

Weak Distribution Detectors Lead to Stronger Generalizability of Vision-Language Prompt Tuning

Kun Ding^{1,3}, Haojian Zhang², Qiang Yu³, Ying Wang^{1,3*}, Shiming Xiang^{1,3}, Chunhong Pan³

¹State Key Laboratory of Multimodal Artificial Intelligence Systems, Institute of Automation, Chinese Academy of Sciences

²Engineering Laboratory for Intelligent Industrial Vision, Institute of Automation, Chinese Academy of Sciences

³Research Center of Aerospace Information, Institute of Automation, Chinese Academy of Sciences

{kun.ding, zhanghaojian2014, qiang.yu}@ia.ac.cn, {ywang, smxiang, chpan}@nlpr.ia.ac.cn

Abstract

We propose a generalized method for boosting the generalization ability of pre-trained vision-language models (VLMs) while fine-tuning on downstream few-shot tasks. The idea is realized by exploiting out-of-distribution (OOD) detection to predict whether a sample belongs to a base distribution or a novel distribution and then using the score generated by a dedicated competition based scoring function to fuse the zero-shot and few-shot classifier. The fused classifier is dynamic, which will bias towards the zero-shot classifier if a sample is more likely from the distribution pre-trained on, leading to improved base-to-novel generalization ability. Our method is performed only in test stage, which is applicable to boost existing methods without time-consuming re-training. Extensive experiments show that even weak distribution detectors can still improve VLMs' generalization ability. Specifically, with the help of OOD detectors, the harmonic mean of CoOp (Zhou et al. 2022b) and ProGrad (Zhu et al. 2023) increase by 2.6 and 1.5 percentage points over 11 recognition datasets in the base-to-novel setting.

Introduction

Many large vision-language models (VLMs) trained on web-scale image-text pairs have emerged continuously in recent years (e.g., CLIP (Radford et al. 2021), ALIGN (Jia et al. 2021)) and they are proved to be very useful on many pure vision and multi-modal tasks, such as few-shot learning (Zhou et al. 2022b), image segmentation (Rao et al. 2022) and image captioning (Mokady, Hertz, and Bermano 2021). The huge parameter count of these models poses extreme challenges for classical fine-tuning techniques, which are both parameter inefficient and of low generalization capability. Vision-language prompt tuning (VLPT) has recently been proposed as an alternative technique (Zhou et al. 2022b,a; Zhu et al. 2023; Yao, Zhang, and Xu 2023), which is demonstrated to be more parameter efficient.

However, it is hard for VLPT to balance the performance between base and novel classes. The first VLPT method CoOp (Zhou et al. 2022b) achieves superior performance on base classes of 11 recognition datasets. Compared to zero-shot CLIP, the mean accuracy on novel classes

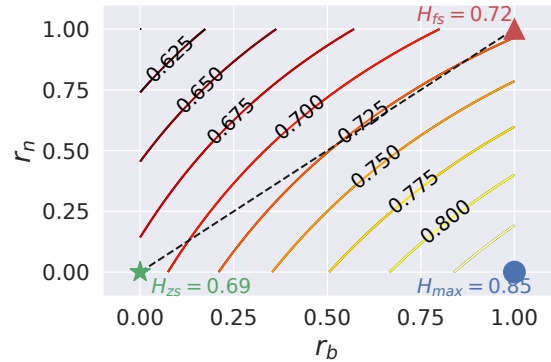


Figure 1: The relation between harmonic mean and ID prediction accuracies when $\alpha = \infty$ (in Eq. 7). r_b and $1 - r_n$ are the in-distribution prediction accuracy of samples in base and novel set. The maximum harmonic mean H_{max} is acquired when $r_b = 1$ and $r_n = 0$. H_{fs} and H_{zs} denote the harmonic mean obtained by applying the few-shot and zero-shot classifier on all test samples, respectively.

drops 11 points (from 74.22% to 63.22%). CoCoOp (Zhou et al. 2022a) alleviates the above problem by conditioning prompts on image feature. CoCoOp improves the novel class accuracy in the sacrifice of base class accuracy and the model design is quite heavy weighted. ProGrad (Zhu et al. 2023) prevents prompt tuning from forgetting the general knowledge in VLMs by only updating the prompt whose gradient is aligned to the general direction. KgCoOp (Yao, Zhang, and Xu 2023) introduces an auxiliary loss to reduce the discrepancy between the textual embeddings generated by learned prompts and hand-crafted prompts.

Different from the above methods, we propose a new viewpoint for addressing the trade-off problem between base and novel classes. This viewpoint is motivated by the fact: zero-shot CLIP excels on novel classes while CoOp excels on base classes (Zhou et al. 2022a). Therefore, if it is possible to know a sample comes from the base distribution or the novel distribution beforehand and choose the right classifier accordingly, such a trade-off problem will be addressed perfectly. The relation between the base-to-novel evaluation metric (harmonic mean) and the in-distribution (ID) prediction accuracies is studied in Fig. 1, which implies that the

*Corresponding author.

harmonic mean of the fused classifier with binary scoring function (in Eq. 7) increases with ID detection accuracies (r_b and $1 - r_n$ on base and novel set respectively). The problem of discriminating a sample from which distribution is similar to the out-of-distribution (OOD) detection problem. As such, we could convert the base-to-novel trade-off problem to an OOD detection problem.

Specifically, we design a competition based scoring function which exploits the ID scores generated by existing OOD detection methods to generate a score between 0 and 1 for a test sample and then use this score to dynamically fuse the pre-trained zero-shot classifier and the few-shot classifier tuned on limited training data. Different from the scoring functions in prior OOD detection methods, the proposed scoring function is based on the comparison between the outputs of two classifiers. The proposed fusion method is dynamic, i.e. the fusion weights vary from different samples, which is better than static method. Also, the fusion is implemented in test stage, making it easy for the integration in existing methods without re-training.

We conduct extensive experiments on 15 recognition datasets under different settings. In the base-to-novel setting, our method improves the harmonic mean accuracy of CoOp and ProGrad by 2.6% and 1.5%, respectively. In the domain generalization setting, the mean accuracies over target domains are competitive, either. The main contributions of this work are listed as follows:

- We propose a new viewpoint for dealing with the base-to-novel generalization problem of VLPT by introducing OOD detection.
- We propose a competition based scoring function for fusing the few-shot and zero-shot classifier dynamically.
- We demonstrate by experiment even weak distribution detectors can still improve the generalizability of existing VLPT methods.

Related Work

Vision-Language Models: Recent years have witnessed lots of vision-language models, such as CLIP (Radford et al. 2021), ALIGN (Jia et al. 2021), BLIP (Li et al. 2022). Benefited from massive image-text pairs, large vision-language models have demonstrated huge potential on many vision tasks, including but not limited to zero-shot learning (Radford et al. 2021; Shu et al. 2022), few-shot learning (Zhou et al. 2022b), image segmentation (Rao et al. 2022), image generation (Ramesh et al. 2022), etc. The key problem of exploiting so many pre-trained VLMs is how to fine-tune them on downstream tasks in a parameter efficient way.

Vision-Language Prompt Tuning: As one of the most popular parameter efficient tuning techniques, prompt tuning is first proposed in NLP and subsequently introduced into vision-language domain by Zhou et al. (2022b). Instead of manually designing prompts, prompt tuning learns soft context embeddings from data, which avoids cumbersome hand-crafts. Based on CoOp, lots of efforts have been made to solve the problems in certain aspects or expand application scope. For example, CoCoOp (Zhou et al. 2022a), ProGrad (Zhu et al. 2023) and KgCoOp (Yao, Zhang, and Xu

2023) are devised to reduce the loss of generalization on unseen tasks. DualCoOp (Sun, Hu, and Saenko 2022) extends CoOp to the multi-label recognition task. SoftCPT (Ding et al. 2022) exploits the cross-task knowledge to further enhance the generalization ability. PLOT (Chen et al. 2023) introduces optimal transport to achieve fine-grained matching between multiple prompts and visual regions. Prompt tuning is also introduced in image segmentation (Rao et al. 2022) and video relation detection (Gao et al. 2023).

OOD Detection: OOD detection is to identify outliers in an open sample space not belonging to any seen training classes, which is especially vital for safety-critical circumstances. According to Yang et al. (2021), OOD detection methods can be categorized into classification-based, density-based, distance-based and reconstruction-based methods. Post-hoc detection is an important member of classification-based methods that does not need extra training data. Given a trained model, post-hoc methods calculate a score to indicate the ID-ness or OOD-ness. Hendrycks and Gimpel (2017) proposed the maximum softmax probabilities (MSP), which is a widely-used baseline. Considering that MSP is problematic for large-scale settings with hundreds or thousands of classes, MaxLogit (Hendrycks et al. 2022) that uses the maximum unnormalized logit as anomaly score was proposed. Liu et al. (2020) proposed an energy based indicator that has good theoretical interpretability. More recently, Zhang and Xiang (2023) decoupled MaxLogit into cosine similarity and logit norm and introduced a tunable factor between them for better performance. Based on these works we propose a novel scoring function tailored for VLPT, which considers two classifiers jointly, i.e. zero-shot and few-shot classifier. This is different from the above methods that only use one trained classifier to design scoring function.

Dynamic Neural Networks: Designing instance-aware dynamic neural networks is an important research topic in the literature. Dynamic convolution (Chen et al. 2020; Hou and Kung 2021; Li, Zhou, and Yao 2022) learns a linear combination of multiple static kernels with input-dependent dynamic attention. Such a design effectively increases model capability without increasing FLOPs a lot. Dynamic quantization (Liu et al. 2022) learns quantization configuration for each image and layer individually, which further reduces computation cost. The Mixture-of-Expert (MoE) layer contains many expert sub-networks and learns a gating function to route individual inputs to some of the experts (Shazeer et al. 2017). It substantially increases the model capacity without introducing large computation overhead. Our proposed method dynamically fuses two classifiers, it can also be seen as a dynamic network. However, it is a test-time method as it does not need re-training.

Background

Radford et al. (2021) demonstrated the contrastive learning based vision-language model CLIP has strong zero-shot recognition ability. The good generability is believed to be benefited from natural language supervision and diverse large-scale training data. CLIP has a two-stream structure:

one image encoder for representing image as visual embedding and one text encoder for representing text as text embedding. Due to the contrastive learning based design, CLIP can be easily adapted for zero-shot and few-shot recognition. **Zero-shot CLIP.** We denote the image and text encoder in CLIP as E_{img} and E_{txt} , respectively. For a zero-shot classification task of N classes, Radford et al. (2021) used the classifier generated by the text encoder for classification. The N class names are assumed to be accessible and the classifier weight (an unnormalized column vector) associated to the i -th class ($i \in \{1, 2, \dots, N\}$) is obtained by encoding its class name, i.e., $\mathbf{w}_i^{zs} = E_{txt}(\text{“[class name]”})$. Adding a prefix string to class name is demonstrated to be useful, that is, $\mathbf{w}_i^{zs} = E_{txt}(\text{“a photo of a [class name]”})$. Collecting the weights of all classes into a matrix results in $\mathbf{W}^{zs} = [\mathbf{w}_1^{zs}, \dots, \mathbf{w}_N^{zs}]^T$. Given an image I , its visual embedding \mathbf{x} is extracted by $\mathbf{x} = E_{img}(I)$, which is an unnormalized vector with a same length to \mathbf{w}_i^{zs} . To classify I , zero-shot CLIP computes the posterior probability of a certain class y as follows:

$$p^{zs}(y|\mathbf{x}) = \exp(o_y^{zs}(\mathbf{x})) / \sum_{i=1}^N \exp(o_i^{zs}(\mathbf{x})), \quad (1)$$

where $o_y^{zs}(\mathbf{x}) = d(\mathbf{x}, \mathbf{w}_y^{zs})/\tau$ is the logit of class y with τ a temperature parameter and $d(\cdot)$ the cosine similarity.

Context Optimization (CoOp). Vision-language prompt tuning (VLPT) methods like CoOp and its successors (ProGrad, KgCoOp, etc.) extend zero-shot CLIP for few-shot classification. Meantime, the fixed prompt prefix is upgraded as a learnable prompt context. They defines a sequence of M vectors as the context denoted as a matrix $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_M]$. For the i -th class, the class token embeddings can be represented as a matrix \mathbf{C}_i . The classifier weight associated to class i is computed by encoding $[\mathbf{V}, \mathbf{C}_i]$ as a text embedding: $\mathbf{w}_i^{fs} = E_{txt}([\mathbf{V}, \mathbf{C}_i])$. Collecting the weights of all classes into a matrix results in $\mathbf{W}^{fs} = [\mathbf{w}_1^{fs}, \dots, \mathbf{w}_N^{fs}]^T$. To classify an image of encoded feature \mathbf{x} , the following posterior probability of a certain class y is calculated:

$$p^{fs}(y|\mathbf{x}) = \exp(o_y^{fs}(\mathbf{x})) / \sum_{i=1}^N \exp(o_i^{fs}(\mathbf{x})), \quad (2)$$

where $o_y^{fs}(\mathbf{x}) = d(\mathbf{x}, \mathbf{w}_y^{fs})/\tau$ is the logit of the y -th class. VLPT methods optimize the context \mathbf{V} while fixing the image and text encoder. To align with upstream task, the negative log likelihood loss is widely adopted.

Method

The flowchart of the proposed method is presented in Fig. 2. For an image, we extract the visual embedding vector. With hand-crafted prompts, we compute the zero-shot classifier. With learnable prompts, we compute the few-shot classifier. Subsequently, two posterior probability distributions, which serve as the inputs of OOD detector, are computed using these two classifiers. The OOD detector outputs in-distribution score for a test sample. The two in-distribution scores, $ids^{zs}(\mathbf{x})$ and $ids^{fs}(\mathbf{x})$, are compared by a scoring function $s(\mathbf{x})$. This final score is used to dynamically weight

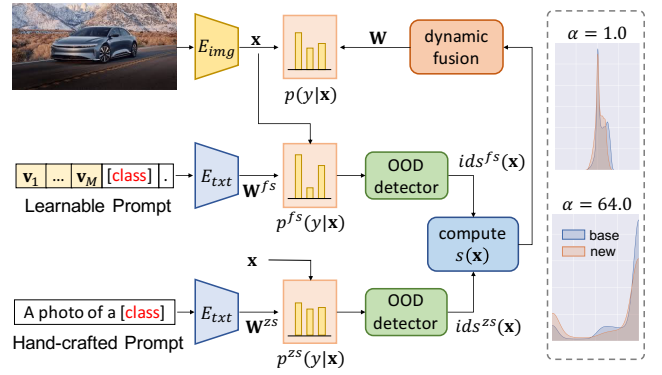


Figure 2: The framework of the proposed method.

the two classifiers. In the following, we detail the key ingredients of the proposed method. Besides, we further analyze the relation between harmonic mean accuracy and OOD detection accuracies.

OOD Detectors

Existing OOD detection methods calculate a score to reflect ID-ness or OOD-ness. By choosing a valid threshold and thresholding the score with it, binary prediction can be obtained. This work adopts post-hoc methods due to their simplicity. The original scores proposed in these method are converted to values reflecting ID-ness. We use three existing methods (MSP, MaxLogit and Energy) and also consider a new entropy based method. Actually, many other OOD detection methods can also be used.

MSP (Hendrycks and Gimpel 2017) uses the maximum softmax probability of a classifier as the ID score, that is

$$ids^*(\mathbf{x}) = \max_i p^*(i|\mathbf{x}), \quad (3)$$

where $*$ is zs or fs . $ids^{zs}(\mathbf{x})$ and $ids^{fs}(\mathbf{x})$ denote the in-distribution score corresponding to the zero-shot and few-shot classifier, respectively.

MaxLogit (Hendrycks et al. 2022) uses the maximum logit of a classifier as the ID score, that is

$$ids^*(\mathbf{x}) = \tau \cdot \max_i o_i^*(\mathbf{x}), \quad (4)$$

where τ is multiplied for normalizing the final score into the range $[-1, 1]$.

In Liu et al. (2020), the free energy function over \mathbf{x} in terms of the denominator of the softmax function is used as the OOD score. Accordingly, the negative free energy is used as the ID score:

$$ids^*(\mathbf{x}) = (\log N + 1/\tau)^{-1} \cdot \tau \cdot \log \sum_i^N \exp(o_i^*(\mathbf{x})), \quad (5)$$

where $(\log N + 1/\tau)^{-1}$ is a normalizing factor.

The entropy of a softmax probability distribution reflects the OOD-ness of a sample. Larger entropy means the sample is more likely from a different distribution. We thus use the negative entropy as the ID score:

$$ids^*(\mathbf{x}) = (\log N)^{-1} \cdot \sum_i^N p^*(i|\mathbf{x}) \log p^*(i|\mathbf{x}), \quad (6)$$

where $(\log N)^{-1}$ is a normalizing factor.

Competition Based Scoring Function

We propose a scoring function to indicate a sample is more likely from the base or novel distribution. The intuition is that if $ids^{fs}(\mathbf{x}) - ids^{zs}(\mathbf{x})$ is larger, it has a higher probability that the sample is from the base distribution (also the distribution fine-tuned on). To obtain a valid weight, the score should be normalized into $[0,1]$, thus, we introduce a sigmoid function $\sigma(x)$ to normalize the above difference. Formally, the scoring function is defined as

$$s(\mathbf{x}) = \sigma(\alpha \cdot (ids^{fs}(\mathbf{x}) - ids^{zs}(\mathbf{x}))), \quad (7)$$

where α is a tunable scaling factor, which controls the dynamic range of the score (ref. Fig. 2). Note that, applying sigmoid function over the difference equals to applying the softmax function over $ids^{fs}(\mathbf{x})$ and $ids^{zs}(\mathbf{x})$, implying the existence of competition. Note that, with a sufficiently large α , the sigmoid function equals to a Heaviside step function. In this case, it actually makes a binary prediction.

Dynamic Fusion

When a sample is more likely from the novel distribution, it is better to use the zero-shot classifier; conversely, when it is more likely from the base distribution, it is better to use the few-shot classifier. A more general method is weighted combining the few-shot and zero-shot classifier with $s(\mathbf{x})$:

$$\mathbf{W} = s(\mathbf{x}) \cdot \mathbf{W}^{fs} + (1 - s(\mathbf{x})) \cdot \mathbf{W}^{zs}. \quad (8)$$

By this way, the fused classifier is instance-conditioned, which dynamically changes with different inputs. Note that other strategies such as weighted combination of the posterior probabilities from the two classifiers are also possible.

Analysis of Harmonic Mean

The harmonic mean performance of the classifier in Eq. 8 has direct relation to OOD detection performance. In the base-to-novel evaluation setting of VLPT, there are two test sets of non-overlapping classes: base set and novel set. Let us denote r_b and r_n as the probability of predicting a sample in the base and novel set as a sample from base distribution, respectively. The classification accuracy of the few-shot classifier \mathbf{W}^{fs} and the zero-shot classifier \mathbf{W}^{zs} on base set are denoted as p_1 and p_0 , respectively. The classification accuracy of the few-shot classifier \mathbf{W}^{fs} and the zero-shot classifier \mathbf{W}^{zs} on novel set are denoted as q_1 and q_0 , respectively. We have the following proposition:

Proposition 1. *The harmonic mean accuracy H of the fused classifier \mathbf{W} with the Heaviside step function used on base and novel set can be represented as:*

$$H = 2P_b P_n / (P_b + P_n), \quad (9)$$

$$P_b = p_1 r_b + p_0 (1 - r_b), \quad (10)$$

$$P_n = q_1 r_n + q_0 (1 - r_n), \quad (11)$$

Proof. The proposition can be proved by the law of total probability. \square

We plot the contour map of H with respect to r_b and r_n in Fig. 1 with $p_0 = 0.6, p_1 = 0.9, q_0 = 0.8, q_1 = 0.6$.

Each point in the figure represents an OOD detector of certain performance. The black dashed line corresponds to random OOD detectors. When moving towards the lower right corner, H increases gradually. A relatively weak classifier (e.g., the classifiers on the 0.750 contour) can still improve the harmonic mean accuracy of the zero-shot and few-shot classifier, i.e., $H_{zs} = 0.69$ and $H_{fs} = 0.72$, respectively. This implies the feasibility of the proposed method.

Experiment

Similar to prior work, we evaluate our method under two settings: 1) base-to-novel generalization within a dataset; 2) domain generalization across datasets with the same classes.

Setup

Dataset. Following prior works (Zhou et al. 2022b), for the base-to-novel experiment, we use 11 image classification datasets: ImageNet (Deng et al. 2009) and Caltech101 (Li, Fergus, and Perona 2007) for generic object classification; Food101 (Bossard, Guillaumin, and Gool 2014), Stanford-Cars (Krause et al. 2013), OxfordPets (Parkhi et al. 2012) and Flowers102 (Nilsback and Zisserman 2008) and FGV-CAircraft (Maji et al. 2013) for fine-grained visual classification; UCF101 (Soomro, Zamir, and Shah 2012) for action recognition; EuroSAT (Helber et al. 2019) for satellite image classification; DTD (Cimpoi et al. 2014) for texture classification; SUN397 (Xiao et al. 2010) for scene recognition. For the domain generalization experiment, we use ImageNet as the source domain and other four datasets, i.e. ImageNetV2 (Recht et al. 2019), ImageNet-Sketch (Wang et al. 2019), ImageNet-A (Gao et al. 2022) and ImageNet-R (Hendrycks et al. 2021) as the target domains.

Compared Methods. We propose to use an OOD detector to enhance the generalization ability of the existing VLPT methods. To validate the versatility of our method, we use four existing methods as the baselines: CoOp (Zhou et al. 2022b), CoCoOp (Zhou et al. 2022a), ProGrad (Zhu et al. 2023) and KgCoOp (Yao, Zhang, and Xu 2023). The results of these methods with and without OOD detector will be compared. Besides, zero-shot CLIP is also compared.

Implementation Details. With the public code of CoOp, CoCoOp, ProGrad and KgCoOp, we reproduce the results on the aforementioned datasets using the suggested parameters. For CoOp, the context length is 16 and random initialization is adopted. The batch size is 32 and 50 epochs are trained. For CoCoOp, the context is initialized by “a photo of a”, batch size is 1 and 10 epochs are trained. For ProGrad, the training setting is identical to CoOp and the two extra hyper-parameters are set to 1.0. For KgCoOp, the context is initialized by “a photo of a” and the additional hyper-parameter λ is set to be 8.0. Due to limited GPU memory, the batch size is set to be 32. Besides, 100 epochs are trained. All methods use the same learning rate scheduler and optimizer. Our reproduced results are slightly different from previously reported, however, the differences are minor. For our method, we insert the proposed test-time technique into these methods and conduct model inference on test sets. The entropy based ID score and $\alpha = 64$ are used as default.

	Base	Novel	H		Base	Novel	H		Base	Novel	H
CLIP	69.16	74.13	71.56	CLIP	72.38	68.10	70.17	CLIP	94.58	94.10	94.34
CoOp	82.29	67.63	74.24	CoOp	76.46	65.47	70.54	CoOp	95.74	93.65	94.68
CoOp+	82.20	72.15	76.85	CoOp+	76.21	68.87	72.35	CoOp+	96.10	94.10	95.09
CoCoOp	80.12	72.36	76.04	CoCoOp	75.86	70.59	73.13	CoCoOp	95.85	94.45	95.14
CoCoOp+	79.26	74.37	76.74	CoCoOp+	75.16	70.65	72.84	CoCoOp+	95.85	94.25	95.04
ProGrad	82.27	70.10	75.70	ProGrad	76.73	67.64	71.90	ProGrad	96.41	93.69	95.03
ProGrad+	80.74	73.95	77.20	ProGrad+	75.79	69.37	72.44	ProGrad+	95.98	94.60	95.29
KgCoOp	81.57	72.87	76.97	KgCoOp	75.80	69.81	72.68	KgCoOp	95.89	94.74	95.31
KgCoOp+	79.81	74.26	76.94	KgCoOp+	74.67	69.57	72.03	KgCoOp+	95.68	94.61	95.14
Pro+Kg	82.43	73.31	77.60	Pro+Kg	76.79	69.84	73.15	Pro+Kg	96.18	94.81	95.49
(a) Average over 11 datasets.			(b) ImageNet.			(c) Caltech101.					
	Base	Novel	H		Base	Novel	H		Base	Novel	H
CLIP	91.05	97.16	94.01	CLIP	63.74	74.89	68.87	CLIP	70.24	77.98	73.91
CoOp	94.99	93.06	94.02	CoOp	78.37	67.49	72.52	CoOp	97.23	62.40	76.02
CoOp+	94.92	96.36	95.63	CoOp+	77.22	72.47	74.77	CoOp+	96.46	73.05	83.14
CoCoOp	95.00	97.81	96.38	CoCoOp	70.73	72.70	71.70	CoCoOp	94.79	68.44	79.49
CoCoOp+	94.99	97.50	96.23	CoCoOp+	69.28	74.50	71.80	CoCoOp+	92.81	75.01	82.97
ProGrad	95.44	95.15	95.29	ProGrad	76.91	70.55	73.59	ProGrad	96.43	69.10	80.51
ProGrad+	94.60	97.42	95.99	ProGrad+	74.16	73.99	74.07	ProGrad+	93.10	76.02	83.70
KgCoOp	95.14	97.62	96.36	KgCoOp	74.21	74.77	74.49	KgCoOp	95.87	72.75	82.72
KgCoOp+	94.39	97.55	95.94	KgCoOp+	70.65	75.54	73.01	KgCoOp+	92.84	75.78	83.45
Pro+Kg	95.50	97.64	96.56	Pro+Kg	76.61	73.54	75.04	Pro+Kg	96.50	74.26	83.93
(d) OxfordPets.			(e) StanfordCars.			(f) Flowers102.					
	Base	Novel	H		Base	Novel	H		Base	Novel	H
CLIP	89.77	91.29	90.52	CLIP	27.44	35.77	31.06	CLIP	69.39	75.53	72.33
CoOp	89.54	90.14	89.84	CoOp	39.44	27.20	32.20	CoOp	81.17	70.39	75.40
CoOp+	90.52	91.61	91.06	CoOp+	39.13	30.81	34.48	CoOp+	81.10	75.30	78.09
CoCoOp	90.55	91.38	90.96	CoCoOp	36.08	33.05	34.50	CoCoOp	79.52	76.62	78.04
CoCoOp+	90.61	91.88	91.24	CoCoOp+	34.73	35.01	34.87	CoCoOp+	78.21	77.97	78.09
ProGrad	90.46	90.54	90.50	ProGrad	39.43	27.90	32.68	ProGrad	81.36	73.95	77.48
ProGrad+	90.55	91.46	91.00	ProGrad+	38.64	31.49	34.70	ProGrad+	79.80	77.29	78.52
KgCoOp	90.40	91.72	91.06	KgCoOp	37.94	31.31	34.31	KgCoOp	80.76	76.05	78.33
KgCoOp+	90.23	91.62	90.92	KgCoOp+	36.49	34.94	35.70	KgCoOp+	78.60	77.29	77.94
Pro+Kg	90.70	91.74	91.22	Pro+Kg	39.14	30.36	34.20	Pro+Kg	81.84	76.98	79.34
(g) Food101.			(h) FGVCAircraft.			(i) SUN397.					
	Base	Novel	H		Base	Novel	H		Base	Novel	H
CLIP	53.94	57.97	55.88	CLIP	57.24	64.08	60.47	CLIP	71.04	78.53	74.60
CoOp	78.74	45.81	57.92	CoOp	89.21	63.08	73.90	CoOp	84.35	65.21	73.56
CoOp+	79.09	47.99	59.73	CoOp+	89.02	66.25	75.97	CoOp+	84.47	76.80	80.45
CoCoOp	74.65	53.66	62.44	CoCoOp	86.02	66.43	74.97	CoCoOp	82.28	70.87	76.15
CoCoOp+	73.77	55.60	63.41	CoCoOp+	85.22	70.38	77.09	CoCoOp+	81.28	75.39	78.22
ProGrad	77.31	48.95	59.94	ProGrad	90.76	65.25	75.92	ProGrad	83.77	68.36	75.28
ProGrad+	75.15	55.88	64.10	ProGrad+	88.36	69.38	77.73	ProGrad+	82.02	76.53	79.18
KgCoOp	80.09	52.98	63.77	KgCoOp	88.11	64.37	74.39	KgCoOp	83.09	75.39	79.05
KgCoOp+	78.20	56.88	65.86	KgCoOp+	84.70	64.86	73.46	KgCoOp+	81.46	78.20	79.80
Pro+Kg	80.05	53.74	64.31	Pro+Kg	89.92	67.82	77.32	Pro+Kg	83.56	75.68	79.43
(j) DTD.			(k) EuroSAT.			(l) UCF101.					

Table 1: Comparison of different methods in the base-to-novel generalization setting using 16 shots. H: Harmonic mean.

We report the final performance averaged over three random seeds. The same evaluation metrics adopted in CLIP are used. All experiments are conducted on RTX A4000 GPU. Without further specification, ViT-B/16 is adopted as the default vision backbone.

Base to Novel Generalization

In this setting, we split the train and test samples in 11 datasets into two groups: base classes (Base) and novel

classes (Novel). The two sets do not share any identical classes. We train VLPT methods on base classes and evaluate them on novel classes. When a VLPT method is enhanced by OOD detection, the corresponding method is annotated by a plus sign (+) in the name. The detailed results of 9 methods are reported in Table 1.

Table 1(a) shows the average results over 11 datasets, where zero-shot CLIP acquires the worst accuracy on base classes. This is because it uses fixed prompts which are sub-

	Source		Target			
	IN	INV2	IN-Sketch	IN-A	IN-R	Avg.
CLIP	66.71	60.89	46.09	47.77	73.98	57.18
CoOp	71.71	64.07	46.96	48.84	74.54	58.60
CoOp+	71.43	64.09	48.09	49.65	75.59	59.36
CoCoOp	71.16	64.35	48.87	50.73	75.90	59.96
CoCoOp+	70.44	63.79	48.72	50.68	76.03	59.81
ProGrad	72.24	64.86	48.18	49.91	75.50	59.61
ProGrad+	71.12	64.05	48.21	49.89	75.78	59.48
KgCoOp	70.67	63.71	48.72	50.39	76.68	59.88
KgCoOp+	69.60	62.91	48.04	49.62	76.04	59.15
Pro+Kg	72.02	64.92	49.06	50.84	76.62	60.36

Table 2: Comparison of methods with and without OOD detection in domain generalization setting. IN=ImageNet.

optimal compared to learned prompts on base set. However, CLIP’s accuracy on novel classes is quite appealing (better than CoOp, CoCoOp, ProGrad and KgCoOp), demonstrating its strong generalization ability across different classes.

CoOp obtains the highest *Base* accuracy, implying its strong fitting capability on base classes. In contrast, its *Novel* accuracy is the worst, showing its low base-to-novel generalization ability. By introducing OOD detection into CoOp, the *Novel* accuracy is improved significantly with sacrificing the *Base* accuracy a little. As such, the harmonic mean increases a lot, as large as 2.6%. As for CoCoOp, CoCoOp+ improves the harmonic mean by 0.7% with a slight drop of *Base* accuracy (i.e. 0.86%) and a remarkable increase of *Novel* accuracy (i.e. 2.01%). Besides, ProGrad+ improves the harmonic mean of ProGrad by 1.5% and the *Novel* accuracy by 3.85% with a slight sacrifice of *Base* accuracy, i.e. 1.53%. It obtains a new state of the art of harmonic mean, i.e. 77.20%, outperforming KgCoOp’s best result, 76.97%. Finally, KgCoOp+ and KgCoOp acquire the similar harmonic mean. One of the reasons the proposed method has little effect on KgCoOp is that it adds an extra loss function to enforce the few-shot and zero-shot classifier approach each other, which is less friendly for model fusion. Another reason is KgCoOp already achieves excellent harmonic mean, i.e. 76.97%, which is very near to the upper bound harmonic mean of fusing KgCoOp and zero-shot CLIP, i.e. 77.67%. This small gap makes further improvement harder. By summarizing the above observations, OOD detection is useful in improving base-to-novel generalization ability of existing VLPT methods.

Table 1(b)-(l) show results on each dataset. First, CoOp+, CoCoOp+, ProGrad+ and KgCoOp+ reduce the *Base* performance compared to CoOp, CoCoOp, ProGrad and KgCoOp on 6/11 datasets, 9/11 datasets, 10/11 datasets and 11/11 datasets, respectively. However, as shown in Table 1(a), the average *Base* performance drops a little especially for CoOp and CoCoOp. Second, they improve the *Novel* performance on 11/11 datasets, 9/11 datasets, 11/11 datasets and 7/11 datasets, respectively, implying boosted performance on new classes. Finally, they improve the har-

Method	Weight	Base	Novel	H
Static	0.95	82.47	68.48	74.83
	0.75	81.07	70.76	75.56
	0.50	78.70	73.13	75.81
	0.25	75.40	74.85	75.12
	0.05	70.51	74.37	72.39
Dynamic	$\alpha=64$	82.20	72.15	76.85

Table 3: Dynamic weighting vs static weighting of CoOp+.

monic mean performance on 11/11 datasets, 8/11 datasets, 11/11 datasets and 4/11 datasets, respectively. Specially, CoOp+ obtains an obvious improvement of harmonic mean of 1.81%, 2.25%, 7.12%, 2.28%, 2.69%, 2.07% and 6.89% on ImageNet, StanfordCars, Flowers102, FGVC Aircraft, SUN397, EuroSAT and UCF101 upon CoOp, respectively. In summary, the results demonstrate OOD detection can improve base-to-novel generalization ability.

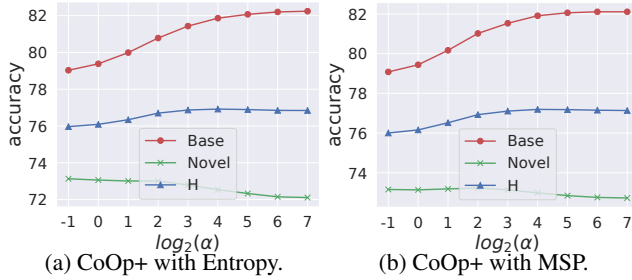
The proposed dynamic fusion is also suitable for two few-shot classifiers with simple modification. By replacing zero-shot CLIP with ProGrad in KgCoOp+, the OOD detection based dynamic fusion (Pro+Kg) can obtain better results compared to both KgCoOp and ProGrad under the base-to-novel generalization setting, as can be seen from Table 1. Actually, Pro+Kg achieves 0.4% higher harmonic accuracy compared to the best result reported above, i.e. 77.20%. Such an effectiveness is due to the strong complementarity between KgCoOp and ProGrad.

Domain Generalization

Domain generalization evaluates the effect of distribution shift under the same classes. The prompts are learned on source domain (i.e., ImageNet) and transferred to target domains (i.e., ImageNetV2, etc.) for evaluation. The results of zero-shot CLIP and four prompt tuning methods with and without OOD detection are reported in Table 2. From the results, OOD detection may reduce the performances on target domains slightly in some cases. Fortunately, the drops are marginal compared to the improvements in Table 1. In Table 2, almost all results of CoOp/CoCoOp/ProGrad/KgCoOp surpass CLIP. This is reasonable as the source and target domains share the classes, learning on source domain has positive effect on improving the recognition accuracy on target domains. As such, when fusing the zero-shot and few-shot classifier, more weights of the latter will gather around one. Meantime, the weaker zero-shot classifier may bring side effect for fusion. These can explain why OOD detection does not demonstrate significant improvements in this setting. Actually, OOD detection also works well in this setting if there is strong complementarity between the two classifiers being fused. From Table 2, KgCoOp and ProGrad have strong complementarity on different datasets. As such, we replace zero-shot CLIP with ProGrad in KgCoOp+. This method is denoted as Pro+Kg. From Table 2, Pro+Kg can achieve a new state of art, i.e. 60.36% average accuracy on target domains. Meantime, the accuracy on source domain is also appealing. The above results suggest that OOD detection is also helpful for domain generalization.

No.	OOD Detector	Base	Novel	H
0	MSP of zero-shot classifier only	78.74	73.61	76.09
1	MSP of few-shot classifier only	81.06	72.30	76.43
2	MSP (best $\alpha = 64$)	82.11	72.76	77.15
3	MSP ($\alpha = \infty$)	82.13	72.70	77.13
4	MaxLogit (best $\alpha = 8$)	78.45	73.52	75.86
5	MaxLogit ($\alpha = \infty$)	76.64	72.53	74.53
6	Energy (best $\alpha = 64$)	78.70	73.19	75.85
7	Energy ($\alpha = \infty$)	75.66	72.37	73.98
8	Entropy (best $\alpha = 64$)	82.20	72.15	76.85
9	Entropy ($\alpha = \infty$)	82.22	72.03	76.79

Table 4: Different OOD detection methods in CoOp+.

Figure 3: Effect of α on accuracies in CoOp+.

Ablation Study

Static Weighting: We can also use static (fixed) weight $s(\mathbf{x})$ for all test samples. The results of using static weights ($s(\mathbf{x}) = 0.05, 0.25, 0.50, 0.75, 0.95$) and dynamic weights ($\alpha = 64$) are compared in Table 3. Using fixed weights also implies the OOD detectors with 50% accuracy are adopted. The improved result of dynamic weighting demonstrates that weak OOD detectors help to boost generalizability across different classes.

Different OOD Detection Methods: The proposed method is flexible, it supports many different OOD detection methods. Besides the aforementioned ID scores, the maximum softmax probability (MSP) of the zero-shot and few-shot classifier can also serve as the weight, i.e., $s(\mathbf{x}) = 1 - ids^{zs}(\mathbf{x})$ and $s(\mathbf{x}) = ids^{fs}(\mathbf{x})$. The results of different methods are compared in Table 4. When using sigmoid function, the results with the best α in $\{0.5, 1, 2, 4, 8, 16, 32, 64, 128\}$ and infinite α are reported. From the table, MSP (Row 2) achieves the best result, slightly better than entropy based method (Row 8). For different ID scores, using sigmoid function with smaller α is better than Heaviside step function ($\alpha = \infty$). The underlying reason is that continuous weights are helpful for model fusion compared to binary weights. We also find that MaxLogit and Energy do not work well, which is contrary to the observations in prior work of OOD detection. We conjecture that, for these methods, the ID scores from different classifiers are not well normalized for comparison.

Effect of Scaling Factor α : The α is the scaling factor of the score difference. Its effect for entropy and MSP based methods are shown in Fig. 3. Larger α damages the perfor-

Method	Backbone	Base-to-Novel			Domain	
		Base	Novel	H	Source	Target
CoOp	ResNet-50	77.17	59.03	66.89	62.99	41.05
CoOp+	ResNet-50	77.25	64.02	70.20	63.08	42.13
			+4.99	+3.31		+1.08
CoCoOp	ResNet-50	74.68	63.48	68.63	63.16	42.95
CoCoOp+	ResNet-50	74.28	66.49	70.17	62.34	42.90
			+3.01	+1.54		-0.05
ProGrad	ResNet-50	77.14	63.37	69.58	63.66	42.11
ProGrad+	ResNet-50	76.09	67.39	71.48	62.92	42.67
			+4.02	+1.90		+0.56
KgCoOp	ResNet-50	76.42	64.93	70.21	62.52	43.18
KgCoOp+	ResNet-50	74.52	68.01	71.12	61.33	42.45
			+3.08	+0.91		-0.73
CoOp	ViT-B/32	79.24	62.77	70.05	66.72	48.32
CoOp+	ViT-B/32	78.99	67.85	73.00	66.68	49.21
			+5.08	+2.95		+0.89
CoCoOp	ViT-B/32	76.26	66.63	71.12	66.07	49.89
CoCoOp+	ViT-B/32	75.55	70.58	72.98	65.39	49.84
			+3.95	+1.86		-0.05
ProGrad	ViT-B/32	78.75	66.19	71.93	66.98	49.33
ProGrad+	ViT-B/32	77.24	70.48	73.71	66.21	49.59
			+4.29	+1.78		+0.26
KgCoOp	ViT-B/32	78.29	69.87	73.84	65.57	49.79
KgCoOp+	ViT-B/32	76.37	71.61	73.91	64.68	49.27
			+1.74	+0.07		-0.52

Table 5: Base-to-novel and domain generalization results with different vision backbones for 16 shots. Mean accuracies over four target domains are reported.

mance on novel classes, but improves the performance on base classes. Thus, it in fact plays a role in balancing the two performances. This makes the balance tunable compared to using Heaviside step function.

Different Backbones: The proposed method also works for other vision backbones. To validate this point, the base-to-novel and domain generalization results with ResNet-50 and ViT-B/32 are presented in Table 5. With OOD detection, the harmonic mean accuracies in base-to-novel setting are improved uniformly (0.07%~3.31%). Meantime, methods with OOD detection can achieve similar or better mean accuracies on target domains in most cases (6 out of 8).

Conclusion

Existing methods for vision-language prompt tuning have poor generalization ability. This work proposes a new viewpoint for balancing the base and novel performance across classes by considering this problem as an OOD detection problem. Based on this idea, we propose a competition based scoring function to generate dynamic weight of zero-shot and few-shot classifier. By extensive experiments, even with weak OOD detectors, fused classifier demonstrates better generalization ability compared to the original one. One limitation of our method is that the effectiveness may be less significant when the two classifiers being fused have little complementarity. Future work will focus on exploiting stronger OOD detectors and addressing the above limitation.

Acknowledgments

This research was supported by the National Natural Science Foundation of China under Grants 62306310 and 62076242, and the National Key Research and Development Program of China under Grant No. 2018AAA0100400.

References

- Bossard, L.; Guillaumin, M.; and Gool, L. V. 2014. Food-101 - Mining Discriminative Components with Random Forests. In *ECCV*, 446–461.
- Chen, G.; Yao, W.; Song, X.; Li, X.; Rao, Y.; and Zhang, K. 2023. PLOT: Prompt Learning with Optimal Transport for Vision-Language Models. In *ICLR*.
- Chen, Y.; Dai, X.; Liu, M.; Chen, D.; Yuan, L.; and Liu, Z. 2020. Dynamic Convolution: Attention Over Convolution Kernels. In *CVPR*, 11027–11036.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing Textures in the Wild. In *CVPR*, 3606–3613.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*, 248–255.
- Ding, K.; Wang, Y.; Liu, P.; Yu, Q.; Zhang, H.; Xiang, S.; and Pan, C. 2022. Prompt Tuning with Soft Context Sharing for Vision-Language Models. *CoRR*, abs/2208.13474.
- Gao, H.; Zhang, H.; Yang, X.; Li, W.; Gao, F.; and Wen, Q. 2022. Generating natural adversarial examples with universal perturbations for text classification. *Neurocomputing*, 471: 175–182.
- Gao, K.; Chen, L.; Zhang, H.; Xiao, J.; and Sun, Q. 2023. Compositional Prompt Tuning with Motion Cues for Open-vocabulary Video Relation Detection. In *ICLR*.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 12(7): 2217–2226.
- Hendrycks, D.; Basart, S.; Mazeika, M.; Zou, A.; Kwon, J.; Mostajabi, M.; Steinhardt, J.; and Song, D. 2022. Scaling Out-of-Distribution Detection for Real-World Settings. In *ICML*, 8759–8773.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; Song, D.; Steinhardt, J.; and Gilmer, J. 2021. The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization. In *ICCV*, 8320–8329.
- Hendrycks, D.; and Gimpel, K. 2017. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In *ICLR*.
- Hou, Z.; and Kung, S. 2021. Parameter Efficient Dynamic Convolution via Tensor Decomposition. In *BMVC*, 107.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.; Parekh, Z.; Pham, H.; Le, Q. V.; Sung, Y.; Li, Z.; and Duerig, T. 2021. Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In *ICML*, volume 139, 4904–4916.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3D Object Representations for Fine-Grained Categorization. In *ICCVW*, 554–561.
- Li, C.; Zhou, A.; and Yao, A. 2022. Omni-Dimensional Dynamic Convolution. In *ICLR*.
- Li, F.; Fergus, R.; and Perona, P. 2007. Learning generative visual models from few training examples: An incremental Bayesian approach tested on 101 object categories. *Comput. Vis. Image Underst.*, 106(1): 59–70.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation. In *ICML*.
- Liu, W.; Wang, X.; Owens, J. D.; and Li, Y. 2020. Energy-based Out-of-distribution Detection. In *NeurIPS*.
- Liu, Z.; Wang, Y.; Han, K.; Ma, S.; and Gao, W. 2022. Instance-Aware Dynamic Neural Network Quantization. In *CVPR*, 12424–12433.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M. B.; and Vedaldi, A. 2013. Fine-Grained Visual Classification of Aircraft. *CoRR*, abs/1306.5151.
- Mokady, R.; Hertz, A.; and Bermano, A. H. 2021. ClipCap: CLIP Prefix for Image Captioning. *CoRR*, abs/2111.09734.
- Nilsback, M.; and Zisserman, A. 2008. Automated Flower Classification over a Large Number of Classes. In *ICVGIP*, 722–729.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. V. 2012. Cats and dogs. In *CVPR*, 3498–3505.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, volume 139, 8748–8763.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *CoRR*, abs/2204.06125.
- Rao, Y.; Zhao, W.; Chen, G.; Tang, Y.; Zhu, Z.; Huang, G.; Zhou, J.; and Lu, J. 2022. DenseCLIP: Language-Guided Dense Prediction with Context-Aware Prompting. In *CVPR*, 18082–18091.
- Recht, B.; Roelofs, R.; Schmidt, L.; and Shankar, V. 2019. Do ImageNet Classifiers Generalize to ImageNet? In *ICML*, volume 97, 5389–5400.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q. V.; Hinton, G. E.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *ICLR*.
- Shu, M.; Nie, W.; Huang, D.; Yu, Z.; Goldstein, T.; Anandkumar, A.; and Xiao, C. 2022. Test-Time Prompt Tuning for Zero-Shot Generalization in Vision-Language Models. In *NeurIPS*.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *CoRR*, abs/1212.0402.
- Sun, X.; Hu, P.; and Saenko, K. 2022. DualCoOp: Fast Adaptation to Multi-Label Recognition with Limited Annotations. *CoRR*, abs/2206.09541.

- Wang, H.; Ge, S.; Lipton, Z. C.; and Xing, E. P. 2019. Learning Robust Global Representations by Penalizing Local Predictive Power. In *NeurIPS*, 10506–10518.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. SUN database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 3485–3492.
- Yang, J.; Zhou, K.; Li, Y.; and Liu, Z. 2021. Generalized Out-of-Distribution Detection: A Survey. *CoRR*, abs/2110.11334.
- Yao, H.; Zhang, R.; and Xu, C. 2023. Visual-Language Prompt Tuning with Knowledge-guided Context Optimization. In *CVPR*.
- Zhang, Z.; and Xiang, X. 2023. Decoupling MaxLogit for Out-of-Distribution Detection. In *CVPR*, 3388–3397.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Conditional Prompt Learning for Vision-Language Models. In *CVPR*, 16816–16825.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Learning to Prompt for Vision-Language Models. *IJCV*, 130(9): 2337–2348.
- Zhu, B.; Niu, Y.; Han, Y.; Wu, Y.; and Zhang, H. 2023. Prompt-aligned Gradient for Prompt Tuning. *ICCV*.