

# Doubly Debiased Test-Time Prompt Tuning for Vision-Language Models

Fei Song<sup>1 2\*</sup>, Yi Li<sup>1 2\*</sup>, Rui Wang<sup>1 2</sup>, Jiahuan Zhou<sup>3</sup>, Changwen Zheng<sup>1 2</sup>, Jiangmeng Li<sup>1 2†</sup>

<sup>1</sup>National Key Laboratory of Space Integrated Information System, Institute of Software, Chinese Academy of Sciences

<sup>2</sup>University of Chinese Academy of Sciences

<sup>3</sup>Wangxuan Institute of Computer Technology, Peking University

{songfei2022, liyi2022, wangrui, changwen, jiangmeng2019}@iscas.ac.cn, jiahuanzhou@pku.edu.cn

## Abstract

Test-time prompt tuning for vision-language models has demonstrated impressive generalization capabilities under zero-shot settings. However, tuning the learnable prompts solely based on unlabeled test data may induce prompt optimization bias, ultimately leading to suboptimal performance on downstream tasks. In this work, we analyze the underlying causes of prompt optimization bias from both the model and data perspectives. In terms of the model, the entropy minimization objective typically focuses on reducing the entropy of model predictions while overlooking their correctness. This can result in overconfident yet incorrect outputs, thereby compromising the quality of prompt optimization. On the data side, prompts affected by optimization bias can introduce misalignment between visual and textual modalities, which further aggravates the prompt optimization bias. To this end, we propose a Doubly Debiased Test-Time Prompt Tuning method, abbreviated as D<sup>2</sup>TPT. Specifically, we first introduce a dynamic retrieval-augmented modulation module that retrieves high-confidence knowledge from a dynamic knowledge base using the test image feature as a query, and uses the retrieved knowledge to modulate the predictions. Guided by the refined predictions, we further develop a reliability-aware prompt optimization module that incorporates a confidence-based weighted ensemble and cross-modal consistency distillation to impose regularization constraints during prompt tuning. Extensive experiments across 15 benchmark datasets involving both natural distribution shifts and cross-datasets generalization demonstrate that D<sup>2</sup>TPT outperforms baselines, validating its effectiveness in mitigating prompt optimization bias.

**Code** — <https://github.com/FF2127/D2TPT>

## Introduction

Benefiting from large-scale pretraining, vision-language models (VLMs), exemplified by CLIP (Radford et al. 2021), have demonstrated remarkable zero-shot generalization capabilities across a wide range of downstream tasks. However, due to the prevalence of domain shift at test time, vision-language models still suffer from performance degradation when deployed in practical downstream scenarios. To

\*Equal contribution.

†Corresponding author.

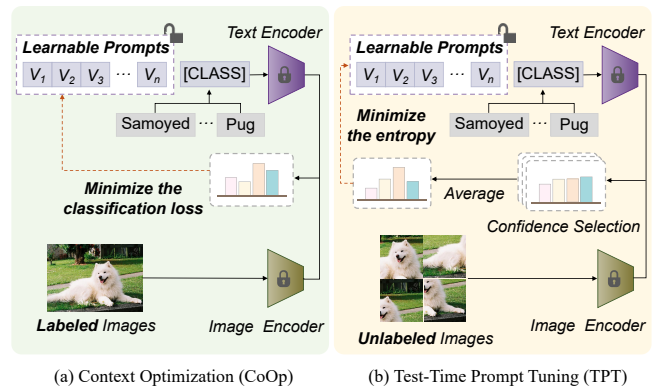


Figure 1: Comparison of the classical prompt tuning CoOp (Zhou et al. 2022) and test-time prompt tuning TPT (Shu et al. 2022). CoOp uses a few labeled samples to optimize the learnable prompt via supervised classification loss, while TPT performs label-free optimization by minimizing the entropy of predictions.

address this issue, prior studies (Zhou et al. 2022; Khatkhat et al. 2023; Li et al. 2025) have investigated a variety of prompt tuning strategies that adapt the models to downstream tasks using labeled training data. Figure 1(a) illustrates a representative method, CoOp (Zhou et al. 2022), which models the context words of the textual prompt using a set of learnable vectors while keeping the entire pre-trained parameters fixed. Through minimizing the classification loss to optimize these learnable vectors, CoOp achieves superb domain generalization performance. Nevertheless, the reliance on annotated data remains a fundamental bottleneck that hinders the practical deployment of VLMs in real-world applications.

As a simple yet effective alternative, Test-Time Prompt Tuning (TPT) (Shu et al. 2022) has attracted widespread attention from researchers in recent years. As shown in Figure 1(b), TPT enables prompt tuning without requiring labeled training data by minimizing entropy to adapt the learnable text prompt for each unlabeled test sample. Although test-time prompt tuning methods (Feng et al. 2023; Zhang et al. 2024) have shown promising results in effectively adapting VLMs to target-domain data, relying solely on un-

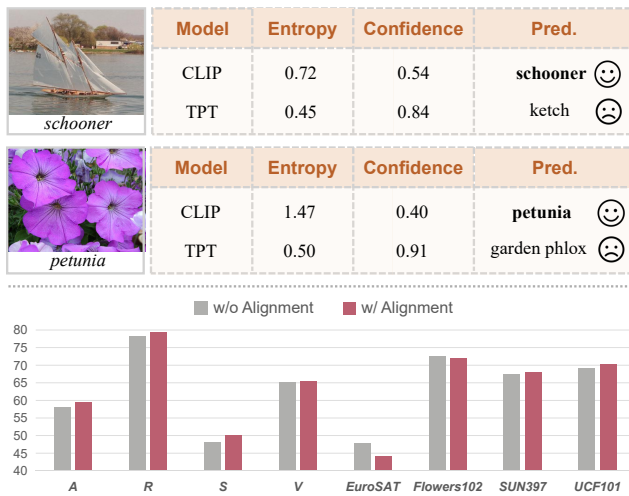


Figure 2: (Top) Examples illustrating that entropy minimization can lead to overconfident predictions. For instance, although TPT’s prediction for petunia has lower entropy and higher confidence, the prediction result is incorrect. (Bottom) Effect of the alignment strategy across different datasets. A, R, S, and V denote the abbreviations of the ImageNet-A, ImageNet-R, ImageNet-Sketch, and ImageNet-V2 datasets, respectively.

labeled test samples to optimize the learnable prompts is intuitively insufficient and can lead to *prompt optimization bias*. We substantiate this claim by analyzing from both the model and data perspectives.

At the model level, existing test-time prompt tuning methods (Shu et al. 2022) optimize prompts by minimizing the entropy of model predictions on test samples. This objective tends to drive the optimization of learnable prompts based on low-entropy predictions. However, as illustrated in Figure 2(top), low-entropy predictions are not necessarily correct. When prompt optimization is guided by such incorrect predictions, the model is likely to produce overconfident yet inaccurate outputs. At the data level, incorporating learnable prompts into image-text inputs is, in principle, expected to enhance semantic alignment between the visual and textual modalities within the shared embedding space. Nevertheless, prompts affected by optimization bias may instead deteriorate this alignment. Typical alignment-based test-time prompt tuning (Samadh et al. 2023) adapts learnable image-text prompts by aligning the statistics of test samples with those of ImageNet. However, this approach is less robust on datasets that differ significantly from ImageNet. As evidenced in Figure 2(bottom), this alignment strategy performs well only on ImageNet-variant datasets, while yielding limited improvements or even performance degradation on other datasets.

To address the aforementioned issues, we propose a Doubly Debiased Test-Time Prompt Tuning method, abbreviated as D<sup>2</sup>TPT, which aims to mitigate prompt optimization bias and enable learnable prompts to better support

model generalization to downstream tasks. D<sup>2</sup>TPT consists of two key ingredients: (1) A dynamic retrieval-augmented modulation module, which introduces a dynamic knowledge base designed to store high-confidence predictions and support continuous updates. When a test image arrives, we use its corresponding feature vector as a query to retrieve the matched class prototype from the knowledge base. The label information associated with the retrieved prototype serves as a high-confidence external supervisory signal to modulate the model’s original prediction for the test image. (2) A reliability-aware prompt optimization module, which imposes two regularization constraints on the optimization of learnable prompts, based on the modulated predictions. On one hand, a confidence-based weighted ensemble strategy is designed to integrate discriminative information from augmented views, thereby suppressing the interference caused by low-quality augmentations. On the other hand, we develop a cross-modal consistency distillation strategy, where the image and text modalities alternately act as teacher and student to mutually learn from each other, encouraging semantic consistency between the two modalities within the shared embedding space. The evaluation results on 15 benchmark datasets involving both natural distribution shifts and cross-datasets generalization demonstrate the effectiveness of D<sup>2</sup>TPT.

Our key contributions can be summarized as follows:

- We reveal that test-time prompt tuning methods solely based on unlabeled test data suffer from prompt optimization bias, and we analyze its underlying causes from both model and data perspectives.
- To mitigate the prompt optimization bias, we propose a Doubly Debiased Test-Time Prompt Tuning method, i.e., D<sup>2</sup>TPT, which consists of a dynamic retrieval-augmented modulation module and a reliability-aware prompt optimization module.
- On representative tasks involving natural distribution shifts and cross-datasets generalization, the proposed method consistently improves the model’s generalization capability.

## Related Work

In this section, we provide an overview of related work on prompting for VLMs and test-time prompt tuning.

### Prompting for VLMs

Leveraging large-scale image-text pairs from the internet, VLMs have demonstrated promising performance across a wide range of downstream tasks. Inspired by the success of prompt learning in natural language processing (Lester, Al-Rfou, and Constant 2021; Li and Liang 2021), researchers have developed a variety of prompt-based methods to effectively adapt VLMs to downstream tasks using only a few labeled examples. Specifically, CoOp (Zhou et al. 2022) introduces a set of learnable vectors into the input of the text branch and optimizes them by minimizing the classification loss. To further enhance the prompt generalization, Bayesian prompt learning (Derakhshani et al. 2023) reformulates prompt learning from a Bayesian perspective,

casting it as a variational inference problem. Nevertheless, prompting only a single branch of CLIP is sub-optimal, as it limits the flexibility to adapt both representation spaces for downstream tasks. Accordingly, Maple (Khattak et al. 2023) introduces multi-modal prompt learning across both the vision and language branches to enhance the alignment between their representations. CoPrompt (Roy and Etemad 2024) imposes a consistency constraint on both the language and vision branches between the trainable and pre-trained models to prevent overfitting on downstream tasks. However, these methods typically rely on labeled training data, thereby limiting their applicability in real-world scenarios.

### Test-time Prompt Tuning

As a simple yet effective approach that adapts vision-language models to downstream tasks at test time without requiring labeled training samples from the target domain, test-time prompt tuning has attracted increasing attention from researchers. TPT (Shu et al. 2022) first introduces test-time prompt tuning, which optimizes prompts by minimizing entropy to encourage consistent predictions across augmented views. Following this paradigm, a series of subsequent studies have been proposed to further enhance TPT. Specifically, DiffTPT (Feng et al. 2023) leverages pre-trained diffusion models to generate diverse and informative data, thereby enriching the diversity of augmented views. C-TPT (Yoon et al. 2024) optimizes prompts during test time via enhanced calibration. To address the challenge of unsupervised prompt optimization without gradients, B<sup>2</sup>TPT (Meng et al. 2025) proposes a black-box test-time prompt tuning that overcomes the gradient-free limitation while reducing complexity. Considering the correlations among test samples, DynaPrompt (Xiao et al. 2025) builds upon an online prompt buffer to adaptively select and optimize relevant prompts for each test sample during tuning. In this work, we focus on the potential prompt optimization bias that arises when tuning learnable prompts solely based on unlabeled test samples. We analyze its underlying causes from both model and data perspectives and propose targeted strategies to effectively mitigate such bias.

### Preliminaries

Before introducing D<sup>2</sup>TPT, we first review the preliminaries of test-time prompt tuning for downstream classification tasks using CLIP.

#### Contrastive Language-Image Pretraining (CLIP)

CLIP is a representative foundation vision-language model (VLM) pre-trained on approximately 400 million image-text pairs. It comprises a visual encoder  $V(\cdot)$  and a textual encoder  $T(\cdot)$ , jointly designed to mitigate the semantic gap between visual inputs and textual descriptions. By leveraging a contrastive learning loss, CLIP aligns the visual and textual modalities into a shared embedding space, enabling the learning of generalizable visual representations and enhancing transferability across downstream tasks. For an image classification task with  $C$  categories, CLIP formulates classification via a textual prompt-based approach. Specifically,

each class label  $c \in \{1 \dots C\}$  is converted into a text prompt  $\{t_c\}_{c=1}^C$  by prepending a template (e.g., “a photo of a”). These text prompts are then encoded by the textual encoder to yield the corresponding text feature vectors  $\{\mathbf{z}_{text}^c\}_{c=1}^C$ , where  $\mathbf{z}_{text}^c = T(t_c)$ . Given a test image  $x_{test}$ , its visual representation is obtained as  $\mathbf{z}_{img} = V(x_{test})$ . The classification is performed by computing the cosine similarity between the image feature and each class-specific text feature. The resulting class probability distribution is given by:

$$P(c|x_{test}) = \frac{\exp(\cos(\mathbf{z}_{img}, \mathbf{z}_{text}^c)/\tau_{temp})}{\sum_{c'=1}^C \exp(\cos(\mathbf{z}_{img}, \mathbf{z}_{text}^{c'})/\tau_{temp})}, \quad (1)$$

where  $\cos(\cdot)$  denotes the cosine similarity function, and  $\tau_{temp}$  is the temperature hyperparameter.

#### Test-time Prompt Tuning for CLIP

Despite the impressive zero-shot generalization capability of CLIP, effectively adapting it to unseen distributions at test time remains a critical challenge. Furthermore, fine-tuning the entire model is computationally prohibitive, particularly for large-scale transformer architectures. Accordingly, test-time prompt tuning has emerged as a lightweight strategy that adapts the input prompts based solely on the test sample itself, without requiring access to labeled training data. Specifically, given a test image  $x_{test}$ , TPT generates  $N$  augmented views using a transformation family  $\mathcal{A}$ , resulting in  $\{\mathcal{A}_n(x_{test})\}_{n=1}^N$ . For each view, CLIP produces a prediction  $P_{\mathbf{p}}(c|\mathcal{A}_n(x_{test}))$  based on a learnable prompt  $\mathbf{p}$ . The optimization objective aims to minimize the entropy of the average prediction distribution over all augmented views:

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \mathcal{H}(\tilde{P}_{\mathbf{p}}(c|x_{test})) \quad (2)$$

where  $\mathcal{H}(\tilde{P}_{\mathbf{p}}(c|x_{test})) = -\sum_{c=1}^C \tilde{P}_{\mathbf{p}}(c|x_{test}) \log \tilde{P}_{\mathbf{p}}(c|x_{test})$ , and  $\tilde{P}_{\mathbf{p}}(c|x_{test}) = \frac{1}{N} \sum_{n=1}^N P_{\mathbf{p}}(c|\mathcal{A}_n(x_{test}))$ .

To mitigate the impact of unreliable augmentations (e.g., random crops that remove discriminative content), TPT employs a confidence selection mechanism that computes the self-entropy  $\mathcal{H}(P^n)$  of each prediction, where  $P^n = P_{\mathbf{p}}(c|\mathcal{A}_n(x_{test}))$ . Then, select only those views with entropy below a threshold  $\tau$ , which is determined as the  $\rho$ -percentile among the entropy values of the  $N$  views, ranked from low to high. Thus, the averaged prediction under confidence selection is computed as follows:

$$\tilde{P}_{\mathbf{p}}(c|x_{test})' = \frac{1}{\rho N} \sum_{n=1}^N \mathbb{I}[\mathcal{H}(P^n) \leq \tau] P^n, \quad (3)$$

where  $\mathbb{I}[\mathcal{H}(P^n) \leq \tau]$  denotes an indicator function that returns 1 if the prediction entropy is below the threshold  $\tau$ , and 0 otherwise.

By minimizing the entropy of predictions averaged over selected high-confidence augmentations, TPT adapts prompts in an unsupervised manner, enabling CLIP to better generalize to out-of-distribution data at test time.

### Methodology

In this section, we present our proposed D<sup>2</sup>TPT method, which introduces modality-specific learnable prompts for

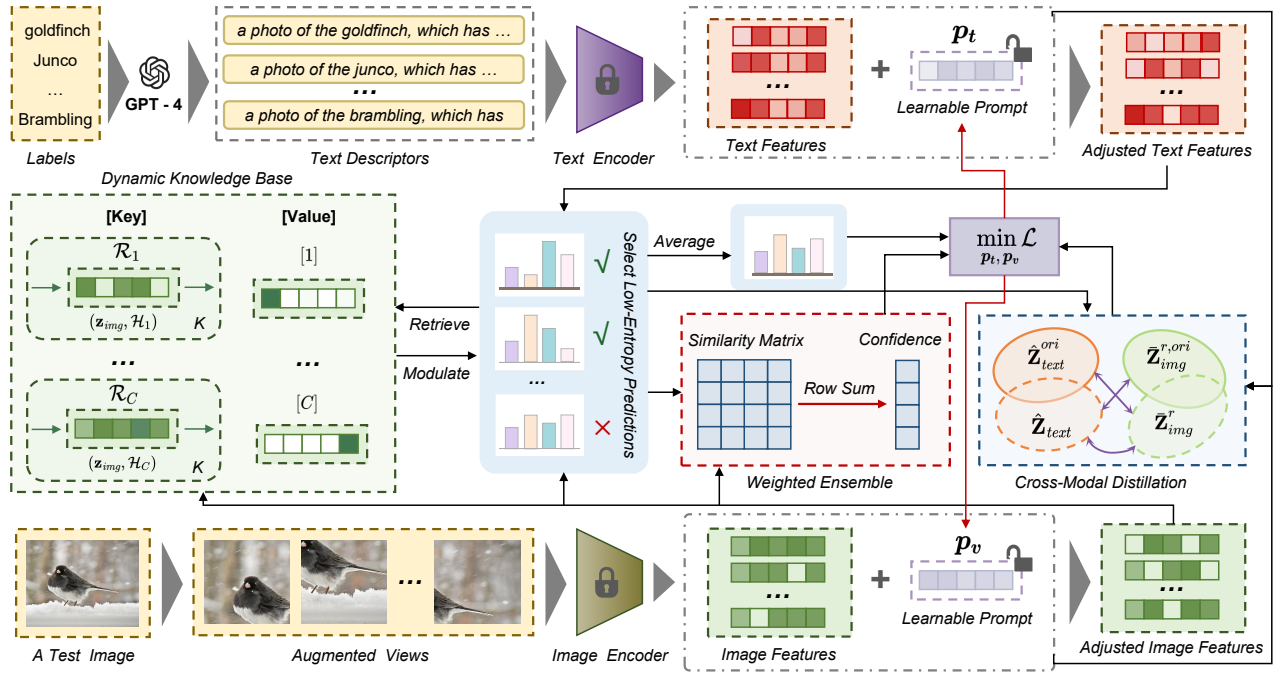


Figure 3: The overall framework of our D<sup>2</sup>TPT method.

image and text inputs, and integrates dedicated modules to jointly mitigate prompt optimization bias. Figure 3 illustrates an overview of D<sup>2</sup>TPT.

### Modality-Specific Prompt Design

In this work, we adopt a prompt ensembling strategy that utilizes multiple contextual prompt templates. Specifically, for each class  $c$ , general prompt templates  $t_c^{gen}$  are constructed following CLIP, while class-specific prompt templates  $t_c^{spe}$  are generated with the assistance of GPT-4 (OpenAI 2023). The textual encoder  $T(\cdot)$  is then applied to obtain the corresponding textual embedding:  $\mathbf{z}_{gen}^c = T(t_c^{gen})$  and  $\mathbf{z}_{spe}^c = T(t_c^{spe})$ . The prototype of these textual prompt templates in the embedding space is calculated as  $\mathbf{z}_{text}^{pro,c} = avg(\mathbf{z}_{gen}^c, \mathbf{z}_{spe}^c)$ , where  $avg(\cdot)$  denotes the mean operation.

Since generating class-specific prompts solely based on class labels can introduce information irrelevant to the input image, we design a learnable textual prompt  $\mathbf{p}_t$  that dynamically adjusts each textual prototype to better align with the visual feature of the test image. For each class  $c$ , the adjusted textual representation is computed as  $\mathbf{z}_{text}^{c'} = \mathbf{z}_{text}^{pro,c} + \mathbf{p}_t$ . For convenience, we use  $\mathbf{Z}_{text} \in \mathbb{R}^{C \times D}$  to represent the learnable text features corresponding to  $C$  class labels, where  $D$  is the feature dimension.

Similarly, while raw images contain rich semantics, they may also include irrelevant or misleading information that impedes alignment with textual features. Therefore, we also design a learnable image prompt  $\mathbf{p}_v$  to adjust the image representation. For each test image  $x_{test}$ , we generate  $N$  augmented views  $\{\mathcal{A}_n(x_{test})\}_{n=1}^N$  using the transformation family  $\mathcal{A}$ , and extract the corresponding features  $\{\mathbf{z}_{img}^n\}_{n=1}^N$  using the visual encoder  $V(\cdot)$ , where  $\mathbf{z}_{img}^n = V(\mathcal{A}_n(x_{test}))$ .

The adjusted image representation for each view is then computed as  $\mathbf{z}_{img}^{n'} = \mathbf{z}_{img}^n + \mathbf{p}_v$ , and we denote the collection of adjusted image features as  $\mathbf{Z}_{img} \in \mathbb{R}^{N \times D}$ .

### Dynamic Retrieval-Augmented Modulation

Retrieval-Augmented Generation (RAG) (Gao et al. 2023) mitigates the limitations of outdated or incomplete model knowledge by retrieving relevant information from external sources and incorporating it into the query. Inspired by this, we propose a dynamic retrieval-augmented modulation module to suppress low-entropy yet incorrect predictions.

Let  $\hat{\mathbf{Z}}_{text}$  and  $\hat{\mathbf{Z}}_{img}$  denote the normalized textual and visual feature matrices, respectively. The similarity logits  $\mathbf{L} = \exp(s) \cdot (\hat{\mathbf{Z}}_{img} \hat{\mathbf{Z}}_{text}^T)$ , where  $s$  is a logit scaling factor and  $\mathbf{L} \in \mathbb{R}^{N \times C}$ . To filter out unreliable predictions, we apply entropy-based confidence selection over the logits and obtain a subset of high-confidence predictions:

$$\mathbf{L}_r = \{\mathbf{L}_i \mid i \in \{1, \dots, N\} \wedge \mathcal{H}(\text{softmax}(\mathbf{L}_i)) \leq \tau\}, \quad (4)$$

where  $\mathbf{L}_i = [L_i^1, \dots, L_i^C] \in \mathbb{R}^{1 \times C}$  denotes the similarity logit for the  $i$ -th augmented view, and  $\text{softmax}(\mathbf{L}_i) = \frac{\exp(L_i^c)}{\sum_{c=1}^C \exp(L_i^c)}$ .

Based on the filtered confident predictions, we construct a dynamic knowledge base that stores test image features, pseudo labels, and corresponding entropy values. The predicted labels are determined as:

$$Y_r = \{\hat{y}_i \mid \mathbf{L}_i \in \mathbf{L}_r \wedge \hat{y}_i = \arg \max_c(\text{softmax}(\mathbf{L}_i))\}, \quad (5)$$

where  $\hat{y}_i$  denotes the predicted label corresponding to  $\mathbf{L}_i$ .

---

**Algorithm 1: D<sup>2</sup>TPT**


---

**Input:** Test image  $x_{test}$ , a set of categories set  $\{1, \dots, C\}$ , visual encoder  $V(\cdot)$ , textual encoder  $T(\cdot)$ , learnable prompts  $\mathbf{p}_v, \mathbf{p}_t$ , entropy threshold  $\tau$ , register capacity  $K$ , hyperparameters  $\alpha, \beta$

**Output:** Predicted category  $c$

- 1: Adapt visual and textual features with learnable prompts:  $\mathbf{z}_{img}^n = \mathbf{z}_{img}^n + \mathbf{p}_v, \mathbf{z}_{text}^c = \mathbf{z}_{text}^{pro,c} + \mathbf{p}_t$ ;
  - 2: Compute similarity logits using normalized vectors:  $\mathbf{L} = \exp(s) \cdot (\hat{\mathbf{Z}}_{img}(\hat{\mathbf{Z}}_{text})^T)$ ;
  - 3: Identify high-confidence predictions and pseudo labels using Equations (4)–(5);
  - 4: Update class-specific register using visual feature  $\mathbf{z}_{img}$ , pseudo label  $\tilde{y}$ , and corresponding entropy  $\mathcal{H}_{\tilde{y}}$  based on Equation (6);
  - 5: Construct key-value pairs  $\mathbf{K}, \mathbf{V}$  from register features and one-hot labels, and compute retrieval-based logits:  $\mathbf{L}_R = \lambda \exp(-\gamma(1 - \mathbf{z}_{img}\mathbf{K}^T))\mathbf{V}$ ;
  - 6: Fuse  $\mathbf{L}$  and  $\mathbf{L}_R$ , and compute the corresponding entropy loss as per Equation (7);
  - 7: Estimate confidence scores for selected images to perform weighted prediction fusion, and compute entropy loss of ensemble logits using Equation (8);
  - 8: Derive distillation logits from both original and adapted features, and compute the corresponding entropy loss via Equation (9);
  - 9: Optimize prompts by minimizing total loss defined in Equation (10) to produce final prediction  $c$ .
- 

We identify the most frequently predicted label  $\tilde{y} = \arg \max_{y \in \{1 \dots C\}} \sum_{i=1}^{|Y_r|} \mathbb{I}(\hat{y}_i = y)$  as the pseudo label, where  $|Y_r|$  denotes the number of predictions in  $Y_r$ . To estimate its confidence, we select the minimum entropy among predictions assigned to  $\tilde{y}$ , computed as  $\mathcal{H}_{\tilde{y}} = \min_{i \in I_{\tilde{y}}} \{\mathcal{H}(\text{softmax}(\mathbf{L}_i))\}$ , where  $I_{\tilde{y}} = \{i \mid \hat{y}_i = \tilde{y}\}$ .

To maintain a high-quality class-wise knowledge base, we adopt a dynamic updating strategy. Given the visual feature  $\mathbf{z}_{img}$  of the test image  $x_{test}$ , the pseudo label  $\tilde{y}$ , and the corresponding minimum entropy  $\mathcal{H}_{\tilde{y}}$ , we construct a tuple  $\mathbf{x} = (\mathbf{z}_{img}, \mathcal{H}_{\tilde{y}})$  for possible inclusion in the register  $\mathcal{R}_{\tilde{y}}$  associated with the pseudo label  $\tilde{y}$  in the knowledge base. The register is updated as follows:

$$\mathcal{R}_{\tilde{y}} = \begin{cases} \mathcal{R}_{\tilde{y}} \cup \{\mathbf{x}\}, & \text{if } 0 < |\mathcal{R}_{\tilde{y}}| < K \\ (\mathcal{R}_{\tilde{y}} \setminus \{\mathbf{x}^m\}) \cup \{\mathbf{x}\}, & \text{if } |\mathcal{R}_{\tilde{y}}| = K \wedge \mathcal{H}_{\tilde{y}} < \mathcal{H}^m, \\ \{\mathbf{x}\}, & \text{if } |\mathcal{R}_{\tilde{y}}| = 0 \end{cases} \quad (6)$$

where  $\mathcal{H}^m = \max\{\mathcal{H}_j \mid \mathcal{H}_j \in \mathcal{R}_{\tilde{y}}\}$ ,  $K$  denotes the maximum capacity of  $\mathcal{R}_{\tilde{y}}$ , and  $\mathbf{x}^m$  is the tuple corresponding to  $\mathcal{H}^m$ . The register is sorted by ascending entropy to retain high-confidence samples for reliable test-time retrieval.

During retrieval, key-value pairs are constructed by aggregating the stored visual features for each class. Specifically, the key for class  $\tilde{y} \in \{1, \dots, C\}$  is computed as the prototype of visual features in its register, i.e.,  $\mathbf{k}_{\tilde{y}} = \frac{1}{|\mathcal{R}_{\tilde{y}}|} \sum_{\mathbf{z}_i \in \mathcal{R}_{\tilde{y}}} \mathbf{z}_i \in \mathbb{R}^{1 \times D}$ , and the corresponding value is

Dataset	#Classes	#Test
ImageNet (Deng et al. 2009)	1,000	50,000
ImageNet-A (Hendrycks et al. 2021b)	200	7,500
ImageNet-R (Hendrycks et al. 2021a)	200	30,000
ImageNet-Sketch (Wang et al. 2019)	1,000	50,889
ImageNet-V2 (Recht et al. 2019)	1,000	10,000
Caltech101 (Fei-Fei, Fergus, and Perona 2004)	100	2,465
OxfordPets (Parkhi et al. 2012)	37	3,669
StanfordCars (Krause et al. 2013)	196	8,041
Flowers102 (Nilsback and Zisserman 2008)	102	2,463
Food101 (Bossard, Guillaumin, and Gool 2014)	101	30,300
Aircraft (Maji et al. 2013)	100	3,333
SUN397 (Xiao et al. 2010)	397	19,850
DTD (Cimpoi et al. 2014)	47	1,692
EuroSAT (Helber et al. 2019)	10	8,100
UCF101 (Soomro, Zamir, and Shah 2012)	101	3,783

Table 1: Overview of datasets used in experiments, including the number of classes and test data.

a one-hot vector  $\mathbf{v}_{\tilde{y}} = \text{one\_hot}(\tilde{y}) \in \mathbb{R}^{1 \times C}$ , with 1 at the  $\tilde{y}$ -th position and 0 elsewhere. Consequently, the key matrix  $\mathbf{K} = [\mathbf{k}_1; \dots; \mathbf{k}_{\check{C}}] \in \mathbb{R}^{\check{C} \times D}$  and value matrix  $\mathbf{V} = [\mathbf{v}_1; \dots; \mathbf{v}_{\check{C}}] \in \mathbb{R}^{\check{C} \times C}$  are formed, where  $\check{C}$  denotes the number of classes stored in the knowledge base.

Inspired by (Zhang et al. 2022), we compute the retrieval-augmented modulation logits for the test image  $x_{test}$  as  $\mathbf{L}_{RAM} = \mathbf{L} + \mathbf{L}_R$ , where  $\mathbf{L}_R = \lambda \exp(-\gamma(1 - \mathbf{z}_{img}\mathbf{K}^T))\mathbf{V}$  represents the retrieval-based logits,  $\lambda$  is the residual ratio, and  $\gamma$  is the sharpness ratio. We then use Equation (4) to select high-confidence logits from  $\mathbf{L}_{RAM}$ , forming  $\mathbf{L}_{RAM}^r \in \mathbb{R}^{M \times C}$ , where  $M$  denotes the number of selected logits. Therefore, the self-entropy loss is computed as:

$$\mathcal{L}_{RAM} = \mathcal{H}\left(\frac{1}{M} \sum_{m=1}^M \text{softmax}(\mathbf{L}_{RAM}^{r,m})\right), \quad (7)$$

where  $\mathbf{L}_{RAM}^{r,m} \in \mathbb{R}^{1 \times C}$  is the  $m$ -th row of  $\mathbf{L}_{RAM}^r$ .

### Reliability-Aware Prompt Optimization

Given  $\mathbf{L}_{RAM}^r \in \mathbb{R}^{M \times C}$ , we obtain the corresponding learnable image feature vectors  $\mathbf{Z}_{img}^r \in \mathbb{R}^{M \times D}$ . To further mitigate the impact of low-quality augmented views, we adopt a confidence-based weighted ensemble strategy to aggregate predictions from different augmented views. Specifically, we first compute the pairwise cosine similarity matrix  $\mathbf{S} = \hat{\mathbf{Z}}_{img}^r (\hat{\mathbf{Z}}_{img}^r)^T \in \mathbb{R}^{M \times M}$ , where  $\hat{\mathbf{Z}}_{img}^r$  denotes the normalized image features. The confidence score for each view is then obtained by summing its similarities with all other views, i.e.,  $\mathbf{s} = [s_1; \dots; s_M] \in \mathbb{R}^{M \times 1}$ , where  $s_i = \sum_{j=1}^M \mathbf{S}_{i,j}$ . Therefore, we obtain the final ensemble weights  $\mathbf{w} = \text{softmax}(\mathbf{s}) \in \mathbb{R}^{M \times 1}$ , which are subsequently used to compute the aggregated prediction  $\mathbf{L}_{EN} \in \mathbb{R}^{1 \times C}$ . Consequently, the corresponding self-entropy loss is computed as:

$$\mathcal{L}_{EN} = \mathcal{H}(\text{softmax}(\mathbf{L}_{EN})), \quad (8)$$

Method	Publication	ImageNet	ImageNet-A	ImageNet-R	ImageNet-S	ImageNet-V	Average	OOD Average
CLIP-ViT-B/16	<i>ICML2021</i>	66.73	47.87	73.98	46.09	60.86	59.11	57.20
TPT	<i>NeurIPS2022</i>	68.98	54.77	77.06	47.94	63.45	62.44	60.81
DiffTPT	<i>ICCV2023</i>	70.30	55.68	75.00	46.80	65.10	62.58	60.65
PromptAlign	<i>NeurIPS2023</i>	-	59.37	79.33	50.23	65.29	-	63.56
TDA	<i>CVPR2024</i>	69.51	60.11	80.24	50.54	64.67	65.01	63.89
SCP	<i>ACMMM2024</i>	68.80	50.50	78.70	46.50	62.60	61.42	59.58
AdaPrompt	<i>AAAI2024</i>	-	47.71	73.98	47.72	59.32	-	57.18
C-TPT	<i>ICLR 2024</i>	69.30	52.90	78.00	48.50	63.40	62.42	60.70
O-TPT	<i>CVPR2025</i>	67.33	49.87	72.55	47.12	61.65	59.70	57.80
DynaPrompt	<i>ICLR2025</i>	69.61	56.17	78.17	48.22	64.67	63.37	61.81
TPS	<i>WACV2025</i>	71.45	60.61	80.20	50.88	64.91	65.61	64.15
D <sup>2</sup> TPT	<i>This Paper</i>	<b>71.85</b>	<b>63.09</b>	<b>80.38</b>	<b>51.75</b>	<b>65.79</b>	<b>66.57</b>	<b>65.25</b>

Table 2: Comparison of top-1 accuracy (%) with baselines under natural distribution shifts. Best performances are in **bold**.

where  $\mathbf{L}_{EN} = \sum_{m=1}^M \mathbf{w}_m \cdot \mathbf{L}_{RAM}^{r,m}$ , and  $\mathbf{w}_m, \mathbf{L}_{RAM}^{r,m}$  represent the weight and prediction of the  $m$ -th augmented view, respectively.

To enhance the alignment between visual and textual modalities after incorporating learnable prompts, we propose a cross-modal consistency distillation strategy. Specifically, let  $\hat{\mathbf{Z}}_{text}^{ori} \in \mathbb{R}^{C \times D}$  and  $\hat{\mathbf{Z}}_{text} \in \mathbb{R}^{C \times D}$  denote the normalized text features before and after prompt adaptation, respectively. In addition, let  $\bar{\mathbf{Z}}_{img}^{r,ori} \in \mathbb{R}^{1 \times D}$  and  $\bar{\mathbf{Z}}_{img}^r \in \mathbb{R}^{1 \times D}$  represent the original and prompt-adapted image feature prototypes obtained after confidence selection. We define the cross-modal distillation logits as  $\mathbf{L}_{MD} \in \mathbb{R}^{1 \times C}$ , and the corresponding self-entropy loss is defined as:

$$\mathcal{L}_{MD} = \mathcal{H}(\text{softmax}(\mathbf{L}_{MD})). \quad (9)$$

where  $\mathbf{L}_{MD} = \mathbf{L}_{v \rightarrow t} + \mathbf{L}_{t \rightarrow v} + \mathbf{L}_{self}$ ,  $\mathbf{L}_{v \rightarrow t} = \exp(s) \cdot (\bar{\mathbf{Z}}_{img}^{r,ori} (\hat{\mathbf{Z}}_{text}^T)^T)$ ,  $\mathbf{L}_{t \rightarrow v} = \exp(s) \cdot (\bar{\mathbf{Z}}_{img}^r (\hat{\mathbf{Z}}_{text}^{ori})^T)$ , and  $\mathbf{L}_{self} = \exp(s) \cdot (\bar{\mathbf{Z}}_{img}^r (\hat{\mathbf{Z}}_{text}^T)^T)$ .

By combining Equations (7), (8), and (9), we derive the final optimization objective:

$$\mathbf{p}_t^*, \mathbf{p}_v^* = \arg \min_{\mathbf{p}_t, \mathbf{p}_v} \mathcal{L}, \quad (10)$$

where  $\mathcal{L} = \mathcal{L}_{RAM} + \alpha \mathcal{L}_{EN} + \beta \mathcal{L}_{MD}$ , and  $\alpha, \beta$  are hyperparameters that balance the influence of  $\mathcal{L}_{EN}$  and  $\mathcal{L}_{MD}$ , respectively. The procedural steps of D<sup>2</sup>TPT are detailed in Algorithm 1.

## Experiments

In this section, we present experimental results to comprehensively evaluate the effectiveness of D<sup>2</sup>TPT.

### Experimental Setup

**Datasets.** Following prior work (Shu et al. 2022), we evaluate our method under two key scenarios: (1) Natural distribution shifts. We use ImageNet (Deng et al. 2009) and four of its out-of-distribution variants, i.e., ImageNet-A (Hendrycks et al. 2021b), ImageNet-R (Hendrycks

et al. 2021a), ImageNet-Sketch (Wang et al. 2019), and ImageNet-V2 (Recht et al. 2019). (2) Cross-datasets generalization. We conduct evaluations across ten diverse classification datasets, including 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 Gool 2014), Aircraft (Maji et al. 2013), SUN397 (Xiao et al. 2010), DTD (Cimpoi et al. 2014), EuroSAT (Helber et al. 2019), and UCF101 (Soomro, Zamir, and Shah 2012). Table 1 presents the detailed statistics of these datasets.

**Baselines.** We compare our D<sup>2</sup>TPT against a comprehensive set of baselines, including both zero-shot and test-time adaptation approaches built upon the CLIP ViT-B/16 (Radford et al. 2021) backbone. Specifically, we consider the following methods: TPT (Shu et al. 2022), DiffTPT (Feng et al. 2023), PromptAlign (Samadh et al. 2023), TDA (Karmanov et al. 2024), SCP (Wang et al. 2024), AdaPrompt (Zhang, Zhou, and Li 2024), C-TPT (Yoon et al. 2024), O-TPT (Sharifdeen et al. 2025), DynaPrompt (Xiao et al. 2025), and TPS (Sui, Wang, and Yeung-Levy 2025). These baselines encompass a diverse range of strategies for adapting VLMs at test time, including prompt tuning, distribution alignment, and other robust adaptation techniques.

**Implementation Details.** We set the register capacity  $K$  for each class in the knowledge base to 3. Following TPT (Shu et al. 2022), each test image is augmented 63 times via random resized cropping and grouped with the original into a batch of 64. Among the predictions, we select the top 10% ( $\tau = 0.1$ ) most confident samples and compute the entropy of their averaged probability. The learnable prompts are then updated for one step by minimizing this entropy using the AdamW optimizer. All experiments are conducted on a single NVIDIA GeForce RTX 4090, and results are averaged over three random seeds.

### Main Results

**Natural Distribution Shifts.** Table 2 presents a comparison of top-1 accuracy across various test datasets under natural distribution shifts. Our proposed method, D<sup>2</sup>TPT, con-

Method	Publication	Caltech101	OxfordPets	StanfordCars	Flowers102	Food101	Aircraft	SUN397	DTD	EuroSAT	UCF101	Average
CLIP-ViT-B/16	<i>ICML2021</i>	93.35	88.25	65.48	67.44	83.65	23.67	62.59	44.27	42.01	65.13	63.58
TPT	<i>NeurIPS2022</i>	94.16	87.79	66.87	68.98	84.67	24.78	65.50	47.75	42.44	68.04	65.10
DiffTPT	<i>ICCV2023</i>	92.49	88.22	67.01	70.10	87.23	25.60	65.74	47.00	43.13	68.22	65.47
PromptAlign	<i>NeurIPS2023</i>	94.01	<b>90.76</b>	68.50	72.39	86.65	24.80	67.54	47.24	47.86	69.47	66.92
TDA	<i>CVPR2024</i>	94.24	88.63	67.28	71.42	86.14	23.91	67.62	47.40	<b>58.00</b>	70.66	67.53
SCP	<i>ACMMM2024</i>	93.90	88.60	65.90	70.00	<b>87.40</b>	24.80	69.10	43.90	47.30	67.80	65.87
AdaPrompt	<i>AAAI2024</i>	94.07	89.64	63.29	72.97	84.72	21.21	65.37	44.75	47.20	67.22	65.04
C-TPT	<i>ICLR 2024</i>	94.10	87.40	66.70	69.90	84.50	23.90	66.00	46.80	48.70	66.70	65.47
O-TPT	<i>CVPR2025</i>	93.95	87.95	64.53	70.07	84.13	23.64	64.23	45.68	42.84	64.16	64.12
DynaPrompt	<i>ICLR2025</i>	94.32	88.28	67.65	69.95	85.42	24.33	66.32	47.96	42.28	68.72	65.52
TPS	<i>WACV2025</i>	95.09	87.35	69.06	71.54	85.23	26.34	68.98	50.47	44.48	71.00	66.95
<b>D<sup>2</sup>TPT</b>	<i>This Paper</i>	<b>95.10</b>	87.55	<b>69.50</b>	<b>75.21</b>	84.95	<b>27.00</b>	<b>69.80</b>	<b>52.88</b>	54.92	<b>72.40</b>	<b>68.93</b>

Table 3: Comparison of top-1 accuracy (%) with baselines under cross-datasets generalization. Best performances are in **bold**.

RAM	RPO		NDS	CDG
	CWE	CMD		
×	×	×	64.57	67.08
✓	×	×	64.82	68.27
✓	✓	×	65.02	68.43
✓	✓	✓	65.25	68.93

Table 4: Ablation study on the effectiveness of each module. NDS and CDG are abbreviations for natural distribution shifts and cross-datasets generalization, respectively.

sistently outperforms all baselines. Specifically, it achieves the highest average accuracy of 66.57%, outperforming all competitors and demonstrating superior generalization capability. In terms of OOD average accuracy, D<sup>2</sup>TPT reaches 65.25%, surpassing the primary baseline TPS by 1.1%. These results validate the effectiveness of our D<sup>2</sup>TPT in handling distribution shifts across diverse visual domains.

**Cross-Datasets Generalization.** Table 3 reports the top-1 accuracy on 10 diverse datasets for evaluating cross-datasets generalization. Our method, D<sup>2</sup>TPT, achieves the highest average accuracy of 68.93%, outperforming all previous baselines. In particular, D<sup>2</sup>TPT obtains the best results on 7 out of 10 datasets. Compared to the primary baseline TPS, D<sup>2</sup>TPT brings a performance gain of 1.98% on average. These results clearly demonstrate the generalization capability of D<sup>2</sup>TPT across diverse visual domains with varying distributions and label semantics.

**Ablation Study.** We assess the impact of removing two modules, i.e., dynamic retrieval-augmented modulation (RAM) and reliability-aware prompt optimization (RPO), on model performance. CWE (confidence-based weighted ensemble) and CMD (cross-modal consistency distillation) are subcomponents of the RPO module. As shown in Ta-




 schooner	<b>Model</b>	<b>Entropy</b>	<b>Confidence</b>	<b>Pred.</b>
	CLIP	0.72	0.54	schooner 😊
	TPT	0.45	0.84	ketch 😞
	<b>Ours</b>	0.12	0.98	schooner 😊
 petunia	<b>Model</b>	<b>Entropy</b>	<b>Confidence</b>	<b>Pred.</b>
	CLIP	1.47	0.40	petunia 😊
	TPT	0.50	0.91	garden phlox 😞
	<b>Ours</b>	0.09	0.98	petunia 😊
 fibrous	<b>Model</b>	<b>Entropy</b>	<b>Confidence</b>	<b>Pred.</b>
	CLIP	2.64	0.30	fibrous 😊
	TPT	1.38	0.61	porous 😞
	<b>Ours</b>	0.41	0.91	fibrous 😊

Figure 4: Case study on prediction confidence and correctness. We use pred. to denote the model’s predicted category.

ble 4, the experimental results highlight the individual contributions of each module. Specifically, removing either the RAM or RPO module leads to a performance drop, while the model achieves the best performance when both modules are used jointly. This confirms the effectiveness of each component in our proposed approach.

## Further Analysis

**Case Study.** Figure 4 presents representative examples where TPT tends to produce low-entropy yet incorrect predictions, reflecting overconfident misclassification. In contrast, our method effectively suppresses such overconfident errors, which can be attributed to the design strategy that modulates the original model predictions during prompt optimization. These examples suggest that our approach mitigates the prompt optimization bias, enabling more confident and accurate predictions.

	NDS	CDG
D <sup>2</sup> TPT vs. TPS	0.046	0.032
D <sup>2</sup> TPT vs. DynaPrompt	0.033	0.021
D <sup>2</sup> TPT vs. O-TPT	0.016	0.003

Table 5: The  $p$ -value for Student’s  $t$ -test across two settings.




Figure 5: Comparison of normalized cross-modal feature distances across different models.

**Significance Test.** To verify that the observed performance improvements are not attributable to random chance, a Student’s  $t$ -test (Mendenhall, Beaver, and Beaver 2020) is conducted between the baseline models and the proposed D<sup>2</sup>TPT. A  $p$ -value below 0.05 is considered statistically significant. As reported in Table 5, all  $p$ -values across the two task settings, i.e., natural distribution shifts and cross-datasets generalization, fall below this threshold. Specifically, the comparisons between TPS and D<sup>2</sup>TPT yield  $p$ -values of 0.046 and 0.032, respectively. These results indicate that D<sup>2</sup>TPT achieves statistically significant improvements over the baselines, further validating the effectiveness of the proposed method.

**Computational Complexity.** To further assess the trade-off between accuracy and computational efficiency, we conduct a complexity analysis under the cross-datasets generalization setting, with results summarized in Table 6. On the Flowers102 dataset, D<sup>2</sup>TPT achieves a Top-1 accuracy of 75.21% with 52736 trainable parameters and an inference throughput of 5.71 FPS. These results are comparable to those of the baseline models. Although D<sup>2</sup>TPT exhibits a slightly lower inference throughput, its relatively low parameter count results in only a minimal increase in computational complexity, demonstrating that the improvements in accuracy come with negligible computational overhead.

**Cross-Modal Alignment Evaluation.** We compute the L2 distances between image and text features to further assess the effectiveness of our method in aligning visual and textual modalities. To facilitate clearer comparisons among models within each dataset, the distances are normalized as proportions. The normalized results are shown in Figure 5, where smaller proportions indicate shorter distances and thus better cross-modal alignment. We observe that our method consistently achieves the lowest proportion of feature distance, outperforming both CLIP and PromptAlign,

Method	Top-1 Accuracy	Params	FPS
TPT	68.98	2048	6.37
PromptAlign	72.39	1185024	7.41
TPS	71.54	52224	6.02
D <sup>2</sup> TPT	75.21	52736	5.71

Table 6: Complexity comparison of D<sup>2</sup>TPT and baselines on the Flowers102 dataset.




Figure 6: Influence of  $\alpha$  and  $\beta$  on model accuracy.

which indicates a stronger alignment between modalities.

**Hyperparameter Analysis.** To systematically investigate the impact of the hyperparameters  $\alpha$  and  $\beta$  on model performance, we conduct an empirical study under the cross-datasets generalization setting. Specifically, we search for the optimal value of  $\alpha$  from the set  $\{1e^0, 1e^{-1}, 1e^{-2}, 1e^{-3}\}$  and  $\beta$  from  $\{1e^{-1}, 1e^{-2}, 1e^{-3}, 1e^{-4}\}$ , using four representative datasets: Caltech101, Flowers102, UCF101, and Aircraft. Figure 6 presents the average performance across the four datasets under different settings of the two hyperparameters. Based on the results, we select  $\alpha = 1e^{-1}$  and  $\beta = 1e^{-3}$  as the final configuration for our model.

## Conclusion

In this work, we propose Doubly Debiased Test-Time Prompt Tuning (D<sup>2</sup>TPT) to mitigate prompt optimization bias from both model and data perspectives. By incorporating a dynamic retrieval-augmented modulation module and a reliability-aware prompt optimization module, D<sup>2</sup>TPT enables more confident and accurate predictions, along with stronger alignment between visual and textual modalities. These improvements jointly alleviate prompt optimization bias and enhance the performance of test-time prompt tuning. Extensive experiments demonstrate that D<sup>2</sup>TPT consistently outperforms existing baselines on benchmarks involving natural distribution shifts and cross-datasets generalization, validating its effectiveness and robustness.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China No. 62406313, 2023 Special Research Assistant Grant Project of the Chinese Academy of Sciences.

## References

- Bossard, L.; Guillaumin, M.; and Gool, L. V. 2014. Food-101 - Mining Discriminative Components with Random Forests. In Fleet, D. J.; Pajdla, T.; Schiele, B.; and Tuytelaars, T., eds., *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VI*, volume 8694 of *Lecture Notes in Computer Science*, 446–461. Springer.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing Textures in the Wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, 3606–3613. IEEE Computer Society.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, 248–255. IEEE Computer Society.
- Derakhshani, M. M.; Sanchez, E.; Bulat, A.; da Costa, V. G. T.; Snoek, C. G. M.; Tzimiropoulos, G.; and Martínez, B. 2023. Bayesian Prompt Learning for Image-Language Model Generalization. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 15191–15200. IEEE.
- 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 *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2004, Washington, DC, USA, June 27 - July 2, 2004*, 178. IEEE Computer Society.
- Feng, C.; Yu, K.; Liu, Y.; Khan, S.; and Zuo, W. 2023. Diverse Data Augmentation with Diffusions for Effective Test-time Prompt Tuning. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 2704–2714. IEEE.
- Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; Guo, Q.; Wang, M.; and Wang, H. 2023. Retrieval-Augmented Generation for Large Language Models: A Survey. *CoRR*, abs/2312.10997.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 12(7): 2217–2226.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; Song, D.; Steinhardt, J.; and Gilmer, J. 2021a. The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, 8320–8329. IEEE.
- Hendrycks, D.; Zhao, K.; Basart, S.; Steinhardt, J.; and Song, D. 2021b. Natural Adversarial Examples. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 15262–15271. Computer Vision Foundation / IEEE.
- Karmanov, A.; Guan, D.; Lu, S.; El-Saddik, A.; and Xing, E. P. 2024. Efficient Test-Time Adaptation of Vision-Language Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, 14162–14171. IEEE.
- Khattak, M. U.; Rasheed, H. A.; Maaz, M.; Khan, S. H.; and Khan, F. S. 2023. MaPLe: Multi-modal Prompt Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, 19113–19122. IEEE.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3D Object Representations for Fine-Grained Categorization. In *2013 IEEE International Conference on Computer Vision Workshops, ICCV Workshops 2013, Sydney, Australia, December 1-8, 2013*, 554–561. IEEE Computer Society.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In Moens, M.; Huang, X.; Specia, L.; and Yih, S. W., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, 3045–3059. Association for Computational Linguistics.
- Li, J.; Mo, W.; Song, F.; Sun, C.; Qiang, W.; Su, B.; and Zheng, C. 2025. Supporting vision-language model few-shot inference with confounder-pruned knowledge prompt. *Neural Networks*, 185: 107173.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, 4582–4597. Association for Computational Linguistics.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M. B.; and Vedaldi, A. 2013. Fine-Grained Visual Classification of Aircraft. *CoRR*, abs/1306.5151.
- Mendenhall, W.; Beaver, R. J.; and Beaver, B. M. 2020. *Introduction to probability and statistics*. Cengage.
- Meng, F.; Cui, C.; Dai, H.; and Gong, S. 2025. Black-Box Test-Time Prompt Tuning for Vision-Language Models. In Walsh, T.; Shah, J.; and Kolter, Z., eds., *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, 6099–6107. AAAI Press.
- Nilsback, M.; and Zisserman, A. 2008. Automated Flower Classification over a Large Number of Classes. In *Sixth Indian Conference on Computer Vision, Graphics & Image Processing, ICVGIP 2008, Bhubaneswar, India, 16-19 December 2008*, 722–729. IEEE Computer Society.
- OpenAI. 2023. GPT-4 Technical Report. *CoRR*, abs/2303.08774.
- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. V. 2012. Cats and dogs. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012*, 3498–3505. IEEE Computer Society.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 8748–8763. PMLR.
- Recht, B.; Roelofs, R.; Schmidt, L.; and Shankar, V. 2019. Do ImageNet Classifiers Generalize to ImageNet? In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 5389–5400. PMLR.
- Roy, S.; and Etemad, A. 2024. Consistency-guided Prompt Learning for Vision-Language Models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Samadh, J. A.; Gani, H.; Hussein, N.; Khattak, M. U.; Naseer, M.; Khan, F. S.; and Khan, S. H. 2023. Align Your Prompts: Test-Time Prompting with Distribution Alignment for Zero-Shot Generalization. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Sharifdeen, A.; Munir, M. A.; Baliah, S.; Khan, S.; and Khan, M. H. 2025. O-TPT: Orthogonality Constraints for Calibrating Test-time Prompt Tuning in Vision-Language Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, 19942–19951. Computer Vision Foundation / IEEE.
- Shu, M.; Nie, W.; Huang, D.; Yu, Z.; Goldstein, T.; Anandkumar, A.; and Xiao, C. 2022. Test-Time Prompt Tuning for Zero-Shot Generalization in Vision-Language Models. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *CoRR*, abs/1212.0402.
- Sui, E.; Wang, X.; and Yeung-Levy, S. 2025. Just Shift It: Test-Time Prototype Shifting for Zero-Shot Generalization with Vision-Language Models. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2025, Tucson, AZ, USA, February 26 - March 6, 2025*, 825–835. IEEE.
- Wang, H.; Ge, S.; Lipton, Z. C.; and Xing, E. P. 2019. Learning Robust Global Representations by Penalizing Local Predictive Power. In Wallach, H. M.; Larochelle, H.; Beygelzimer, A.; d’Alché-Buc, F.; Fox, E. B.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 10506–10518.
- Wang, R.; Zuo, H.; Fang, Z.; and Lu, J. 2024. Towards Robustness Prompt Tuning with Fully Test-Time Adaptation for CLIP’s Zero-Shot Generalization. In Cai, J.; Kankanhalli, M. S.; Prabhakaran, B.; Boll, S.; Subramanian, R.; Zheng, L.; Singh, V. K.; César, P.; Xie, L.; and Xu, D., eds., *Proceedings of the 32nd ACM International Conference on Multimedia, MM 2024, Melbourne, VIC, Australia, 28 October 2024 - 1 November 2024*, 8604–8612. ACM.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. SUN database: Large-scale scene recognition from abbey to zoo. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, 3485–3492. IEEE Computer Society.
- Xiao, Z.; Yan, S.; Hong, J.; Cai, J.; Jiang, X.; Hu, Y.; Shen, J.; Wang, C.; and Snoek, C. G. M. 2025. DynaPrompt: Dynamic Test-Time Prompt Tuning. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Yoon, H. S.; Yoon, E.; Tee, J. T. J.; Hasegawa-Johnson, M. A.; Li, Y.; and Yoo, C. D. 2024. C-TPT: Calibrated Test-Time Prompt Tuning for Vision-Language Models via Text Feature Dispersion. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Zhang, C.; Stepputtis, S.; Sycara, K. P.; and Xie, Y. 2024. Dual Prototype Evolving for Test-Time Generalization of Vision-Language Models. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J. M.; and Zhang, C., eds., *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Zhang, D.; Zhou, Z.; and Li, Y. 2024. Robust Test-Time Adaptation for Zero-Shot Prompt Tuning. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024, Vancouver, Canada*, 16714–16722. AAAI Press.
- Zhang, R.; Zhang, W.; Fang, R.; Gao, P.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2022. Tip-Adapter: Training-Free Adaptation of CLIP for Few-Shot Classification. In Avidan, S.; Brostow, G. J.; Cissé, M.; Farinella, G. M.; and Hassner, T., eds., *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXV*, volume 13695 of *Lecture Notes in Computer Science*, 493–510. Springer.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to Prompt for Vision-Language Models. *Int. J. Comput. Vis.*, 130(9): 2337–2348.