

# Text and Image Are Mutually Beneficial: Enhancing Training-Free Few-Shot Classification with CLIP

Yayuan Li<sup>1</sup>, Jintao Guo<sup>1</sup>, Lei Qi<sup>2</sup>, Wenbin Li<sup>1</sup>, Yinghuan Shi<sup>1\*</sup>

<sup>1</sup>National Key Laboratory for Novel Software Technology, Nanjing University, China

<sup>2</sup>School of Computer Science and Engineering, Southeast University, China

{liyayuan, guojintao}@smail.nju.edu.cn, qilei@seu.edu.cn, {liwenbin, syh}@smail.nju.edu.cn

## Abstract

Contrastive Language-Image Pretraining (CLIP) has been widely used in vision tasks. Notably, CLIP has demonstrated promising performance in few-shot learning (FSL). However, existing CLIP-based methods in training-free FSL (*i.e.*, without the requirement of additional training) mainly learn different modalities independently, leading to two essential issues: 1) severe anomalous match in image modality; 2) varying quality of generated text prompts. To address these issues, we build a mutual guidance mechanism, that introduces an Image-Guided-Text (IGT) component to rectify varying quality of text prompts through image representations, and a Text-Guided-Image (TGI) component to mitigate the anomalous match of image modality through text representations. By integrating IGT and TGI, we adopt a perspective of Text-Image Mutual guidance Optimization, proposing TIMO. Extensive experiments show that TIMO significantly outperforms the state-of-the-art (SOTA) training-free method. Additionally, by exploring the extent of mutual guidance, we propose an enhanced variant, TIMO-S, which even surpasses the best training-required methods by 0.33% with approximately  $\times 100$  less time cost.

**Code** — <https://github.com/lyymuwu/TIMO>

**Extended version** — <https://arxiv.org/abs/2412.11375>

## Introduction

Given its strong zero-shot capability, Contrastive Language-Image Pretraining (CLIP) (Radford et al. 2021) has recently been adopted for various downstream tasks (Shao et al. 2024; Tang et al. 2024; Martin et al. 2024; Zhou et al. 2023). Among these, CLIP-based few-shot learning (FSL) has garnered significant attention due to its ability to quickly adapt to downstream tasks with few training samples (Wang et al. 2020; Sanghi et al. 2022; Zhao et al. 2024; Mu et al. 2022; Wang et al. 2024a). However, as models and datasets grow, traditional training-required FSL methods face significant time costs due to backpropagation and are prone to overfitting, diminishing the rapid adaptability of FSL. Consequently, training-free FSL methods (Zhang et al. 2022; Zhu

\*Corresponding author.

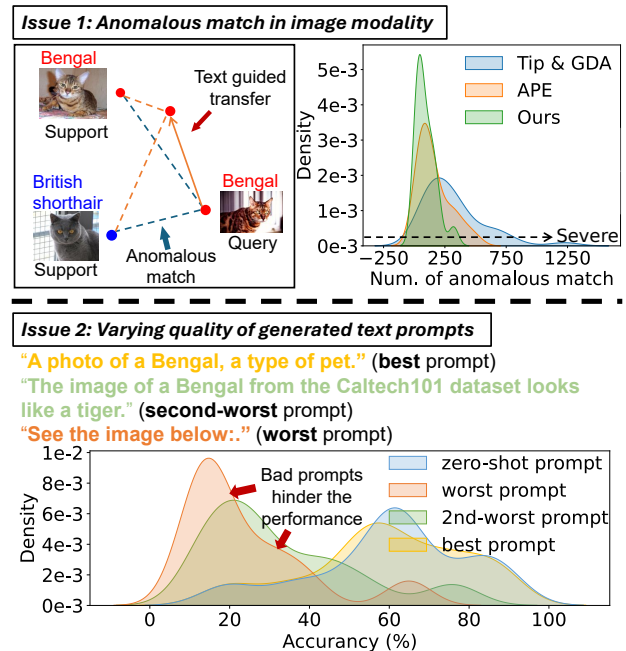


Figure 1: Defects of independent modelling. The top of the figure illustrates how the anomalous match left from the CLIP pretraining stage results in high similarity between images from distinct categories; The bottom shows the varying quality of the automatically generated prompts, where poor prompts can negatively impact the final performance.

et al. 2023; Wang et al. 2024b) have emerged, aiming to improve model performance while avoiding iterative optimization methods like gradient descent to enhance efficiency.

Despite the promising performance of previous state-of-the-art (SOTA) training-free FSL methods, they generally model different modalities separately, ignoring the importance of effective cross-modal interaction (*e.g.*, modality alignment). This is particularly challenging to achieve in training-free settings, as we cannot update model parameters as in conventional methods. Consequently, previous related works either only focus on designing sophisticated architectures (Wang et al. 2024b; Zhang et al. 2022; Udandarao, Gupta, and Albanie 2023; Zhang et al. 2023), or only con-

centrate designing prompts to make them more adaptable to image classification tasks (Zhou et al. 2022b; Lu et al. 2022; Pratt et al. 2023; Roth et al. 2023). This widespread lack of inter-modal interaction leads to previous SOTA models failing to capture the interrelation from various sources, resulting in suboptimal classifiers (Gong 2017; Tan et al. 2021). This observation deeply motivates us to investigate the consequences of inadequate cross-modal interaction and explore potential solutions for the following two bottlenecks.

**1) Severe anomalous match in image modality.** Prior works have highlighted that different modalities possess inherent strengths and weaknesses (Fan et al. 2023; Zhang et al. 2024a). Notably, we found that the information contained in a single image is significantly more abundant than that in text, which makes the CLIP’s inherent flaws (caused at the pretraining stage) of the image modality more pronounced. Besides, this issue is also particularly pronounced in FSL with scarce training samples, as it is often masked by the large training datasets in fully supervised settings. For instance, limited support images tend to lead to high confidence in incorrect predictions, which is a specific manifestation of the inherent weakness. This issue was observed in our task, as shown in Fig. 1, where a 1-shot classification mispredicts a *Bengal* as a *British Shorthair*. To investigate further, we conducted a statistical analysis of the *OxfordPets* dataset. We define an anomalous match as a high degree of similarity between the prototype feature of one category and the image feature of another category. Our findings reveal that anomalous matches are prevalent, as illustrated in the density plot at the top of Fig. 1, where the prototype features of most categories have hundreds of anomalous matches. This issue is particularly severe for methods like Tip-Adapter (Zhang et al. 2022) and GDA-CLIP (Wang et al. 2024b), which have not considered inter-modal alignment. In contrast, APE (Zhu et al. 2023) mitigates this problem to some extent through dimensionality reduction. Ultimately, this phenomenon significantly increases the likelihood of misclassification in subsequent tasks. Here, we raise a question: *Could the text be well utilized to mitigate the anomalous match in image modality?*

**2) Varying quality of generated prompts.** Previous training-free methods produce prompts by independently generating descriptive and structured prompts (Roth et al. 2023). However, we found that these independently modelled prompts exhibit quality variations, leading to a larger semantic gap between text and image modalities and suboptimal classifiers. To quantify this issue, we use the similarity between prompts and image prototypical features to evaluate the prompt quality, identifying the prompt with the lowest similarity as the “worst”. Taking CuPL (Pratt et al. 2023) as an example, the bottom density plot of Fig. 1 shows that the worst prompts selected from the entire set perform significantly worse than the others. Notably, the best-selected prompt is weaker than the zero-shot prompt obtained by averaging all prompts, indicating that classification performance improves with the inclusion of more prompts. This demonstrates that suboptimal prompts often arise because multiple independently obtained prompts in the training-free condition may not all suit downstream tasks. Without the

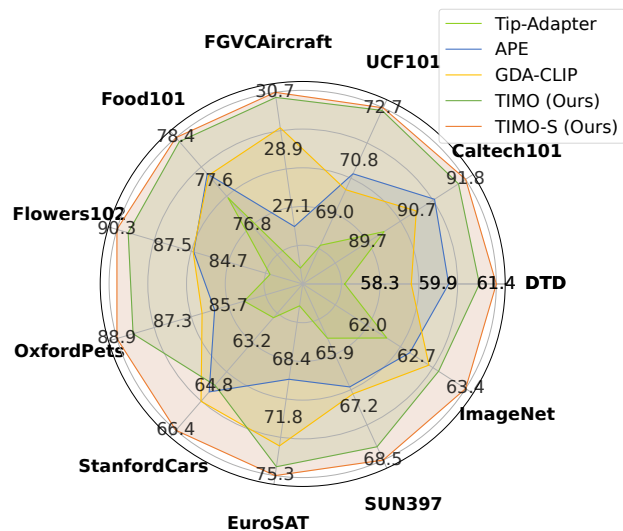


Figure 2: Comparison of the average Top-1 accuracy of various FSL methods across 11 datasets. TIMO-S outperforms the previous best method by an average of 1.76%.

guidance of image modality, suboptimal prompts limit the final classification performance. Here, we raise a question: *Could the image be well utilized to mitigate the varying quality of text generated prompts?*

To jointly solve the 2 aforementioned issues, we explore the explicitly mutual guidance of images and text. To address the first issue (*i.e.*, anomalous match), we observe that textual representation is generally more robust than image representation and that the knowledge from the two modalities is usually complementary. We consider incorporating textual features into the construction of the image classifier to alleviate the anomalous match of image modality. To address the second issue (*i.e.*, varying quality of prompts), we find that the suitability of prompts for specific tasks can be effectively assessed using image representations. Thus, we can use the image modality to guide a linear combination among prompts, enhancing the knowledge of text modality. Ultimately, we propose **Text-Image Mutual guidance Optimization (TIMO)** to address these two issues together.

Our method can integrate with previous training-free methods, enhancing modality interaction. Combining Text-Guided-Image (TGI) and Image-Guided-Text (IGT) with Tip-Adapter (Tip-MG) improves its performance by 1.20% across 11 datasets. Combining them with GDA-CLIP (Wang et al. 2024b) (TIMO) achieves a 1.32% improvement. Our enhanced variant, TIMO-S, outperforms the second-best training-free method by 1.76%. It even surpasses the best training-required method by 0.33% while reducing time cost by approximately  $\times 100$ . Performance comparisons of different training-free methods in Fig. 2 show our method achieves the best results across all datasets. The research contributions of this paper are summarized as follows:

1) We recognize that independently modelling different modalities leads to severe anomalous match and varying prompt quality.

2) From the perspective of text-guided-image, we employ an instance-based transfer method to mitigate the severe anomalous matches in the image modality.

3) From the perspective of image-guided-text, we model an optimization task for the linear combination of prompts to address varying prompt quality in the text modality.

## Related Work

**CLIP-based Multimodal Fusion.** Vision-Language Models (VLMs), such as CLIP, have emerged as a new paradigm for foundational models by seamlessly integrating visual and textual information processing. In tasks like image classification, VLMs transfer textual knowledge to the visual modality to enhance performance. Previous training-free FSL methods like CLIP-Adapter (Gao et al. 2021), Tip-Adapter (Zhang et al. 2022), APE (Zhu et al. 2023), and GDA-CLIP (Wang et al. 2024b) have generally overlooked the importance of inter-modal interaction, resulting in limited knowledge transfer between modalities and consequently restricting the performance on downstream tasks.

Some works have recognized the importance of inter-modal interactions. Tip-X (Udandarao, Gupta, and Albanie 2023) uses inter-modal distances to calibrate intra-modal distances. VT-CLIP (Qiu et al. 2023) establishes semantic correlations between text and images through image-guided attention. CALIP (Guo et al. 2023) employs a parameter-free attention module to improve zero-shot classification. Cross-Modal (Lin et al. 2023) repurposes class names as additional support samples during the training stage to combine knowledge at the classification head. However, these methods either ignore that multiple prompts can represent the same class or lack in-depth interaction with multimodal fusion. We solve this by explicitly image-text mutual guidance and ultimately improving performance.

**Prompts Synthesis.** VLMs have excellent performance in zero-shot classification, but recent studies show that task-specific prompt adjustments enhance performance (Liu et al. 2023; Huang, Chu, and Wei 2022). Methods like CoOp (Zhou et al. 2022b), CoCoOp (Zhou et al. 2022a), ProDA (Lu et al. 2022), PLOT (Chen et al. 2023), and CPL (Zhang et al. 2024b) optimize prompts via gradient descent. Alternatively, studies reveal that prompts from LLMs contain crucial discriminatory characteristics of image classification tasks (Pratt et al. 2023). CuPL (Pratt et al. 2023) and DCLIP (Menon and Vondrick 2022) utilise LLMs to acquire fine-grained prompts, while WaffleCLIP (Roth et al. 2023) suggests LLM-generated semantic descriptors enhance classification in a complementary way. Additionally, some works use other external resources like Wikipedia (Tahmasebzadeh et al. 2024; Paz-Argaman et al. 2020; Naeem et al. 2022) and WordNet (Fellbaum 1998; Roth, Vinyals, and Akata 2022; Shen et al. 2022) to help prompts construction.

This paper finds that prompts generated by LLMs may not be fully exploited in FSL due to the modality gap with the image. Therefore, we derive prompt weights by modelling the optimization task of linear combination among prompts, obtaining rectified textual feature representations. This enhancement bridges the gap between different modalities, thereby improving image classification.

## Method

We introduce the proposed technical details of image-text mutual guidance in CLIP-based few-shot classification. Firstly, we formulate the problem and establish relevant notations. Subsequently, two subsections elucidate how Image-Guide-Text (IGT) and Text-Guide-Image (TGI) are implemented, respectively. Then, we provide a detailed explanation of how TIMO and TIMO-S are constructed.

### Preliminaries

We begin with a brief introduction to CLIP (Radford et al. 2021), which is trained on paired text and image data, enabling it to map features from different modalities into a joint feature space. In the zero-shot setting, CLIP encodes the text with prompt templates (*e.g.*, “a photo of a [CLASS]”) into the text feature  $\mathbf{F}_t = [\mathbf{F}_t^1, \dots, \mathbf{F}_t^i, \dots, \mathbf{F}_t^N] \in \mathbb{R}^{N \times P \times D}$ , where  $i$  represents  $i$ -th of the  $N$  categories,  $P$  is the number of prompts per category and  $D$  is the feature dimension. Then following Snell, Swersky, and Zemel (2017), we average all  $P$  prompts yields prototypical features  $\mathbf{W}_t \in \mathbb{R}^{N \times D}$ . Similarly, CLIP encodes the support images into the feature  $\mathbf{F}_v \in \mathbb{R}^{N \times K \times D}$  with prototypical features  $\mathbf{W}_v = [\mathbf{W}_v^1, \dots, \mathbf{W}_v^i, \dots, \mathbf{W}_v^N] \in \mathbb{R}^{N \times D}$ , where  $K$  is the number of shots in FSL. For a query image, it is encoded by CLIP as a visual feature  $\mathbf{f}_v \in \mathbb{R}^D$ .

In zero-shot classification, CLIP uses cosine similarity to build the zero-shot classifier and calculates the logits, which denotes the prediction probabilities for all  $N$  categories.

$$\text{logits}_t = \frac{\mathbf{W}_t \mathbf{f}_v}{\|\mathbf{W}_t\| \|\mathbf{f}_v\|} \in \mathbb{R}^N. \quad (1)$$

In FSL, support image samples are often used to build an additional classifier (*i.e.*,  $\text{Classifier}_{\text{Image}}$ ) (Zhang et al. 2022; Wang et al. 2024b), independent of the aforementioned zero-shot classifier. Then, predictions from both image and text classifiers are combined by weighted summation.

$$\text{logits}_v = \text{Classifier}_{\text{Image}}(\mathbf{f}_v) \in \mathbb{R}^N, \quad (2)$$

$$\text{logits}_o = \text{logits}_v + \alpha \text{logits}_t, \quad (3)$$

where the hyperparameter  $\alpha$  balances the weights of image and text predictions. However, previous SOTA methods build classifiers independently for each modality, ignoring inter-modal correlations. This oversight leads to a semantic gap and leaves single-modality weaknesses unaddressed. We propose a mutual guidance method between image and text modalities, detailed in the following subsection.

### Text-Guide-Image (TGI)

To fully exploit the interrelation between multimodality and alleviate the weakness of each single modality. Here, we discuss the details of the TGI classifier. Since CLIP encodes image and text modalities into homogeneous features located in the same feature space but with different data distributions, we propose using an instance-based transfer learning method. This involves computing a weight matrix  $\mathbf{s}^i \in \mathbb{R}^P$  for each prompt with label  $i$ , and subsequently reordering  $\mathbf{F}_t^i$  based on the descending sequence of  $\mathbf{s}^i$ . Here, we introduce

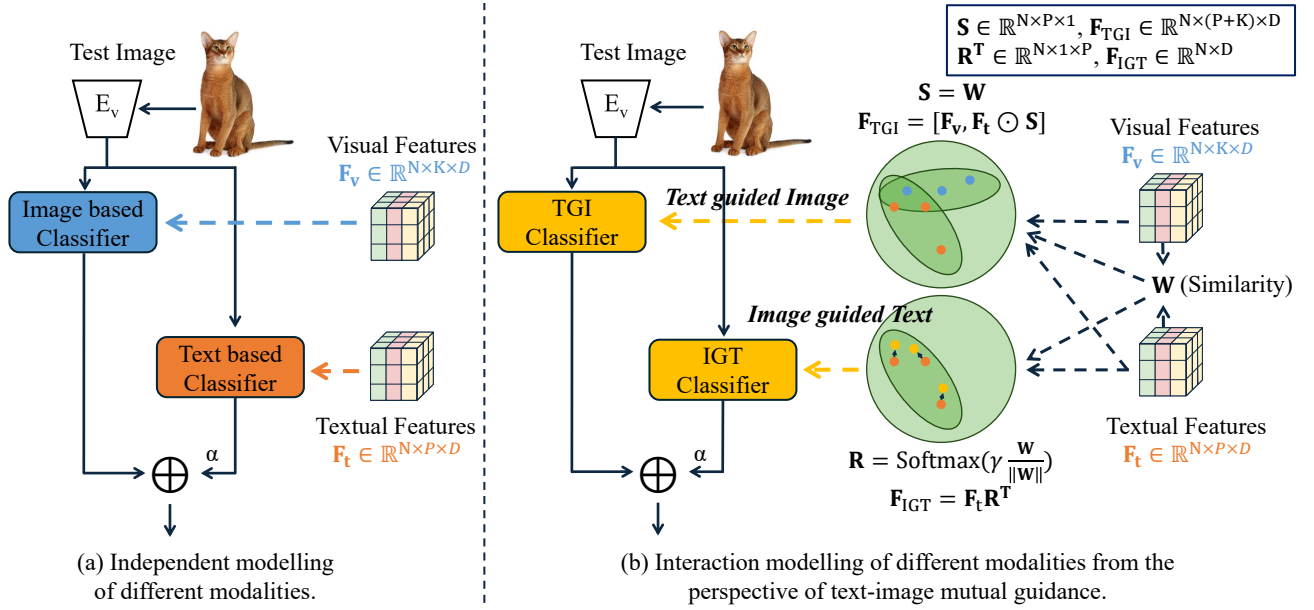


Figure 3: Comparison between existing CLIP-based few-shot methods and ours. (a) indicates the general architecture of previous work. (b) shows how text-image mutual guidance is applied in this work to interact deeply with different modalities.

a hyper-parameter  $\beta$  to adjust the degree of text guidance for the image modality by controlling the number of text prompts involved in the instance-based transfer. The default value for  $\beta$  is set to  $P$ , but by applying a grid search on the validation set, we can optimize this parameter, leading to our TIMO-S variant which will be discussed later.

$$s^i = \frac{\mathbf{F}_t^i \mathbf{W}_v^i}{\|\mathbf{F}_t^i\| \|\mathbf{W}_v^i\|} \in \mathbb{R}^P, \quad (4)$$

$$s^i = \text{diag}(\mathbf{I}_\beta, \mathbf{0}_{P-\beta, P-\beta}) s^i \in \mathbb{R}^P. \quad (5)$$

Here,  $\text{diag}$  is defined as the operation that combines multiple block matrices into a diagonal matrix. Then the obtained weights  $\mathbf{S} = [s^1, \dots, s^i, \dots, s^N] \in \mathbb{R}^{N \times P}$  are element-wise multiplied with the textual features. Subsequently, the Concat operation combines the image features with the non-zero weighted text features to achieve instance-based transfer. This results in the feature  $\mathbf{F}_{TGI}$ , which can be used to construct the TGI classifier.

$$\mathbf{F}_{TGI} = \text{Concat}(\mathbf{F}_v, \mathbf{F}_t \odot \mathbf{S}) \in \mathbb{R}^{N \times (K+\beta) \times D}. \quad (6)$$

After obtaining the refined features  $\mathbf{F}_{TGI}$ , we use the previously established image classifier construction method (*i.e.*, ImageClassifierBuilder) to build the classifier  $\text{Classifier}_{TGI}$ . By substituting ImageClassifierBuilder, our method can achieve the plug-and-play capability with previous SOTA methods. Then, using the  $\text{Classifier}_{TGI}$ , we obtain the logits corresponding to the query image feature  $\mathbf{f}_v$ .

$$\text{logits}_{TGI} = \text{Classifier}_{TGI}(\mathbf{f}_v) \in \mathbb{R}^N. \quad (7)$$

### Image-Guide-Text (IGT)

After the discussion of TGI, we now further discuss how to implement an image-guide-text (IGT) classifier to rectify biased prompts that are independently obtained. From the perspective of ensemble learning, a good IGT classifier needs to satisfy two targets (Krogh and Vedelsby 1994). **Target I:** the IGT Classifier should classify the data correctly; **Target II:** the IGT classifier should hold the individuality information (*i.e.*, keep diversity) from the TGI classifier.

For computational convenience, we first normalize all features using the L2 norm. To achieve **Target I** in a training-free setting, we define a vector  $\mathbf{r}^i \in \mathbb{R}^P$  to maximize the multiplication result between the transformed matrix  $\mathbf{r}^{i\top} \mathbf{F}_t^i$  and the corresponding prototype visual features  $\mathbf{W}_v^i$  in category  $i$ . This operation computes the cosine similarity between each vector in  $\mathbf{F}_t^i$  and the prototype  $\mathbf{W}_v^i$ , with  $\mathbf{r}^i$  assigning higher weights to more similar vectors. However, the vector's magnitude also affects the optimisation objective. Thus, to meet **Target II**'s requirements, we constrain the norm of  $\mathbf{r}^i$  to prevent excessive loss of original text features in the transformed matrix  $\mathbf{r}^{i\top} \mathbf{F}_t^i$ . Here, we introduce a hyperparameter  $\gamma$  to adjust the degree of image guidance for the text modality by controlling the norm of  $\mathbf{r}^i$ . The final optimization task is formulated as follows:

$$\max_{\mathbf{r}^i} \mathbf{r}^{i\top} \mathbf{F}_t^i \mathbf{W}_v^i, \quad \text{s.t.} \|\mathbf{r}^i\| = \gamma. \quad (8)$$

Using the Lagrange multipliers method, Eq. (8) can be transformed into an unconstrained problem:

$$\min_{\mathbf{r}^i} \mathcal{L} = -\mathbf{r}^{i\top} \mathbf{F}_t^i \mathbf{W}_v^i + \lambda(\mathbf{r}^{i\top} \mathbf{r}^i - \gamma). \quad (9)$$

By solving the system of two equations where the partial derivatives  $\nabla_{\mathbf{r}^i} \mathcal{L}$  and  $\nabla_{\lambda} \mathcal{L}$  are both set to zero, we obtain two sets of solutions.

$$\begin{cases} \lambda = \pm \frac{1}{2\gamma} \|\mathbf{F}_t^i \mathbf{W}_v^i\|, \\ \mathbf{r}^i = \pm \frac{\gamma \mathbf{F}_t^i \mathbf{W}_v^i}{\|\mathbf{F}_t^i \mathbf{W}_v^i\|}. \end{cases} \quad (10)$$

Substituting  $\lambda$  back into the original equation, when  $\lambda$  is positive, the original function with respect to  $\mathbf{r}^i$  is convex, yielding the minimum value of the function  $\mathcal{L}$ . Otherwise, the maximum  $\mathcal{L}$  is obtained. Hence, the combination  $\lambda = \frac{1}{2\gamma} \|\mathbf{F}_t^i \mathbf{W}_v^i\|$  and  $\mathbf{r}^i = \frac{\gamma \mathbf{F}_t^i \mathbf{W}_v^i}{\|\mathbf{F}_t^i \mathbf{W}_v^i\|}$  is chosen. Since both  $\mathbf{F}_t^i$  and  $\mathbf{W}_v^i$  are pre-normalised,  $\mathbf{r}^i$  is essentially equivalent to the normalisation of  $\mathbf{s}^i$  (discussed in the previous section). In the end, the logits<sub>IGT</sub> are defined as below:

$$\mathbf{R} = [\dots, \text{SoftMax}(\mathbf{r}^i), \dots] \in \mathbb{R}^{N \times P}, \quad (11)$$

$$\mathbf{F}_{\text{IGT}} = \mathbf{F}_t \mathbf{R}^\top \in \mathbb{R}^{N \times D}. \quad (12)$$

Similar to the construction of Classifier<sub>TGI</sub>, we can use the  $\mathbf{F}_{\text{IGT}}$  and TextClassifierBuilder to construct a classifier Classifier<sub>IGT</sub> and then achieve the prediction of  $\mathbf{f}_v$ .

$$\text{logits}_{\text{IGT}} = \text{Classifier}_{\text{IGT}}(\mathbf{f}_v). \quad (13)$$

## Image-Text Mutual Guidance

As aforementioned, TGI and IGT guide the image and text modalities separately. Additionally, their ability to seamlessly integrate with existing methods (*i.e.*, plug-and-play) enables them to effectively enhance the performance of previous approaches. We use the latest and top-performing method, GDA-CLIP (Wang et al. 2024b), as an example to illustrate how TGI and IGT can be integrated (Note that TGI and IGT can easily combine with other methods as well), leading to the proposed TIMO and TIMO-S.

In general, as shown in Fig. 3, previous methods lack an Image-Text Mutual Guidance module and thus construct independent classifiers using the encoded image features  $\mathbf{F}_v$  and text features  $\mathbf{F}_t$ . We found this independent modelling leads to suboptimal classifiers. To address this, we introduced TGI and IGT components, transforming the original features into refined features  $\mathbf{F}_{\text{TGI}}$  and  $\mathbf{F}_{\text{IGT}}$  via Eq. (6) and (12). Replacing the original unimodal features in GDA-CLIP with our refined features, while keeping the classifier construction process unchanged, allows for easy integration of our method with GDA-CLIP or other methods.

By incorporating TGI and IGT into existing methods, we proposed TIMO. Its prediction can be obtained by the ensemble logits<sub>TGI</sub> from Eq. (7) and logits<sub>IGT</sub> from Eq. (13).

$$\text{logits}_{\text{TIMO}} = \text{logits}_{\text{TGI}} + \alpha \text{logits}_{\text{IGT}}. \quad (14)$$

A grid search over  $\alpha \in \{10^{-4}, 10^{-3}, \dots, 10^4\}$  is performed in the validation set to determine its optimal value, following established methodologies. By integrating TGI and IGT, we achieved mutual guidance between the image and text modalities, effectively overcoming individual weaknesses and narrowing the semantic gap between them.

**Construction of TIMO-S.** Experiments reveal that adjusting the hyperparameters  $\beta$  and  $\gamma$ , which control the degree of mutual guidance, can further enhance TIMO’s performance. Thus, we suggest an additional grid search on the validation set for  $\beta$  and  $\gamma$ , resulting in a more robust model, TIMO-S. Note that, all compared training-free methods employ similar grid searches, ensuring fair comparison.

## Experiments

### Setting and Implementation

**Dataset.** According to previous works (Wang et al. 2024b; Zhang et al. 2022), we use 11 image classification benchmarks to evaluate few-shot image classification. The specific datasets and their licenses are detailed in the Appendix.

**Baselines.** Our method is compared with 2 categories of existing methods. 1) For training-free methods, Zero-Shot CLIP (Radford et al. 2021), Tip-Adapter (Zhang et al. 2022), APE (Zhu et al. 2023) and GDA-CLIP (Wang et al. 2024b) are compared; 2) For training-required methods, Tip-Adapter-F (Zhang et al. 2022), Cross-Modal (Lin et al. 2023) APE-T (Zhu et al. 2023) are compared.

**Implementation Detail.** This paper follows the data pre-processing protocol of CoOp (Zhou et al. 2022b) and Tip-Adapter (Zhang et al. 2022). For TIMO, we set the hyperparameter  $\gamma$  to 50 for all datasets, except 1 for ImageNet and 100 for Flowers102. To evaluate our method, we follow (Zhu et al. 2023; Wang et al. 2024b; Zhang et al. 2022) using ResNet-50 (RN50) (He et al. 2016) as the default visual encoder of CLIP, compare 1, 2, 4, 8, and 16 few-shot training sets, and use textual prompts from APE (Zhu et al. 2023) and CuPL (Pratt et al. 2023). All experiments use PyTorch (Paszke et al. 2019) and are conducted on a single NVIDIA GeForce RTX 3090Ti GPU.

### Main Result

**Performance Analysis.** Tab. 1 shows few-shot classification results for both training-free and training-required methods, with bolded and underlined results representing optimal and suboptimal results separately. Linear Probe CLIP (Radford et al. 2021) simply fine-tunes the classifier but results in performance degradation in low-shot settings than zero-shot CLIP (Radford et al. 2021). Subsequent works strive to preserve the model’s zero-shot capabilities, where our TIMO enhance and spreads these capabilities through cross-modal mutual guidance mechanism, surpassing all previous training-free methods (especially in low-shot settings). Additionally, our enhanced variant, TIMO-S, even outperforms all training-required methods under training-free conditions.

**Robustness Analysis.** To evaluate the robustness of our method under varying conditions, we test its performance using different CLIP backbones. The results are presented in Tab. 2, indicating that TIMO-S consistently outperforms prior methods on all visual encoders. Consistent with prior research (Wang et al. 2024b), we also use 16-shot ImageNet as training data and assess domain generalization on out-of-distribution datasets, including ImageNet-V2, ImageNet-Sketch, ImageNet-A, and ImageNet-R, using a RN50-based CLIP. Fig. 4 shows that TIMO-S outperforms GDA-CLIP

Methods	Training-Free	Number of shots					Avg.
		1	2	4	8	16	
Linear Probe CLIP <sub>ICML'21</sub> (Radford et al. 2021)	✗	36.10	46.99	56.72	64.66	70.56	55.01
CoOp <sub>IJCV'22</sub> (Zhou et al. 2022b)	✗	59.56	61.78	66.47	69.85	73.33	66.20
Tip-Adapter-F <sub>ECCV'22</sub> (Zhang et al. 2022)	✗	65.96	67.86	70.28	72.99	75.87	70.59
Cross-Modal Partial <sub>CVPR'23</sub> (Lin et al. 2023)	✗	65.97	68.37	71.14	<u>73.96</u>	<b>77.05</b>	<u>71.30</u>
APE-T <sub>ICCV'23</sub> (Zhu et al. 2023)	✗	65.96	68.56	71.27	73.88	76.69	71.27
LP++ <sub>CVPR'24</sub> (Huang et al. 2024)	✗	64.56	66.32	69.41	71.79	74.70	69.36
CLAP <sub>CVPR'24</sub> (Silva-Rodríguez et al. 2024)	✗	66.10	<b>69.10</b>	<b>72.10</b>	73.50	75.10	71.18
Zero-Shot CLIP <sub>ICML'21</sub> (Radford et al. 2021)	✓	62.11	62.11	62.11	62.11	62.11	62.11
Tip-Adapter <sub>ECCV'22</sub> (Zhang et al. 2022)	✓	64.24	65.51	67.31	69.00	70.61	67.33
Tip-MG (Ours)	✓	66.14	67.35	68.86	70.58	71.66	68.92
APE <sub>ICCV'23</sub> (Zhu et al. 2023)	✓	65.78	67.24	69.32	71.16	72.99	69.30
GDA-CLIP <sub>ICLR'24</sub> (Wang et al. 2024b)	✓	63.69	66.39	69.86	73.17	76.23	69.87
TIMO (Ours)	✓	<u>66.43</u>	68.15	71.20	73.55	76.61	71.19
TIMO-S (Ours)	✓	<b>66.78</b>	<u>68.69</u>	<u>71.63</u>	<b>74.20</b>	<u>76.86</u>	<b>71.63</b>

Table 1: Performance (%) of few-shot classification on 11 datasets, showing top-1 accuracy across 3 random seeds. TIMO surpasses all previous training-free approaches. Furthermore, its enhanced variant, TIMO-S, outperforms all previous training-required methods. Results of Linear Probe CLIP and CLAP are sourced from Huang et al. (2024) and Wang et al. (2024b).

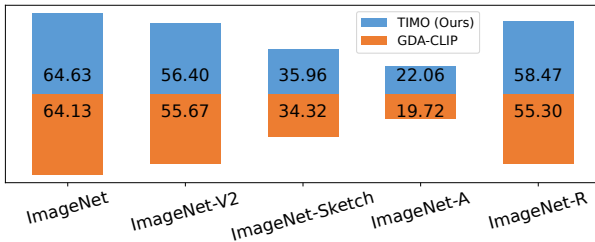


Figure 4: Domain Generalization Performance (%) of TIMO and GDA-CLIP.

across all target datasets, likely due to the more effective engagement of robust textual features in image classification.

## Ablation Study

**Effectiveness of TGI and IGT.** We investigate whether the developed TGI and IGT classifiers effectively address the previously noted issues: anomalous match and varying quality of generated prompts.

For the first issue, as shown in the bottom of Fig. 1, our method effectively reduces the number of anomalous matches compared to all previous approaches. Furthermore, our method’s effectiveness has been extensively validated across multiple datasets through experiments.

For the second issue, we evaluated the performance using all prompts, only the best half of the prompts, only the worst half of the prompts, and rectified prompts by IGT, as shown in Fig. 5. The results indicate that our method achieves overall better task-relevant textual representations.

Additionally, we conducted ablation experiments on the performance of TGI and IGT classifiers and their combinations. Tab. 3 shows that integrating either TGI or IGT into

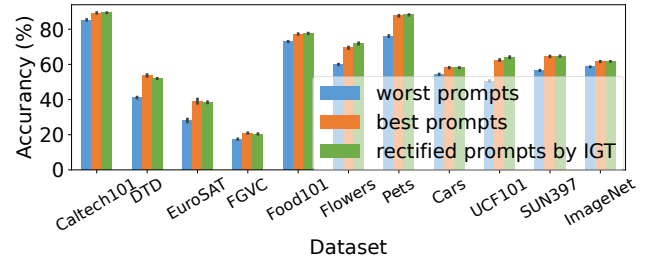


Figure 5: Ablation of prompt selection.

previous methods enhances the model’s discriminative ability. The best performance improvement is achieved when both TGI and IGT are applied together.

**Efficacy of hyperparameters.** Our method involves 3 hyperparameters. We conducted experiments on the DTD dataset to evaluate their impact on the model’s performance.

$\alpha$  balances the weight of the TGI and IGT classifiers. As  $\alpha = 0$ , only TGI is active. As  $\alpha \rightarrow \infty$ , only IGT is active. Fig. 6 shows its ablation study, indicating that as  $\alpha$  increases, performance under different shots starts to converge.

$\beta$  and  $\gamma$  adjust the degree of text/image guidance for the image/text modality. As shown in Fig. 6, their overall effect on performance is relatively stable, but in rare cases, they may cause abrupt performance changes.

**Different Prompt Templates.** We investigate how different prompt methods may affect the effectiveness of the IGT. In this regard, we primarily focus 4 prompt methods: CuPL (Pratt et al. 2023), DCLIP (Menon and Vondrick 2022) WaffleCLIP (Roth et al. 2023) and use them all together.

From Fig. 7, it is evident that different prompt templates impact the efficacy of the IGT classifier. Another interesting

Methods	RN50	RN101	ViT-B/32	ViT-B/16	Avg.
Tip-Adapter (Zhang et al. 2022)	67.33	68.13	70.07	74.20	69.92
APE (Zhu et al. 2023)	69.30	70.41	71.78	75.89	71.84
GDA-CLIP (Wang et al. 2024b)	69.87	71.60	72.70	76.69	72.71
TIMO (Ours)	71.19	72.24	73.71	77.61	73.69
TIMO-S (Ours)	<b>71.63</b>	<b>72.98</b>	<b>74.05</b>	<b>77.97</b>	<b>74.16</b>

Table 2: The average performance (%) of few-shot classifications using different visual encoders across 11 datasets and 3 random seeds.

Methods	TGI	IGT	Avg (%)
	—	—	67.28
Tip-Adapter (Zhang et al. 2022)	+TGI	—	68.35
	—	+IGT	68.43
	+TGI	+IGT	68.92
GDA-CLIP (Wang et al. 2024b)	—	—	69.87
	+TGI	—	70.62
	—	+IGT	70.65
	+TGI	+IGT	71.19

Table 3: Ablation study on different modules and their combinations.

Methods	Acc. (%)	Time (s)	Param.
Tip-Adapter (Zhang et al. 2022)	62.88	257.4	<i>NKD</i>
APE (Zhu et al. 2023)	63.43	787.5	<i>NKD</i>
GDA-CLIP (Wang et al. 2024b)	64.46	1.1	<i>ND</i>
TIMO (Ours)	64.71	1.1	<i>ND</i>
TIMO-S (Ours)	64.85	6.0	<i>ND</i>

Table 4: Complexity Analysis on 16-shots ImageNet.

observation is that IGT improves more on CuPL than using all prompts. This indicates that the FSL could perform better when employing prompts containing richer and more diverse information, such as CuPL.

### Complexity Discussion

From Tab. 4, our method focuses solely on image-text mutual guidance, which does not introduce extra parameters or increase space complexity. Regarding time complexity, TIMO adds an  $O(1)$  operation compared to previous methods, maintaining the overall time complexity. However, TIMO includes additional hyperparameter search, resulting in better performance and relatively higher time cost.

### Conclusion

This paper introduces TIMO, a novel model for CLIP-based few-shot image classification. Unlike prior training-free methods that independently model different modalities, TIMO introduces a unique perspective by incorporating the IGT classifier and the TGI classifier to address the problem of the modality’s inherent weakness and suboptimal prompt. Extensive experiments demonstrate the effectiveness and universality of the cross-modal mutual guidance,

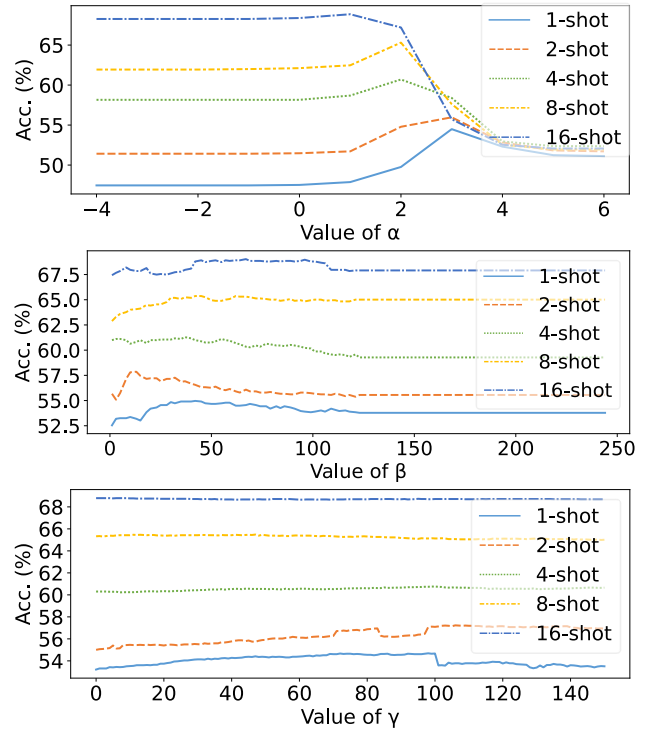


Figure 6: Ablation of hyperparameters.




Figure 7: Performance improvement brought by IGT compared with zero-shot classification.

showing that TIMO-S outperforms previous best training-required methods under training-free conditions. However, our method does not employ novel training techniques to fully exploit CLIP’s potential, and our multimodal mutual guidance mechanism could potentially be extended to other fields (*e.g.*, image retrieval). These challenges remain open for future exploration.

## Acknowledgments

This work was supported by the National Key R&D Program of China (2023ZD0120700, 2023ZD0120701), NSFC Project (62222604, 62206052, 624B2063), China Postdoctoral Science Foundation (2024M750424), the Fundamental Research Funds for the Central Universities (020214380120), the State Key Laboratory Fund (ZZKT2024A14), the Postdoctoral Fellowship Program of CPSF (GZC20240252), and the Jiangsu Funding Program for Excellent Postdoctoral Talent (2024ZB242).

## References

- Chen, G.; Yao, W.; Song, X.; Li, X.; Rao, Y.; and Zhang, K. 2023. PLOT: Prompt Learning with Optimal Transport for Vision-Language Models. (arXiv:2210.01253).
- Fan, Y.; Xu, W.; Wang, H.; Wang, J.; and Guo, S. 2023. PMR: Prototypical Modal Rebalance for Multimodal Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20029–20038.
- Fellbaum, C. 1998. *WordNet: An Electronic Lexical Database*. MIT Press. ISBN 978-0-262-06197-1.
- Gao, P.; Geng, S.; Zhang, R.; Ma, T.; Fang, R.; Zhang, Y.; Li, H.; and Qiao, Y. 2021. CLIP-Adapter: Better Vision-Language Models with Feature Adapters. (arXiv:2110.04544).
- Gong, C. 2017. Exploring Commonality and Individuality for Multi-Modal Curriculum Learning. In *AAAI*, volume 31.
- Guo, Z.; Zhang, R.; Qiu, L.; Ma, X.; Miao, X.; He, X.; and Cui, B. 2023. CALIP: Zero-Shot Enhancement of CLIP with Parameter-Free Attention. In *AAAI*, volume 37, 746–754.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Huang, T.; Chu, J.; and Wei, F. 2022. Unsupervised Prompt Learning for Vision-Language Models. (arXiv:2204.03649).
- Huang, Y.; Shakeri, F.; Dolz, J.; Boudiaf, M.; Bahig, H.; and Ayed, I. B. 2024. LP++: A Surprisingly Strong Linear Probe for Few-Shot CLIP. (arXiv:2404.02285).
- Krogh, A.; and Vedelsby, J. 1994. Neural Network Ensembles, Cross Validation, and Active Learning. In *Advances in Neural Information Processing Systems*, volume 7.
- Lin, Z.; Yu, S.; Kuang, Z.; Pathak, D.; and Ramanan, D. 2023. Multimodality Helps Unimodality: Cross-Modal Few-Shot Learning With Multimodal Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19325–19337.
- Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2023. Pre-Train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *ACM Computing Surveys*, 55(9): 195:1–195:35.
- Lu, Y.; Liu, J.; Zhang, Y.; Liu, Y.; and Tian, X. 2022. Prompt Distribution Learning. (arXiv:2205.03340).
- Martin, S.; Huang, Y.; Shakeri, F.; Pesquet, J.-C.; and Ben Ayed, I. 2024. Transductive Zero-Shot and Few-Shot CLIP. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28816–28826.
- Menon, S.; and Vondrick, C. 2022. Visual Classification via Description from Large Language Models. (arXiv:2210.07183).
- Mu, N.; Kirillov, A.; Wagner, D.; and Xie, S. 2022. SLIP: Self-supervision Meets Language-Image Pre-training. 529–544.
- Naeem, M. F.; Xian, Y.; Gool, L. V.; and Tombari, F. 2022. I2DFormer: Learning Image to Document Attention for Zero-Shot Image Classification. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 12283–12294.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, volume 32.
- Paz-Argaman, T.; Atzmon, Y.; Chechik, G.; and Tsarfaty, R. 2020. ZEST: Zero-shot Learning from Text Descriptions Using Textual Similarity and Visual Summarization. (arXiv:2010.03276).
- Pratt, S.; Covert, I.; Liu, R.; and Farhadi, A. 2023. What Does a Platypus Look Like? Generating Customized Prompts for Zero-Shot Image Classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15691–15701.
- Qiu, L.; Zhang, R.; Guo, Z.; Zeng, Z.; Guo, Z.; Li, Y.; and Zhang, G. 2023. VT-CLIP: Enhancing Vision-Language Models with Visual-guided Texts. (arXiv:2112.02399).
- 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 *Proceedings of the 38th International Conference on Machine Learning*, 8748–8763. PMLR.
- Roth, K.; Kim, J. M.; Koepke, A. S.; Vinyals, O.; Schmid, C.; and Akata, Z. 2023. Waffling Around for Performance: Visual Classification with Random Words and Broad Concepts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15746–15757.
- Roth, K.; Vinyals, O.; and Akata, Z. 2022. Integrating Language Guidance Into Vision-Based Deep Metric Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16177–16189.
- Sanghi, A.; Chu, H.; Lambourne, J. G.; Wang, Y.; Cheng, C.-Y.; Fumero, M.; and Malekshan, K. R. 2022. CLIP-Forge: Towards Zero-Shot Text-To-Shape Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18603–18613.
- Shao, S.; Bai, Y.; Wang, Y.; Liu, B.; and Zhou, Y. 2024. DeIL: Direct-and-Inverse CLIP for Open-World Few-Shot Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28505–28514.

- Shen, S.; Li, C.; Hu, X.; Xie, Y.; Yang, J.; Zhang, P.; Gan, Z.; Wang, L.; Yuan, L.; Liu, C.; Keutzer, K.; Darrell, T.; Rohrbach, A.; and Gao, J. 2022. K-LITE: Learning Transferable Visual Models with External Knowledge. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 15558–15573.
- Silva-Rodríguez, J.; Hajimiri, S.; Ben Ayed, I.; and Dolz, J. 2024. A Closer Look at the Few-Shot Adaptation of Large Vision-Language Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23681–23690.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical Networks for Few-shot Learning. In *Advances in Neural Information Processing Systems*, volume 30.
- Tahmasebzadeh, G.; Springstein, M.; Ewerth, R.; and Müller-Budack, E. 2024. Few-Shot Event Classification in Images Using Knowledge Graphs for Prompting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 7286–7295.
- Tan, Q.; Yu, G.; Wang, J.; Domeniconi, C.; and Zhang, X. 2021. Individuality- and Commonality-Based Multiview Multilabel Learning. *IEEE Transactions on Cybernetics*, 51(3): 1716–1727.
- Tang, Y.; Lin, Z.; Wang, Q.; Zhu, P.; and Hu, Q. 2024. AMU-Tuning: Effective Logit Bias for CLIP-based Few-shot Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23323–23333.
- Udandarao, V.; Gupta, A.; and Albanie, S. 2023. SuS-X: Training-Free Name-Only Transfer of Vision-Language Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2725–2736.
- Wang, Y.; Yao, Q.; Kwok, J.; and Ni, L. M. 2020. Generalizing from a Few Examples: A Survey on Few-Shot Learning. (arXiv:1904.05046).
- Wang, Z.; Liang, J.; He, R.; Wang, Z.; and Tan, T. 2024a. Connecting the Dots: Collaborative Fine-tuning for Black-Box Vision-Language Models. arXiv:2402.04050.
- Wang, Z.; Liang, J.; Sheng, L.; He, R.; Wang, Z.; and Tan, T. 2024b. A Hard-to-Beat Baseline for Training-free CLIP-based Adaptation. (arXiv:2402.04087).
- Zhang, R.; Hu, X.; Li, B.; Huang, S.; Deng, H.; Qiao, Y.; Gao, P.; and Li, H. 2023. Prompt, Generate, Then Cache: Cascade of Foundation Models Makes Strong Few-Shot Learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15211–15222.
- 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.; Cissé, M.; Farinella, G. M.; and Hassner, T., eds., *Proceedings of the European Conference on Computer Vision (ECCV)*, 493–510. Cham: Springer Nature Switzerland. ISBN 978-3-031-19833-5.
- Zhang, X.; Yoon, J.; Bansal, M.; and Yao, H. 2024a. Multimodal Representation Learning by Alternating Unimodal Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27456–27466.
- Zhang, Y.; Zhang, C.; Yu, K.; Tang, Y.; and He, Z. 2024b. Concept-Guided Prompt Learning for Generalization in Vision-Language Models. (arXiv:2401.07457).
- Zhao, Z.; Liu, Y.; Wu, H.; Li, Y.; Wang, S.; Teng, L.; Liu, D.; Cui, Z.; Wang, Q.; and Shen, D. 2024. CLIP in Medical Imaging: A Comprehensive Survey. (arXiv:2312.07353).
- 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, Z.; Lei, Y.; Zhang, B.; Liu, L.; and Liu, Y. 2023. Zeg-CLIP: Towards Adapting CLIP for Zero-Shot Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11175–11185.
- Zhu, X.; Zhang, R.; He, B.; Zhou, A.; Wang, D.; Zhao, B.; and Gao, P. 2023. Not All Features Matter: Enhancing Few-shot CLIP with Adaptive Prior Refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2605–2615.