

A Similarity Paradigm Through Textual Regularization Without Forgetting

Fangming Cui¹, Jan Fong², Rongfei Zeng³, Xinmei Tian⁴, Jun Yu^{5*}

¹Shanghai Jiao Tong University

²Hong Kong Baptist University

³Northeastern University

⁴University of Science and Technology of China

⁵Harbin Institute of Technology (Shenzhen)

cuifangming@sjtu.edu.cn, zengrf@swc.neu.edu.cn, xinmei@ustc.edu.cn, yujun@hit.edu.cn

Abstract

Prompt learning has emerged as a promising method for adapting pre-trained visual-language models (VLMs) to a range of downstream tasks. While optimizing the context can be effective for improving performance on specific tasks, it can often lead to poor generalization performance on unseen classes or datasets sampled from different distributions. It may be attributed to the fact that textual prompts tend to overfit downstream data distributions, leading to the forgetting of generalized knowledge derived from hand-crafted prompts. In this paper, we propose a novel method called Similarity Paradigm with Textual Regularization (SPTR) for prompt learning without forgetting. SPTR is a two-pronged design based on hand-crafted prompts that is an inseparable framework. 1) To avoid forgetting general textual knowledge, we introduce the optimal transport as a textual regularization to finely ensure approximation with hand-crafted features and tuning textual features. 2) In order to continuously unleash the general ability of multiple hand-crafted prompts, we propose a similarity paradigm for natural alignment score and adversarial alignment score to improve model robustness for generalization. Both modules share a common objective in addressing generalization issues, aiming to maximize the generalization capability derived from multiple hand-crafted prompts. Four representative tasks (i.e., non-generalization few-shot learning, base-to-novel generalization, cross-dataset generalization, domain generalization) across 11 datasets demonstrate that SPTR outperforms existing prompt learning methods.

Introduction

Large vision-language models (VLMs) like pre-trained CLIP (Radford et al. 2021) has demonstrated remarkable generalization capabilities across a wide range of downstream tasks (Bangalath et al. 2022; Cao et al. 2023b; Zhang et al. 2021a). By learning to associate textual descriptions with corresponding visual content, these models can perform tasks such as image classification, object detection, image generation, and more (Cao et al. 2024a, 2023a; Li et al. 2024b; Liu et al. 2024a, 2023, 2024b; Shi et al. 2024; Xie et al. 2023; Zhang et al. 2022).

The generalization capabilities of these VLMs have been validated through their impressive performance on various

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Method	4 tasks (Avg.)	
CLIP (ICML2021)	67.69	Δ
CoOp (IJCV2022)	68.31	+0.62
CoCoOp (CVPR2022)	68.63	+0.94
KgCoOp (CVPR2023)	70.71	+3.02
PLOT (ICLR2023)	71.24	+3.55
MaPLe (CVPR2023)	71.58	+3.89
PromptSRC (ICCV2023)	72.09	+4.40
TCP (CVPR2024)	71.79	+4.10
Ours (AAAI2025)	72.89	+5.20

Table 1: We test our method on 4 tasks: non-generalization few-shot learning, base-to-novel generalization, domain generalization, and cross-dataset evaluation. Ours (AAAI2025) overall shows competitive results for average performance compared to the previous prompting methods. The symbol Δ represents the value that increases in comparison to CLIP.

benchmarks and tasks. They have shown strong performance in both seen and unseen classes, indicating their ability to generalize well to novel concepts and data distributions. CLIP utilizes pre-defined text inputs or prompts during inference to generate classification weights. These prompts can include hand-engineered text, such as “a photo of a [class]”, which guides the text encoder of CLIP. As a zero-shot model, CLIP does not require task-specific fine-tuning on specific downstream tasks. This makes it highly versatile and capable of performing reasonably well on a wide range of tasks without any fine-tuning. It can leverage the alignment between two modalities learned during pre-training to make predictions and understand the semantic relationships between them.

Advanced techniques have been proposed to introduce learnable parameters for few-shot textual prompting to enhance the performance of pre-trained CLIP in downstream classification tasks and automate prompts. One significant method in this area is CoOp, introduced in the seminal work (Zhou et al. 2022b). CoOp leverages the differentiable nature of neural networks to transform context words in a prompt into a set of learnable embeddings. By learning these textual embeddings of prompts, CoOp achieves substantial improvements over manually tuned prompts on various image recognition datasets, even with limited labeled images (few shots) available for training.

Although textual embedding tuning methods perform well on classification tasks for seen classes, they may not generalize effectively to unseen classes. As shown in Table 2 (left-top), the average performance of these textual embedding methods (Chen et al. 2022; Yao, Zhang, and Xu 2023b; Zhou et al. 2022a,b) on unseen classes is lower than that of hand-crafted CLIP (74.22%). This suggests that textual embeddings may not adequately capture the semantic relationships and visual characteristics specific to the unseen classes. One possible reason for the degradation in performance for unseen classes is that textual prompts tend to overfit downstream data distributions (Zhou et al. 2022a). This overfitting can lead to the forgetting of hand-crafted general textual knowledge, which possesses strong generalization abilities.

To address the degradation (Lu et al. 2022; Zhu et al. 2023) in performance for unseen classes, we propose a novel method called Similarity Paradigm with Textual Regularization (SPTR) for prompt learning. SPTR is a two-pronged design based on multiple hand-crafted prompts that is an inseparable prompting framework. We attempt to awaken and maintain pre-trained general knowledge, as general knowledge has excellent generalization ability. We use multiple hand-crafted prompts to create multiple pre-trained text features, fully leveraging the generalized capabilities of the CLIP. To bring the tuning features closer to the multiple pre-trained features, we apply the optimal transport as a textual regularization to finely ensure approximation with hand-crafted features of pre-trained CLIP and tuning textual features that avoid forgetting essential general textual knowledge. The optimal transport can finely calculate the distance between two distributions under the form of multiple sampling, which achieves fine-grained matching across two distributions. In order to continuously unleash the general ability of multiple hand-crafted prompts, we propose a similarity paradigm for natural alignment score and adversarial alignment score to improve model robustness for generalization.

As shown in Table 1, SPTR shows competitive results for average performance compared to the representative prompting methods, it demonstrates the robustness of SPTR in different types of classification tasks. Our main contributions can be summarized as follows:

- To avoid forgetting essential general textual knowledge, we propose to introduce the optimal transport plan as a textual regularization to finely ensure approximation with multiple hand-crafted features and tuning textual features.
- Based on multiple hand-crafted prompts, we propose a novel similarity paradigm for natural alignment score and adversarial alignment score to improve generalization.
- Extensive experiments show that SPTR performs favorably well on 4 types of representative tasks across 11 datasets compared to the existing prompt learning methods, achieving state-of-the-art performance.

Methodology

Preliminaries

Adversarial Training. Adversarial attacks (Tang et al. 2022; Zhang et al. 2021b, 2020) refer to the creation of maliciously crafted inputs, called adversarial examples, that are

designed to deceive machine learning models (Jia et al. 2020, 2022; Miyato, Dai, and Goodfellow 2016). A common method (Madry et al. 2017) with crafting adversarial examples for VLMs involves seeking a perturbation \mathbf{r} for an input natural image \mathbf{x} with label y to maximize the dissimilarity, typically cosine dissimilarity. To optimize (minimize) the model after the maximum disturbance, the design of adversarial training for vision-language cross-entropy loss can be formalized as follows:

$$\min_{\theta} \max_{\|\mathbf{r}\|_p \leq \epsilon} \mathcal{L}_{\text{CE}}(\mathbf{x} + \mathbf{r}, y; \theta), \quad (1)$$

where $\|\cdot\|_p$ denotes the ℓ_p -norm, θ denotes the tuning parameters, and ϵ controls the strength of adversarial perturbations. This design involves training models on adversarial examples generated to maximize prediction errors.

Pre-Trained CLIP. Our method is based on a pre-trained vision-language model, hand-crafted CLIP, which is a zero-shot learning method (Li et al. 2024c; Wortsman et al. 2022). The pre-trained CLIP has two encoders: a frozen text encoder and a frozen image encoder, which separately map a hand-crafted textual input p and training images \mathbf{x} into a shared feature space through transformer blocks. In the pre-training stage of CLIP, text encoder $\mathcal{F}_t(\cdot)$ and image encoder $\mathcal{F}_v(\cdot)$ undergo simultaneous training on extensive datasets containing text-image pairs. To achieve this, a contrastive loss function \mathcal{L}_{CE} is utilized to maximize the cosine similarity between text-image pairs while minimizing the cosine similarity between unmatched pairs within the feature space:

$$\mathcal{L}_{\text{CE}} = -\log \frac{\exp\{\text{sim}(\mathcal{F}_t(p_y), \mathcal{F}_v(\mathbf{x})) / \tau\}}{\sum_{p_i \in \mathcal{P}} \exp\{\text{sim}(\mathcal{F}_t(p_i), \mathcal{F}_v(\mathbf{x})) / \tau\}}, \quad (2)$$

where $y \in \mathcal{Y}$ is the label of $\mathbf{x} \in \mathcal{X}$, $\mathcal{P} = \{p_i\}_{i=1}^K$ presents the set of K hand-crafted prompts, $\text{sim}(\cdot, \cdot)$ stands for cosine similarity between textual features and visual features, and τ denotes a temperature parameter. To this end, the classifier consists of K textual features derived from hand-crafted prompts $\mathcal{P} = \{p_i\}_{i=1}^K$, where the prompt p_i for the i -th class may have the form of “a picture of a”, “a bad photo of a”, “a photo of many”, etc (Radford et al. 2021).

Textual Prompting. Recent studies suggest that fine-tuning prompts for specific images may outperform the use of manually crafted prompts. More precisely, the class name is preserved as prior knowledge to guarantee that the acquired prompts can construct a classifier, whereas the word (also known as context) embeddings of prompts are treated as adjustable parameters. The learnable words in the above prompt are initialized using “a photo of a [classname]”. In the inference phase, the prompts equipped with the acquired context can generate textual features for classification purposes.

Textual Regularization. The textual regularization techniques (Khattak et al. 2023b; Yao, Zhang, and Xu 2023a,b) refer to the regularization of pre-trained and fine-tuned features limited to the text branch through metric function to prevent forgetting essential general knowledge. The key concept for this component is to awaken and sustain pre-trained general knowledge (Akinwande et al. 2023), given its exceptional generalization capabilities, which address the degradation in performance for unseen classes.

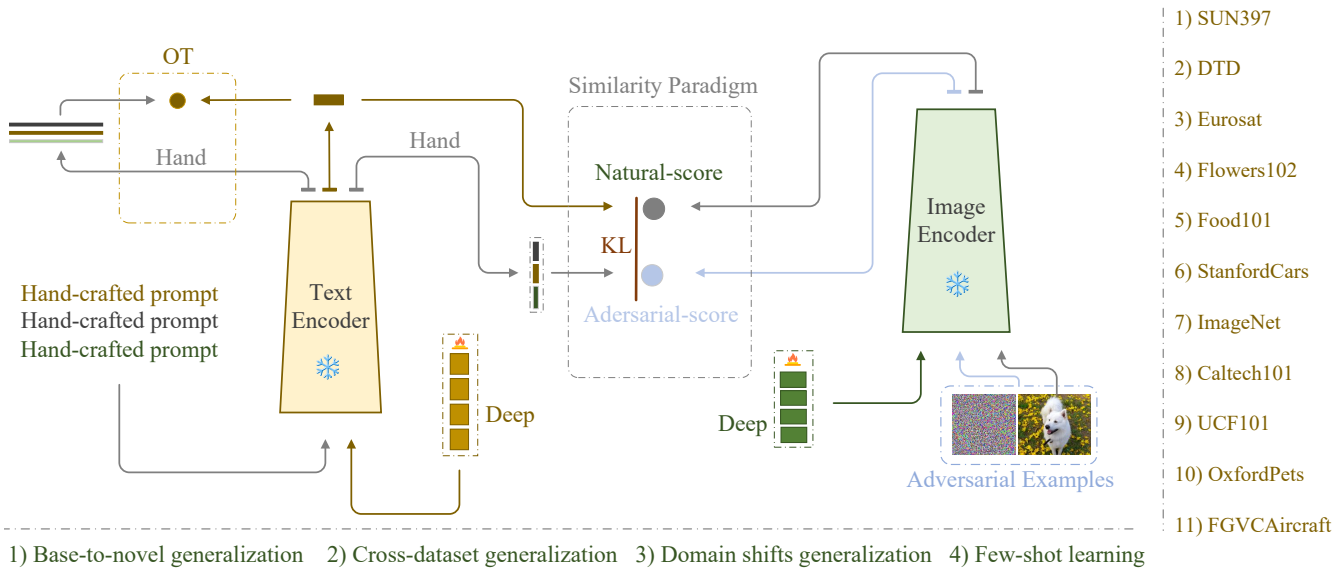


Figure 1: We propose a novel similarity paradigm with a textual regularization of the OT for prompting called SPTR. SPTR performs favorably well on 4 types of representative tasks across 11 datasets, achieving state-of-the-art performance.

For the first time, we employ optimal transport (OT) (Vilani et al. 2009) as a textual regularization, thereby preventing the loss of essential general textual knowledge. The optimal transport facilitates the computation of distances between two distributions (multiple pre-trained textual features and single-tuning textual features). Unlike traditional distance metrics (Yao, Zhang, and Xu 2023b, 2024), such as L1 and MSE, the optimal transport adopts a detailed matching strategy with multiple hand-crafted prompts to measure distances between distributions, which is different from single hand-crafted prompts (Yao, Zhang, and Xu 2023b, 2024). Specifically, let T_{tun}^{sets} denote the tuning textual feature sets, and let T^{sets} denote the pre-trained hand-crafted textual feature sets. Further, we can obtain M tuning local textual feature belongs to T_{tun}^{sets} , and N pre-trained local textual features with N hand-crafted prompts belong to T^{sets} . Note that $M = 1$ because the learnable textual prompt is a single prompt, which is different from PLOT with multiple learnable prompts. Following the OT algorithm, our method learns the plan by minimizing the following distance Dis to push T^{sets} to T_{tun}^{sets} for fine-grained alignment. To this end, we can obtain the textual regularization distance Dis .

Proposed Similarity Paradigm

In order to continuously unleash the general ability of multiple hand-crafted prompts, we propose a similarity paradigm for natural alignment score and adversarial alignment score to improve model robustness for generalization. When training models for vision and language tasks, it is crucial to consider both natural and adversarial scenarios. By narrowing the distribution gap between natural text-image similarity and adversarial text-image similarity, the model is encouraged to make more consistent and robust predictions in the presence of adversarial inputs.

Inspired by distributional robust optimization, we propose to perform distributional exploration for the image distribution, which would identify the worst-case distribution for tuning prompts. We optimize the learnable visual prompt \mathcal{P}_v and learnable textual prompt \mathcal{P}_t using the adversarial images as the identified worst-case distribution for robustness improvement. The adversarial examples are obtained by adding adversarial perturbations \mathbf{r} to natural images \mathbf{x} . Following previous adversarial attack methods, we apply the projected gradient descent (PGD) algorithm (Madry et al. 2017) to search for adversarial perturbation \mathbf{r} iteratively:

$$\mathbf{r}^{(t+1)} = \prod_{\mathcal{C}(\mathbf{r}, \epsilon, p)} (\mathbf{r}^t + \mu \delta_t^t), \quad (3)$$

where \prod designates a projection function, $\mathcal{C}(\mathbf{r}, \epsilon, p)$ stands for the ℓ_p -norm ball centered at the embedding \mathbf{r} through a radius ϵ , μ is the step size, δ^t represents the update direction at the t -th iteration. To solve the constrained optimization problem, we can derive the update direction as follows:

$$\delta^t = \epsilon \text{sign}(\nabla_{\mathbf{x}} \mathcal{L}_{CE}(\mathbf{x}, y, \mathbf{r})) \frac{|\nabla_{\mathbf{x}} \mathcal{L}_{CE}(\mathbf{x}, y, \mathbf{r})|^{q-1}}{(\|\nabla_{\mathbf{x}} \mathcal{L}_{CE}(\mathbf{x}, y, \mathbf{r})\|_q^{1/p})^{q-1}}, \quad (4)$$

where $\text{sign}(\cdot)$ designates the sign function, $\mathcal{L}_{CE}(\cdot)$ designates the cross-entropy loss, q is the dual of p .

To this end, we propose a similarity paradigm for minimizing the Kullback-Leibler (KL) divergence between natural Vision-Language (V-L) alignment score (cosine similarity) and adversarial V-L alignment score, as shown in Figure 1. The similarity paradigm aims to mitigate the potential vulnerability to robustness issues stemming from disparities in generalizing natural samples. Specifically, let \mathbf{v}_n denote the output visual features of natural images, \mathbf{v}_{adv} denote the output visual features of adversarial examples, \mathbf{t}_{tun} denote tuning textual feature. Further, we let \mathbf{t}_{hand} denote the

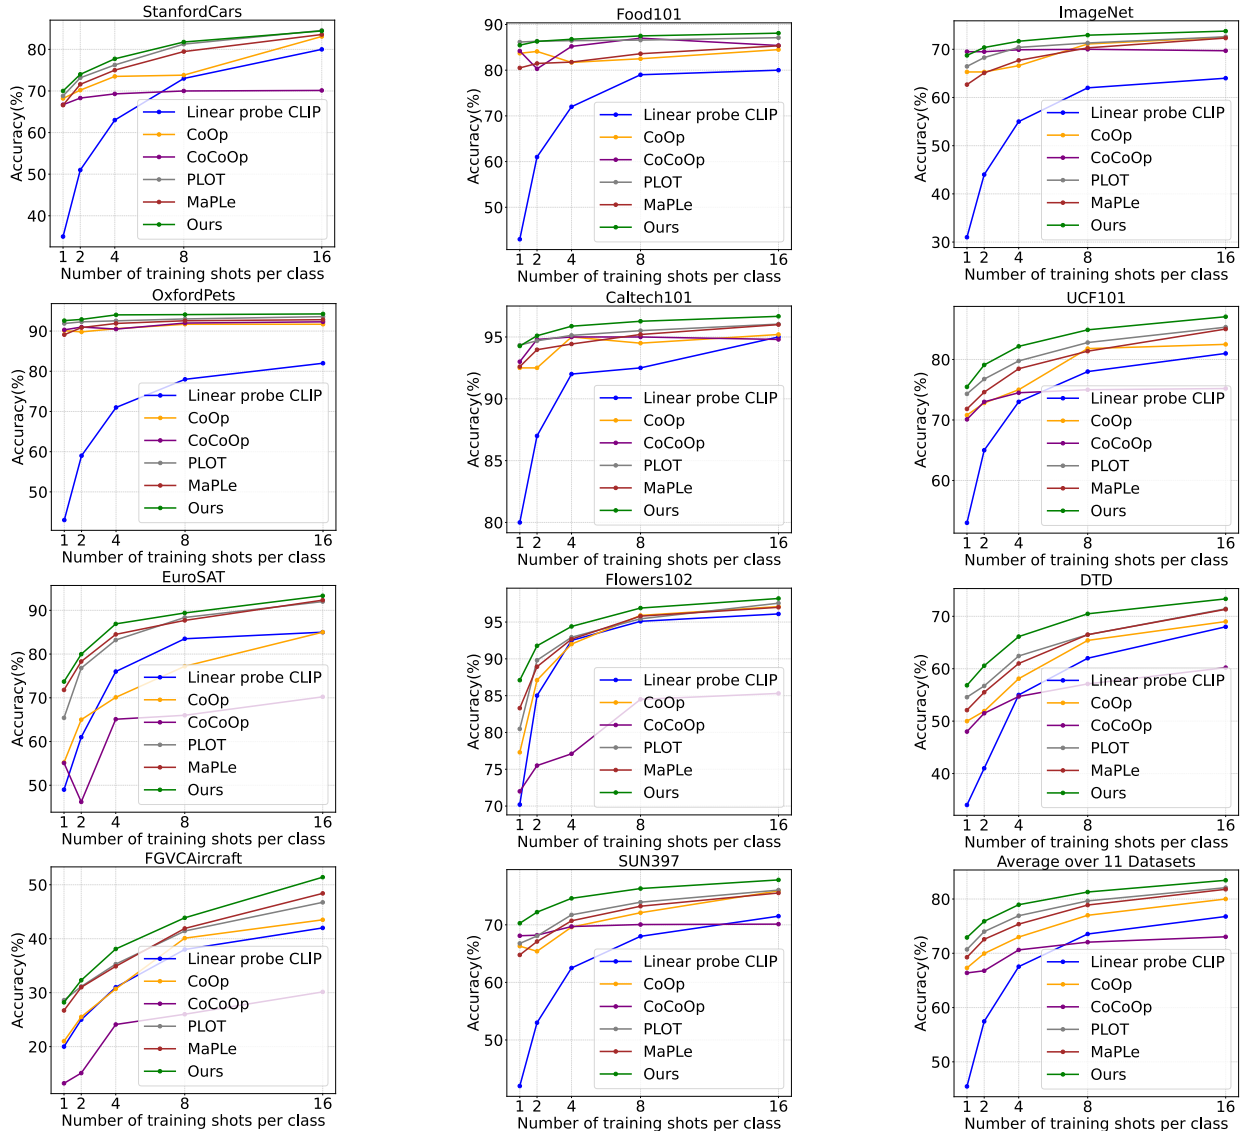


Figure 2: Non-generalization few-shot learning for all shots ($K = 1, 2, 4, 8, 16$).

compressed hand-crafted feature, which is the average of N templates hand-crafted features (Radford et al. 2021), leading to learning more diverse generalized knowledge from the frozen CLIP. The $\text{sim}(\cdot, \cdot)$ stands for cosine score (logit) between visual features and textual features. The objective of the Similarity Paradigm (SP) can be stated as:

$$\mathcal{L}_{\text{SP}} = \mathcal{KL}(\text{sim}(\mathbf{t}_{\text{tun}}, \mathbf{v}_n), \text{sim}(\mathbf{t}_{\text{hand}}, \mathbf{v}_{\text{adv}})). \quad (5)$$

Accordingly, the final objective $\mathcal{L}_{\text{total}}$ of SPTR with a textual hyper-parameter α can be formulated as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CE}} + \mathcal{L}_{\text{SP}} + \alpha \text{Dis}, \quad (6)$$

where \mathcal{L}_{CE} represents two-modalities alignment of cross-entropy loss. The Dis avoids forgetting essential general textual knowledge of hand-crafted CLIP having a strong generalization ability. Based on multiple hand-crafted prompts,

\mathcal{L}_{SP} as a KL constraint for natural alignment score and adversarial alignment score that mitigates the potential failure in robustness caused by discrepancies in natural generalization.

Experiments

Experimental Settings

Baselines. We compare the experimental performance based on ViT-B/16 CLIP with these methods: CLIP (Radford et al. 2021), CoOp (Zhou et al. 2022b), CoCoOp (Co) (Zhou et al. 2022a), KgCoOp (Kg) (Yao, Zhang, and Xu 2023b), PLOT (Chen et al. 2022), MaPLe (Khattak et al. 2023a), PromptSRC (SRC) (Khattak et al. 2023b) and TCP (Yao, Zhang, and Xu 2024). Differing from the textual single-modal regularization of our method, PLOT employs the optimal transport plan in the multi-modal regularization for alignment with visual and textual prompting branches.

(a) Average over 11 datasets				(b) ImageNet				(c) Caltech101				(d) OxfordPets					
Base		Novel		Base		Novel		Base		Novel		Base		Novel		HM	
CLIP	69.34	74.22	71.70	CLIP	72.43	68.14	70.22	CLIP	96.84	94.00	95.40	CLIP	91.17	97.26	94.12		
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		
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		
Kg	80.73	73.60	77.00	Kg	75.83	69.96	72.78	Kg	97.72	94.39	96.03	Kg	94.65	98.10	96.38		
PLOT	81.24	72.98	76.89	PLOT	75.33	70.48	72.83	PLOT	97.86	93.99	95.92	PLOT	95.7	98.09	96.89		
MaPLe	82.28	75.14	78.55	MaPLe	76.66	70.54	73.47	MaPLe	97.74	94.36	96.02	MaPLe	95.43	97.76	96.58		
SRC*	84.12	75.02	79.31	SRC*	77.60	70.73	74.01	SRC*	98.10	93.90	95.97	SRC*	95.50	97.40	96.44		
TCP	84.13	75.36	79.51	TCP	77.27	69.87	73.38	TCP	98.23	95.33	96.42	TCP	94.67	97.20	95.92		
SPTR	84.85	76.76	80.61	SPTR	77.60	71.75	74.56	SPTR	98.01	93.83	95.87	SPTR	94.85	97.10	95.97		
(e) EuroSAT				(f) UCF101				(g) StanfordCars				(h) Flowers102					
Base		Novel		Base		Novel		Base		Novel		Base		Novel		HM	
CLIP	56.48	64.05	60.03	CLIP	70.53	77.50	73.85	CLIP	63.37	74.89	68.65	CLIP	72.08	77.80	74.83		
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		
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		
Kg	85.64	64.34	73.48	Kg	82.89	76.67	79.65	Kg	71.76	75.04	73.36	Kg	95.00	74.73	83.65		
PLOT	90.2	63.5	74.54	PLOT	82.56	75.56	78.92	PLOT	71.5	73.77	72.62	PLOT	95.10	72.20	82.10		
MaPLe	94.07	73.23	82.35	MaPLe	83.00	78.66	80.77	MaPLe	72.94	74.00	73.47	MaPLe	95.92	72.46	82.56		
SRC*	92.40	68.43	78.63	SRC*	86.93	78.33	82.41	SRC*	78.40	74.73	75.52	SRC*	97.90	76.77	86.06		
TCP	91.63	74.73	82.32	TCP	87.13	80.77	83.83	TCP	80.80	74.13	77.32	TCP	97.73	75.57	85.23		
SPTR	93.15	72.80	81.73	SPTR	88.17	81.00	84.47	SPTR	81.85	75.43	78.53	SPTR	98.56	77.59	86.86		
(i) Food101				(j) FGVCaircraft				(k) SUN397				(l) DTD					
Base		Novel		Base		Novel		Base		Novel		Base		Novel		HM	
CLIP	90.10	91.22	90.66	CLIP	27.19	36.29	31.09	CLIP	69.36	75.35	72.23	CLIP	53.24	59.90	56.37		
CoOp	88.33	82.26	85.19	CoOp	40.44	22.30	28.75	CoOp	80.60	65.89	72.51	CoOp	79.44	41.18	54.24		
Co	90.70	91.29	90.99	Co	33.41	23.71	27.74	Co	79.74	76.86	78.27	Co	77.01	56.00	64.85		
Kg	90.5	91.7	91.09	Kg	36.21	33.55	34.83	Kg	80.29	76.53	78.36	Kg	77.55	54.99	64.35		
PLOT	90.98	91.54	91.28	PLOT	35.6	28.5	31.66	PLOT	79.96	77.33	78.64	PLOT	78.9	57.9	66.8		
MaPLe	90.71	92.05	91.38	MaPLe	37.44	35.61	36.50	MaPLe	80.82	78.70	79.75	MaPLe	80.36	59.18	68.16		
SRC*	90.63	91.50	91.06	SRC*	42.30	36.97	39.46	SRC*	82.83	79.00	80.87	SRC*	82.60	57.50	67.80		
TCP	90.57	91.37	90.97	TCP	41.97	34.43	37.83	TCP	82.63	78.20	80.35	TCP	82.77	58.07	68.25		
SPTR	91.11	92.74	91.94	SPTR	44.26	40.18	42.09	SPTR	82.56	79.26	80.90	SPTR	83.35	62.58	71.50		

Table 2: Base-to-novel generalization experiments. “*” denote the performance obtained by our re-implementation.

Datasets. The datasets cover multiple scenes including ImageNet (Deng et al. 2009) and Caltech101 (Fei-Fei, Fergus, and Perona 2004), OxfordPets (Parkhi et al. 2012), StanfordCars (Krause et al. 2013), Flowers102 (Nilsback and Zisserman 2008), Food101 (Bossard, Guillaumin, and Van Gool 2014), FGVCaircraft (Maji et al. 2013), SUN397 (Xiao et al. 2010) UCF101 (Soomro, Zamir, and Shah 2012), DTD (Cimpoi et al. 2014) and EuroSAT (Helber et al. 2019). We use ImageNetA (Hendrycks et al. 2021b), ImageNet-R (Hendrycks et al. 2021a), ImageNet-Sketch (Wang et al. 2019) and ImageNetV2 (Recht et al. 2019) for domain generalization.

Implementation Details. We set the embeddings length to 4. We train 50 epochs for few-shot learning tasks and 20 epochs for other tasks. The learning rate is 0.0025 via SGD optimizer on a single GPU. We use the ViT-B/16 model-based CLIP and set α to 0.3. For domain generalization and cross-dataset evaluation, we train the ImageNet source model on all classes in the first 3 layers of encoders. For the base-to-novel settings

and few-shot learning, we set learning depth to 9. We set N to 60, which is consistent with pre-trained CLIP. Attacks are generated with a perturbation boundary $\epsilon = 1/255$.

Experimental Results

Few-Shot Learning. The evaluation of the model is performed at different K -shot levels per class, encompassing values of 1, 2, 4, 8, and 16. In Figure 2, our method consistently delivers enhancements across all few-shot settings when compared to existing methods. SPTR achieves improved performance in non-generalization classification, even when training resources are limited ($K = 2$).

Base-to-Novel Generalization. The datasets are split into base and novel classes. The model is trained only on the base classes in the 16 shots setting and evaluated on base and novel classes. In Table 2, SPTR provides the best-averaged results of 84.85%, 76.76%, and 80.61% on the base classes, novel classes, and harmonic mean, respectively.

	Source					Target							
	ImageNet	Caltech101	OxfordPets	StanfordCars	Flowers102	Food101	Aircraft	SUN397	DTD	EuroSAT	UCF101	Average	
CLIP	66.72	92.94	89.07	65.29	71.30	86.11	24.87	62.62	44.56	47.69	66.77	65.12	Δ
CoOp	71.51	93.70	89.14	64.51	68.71	85.30	18.47	64.15	41.92	46.39	66.55	63.88	-2.3
Co	71.02	94.43	90.14	65.32	71.88	86.06	22.94	67.36	45.73	45.37	68.21	65.74	+0.6
Kg	70.66	93.92	89.83	65.41	70.01	86.36	22.51	66.16	46.35	46.04	68.50	65.51	+0.4
PLOT	70.15	94.60	90.23	65.41	71.97	86.32	22.87	67.22	44.99	46.57	68.32	65.85	+0.7
MaPLe	70.72	93.53	90.49	65.57	72.23	86.20	24.74	67.01	46.49	48.06	68.69	66.30	+1.2
SRC	71.27	93.60	90.25	65.70	70.25	86.15	23.90	67.10	46.87	45.50	68.75	65.81	+0.7
TCP	71.40	93.97	91.25	64.69	71.21	86.69	23.45	67.15	44.35	51.45	68.73	66.26	+1.2
Ours	70.05	94.45	90.70	65.83	72.22	86.13	23.60	67.00	49.13	51.33	69.10	66.95	+1.8

Table 3: Cross-dataset benchmark evaluation. The symbol Δ represents the value that increases in comparison to CLIP.

	Source	Target					
	ImageNet	-V2	-S	-A	-R	Avg.	
CLIP	66.73	60.83	46.15	47.77	73.96	57.18	Δ
CoOp	71.51	64.2	47.99	49.71	75.21	59.28	+2.1
Co	71.02	64.07	48.75	50.63	76.18	59.91	+2.8
Kg	70.66	64.1	48.97	50.69	76.7	60.11	+2.9
PLOT	70.15	64.17	49.15	50.83	76.5	60.16	+2.9
MaPLe	70.72	64.07	49.15	50.9	76.98	60.27	+3.1
SRC*	71.27	64.35	49.55	50.90	76.80	60.40	+3.2
TCP	71.40	64.05	49.10	50.55	77.8	60.37	+3.2
Ours	70.05	64.40	48.78	51.30	77.9	60.59	+3.4

Table 4: Domain generalization experiments. The symbol Δ represents the value that increases in comparison to CLIP. ‘*’ denotes the performance obtained by our re-implementation.

Domain Generalization. We train our model using the ImageNet in 16 shots and evaluate its performance directly on 4 different variants of the ImageNet. In Table 4, SPTR outperforms all existing methods on target datasets, with an overall highest average accuracy of 60.59%. Compared with TCP (CVPR2024), SPTR shows improved performance in ImageNet-A with a large margin.

Cross-Dataset Generalization. We evaluate our method on 10 unseen datasets. In Table 3, in comparison with TCP (CVPR2024), SPTR achieves better generalization on average for 10 datasets. This indicates that SPTR favors better generalization for a wide range of unseen datasets.

Further Analysis

Training and Inference Compute Cost Analysis

In Table 5, we show the computed cost analysis (Wang et al. 2021; Yang et al. 2022) of SPTR and compare it with prompt learning methods. Compared to MaPLe, SPTR has fewer parameters. CoCoOp is significantly slower, due to the single output design of the image encoder. CoCoOp transfers the output features of the image encoder to the text embeddings,

which is more inference time consumption. Due to the presence of adversarial designs in SPTR, the inference speed is relatively slow (Madry et al. 2017).

Method	Params	Train time (min)	HM
CoOp	2048	10.88	71.65
Co	35360	39.53	75.83
Kg	2048	11.25	77.00
PLOT	8192	10.85	76.48
MaPLe	3.55 M	10.58	79.68
SRC	5120	13.13	79.97
TCP	2048	11.85	80.11
SPTR (Ours)	5120	14.50	80.31

Table 5: The computational cost comparison using SUN397 dataset. Training time for all methods is calculated for 10 epochs on a single GPU.

Embeddings Length and Learning Depth

In Table 6, our findings indicate that the performance reaches its peak when the length of embeddings is set to 4 on base-to-novel generalization for HM. In Table 7, we note that increasing the learning depth generally increases the performance. As the number of layers increases to 11, the HM value decreases. It indicates excessive fine-tuning of the model, causing it to lose CLIP generalization and the ability of V-L alignment (Zhou et al. 2022a).

Embeddings Length	1	2	4	6	8
Caltech101	95.13	95.31	95.87	93.10	95.33
Food101	89.11	91.09	91.94	88.19	90.92

Table 6: Analysis of embeddings length (HM).

Distance Measurement for Textual Branch

The analysis in Table 9 suggests that, among the various textual regularization techniques explored, the OT (ours) leads

Learning Depth	1	3	5	7	9	11
Caltech101	91.83	94.12	95.11	95.19	95.87	95.33
Food101	86.53	88.00	89.71	91.01	91.94	91.10

Table 7: Analysis of learning depth (HM).

	ViT-B/32	ViT-B/16	EVA-CLIP
HM	78.55	80.61	83.12

Table 8: Applying to ViT instances for Avg. (11 datasets).

Design	Base Acc.	Novel Acc.	HM
MSE	84.81	74.19	79.17
L1	84.60	76.00	80.09
Cosine	84.79	75.51	79.90
OT (Ours)	84.85	76.76	80.61

Table 9: Distance measurement for Avg. (11 datasets).

	Caltech101			Food101		
ϵ	1/255	4/255	8/255	1/255	4/255	8/255
HM	95.87	94.00	93.58	91.94	90.88	89.45

Table 10: Ablation of perturbation boundary ϵ .

	Caltech101			Food101		
α	0.1	0.3	0.5	0.1	0.3	0.5
HM	93.00	95.87	95.18	88.22	91.94	90.32

Table 11: Ablation of textual regularization α .

to the best performance on base-to-novel generalization for 11 datasets. Cosine similarity primarily focuses on the angular similarity between vectors, MSE and L1 loss are designed to measure the distance between individual data points, rather than the overall distribution of the data. The optimal transport can calculate the distance between two distributions with a fine-grained matching to align features.

Analysis of Perturbation Boundary and Textual α

In Table 10 and Table 11, our findings indicate that the performance reaches its peak when the perturbation boundary ϵ is set to 1/255 and α to 0.3 on base-to-novel generalization.

Applying to Other ViT Instances

Our framework was discovered to be adaptable to various types of Vision Transformer (ViT) instances, such as EVA-CLIP (Fang et al. 2023) and ViT-B/32 in Table 8. In the future, we will validate our approach on more ViT architectures.

Analysis of Integrated Contribution

We employ an independent V-L prompting framework as our baseline model (row-1). By enforcing the optimal trans-

Component	Base Acc.	Novel Acc.	HM
1: Baseline model	84.32	71.38	77.34
2: + OT	84.39	74.90	79.36
3: + OT + Similarity paradigm	84.85	76.76	80.61

Table 12: Analysis of our components for Avg. (11 datasets).

port (OT) for textual regularization based on row-1, row-2 is higher than row-1 on novel performance in Table 12. This is because OT plays a crucial role in mitigating the forgetting of general textual knowledge with a fine-grained perspective. By integrating a similarity paradigm based on row-2, we observed that row-3 outperforms row-2 on novel classes while maintaining base class gains. It suggests that the similarity paradigm enhances the ability of generalized robustness.

Related Work

Prompt Learning for Vision-Language Models

Recently, the strong generalization capability of CLIP has made it a foundation for many methods in adapting pre-trained Vision-Language models (VLMs) for downstream tasks. Prompt tuning is a widely used technique in VLMs for learning downstream tasks. The use of text prompts, which are instructions provided to the language branch of VLMs, is a common practice to enhance task understanding. Full fine-tuning (Cao et al. 2024b) and linear probe are two common methods used to adapt VLMs to downstream tasks. Some methods (Yao, Zhang, and Xu 2023a; Zhou et al. 2022a,b) fine-tune the CLIP model specifically for few-shot image recognition by optimizing a continuous set of token embeddings in the language branch. Moreover, some methods constrain the prompts to contain the essential general knowledge and prior distribution learning (Lu et al. 2022). In addition to multi-modal prompt tuning (Khattak et al. 2023a,b; Lee et al. 2023), some methods (Li et al. 2024a; Zhou et al. 2024) apply adversarial training perspective for prompt tuning to effectively align its V-L representations for pre-trained CLIP.

Conclusion

Prompt learning involves optimizing the embeddings for adaptation to downstream tasks. Advanced prompt learning methods usually initialize and optimize the context, e.g., “a photo of a []”, for downstream task adaptation. While optimizing the embeddings can enhance performance on specific tasks, it may result in subpar generalization performance on unseen classes or datasets drawn from diverse distributions. To mitigate this issue, we propose SPTR, which is a two-pronged design based on hand-crafted prompts. Four representative tasks (i.e., non-generalization few-shot learning, base-to-novel generalization, cross-dataset generalization, domain generalization) across 11 datasets demonstrate that SPTR outperforms existing prompt learning methods.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (62125201).

References

- Akinwande, V.; Jiang, Y.; Sam, D.; and Kolter, J. Z. 2023. Understanding prompt engineering may not require rethinking generalization. *arXiv preprint arXiv:2310.03957*.
- Bangalath, H.; Maaz, M.; Khattak, M. U.; Khan, S. H.; and Shahbaz Khan, F. 2022. Bridging the gap between object and image-level representations for open-vocabulary detection. *Advances in Neural Information Processing Systems*, 35: 33781–33794.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. 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*, 446–461. Springer.
- Cao, J.; Guan, S.; Ge, Y.; Li, W.; Yang, X.; and Ma, C. 2024a. Neural Material Adaptor for Visual Grounding of Intrinsic Dynamics. *arXiv preprint arXiv:2410.08257*.
- Cao, Q.; Chen, Y.; Ma, C.; and Yang, X. 2023a. Break the bias: Delving semantic transform invariance for few-shot segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Cao, Q.; Chen, Y.; Ma, C.; and Yang, X. 2023b. Few-shot rotation-invariant aerial image semantic segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–13.
- Cao, Q.; Xu, Z.; Chen, Y.; Ma, C.; and Yang, X. 2024b. Domain-controlled prompt learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 936–944.
- Chen, G.; Yao, W.; Song, X.; Li, X.; Rao, Y.; and Zhang, K. 2022. Plot: Prompt learning with optimal transport for vision-language models. *arXiv preprint arXiv:2210.01253*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3606–3613.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Fang, Y.; Wang, W.; Xie, B.; Sun, Q.; Wu, L.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2023. EVA: Exploring the Limits of Masked Visual Representation Learning at Scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 19358–19369.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. 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*, 178–178. IEEE.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7): 2217–2226.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; et al. 2021a. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8340–8349.
- Hendrycks, D.; Zhao, K.; Basart, S.; Steinhardt, J.; and Song, D. 2021b. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15262–15271.
- Jia, S.; Ma, C.; Song, Y.; and Yang, X. 2020. Robust tracking against adversarial attacks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX 16*, 69–84. Springer.
- Jia, S.; Ma, C.; Yao, T.; Yin, B.; Ding, S.; and Yang, X. 2022. Exploring frequency adversarial attacks for face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4103–4112.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023a. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19113–19122.
- Khattak, M. U.; Wasim, S. T.; Naseer, M.; Khan, S.; Yang, M.-H.; and Khan, F. S. 2023b. Self-regulating prompts: Foundational model adaptation without forgetting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15190–15200.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, 554–561.
- Lee, Y.-L.; Tsai, Y.-H.; Chiu, W.-C.; and Lee, C.-Y. 2023. Multimodal prompting with missing modalities for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14943–14952.
- Li, L.; Guan, H.; Qiu, J.; and Spratling, M. 2024a. One prompt word is enough to boost adversarial robustness for pre-trained vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24408–24419.
- Li, W.; Guo, Y.; Zheng, J.; Lin, H.; Ma, C.; Fang, L.; and Yang, X. 2024b. SparseFormer: Detecting Objects in HRW Shots via Sparse Vision Transformer. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 4851–4860.
- Li, X.; Ma, C.; Yang, X.; and Yang, M.-H. 2024c. VidToMe: Video Token Merging for Zero-Shot Video Editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7486–7495.
- Liu, C.; Wang, X.; Fan, Y.; Li, S.; and Qian, X. 2024a. Decoupling degradations with recurrent network for video restoration in under-display camera. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3558–3566.
- Liu, C.; Wang, X.; Li, S.; Wang, Y.; and Qian, X. 2023. Fsi: Frequency and spatial interactive learning for image restoration in under-display cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12537–12546.

- Liu, C.; Wang, X.; Xu, X.; Tian, R.; Li, S.; Qian, X.; and Yang, M.-H. 2024b. Motion-adaptive Separable Collaborative Filters for Blind Motion Deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25595–25605.
- Lu, Y.; Liu, J.; Zhang, Y.; Liu, Y.; and Tian, X. 2022. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5206–5215.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M.; and Vedaldi, A. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*.
- Miyato, T.; Dai, A. M.; and Goodfellow, I. 2016. Adversarial training methods for semi-supervised text classification. *arXiv preprint arXiv:1605.07725*.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, 722–729. IEEE.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. 2012. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, 3498–3505. IEEE.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Recht, B.; Roelofs, R.; Schmidt, L.; and Shankar, V. 2019. Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*, 5389–5400. PMLR.
- Shi, C.; Wang, X.; Shi, S.; Wang, X.; Zhu, M.; Wang, N.; and Gao, X. 2024. Foodfusion: A novel approach for food image composition via diffusion models. *arXiv preprint arXiv:2408.14135*.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Tang, Z.; Zhang, Y.; Shi, S.; He, X.; Han, B.; and Chu, X. 2022. Virtual homogeneity learning: Defending against data heterogeneity in federated learning. In *International Conference on Machine Learning*, 21111–21132. PMLR.
- Villani, C.; et al. 2009. *Optimal transport: old and new*, volume 338. Springer.
- Wang, H.; Ge, S.; Lipton, Z.; and Xing, E. P. 2019. Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32.
- Wang, X.; Lai, S.; Chai, Z.; Zhang, X.; and Qian, X. 2021. SPGNet: Serial and parallel group network. *IEEE Transactions on Multimedia*, 24: 2804–2814.
- Wortsman, M.; Ilharco, G.; Kim, J. W.; Li, M.; Kornblith, S.; Roelofs, R.; Lopes, R. G.; Hajishirzi, H.; Farhadi, A.; Namkoong, H.; et al. 2022. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7959–7971.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, 3485–3492. IEEE.
- Xie, F.; Chu, L.; Li, J.; Lu, Y.; and Ma, C. 2023. Videotrack: Learning to track objects via video transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22826–22835.
- Yang, J.; Lai, S.; Wang, X.; Wang, Y.; and Qian, X. 2022. Diversity-learning block: conquer feature homogenization of multibranch. *IEEE Transactions on Neural Networks and Learning Systems*.
- Yao, H.; Zhang, R.; and Xu, C. 2023a. TCP: Textual-based Class-aware Prompt tuning for Visual-Language Model. *arXiv preprint arXiv:2311.18231*.
- Yao, H.; Zhang, R.; and Xu, C. 2023b. Visual-Language Prompt Tuning With Knowledge-Guided Context Optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6757–6767.
- Yao, H.; Zhang, R.; and Xu, C. 2024. TCP: Textual-based Class-aware Prompt tuning for Visual-Language Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23438–23448.
- Zhang, R.; Fang, R.; Zhang, W.; Gao, P.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2021a. Tip-adapter: Training-free clip-adapter for better vision-language modeling. *arXiv preprint arXiv:2111.03930*.
- Zhang, W.; Lyu, S.; Xue, F.; Yao, C.; Zhu, Z.; and Jia, Z. 2022. Predict the rover mobility over soft terrain using articulated wheeled bevameter. *IEEE Robotics and Automation Letters*, 7(4): 12062–12069.
- Zhang, Y.; Gong, M.; Liu, T.; Niu, G.; Tian, X.; Han, B.; Schölkopf, B.; and Zhang, K. 2021b. Causaladv: Adversarial robustness through the lens of causality. *arXiv preprint arXiv:2106.06196*.
- Zhang, Y.; Tian, X.; Li, Y.; Wang, X.; and Tao, D. 2020. Principal component adversarial example. *IEEE Transactions on Image Processing*, 29: 4804–4815.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16816–16825.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.
- Zhou, Y.; Xia, X.; Lin, Z.; Han, B.; and Liu, T. 2024. Few-Shot Adversarial Prompt Learning on Vision-Language Models. *arXiv preprint arXiv:2403.14774*.
- Zhu, B.; Niu, Y.; Han, Y.; Wu, Y.; and Zhang, H. 2023. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15659–15669.