

# TCP: Textual-based Class-aware Prompt tuning for Visual-Language Model

Hantao Yao<sup>1</sup>, Rui Zhang<sup>2</sup>, Changsheng Xu<sup>1,3</sup>

<sup>1</sup> State Key Laboratory of Multimodal Artificial Intelligence Systems, Institute of Automation, Chinese Academy of Sciences, China

<sup>2</sup> State Key Lab of Processors, Institute of Computing Technology, Chinese Academy of Sciences, China;

<sup>3</sup> University of Chinese Academy of Sciences, China

hantao.yao@nlpr.ia.ac.cn; zhangrui@ict.ac.cn; csxu@nlpr.ia.ac.cn

## Abstract

Prompt tuning represents a valuable technique for adapting pre-trained visual-language models (VLM) to various downstream tasks. Recent advancements in CoOp-based methods propose a set of learnable domain-shared or image-conditional textual tokens to facilitate the generation of task-specific textual classifiers. However, those textual tokens have a limited generalization ability regarding unseen domains, as they cannot dynamically adjust to the distribution of testing classes. To tackle this issue, we present a novel Textual-based Class-aware Prompt tuning (TCP) that explicitly incorporates prior knowledge about classes to enhance their discriminability. The critical concept of TCP involves leveraging Textual Knowledge Embedding (TKE) to map the high generalizability of class-level textual knowledge into class-aware textual tokens. By seamlessly integrating these class-aware prompts into the Text Encoder, a dynamic class-aware classifier is generated to enhance discriminability for unseen domains. During inference, TKE dynamically generates class-aware prompts related to the unseen classes. Comprehensive evaluations demonstrate that TKE serves as a plug-and-play module effortlessly combinable with existing methods. Furthermore, TCP consistently achieves superior performance while demanding less training time<sup>1</sup>.

## 1. Introduction

Large-scale image-text pairs have the ability to train a Visual-language model (VLM) with a powerful generalization capacity for various downstream tasks [2, 33]. However, training these models from scratch requires a huge dataset with labeled images, making it difficult to apply them directly to downstream tasks with fewer images. To address this issue, three common techniques have been pro-

<sup>1</sup>[https://github.com/htyao89/Textual-based\\_Class-aware\\_prompt\\_tuning](https://github.com/htyao89/Textual-based_Class-aware_prompt_tuning)

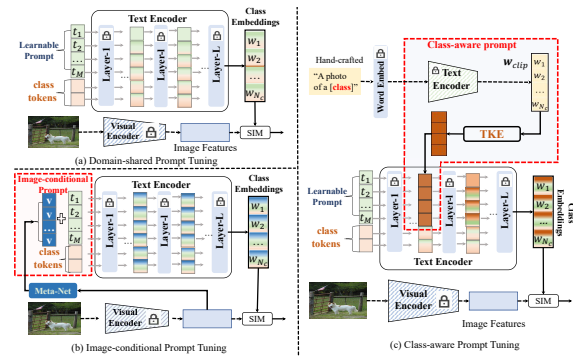


Figure 1. Comparison with existing frameworks. (a) Domain-shared prompt tuning applies the same learnable prompt between the training and testing domains. (b) Image-conditional prompt tuning combines the image embedding with the learnable prompt; (c) Class-aware Prompt tuning injects the class-level textual embedding into the Text Encoder with the class-aware prompt.

posed: fine-tuning [30], prompt tuning [50], adapter [12] and LoRA [15]. Among them, prompt tuning is a simple and efficient framework that transfers the essential general knowledge of VLM to the downstream tasks.

Prompt tuning<sup>2</sup> is a technique that combines learnable textual tokens with class tokens to generate a discriminative textual classifier, known as Context Optimization (CoOp) [50]. Recently, various CoOp-based methods [4, 5, 17, 18, 22, 31, 40, 43, 50, 52] infer the domain-shared prompt tokens between training and testing domains (Figure 1(a)). Nevertheless, as the domain-shared prompt tokens are derived from labeled training images, their performance is suboptimal when confronted with unseen test classes. To enhance the generalization capacity of learnable prompt tokens, image-conditional prompts have been proposed in [45, 49] by fusing image features and learnable tokens (Figure 1(b)). Notably, image-conditional

<sup>2</sup>For the visual-language model, two types of prompt-tuning strategies exist: visual prompt tuning and textual prompt tuning. In this work, we focus on the textual prompt tuning and do not consider the visual prompt tuning.

textual tokens encapsulate specific knowledge of each image, particularly for testing images, thereby rendering it more apt for generalization to unseen testing images. However, image-conditional prompt tokens with image-specific knowledge have less ability in elevating the distribution of class embeddings. In summary, classifiers generated by domain-shared and image-conditional textual tokens exhibit suboptimal performance for the unseen classes, primarily due to their inability to explicitly model the class distribution. Consequently, it is imperative to establish a dynamic relationship between the learnable prompt and the textual knowledge of each class to augment its discriminative prowess.

The frozen CLIP, coupled with a hand-crafted prompt, exhibits robust generalization capabilities to novel classes, rendering it a valuable source of prior textual knowledge for each class. By associating the class-level textual knowledge with the learnable prompts, a class-aware prompt can be formulated for enhancing the discriminative capacity of the textual classifier. To achieve this, we employ an embedding module to project the class-aware textual knowledge into the class-aware prompt tokens, as shown in Figure 1(c). The resultant class-aware prompt incorporates prior textual knowledge specific to each class, endowing the generated textual classifier with heightened discriminative prowess. Furthermore, the class-aware prompt facilitates the generation of classifiers for both seen and unseen classes by leveraging textual knowledge from both categories. To sum up, the trained embedding module can generate a class-aware prompt for each class based on its description ('class name'), thereby enhancing generalization and discriminative capabilities in class-level textual embeddings.

Therefore, we propose a novel Textual-based Class-aware Prompt tuning (TCP) based on the framework of CoOp, shown in Figure 2. In addition to the domain-share textual tokens introduced in CoOp, TCP contributes a novel Textual Knowledge Embedding (TKE) to map the class-level textual knowledge into class-aware prompt tokens. Moreover, a class-aware textual classifier is generated by inserting the class-aware prompt tokens into the middle layer of the Text Encoder. We use standard contrastive loss and knowledge-guided consistency [43] to optimize the TKE and the learnable prompt tokens. During inference, TCP generates a class-aware classifier for unseen classes by feeding the domain-shared prompt tokens and the class-aware prompt tokens generated by the TKE into the frozen Text Encoder.

Overall, the proposed TCP explicitly steers prompts to learn a class-aware knowledge that maximizes the generalization and discriminative of the downstream tasks. Evaluation over 11 image classification datasets on base-to-new generalization, cross-dataset generalization, and few-shot learning verify shows that TCP is an efficient method that

obtain a higher performance with less training time. In summary, the proposed Textual-based Class-aware Prompt tuning (TCP) has the following main contributions:

1. An effectively Textual-based Class-aware Prompt tuning is proposed by injecting the textual class-aware prompts generated by Textual Knowledge Embedding (TKE) into the Text Encoder.
2. We demonstrate that explicitly incorporating the prior knowledge of each class into the learnable prompt tokens can enhance the discriminative of the class distribution.
3. Textual Knowledge Embedding (TKE) is a plug-and-play module that can quickly insert existing prompt tuning methods to improve their performance further.

## 2. Related Works

### 2.1. Vision-Language Models

Recently, researchers have shown that the Visual-Language Models (VLM) [2, 33] consisting of visual and textual modalities trained on the large-scale of image-text pairs has a powerful generalization and discriminativeion. To further boost the descriptive ability of VLM, the VLM models are boosted from the following aspect: 1) using a stronger text encoder or visual encoder [25, 41, 46]; 2) deeply fusing the visual and text knowledges [23, 38]; 3) using more images [16, 33, 35, 36]. To boost the diversity of text description, Masked Language Modeling (MLM) [20] [26] randomly erases the words in the text description used for representation learning. Unlike MLM, Masked autoencoder-based methods [13] are proposed to boost the descriptive ability by randomly masking the image patches. Among existing VLM models, CLIP is a representative and straightforward framework for inferring the independent visual and text encoder using the contrastive loss based on 400 million image-text association pairs. As CLIP has a good generalization, most existing CoOp-based methods are proposed based on CLIP for adapting the pre-trained VLM into the downstream task. Similar to existing methods, we conduct the prompt-tuning strategy on the TextEncoder of CLIP to obtain a task-specific textual embedding for prediction.

### 2.2. Prompt Tuning

To adapt the pretrained VLM to the downstream tasks, the prompt tuning [10, 22, 24, 31, 33, 45] always applies task-related textual tokens to infer the task-specific textual knowledge. In CLIP[33], the hand-crafted template "a photo of a [CLASS]" is used to embed the textual embedding for zero-shot prediction. However, the hand-crafted prompts have a poor ability to describe the downstream task. Textual prompt tuning is applied to boost the textual embedding by inferring a set of learnable textual tokens combined with the class tokens. For example, Context Optimization (CoOp) [50] replaces the hand-crafted prompts

with the learnable soft prompts. To improve the generalization of the learnable textual prompt in CoOp, Conditional Context Optimization(CoCoOp) [49] and VPT [45] generates an image-conditional prompt fusing the image feature and the learnable textual prompt. Moreover, Knowledge-Guided Context Optimization(KgCoOp) [43], ProGrad [51], and Prompt Regularization(ProReg) [52] constrain the proposed learnable prompts contain the essential general knowledge. Unlike the above methods, which consider textual prompts, Ensembling Context Optimization(ECO) [1] employs prompt ensembling to combine multiple prompts. To obtain high-quality task-related tokens, ProDA [27] considers the prompt’s prior distribution learning, and Distribution-Aware Prompt Tuning (DAPT) [5] optimizes the learnable prompt by maximizing inter-dispersion. Besides the textual knowledge from “class-name”, Knowledge-Aware Prompt Tuning(KAPT) [17] employs the external knowledge to generate the discriminative knowledge-aware prompt for the novel categories. PLOT [4] applies optimal transport to match the vision and text modalities for generating the discriminative and visual-aligned local textual prompt tokens. Besides the textual prompt tuning, Multi-modal Prompt Learning (MaPLe) [18] and PromptSRC [19] conduct the visual-textual prompt tuning by jointly conducting the prompt tuning on the visual and text encoders. Multitask Vision-Language Prompt Tuning(MVLPT) [37]incorporates cross-task knowledge into prompt tuning for vision-language models. DenseCLIP [34] uses the context-aware prompt strategy to generate dense prediction tasks, and CLIP-Adapter [11] applies an adapter to adjust the visual or text embeddings.

Existing methods commonly infer two types of prompt tokens: domain-share and image-conditional. However, the textual classifiers generated with these tokens tend to perform poorly on unseen classes. To mitigate this limitation, we propose a novel Textual-based Class-aware Prompt tuning (TCP) that employs a dynamic class-aware token to enhance the generalization and discriminative capabilities of the learnable textual prompt. Moreover, we introduce a Textual Knowledge Embedding to project class-level textual knowledge into the class-aware prompts, which are then inserted into the Text Encoder to generate a discriminative class-aware classifier. The evaluation results demonstrate that integrating class-level prior knowledge into the prompt tokens significantly enhances the discriminative ability of the prompt tuning process.

### 3. Methodology

As the Textual-based Class-aware Prompt tuning(TCP) is proposed based on Context Optimization (CoOp), we first briefly review CoOp and then introduce the proposed TCP.

### 3.1. Preliminaries

Existing CoOp-based methods are proposed based on the powerful Contrastive Language-Image Pre-training(CLIP). Given an image along with its corresponding textual description, CLIP uses visual and text encoders to extract the visual and text embeddings. After that, the constrastive loss between the visual and textual embeddings is calculated to align those two embeddings. To effectively adapt CLIP for the downstream task, CLIP applies the hand-crafted template “a photo of a { }” to extract the general class-level textual embedding, defined as  $\mathbf{W}^{clip} = \{\mathbf{w}_i^{clip}\}_{i=1}^{N_c}$ , where  $\mathbf{w}_i^{clip}$  is the textual embedding of  $i$ -th class, and  $N_c$  is the number of classes. Given the ‘class-name’ of  $i$ -th class, Word Embedded  $e(\cdot)$  firstly embeds the hand-crafted description into a vectorized textual tokens:  $\mathbf{t}_i^{clip} = e(\text{“a photo of a {class-name}”})$ . After that, Text Encoder  $\theta$  maps the vectorized textual tokens  $\mathbf{t}_i^{clip}$  into the class-level embedding:  $\mathbf{w}_i^{clip} = \theta(\mathbf{t}_i^{clip})$ .

To improve the discriminative of the class-level embedding, the prompt tuning methods of Context Optimization(CoOp) replaces the hand-crafted textual tokens with a set of learnable textual tokens  $\mathbb{T} = \{\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_M\}$ , where  $M$  is the length of tokens. Similar to CLIP, the corresponding class token  $\mathbf{c}_i$  is concatenated with the learnable tokens  $\mathbb{T}$  for generating the textual token  $\mathbf{t}_i^{coop} = \{\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_M, \mathbf{c}_i\}$ . Next, the textual embedding  $\mathbf{w}_i^{coop}$  is obtained by feeding the textual tokens  $\mathbf{t}_i^{coop}$  into Text Encoder  $\theta$ , *i.e.*,  $\mathbf{w}_i^{coop} = \theta(\mathbf{t}_i^{coop})$ . Finally, the textual embeddings of all classes are defined as  $\mathbf{W}^{coop} = \{\mathbf{w}_i^{coop}\}_{i=1}^{N_c}$ . CoOp infers the learnable textual tokens  $\mathbb{T}$  by minimizing the contrastive loss between the image’s embedding  $\mathbf{x}$  and its class embedding  $\mathbf{w}_y^{coop}$ :

$$\mathcal{L}_{ce} = \frac{1}{N} \sum_{(\mathbf{x}, y) \in \mathcal{D}_s} \frac{\exp(d(\mathbf{x}, \mathbf{w}_y^{coop})/\tau)}{\sum_{i=1}^{N_c} \exp(d(\mathbf{x}, \mathbf{w}_i^{coop})/\tau)}, \quad (1)$$

where  $\mathcal{D}_s$  is the seen dataset, and  $d(\cdot)$  is the cosine distance.  $\tau$  is a temperature factor defined in CLIP, and  $N$  is the number of training images.

As the generated textual embedding has a good generalization ability for the novel classes, KgCoOp further adds an efficient consistency  $\mathcal{L}_{kg}$  between the generated embedding  $\mathbf{W}^{coop}$  and the general embedding  $\mathbf{W}^{clip}$ ,

$$\mathcal{L}_{kg} = \|\mathbf{W}^{clip} - \mathbf{W}^{coop}\|_2^2. \quad (2)$$

Therefore, a robust objective for prompt tuning is:

$$\mathcal{L} = \mathcal{L}_{ce} + \omega \mathcal{L}_{kg}, \quad (3)$$

where  $\omega$  is set as 8.0 same as in KgCoOp [43].

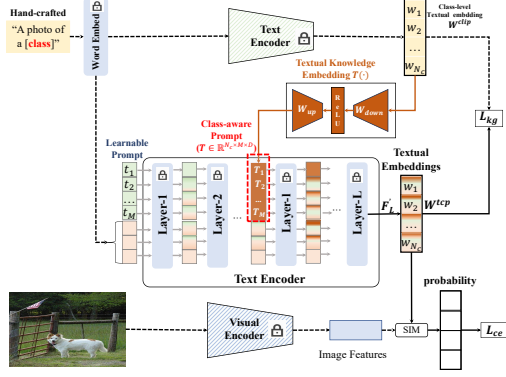


Figure 2. The framework of the proposed TCP.

### 3.2. Textual-based Class-aware Prompt tuning

Based on the pre-trained Text Encoder in CLIP, the textual prompt tuning aims to infer a set of the domain-shared or image-conditional textual tokens combined with the general class tokens for generating the specific class embeddings. However, the textual classifier generated by those textual tokens performs a worse generalization on the unseen classes because they can not model the distribution of the testing classes. Research has shown that utilizing the general textual knowledge extracted by frozen CLIP can create a discriminative prior knowledge of novel classes, enhancing the discriminative and generalization of the learnable prompt. Drawing on the benefit of the general textual knowledge of seen and unseen classes, we propose a Textual-based Class-aware Prompt tuning(TCP) to adapt the pretrained CLIP to the downstream task. As shown in Figure 2, TCP employs a Textual Knowledge Embedding(TKE) to transfer the general class-level textual embedding into the class-aware prompt, which is then combined with the learnable textual tokens for generating the class-aware classifiers. TKE is advantageous for unseen classes, as it generates class-specific prompts to obtain an unseen class-aware textual classifier with better discriminative abilities. Moreover, explicitly incorporating seen class-aware prompts can enhance the discriminative power of seen classes.

Given the general class-level textual embedding  $\mathbf{W}^{clip} \in \mathbb{R}^{N_c \times D_t}$  with  $N_c$  training classes, Textual Knowledge Embedding(TKE)  $\mathcal{T}(\cdot)$  is proposed to project the class-level embedding  $\mathbf{W}^{clip}$  into the class-aware prompt  $\mathbf{T} = \mathcal{T}(\mathbf{W}^{clip})$ . As shown in Figure 2, TKE consists of two layers: down-project layer and up-project layer. The down-project layer uses the weight  $\mathbf{W}_{down} \in \mathbb{R}^{D_t \times D_{mid}}$  to project the textual embedding into the low-dimension feature with the dimension of  $D_{mid}$ . Next, the weight  $\mathbf{W}_{up} \in \mathbb{R}^{D_{mid} \times D'}$  of the up-project layer maps the low-dimension feature into the high-dimension feature with the dimension of  $D'$ . Note that  $D'$  is determined by the prompt's length  $M$  and dimension  $D$ :  $D' = M \times D$ . In

summary, the general textual embedding  $\mathbf{W}^{clip} \in \mathbb{R}^{N_c \times D_t}$  can be projected into the class-aware textual tokens  $\mathbf{T} \in \mathbb{R}^{N_c \times D'}$ , which is further reshaped into the shape of  $\mathbf{T} \in \mathbb{R}^{N_c \times M \times D}$  for inserting into the middle layer of Text Encoder  $\theta$ .

Assuming we insert the class-aware prompt  $\mathbf{T}$  into  $l$ -th layer of Text Encoder  $\theta$ . We will give a detailed analysis of the hyperparameter  $l$  in the following. Similar to CoOp, by combining the domain-shared learnable textual tokens  $\mathbb{T} = [\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_M]$  and the pre-trained class tokens  $\mathbf{C}$  of all classes, we can obtain the input textual tokens  $\mathbf{F}_0 = \{\mathbb{T}, \mathbf{C}\}$  of the Text Encoder, where  $\mathbf{C} = \{\mathbf{c}_i\}_{i=1}^{N_c}$ , and  $\mathbf{c}_i$  is the vectorized textual tokens of  $i$ -th class. The textual token  $\mathbf{F}_0$  is fed into the first  $l$  layers of Text Encoder for obtaining the middle-level textual embedding  $\mathbf{F}_l$ . Formally, the textual token  $\mathbf{F}_i (i \leq l)$  of  $i$ -th layer is defined as:

$$\mathbf{F}_i = \theta_i(\mathbf{F}_{i-1}), i \in [1, l], \quad (4)$$

where  $\theta_i$  is the  $i$ -th layer of Text Encoder.

For the textual tokens  $\mathbf{F}_l \in \mathbb{R}^{N_c \times N_t \times D}$  and the class-aware prompt tokens  $\mathbf{T} \in \mathbb{R}^{N_c \times M \times D}$ , the first dimension is related to the number of classes. Therefore, the learnable prompt tokens are always inserted into the second dimension of  $\mathbf{F}_l$ , the same as CoOp. Formally, the class-aware prompt  $\mathbf{T}$  is inserted into  $\mathbf{F}_l$  to generate the class-aware enhanced tokens  $\mathbf{F}'_l$ ,

$$\mathbf{F}'_l = [\mathbf{T}_1, \mathbf{T}_2, \dots, \mathbf{T}_M, \mathbf{F}_{l,M+1}, \mathbf{F}_{l,M+2}, \dots, \mathbf{F}_{l,N_t}], \quad (5)$$

where  $\mathbf{T}_i$  denotes the  $i$ -th index of  $\mathbf{T}$  in the second dimension, and  $\mathbf{F}_{l,j}$  denote the correspond  $j$ -th index of  $\mathbf{F}_l$  in the second dimension, *i.e.*,  $\mathbf{T}_i = \mathbf{T}[:, i, :]$ , and  $\mathbf{F}_{l,j} = \mathbf{F}_l[:, j, :]$ .

After that, the class-enhanced textual tokens  $\mathbf{F}'_l$  are fed into the rest layers for generating the class-aware textual embedding,

$$\mathbf{F}'_i = \theta_i(\mathbf{F}'_{i-1}), i \in [l+1, L]. \quad (6)$$

The output of the last layer  $\mathbf{F}'_L$  is treated as the class embedding  $\mathbf{W}^{tcp}$  used for optimization with contrastive loss and knowledge-guided consistency loss in Eq. (3).

## 4. Experiments

Similar to CoOp [50], we evaluate the effectiveness of TCP from the three types of tasks: 1) generalization from base-to-new classes within a dataset; 2) few-shot learning with  $K$ -shot labeled images; 3) cross-dataset generalization from the imagenet to other datasets. More detailed results will be provided in the Supplementary materials.

### 4.1. Experimental Setup

**Dataset.** We conduct the base-to-new generaliation, few-shot learning, and cross-dataset generalization on 11 benchmarks, *i.e.*, ImageNet [7], Caltech [9], OxfordPets [32],



Table 1. Comparison on the base-to-new generalization setting with 16-shot samples. ‘tp’, ‘dtp’, ‘vp’, and ‘dvp’ denote the ‘textual prompt’, ‘deep textual prompt’, ‘visual prompt’, and ‘deep visual prompt’, respectively. PromptSRC are based on deep visual-textual prompt tuning(‘dvp+dtp’). ‘\*’ denote the performance obtained by our re-implementation.

Datasets	Sets	CoOp* (ICCV22)	CoCoOp (CVPR22)	DAPT* (ICCV23)	ProGrad* (ICCV23)	ProDA (CVPR22)	KgCoOp (CVPR23)	RPO (ICCV23)	PLOT* (ICLR23)	LFA (ICCV23)	MaPLe (CVPR23)	DePT (Arxiv23)	PromptSRC* (ICCV23)	TCP
		tp	tp	tp+vp	tp	tp	tp	dtp+dvp	tp+vp	–	dtp+dvp	tp	dtp+dvp	tp
Average	Base	82.38	80.47	83.18	82.48	81.56	80.73	81.13	83.98	83.62	82.28	83.62	84.12	<b>84.13</b>
	New	67.96	71.69	69.27	70.75	72.30	73.6	75.00	71.72	74.56	75.14	75.04	75.02	<b>75.36</b>
	H	74.48	75.83	75.59	76.16	76.65	77.0	77.78	77.37	78.83	78.55	79.10	79.31	<b>79.51</b>
ImageNet	Base	76.46	75.98	76.83	77.02	75.40	75.83	76.60	77.30	76.89	76.66	77.03	77.75	77.27
	New	66.31	70.43	69.27	66.66	70.23	69.96	<b>71.57</b>	69.87	69.36	70.54	70.13	70.70	69.87
	H	71.02	73.10	72.85	71.46	72.72	72.78	74.00	73.40	72.93	73.47	73.42	<b>74.06</b>	73.38
Caltech101	Base	97.80	97.96	97.83	98.02	98.27	97.72	97.97	<b>98.53</b>	98.41	97.74	98.30	98.13	98.23
	New	93.27	93.81	93.07	93.89	93.23	94.39	94.37	92.80	93.93	94.36	94.60	93.90	<b>94.67</b>
	H	95.48	95.84	95.39	95.91	95.68	96.03	96.03	95.58	96.13	96.02	96.41	95.97	<b>96.42</b>
OxfordPets	Base	94.47	95.20	95.00	95.07	95.43	94.65	94.63	94.50	95.13	95.43	94.33	<b>95.50</b>	94.67
	New	96.00	97.69	95.83	97.63	<b>97.83</b>	97.76	97.50	96.83	96.23	97.76	97.23	97.40	97.20
	H	95.23	96.43	95.41	96.33	<b>96.62</b>	96.18	96.05	95.65	95.68	96.58	95.76	96.44	95.92
Cars	Base	75.67	70.49	75.80	77.68	74.70	71.76	73.87	79.07	76.32	72.94	79.13	78.40	<b>80.80</b>
	New	67.53	73.59	63.93	68.63	71.20	75.04	<b>75.53</b>	74.80	74.88	74.00	75.47	74.73	74.13
	H	71.37	72.01	69.36	72.88	72.91	73.36	74.69	76.88	75.59	73.47	77.26	75.52	<b>77.32</b>
Flowers	Base	97.27	94.87	96.97	95.54	97.70	95.00	94.13	97.93	97.34	95.92	<b>98.00</b>	97.90	97.73
	New	67.13	71.75	60.90	71.87	68.68	74.73	76.67	73.53	75.44	72.46	76.37	<b>76.77</b>	75.57
	H	79.44	81.71	74.81	82.03	80.66	83.65	84.50	83.99	85.00	82.56	85.84	<b>86.06</b>	85.23
Food101	Base	89.37	<b>90.70</b>	90.37	90.37	90.30	90.5	90.33	89.80	90.52	90.71	90.50	90.63	90.57
	New	88.77	91.29	91.30	89.59	88.57	91.7	90.83	91.37	91.48	<b>92.05</b>	91.60	91.50	91.37
	H	89.07	90.99	90.83	89.98	89.43	91.09	90.58	90.58	91.00	<b>91.38</b>	91.05	91.06	90.97
Aircraft	Base	39.67	33.41	39.97	40.54	36.90	36.21	37.33	42.13	41.48	37.44	<b>43.20</b>	42.30	41.97
	New	31.23	23.71	29.80	27.57	34.13	33.55	34.20	33.73	32.29	35.61	<b>34.83</b>	<b>36.97</b>	34.43
	H	34.95	27.74	34.14	32.82	35.46	34.83	35.70	37.46	36.31	36.50	38.57	<b>39.46</b>	37.83
SUN397	Base	80.85	79.74	80.97	81.26	78.67	80.29	80.60	82.20	82.13	80.82	82.33	<b>82.83</b>	82.63
	New	68.34	76.86	76.97	74.17	76.93	77.80	76.53	73.63	77.20	78.70	77.80	<b>79.00</b>	78.20
	H	74.07	78.27	78.92	77.55	77.79	78.36	79.18	77.68	79.59	79.75	80.00	<b>80.87</b>	80.35
DTD	Base	79.97	77.01	82.23	77.35	80.67	77.55	76.70	81.97	81.29	80.36	82.20	82.60	<b>82.77</b>
	New	48.60	56.00	54.23	52.35	56.48	54.99	62.13	43.80	60.63	<b>59.18</b>	59.13	57.50	58.07
	H	60.46	64.85	65.36	62.45	66.44	64.35	68.61	57.09	<b>69.46</b>	68.16	68.78	67.80	68.25
EuroSAT	Base	90.10	87.49	<b>94.73</b>	90.11	83.90	85.64	86.63	93.70	93.40	94.07	89.03	92.40	91.63
	New	53.00	60.04	50.33	60.89	66.00	64.34	68.97	62.67	71.24	73.23	71.07	68.43	<b>74.73</b>
	H	66.74	71.21	65.74	72.67	73.88	73.48	76.79	75.11	80.83	82.3	79.04	78.63	<b>82.32</b>
UCF101	Base	84.53	82.33	84.30	84.33	85.23	82.89	83.67	86.60	86.97	83.00	85.80	86.93	<b>87.13</b>
	New	67.37	73.45	76.33	74.94	71.97	82.87	75.43	75.90	77.48	78.66	77.23	78.33	<b>80.77</b>
	H	74.98	77.67	80.12	79.35	78.04	79.65	79.34	80.90	81.95	80.77	81.29	82.41	<b>83.83</b>

StanfordCars [21], Flowers [29], Food101 [3], FGVC-Aircraft [28], EuroSAT [14], UCF101 [39], DTD [6], and SUN397 [42]. A detailed introduction of those dataset is given in the Supplementary materials.

**Training Details.** Our implementation is based on CoOp’s code [50]. All experiments are conducted based on the CLIP with the backbone of ViT-B/16 [8]. The prompt’s length  $M$  is set as 4 with random initialization of “X X X X { }”. The final performance is averaged over three random seeds(1/2/3)for a fair comparison. Adam optimizer is applied for optimization with the learning rate of  $2e-3$  and the batch size of 32, and the training epochs is 50. All experiments are conducted on RTX 3090.

**Baselines.** Recently CoOp-based methods are used for comparison, e.g., CoOp [50], CoCoOp [49], ProGrad [51], ProDA [27], KgCoOp [43], PromptSRC [19], MaPLe [18], LFA [31], DePT [47], DAPT [5], PLOT [4], TaskRes [44], RPO [22], VPT [45], and TIP-Adapter-F [48].

## 4.2. Base-to-New Class Generalization

To verify the generalization ability of the prompt tuning, the base-to-new class generalization splits each dataset into two

disjoint subsets: *Base* and *New* classes, where *Base* classes are used to infer the learnable prompt and TKE, and *New* classes are used for evaluation. Note that each *New* class contains 16-shot samples. We summarize the comparison between the proposed TCP and existing methods in Table 1 from the following two aspects: Textual Prompt Tuning and Visual-Textual Prompt Tuning.

*Textual Prompt Tuning:* Based on the visual feature extracted from the frozen visual encoder, the textual prompt tuning (e.g., CoOp, CoCoOp, ProGrad, ProDA, and KgCoOp) merely conducts the textual prompt tuning by inserting a set of learnable textual prompts into the Text Encoder for obtaining a discriminative textual classifier, which is same as our proposed TCP. Among all textual prompt tuning methods, KgCoOp [43] is a more substantial baseline, which is also the baseline of our TCP. Compared to KgCoOp, TCP improves the metric term of *Base/New/H* from 80.73%/73.6%/77.0% to 84.35%/75.28%/79.56%. We can observe that a noticeable improvement of 3.62% and 1.68% are obtained in the *Base* class and *New* class, respectively. The superior performance proves that TCP using the class-aware prompt to inject the class knowledge can generate

Table 2. Comparison of cross-dataset evaluation. ‘tp’, ‘dtp’, ‘vp’, and ‘dvp’ denote the ‘textual prompt’, ‘deep textual prompt’, ‘visual prompt’, and ‘deep visual prompt’, respectively. Note that DAPT and MaPLe are based on visual-textual prompt tuning(‘vp+tp’).

Datasets	CLIP	CoOp	ProGrad	KgCoOp	DePT	VPT	PLOT	PromptSRC	MaPLe	DAPT	TCP
		tp	tp	tp	tp	tp+vp	tp+vp	dtp+vp	dtp+dvp	tp+vp	tp
ImageNet	66.70	71.51	72.24	70.66	<b>72.77</b>	69.73	71.60	71.27	70.72	71.60	71.40
Caltech101	93.30	93.70	91.52	93.92	<b>94.23</b>	93.67	92.07	93.60	93.53	93.50	93.97
OxfordPets	89.10	89.14	89.64	89.83	90.03	89.27	90.10	90.25	90.49	90.67	<b>91.25</b>
StanfordCars	65.70	64.51	62.39	65.41	65.57	65.5	65.70	65.70	65.57	<b>65.93</b>	64.69
Flowers	70.70	68.71	67.87	70.01	70.57	70.2	69.23	70.25	72.20	71.70	<b>71.21</b>
Food101	85.90	85.30	85.40	86.36	86.37	86.27	86.23	86.15	86.20	86.10	<b>86.69</b>
FGVCAircraft	24.90	18.47	20.16	22.51	23.27	22.13	<b>25.00</b>	23.90	24.74	23.03	23.45
SUN397	62.60	64.15	62.47	66.16	66.67	66.57	61.67	67.10	67.01	67.00	<b>67.15</b>
DTD	44.30	41.92	39.42	46.35	45.97	46.93	38.60	<b>46.87</b>	46.49	44.00	44.35
EuroSAT	48.30	46.39	43.46	46.04	43.53	47.43	47.83	45.50	48.06	<b>52.47</b>	51.45
UCF101	67.60	66.55	64.29	68.50	<b>69.30</b>	67.20	67.00	68.75	68.69	68.73	68.73
Avg.	65.24	63.88	62.71	65.51	65.55	65.52	64.34	65.81	66.30	<b>66.31</b>	66.29

the discriminative textual classifier for both seen and unseen classes. DePT [47] is another work that also treats KgCoOp as the baseline and obtains the best  $H$  of 79.10% among all textual prompt tuning methods. Compared to DePT, the proposed TCP obtains the performance of 84.35%/75.28%/79.56% in the metric term of  $Base/New/H$ , achieving the improvement of 0.51%/0.32%/0.41%.

*Visual-Textual Prompt Tuning:* Different from the textual prompt tuning merely adapts Text Encoder, DAPT [5], PLOT [4], RPO [22], MaPLe [18], and PromptSRC [19] all jointly performs prompt tuning on the text and visual encoders for both improving the discriminative of textual embedding and visual embeddings. Unlike CoOp and TCP, which conduct the prompt tuning on one or two layers, those methods conduct the visual-textual prompt tunings on several layers of visual and text encoders. Compared to them, the proposed TCP with shallow textual prompt tuning obtains a higher performance, *e.g.*, 79.51% vs 79.31% for PromptSRC in terms of  $H$ . Compared to deep visual-textual prompt tuning, the superior performance demonstrates that class-aware prompt tuning can generate discriminative and generalization textual-level classifiers.

### 4.3. Cross-Dataset Generalization

In the *Base-to-New* generalization setting, *New* classes always have a similar data distribution as the *Base* classes. To further verify the generalization of the proposed TCP, in the Cross-Dataset generalization, TCP is trained from the ImageNet and directly evaluated on the unrelated datasets, *e.g.*, the rest ten datasets. The comparison between the proposed TCP and existing methods is summarized in Table 2. From Table 2, we can observe that the proposed TCP obtains the highest average performance among all textual prompt tuning methods (66.29% vs 65.55% of DePT [47]), and obtains a comparable performance with the visual-textual prompt tuning methods (66.29% vs 66.31% of DAPT [5]), demonstrating the effectiveness of TCP in learning the generalization knowledge.

Table 3. Effect of TCP on CoOp, KgCoOp, ProGrad, PromptSRC, and DAPT in the base-to-new generalization setting regarding final average performance.

Methods	Base	New	H
CoOp [50]	82.63	67.99	74.6
CoOp+TKE	83.10 (↑ 0.47)	70.17 (↑ 2.18)	76.09 (↑ 1.49)
KgCoOp [43]	80.73	73.6	77.00
KgCoOp+TKE	84.13 (↑ 3.40)	75.36 (↑ 1.76)	79.51 (↑ 2.51)
ProGrad [51]	82.48	70.75	76.16
ProGrad+TKE	82.61 (↑ 0.13)	72.91 (↑ 2.16)	77.46 (↑ 1.30)
PromptSRC [19]	82.07	74.83	78.03
PromptSRC+TKE	83.74 (↑ 1.67)	75.85 (↑ 1.02)	79.60 (↑ 1.57)
DAPT [5]	83.18	69.27	75.59
DAPT+TKE	84.15 (↑ 0.97)	74.87 (↑ 5.60)	79.24 (↑ 3.65)

### 4.4. Few-shot Classification

To verify that the proposed TCP can infer the class-aware knowledge, the few-shot classification is conducted on all 11 datasets using K-shot labeled source images and evaluated on the standard testing domain with the same class space as the training classes. The 4-shot setting comparison between the proposed TCP and existing methods is summarized in Table 4. We can observe that the proposed TCP achieves a higher average performance than existing methods, *i.e.*, obtaining the average performance of 76.32%.

### 4.5. Ablation Study

**Effect of TKE’s generalization.** The critical component of TCP is the class-aware prompt generated by TKE, which embeds the prior class knowledge and can be easily integrated with existing CoOp-based methods. As shown in Table 3, CoOp, KgCoOp, ProGrad, PromptSRC, and DAPT have all demonstrated improved performance when TKE is used. Specifically, improvements of 1.49%, 2.51%, 1.30%, 1.57%, and 3.65% were achieved for CoOp, KgCoOp, ProGrad, PromptSRC<sup>3</sup>, and DAPT, respectively. This superior performance highlights TCP’s plug-and-play functionality,

<sup>3</sup>PromptSRC is re-implemented with the shallow visual-textual prompt tuning.

Table 4. Comparison of few-shot learning with 4-shot samples.

	CLIP	CoOp	CoCoOp	ProGrad	KgCoOp	MaPLE	TIP-Adapter-F	DAPT	PromptSRC	PLOT	TaskRes	TCP
ImageNet	66.70	69.37	70.55	70.21	70.19	70.67	70.78	70.80	<b>70.80</b>	70.40	62.87	70.48
Caltech101	93.30	94.44	94.98	94.93	94.65	94.30	94.77	94.23	94.77	<b>95.13</b>	94.67	95.00
OxfordPets	89.10	91.30	<b>93.01</b>	93.21	93.20	92.05	92.26	92.17	93.23	92.55	92.00	91.90
StanfordCars	65.70	72.73	69.10	71.75	71.98	68.70	74.42	74.40	71.83	74.93	75.90	<b>76.30</b>
Flowers	70.70	91.14	82.56	89.98	90.69	<b>80.80</b>	92.98	92.37	91.31	92.93	91.50	<b>94.40</b>
Food101	85.90	82.58	86.64	85.77	86.59	<b>86.90</b>	86.18	83.60	86.06	86.46	86.03	85.3
FGVCAircraft	24.90	33.18	30.87	32.93	32.47	29.03	35.49	32.47	32.80	35.29	33.80	<b>36.20</b>
SUN397	62.60	70.13	70.50	71.17	71.79	71.47	70.65	72.20	<b>72.80</b>	70.42	72.70	72.11
DTD	44.30	58.57	54.79	57.72	58.31	54.73	61.70	61.37	60.64	62.43	59.57	<b>63.97</b>
EuroSAT	48.30	68.62	63.83	70.84	71.06	54.87	78.27	72.73	75.02	<b>80.70</b>	72.87	77.43
UCF101	67.60	77.41	74.99	77.82	78.40	73.70	79.73	79.40	79.35	79.76	76.10	<b>80.83</b>
Avg.	65.37	73.59	71.98	74.21	74.48	70.66	76.11	75.07	75.33	76.45	74.36	<b>76.72</b>

Table 5. Domain-shared prompt vs Class-aware prompt.

Baseline	Domain-Shared	Class-Aware	Base	New	H
✓	✓		80.73	73.6	77
✓		✓	84.05	75.18	79.36
✓	✓	✓	<b>84.13</b>	<b>75.36</b>	<b>79.51</b>




Figure 3. Effect of  $M$  on the base-to-new generalization setting.

allowing easy integration with existing methods.

**Domain-shared prompt vs the class-aware prompt.** It is common for existing methods to infer the domain-shared prompt tokens. However, TCP differs significantly from them in using a class-aware prompt to capture class-aware knowledge in the Text Encoder. We thus compare the domain-shared and class-aware prompts, and summarize the related results in Table 5. We can observe that simply using the class-aware prompt yields a fantastic performance of 79.36% in terms of  $H$ , higher than existing methods. Additionally, as shown in Table 5, using the class-aware prompt results in a noticeable improvement over the domain-shared prompt. For instance, the *Base/New* performance is improved from 80.73%/73.6% to 84.05%/75.18%. The superior performance proves the reasonableness and effectiveness of injecting the class-aware knowledge into the prompt tuning. By further considering the domain-shared and class-aware prompts, the final performance is reached at 79.51%

**Effect of Prompt length  $M$ .** In this study, we examine the impact of prompt length ( $M$ ) on the *Base-to-New* generalization setting. We compared prompt lengths of 1, 2, 4, 8, and 16, and the results are presented in Figure 3. The

Table 6. Comparison of different templates.

Templates	Base	New	H
'X X X X '	<b>84.13</b>	<b>75.36</b>	<b>79.51</b>
'a photo of a '	83.94	74.94	79.18
'this is a picture '	84.00	74.66	79.06




Figure 4. Effect of dimension  $D_{mid}$  in TKE.

Figure 5. Effect of insert layer  $l$ .

findings indicate that longer prompt lengths produce better performance, with  $M = 8$  producing the best results. However, the performance gap between the highest and lowest values of  $H$ , which were 79.63% and 79.26%, respectively, is only 0.37%. Therefore, the proposed TCP is relatively insensitive to prompt length ( $M$ ).

**Effect of different templates.** Unlike most CoOp-based methods using the handcrafted template “a photo of a {}”, the proposed TCP utilizes the random initialization template (“X X X X {}”). Therefore, we conducted a comparison among different templates. Table 6 shows that using the random template performs better than the handcrafted template. This is because TCP uses the class-aware prompt without considering the domain-share prompt, which has achieved better performance. For instance, in Table 5, TCP with the class-aware prompt achieves comparable performance of 84.05%/75.18%/79.36% in terms of *Base/New/H* compared to existing methods. Moreover, the domain-shared prompt initialized as the random template provides complementary knowledge to the class-aware prompt.

**Effect of  $D_{mid}$  in TAK.** We analyze the effect of the dimension  $D_{mid}$  in TKE, and show the comparison in Figure 4. Setting  $D_{mid}=128$  obtains the best performance, and a smaller  $D_{mid}$  would get a worse performance.

**Insert which layer?** TCP inserts the class-aware textual prompts into a particular layer of the Text Encoder. The im-

Table 7. Comparison with single-layer vs multiple layers prompt tuning.

Modes	Layers	Base	New	H
TCP-Shallow	{8}	84.13	<b>75.36</b>	<b>79.51</b>
TCP-Deeper	{4:8}	84.02	75.13	79.33
TCP-Deeper	{8:10}	<b>84.24</b>	75.29	79.51
TCP-Deeper	{4:8:10}	84.11	74.51	79.02
TCP-Deeper	{4:6:8:10}	84.17	74.66	79.13

Dataset	TCP	CLIP-Adapter	CoOp-Adapter
ImageNet	74.0	73.0	73.0
SUN397	84.0	83.0	83.0
Caltech101	94.0	93.0	93.0
Oxford102	94.0	93.0	93.0
Sun397+Cal	74.0	73.0	73.0
Flowers	84.0	83.0	83.0
Fashion101	84.0	83.0	83.0
FGVC Aircraft	34.0	34.0	34.0
DTD	74.0	73.0	73.0
EuroSAT	84.0	83.0	83.0
UCF101	84.0	83.0	83.0
Avg	84.0	83.0	83.0

Figure 6. Comparison with TCP and Adapter-based methods.

part of inserting prompts into different layers was analyzed and shown in Figure 5. The best performance was achieved when the class-aware prompt tuning was done on the 8th layer. This is because doing so on earlier layers ( $l < 8$ ) does not transfer class-aware knowledge to the textual embedding, and doing it on later layers can easily affect the textual embedding.

**Single-layer vs Multiple-layers?** The results of existing MaPLe and PromptSRC show that deep prompt learning applied on multiple layers can perform better than shallow prompt tuning with a single layer. Therefore, we compared the performance of shallow prompt tuning with deep prompt tuning. The results in Table 7 indicate that the shallow class-aware prompt tuning with a single layer provides better performance. For instance, TCP-Shallow with 8-layer achieves the highest  $H$  of 79.51%, which surpasses all variants of TCP-deeper that have two, three, and four layers. Additionally, we observed that TCP-shallow obtains the best performance of 75.36% for the *New* class, and using more layers results in lower performance.

**Compared to Adapter-based methods:** In TCP, TKE acts as a special adapter, leading to the proposed TCP is similar to Adapter-based methods. We compared TCP with several Adapter-based methods such as CLIP-Adapter and CoOp-Adapter, which apply the adapter on the general textual-level class embeddings. However, TCP applies TKE to transfer the textual-level class embedding into the class-aware prompt combined with the class tokens in the middle layer of the Text Encoder. As shown in Figure 6, TCP with TKE performs better than the existing CLIP-Adapter and CoOp-Adapter on all 11 datasets.

**Training time analysis:** We further analyze the training time of the proposed TCP. The training time was determined based on the average time to process one image from the ImageNet with a 16-shot setting. According to Table 8, the proposed TCP is an efficient method that performs better

Table 8. Training time comparison(ms/image).

	CoOp	ProGrad	KgCoOp	PLOT	RPO	MaPLe	PromptSRC	TCP
ms/image	6ms	22ms	6ms	78.8ms	190ms	90.7ms	43.2ms	6.4ms
H	74.48	76.16	77.00	77.37	77.78	78.55	79.31	79.51




Figure 7. Visualization of the probability of CoOp and TCP.

with less running time. This is because TCP adds an efficient TKE component for generating class-aware prompt upon CoOp and KgCoOp, resulting in slightly more time consumption than CoOp and KgCoOp.

**Visualization.** To prove that the proposed TCP can generate the discriminative textual classifier for prediction, we further visualize the prediction probability among all classes. As shown in Figure 7, the proposed TCP has an obvious inter-class classes than existing methods.

## 5. Conclusion

To improve the generalization and discriminative abilities of the learnable prompt, we introduce a Textual-based Class-aware Prompt tuning method that takes advantage of the benefits of general class-level textual knowledge. To achieve this, we propose a Textual Knowledge Embedding (TKE) that transfers the class-level textual embedding into the class-aware prompt. This is combined with pre-trained class tokens to generate task-specific textual knowledge. Several benchmarks and tasks have shown that the class-aware prompt is effective for prompt tuning.

However, the class-aware prompt in TCP heavily relies on the discriminative ability of the general textual embedding. On the other hand, a weaker textual embedding will produce a weaker textual classifier. For example, TCP does not perform well on the FGVC Aircraft dataset. Therefore, in the future, we plan to explore ways to use weaker textual knowledge to obtain a discriminative textual classifier.

**Acknowledgements:** This work was supported by National Science and Technology Major Project(2021ZD0112202), National Natural Science Foundation of China (62376268,U23A20387,62036012,U21B2044,62376196), and Beijing Natural Science Foundation (4222039).



## References

- [1] Lorenzo Agnolucci, Alberto Baldrati, Francesco Todino, Federico Becattini, Marco Bertini, and Alberto Del Bimbo. Eco: Ensembling context optimization for vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2811–2815, 2023. [3](#)
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. [1](#), [2](#)
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 - mining discriminative components with random forests. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VI*, pages 446–461. Springer, 2014. [5](#)
- [4] Guangyi Chen, Weiran Yao, Xiangchen Song, Xinyue Li, Yongming Rao, and Kun Zhang. PLOT: prompt learning with optimal transport for vision-language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. [1](#), [3](#), [5](#), [6](#)
- [5] Eulrang Cho, Jooyeon Kim, and Hyunwoo J Kim. Distribution-aware prompt tuning for vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22004–22013, 2023. [1](#), [3](#), [5](#), [6](#)
- [6] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 3606–3613. IEEE Computer Society, 2014. [5](#)
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pages 248–255. IEEE Computer Society, 2009. [4](#)
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [5](#)
- [9] Li Fei-Fei, Robert Fergus, and Pietro Perona. 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, 2007. [4](#)
- [10] Zhe Gan, Linjie Li, Chunyuan Li, Lijuan Wang, Zicheng Liu, Jianfeng Gao, et al. Vision-language pre-training: Basics, recent advances, and future trends. *Foundations and Trends® in Computer Graphics and Vision*, 14(3–4):163–352, 2022. [2](#)
- [11] 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. *CoRR*, abs/2110.04544, 2021. [3](#)
- [12] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. [1](#)
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 15979–15988. IEEE, 2022. [2](#)
- [14] 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 J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 12(7):2217–2226, 2019. [5](#)
- [15] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. [1](#)
- [16] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 4904–4916. PMLR, 2021. [2](#)
- [17] Baoshuo Kan, Teng Wang, Wenpeng Lu, Xiantong Zhen, Weili Guan, and Feng Zheng. Knowledge-aware prompt tuning for generalizable vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15670–15680, 2023. [1](#), [3](#)
- [18] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19113–19122, 2023. [1](#), [3](#), [5](#), [6](#)
- [19] Muhammad Uzair Khattak, Syed Talal Wasim, Muzammal Naseer, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Self-regulating prompts: Foundational model adaptation without forgetting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15190–15200, 2023. [3](#), [5](#), [6](#)
- [20] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 5583–5594. PMLR, 2021. [2](#)
- [21] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *2013 IEEE International Conference on Computer Vision*

- Workshops, ICCV Workshops 2013, Sydney, Australia, December 1-8, 2013*, pages 554–561. IEEE Computer Society, 2013. 5
- [22] Dongjun Lee, Seokwon Song, Jihee Suh, Joonmyeong Choi, Sanghyeok Lee, and Hyunwoo J Kim. Read-only prompt optimization for vision-language few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1401–1411, 2023. 1, 2, 5, 6
- [23] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2
- [24] Juncheng Li, Minghe Gao, Longhui Wei, Siliang Tang, Wenqiao Zhang, Mengze Li, Wei Ji, Qi Tian, Tat-Seng Chua, and Yueting Zhuang. Gradient-regulated meta-prompt learning for generalizable vision-language models. *arXiv preprint arXiv:2303.06571*, 2023. 2
- [25] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 2
- [26] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23, 2019. 2
- [27] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5196–5205. IEEE, 2022. 3, 5
- [28] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *CoRR*, abs/1306.5151, 2013. 5
- [29] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Sixth Indian Conference on Computer Vision, Graphics & Image Processing, ICVGIP 2008, Bhubaneswar, India, 16-19 December 2008*, pages 722–729. IEEE Computer Society, 2008. 5
- [30] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 1717–1724. IEEE Computer Society, 2014. 1
- [31] Yassine Ouali, Adrian Bulat, Brais Matinez, and Georgios Tzimiropoulos. Black box few-shot adaptation for vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15534–15546, 2023. 1, 2, 5
- [32] Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012*, pages 3498–3505. IEEE Computer Society, 2012. 4
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 8748–8763. PMLR, 2021. 1, 2
- [34] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 18061–18070. IEEE, 2022. 3
- [35] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 2
- [36] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 2
- [37] Sheng Shen, Shijia Yang, Tianjun Zhang, Bohan Zhai, Joseph E. Gonzalez, Kurt Keutzer, and Trevor Darrell. Multi-task vision-language prompt tuning. *CoRR*, abs/2211.11720, 2022. 3
- [38] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15638–15650, 2022. 2
- [39] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012. 5
- [40] Jingchen Sun, Jiayu Qin, Zihao Lin, and Changyou Chen. Prompt tuning based adapter for vision-language model adaption. *CoRR*, abs/2303.15234, 2023. 1
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. 2
- [42] Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. SUN database: Large-scale scene recognition from abbey to zoo. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, pages 3485–3492. IEEE Computer Society, 2010. 5
- [43] Hantao Yao, Rui Zhang, and Changsheng Xu. Visual-language prompt tuning with knowledge-guided context optimization. In *Proceedings of the IEEE/CVF Conference*

- on *Computer Vision and Pattern Recognition*, pages 6757–6767, 2023. [1](#), [2](#), [3](#), [5](#), [6](#)
- [44] Tao Yu, Zhihe Lu, Xin Jin, Zhibo Chen, and Xinchao Wang. Task residual for tuning vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 10899–10909. IEEE, 2023. [5](#)
- [45] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. *CoRR*, abs/2210.07225, 2022. [1](#), [2](#), [3](#), [5](#)
- [46] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12104–12113, 2022. [2](#)
- [47] Ji Zhang, Shihan Wu, Lianli Gao, Hengtao Shen, and Jingkuan Song. Dept: Decoupled prompt tuning. *arXiv preprint arXiv:2309.07439*, 2023. [5](#), [6](#)
- [48] 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 *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXV*, pages 493–510. Springer, 2022. [5](#)
- [49] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 16795–16804. IEEE, 2022. [1](#), [3](#), [5](#)
- [50] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *Int. J. Comput. Vis.*, 130(9):2337–2348, 2022. [1](#), [2](#), [4](#), [5](#), [6](#)
- [51] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15659–15669, 2023. [3](#), [5](#), [6](#)
- [52] Beier Zhu, Yulei Niu, Saeil Lee, Minhoe Hur, and Hanwang Zhang. Debaised fine-tuning for vision-language models by prompt regularization. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 3834–3842. AAAI Press, 2023. [1](#), [3](#)