

Decoupling Template Bias in CLIP: Harnessing Empty Prompts for Enhanced Few-Shot Learning

Zhenyu Zhang¹, Guangyao Chen^{2*}, Yixiong Zou^{1*}, Zhimeng Huang², Yuhua Li^{1*}

¹School of Computer Science and Technology, Huazhong University of Science and Technology

²National Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

Abstract

The Contrastive Language-Image Pre-Training (CLIP) model excels in few-shot learning by aligning visual and textual representations. Our study shows that template-sample similarity (TSS), defined as the resemblance between a text template and an image sample, introduces bias. This bias leads the model to rely on template proximity rather than true sample-to-category alignment, reducing both accuracy and robustness in classification. We present a framework that uses empty prompts, textual inputs that convey the idea of “emptiness” without category information. These prompts capture unbiased template features and offset TSS bias. The framework employs two stages. During pre-training, empty prompts reveal and reduce template-induced bias within the CLIP encoder. During few-shot fine-tuning, a bias calibration loss enforces correct alignment between images and their categories, ensuring the model focuses on relevant visual cues. Experiments across multiple benchmarks demonstrate that our template correction method significantly reduces performance fluctuations caused by TSS, yielding higher classification accuracy and stronger robustness.

Introduction

CLIP (Contrastive Language-Image Pre-Training) (Radford et al. 2021) is a multimodal pre-trained neural network designed to align images and text using large-scale paired image-text data. The model consists of two branches: a text encoder and an image encoder, each mapping textual descriptions and visual samples into low-dimensional vector representations. During pre-training, CLIP learns to perform a wide range of tasks, including OCR (Materzyńska, Torralba, and Bau 2022), geolocation (Vivanco Cepeda, Nayak, and Shah 2024), and action recognition (Ke et al. 2018). In the prediction phase, CLIP generates predictions by calculating the cosine similarity between text and image vectors. This model is particularly effective for zero-shot learning tasks, where it can make predictions without having seen training examples of new images or text.

In downstream few-shot learning tasks, CLIP classifies an image by encoding text descriptions such as “a photo of a cat” and comparing the resulting embeddings to im-

age features. We define three cosine-based similarity measures: *template-sample similarity* (TSS), between an image and the embedding of the blank template “a photo of a {}”; *class-sample similarity* (CSS), between an image and the embedding of the bare class name (e.g., “cat”); and *prompt-sample similarity* (PSS), between an image and the embedding of the full prompt (e.g., “a photo of a cat”). Prior work has investigated the impact of classname variations on CLIP’s zero-shot performance (Zhang and Ré 2022; Zhu et al. 2023b; Shin et al. 2024; Cabello et al. 2023) and demonstrated that carefully engineered prompts can yield substantial accuracy gains (Yang et al. 2023; You et al. 2024; Wen et al. 2024; Chen et al. 2024; Fabbrizzi et al. 2022; Seth, Hemani, and Agarwal 2023). In parallel, several studies have shown that different template choices can induce large fluctuations in performance (Liusie, Manakul, and Gales 2023; He et al. 2024). As illustrated in the upper panel of Figure 1, encoding each image with the generic template “a photo of a {}” produces systematically varying Template-Sample Similarity (TSS) scores, and appending the true class name yields different Prompt-Sample Similarity (PSS) values. Concretely, images with higher TSS scores consistently generate higher PSS across all class prompts. Most importantly, **the model exhibits an attention bias even when the input is a blank template, and this bias persists with a full prompt (including category information), resulting in incorrect final predictions.** These findings pose an important question: *although templates supply crucial contextual cues, might they also introduce systematic biases that limit CLIP’s ability to discriminate among classes and thereby constrain its overall effectiveness?*

To elucidate the effects of template-sample similarity (TSS) and optimization strategies on few-shot performance, we performed a detailed analysis of how different prompt templates influence classification accuracy. We find that, in low-data regimes, there is a strong correlation between an image’s TSS score and its likelihood of being correctly classified, suggesting that models often exploit template similarity instead of truly aligning samples with their correct labels. Existing few-shot methods largely ignore this template-induced distortion. When only a handful of examples are available, such biases cannot be overcome by data augmentation or extensive training alone. Consequently, explicitly accounting for TSS is essential to improve both the robust-

*Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

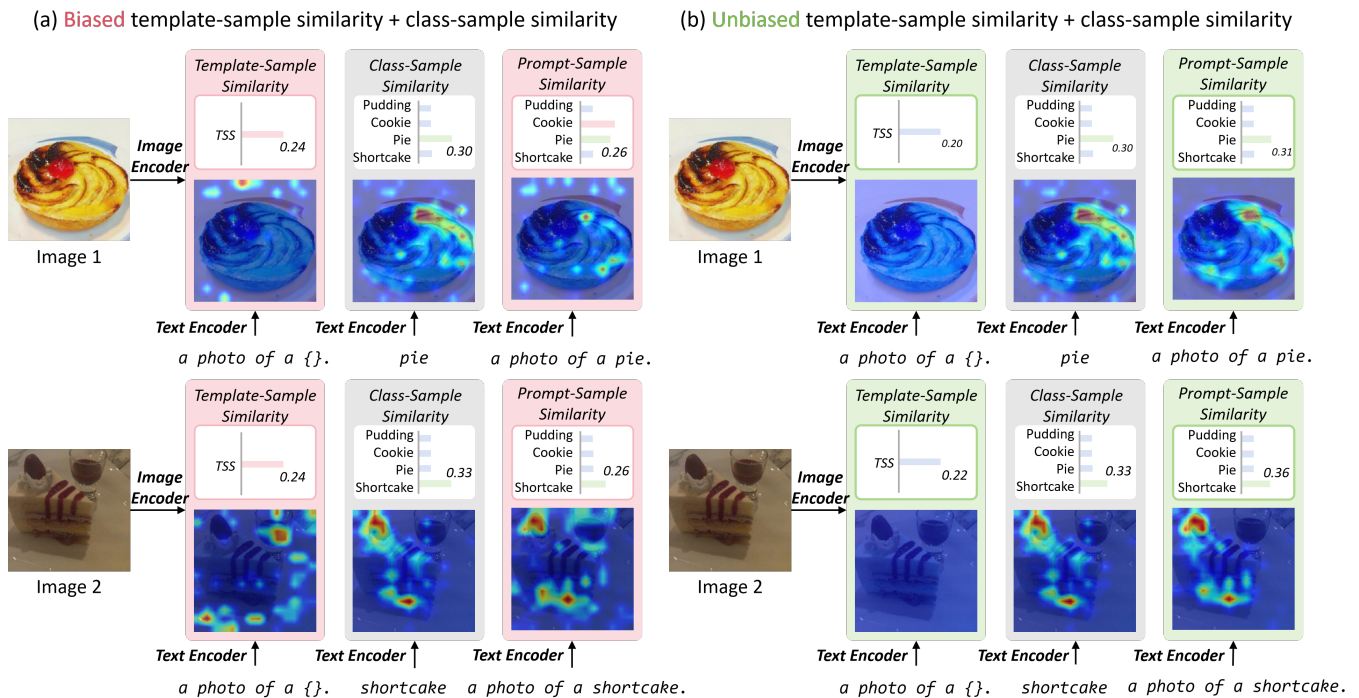


Figure 1: Impact of template bias on CLIP classification. (a) A biased prompt (“a photo of a ”) introduces extra template–sample similarity that skews attention toward irrelevant regions and inflates incorrect class scores—even when the class name matches the object. (b) An unbiased prompt removes template bias, so the model focuses on genuinely discriminative features and the model’s prediction relies solely on class–sample similarity and correct prompt–sample similarity. .

ness and the accuracy of few-shot learning systems.

Based on these findings, we propose a method to eliminate the bias caused by the relationship between templates and samples, thereby further improving few-shot learning. As shown in Figure 3, we first construct multiple empty prompts by combining the template with words that represent “emptiness”. This captures the feature representations of the template’s textual space when category information is absent. We address the template bias by ensuring that these empty prompts exhibit consistent similarity with the samples, forcing the template-sample similarity to be uniform across all samples and correcting the bias introduced by the template. Second, during the pretraining phase, these empty prompts are used to detect and correct biases within the CLIP model, ensuring that the initial model parameters are less influenced by template similarities. Finally, we perform few-shot fine-tuning using limited labeled samples, simultaneously applying a bias calibration loss to maintain the model’s focus on sample-category alignment rather than template similarity. This two-stage training approach effectively reduces the dependency on TSS, leading to enhanced classification accuracy and robustness across various templates and categories. Our method ensures that the model relies more on the intrinsic alignment between samples and their true categories, thereby improving the fairness and performance of few-shot learning models.

In summary, our contributions are as follows:

- We uncover and analyze significant performance fluctua-

tions in CLIP-based models caused by the similarity between text templates and image samples. This indicates that while templates improve the overall performance, they also introduce a non-negligible prior bias.

- We construct multiple “empty” prompts representing the concept of “emptiness” capturing the textual feature representations of templates without category information.
- Based on these insights, we introduce a two-stage training strategy that effectively reduces dependency on template-sample similarity, enhancing accuracy and robustness in few-shot learning.
- Extensive experiments on multiple datasets demonstrate that our template correction method significantly reduces performance fluctuations caused by templates and improves the overall performance of CLIP-based few-shot learning models.

Template-Sample Similarity

What is Template-Sample Similarity?

CLIP-based image classification compares visual features with text embeddings formed by inserting class names into templates. For a given image s , CLIP computes cosine similarities between its visual embedding $\theta_v(s)$ and three text embeddings: the blank template $t_0 =$ “a photo of a {}”, the bare class name c , and the full prompt $p_c =$ “a photo of a c”. We denote these met-




Figure 2: (Left) Correlation between Template-Sample Similarity and classification accuracy on EuroSAT (correlation coefficients is included in the legend). (Middle) The evolution of the absolute value of correlation coefficients between classification accuracy and template-sample similarity over the course of 1-shot training on EuroSAT. (Right) The effect of different templates on overall performance, shown both before and after applying template correction on EuroSAT.

	ImagNet	SUN	Aircraft	EuroSAT	Cars	Food	Pets	Flowers	Caltech	DTD	UCF	Avg.
Only class	64.08	60.81	20.97	46.38	63.67	83.94	81.92	64.67	87.22	45.80	63.83	62.12
w/ Template	66.68	62.49	23.25	42.38	65.40	85.27	88.36	67.35	93.38	44.32	65.15	64.00
Misclassification Ratio	4.82	6.49	4.32	9.79	5.19	2.39	1.47	4.1	1.29	8.21	6.18	4.93

Table 1: Comparison of classification performance using bare class names versus full templates. We define the *misclassification ratio* as **(samples correctly classified using only class names but misclassified after introducing templates) / (total samples)**. Although templates improve overall accuracy, they also induce errors on a subset of samples that were previously classified correctly.

rics as

$$\text{TSS}(s) = \cos(\theta_t(t_0), \theta_v(s)), \quad (1)$$

$$\text{CSS}(c, s) = \cos(\theta_t(c), \theta_v(s)), \quad (2)$$

$$\text{PSS}(c, s) = \cos(\theta_t(p_c), \theta_v(s)). \quad (3)$$

where θ_t and θ_v are the text and image encoders, respectively. In particular, Template-Sample Similarity (TSS) characterizes the intrinsic compatibility between a visual sample and the general context suggested by the blank template, independent of specific semantic content. Class-Sample Similarity (CSS), on the other hand, directly quantifies the alignment between the visual sample and the semantic embedding of the category name itself. Finally, Prompt-Sample Similarity (PSS) encapsulates both template context and semantic specificity, combining their joint effects into a single metric.

Here, TSS captures how well the image aligns with the generic template context, CSS measures the image’s alignment with the class name itself, and PSS reflects their combined effect. As Figure 1(a) shows, samples with higher TSS consistently achieve higher PSS across classes, indicating a systematic template-driven bias in CLIP’s predictions. By separating TSS (template influence) from CSS (category-specific signal), we can more precisely analyze bias and classification performance in multimodal representation learning.

Impact of Template-Sample Similarity

In classification, the essential determinant is the alignment between each visual sample and its true category. We further

ask whether the Template-Sample Similarity Score (TSS) also influences the outcome. Table 1 shows that appending full templates to the CLIP pipeline consistently raises accuracy across all eleven datasets—boosting average accuracy from 62.12% (using only class names) to 64.00% (with templates). Yet this gain carries a penalty: on average, 4.93% of samples that CLIP correctly classified by class names alone are misclassified when templates are applied. The misclassification ratio peaks at 9.79% on EuroSAT and exceeds 6% for SUN and UCF, while Caltech exhibits the lowest rate at 1.29%. Thus, Table 1 highlights a crucial trade-off—**templates enhance overall accuracy but can also introduce bias**.

To investigate this effect more deeply, we use an unfine-tuned CLIP model to compute the TSS for each sample in EuroSAT (Helber et al. 2019) with templates that omit category names. The TSS values are grouped into bins of 400 samples each (e.g., samples with TSS values between 0.19 and 0.20). We then compute the classification accuracy for each bin. Figure 2(Left) shows a strong correlation (correlation coefficient 0.9) between TSS and the probability of correct classification, indicating that TSS significantly affects whether a sample is accurately classified using a given template.

This phenomenon can be attributed to template-induced bias. As illustrated in Figure 1, when textual features of a template differ significantly from visual characteristics of a sample, the template becomes unsuitable. For example, “a centered satellite photo of” aligns well with satellite images’

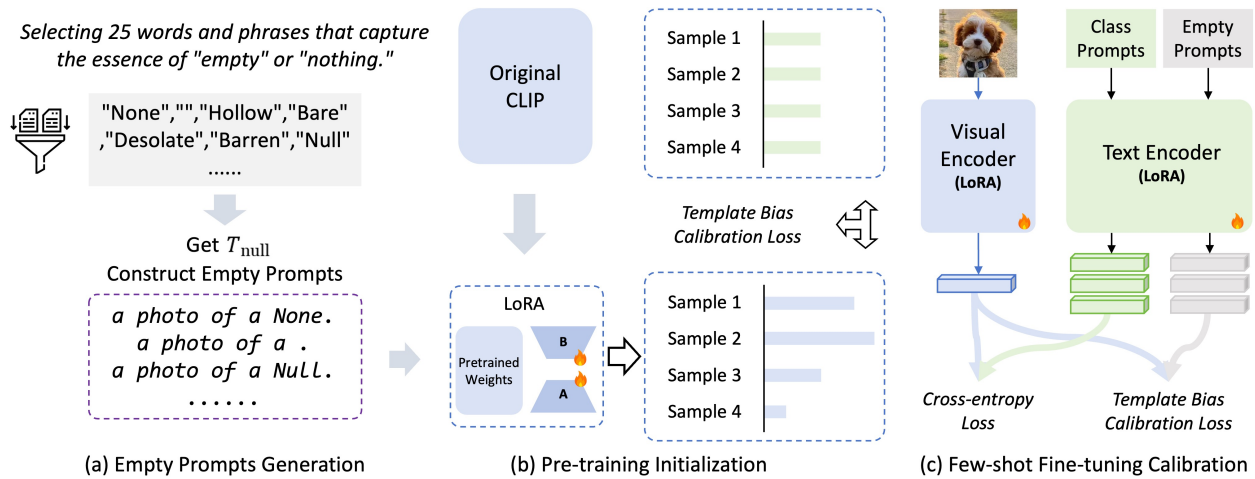


Figure 3: The overall framework for template correction involves three main stages. (a) **Empty Prompts Generation:** A diverse set of empty prompts is manually curated to help identify potential template-induced biases. (b) **Pretraining Initialization:** These empty prompts are used to detect and correct biases within the CLIP model during the pretraining phase. (c) **Few-shot Fine-tuning Calibration:** Finally, the model undergoes fine-tuning with few-shot samples to calibrate its performance and improve classification accuracy.

features and help classification, while “a cartoon photo of” generates mismatched features and leads to errors. Based on these findings, we draw the following observations.

Observation 0.1. *A sample’s classification accuracy depends jointly on how well it matches the chosen template and how well it aligns with its true class. In practice, samples that score highly against a template tend to be classified correctly more often when the correct class name is appended. This demonstrates that TSS and CSS act together to determine performance.*

As shown in Figure 1(a), although the TSS captures rich contextual cues, it can inadvertently align with irrelevant visual features during the pre-training alignment phase, causing the model to attend to non-target regions instead of the object of interest. While the precise mechanism by which the template influences performance remains unclear, this observation highlights an important design consideration for few-shot learning algorithms.

Can few-shot training eliminate template bias?

In an ideal classification framework, decision outcomes should depend solely on the degree of alignment between each sample and its true category, remaining invariant to extraneous factors such as TSS. However, as shown in Figure 2(Right), using different templates to construct prompts during training causes substantial fluctuations in model performance, demonstrating that template-induced bias endures despite fine-tuning. To further quantify this effect, we measured the variance in classification accuracy across all categories for each template. Figure 2(Middle) reveals that fine-tuning techniques like LoRA effectively reduce—but do not entirely eliminate—the correlation between classification accuracy and TSS in the CLIP model. Specifically, the correlation coefficient declines after adaptation but stabi-

lizes at approximately 0.2, indicating a residual dependence on template similarity. These findings underscore that direct fine-tuning alone is insufficient to fully remove bias introduced by template prompts.

Observation 0.2. *Direct fine-tuning alone cannot completely eliminate the bias induced by template prompts, as a non-negligible correlation between classification accuracy and TSS remains.*

Template Correction for Few-Shot Learning

Task Formulation

In classification tasks using VLMs, we are provided with a set of K candidate classes. For each class, a textual description, known as a prompt, is created using a predefined template. For example, the template “a photo of a { }” can generate prompts like “a photo of a cat” or “a photo of a dog.” This template guides the text encoder to extract features that align with the visual characteristics of each class. Let c_k denote the tokenized prompt for the k -th class. The language encoder θ_t processes c_k to produce the normalized textual embedding $\mathbf{t}_k = \theta_t(c_k)$, which lies on a unit hypersphere. Similarly, each image \mathbf{x}_i is processed by the visual encoder θ_v to obtain the normalized visual embedding $\mathbf{f}_i = \theta_v(\mathbf{x}_i)$.

Zero-shot prediction is the simplest method for adapting a VLM to downstream tasks, leveraging the model’s pretraining (Radford et al. 2021). For a test image \mathbf{x}_i , the cosine similarity between its visual embedding \mathbf{f}_i and each class’s textual embedding \mathbf{t}_k is calculated to obtain the prediction logit: $l_{i,k} = \mathbf{f}_i^\top \mathbf{t}_k$. These logits are then transformed into posterior probabilities using the softmax function with temperature scaling τ : $p_{i,k} = \exp(l_{i,k}/\tau) / \sum_{j=1}^K \exp(l_{i,j}/\tau)$.

The class with the highest posterior probability $p_{i,k}$ is selected as the predicted label for the image \mathbf{x}_i . This approach

allows the VLM to perform classification without requiring additional training, relying solely on the alignment between textual and visual embeddings.

During fine-tuning on downstream tasks, LoRA (Hu et al. 2021) models the incremental update of pre-trained weights as the product of two small matrices, following the concept of the ‘‘intrinsic rank’’ of a downstream task. In our method, we apply LoRA to both text and visual branches of the CLIP model to fine-tune the original CLIP model, following the settings defined in (Zanella and Ben Ayed 2024).

Empty Prompts Generation

To examine CLIP’s inherent biases toward specific template structures, we generate empty prompts by combining standard templates with inputs that lack meaningful content. These empty prompts enable us to identify and understand the biases introduced by the templates. The bias information extracted from these prompts is then used to calibrate the CLIP model, thereby enhancing its fairness and accuracy.

Creating a diverse set of meaningless inputs, denoted as T_{null} , is crucial to prevent overfitting to particular instances and to ensure effective calibration. First, we select common template structures used in CLIP-based classification tasks, such as ‘‘a photo of a {’’’, then generate 100 words/phrases representing ‘‘empty’’ using GPT-4. For each candidate, we compute the average similarity to all category names and select the top-25 with highest similarity. This ensures selected words lie in the same feature space as category names and remain semantically ambiguous.

The selected 25 words are used to fill these templates, resulting in prompts like ‘‘a photo of nothing’’ or ‘‘a photo of emptiness’’. We then input these empty prompts into the CLIP model and analyze its responses to detect any inherent biases related to the template structures. Finally, the bias information derived from the empty prompts is used to adjust and calibrate the CLIP model, thereby reducing unwanted template-induced biases.

Template Bias Calibration Loss

As shown in Figure 1, the final classification results are influenced by both the template and the class name because the similarity between the template and the samples varies. To reduce this bias, we introduce a constraint that aligns the similarity relationships between the samples and the template during the classification process.

Figure 3 illustrates our proposed method for removing this bias. We define \mathcal{M} as the current model, and \mathbf{t}_i^E as the normalized textual embedding for the empty prompt (‘‘a photo of a $T_{null,i}$ ’’). Ideally, the similarity between \mathbf{t}_i^E and different samples should be consistent. Specifically, if each sample is treated as a separate class in the class set \mathcal{Y}_x , using the samples’ visual features as classifier weights and \mathbf{t}_i^E as the query feature, the classification should satisfy the following condition:

$$\mathbb{E}_{\mathbf{t}_i^E} [P(y | \mathbf{t}_i^E, \mathcal{M}; \forall y \in \mathcal{Y}_x)] = \frac{1}{|\mathcal{Y}_x|} \quad (4)$$

This means that the expected output distribution for \mathbf{t}_i^E should be uniform across all labels. Let $p_{i,k}^E =$

$P(y_k | \mathbf{t}_i^E, \mathcal{M}; \forall y_k \in \mathcal{Y}_x)$. To correct for the template bias, we design a loss function L_{tb} based on the above equation:

$$L_{tb} = -\frac{1}{|T_{null}|} \sum_{i=1}^{|T_{null}|} \sum_{k=1}^{|\mathcal{Y}_x|} \frac{1}{|\mathcal{Y}_x|} \ln p_{i,k}^E \quad (5)$$

L_{tb} ensures that each empty prompt is classified with equal probability for each category associated with the sample. This approach guarantees that the similarity between the instances and the template remains uniform.

Training Process

The training of our model is carried out in two distinct stages: pretraining initialization and few-shot finetuning calibration.

Pretraining Initialization. In the first stage, we perform unsupervised pre-training using the loss function $L_{pre} = L_{tb}$. The main objective of this stage is to initialize the model parameters without any bias. By minimizing L_{tb} , we ensure that the model starts with unbiased initial parameters, which is crucial for preventing template-induced biases from affecting the learning process.

Few-shot Finetuning calibration. In the second stage, we fine-tune the model using a few labeled samples. During this phase, we apply the cross-entropy loss function, as commonly used in (Zhou et al. 2022b,a; Zanella and Ben Ayed 2024), to adapt the model to specific tasks:

$$L_{ce} = -\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K y_{ik} \ln p_{i,k} \quad (6)$$

Here, N is the number of samples, K is the number of classes, y_{ik} is the ground truth label, and $p_{i,k}$ is the predicted probability for class k of sample i .

In addition to L_{ce} , we retain the template bias calibration loss L_{tb} during fine-tuning to maintain the model’s unbiasedness with respect to the template. The total loss function for fine-tuning is therefore defined as:

$$L_{fine} = L_{ce} + \alpha L_{tb} \quad (7)$$

where α is a hyperparameter that controls the weight of the bias calibration loss relative to the cross-entropy loss.

This two-stage training process ensures that the model not only learns to perform accurately on the target task with limited labeled data but also remains free from biases introduced by the template structures. By first establishing unbiased initial parameters and then carefully fine-tuning with a balanced loss function, our approach achieves both fairness and high performance in classification tasks. Please refer to the Appendix for the pseudocode of the algorithm.

Experiments

Experimental Setup

Dataset. Following the setup of previous work (Zhou et al. 2022b; Zanella and Ben Ayed 2024), we evaluate our approach on 11 datasets covering various fine-grained classification tasks: scenes (SUN397 (Xiao et al. 2010)), aircraft types (Aircraft (Maji et al. 2013)), satellite imagery

Shots	Method	ImageNet	SUN	Aircraft	EuroSAT	Cars	Food	Pets	Flowers	Caltech	DTD	UCF	Average
1	CoOp (4)	68.0	67.3	26.2	50.9	67.1	82.6	90.3	72.7	93.2	50.1	70.7	67.2
	CoOp (16)	65.7	67.0	20.8	56.4	67.5	84.3	90.2	78.3	92.5	50.1	71.2	67.6
	CoCoOp	69.4	68.7	28.1	55.4	67.6	84.9	91.9	73.4	94.1	52.6	70.4	68.8
	TIP-Adapter-F	69.4	67.2	28.8	67.8	67.1	85.8	90.6	83.8	94.0	51.6	73.4	70.9
	CLIP-Adapter	67.9	65.4	25.2	49.3	65.7	86.1	89.0	71.3	92.0	44.2	66.9	65.7
	PLOT++	66.5	66.8	28.6	65.4	68.8	86.2	91.9	80.5	94.3	54.6	74.3	70.7
	KgCoOp	68.9	68.4	26.8	61.9	66.7	86.4	92.1	74.7	94.2	52.7	72.8	69.6
	TaskRes	69.6	68.1	31.3	65.4	68.8	84.6	90.2	81.7	93.6	53.8	71.7	70.8
	MaPLe	69.7	69.3	28.1	29.1	67.6	85.4	91.4	74.9	93.6	50.0	71.1	66.4
	ProGrad	67.0	67.0	28.8	57.0	68.2	84.9	91.4	80.9	93.5	52.8	73.3	69.5
	LoRA-CLIP	70.4	70.4	30.2	72.3	70.1	84.3	92.3	83.2	93.7	54.3	76.3	72.5
OURS	70.5	70.5	32.2	78.4	70.6	85.7	92.9	83.3	94.8	54.5	76.0	73.6	
4	CoOp (4)	69.7	70.6	29.7	65.8	73.4	83.5	92.3	86.6	94.5	58.5	78.1	73.0
	CoOp (16)	68.8	69.7	30.9	69.7	74.4	84.5	92.5	92.2	94.5	59.5	77.6	74.0
	CoCoOp	70.6	70.4	30.6	61.7	69.5	86.3	92.7	81.5	94.8	55.7	75.3	71.7
	TIP-Adapter-F	70.7	70.8	35.7	76.8	74.1	86.5	91.9	92.1	94.8	59.8	78.1	75.6
	CLIP-Adapter	68.6	68.0	27.9	51.2	67.5	86.5	90.8	73.1	94.0	46.1	70.6	67.7
	PLOT++	70.4	71.7	35.3	83.2	76.3	86.5	92.6	92.9	95.1	62.4	79.8	76.9
	KgCoOp	69.9	71.5	32.2	71.8	69.5	86.9	92.6	87.0	95.0	58.7	77.6	73.9
	TaskRes	71.0	72.7	33.4	74.2	76.0	86.0	91.9	85.0	95.0	60.1	76.2	74.7
	MaPLe	70.6	71.4	30.1	69.9	70.1	86.7	93.3	84.9	95.0	59.0	77.1	73.5
	ProGrad	70.2	71.7	34.1	69.6	75.0	85.4	92.1	91.1	94.4	59.7	77.9	74.7
	LoRA-CLIP	71.4	72.8	37.9	84.9	77.4	82.7	91.0	93.7	95.2	63.8	81.1	77.4
OURS	71.5	74.0	40.2	87.6	77.7	87.1	93.7	94.8	95.7	65.6	82.2	79.1	
8	CoOp(4)	70.8	72.4	37.0	74.7	76.8	83.3	92.1	95.0	94.7	63.7	79.8	76.4
	CoOp(16)	70.6	71.9	38.5	77.1	79.0	82.7	91.3	94.9	94.5	64.8	80.0	76.8
	CoCoOp	70.8	71.5	32.4	69.1	70.4	87.0	93.3	86.3	94.9	60.1	75.9	73.8
	TIP-Adapter-F	71.7	73.5	39.5	81.3	78.3	86.9	91.8	94.3	95.2	66.7	82.0	78.3
	CLIP-Adapter	69.1	71.7	30.5	61.6	70.7	86.9	91.9	83.3	94.5	50.5	76.2	71.5
	PLOT++	71.3	73.9	41.4	88.4	81.3	86.6	93.0	95.4	95.5	66.5	82.8	79.6
	KgCoOp	70.2	72.6	34.8	73.9	72.8	87.0	93.0	91.5	95.1	65.6	80.0	76.0
	TaskRes	72.3	74.6	40.3	77.5	79.6	86.4	92.0	96.0	95.3	66.7	81.6	78.4
	MaPLe	71.3	73.2	33.8	82.8	71.3	87.2	93.1	90.5	95.1	63.0	79.5	76.4
	ProGrad	71.3	73.0	37.7	77.8	78.7	86.1	92.2	95.0	94.8	63.9	80.5	77.4
	CLIP-LoRA	72.3	74.7	45.7	89.7	82.1	83.1	91.7	96.3	95.6	67.5	84.1	80.3
OURS	72.3	75.5	49.7	91.2	82.5	87.3	94.2	96.9	95.9	69.2	85.0	81.8	

Table 2: Detailed results for 11 datasets with the ViT-B/16 as visual backbone. Top-1 accuracy averaged over 3 random seeds is reported. Highest value is high lighted in bold.

(EuroSAT (Helber et al. 2019)), automobiles (Stanford-Cars (Krause et al. 2013)), food items (Food101 (Bossard, Guillaumin, and Van Gool 2014)), pet breeds (Oxford-Pets (Parkhi et al. 2012)), flowers (Flower102 (Nilsback and Zisserman 2008)), general objects (Caltech101 (Fei-Fei, Fergus, and Perona 2004)), textures (DTD (Cimpoi et al. 2014)), and human actions (UCF101 (Soomro 2012)), in addition to ImageNet (Deng et al. 2009). These datasets provide a comprehensive benchmarking framework for evaluating few-shot visual classification tasks.

Comparative Methods. We compare our model against several prompt-based methods, including CoOp (Zhou et al. 2022b) with 4 learnable tokens, CoOp (Zhou et al. 2022b) with 16 learnable tokens, CoCoOp (Zhou et al. 2022a), PLOT++ (Chen et al. 2022) (an adaptation of the original PLOT designed for transformer architectures), Kg-CoOp (Yao, Zhang, and Xu 2023), MaPLe (Khattak et al. 2023), and ProGrad (Zhu et al. 2023a) with 16 tokens. Additionally, we evaluate adapter-based methods such as Tip-Adapter-F (Zhang et al. 2022) and TaskRes (Yu et al. 2023), as well as the LoRA-based method CLIP-LoRA (Zanella and Ben Ayed 2024). These comparative methods provide a

comprehensive benchmark to assess the effectiveness of our proposed approach in few-shot visual classification tasks.

Few-shot Learning Result

Table 2 presents the few-shot learning performance of various models across 11 datasets. Among the baselines, LoRA is the most competitive, achieving the second-best overall performance. Our model, which adds template bias correction to LoRA, shows superior or comparable results on most datasets, with significant improvements on EuroSAT, Food101, OxfordPet, DTD, and Aircraft. On average, our model improves performance by approximately 1.5 points over LoRA across the 11 datasets.

Ablation Study

We conduct ablation study to evaluate the contribution of each module, as shown in Table 3. The modules are defined as follows: **Pretrain**: The pretraining initialization phase; **Multiple**: The use of multiple empty prompts to create a diverse set of prompts; **+ L_{tb}** : The incorporation of the L_{tb} loss during the few-shot fine-tuning phase. Our results demonstrate that each module enhances performance.

Pretrain	-	-	-	✓	✓	✓
Multiple	-	-	✓	-	✓	✓
+ L_{tb}	-	✓	✓	-	-	✓
1-shot	72.3	73.9	76.2	73.3	74.2	78.4
4-shot	84.9	85.8	86.9	86.4	87.9	87.6

Table 3: Ablating main components of our model on EuroSAT.

Method	ImageNet	SUN	Aircraft	EuroSAT	DTD	Average
Baseline	70.4	70.4	30.2	72.3	54.3	59.5
Pull Closer	69.5	69.3	30.0	71.6	51.9	58.5
Push Away	68.4	65.8	7.9	62.9	48.4	50.7
Ours	70.5	70.5	32.2	78.3	54.5	61.2

Table 4: Evaluation of three TSS modification strategies.

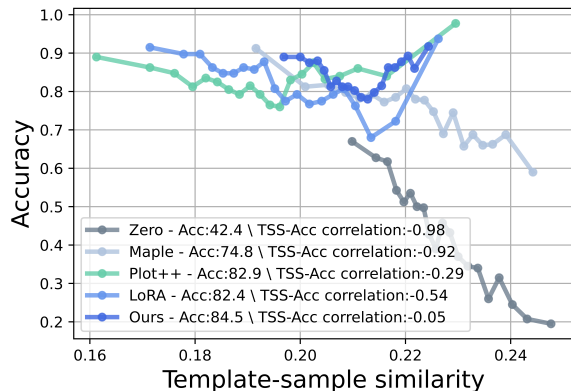


Figure 4: The relationship between template-sample similarity and classification accuracy under different training methods. Result for 4-shot finetuning on EuroSAT, seed 1.

In our method, we eliminate the bias introduced by the template by ensuring a consistent distance between each sample and the “empty template” (i.e., enforcing a consistent TSS). To demonstrate the effectiveness of this design, we further compared two alternative strategies for modifying TSS: **Pull Closer**, which reduces the distance between each sample and the “empty template” (increasing TSS), and **Push Away**, which increases the distance (decreasing TSS). The results, shown in Table 4, indicate that neither the Pull Closer nor the Push Away strategies lead to performance improvements. In contrast, our method effectively enhances the model’s performance.

Mitigating Template-Induced Bias

As discussed in the introduction, template-induced bias can significantly impact the performance of few-shot learning models. In this section, we provide a detailed analysis to evaluate how effectively our approach mitigates this bias. As shown in Figure 4, the zero-shot CLIP model exhibits a strong negative correlation (-0.98) between TSS and classification accuracy. Training with LoRA reduces this correlation to -0.54 while improving performance. In contrast,

our method nearly eliminates this correlation (-0.05) and achieves the highest classification accuracy among the compared models. Figure 4 shows that the zero-shot CLIP model has a strong negative correlation (-0.98) between TSS and accuracy. LoRA training reduces this correlation to -0.54 and improves performance. Our method almost eliminates this correlation (-0.05) and achieves the highest accuracy. Figure 2 (Middle) shows our method provides a lower-bias initialization, with a TSS-accuracy correlation of 0.38 at iteration 0, compared to the original model’s 0.85. Throughout training, our model maintains a lower correlation, showing that the constraint loss reduces template bias and focuses on sample-specific information, leading to a 2-point improvement over LoRA-only training. Figure 2 (Right) illustrates the impact of templates on performance. LoRA-trained models are affected by template bias, while our model remains robust to template variations and consistently outperforms LoRA-trained models. Our method effectively mitigates template-induced bias, improving performance.

Related Work

VLMs combine visual and textual data, making bias evaluation particularly challenging. Models like CLIP (Radford et al. 2021) and LLaVA (Liu et al. 2024) can inherit biases from the images and texts used during training. These biases arise from both visual and textual sources, affecting how the model behaves. Several studies (Fabrizzi et al. 2022; Cabello et al. 2023; Seth, Hemani, and Agarwal 2023; Zhang et al. 2024b) have examined biases in VLMs caused by imbalanced datasets and stereotypes, leading to biased representations of gender, race, and professions. To evaluate these biases, datasets such as PAIRS (Fraser and Kiritchenko 2024), which includes images of individuals from diverse racial and gender backgrounds, and VisoGender (Hall et al. 2024), which looks at gender bias in pronoun resolution, have been created. Other research uses specially designed inputs to detect biases in VLMs. For instance, (Leng et al. 2024) investigates bias in the text branch by introducing noisy samples to the visual branch. Similarly, (Zhang et al. 2024a) examines bias in the text branch by providing empty visual inputs. Additionally, (Chuang et al. 2023) uses prompts containing biased terms to identify biased directions in text embeddings. While these studies focus on the inherent biases within VLMs, few have explored biases that may arise when VLMs are applied to downstream tasks (Zhang et al. 2024c,b; Zou et al. 2022, 2024b,a, 2021; Xiao et al. 2025; Zeng et al. 2025; Chen et al. 2025; Zhang et al. 2024c), such as those introduced by templates. To our knowledge, this study is the first to identify and investigate the biases caused by the template.

Conclusion

In conclusion, we reveal that TSS introduces significant bias in CLIP. To address this issue, we propose a novel method that uses “empty” prompts to correct the bias introduced by TSS. Extensive experiments demonstrate the effectiveness of our method, highlighting the importance of addressing template-induced biases for improved few-shot learning.

Acknowledgments

This work is supported by the National Key Research and Development Program of China under grant 2024YFC3307900; the National Natural Science Foundation of China under grants 62436003, 62206102, 62376103, 62302184, and 62402015; the Major Science and Technology Project of Hubei Province under grant 2025BAB011 and 2024BAA008; the Hubei Science and Technology Talent Service Project under grant 2024DJC078; Ant Group through the CCF–Ant Research Fund; the Postdoctoral Fellowship Program of the China Postdoctoral Science Foundation under grant GZB20230024; and the China Postdoctoral Science Foundation under grant 2024M750100. Computations were performed on the HPC Platform of Huazhong University of Science and Technology.

References

- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101—mining discriminative components with random forests. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part VI 13*, 446–461. Springer.
- Cabello, L.; Bugliarello, E.; Brandl, S.; and Elliott, D. 2023. Evaluating bias and fairness in gender-neutral pretrained vision-and-language models. *arXiv preprint arXiv:2310.17530*.
- Chen, G.; Horstmann, K.; Wang, Z.; and You, F. 2025. Automated Essential Concept Discovery for Few-Shot Out-of-Distribution Detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 3964–3974.
- Chen, G.; Yao, W.; Song, X.; Li, X.; Rao, Y.; and Zhang, K. 2022. Plot: Prompt learning with optimal transport for vision-language models. *arXiv preprint arXiv:2210.01253*.
- Chen, H.; Raj, B.; Xie, X.; and Wang, J. 2024. On Catastrophic Inheritance of Large Foundation Models. *Journal of Data-centric Machine Learning Research*.
- Chuang, C.-Y.; Jampani, V.; Li, Y.; Torralba, A.; and Jegelka, S. 2023. Debiasing vision-language models via biased prompts. *arXiv preprint arXiv:2302.00070*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3606–3613.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Fabbrizzi, S.; Papadopoulos, S.; Ntoutsis, E.; and Kompatsiaris, I. 2022. A survey on bias in visual datasets. *Computer Vision and Image Understanding*, 223: 103552.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, 178–178. IEEE.
- Fraser, K. C.; and Kiritchenko, S. 2024. Examining Gender and Racial Bias in Large Vision-Language Models Using a Novel Dataset of Parallel Images. *arXiv preprint arXiv:2402.05779*.
- Hall, S. M.; Gonçalves Abrantes, F.; Zhu, H.; Sodunke, G.; Shtedritski, A.; and Kirk, H. R. 2024. Visogender: A dataset for benchmarking gender bias in image-text pronoun resolution. *Advances in Neural Information Processing Systems*, 36.
- He, J.; Rungta, M.; Koleczek, D.; Sekhon, A.; Wang, F. X.; and Hasan, S. 2024. Does Prompt Formatting Have Any Impact on LLM Performance? *arXiv preprint arXiv:2411.10541*.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7): 2217–2226.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Ke, Q.; Bennamoun, M.; An, S.; Sohel, F.; and Boussaid, F. 2018. Learning clip representations for skeleton-based 3D action recognition. *IEEE Transactions on Image Processing*, 27(6): 2842–2855.
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19113–19122.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, 554–561.
- Leng, S.; Zhang, H.; Chen, G.; Li, X.; Lu, S.; Miao, C.; and Bing, L. 2024. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13872–13882.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liusie, A.; Manakul, P.; and Gales, M. 2023. Mitigating Word Bias in Zero-shot Prompt-based Classifiers. In *Findings of the Association for Computational Linguistics: IJCNLP-AAACL 2023 (Findings)*, 327–335.
- Maji, S.; Rahtu, E.; Kannala, J.; Blaschko, M.; and Vedaldi, A. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*.
- Materzyńska, J.; Torralba, A.; and Bau, D. 2022. Disentangling visual and written concepts in CLIP. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16410–16419.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, 722–729. IEEE.

- Parkhi, O. M.; Vedaldi, A.; Zisserman, A.; and Jawahar, C. 2012. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, 3498–3505. IEEE.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Seth, A.; Hemani, M.; and Agarwal, C. 2023. Dear: Debiasing vision-language models with additive residuals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6820–6829.
- Shin, C.; Zhao, J.; Crompton, S.; Vishwakarma, H.; and Sala, F. 2024. OTTER: Effortless Label Distribution Adaptation of Zero-shot Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Soomro, K. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Vivanco Cepeda, V.; Nayak, G. K.; and Shah, M. 2024. Geoclip: Clip-inspired alignment between locations and images for effective worldwide geo-localization. *Advances in Neural Information Processing Systems*, 36.
- Wen, X.; Zhao, B.; Chen, Y.; Pang, J.; and Qi, X. 2024. What Makes CLIP More Robust to Long-Tailed Pre-Training Data? A Controlled Study for Transferable Insights. *Advances in Neural Information Processing Systems*, 37: 36567–36601.
- Xiao, D.; Chen, G.; Peng, P.; Huang, Y.; Zhao, Y.; Dai, Y.; and Tian, Y. 2025. When Every Millisecond Counts: Real-Time Anomaly Detection via the Multimodal Asynchronous Hybrid Network. In *Forty-second International Conference on Machine Learning*.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, 3485–3492. IEEE.
- Yang, Y.; Nushi, B.; Palangi, H.; and Mirzasoleiman, B. 2023. Mitigating spurious correlations in multi-modal models during fine-tuning. In *International Conference on Machine Learning*, 39365–39379. PMLR.
- Yao, H.; Zhang, R.; and Xu, C. 2023. Visual-language prompt tuning with knowledge-guided context optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6757–6767.
- You, C.; Mint, Y.; Dai, W.; Sekhon, J. S.; Staib, L.; and Duncan, J. S. 2024. Calibrating multi-modal representations: A pursuit of group robustness without annotations. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 26140–26150. IEEE.
- Yu, T.; Lu, Z.; Jin, X.; Chen, Z.; and Wang, X. 2023. Task residual for tuning vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10899–10909.
- Zanella, M.; and Ben Ayed, I. 2024. Low-Rank Few-Shot Adaptation of Vision-Language Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1593–1603.
- Zeng, Y.; Wu, H.; Nie, W.; Chen, G.; Zheng, X.; Shen, Y.; Peng, J.; Tian, Y.; and Ji, R. 2025. From Objects to Events: Unlocking Complex Visual Understanding in Object Detectors via LLM-guided Symbolic Reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 24380–24391.
- Zhang, M.; and Ré, C. 2022. Contrastive adapters for foundation model group robustness. *Advances in Neural Information Processing Systems*, 35: 21682–21697.
- Zhang, R.; Zhang, W.; Fang, R.; Gao, P.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2022. Tip-adapter: Training-free adaption of clip for few-shot classification. In *European conference on computer vision*, 493–510. Springer.
- Zhang, Y.-F.; Yu, W.; Wen, Q.; Wang, X.; Zhang, Z.; Wang, L.; Jin, R.; and Tan, T. 2024a. Debiasing multimodal large language models. *CoRR*.
- Zhang, Z.; Chen, G.; Zou, Y.; Huang, Z.; Li, Y.; and Li, R. 2024b. Micm: Rethinking unsupervised pretraining for enhanced few-shot learning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7686–7695.
- Zhang, Z.; Chen, G.; Zou, Y.; Li, Y.; and Li, R. 2024c. Learning Unknowns from Unknowns: Diversified Negative Prototypes Generator for Few-Shot Open-Set Recognition. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 6053–6062.
- 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.
- Zhu, B.; Niu, Y.; Han, Y.; Wu, Y.; and Zhang, H. 2023a. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15659–15669.
- Zhu, B.; Tang, K.; Sun, Q.; and Zhang, H. 2023b. Generalized logit adjustment: Calibrating fine-tuned models by removing label bias in foundation models. *Advances in Neural Information Processing Systems*, 36: 64663–64680.
- Zou, Y.; Liu, Y.; Hu, Y.; Li, Y.; and Li, R. 2024a. Flatten long-range loss landscapes for cross-domain few-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23575–23584.
- Zou, Y.; Yi, S.; Li, Y.; and Li, R. 2024b. A closer look at the CLS token for cross-domain few-shot learning. *Advances in Neural Information Processing Systems*, 37: 85523–85545.
- Zou, Y.; Zhang, S.; Chen, G.; Tian, Y.; Keutzer, K.; and Moura, J. M. 2021. Annotation-efficient untrimmed video action recognition. In *Proceedings of the 29th ACM International Conference on Multimedia*, 487–495.
- Zou, Y.; Zhang, S.; Li, Y.; and Li, R. 2022. Margin-based few-shot class-incremental learning with class-level overfitting mitigation. *Advances in neural information processing systems*, 35: 27267–27279.