers have proposed efficient tuning methods, such as adapter tuning [12, 43] or prompt tuning [47, 39, 24]. These techniques incorporate a small number of task-specific parameters and train them solely on the downstream task, thus significantly improving the performance and reducing computational requirements.

The essence of efficient tuning methods lies in two fundamental components, *i.e.* leveraging the well-learned knowledge structure of VLMs and efficiently exploring the task-specific knowledge given few-shot data. Despite its potential, however, most existing efficient transfer learning approaches directly utilize all parameters of pre-trained VLMs and do not consider further unleashing the potential

of the well-learned knowledge of VLMs. Specifically, prompt tuning methods [47] use the frozen CLIP model and add the extra learnable parameters from the input side as shown in Fig. 1a. Adapter modules [12, 43] consist of a small set of the learnable module, further inserted into the frozen pre-trained model for adaptation as in Fig. 1b. Despite the considerable efforts in efficient tuning methods from the prompt or adapter side, these methods do not explore the frozen CLIP parameters at all, choosing instead to add additional modules to learn task-specific knowledge. Thus, as shown in Fig. 1c, we adopt **mask tuning** to explore the well-learned knowledge structure of VLMs and uncover the hidden knowledge in them for task-specific domains.

In the field of neural physiology [16, 9, 40], it has been discovered that neurons in the brain cortex exhibit diverse knowledge of various visual features such as shape, color, and depth. The knowledge is distributed in distinct neurons that have specific functions and work in conjunction with one another, termed neural pathways. When there is knowledge of a new environment coming, the neurons will compare it with the old knowledge learned in the past and then pick new conjunctions (*i.e.*, neural pathways) to adapt to the new environment. Analogous to VLMs, parameters act as a manifestation of neurons and are responsible for memorizing knowledge from data. Thus, selecting suitable parameters as parameter pathways is beneficial for uncovering the key knowledge of downstream tasks.

Inspired by the neural pathways, we propose an efficient *Regularized Mask Tuning (R-MT)* method to mask the parameters of the pre-trained VLMs under a learnable selection. Specifically, we first identify a subset of the parameters (*e.g.*, multi-head self-attentive layer) based on the magnitude of the gradient changes as sensitive network parameters for downstream tasks. Then, we introduce a binary mask equipped with gradient dropout regularization to the selected parameters. Because few-shot training tends to cause overfitting, we introduce the logits from pre-trained VLMs as the general knowledge to prevent mask tuning from forgetting. Concretely, the gradient dropout regularity as an effective regularizer introduces the probabilistic masking strategy that samples gradients based on the level of consistency of the downstream-related knowledge and the general knowledge, which can reject weak loss minima that may lead to overfitting. Intrinsic evaluations reveal that representations generated by our mask tuning can encode knowledge preferences for specific downstream tasks. The findings also indicate that selecting well-placed parameters is crucial for achieving successful transfer settings. Moreover, our method is orthogonal to most existing parameter-efficient adaption methods (*e.g.*, adapter and prompt) and endows them the ability to customization on downstream needs. Extensive experiments on 11 datasets demonstrate the effectiveness of the proposed method.

2. Related Work

Vision-language models [30, 37, 36, 22] achieve cross-modality alignment by learning a joint embedding space for text and image representation. A typical VLM consists of three components: text encoder, image encoder, and alignment function. The text and image encoder is trained separately before being connected by the alignment function in the early stage [11]. Recent VLMs such as CLIP [30] and Align [17] jointly optimize text and image encoder through contrastive learning. Benefiting from the millions of text-image pairs from the web and the multi-modality structure, these VLMs achieve exceptional zero-shot transfer capacity in the downstream tasks. Toward better transfer ability, researchers propose a series of parameter-efficient methods to adapt CLIP to downstream tasks, such as image recognition [47, 39, 12, 43, 14].

Parameter-efficient adaption methods for CLIP can be coarsely divided into two categories: prompt tuning [47, 42, 39, 4] and adapter tuning [12, 43]. Inspired by the success of prompt learning in NLP [3, 20, 13], some researchers involve prompt learning methods in CLIP to improve the few-shot transfer capacity. Zhou *et al.* [47] first introduce learnable text prompts to adapt CLIP to visual recognition tasks, which brings a great improvement over Zero-shot CLIP. Zang *et al.* [39] propose a unified prompt learning strategy for text and image encoder, which simultaneously refine the text and image representation for adapting to downstream tasks. Adapter modules consist of a small set of learnable parameters, which are further inserted into the frozen pre-trained model for adaptation. Gao *et al.* [12] add adapters after text and image branches through residual connection. Zhang *et al.* [43] employ a training-free adapter module following the image encoder, which is initialized using the knowledge extracted from the downstream training set. However, existing methods mainly focus on changing the input of CLIP (*i.e.*, text prompt tuning [47] and visual prompt tuning [39]) or adding extra modules out of CLIP (*i.e.*, adapter tuning [12, 43]), which neglects to excavate the inner power of CLIP.

Binary mask is commonly used to find a subnetwork structure from the model, which can be viewed as a way of network pruning. It can be achieved through a straight-through estimator [1, 31]. Csordás *et al.* [6] learn binary masks to identify subnetworks responsible for different tasks. Zhang *et al.* [41] search subnetworks with binary mask to achieve better out-of-distribution (OOD) performance. These works focus on finding a functional subpart of weights inside a given pre-trained neural network, which can be retrained for new tasks. However, Zhou *et al.* [45] find that applying binary mask with a model is also a way to train the model by investigating the lottery ticket hypothesis of network pruning. Recently, researchers [23, 44] propose that training binary mask for

a pre-trained language model is similar to finetuning and is more parameter-efficient. Moreover, Mallya *et al.* [26] train binary mask with fixed convolutional neural network for image classification, which achieves good performance. These works demonstrate the capacity of binary masks in parameter-efficient training in natural language processing and computer vision. Different from these methods, we propose a regularized mask tuning to search an important subset of weights in the image encoder of fixed CLIP for downstream tasks. Moreover, the regularized mask tuning can be further combined with other parameter-efficient methods presuming better few-shot performance.

3. Method

In this section, we introduce the Regularized Mask Tuning (RMT) method in detail, which aims to better adapt CLIP to downstream tasks.

3.1. Preliminaries of CLIP

CLIP [30] mainly consists of two components: image encoder $G_I(\theta)$ and text encoder $G_T(\beta)$, which are designed to project image and text into the same feature embedding space. Specifically, given the images $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$ and a set of corresponding categories, the image recognition task aims to classify an image to a specific category. Here, the m denotes the number of images in the dataset. To zero-shot adapt CLIP to the image recognition task, the name of category y_i is filled into a set of words, *e.g.* “a photo of a [class]”, to construct a hand-craft text prompt t_i as the input of text encoder. The possibility of an image \mathbf{x}_j being assigned to class y_i is formulated as following:

$$\mathbf{g}_i = G_T(t_i; \theta), \mathbf{f}_j = G_I(\mathbf{x}_j; \beta), \quad (1)$$

$$p(\mathbf{y} = i | \mathbf{x}_j) = \frac{\exp(\cos(\mathbf{g}_i, \mathbf{f}_j) / \tau)}{\sum_{n=1}^k \exp(\cos(\mathbf{g}_n, \mathbf{f}_j) / \tau)}, \quad (2)$$

where the $\cos(\cdot, \cdot)$ denotes the cosine similarity between two inputs and τ is a learnable temperature parameter.

3.2. Mask Tuning

Although CLIP has strong zero-shot performance, we argue that the knowledge required by a downstream task already exists in the pre-trained weights but may get concealed by some unnecessary information emerging in the upstream pre-training. To uncover the hidden valuable knowledge in pre-trained weights, we aim to identify the weights required by a downstream task, termed as a neural pathway, to facilitate few-shot learning [47]. Concretely, we take the parameters θ of the image encoder as an example for analysis. Given $\theta = (\theta_1, \dots, \theta_n)^T \in \Theta \subset \mathbb{R}^N$ where N is the parameter volume, our aim is to learn a binary mask

matrix M^{bin} as a downstream-related neural pathway:

$$\theta_M := \theta \odot M^{bin}, \quad (3)$$

where the \odot refers to Hadamard product and $\theta_{\mathcal{N}}^T \in \Theta_{\mathcal{N}} \subset \mathbb{R}^n$, $n \ll N$, refers to a small subset of pre-trained weights. By utilizing solely the parameters of the subset of pre-trained weights, it is adequate to transfer to the downstream domain. Since the binarized function shown in Eq. (3) is non-differentiable, we conduct a real-valued mask weight M and pass it through an element-wise thresholding binary function to obtain M^{bin} . Meanwhile, we use the gradient $\frac{\partial \mathcal{L}_{ce}}{\partial M^{bin}}$ as a noisy estimator of $\frac{\partial \mathcal{L}_{ce}}{\partial M}$ to update the M , following the previous work [44, 21]. The optimization can be formulated as:

$$M \leftarrow M - \gamma \frac{\partial \mathcal{L}_{ce}}{\partial M^{bin}}, \quad (4)$$

where the γ denotes the learning rate that controls the sparsity of mask, and \mathcal{L}_{ce} denotes the Cross-Entropy (CE) loss. The binary mask $M^{bin} = \mathcal{I}[M > \alpha]$, where α is a hard threshold. An astonishing discovery is that setting **0.16%** parameters of CLIP image encoder to 0 results in a performance improvement of **44.40%** compared to the zero-shot performance in the EuroSAT. This discovery supports the notion that certain parameters contain valuable knowledge for downstream tasks, which are also duplicated in redundant parameters. Consequently, selecting an efficient neural pathway from pre-trained weights significantly influences performance.

Which layers to apply binary mask? While this method is capable of efficiently identifying the parameter pathway that is most suitable for the downstream task, the sheer number of mask parameters to be trained can be overwhelming. As a result, it is crucial to devise a means of assessing parameter importance, by identifying the relevant neural pathway based on these significant parameters to mask. This method represents a more balanced approach, one that strikes a delicate balance between computational effort and overall performance. Our goal is to identify a subset of weights that can be effectively transferred to the downstream task while retaining important general information about the model.

To achieve this, we analyze the change in mask weight M for each layer after training on the target dataset with the CE loss, *i.e.*, $\Delta = \sum \gamma * \frac{\partial \mathcal{L}_{ce}}{\partial M}$. These parameters come from two types of layers – (1) multi-head self-attention (MSHA) layers and (2) multilayer perception (MLP) layers. We present the mean Δ when the mask weight is training on the 11 datasets in Fig. 3. As we see, the MSHA layer parameters have relatively higher Δ compared to the MLP layer. Moreover, MSHA layers are 20% of the total parameter count in the model and achieve the same

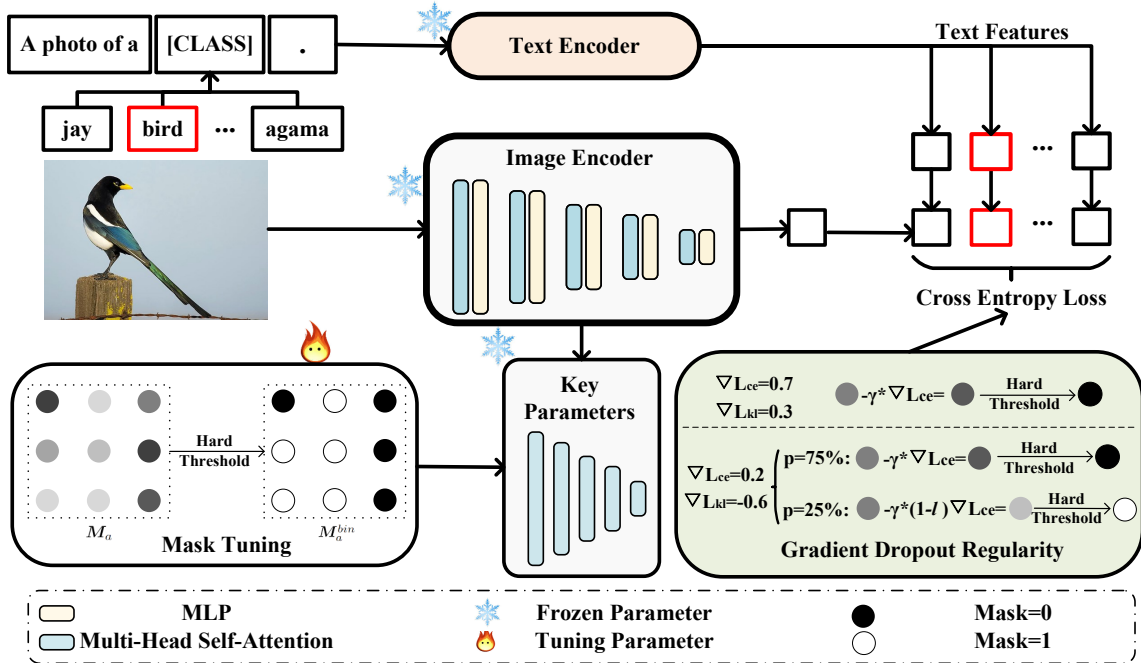


Figure 2. Overview of the proposed regularized mask tuning for frozen CLIP. During training, we maintain a set of mask weights M_a which are passed through a threshold function to obtain binary masks M_a^{bin} . There, we select the MHSA as the key parameter. When updating the mask, we introduce a novel gradient dropout strategy to regularize the parameter selection, in order to prevent the model from forgetting general knowledge from pre-trained CLIP and overfitting the downstream data.




Figure 3. Analysis of change in mask weights M when fine-tuning to downstream tasks with $\nabla\mathcal{L}_{ce}$. Over 11 datasets, the mean change in the MHSA layers is significantly higher than MLP over.

performance (e.g., 83.96% vs. 83.96% on average over 11 datasets). This suggests binary attention mask M_a^{bin} plays a significant role during fine-tuning, and we leverage that in our method as shown in Fig. 2. We name this method **Attention Mask Tuning (AMT)**. Moreover, we term binary mask on all layers as **Parameter Mask Tuning (PMT)** and MLP layers as **Multilayer perception Mask Tuning (MMT)**, respectively, for distinction.

3.3. Gradient Dropout Regularity

Stochastic Gradient Descent (SGD) [33] is a popular optimization algorithm used in machine learning to minimize a loss function during training. SGD works by randomly selecting a small subset of training examples to compute the gradient of the loss function. It can help to avoid overfitting, as it adds some level of randomness to the gradient updates. This helps to prevent the algorithm from getting stuck in local minima and encourages exploration of the solution

space. But in few-shot learning scenarios, particularly in 1-shot or 2-shot learning, the mini-batch data is typically derived from the entire training set to compute the gradient. Thus, this approach lacks the stochastic property of traditional SGD, which can lead to the overfitting of the model to the training data.

In order to make our binary tuning method better suited for few-shot scenarios, we develop Gradient Dropout Regularity formalism that randomly introduces the gradient regularity to reduce the amount of overfitting that occurs and help the model generalize better to new domains. We deem the zero-shot CLIP predictions as the general knowledge and the label from the downstream task as the target-specific knowledge. Then we introduce the Kullback-Leibler (KL) divergence between them to regularize the gradient. To implement Gradient Dropout Regularity, we first define the Gradient Retaining Purity \mathcal{P} as follows

$$\mathcal{P} = \frac{1}{2} \left(1 + \frac{\text{sgn}(\nabla\mathcal{L}_{ce})(\nabla\mathcal{L}_{ce} + \nabla\mathcal{L}_{kl})}{|\nabla\mathcal{L}_{ce}| + |\nabla\mathcal{L}_{kl}|} \right), \quad (5)$$

where \mathcal{P} is bounded by $[0, 1]$. There are two ways to describe the relationship between $\nabla\mathcal{L}_{ce}$ and $\nabla\mathcal{L}_{kl}$. Firstly, their sign is the same target-specific, implying that the optimization direction of few-shot downstream knowledge is compatible with general knowledge. Thus, we can safely update the gradient as $\nabla\mathcal{L}_{ce}$. Secondly, their signs are different at the updated position, indicating that optimizing

the binary mask with $\nabla\mathcal{L}_{ce}$ will result in the loss of pre-trained general knowledge. This implies that few-shot downstream knowledge conflicts with general knowledge. In this case, we regularize the $\nabla\mathcal{L}_{ce}$ via random gradient dropout strategy under the guidance of $\nabla\mathcal{L}_{kl}$ to optimize the model for classification. We mathematically rewrite the Eq. (5) formulated as:

$$\mathcal{P} = \begin{cases} 1, & \text{if } \nabla\mathcal{L}_{ce} \cdot \nabla\mathcal{L}_{kl} \geq 0 \\ (1 + \frac{\nabla\mathcal{L}_{ce} + \nabla\mathcal{L}_{kl}}{|\nabla\mathcal{L}_{ce}| + |\nabla\mathcal{L}_{kl}|})/2, & \text{if } \nabla\mathcal{L}_{ce} > 0 \text{ and } \nabla\mathcal{L}_{kl} < 0 \\ (1 - \frac{\nabla\mathcal{L}_{ce} + \nabla\mathcal{L}_{kl}}{|\nabla\mathcal{L}_{ce}| + |\nabla\mathcal{L}_{kl}|})/2, & \text{if } \nabla\mathcal{L}_{ce} < 0 \text{ and } \nabla\mathcal{L}_{kl} > 0. \end{cases} \quad (6)$$

Thus, \mathcal{P} is a measure of the agreement of general and target-specific knowledge at the updated position. We formulate a gradient dropout function \mathcal{M}_{ce} as: $\mathcal{M}_{ce} = \mathcal{I}[\mathcal{P} > U]$, where \mathcal{I} the standard indicator function. U is a tensor composed of i.i.d $U(0, 1)$ random variables. The optimization can be formulated as:

$$\mathbf{M}_a \leftarrow \mathbf{M}_a - \gamma * (1 - l + l * \mathcal{I}[\mathcal{P} > U]) * \frac{\partial\mathcal{L}_{ce}}{\partial\mathbf{M}_a^{bin}}, \quad (7)$$

where $l \in [0, 1]$ is a leak parameter. $l < 1$ means we allow $\nabla\mathcal{L}_{ce}$ leak through. The complete Gradient Dropout Regularity technique involves computing the purity metric \mathcal{P} at each gradient point and building a gradient consistency framework for cross-entropy loss with the help of KL divergence guidance. The steps for this are outlined in Algorithm 1. We name the AMT with the Gradient Dropout Regularity technique as R-AMT. Similarly, PMT and MMT with the Gradient Dropout Regularity technique are named R-PMT and R-MMT, respectively.

4. Experiments

4.1. Experimental Settings

Datasets. We conduct experiments on 11 publicly available image classification datasets following CoOP [47]. The datasets including ImageNet [7], FGVCaircraft [25], StanfordCars [19], Flowers102 [28], Caltech101 [10], DTD [5], EuroSAT [15], Food101 [2], UCF101 [34], OxfordPets [29], and SUN397 [38].

Implementation Details. We transfer CLIP to the few-shot image classification task with AMT and R-AMT. Specifically, we use 1, 2, 4, 8, and 16-shot training sets to optimize the model and evaluate it on the full test set, following [30]. For n -shot image classification, we random sample n images per category for training. All results reported below are the average of three runs with different random seeds. All images are resized to 224×224 . Random cropping, resizing, and random horizontal flipping strategy are used for data augmentation. We utilize ViT-B/16 as the visual backbone of CLIP. For a fair comparison, all experiments only use single text prompt, except learnable

Algorithm 1 Regularized Mask Tuning

Input: The image encoder G_I and text encoder G_T of CLIP, data \mathcal{D}_{train} for downstream task, hard threshold α , and leak parameter l .

Result: Mask \mathbf{M}_a for image encoder

- 1 Construct hand-craft text prompt set $\mathcal{T} = \{t_c\}_{c=1}^{\mathcal{C}}$ with the label set of \mathcal{D}_{train}
 - 2 Extract text features $\mathbf{g}_c = G_T(t_c), c = 1, 2, \dots, \mathcal{C}$
 - 3 Initialize the learnable mask weight \mathbf{M}_a according the the weight choices
 - 4 **for** $e \in [1, epoch]$ **do**
 - 5 Sample a mini-batch $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$ from \mathcal{D}_{train}
 - 6 Apply hard threshold α on the mask weight \mathbf{M}_a to calculate CE loss \mathcal{L}_{ce} and KL loss \mathcal{L}_{kl}
 - 7 Calculate $\mathcal{P} = \frac{1}{2} \left(1 + \frac{\text{sgn}(\nabla\mathcal{L}_{ce})(\nabla\mathcal{L}_{ce} + \nabla\mathcal{L}_{kl})}{|\nabla\mathcal{L}_{ce}| + |\nabla\mathcal{L}_{kl}|} \right)$
 - 8 Sample U from Uniform Distribution $U(0, 1)$
 - 9 Set final gradient
 $\nabla_{final} = (1 - l * (1 - \mathcal{I}[\mathcal{P} > U])) \nabla\mathcal{L}_{ce}$
 - 10 Optimize the learnable matrix \mathbf{M}_a with ∇_{final} by gradient descent: $\mathbf{M}_a \leftarrow \mathbf{M}_a - \gamma \nabla_{final}$
 - 11 **end**
-

text prompt methods, *e.g.*, for ImageNet and SUN397, the text prompt is set to be “a photo of a [class].” We adopt Adam optimizer for optimization. The mask weights are initialized element-wise with 10^{-2} . The threshold α is set to be 5×10^{-3} . The l in Eq. (7) is set to 0.3 for datasets except for ImageNet, SUN397, and Food101 in 16-shot experiments. And l is set to 1.0 in other experiments.

4.2. Comparison to State-of-the-Art Methods

Main Results on 11 Datasets. We compare AMT and R-AMT with Zero-shot CLIP and five state-of-the-art methods on the 11 datasets as mentioned above, demonstrated in Fig. 4. Zero-shot CLIP directly transfers to the downstream task without training. The state-of-the-art methods include prompt tuning methods, *i.e.*, CoOP [47], VPT [18], UPT [39], ProGrad [48], and adapter tuning method TIP-Adapter [43]. According to Fig. 4a, the AMT and R-AMT outperform these methods on average over 11 datasets, which approves the ability of AMT and R-AMT to transfer CLIP to the downstream tasks. R-AMT achieves better performance compared with AMT. It indicates the gradient dropout regularity formalism is able to enhance the transfer ability of mask tuning in few-shot scenarios.

Results on base-to-new generalization setting. Following CoCoOP [46], we conduct experiments on base-to-new generalization setting. Concretely, the classes are split equally into the base and new classes on each dataset. The base classes are used for training. The l is set to 1 in all base-to-new generalization experiments. The averaged results over 11 datasets are shown in Tab. 1. The numerical

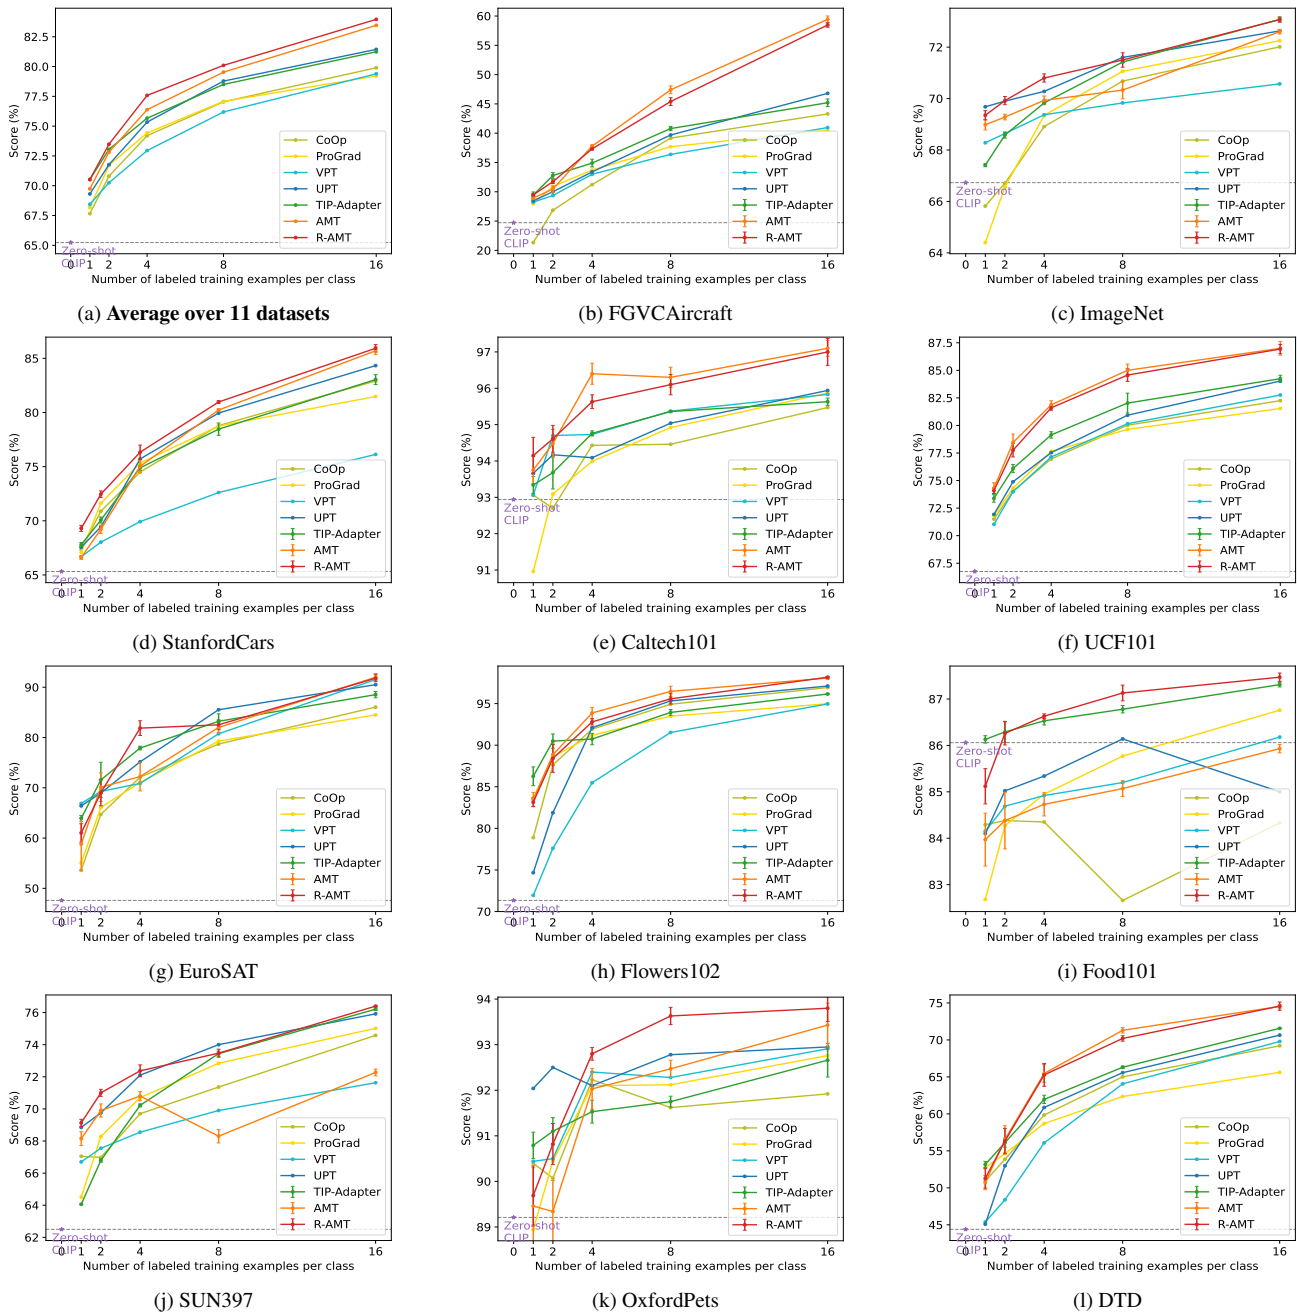


Figure 4. Accuracy (%) of few-shot learning, *i.e.*, 16/8/4/2/1-shot, on the 11 datasets. We report the average accuracy over three runs. For AMT, R-AMT, and TIP-Adapter, we demonstrate the *error bar* in all figures.

experimental results on each dataset are shown in Supplementary. Overall, R-AMT reaches the best performance, which surpasses the second best method CLIP-Adapter [12] 2.00% on the harmonic mean on average. Notably, the AMT achieves quite high performance on the base classes. But the accuracy has significantly degraded (5.11% on average) on the new classes compared with Zero-shot CLIP. We deem the degradation to be the result of overfitting since

the amount of training data is too small for some datasets, *e.g.*, Eurosat. The R-AMT achieves competitive results with AMT on base classes. However, the performance of R-AMT improves AMT by 3.04% on average in new classes, which demonstrates the anti-overfitting ability of the proposed gradient dropout regularity formalism.

The robustness to distribution shift. We evaluate the out-of-distribution (OOD) ability of AMT and R-AMT by

Table 1. Comparison on the base-to-new generalization setting on the average over 11 datasets with 16 shots. “H” denotes the harmonic mean of the accuracy on base and new classes. Thanks to gradient dropout regularity, R-AMT can efficiently maintain the knowledge of new classes while improving the anti-overfitting ability of the model to base classes. The error bar and performance of each dataset are provided in the supplementary materials.

Method	Base	New	H
Zero-shot CLIP	69.34	74.22	71.70
CoCoOP [46]	80.47	71.69	75.83
ProGrad [48]	82.79	68.55	75.00
CLIP-adapter [12]	82.62	70.97	76.35
AMT	86.17	69.11	<u>76.70</u>
R-AMT	<u>85.71</u>	<u>72.15</u>	78.35

training them on ImageNet and evaluating on ImageNet-V2 [32] and Imagenet-Sketch [35], following [43]. The evaluating datasets have compatible categories with the training set. But the three datasets are different in semantics. The OOD experimental results are shown in Tab. 2. R-AMT achieves the best performance, which surpasses TIP-Adapter [43] 1.06% on ImageNet-V2 and surpasses CoOP [47] 0.04% on Imagenet-Sketch. This indicates the R-AMT is also capable of OOD tasks. Moreover, R-AMT boosts AMT 0.47%, 0.94%, and 0.91% on ImageNet, ImageNet-V2, and Imagenet-Sketch, respectively. It further proves that R-AMT benefits from the gradient dropout regularity technique in terms of enhancing transfer and anti-overfitting ability.

4.3. Combination with State-of-the-Art Methods

To prove the R-AMT is synergistic to existing parameter-efficient methods, we combine it with CoOP [47] and TIP-Adapter [43] on 11 datasets with 16 shots, as shown in Tab. 3. Concretely, we first load the binary masks trained with R-AMT and multiply them with the original parameters of the image encoder of CLIP. Then we train the learnable contextual prompt or adapter following CoOP and TIP-Adapter. Particularly, the few-shot training set for R-AMT, CoOP+R-AMT, and TIP-Adapter+R-AMT is the same. For CoOP+R-AMT, the learned text prompt is randomly initialized and the length of the text prompt is set to 16, using the same training details as CoOP [47]. The CoOP+R-AMT boosts the performance of CoOP by 3.26% on average. This indicates the R-AMT provides a more reliable image encoder for learning better text prompts using CoOP. In addition, this combination approach directly uses a mask that is optimized by hand-craft text and does not update this mask for the learnable text prompts from CoOP, resulting in not completely unleashing the potential of the mask for downstream tasks (*i.e.*, not surpass the R-AMT). For TIP-Adapter+R-AMT, the training details are also the same as the TIP-Adapter [43]. The TIP-Adapter+R-AMT improves TIP-Adapter 3.13% on average with 16

Table 2. Comparison on robustness to distribution shift.

Method	Source ImageNet	Target		Average
		-V2	-Sketch	
Zero-shot CLIP	66.73	60.83	46.15	57.90
Linear probe	65.85	56.26	34.77	52.29
CoOP	71.73	64.56	47.89	61.39
CLIP-adapter	71.77	63.97	46.27	60.67
TIP-adapter	73.08	64.85	46.76	<u>61.56</u>
AMT	<u>72.60±0.12</u>	<u>64.97±0.11</u>	<u>47.02±0.13</u>	61.53
R-AMT	<u>73.07±0.10</u>	65.91±0.34	47.93±0.26	62.30

shots. This verifies the ability of R-AMT to endow existing parameter-efficient methods with the ability to better adapt to the downstream task.

4.4. Ablation Studies

Analysis of Masking different layers. We conduct ablation studies on masking different layers of the image encoder. Tab. 4 shows the results on 16 shots over 11 datasets. Concretely, we apply the binary masks on all weight matrices of convolutional layers and fully connected layers when training R-PMT. For R-AMT, the binary masks are applied on the multi-head self-attention (MHSA) layers, while for R-MMT, the binary masks are applied on the multilayer perceptron (MLP) layers. R-AMT achieves equal performance with R-PMT on the average of 11 datasets, which surpasses R-MMT 0.57%. But the R-AMT only uses 6.7M for storing the trained model, which is 12.3M less than R-PMT. Moreover, we find the R-AMT achieves superior performance when there are limited training classes, *e.g.*, EuroSAT. Thus, we deem the R-AMT to be a more practical method.

Influence of gradient dropout regularity. We explore the influence of gradient dropout regularity with 16 shots ImageNet. The experimental results are shown in Tab. 5. The proposed gradient dropout regularity requires the guidance of KL divergence. Thus, we conduct an ablation study by directly adding the KL loss \mathcal{L}_{kl} with the cross-entropy loss \mathcal{L}_{ce} to training the binary mask, which is termed as AMT+KL loss. It shows that if we directly add these two losses, the accuracy drops by 0.68% on 16-shot ImageNet. Because the \mathcal{L}_{ce} aims to transfer the model to downstream tasks, while the \mathcal{L}_{kl} requires the disparity between the classification logits of AMT and CLIP is not large. Directly adding \mathcal{L}_{kl} with \mathcal{L}_{ce} limits the transfer ability of AMT. AgreeGrad [27] adopts gradient surgery to solve the domain conflict, but it excessively believes in previous knowledge from KL loss, resulting in performance degradation. Recently, ProGrad [48] proposes a gradient projection method for training text prompts. This gradient projection method and our gradient dropout regularity technique both require the guidance of KL divergence. Thus, we employ the gradient projection method to train our AMT

Table 3. **Combining with the state-of-the-art methods on 16 shots.** Our mask tuning is synergistic with most existing parameter-efficient tuning methods (*e.g.*, adapter tuning [43] and prompt tuning [47]) and can boost about 3% performance on top of them.

Method	ImageNet	Caltech101	FGVCAircraft	StanfordCars	Flowers102	OxfordPets	Food101	DTD	EuroSAT	UCF101	SUN397	Average	Gain
Zero-shot CLIP	66.73	92.94	24.72	65.32	71.34	89.21	86.06	44.39	47.60	66.75	62.50	65.23	-
R-AMT	73.07	97.00	58.47	85.93	98.17	93.80	87.47	74.57	91.80	86.93	76.40	83.96	+18.73
CoOP [47]	72.01	95.47	43.29	82.91	96.93	91.92	84.33	69.21	86.05	82.25	74.58	79.90	-
CoOP+R-AMT	73.35	96.70	56.37	85.63	97.83	93.20	86.13	73.03	90.20	86.87	75.45	83.16	+3.26
TIP-Adapter [43]	73.08	95.63	45.20	83.04	96.15	92.66	87.31	71.57	88.53	84.24	76.21	81.24	-
TIP-Adapter+R-AMT	74.28	96.97	61.07	86.27	97.80	94.07	87.43	74.77	91.50	86.93	76.97	84.37	+3.13

Table 4. **Effect of performing masking on different layers.** Attaching a binary mask to the multi-head self-attention layer (*i.e.*, R-AMT) achieves the same performance as R-PMT but with lower computational effort.

Method	ImageNet	Caltech101	FGVCAircraft	StanfordCars	Flowers102	OxfordPets	Food101	DTD	EuroSAT	UCF101	SUN397	Average	Storage Space
R-AMT	73.07	97.00	58.47	85.93	98.17	93.80	87.47	74.57	91.80	86.93	76.40	83.96	6.7M
R-MMT	73.52	96.77	59.57	86.43	98.07	93.83	87.40	75.73	84.07	87.70	74.23	83.39	14M
R-PMT	73.48	96.63	60.30	86.33	98.27	93.77	87.50	75.60	88.20	87.33	76.12	83.96	19M

for comparison, named AMT+ProGrad. AMT+ProGrad surpasses AMT by 0.10%, but is 0.37% lower than R-AMT ($l=1.0$). It indicates the gradient projection technique can help mask tuning. But the improvement is limited since all conflict gradients are forced to be projected in the vertical direction. The gradient dropout regularity adds some level of randomness to the gradient guided by the KL divergence, which helps the model generalize better to downstream tasks. Moreover, we analyze the influence of l in the gradient dropout regularity technique. A smaller l implies a higher probability of CE-related gradient maintenance, which divers the binary masks more sparse. The R-AMT achieves the best performance on 16-shot ImageNet when $l = 1.0$. When l is small than 1.0, The performance degradation is caused by the leak through gradients of \mathcal{L}_{ce} , which conflicts with the general knowledge of CLIP.

Sparsity and performance comparison with zero-shot CLIP. In Fig. 6, we demonstrate the absolute improvement of R-AMT compared with Zero-shot CLIP and the sparsity of binary masks on the 16-shot setting. The R-AMT boosts the performance of Zero-shot CLIP on all datasets. Significant improvements are achieved on the EuroSAT and FGVCAircraft dataset, which reach 44.20% and 33.75%, respectively. The most sparse binary mask is obtained on StanfordCars. After setting 4.77% parameters to 0, the R-AMT improves Zero-shot CLIP 20.61% on accuracy. In total, we manage to deliver **18.73%** performance improvement compared to the zero-shot CLIP via masking an average of only **2.56%** parameters. It proves that the pre-trained weights contain some unnecessary information for the downstream task, which may harm the transfer ability

Table 5. **Ablation studies on gradient dropout regularity strategy.** The proposed gradient dropout regularity can make better use of general knowledge of KL loss while exploring the knowledge of downstream data. “ l ” controls the level of agreement in CE Loss.

Method	l	Accuracy	Gain	Sparsity
Zero-shot CLIP	-	66.73	-	-
AMT	-	72.60±0.12	-	2.64
AMT+KL loss	-	71.92±0.06	-0.68	2.58
AMT+AgreeGrad [27]	-	68.82±0.09	-3.78	1.73
AMT+ProGrad [48]	-	72.70±0.22	+0.10	2.67
R-AMT	1.0	73.07±0.10	+0.47	2.45
R-AMT	0.8	72.97±0.13	+0.37	2.50
R-AMT	0.5	72.95±0.19	+0.35	2.56
R-AMT	0.3	72.87±0.05	+0.27	2.59
R-AMT	0.1	72.67±0.12	+0.07	2.61

of the pre-trained model.

Analysis of hard threshold α . We conduct ablation study on the hard threshold α with the initial value of mask weights fixed in Tab. 6. It shows the binary masks are sparser as α gets larger. The R-AMT achieves the best accuracy when $\alpha = 5 \times 10^{-3}$. We deem that some redundant information still has not been masked when $\alpha = 4 \times 10^{-3}$. Thus, this information still influences the performance of the model in the downstream task. When $\alpha = 6 \times 10^{-3}$, some valuable parameters are moved out by the binary masks. It causes performance degradation.

Applying mask tuning on text/image encoders of CLIP. To further validate the effectiveness of mask tuning on different encoders, we also adopt R-AMT on the text/image encoders of CLIP, as shown in Fig. 5. Tab. 7 demonstrates that the performance of R-AMT on the image encoder is

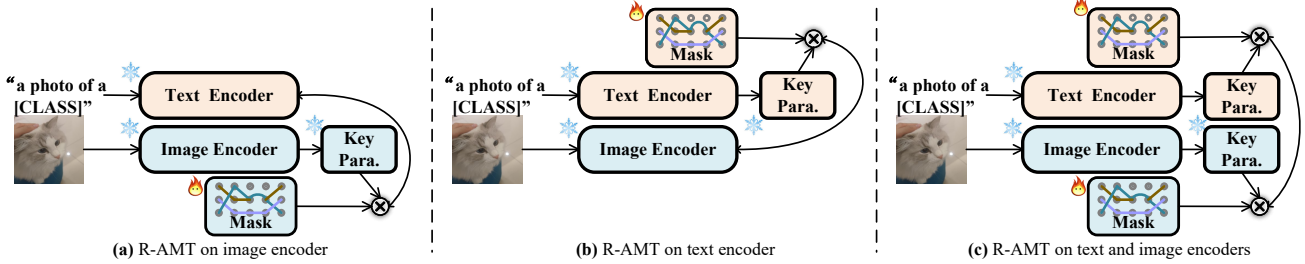


Figure 5. Applying the regularized mask tuning on different encoders of CLIP.

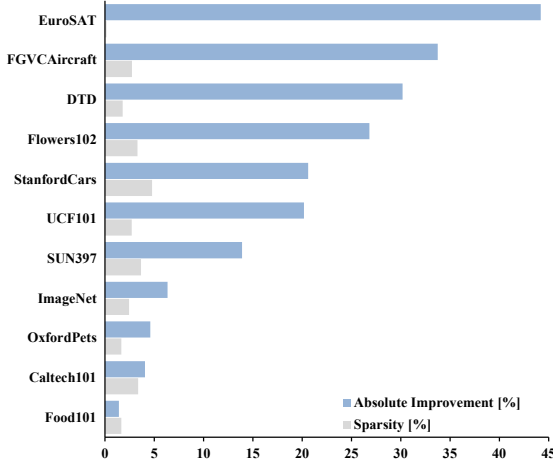


Figure 6. Comparison with Zero-shot CLIP in terms of accuracy and sparsity on 16-shot datasets. The sparsity means the percentage of the number of discarded parameters (mask=0).

Table 6. Effect of hard threshold α on 16-shot ImageNet. The threshold determines the sparsity of the model.

α	4×10^{-3}	5×10^{-3}	6×10^{-3}
Accuracy	72.87 ± 0.06	73.07 ± 0.10	72.91 ± 0.14
Sparsity	1.99	2.45	3.12

comparable to that of R-AMT on the text encoder while requiring less training time. Notably, the best performance is achieved when R-AMT is applied to both the image and text encoders. Considering the balance between training time and performance, we have chosen to adopt R-AMT on the image encoder as our method.

Table 7. Influence of applying the regularized mask tuning on different encoders of CLIP on 16-shot ImageNet.

Methods	Zero-shot CLIP	R-AMT		
		Image	Text	Image + Text
Accuracy	66.73	73.07	73.05	74.00
Error Bar	-	± 0.10	± 0.07	± 0.14
FPS	-	14.3	1.0	0.6

Different Vision Backbones. In Tab. 8, we report the results of implementing R-AMT and R-PMT with different vision backbones of CLIP on 16-shot ImageNet, including ResNet50, ResNet101, ViT-B-16, and ViT-B-32. Con-

cretely, for R-PMT, we apply the binary masks on all weight matrices of convolutional layers and fully connected layers. We observe that R-PMT achieves the best accuracy on all kinds of vision backbones on 16 shots ImageNet. When utilizing ResNet50, ResNet101, ViT-B-16, and ViT-B-32 as vision backbones, the R-PMT outperforms the second-best method by 0.68%, 1.13%, 0.40%, and 0.88%, respectively. These results demonstrate that the binary tuning is superior to the prompts tuning and adapter tuning. When using ViT as the visual backbone, R-AMT achieves competitive results with R-PMT but introduces fewer parameters. Thus, we still recommend R-AMT when ViT is the visual backbone of CLIP.

Table 8. Comparison with the state-of-the-art methods with different vision backbones on 16-shot ImageNet.

Method	ResNet50	ResNet101	ViT-B/16	ViT-B/32
Zero-shot CLIP	58.18	61.62	66.73	62.05
VPT [18]	-	-	70.57	-
CoOP [47]	62.90	66.60	71.92	66.85
CLIP-Adapter [12]	63.59	65.39	71.13	66.19
TIP-Adapter [43]	<u>64.17</u>	66.42	<u>73.08</u>	<u>67.12</u>
UPT [39]	-	-	72.63	-
PLOT [4]	63.01	-	-	-
R-AMT	-	-	73.07	67.84
R-PMT	64.85 (+0.68)	67.73 (+1.13)	73.48 (+0.40)	68.00 (+0.88)

5. Conclusion

In this work, we design a new type of tuning method, termed regularized mask tuning, that masks the network parameters under a learnable selection. Specifically, we first identify a set of parameters that are key to a given downstream task, then attach a binary mask to this parameter, and finally optimize these masks on the downstream data with the parameters frozen. When updating the mask, we introduce a novel gradient dropout strategy to regularize the parameter selection, to prevent the model from forgetting and overfitting. Extensive experiments demonstrate that our method consistently outperforms existing methods and is synergistic with them. Future work will explore applying mask tuning to other visual tasks such as segmentation.

6. Acknowledgements

Wei Chen is supported by the Fundamental Research Funds for the Central Universities (226-2022-00235).

References

- [1] Yoshua Bengio. Estimating or propagating gradients through stochastic neurons. *arXiv preprint arXiv:1305.2982*, 2013. **2**
- [2] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *Eur. Conf. Comput. Vis.*, pages 446–461, 2014. **5**
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Adv. Neural Inform. Process. Syst.*, pages 1877–1901, 2020. **2**
- [4] Guangyi Chen, Weiran Yao, Xiangchen Song, Xinyue Li, Yongming Rao, and Kun Zhang. Prompt learning with optimal transport for vision-language models. *arXiv preprint arXiv:2210.01253*, 2022. **2, 9**
- [5] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3606–3613, 2014. **5**
- [6] Róbert Csordás, Sjoerd van Steenkiste, and Jürgen Schmidhuber. Are neural nets modular? inspecting functional modularity through differentiable weight masks. *Int. Conf. Learn. Represent.*, 2020. **2**
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255, 2009. **5**
- [8] Yu Du, Fangyun Wei, Zihe Zhang, Miaojing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 14084–14093, 2022. **1**
- [9] Stephen A Engel, Gary H Glover, and Brian A Wandell. Retinotopic organization in human visual cortex and the spatial precision of functional mri. *Cerebral cortex (New York, NY: 1991)*, pages 181–192, 1997. **2**
- [10] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 178–178, 2004. **5**
- [11] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Adv. Neural Inform. Process. Syst.*, 26, 2013. **2**
- [12] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. **1, 2, 6, 7, 9**
- [13] Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020. **2**
- [14] Ziyu Guo, Renrui Zhang, Longtian Qiu, Xianzheng Ma, Xupeng Miao, Xuming He, and Bin Cui. Calip: Zero-shot enhancement of clip with parameter-free attention. *arXiv preprint arXiv:2209.14169*, 2022. **2**
- [15] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, pages 2217–2226, 2019. **5**
- [16] David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, page 106, 1962. **2**
- [17] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Int. Conf. Mach. Learn.*, pages 4904–4916, 2021. **2**
- [18] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. *arXiv preprint arXiv:2203.12119*, 2022. **5, 9**
- [19] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Int. Conf. Comput. Vis. Worksh.*, June 2013. **5**
- [20] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. **2**
- [21] Tao Lin, Sebastian U Stich, Luis Barba, Daniil Dmitriev, and Martin Jaggi. Dynamic model pruning with feedback. *arXiv preprint arXiv:2006.07253*, 2020. **3**
- [22] Xuejing Liu, Liang Li, Shuhui Wang, Zheng-Jun Zha, Dechao Meng, and Qingming Huang. Adaptive reconstruction network for weakly supervised referring expression grounding. In *Int. Conf. Comput. Vis.*, pages 2611–2620, 2019. **2**
- [23] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018. **2**
- [24] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5206–5215, 2022. **1**
- [25] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. **5**
- [26] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Eur. Conf. Comput. Vis.*, pages 67–82, 2018. **3**
- [27] Lucas Mansilla, Rodrigo Echeveste, Diego H Milone, and Enzo Ferrante. Domain generalization via gradient surgery. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6630–6638, 2021. **7, 8**
- [28] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729, 2008. **5**
- [29] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3498–3505, 2012. **5**

- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Int. Conf. Mach. Learn.*, pages 8748–8763, 2021. [1](#), [2](#), [3](#), [5](#)
- [31] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Eur. Conf. Comput. Vis.*, pages 525–542, 2016. [2](#)
- [32] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *Int. Conf. Mach. Learn.*, pages 5389–5400. PMLR, 2019. [7](#)
- [33] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951. [4](#)
- [34] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [5](#)
- [35] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. *Adv. Neural Inform. Process. Syst.*, 2019. [7](#)
- [36] Hao Wang, Zheng-Jun Zha, Liang Li, Xuejin Chen, and Jiebo Luo. Semantic and relation modulation for audio-visual event localization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022. [2](#)
- [37] Hao Wang, Zheng-Jun Zha, Liang Li, Dong Liu, and Jiebo Luo. Structured multi-level interaction network for video moment localization via language query. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7026–7035, 2021. [2](#)
- [38] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492, 2010. [5](#)
- [39] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. *arXiv preprint arXiv:2210.07225*, 2022. [1](#), [2](#), [5](#), [9](#)
- [40] Semir Zeki and Stewart Shipp. The functional logic of cortical connections. *Nature*, pages 311–317, 1988. [2](#)
- [41] Dinghui Zhang, Kartik Ahuja, Yilun Xu, Yisen Wang, and Aaron Courville. Can subnetwork structure be the key to out-of-distribution generalization? In *Int. Conf. Mach. Learn.*, pages 12356–12367, 2021. [2](#)
- [42] Renrui Zhang, Xiangfei Hu, Bohao Li, Siyuan Huang, Hanqiu Deng, Hongsheng Li, Yu Qiao, and Peng Gao. Prompt, generate, then cache: Cascade of foundation models makes strong few-shot learners. *arXiv preprint arXiv:2303.02151*, 2023. [2](#)
- [43] Renrui Zhang, Wei Zhang, Rongyao Fang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free adaption of clip for few-shot classification. In *Eur. Conf. Comput. Vis.*, pages 493–510, 2022. [1](#), [2](#), [5](#), [7](#), [8](#), [9](#)
- [44] Mengjie Zhao, Tao Lin, Fei Mi, Martin Jaggi, and Hinrich Schütze. Masking as an efficient alternative to fine-tuning for pretrained language models. *arXiv preprint arXiv:2004.12406*, 2020. [2](#), [3](#)
- [45] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. *Adv. Neural Inform. Process. Syst.*, 32, 2019. [2](#)
- [46] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16816–16825, 2022. [5](#), [7](#)
- [47] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *Int. J. Comput. Vis.*, pages 2337–2348, 2022. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#), [9](#)
- [48] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. *arXiv preprint arXiv:2205.14865*, 2022. [5](#), [7](#), [8](#)