Enhancing Target-unspecific Tasks through a Features Matrix

Fangming Cui¹ Yonggang Zhang² Xuan Wang³ Xinmei Tian⁴ Jun Yu^{† 5}

Abstract

Recent developments in prompt learning of large Vision-Language Models (VLMs) have significantly improved performance in target-specific tasks. However, these prompting methods often struggle to tackle the target-unspecific or generalizable tasks effectively. It may be attributed to the fact that overfitting training causes the model to forget its general knowledge. The general knowledge has a strong promotion on target-unspecific tasks. To alleviate this issue, we propose a novel Features Matrix (FM) approach designed to enhance these models on target-unspecific tasks. Our method extracts and leverages general knowledge, shaping a Features Matrix (FM). Specifically, the FM captures the semantics of diverse inputs from a deep and fine perspective, preserving essential general knowledge, which mitigates the risk of overfitting. Representative evaluations demonstrate that: 1) the FM is compatible with existing frameworks as a generic and flexible module, and 2) the FM significantly showcases its effectiveness in enhancing target-unspecific tasks (base-to-novel generalization, domain generalization, and cross-dataset generalization), achieving state-of-the-art performance.

1. Introduction

Large vision-language models such as CLIP (Radford et al., 2021) have attracted increasing attention for remarkable generalization performance. Vision-Language Models (VLMs) are trained to align textual and visual modalities by leveraging extensive datasets. For instance, a prominent example of such models is CLIP, which has achieved remarkable success across a wide range of downstream tasks (Greer et al., 2024; Sanghi et al., 2022; Etchegaray et al., 2024). CLIP utilizes a large collection of 400 million text-image pairs. One of CLIP's most appealing features is its ability to perform zero-shot inference. During inference, CLIP utilizes hand-crafted text inputs, known as prompts, to generate classification weights to predict image features, all without requiring any target-specific parameters training.

In contrast to hand-crafted prompts, a model-parameter tuning method has been proposed as prompt learning to automatically learn prompt embeddings (Zhou et al., 2022a). For instance, CoOp (Zhou et al., 2022b) represents the first method that specifically focuses on learning the text embeddings of prompts with few-shot samples training while keeping the CLIP model frozen. Although CoOp (Zhou et al., 2022b) has shown significant performance improvements over hand-crafted prompts in target-specific base classes, it may yield inferior performance compared to the handcrafted prompt CLIP in generalization tasks, e.g., the novel class of generalization from base-to-novel (see Table 1). To overcome this challenge, some effective methods (Lu et al., 2022; Zhu et al., 2023; Zhou et al., 2022a;b; Yao et al., 2023a; Chen et al.; Li et al., 2024a) with tuning textual embeddings have been proposed. These methods aim to enhance the performance of novel classes, surpassing the novel performance of the previous CoOp. Although the novel ability of these methods surpasses that of CoOp (67.96%), it is unfortunate that the novel ability of these methods still falls short compared to the hand-crafted prompt method (CLIP: 74.22%), as shown in Table 1 (Novel). A possible reason (Zhou et al., 2022a) may be that these methods with tuning text embeddings tend to overfit the downstream data distributions. This overfitting can result in the model losing its inherent generalization capabilities (Yao et al., 2023a) obtained from hand-crafted prompts (Radford et al., 2021).

To alleviate this problem, we propose a novel Features Matrix (FM) with CLIP for target-unspecific tasks. Our method incorporates multiple hand-crafted prompts to extract general information as a pre-trained matrix from frozen CLIP to enhance generalization. The matrix of pre-trained features can delve into the semantics of different hand-crafted prompts finely and deeply, which reduces the risk of forgetting the essential general knowledge of pre-trained CLIP. Importantly, our method is compatible with current prompt learning frameworks for textual or multi-modal prompt learning and serves as a flexible and generic module. As

¹Shanghai Jiao Tong University ²Hong Kong Baptist University ³Meituan Inc. ⁴University of Science and Technology of China ⁵Harbin Institute of Technology (Shenzhen). Correspondence to: Jun Yu[†] <yujun@hit.edu.cn>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).



Figure 1. Our method is orthogonal to representative architectures, such as MaPLe (M) and PromptSRC (S), surpassing the existing easy-to-use DePT (Zhang et al., 2024a) by a significant margin.

demonstrated in Figure 1, our method, integrated with the representative MaPLe (Khattak et al., 2023a) and Prompt-SRC (Khattak et al., 2023b), exhibits competitive results in base-to-novel generalization, domain generalization, and cross-dataset generalization across 11 datasets. Meanwhile, our method surpasses the existing easy-to-use DePT (Zhang et al., 2024a) by a significant margin. This underscores the robustness of our method across various tasks.

Our contributions can be summarized as follows:

- The proposed method incorporates multiple handcrafted prompts with classes to extract general knowledge as a pre-trained Features Matrix (FM).
- We propose an easy-to-use design, which is compatible with representative textual or multi-modal prompt learning frameworks for adapting CLIP.
- Various target-unspecific tasks (base-to-novel generalization, cross-dataset generalization, and domain generalization) across 11 datasets demonstrate that our method demonstrates its effectiveness.

2. Preliminaries

Notations. Considering a pre-trained VLM, let $\mathcal{E}_v(\cdot)$ be its image encoder and $\mathcal{E}_t(\cdot)$ be its text encoder. The image encoder transforms input images into feature embeddings, capturing the visual information within the images. The text encoder generates representations for word embedding sequences, capturing the semantic information conveyed by the text prompts **p**. Generally, a hand-crafted prompt **p** may have the form of "a photo of a [Class]". In this paper, **x** represents an arbitrary image, and *l* denotes the label.

Hand-Crafted CLIP. During the pre-training phase of CLIP (Radford et al., 2021), the image and text encoders are jointly trained on large-scale text-image pairs using a contrastive loss. This loss maximizes the cosine similarity between matching pairs and minimizes it between non-matching pairs, enabling the encoders to learn aligned visual and textual representations effectively. The final prediction score between the image x and text prompt **p** is computed using contrastive learning. The final prediction probability of alignment is computed by the matching score as follows:

$$p(l = k \mid \mathbf{x}) = \frac{\exp\left\{\cos\left(\mathcal{E}_{t}\left(\mathbf{p}_{k}\right), \mathcal{E}_{v}(\mathbf{x})\right)/\tau\right\}}{\sum_{k'=1}^{K}\exp\left\{\cos\left(\mathcal{E}_{t}\left(\mathbf{p}_{k'}\right), \mathcal{E}_{v}(\mathbf{x})\right)/\tau\right\}},$$
(1)

where *l* is the label of \mathbf{x} , $\cos(\cdot, \cdot)$ stands for cosine similarity between two vectors, and $\tau > 0$ represents a temperature parameter. Here, the classifier consists of *K* textual features derived from prompts $\{\mathbf{p}_{k'}\}_{k'=1}^{C}$, where the prompt $\mathbf{p}_{k'}$ for the *k'*-th class may have the form of "a photo of a".

Textual Prompting. In textual prompting, the class name is retained as prior knowledge, while the word embeddings (referred to as context) of prompts are treated as learnable parameters, as shown in CoOp (Zhou et al., 2022b) and Co-CoOp (Zhou et al., 2022a) of Figure 2. By modeling these context embeddings as trainable parameters, the model can optimize the prompts based on the specific requirements, enhancing the alignment between images and prompts. For class k, the tuning feature of the text encoder is denoted as t_k^{tun} in a dataset with a total of K classes. The final prediction probability of cross-entropy loss for two-modalities alignment is computed as:

$$p(l = k \mid \mathbf{x}) = \frac{\exp\left\{\cos\left(t_k^{tun}, \mathcal{E}_v(\mathbf{x})\right) / \tau\right\}}{\sum_{k'=1}^{K} \exp\left\{\cos\left(t_{k'}^{tun}, \mathcal{E}_v(\mathbf{x})\right) / \tau\right\}}, \quad (2)$$

where $\cos(\cdot, \cdot)$ stands for cosine similarity, and τ denotes a temperature parameter. In some learned text embedding methods, the method aligns with the CoOp methodology by setting the embeddings to be shared across different classes. During the inference phase, the prompts, along with the learned embeddings, can generate textual features that are used for classification purposes. By leveraging these learned embeddings, the model can produce more effective representations, leading to improved target-specific classification performance (Zhou et al., 2022b).

Multi-Modal Prompting. Based on textual deep prompting, this design uses only a limited number of trainable parameters based on the image encoder, as shown in MaPLe (Khattak et al., 2023a) and PromptSRC (Khattak et al., 2023b) of Figure 2. The visual prompts are introduced at every transformer layer's input space. In the visual branch, the input image $\mathbf{x} \in \mathbb{R}^{C \times H \times W}$ is divided into M patches. And, a tuning embedding of class i_{class} is appended with the input

Enhancing Target-unspecific Tasks through a Features Matrix



Figure 2. Illustration of representative textual prompting frameworks (CoOp and CoCoOp) and multi-modal prompting frameworks (MaPLe and PromptSRC). We propose a flexible and generic design, which is compatible with these representative architectures. In the figure, "snowflake pattern" represents parameter freezing, "flame pattern" represents learnable pattern, "Deep" represents learnable tokens embedded in several layers of the encoder, where "T" represents learnable text embedding, "V" represents learnable visual embedding, "Class priors" represents which category it belongs to, "Tuning Similarity" represents the calculation of cosine similarity for fine-tuning architecture, and the light gray "Similarity" represents the calculation of cosine similarity for frozen architecture. In the MaPLe architecture diagram, the "e" represents the matrix function that connects the encoders of two modalities in several layers.

patches. The foundation of vision prompting is ViT (Vision Transformer) (Dosovitskiy et al., 2020) backbone, which shares the same image encoder as CLIP. Let \mathcal{P}_v denote the visual embeddings of prompts, and the image encoder processes the input tokens to generate tuning visual features, as follows (Jia et al., 2022):

$$\tilde{\mathbf{x}}_p = \{\mathcal{P}_v, \boldsymbol{i}_{class}, \boldsymbol{i}_1, \boldsymbol{i}_2, \cdots, \boldsymbol{i}_M\}.$$
(3)

The tuning of visual embeddings are introduced in the image encoder for deep prompt tuning.

3. Methodology

3.1. Current Challenge

As shown in Table 1, such as textual prompting (Zhou et al., 2022a;b; Yao et al., 2023a; Lu et al., 2022) and multi-modal prompting (Chen et al.), have shown significant improvements in target-specific base classes. However, these methods often suffer from overfitting, leading to poor generalization in novel classes. To improve the novel performance of models, prompting regularization involves constraining the pre-trained and fine-tuning features in the text branch through computed loss. This constraint helps prevent forgetting general knowledge and ensures that the model retains essential information for downstream tasks. The objective of this design is to activate and maintain the pre-trained generalization

abilities to mitigate performance degradation when dealing with target-unspecific classes.

Table 1. Our method is an easy-to-use design, integrated into CoOp, obtaining a higher average performance (11 datasets) on novel (target-unspecific) classes.

Method	Hand Features	Base	Novel
CLIP (ICML2021)	Single	69.34	74.22
CoCoOp (CVPR2022)	No	80.47	71.69
ProDA (CVPR2022)	No	81.56	72.30
PLOT (ICLR2023)	No	81.24	72.98
ProGrad (ICCV2023)	No	82.48	70.75
CoOp (IJCV2022)	No	82.38	67.96
+ DePT (CVPR2024)	No	83.66	71.82
+ Kg (CVPR2023)	Single	80.73	72.70
+ Ours	Matrix	81.15	74.66

It is important to recognize that the fixed CLIP's fundamental pre-trained characteristics demonstrate robust generalization capabilities. Recently, KgCoOp (Yao et al., 2023a) incorporates the constraint with L1 loss on fine-tuning feature with the pre-trained feature for the textual branch, avoiding the loss of CLIP's original generalization capability.

Motivation. However, the novel class performance of handcrafted-based KgCoOp (72.70%) is still lower than pretrained CLIP (74.22%) (Radford et al., 2021) of Table 1. It



Figure 3. Illustration of our easy-to-use method. We propose a novel Features Matrix (FM) for enhancing target-unspecific tasks. Our method incorporates multiple hand-crafted prompts with classes to extract general knowledge as a pre-trained features matrix. Various generalization tasks across 11 datasets demonstrate that our method outperforms existing prompt learning methods.

may be attributed to the fact that this single hand-crafted prompting regularization with CLIP cannot fully release and excavate the diverse semantics for utilizing essential general knowledge (Khattak et al., 2023b).

3.2. Proposed Features Matrix (FM)

To tackle this challenge and problem, we propose a novel design called Features Matrix (FM) for enhancing targetunspecific tasks. Our method incorporates multiple handcrafted prompts with classes of the same datasets to extract general information as a pre-trained features matrix from CLIP to enhance generalization. In this way, we can explore and excavate the semantic information brought by each hand-crafted prompt more finely and deeply. By specifically aligning these multiple pre-trained hand-crafted unexpected features and tuning image features, our method aims to enhance the model's generalization and robustness. This involves focusing on indistinguishable features that pose a challenge to the model during training. As shown in Figure 3, we employ a set of hand-crafted prompt templates (P1, P2, P3, P4, etc) as inputs in the text encoder of pretrained frozen CLIP. These prompt templates, such as "a photo of one", "a photo of a ", "a picture of a", "a drawing of a", etc. The number of prompt templates (Radford et al., 2021) is set to 60. In a dataset containing K classes, the input to the text encoder is the frozen embeddings of handcrafted prompts, consisting of different classes. The output of the frozen text encoder can be viewed as a features matrix, as depicted in Figure 3. Specifically, we let t denote the sets of pre-trained text features extracted by a frozen text encoder

through a set of hand-crafted prompt templates (P1, P2, P3, P4, etc) with classes. In a dataset with a total of K classes, features belonging to the current label k of image samples are referred to as designated features t_k , while features from other classes are considered non-designated features t_k .

Our method computes the matching scores with the cosine distance for each visual output feature v^{tun} and designated pre-trained text features of the features matrix as $\cos(t_k, v^{tun})$. Similarly, we compute the matching scores between each visual output feature and the non-designated pre-trained text features, denoted as $\cos(t_{\hat{k}}, v^{tun})$. In correspondence with the features matrix, these matching scores can be viewed as a scores matrix, as depicted in Figure 3. Subsequently, we identify the textual feature with lower scores (low- β) among the designated features, which is referred to as the unexpected designated features set F_{un}^k . The low- β is that there is a list of scores sorted from low to high, with β values assigned from the front of the list. Similarly, we identify textual features with higher scores (top- β) among the non-designated features, which are referred to as unexpected non-designated feature sets $F_{un}^{\hat{k}}$. Accordingly, the contrastive loss \mathcal{L}_{CL} of unexpected features similarity is denoted as follows,

$$-\log\frac{\exp\left\{\cos\left(t_{k}, v^{tun}\right)\right\}}{\exp\left\{\cos\left(t_{k}, v^{tun}\right)\right\} + \exp\left\{\cos\left(t_{\hat{k}}, v^{tun}\right)\right\}}, \quad (4)$$

where the $t_k \in F_{un}^k$ and $t_{\hat{k}} \in F_{un}^{\hat{k}}$ denote the selected unexpected features for designated unexpected features and non-designated unexpected features. The objective is to explore and excavate the unexpected features of general information. We train the objective for aligning the unexpected text features and tuning visual features, and the contrastive loss \mathcal{L}_{CL} optimizes visual embeddings. For twomodal alignment, let $\mathcal{L}_{CE}(\cdot, \cdot)$ denote the cross-entropy loss with prompts **p** for samples S, as follows,

$$\mathcal{L}_{CE} = \arg\min_{\mathbf{p}} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{S}} \mathcal{L} \left\{ \cos\left(t^{tun}, v^{tun}\right), l \right\}.$$
 (5)

Accordingly, the objective $\mathcal{L}_{\text{total}}$ with a hyper-parameter γ of our method can be formulated as follows,

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CE}} + \gamma \mathcal{L}_{\text{CL}}.$$
 (6)

Consequently, optimizing models with \mathcal{L}_{CL} between the pre-trained textual unexpected features and image tuning features. The \mathcal{L}_{CE} represents two modalities of alignment of deep vision-language prompting. Importantly, our method is compatible with representative prompt learning frameworks as a generic and flexible module for textual prompt learning, such as CoOp and CoCoOp, or multi-modal prompt learning, such as MaPLe and PromptSRC.

4. Main Generalization Tasks

We evaluate our method on generalization tasks. As evidenced by systematic benchmarking in Figure 1, our framework, when synergistically integrated with MaPLe (Khattak et al., 2023a) and PromptSRC (Khattak et al., 2023b), achieves state-of-the-art performance across three critical generalization dimensions (base-to-novel generalization, domain generalization, and cross-dataset generalization) spanning 11 heterogeneous benchmarks. Notably, our method surpasses the existing easy-to-use DePT (Zhang et al., 2024a) by a significant margin in generalization scenarios through low-resource training.

4.1. Benchmark Setting

Compared Methods. In three generalization experiments, we compare our method with CoOp (IJCV2022) (Zhou et al., 2022b), CoCoOp (CVPR2022) (Zhou et al., 2022a), MaPLe (CVPR2023) (Khattak et al., 2023a), Prompt-SRC (ICCV2023) (Khattak et al., 2023b), and DePT (CVPR2024) (Zhang et al., 2024a). The DePT is an easy-to-use method. We process frozen CLIP using a linear probe method. Comparison methods are based on the ViT-B/16 architecture for fair comparison.

Implementation Details. We employ CLIP (Radford et al., 2021) model based on the ViT-B/16 architecture. For the PromptSRC-based and MaPLe-based frameworks, we set the visual and textual embedding length to 4. We set the easy-to-use module γ to 0.1 and matching scores (top and low) β to 5. Training for 30 epochs for a base-to-novel setting in the first 9 transformer layers. Training for 20

epochs for the domain generalization setting and the crossdataset evaluation setting in the first 3 transformer layers. We use a SGD optimizer with a learning rate of 0.0025 on a single GPU. Following the hand-crafted CLIP, we use the hand-crafted template sets from (Radford et al., 2021).

Datasets. In Table 10, the datasets cover multiple recognition tasks including ImageNet (Deng et al., 2009) and Caltech101 (Fei-Fei et al., 2004) which consists of generic objects, OxfordPets (Parkhi et al., 2012), StanfordCars (Krause et al., 2013), Flowers102 (Nilsback & Zisserman, 2008), Food101 (Bossard et al., 2014), and FGVCAircraft (Maji et al., 2013) for fine-grained classification, SUN397 (Xiao et al., 2010) for scene recognition, UCF101 (Soomro et al., 2012) for action recognition, DTD (Cimpoi et al., 2014) for texture classification, and EuroSAT (Helber et al., 2019) which consists of satellite images. We leverage datasets such as ImageNetA (Hendrycks et al., 2021b), ImageNet-R (Hendrycks et al., 2021a), ImageNet-Sketch (Wang et al., 2019), and ImageNetV2 (Recht et al., 2019) to assess the model's performance across different domain distributions.

4.2. Base-to-Novel Generalization Task

In the base-to-novel generalization task, the datasets are divided into base and novel classes. The model is trained on the base classes in a 16-shot setting, and tested on both the base and novel classes across 11 different datasets. The number of classes for base and novel is the same, which means that all classes in a dataset are evenly divided into two groups of classes. The process of dividing all classes in the dataset is randomly selected. HM refers to harmonic mean. The HM evaluates the generalizable and non-generalizable ability of our methods.

Discussion. In Table 2, our method demonstrates significant improvements on all 11 datasets on HM. We found that our performance is higher than DePT (Zhang et al., 2024a) (CVPR2024) when integrated with different methods. Compared to DePT, our method significantly enhances the performance of base classes while also improving the accuracy on novel classes. In terms of average performance, our method based on PromptSRC (S) attains 85.70% accuracy on the base classes, and 77.35% accuracy on the novel classes. Our method, as an easy-to-use module, achieves performance improvements for various baseline frameworks, such as CoOp, CoCoOp, MaPLe (M), and PromptSRC (S).

4.3. Domain Generalization Task

We train our model using the ImageNet in 16 shots and we leverage ImageNetA, ImageNet-R, ImageNet-Sketch, and ImageNetV2 to assess the model's performance. The training ImageNet is a non-generalizable setting in Table 3 (column-1). This experiment aims to validate the potential of our method in domain shifts (Fang et al., 2023). *Table 2.* Base-to-novel generalization. In the base-to-novel generalization task, the datasets are divided into base and novel classes. The model is trained on the base classes in a 16-shot setting, and tested on both the base and novel classes across 11 different datasets. Compared with easy-to-use DePT (CVPR2024) in detail, our method achieves consistent average performance improvement over different representative baselines (CoOp, CoCoOp, MaPLe, PromptSRC).

(a) Average over 11 datasets			(1	b) Imag	eNet		(c) Caltee	ch101		(d) OxfordPets				
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CoOp	82.69	63.22	71.66	CoOp	76.47	67.88	71.92	CoOp	98.00	89.81	93.73	CoOp	93.67	95.29	94.47
+DePT	83.66	71.82	77.29	+DePT	77.13	70.10	73.45	+DePT	98.33	94.33	96.29	+DePT	94.70	97.63	96.14
+ Ours	81.15	74.66	77.79	+ Ours	75.85	71.33	73.53	+ Ours	97.58	96.60	97.13	+ Ours	93.78	97.80	95.78
Co	80.47	71.69	75.83	Co	75.98	70.43	73.10	Co	97.96	93.81	95.84	Co	95.20	97.69	96.43
+DePT	83.80	72.89	77.97	+DePT	76.87	70.47	73.53	+DePT	98.37	93.87	96.06	+DePT	94.03	97.20	95.59
+ Ours	81.68	75.55	78.52	+ Ours	77.35	72.36	74.79	+ Ours	98.61	96.75	97.70	+ Ours	95.33	98.15	96.76
M	82.28	75.14	78.55	M	76.66	70.54	73.47	M	97.74	94.36	96.02	M	95.43	97.76	96.58
+DePT	84.85	74.82	79.52	+DePT	77.87	70.23	73.85	+DePT	98.53	95.03	96.75	+DePT	95.03	97.83	96.41
+ Ours	84.45	76.53	80.32	+ Ours	78.18	71.38	74.62	+ Ours	98.35	96.11	97.22	+ Ours	95.85	98.22	97.04
S	84.26	76.10	79.97	S	77.60	70.73	74.01	S	98.10	94.03	96.02	S	95.33	97.30	96.30
+ DePT	85.19	76.17	80.43	+ DePT	78.20	70.27	74.02	+ DePT	98.57	94.10	96.28	+ DePT	95.43	97.33	96.37
+ Ours	85.70	77.35	81.32	+ Ours	78.90	71.58	75.07	+ Ours	98.62	95.88	97.27	+ Ours	95.95	97.92	96.95

(e) Euro	SAT	(f) UCF101				(g) StanfordCars				(h) Flowers102				
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CoOp	92.19	54.74	68.69	CoOp	84.69	56.05	67.46	CoOp	78.12	60.40	68.13	CoOp	97.60	59.67	74.06
+DePT	88.27	66.27	75.70	+ DePT	85.43	72.17	78.24	+ DePT	79.67	72.40	75.86	+ DePT	98.20	72.00	83.08
+ Ours	88.35	65.33	75.13	+ Ours	83.10	78.85	80.94	+ Ours	74.32	76.87	75.58	+ Ours	96.22	72.32	82.57
Co	87.49	60.04	71.21	Co	82.33	73.45	77.64	Co	70.49	73.59	72.01	Co	94.87	71.75	81.71
+DePT	90.27	66.87	76.82	+ DePT	85.70	72.80	78.73	+ DePT	79.87	73.33	76.46	+ DePT	98.33	72.57	83.51
+ Ours	88.15	70.11	78.12	+ Ours	82.95	77.21	80.00	+ Ours	72.88	76.10	74.46	+ Ours	95.61	74.93	84.03
M	94.07	73.23	82.35	M	83.00	78.66	80.77	M	72.94	74.00	73.47	M	95.92	72.46	82.56
+DePT	94.43	76.23	84.36	+ DePT	86.87	78.10	82.25	+ DePT	80.93	71.73	76.06	+ DePT	98.03	73.17	83.79
+ Ours	94.22	75.65	83.93	+ Ours	87.33	79.10	83.02	+ Ours	78.66	75.13	76.86	+ Ours	98.21	75.00	85.07
S	92.90	73.90	82.32	S	87.10	78.80	82.74	S	78.27	74.97	76.58	S	98.07	76.50	85.95
+ DePT	92.23	77.90	84.88	+ DePT	87.73	77.70	82.46	+ DePT	80.80	75.00	77.79	+ DePT	98.40	77.10	86.46
+ Ours	95.50	76.85	85.17	+ Ours	88.80	79.50	83.93	+ Ours	80.91	76.51	78.68	+ Ours	98.81	78.10	87.26

(i) Food101	(j) FGVCAircraft	(k) SUN397	(l) DTD			
Base Novel HM	Base Novel HM	Base Novel HM	Base Novel HM			
CoOp 88.33 82.26 85.19 + DePT 90.43 91.33 90.8 + Ours 89.98 92.85 91.4	CoOp 40.44 22.30 28.75 + DePT 42.53 22.53 29.46 + Ours 37.32 34.61 35.92	CoOp80.6065.8972.51+ DePT82.3775.0778.55+ Ours79.1278.3878.77	CoOp 79.44 41.18 54.24 + DePT 83.20 56.13 67.04 + Ours 77.10 56.38 65.15			
Co90.7091.2990.94+ DePT90.3091.3090.86+ Ours90.6191.9391.23	Co 33.41 23.71 27.74 + DePT 43.07 31.30 36.25 + Ours 37.87 34.91 36.33	Co 79.74 76.86 78.27 + DePT 82.20 76.73 79.37 + Ours 80.32 79.00 79.68	Co77.0156.0064.85+ DePT82.7755.4066.37+ Ours78.9059.6167.93			
M 90.71 92.05 91.33 + DePT 90.33 91.53 90.93 + Ours 90.31 92.81 91.55	B M 37.44 35.61 36.50 + DePT 44.53 32.80 37.78 + Ours 42.46 37.62 39.89	M 80.82 78.70 79.75 + DePT 82.90 76.40 79.52 + Ours 82.35 79.81 81.07	M 80.36 59.18 68.16 + DePT 83.87 59.93 69.91 + Ours 83.01 60.98 70.32			
S 90.67 91.53 91.10 + DePT 90.87 91.57 91.22 + Ours 90.61 92.30 91.43	S 42.73 37.87 40.15 + DePT 45.70 36.73 40.73 5 + Ours 45.81 39.11 42.20	S82.6778.4780.52+ DePT83.2778.9781.06+ Ours83.9080.5182.20	S 83.37 62.97 71.75 + DePT 84.80 61.20 71.09 + Ours 84.90 62.58 72.07			

Table 3. Domain generalization. We train our model using the ImageNet in 16 shots and test its performance on 4 different variants of the ImageNet. When compared to DePT, our method consistently outperforms it across all ImageNet variant datasets.

	Source					
	ImageNet	-V2	-S	-A	-R	Average
CoOp	71.51	64.2	47.99	49.71	75.21	59.28
+ DePT	72.63	64.80	48.05	50.00	75.50	59.58
+ Ours	71.82	65.13	48.10	50.15	76.13	59.87
CoCoOp	71.02	64.07	48.75	50.63	76.18	59.91
+ DePT	72.77	65.10	49.10	51.00	76.85	60.51
+ Ours	72.10	65.25	50.13	52.11	77.18	61.16
MaPLe	70.72	64.07	49.15	50.9	76.98	60.27
+ DePT	73.27	65.33	49.05	51.25	77.50	60.78
+ Ours	71.56	65.45	50.33	52.32	78.10	61.55
PromptSRC	71.27	64.35	49.55	50.90	77.80	60.65
+ DePT	71.60	64.51	50.15	51.88	77.18	60.93
+ Ours	71.41	65.50	51.33	52.00	78.85	61.92

Discussion. In Table 3, our method based on PromptSRC demonstrates the highest average improvement of 1.27% over the PromptSRC method. Additionally, when compared to DePT, our method consistently outperforms it across all ImageNet variant datasets.

4.4. Cross-Dataset Generalization Task

We train our model with 16 shots on the ImageNet dataset and test the model on 10 other unseen datasets. The training ImageNet is a non-generalizable setting in Table 4 (column-1). This experiment aims to validate the potential of our method in a wide range of cross-dataset transfers.

Discussion. As shown in Table 4, our method based on PromptSRC shows competitive performance in 10/10 over the generic and flexible method DePT. These findings suggest that our method excels in achieving better generalization across a diverse range of unseen datasets. However, our method for training ImageNet is lower than that of DePT.

5. Further Studies

In this section, we discuss further studies, focusing on the impact of embedding length and depth, computational cost, applying to other ViT instances, analysis of γ for \mathcal{L}_{CL} , and analysis of Top- β and Low- β for pre-trained features matrix.

5.1. Learning Depth

In Table 5, we note that increasing the learning depth generally increases the performance based on PromptSRC. As the number of layers increases to 11, the HM value decreases.

5.2. Embeddings Length

In Table 7, our findings indicate that the performance based on PromptSRC reaches its peak when the length of embeddings is set to 4 on HM for an average of 11 datasets. Our ablation studies are based on keeping all other settings unchanged. It indicates excessive fine-tuning of the model, causing it to lose CLIP generalization.

5.3. Analysis of γ for $\mathcal{L}_{\mathrm{CL}}$

In Table 9, it is observed that as the number of γ increases to 0.1, there is a peak in HM for an average of 11 datasets. We conduct analysis based on the PromptSRC model.

5.4. Analysis of Top- β and Low- β

In Table 8, it is observed that as the number of Top- β and Low- β increases to 5, there is a peak in HM. We conduct analysis based on the PromptSRC model for an average of 11 datasets. We keep other hyper-parameters unchanged.

5.5. Applying to Other ViT Instances

In Table 6, we apply our method in various types of ViT instances. Our experimental performance is based on the average performance of 11 datasets. Our generic and flexible module based on PromptSRC has the lowest performance of ViT-B/32, with a value of 79.70%. And, our method, which leverages PromptSRC, achieves the best performance with ViT-L/14, with a value of 84.20%.

5.6. Inference Stage Computational Cost

In Table 11, the compute cost analysis is performed using the SUN397 dataset over 10 epochs on a single GPU. Our method (row 3) may have a slower inference speed as a result of multiple cosine similarity calculations. The number of parameters we introduced has remained relatively consistent and stable in training. This implies that usually, our parameter quantity is lower than that of DePT.

6. Related Work

Vision-Language Models (VLMs). Recently, researchers have demonstrated strong generalization capability of Vision-Language Models (VLMs) (Alayrac et al., 2022; Shao et al., 2025), which involve training on large-scale datasets of image-text pairs. Using a large number of samples for training is indeed one of the most effective methods (Yang et al., 2021; Wang et al., 2023). Such as CLIP (Radford et al., 2021), which is a prominent and straightforward framework among existing VLMs. The strong generalization capability of CLIP has made it a foundation for many methods in adapting pre-trained VLMs for downstream tasks (Sanghi et al., 2022; Maaz et al., 2022;

	Source						Target					
	ImageNet	Caltech101	OxfordPets	StanfordCars	Flowers102	Food101	Aircraft	SUN397	DTD	EuroSAT	UCF101	Average
CoOp (Zhou et al., 2022b)	71.51	93.70	89.14	64.51	68.71	85.30	18.47	64.15	41.92	46.39	66.55	63.88
+ DePT	72.63	93.30	90.00	65.53	70.50	85.97	21.90	66.07	43.17	44.97	68.80	65.02
+ Ours	71.82	94.10	90.33	65.82	70.01	86.10	20.71	65.11	43.98	46.55	68.00	65.07
CoCoOp (Zhou et al., 2022a)	71.02	94.43	90.14	65.32	71.88	86.06	22.94	67.36	45.73	45.37	68.21	65.74
+ DePT	72.77	94.10	90.63	66.23	72.17	86.27	22.90	67.30	45.50	44.17	69.53	65.88
+ Ours	72.10	94.88	90.57	65.80	72.15	87.00	22.78	68.12	45.98	46.15	69.10	66.25
MaPLe (Khattak et al., 2023a)	70.72	93.53	90.49	65.57	72.23	86.20	24.74	67.01	46.49	48.06	68.69	66.30
+ DePT	73.27	92.53	90.10	64.60	70.10	85.57	23.63	66.40	45.03	40.13	67.53	64.56
+ Ours	71.56	94.00	90.91	65.92	73.10	86.78	24.33	68.35	46.13	49.01	68.81	66.73
PromptSRC (Khattak et al., 2023b)	71.27	93.60	90.25	65.70	70.25	86.15	23.90	67.10	46.87	45.50	68.75	65.81
+ DePT	71.60	93.80	90.13	66.00	70.93	86.27	24.30	67.23	46.60	45.83	69.10	66.02
+ Ours	71.41	94.95	92.14	66.50	73.25	87.53	25.81	68.10	49.01	49.20	69.74	67.62

Table 4. Cross-dataset generalization. Our method achieves overall favorable performance.

Table 5. Analysis of learning depth based on PromptSRC for an average of 11 datasets. HM refers to harmonic mean.

Learning Depth	1	3	5	7	9	11
HM	77.11	78.03	79.12	80.01	81.32	80.21

Table 6. Applying to other ViT instances.

	Avg. (11 datasets)							
	B/32	L/14						
HM (ViT)	79.70	81.32	84.20					

Bangalath et al., 2022; Zhang et al., 2021; Wang et al., 2022; Xu et al., 2025). To enhance the generalization ability of VLMs, researchers have explored various approaches (Zhou et al., 2025b;a). Some approaches involve enhancing the text encoder or the visual encoder (Vaswani et al., 2017). By improving the capabilities of the text encoder, the model can better capture the semantics and contextual information in the textual input (Zhang et al., 2024b). Similarly, enhancing the visual encoder allows the model to extract discriminative features from the visual input.

Prompt Tuning in VLMs. Prompt tuning (Zhang & Tian, 2025; Pan et al., 2024; Cao et al., 2025; Zhou et al., 2024; Li et al., 2024a) is a commonly employed technique in the field of Natural Language Processing (NLP) for training on downstream tasks (Yu et al., 2025; 2024b;a). Leveraging text prompts, which are instructions given to the language model component of VLMs, is a prevalent practice to improve

Table 7. Analysis of embeddings length based on PromptSRC for an average of 11 datasets. HM refers to harmonic mean.

Prompt Length	1	2	4	6	8	10
HM	77.01	80.05	81.32	80.30	79.11	78.54

Table 8. Identifying scores of Top- β and Low- β for an average of 11 datasets. We conduct analysis based on the PromptSRC (Khat-tak et al., 2023b) model. HM refers to harmonic mean.

	Avg. (11 datasets)								
	3	4	5	6					
HM (β)	78.21	79.81	81.32	80.51					

Table 9. Analysis of γ for \mathcal{L}_{CL} in pre-trained features matrix. We conduct analysis based on the PromptSRC (Khattak et al., 2023b) model for an average of 11 datasets. HM refers to harmonic mean.

	Avg. (11 datasets)							
	0.05	0.1	0.5	0.9				
HM (γ)	80.13	81.32	81.00	79.88				

task comprehension (Cui et al., 2025b). Full fine-tuning and linear probes are two commonly employed approaches for adapting VLMs to downstream tasks (Fang et al., 2025; Liu et al., 2025). The constraints of both methods have prompted research into innovative techniques influenced by prompt tuning within the realm of VLMs (Yin et al., 2025). CoOp (Zhou et al., 2022b), ProDA (Lu et al., 2022), and

Dataset	Classes	Train	Val	Test
ImageNet (Deng et al., 2009)	1,000	1.28 M	N/A	50,000
Caltech101 (Fei-Fei et al., 2004)	100	4,128	1,649	2,465
OxfordPets (Parkhi et al., 2012)	37	2,944	736	3,669
StanfordCars (Krause et al., 2013)	196	6,509	1,635	8,041
Flowers102 (Nilsback & Zisserman, 2008)	102	4,093	1,633	2,463
Food101 (Bossard et al., 2014)	101	50,500	20,200	30,300
FGVCAircraft (Maji et al., 2013)	100	3,334	3,333	3,333
SUN397 (Xiao et al., 2010)	397	15,880	3,970	19,850
DTD (Cimpoi et al., 2014)	47	2,820	1,128	1,692
EuroSAT (Helber et al., 2019)	10	13,500	5,400	8,100
UCF101 (Soomro et al., 2012)	101	7,639	1,898	3,783
-V2 (Recht et al., 2019)	1,000	N/A	N/A	10,000
-Sketch (Wang et al., 2019)	1,000	N/A	N/A	50,889
-A (Hendrycks et al., 2021b)	200	N/A	N/A	7,500
-R (Hendrycks et al., 2021a)	200	N/A	N/A	30,000

Table 10. Training and testing datasets. The dataset consists of 11 image classification datasets and four variant datasets of ImageNet.

Table 11. The compute cost analysis is performed using the SUN397 (Xiao et al., 2010) dataset over 10 epochs on a single GPU. 'N': the num of classes in the base task (Zhang et al., 2024a).

Method	Train time	Learnable para.	HM
CoOp	10.88min	8K	71.65
+ DePT	10.91min	+ (2+N/2)K	77.30
+ Ours	13.56min	+ 0K	78.66

CoCoOp (Zhou et al., 2022a) fine-tune the CLIP model for few-shot image recognition by optimizing a continuous set of embeddings within the textual branch. The imageconditional prompt utilized in CoCoOp significantly contributes to improving generalization to unseen classes (Peng et al., 2025; 2023). By conditioning prompts on visual features, CoCoOp (Zhou et al., 2022a) ensures that the language model focuses on pertinent visual information when making predictions. Moreover, some approaches (Yao et al., 2023b;a; Cui et al., 2024; 2025a) constrain the learnable prompts to contain the essential general knowledge. In addition, the approach (Zhang et al., 2024a) introduces a flexible approach to align its vision and language representations.

7. Future Works and Limitations

In the future, we plan to investigate the potential of our method in other tasks (Tian et al., 2024; Yang et al., 2023; Li et al., 2019; Yu et al., 2023; Li et al., 2024c; Zhang et al., 2022; Wang et al., 2021; Li et al., 2025a; Jia et al., 2025; Liu et al., 2023b; Li et al., 2024b) and downstream scenarios (Pan et al., 2025; Jin et al., 2025; Li et al., 2025b; Liu et al., 2024; Wu et al., 2024; Liu et al., 2023c; d). Due to multiple processes of our method, the training speed is slower. And, the ImageNet training in Table 4 is lower.

8. Conclusion

Prompt learning is a promising method for adapting pretrained Visual-Language Models (VLMs) for target-specific classification tasks. However, these optimization frameworks exhibit limited efficacy when applied to targetunspecific or generalizable scenarios. This performance degradation may stem from overfitting-induced catastrophic forgetting, where the model loses generalizable knowledge critical for adapting to unseen tasks during the fine-tuning process. In this paper, we propose a Features Matrix (FM) for vision-language models. Our method incorporates multiple hand-crafted prompts to extract general information as a pre-trained features matrix from CLIP to enhance generalization. The matrix of pre-trained features can delve into the semantics of different hand-crafted prompts at a more profound level, which reduces the risk of forgetting the essential general knowledge of pre-trained CLIP. Our method focuses on the challenge of various target-unspecific tasks and scenarios across a wide range of real datasets. Importantly, our method is compatible with representative prompt learning frameworks for textual or multi-modal prompts as a generic and flexible module.

Acknowledgement

XM T was supported in part by the NSFC No. 62222117. J Y was supported in part by the NSFC No. 62125201 and U24B20174.

Impact Statement

The societal implications of our work involve democratizing the availability of potent AI resources, as our plug-and-play method can achieve impressive performance even in the absence of extensive labeled data.

References

- Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022.
- Bangalath, H., Maaz, M., Khattak, M. U., Khan, S. H., and Shahbaz Khan, F. Bridging the gap between object and image-level representations for open-vocabulary detection. Advances in Neural Information Processing Systems, 35:33781–33794, 2022.
- Bossard, L., Guillaumin, M., and Van Gool, L. 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 13, pp. 446–461. Springer, 2014.
- Cao, Q., Chen, Y., Lu, L., Sun, H., Zeng, Z., Yang, X., and Zhang, D. Generalized domain prompt learning for accessible scientific vision-language models. *Nexus*, 2(2), 2025.
- Chen, G., Yao, W., Song, X., Li, X., Rao, Y., and Zhang, K. Plot: Prompt learning with optimal transport for visionlanguage models. In *The Eleventh International Conference on Learning Representations*.
- Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3606–3613, 2014.
- Cui, F., Yang, X., Wu, C., Xiao, L., and Tian, X. Advancing prompt learning through an external layer. In *Proceedings of the 32nd ACM International Conference* on Multimedia, pp. 807–816, 2024.
- Cui, F., Fong, J., Zeng, R., Tian, X., and Yu, J. A similarity paradigm through textual regularization without forgetting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 16100–16108, 2025a.
- Cui, F., Zhang, Y., Wang, X., Wang, X., and Xiao, L. Generalizable prompt learning of clip: A brief overview. arXiv preprint arXiv:2503.01263, 2025b.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pp. 248–255. Ieee, 2009.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020.

- Etchegaray, D., Huang, Z., Harada, T., and Luo, Y. Find n'propagate: Open-vocabulary 3d object detection in urban environments. *arXiv preprint arXiv:2403.13556*, 2024.
- Fang, Z., Lu, J., and Zhang, G. An extremely simple algorithm for source domain reconstruction. *IEEE Transactions on Cybernetics*, 54(3):1921–1933, 2023.
- Fang, Z., Yu, H., Xu, C., Li, Z., Ying, J., and Xie, S. Contrastive zero-shot relational learning for knowledge graph completion. *Knowledge-Based Systems*, pp. 113425, 2025.
- Fei-Fei, L., Fergus, R., and Perona, P. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In 2004 conference on computer vision and pattern recognition workshop, pp. 178–178. IEEE, 2004.
- Greer, R., Antoniussen, B., Andersen, M. V., Møgelmose, A., and Trivedi, M. M. The why, when, and how to use active learning in large-data-driven 3d object detection for safe autonomous driving: An empirical exploration. *arXiv preprint arXiv:2401.16634*, 2024.
- Helber, P., Bischke, B., Dengel, A., and Borth, D. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019.
- Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8340–8349, 2021a.
- Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., and Song, D. Natural adversarial examples. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 15262–15271, 2021b.
- Jia, M., Tang, L., Chen, B.-C., Cardie, C., Belongie, S., Hariharan, B., and Lim, S.-N. Visual prompt tuning. In *European Conference on Computer Vision*, pp. 709–727. Springer, 2022.
- Jia, Y., Wang, Z., Han, Z., Li, J., Zhang, M., Lao, Z., Han, Y., Gao, R., Gao, J., Zheng, Z., et al. Variable and intelligent catalyst design based on local chemical environments in sulfur redox reactions. *Joule*, 9(5), 2025.
- Jin, K., Liu, Z., Wang, J., and Wang, H. Unmanned surface vehicle navigation under disturbances: World model enhanced reinforcement learning. *IEEE/ASME Transactions on Mechatronics*, 2025.

- Khattak, M. U., Rasheed, H., Maaz, M., Khan, S., and Khan, F. S. Maple: Multi-modal prompt learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 19113–19122, 2023a.
- Khattak, M. U., Wasim, S. T., Naseer, M., Khan, S., Yang, M.-H., and Khan, F. S. Self-regulating prompts: Foundational model adaptation without forgetting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15190–15200, 2023b.
- Krause, J., Stark, M., Deng, J., and Fei-Fei, L. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pp. 554–561, 2013.
- Li, H., Huang, B., Chuai, M., Zheng, Z., Chen, H., Piao, Z., Zhou, G., and Fan, H. J. Dual-plating aqueous zn–iodine batteries enabled via halogen-complexation chemistry for large-scale energy storage. *Energy & Environmental Science*, 18(7):3160–3168, 2025a.
- Li, H., Mai, Y., Gao, M., He, J., Liu, Z., and Wang, H. Large-scale lidar-based loop closing via combination of equivariance and invariance on se (3). *IEEE/ASME Transactions on Mechatronics*, 2025b.
- Li, H., Zhang, M., Zheng, Z., Wu, X., Xiao, X., Piao, Z., Li, C., Jia, Y., Yang, J., and Zhou, G. High-capacity and long-life aqueous zn-span batteries with tandem catalysis. *Advanced Materials*, 37(5):2409771, 2025c.
- Li, J., Zhang, Y., Liu, X., Liu, P., Yin, H., and Liu, D. C. A robust automatic ultrasound spectral envelope estimation. *Information*, 10(6):199, 2019.
- Li, J., Li, H., Erfani, S. M., Feng, L., Bailey, J., and Liu, F. Visual-text cross alignment: Refining the similarity score in vision-language models. In *International Conference* on Machine Learning, pp. 28018–28039. PMLR, 2024a.
- Li, Q., Hu, Z., Wang, Y., Li, L., Fan, Y., King, I., Jia, G., Wang, S., Song, L., and Li, Y. Progress and opportunities of foundation models in bioinformatics. *Briefings in Bioinformatics*, 25(6):bbae548, 2024b.
- Li, W., Guo, Y., Zheng, J., Lin, H., Ma, C., Fang, L., and Yang, X. Sparseformer: Detecting objects in hrw shots via sparse vision transformer. In *Proceedings of the 32nd* ACM International Conference on Multimedia, pp. 4851– 4860, 2024c.
- Liu, J., Zhuo, D., Feng, Z., Zhu, S., Peng, C., Liu, Z., and Wang, H. Dvlo: Deep visual-lidar odometry with local-to-global feature fusion and bi-directional structure alignment. In *European Conference on Computer Vision*, pp. 475–493. Springer, 2024.

- Liu, S., Chen, Z., Wu, M., Liu, C., and Chen, L. Wisr: Wireless domain generalization based on style randomization. *IEEE Transactions on Mobile Computing*, 23(5): 4520–4532, 2023a.
- Liu, S., Chen, Z., Wu, M., Wang, H., Xing, B., and Chen, L. Generalizing wireless cross-multiple-factor gesture recognition to unseen domains. *IEEE Transactions on Mobile Computing*, 23(5):5083–5096, 2023b.
- Liu, Z., Li, K., Hao, T., and Wang, H. Visual servoing of rigid-link flexible-joint manipulators in the presence of unknown camera parameters and boundary output. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 53(8):5096–5105, 2023c.
- Liu, Z., Zhai, Y., Li, J., Wang, G., Miao, Y., and Wang, H. Graph relational reinforcement learning for mobile robot navigation in large-scale crowded environments. *IEEE Transactions on Intelligent Transportation Systems*, 24 (8):8776–8787, 2023d.
- Liu, Z., Yu, H., and Luo, X. Federated graph anomaly detection via disentangled representation learning. In *Proceedings of the ACM on Web Conference 2025*, pp. 1216–1224, 2025.
- Lu, Y., Liu, J., Zhang, Y., Liu, Y., and Tian, X. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5206–5215, 2022.
- Maaz, M., Rasheed, H., Khan, S., Khan, F. S., Anwer, R. M., and Yang, M.-H. Class-agnostic object detection with multi-modal transformer. In *European Conference on Computer Vision*, pp. 512–531. Springer, 2022.
- Maji, S., Rahtu, E., Kannala, J., Blaschko, M., and Vedaldi, A. Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151, 2013.
- Nilsback, M.-E. and Zisserman, A. Automated flower classification over a large number of classes. In 2008 Sixth Indian conference on computer vision, graphics & image processing, pp. 722–729. IEEE, 2008.
- Pan, B., Li, Q., Tang, X., Huang, W., Fang, Z., Liu, F., Wang, J., Yu, J., and Shi, Y. Nlprompt: Noise-label prompt learning for vision-language models. *arXiv preprint* arXiv:2412.01256, 2024.
- Pan, Y., Xu, Y., Liu, Z., and Wang, H. Planning from imagination: Episodic simulation and episodic memory for vision-and-language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 6345–6353, 2025.

- Parkhi, O. M., Vedaldi, A., Zisserman, A., and Jawahar, C. Cats and dogs. In 2012 IEEE conference on computer vision and pattern recognition, pp. 3498–3505. IEEE, 2012.
- Peng, Z., Tian, Q., Xu, J., Jin, Y., Lu, X., Tan, X., Xie, Y., and Ma, L. Generalized category discovery in semantic segmentation. arXiv preprint arXiv:2311.11525, 2023.
- Peng, Z., Ma, J., Sun, Z., Yi, R., Song, H., Tan, X., and Ma, L. Mos: Modeling object-scene associations in generalized category discovery. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 15118– 15128, 2025.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*, pp. 5389–5400. PMLR, 2019.
- Sanghi, A., Chu, H., Lambourne, J. G., Wang, Y., Cheng, C.-Y., Fumero, M., and Malekshan, K. R. Clip-forge: Towards zero-shot text-to-shape generation. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 18603–18613, 2022.
- Shao, Z., Wang, M., Yu, Z., Pan, W., Yang, Y., Wei, T., Zhang, H., Mao, N., Chen, W., and Yu, J. Growing a twig to accelerate large vision-language models. *arXiv* preprint arXiv:2503.14075, 2025.
- Soomro, K., Zamir, A. R., and Shah, M. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- Tian, F., Li, Y., Yan, Y., Guan, S., Ge, Y., and Yang, X. Postedit: Posterior sampling for efficient zero-shot image editing. arXiv preprint arXiv:2410.04844, 2024.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information* processing systems, 30, 2017.
- Wang, H., Ge, S., Lipton, Z., and Xing, E. P. Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32, 2019.
- Wang, Y., Liu, C., Li, S., Wang, T., Lin, W., and Yu, X. A multi-target tracking algorithm for fast-moving workpieces based on event camera. In *IECON 2021–47th*

Annual Conference of the IEEE Industrial Electronics Society, pp. 1–5. IEEE, 2021.

- Wang, Y., Yan, Y., Shi, D., Zhu, W., Xia, J., Jeff, T., Jin, S., Gao, K., Li, X., and Yang, X. Nerf-ibvs: visual servo based on nerf for visual localization and navigation. *Advances in Neural Information Processing Systems*, 36: 8292–8304, 2023.
- Wang, Z., Xu, Q., Yang, Z., He, Y., Cao, X., and Huang, Q. Openauc: Towards auc-oriented open-set recognition. *Advances in Neural Information Processing Systems*, 35: 25033–25045, 2022.
- Wu, W., Wang, Q., Wang, G., Wang, J., Zhao, T., Liu, Y., Gao, D., Liu, Z., and Wang, H. Emie-map: Large-scale road surface reconstruction based on explicit mesh and implicit encoding. In *European Conference on Computer Vision*, pp. 370–386. Springer, 2024.
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. Sun database: Large-scale scene recognition from abbey to zoo. In 2010 IEEE computer society conference on computer vision and pattern recognition, pp. 3485–3492. IEEE, 2010.
- Xu, Y., Pan, Y., Liu, Z., and Wang, H. Flame: Learning to navigate with multimodal llm in urban environments. In *Proceedings of the AAAI Conference on Artificial Intelli*gence, volume 39, pp. 9005–9013, 2025.
- Yang, Z., Xu, Q., Bao, S., Cao, X., and Huang, Q. Learning with multiclass auc: Theory and algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44 (11):7747–7763, 2021.
- Yang, Z., Xu, Q., Bao, S., He, Y., Cao, X., and Huang, Q. Optimizing two-way partial auc with an end-to-end framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):10228–10246, 2023. doi: 10.1109/TPAMI.2022.3185311.
- Yao, H., Zhang, R., and Xu, C. Visual-language prompt tuning with knowledge-guided context optimization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 6757–6767, June 2023a.
- Yao, H., Zhang, R., and Xu, C. Visual-language prompt tuning with knowledge-guided context optimization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 6757–6767, 2023b.
- Yin, H., Ren, Y., Yan, K., Ding, S., and Hao, Y. Rod-mllm: Towards more reliable object detection in multimodal large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 14358– 14368, 2025.

- Yu, H., Liu, W., Zhu, N., Li, P., and Luo, X. In-gfd: An interpretable graph fraud detection model for spam reviews. *IEEE Transactions on Artificial Intelligence*, 2024a.
- Yu, H., Wen, J., Sun, Y., Wei, X., and Lu, J. Ca-gnn: A competence-aware graph neural network for semisupervised learning on streaming data. *IEEE Transactions* on Cybernetics, 2024b.
- Yu, H., Wen, J., and Zheng, Z. Camel: Cross-modality adaptive meta-learning for text-based person retrieval. *IEEE Transactions on Information Forensics and Security*, 2025.
- Yu, Z., Zheng, L., Zhao, Z., Wu, F., Fan, J., Ren, K., and Yu, J. Anetqa: a large-scale benchmark for fine-grained compositional reasoning over untrimmed videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision* and Pattern Recognition, pp. 23191–23200, 2023.
- Zhang, J., Wu, S., Gao, L., Shen, H. T., and Song, J. Dept: Decoupled prompt tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12924–12933, 2024a.
- Zhang, R., Fang, R., Zhang, W., Gao, P., Li, K., Dai, J., Qiao, Y., and Li, H. Tip-adapter: Training-free clip-adapter for better vision-language modeling. *arXiv preprint arXiv:2111.03930*, 2021.
- Zhang, W., Lyu, S., Xue, F., Yao, C., Zhu, Z., and Jia, Z. Predict the rover mobility over soft terrain using articulated wheeled bevameter. *IEEE Robotics and Automation Letters*, 7(4):12062–12069, 2022.
- Zhang, W., Wu, L., Zhang, Z., Yu, T., Ma, C., Jin, X., Yang, X., and Zeng, W. Unleash the power of visionlanguage models by visual attention prompt and multimodal interaction. *IEEE Transactions on Multimedia*, 2024b.
- Zhang, Y. and Tian, X. Consistent prompt learning for vision-language models. *Knowledge-Based Systems*, 310: 112974, 2025.
- Zhou, C., Han, S., Yang, L., Zhou, Y., Cheng, X., Wang, Y., and Li, H. Rankadaptor: Hierarchical rank allocation for efficient fine-tuning pruned llms via performance model. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 5781–5795, 2025a.
- Zhou, C., Zhou, Y., Zhang, S., Wang, Y., and Liu, Z. Dynamic operator optimization for efficient multi-tenant lora model serving. In *Proceedings of the AAAI Conference* on Artificial Intelligence, volume 39, pp. 22910–22918, 2025b.

- Zhou, K., Yang, J., Loy, C. C., and Liu, Z. Conditional prompt learning for vision-language models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16816–16825, 2022a.
- Zhou, K., Yang, J., Loy, C. C., and Liu, Z. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022b.
- Zhou, Y., Xia, X., Lin, Z., Han, B., and Liu, T. Few-shot adversarial prompt learning on vision-language models. *Advances in Neural Information Processing Systems*, 37: 3122–3156, 2024.
- Zhu, B., Niu, Y., Han, Y., Wu, Y., and Zhang, H. Promptaligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15659–15669, 2023.