

PromptSync: Bridging Domain Gaps in Vision-Language Models through Class-Aware Prototype Alignment and Discrimination

Anant Khandelwal

Glance AI

anant.iitd.2085@gmail.com

Abstract

The potential for zero-shot generalization in vision-language (V-L) models such as CLIP has spurred their widespread adoption in addressing numerous downstream tasks. Previous methods have employed test-time prompt tuning to adapt the model to unseen domains, but they overlooked the issue of imbalanced class distributions. In this study, we explicitly address this problem by employing class-aware prototype alignment weighted by mean class probabilities obtained for the test sample and filtered augmented views. Additionally, we ensure that the class probabilities are as accurate as possible by performing prototype discrimination using contrastive learning. The combination of alignment and discriminative loss serves as a geometric regularizer, preventing the prompt representation from collapsing onto a single class and effectively bridging the distribution gap between the source and test domains. Our method, named PromptSync, synchronizes the prompts for each test sample on both the text and vision branches of the V-L model. In empirical evaluations on the domain generalization benchmark, our method outperforms previous best methods by 2.33% in overall performance, by 1% in base-to-novel generalization, and by 2.84% in cross-dataset transfer tasks.

1. Introduction

Training Vision-Language Models (VLMs) with large-scale image-text pairs is known for imparting robust generalization capabilities across diverse downstream tasks [1, 20, 31, 41, 42, 44]. However, training these models from scratch for each downstream task is very time-consuming. Moreover, the essence of pre-training with a large-scale dataset is lost when the pre-trained model is not generalizable across downstream tasks. This is due to unexpected changes in data distribution, and the sensitivity to these shifts leads to a decline in performance [16, 30, 32]. To tackle this, there exist three most commonly used techniques: fine-tuning [28],

prompt tuning [47], adapter [13], and LoRA [19]. Among these, prompt tuning is the simple, recent, and most widely used technique for foundation models [21, 22, 43, 46, 47]. However, prompt learning/tuning approaches are used during the training phase to learn representative prompts based on the training data for the downstream task. This approach does not specifically address the distribution shift present in the dataset. Recent methods, TPT [35] and PromptAlign [33], adjust the learnable prompt tokens dynamically during testing to enable test-time adaptation and align the context of the test sample as per the seen distribution by the model. Specifically, TPT [35] updates the learnable prompt tokens (keeping the model parameters frozen) by minimizing the entropy of top- N_k confidently predicted samples, acquired through diverse augmented views of the incoming test sample. Additionally, PromptAlign [33] aligns token distribution of the test sample in the visual branch with the pre-computed statistics of the complete proxy source dataset irrespective of the fact that one class distribution may have different mean and variance than the other classes.

In this work, we demonstrate the multi-modal test-time adaptation of prompts. In contrast to PromptAlign, which aligns the distribution for the complete source dataset with test sample, we propose class-aware prototype alignment to address the distributional shift on a class-wise basis. For instance, in an open world there are 360 different breeds of dogs compared to only 71 for cats, leading to one class having higher variance than the others. For each test sample, we obtain randomly augmented views (for both text and image) that are fed to the model for prompt tuning on both the textual and visual branches. We adapt the learnable prompt tokens by aligning the prototype for test sample and confident augmented views with the pre-computed class prototypes (obtained from the proxy source dataset) weighted by the mean probability of each class obtained from confident augmented views. Before alignment, we update the prompt tokens on both the text and visual branches using prototype discrimination and then use updated prompts to align the test sample and augmented views with class prototypes using mean class probabilities. This is based on the idea that pro-

prototype vector can capture the complete information of mean and variance for each class distribution and hence it mitigates the class collapse (during test time adaptation) due to high variance of particular classes. Empirical evaluation of our methods shows state-of-the-art Top-1 accuracy for three tasks: domain generalization, base-to-novel generalization, and cross-dataset transfer. This validates the effectiveness of our method in enhancing zero-shot generalization. Our contributions can be summarized as follows:

- We propose a class-aware prototype alignment technique for individual test samples to align the context of each test sample with the source distribution on a class-wise basis, thereby mitigating the effects of distributional shift between classes.
- We propose class-aware prototype discrimination to discover the class distribution for efficient alignment. Additionally, we propose the offline computation of class prototypes from a proxy source dataset for foundation V-L models.
- We propose multi-modal test-time prompt tuning for both text and visual branches. Empirical evaluation on base-to-novel generalization, domain generalization, and cross-dataset transfer shows the efficiency of our method over existing methods.

2. Related Work

Vision-Language (V-L) foundation models like CLIP [31] and ALIGN [20] have emerged as robust zero-shot generalizable models. They integrate image and text modalities through pre-training on extensive image-text pairs. However, adapting these models to specific downstream tasks with limited data remains challenging. Recent methods explore prompt tuning in CLIP-like models, treating prompts as continuous learnable vectors and fine-tuning them while keeping the model parameters frozen. CoOp [47] proposed fine-tuning CLIP by learning a set of prompts in the text encoder. CoCoOp [46], an improvement over CoOp, dynamically conditions the text prompts by the image embeddings. MaPLe [21] is a deep prompting baseline that tunes prompts on both text and image branches, further conditioning image prompts on text prompts using a V-L coupling function. However, these approaches necessitate training data for prompt learning, limiting adaptation to novel datasets during test time. Recent approaches like TPT [35] aim to learn prompts exclusively at test time but encounter challenges in handling distribution misalignment between CLIP’s pre-training data and downstream test data. PromptAlign [33] addresses this by introducing token distribution alignment in the image branch. However, it does not account for the potential variance in class distributions. In contrast, our method, inspired by a multi-modal prompting variant [21], actively aligns class prototypes by leveraging a proxy dataset as a substitute for unavailable CLIP pre-training data. To

our knowledge, our approach is the first to explicitly address class-aware distribution misalignment in V-L foundational models during test time.

3. Methodology

Revisiting CLIP: Our approach is based on the pre-trained V-L model: Contrastive Language-Image Pre-Training (CLIP). It consists of a text and visual encoder (denoted by \mathcal{F}_t and \mathcal{F}_v , respectively, and their pre-trained parameters are represented by $\theta_{\text{CLIP}} = \{\theta_t, \theta_v\}$, respectively), used for mapping the text and image to the vector representation, respectively. The input image is \mathbf{X} , which is divided into M patches, and the [CLS] token is prepended to these M patch tokens that are projected to produce $\tilde{\mathbf{X}}_v = \{e_{[\text{CLS}]}, e_1, e_2, \dots, e_M\}$, where e_i is the embedding for the corresponding patch token in \mathbf{X} . The image encoder produces latent visual feature representation $\tilde{\mathbf{f}}_v = \mathcal{F}_v(\tilde{\mathbf{X}}_v, \theta_v)$ with transformer blocks from $\tilde{\mathbf{X}}_v$. The class label y is embedded within a text template, such as “a photo of a <CLS>” resulting in $\tilde{\mathbf{X}}_t = \{SOS, \mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_L, \mathbf{c}_k, EOS\}$, where SOS and EOS are the start and end token embeddings and \mathbf{t}_l and \mathbf{c}_k are the token embeddings corresponding to the text template and the class label, respectively. Similarly, the text encoder \mathcal{F}_t encodes $\tilde{\mathbf{X}}_t$ with transformer blocks to produce latent text feature representation $\tilde{\mathbf{f}}_t = \mathcal{F}_t(\tilde{\mathbf{X}}_t, \theta_t)$. For zero-shot inference, each text feature for class labels $y = \{1, 2, \dots, C\}$ is paired with an image feature to compute the similarity score $s_i = \text{sim}(\tilde{\mathbf{f}}_{t_i} \cdot \tilde{\mathbf{f}}_v)$ where $\text{sim}(\cdot)$ denotes cosine similarity. The predicted probability on \mathbf{X} for each y_i is given as $p(y_i|\mathbf{X}) = \frac{e^{\text{sim}(\tilde{\mathbf{f}}_{t_i} \cdot \tilde{\mathbf{f}}_v)/\tau}}{\sum_{j=1}^C e^{\text{sim}(\tilde{\mathbf{f}}_{t_j} \cdot \tilde{\mathbf{f}}_v)/\tau}}$, where τ is the temperature of softmax.

Prompt Tuning: CLIP integrates a considerable pool of knowledge derived from its training on millions of image-text pairs characterized by varying degrees of noise. Prompt tuning methods aim to extract the rich features learned by the CLIP model. Recent approaches [3, 21, 43, 46, 47] append extra learnable prompts to the input of image and text encoders while keeping them frozen. Modified input prompts with frozen encoders generate undistorted and rich CLIP features, where prompt tuning tries to map the context to the source distribution, i.e., the CLIP pre-training dataset. In our work, we use a recent multi-modal prompting baseline [21] where prompt tuning is performed on both the text and image encoders. Specifically, the image and text encoders process the input $\tilde{\mathbf{X}}_v = \{e_{[\text{CLS}]}, \mathbf{p}_v, e_1, e_2, \dots, e_M\}$ and $\tilde{\mathbf{X}}_t = \{SOS, \mathbf{p}_t, \mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_L, \mathbf{c}_k, EOS\}$ respectively. The learnable prompts \mathbf{p}_v and \mathbf{p}_t represent the V visual and T textual tokens, respectively. We will call prompts \mathbf{p}_t and \mathbf{p}_v as \mathbf{p} only. Our approach is based on deep prompting, as in [21], along with text and image prompts at subsequent transformer blocks. We suggest referring to [21] for more details on baseline architecture.

Test Time Adaptation: Test-time adaptation aims to boost generalisation in a zero-shot manner. Existing methods, Test time prompt tuning (TPT)[35] and PromptAlign[33], both are introduced to provide the model context that is customized for each individual test sample in order to extract rich knowledge from CLIP. For both methods, several augmented views $\mathcal{H}(\mathbf{X}_{test})$ are generated from the given test sample \mathbf{X}_{test} . The average entropy for the filtered views (selected using a confidence threshold) is then used to update the prompts \mathbf{p} using the following unsupervised objective:

$$\mathcal{L}_{ent} = \operatorname{argmin}_{\mathbf{p}} - \sum_{i=1}^C \tilde{p}_{\mathbf{p}}(y_i | \mathbf{X}_{test}) \log \tilde{p}_{\mathbf{p}}(y_i | \mathbf{X}_{test}) \quad (1)$$

where $\tilde{p}_{\mathbf{p}}(y_i | \mathbf{X}_{test})$ is the average of vector class probabilities (over the filtered augmented views) produced by the model. Additionally, PromptAlign uses distribution alignment loss, which aligns the mean and variance of filtered augmented views of the test sample with source statistics across layers of the model.

3.1. Proposed Method: PromptSync

The multi-modal test-time prompt tuning method, PromptAlign [33], updates text and visual prompts using entropy loss and distribution alignment loss with highly confident augmented views (obtained from a test sample \mathbf{X}_{test}). PromptAlign, despite considering the distribution, does not take into account the fact that the distribution of each class/domain can be entirely different from other classes/domains, and hence using the source statistics of mean and variance for distribution alignment can still be suboptimal. Inspired by prototype learning [38] and Extreme-Multi-PatchSSL (EMPSSL) [39], which establish a prototype/benchmark for each class/sample, we propose class-wise prototype alignment between original and augmented views for both source and test samples. The architecture of PromptSync is shown in Figure 1. We use the parameter update from prototype discrimination to generate the class probabilities for the test sample and its augmented views. We accumulate the average of gradients from prototype alignment loss weighted by class probabilities for confident augmented views. The accumulated gradient over multiple iterations is then applied for prompt tuning during test-time adaptation.

3.2. Class-aware Prototype Generation

We generated prototypes for each class for both text and visual branches. The prototype for each class is computed using proxy source dataset. For a test sample, \mathbf{X}_{test} and its N_k random views (generated using a set of augmentations \mathcal{H} on \mathbf{X}_{test}) the prototype vector is generated. Let's denote the token e_i features of a sample $x \in \{\mathbf{X}_{test} + \mathcal{H}(\mathbf{X}_{test})\}$, at the output of the text encoder and visual encoder as $E_T(x, e_i)$ and $E_V(x, e_i)$, respectively. The prototype for a sample

from text and visual branches is given as:

$$h_x^{\{t,v\}} = \frac{1}{|P|} \sum_{i=1}^{|P|} E_{\{T,V\}}(x, e_i) \quad (2)$$

$$h_{cls,x}^v = E_V(x, e_{cls})$$

where $|P|$ represents the total number of tokens (learnable and non-learnable for both text and visual), excluding EOS, SOS, and CLS. t, v represents textual branch and visual branch respectively. For the proxy source dataset, the class-aware prototype is obtained as:

$$h_{c_k}^t = \frac{1}{|\mathcal{D}(c_k)|} \sum_{x \in \mathcal{D}(c_k)} h_x^t, \quad (3)$$

$$h_{c_k}^v = \frac{1}{|\mathcal{D}(c_k)|} \sum_{x \in \mathcal{D}(c_k)} h_x^v, \quad (4)$$

$$h_{cls,c_k}^v = \frac{1}{|\mathcal{D}(c_k)|} \sum_{x \in \mathcal{D}(c_k)} h_{cls,x}^v \quad (5)$$

where $\mathcal{D}(c_k)$ contains all samples for class c_k . The prototypes for augmented views are calculated using the augmented samples for each class c_k denoted as $\mathcal{D}^{aug}(c_k)$ and the corresponding prototypes are denoted as $h_{c_k}^{aug,t}$, $h_{c_k}^{aug,v}$ and $h_{cls,c_k}^{aug,v}$ respectively.

3.3. Prototype Discriminating Loss

The discriminating loss is responsible for training learnable prompts to distinguish the context of samples from one class compared to other classes. This goal is achieved by pushing the class prototype $h_{c_k}^m$, $m \in \{t, v\}$ for both text and visual branches away from the prototype of class $h_{c_j}^m$, $m \in \{t, v\}$ where $c_k \neq c_j$. Likewise, we pull prototypes $h_{c_k}^m$ and $h_{c_k}^{aug,m}$ for same class and push away augmented ones for $c_j \neq c_k$. In this regard, contrastive learning [7, 8, 14, 23] offers a solution to pull prototypes of positive pairs and push away negative pairs. We refer to [34] to propose our discriminating loss \mathcal{L}_D , formally expressed as:

$$\mathcal{L}_{pos}(c_k) = \frac{1}{|\mathcal{H}|} \sum_{aug \in \mathcal{H}} e^{\operatorname{sim}(h_{c_k}^m, h_{c_k}^{aug,m})/\tau} \quad (6)$$

$$\begin{aligned} \mathcal{L}_{neg}(c_k) = & \frac{1}{|\mathcal{H}|} \sum_{aug \in \mathcal{H}} \sum_{c=1, c \neq c_k}^C e^{\operatorname{sim}(h_{c_k}^m, h_c^m)/\tau} \\ & + e^{\operatorname{sim}(h_{c_k}^{aug,m}, h_c^m)/\tau} + e^{\operatorname{sim}(h_{c_k}^m, h_c^{aug,m})/\tau} \end{aligned} \quad (7)$$

$$\mathcal{L}_D = -\frac{1}{|m| * C} \sum_{\forall m} \sum_{c=1}^C \log \frac{\mathcal{L}_{pos}}{\mathcal{L}_{neg}} \quad (8)$$

where $c_k \in [1, C]$ and $m \in \{t, v\}$, i.e. $|m| = 2$. The prototypes $h_{c_k}^m$ and $h_{c_k}^{aug,m}$ additionally contains h_{cls,c_k}^v , $h_{cls,c_k}^{aug,v}$




Figure 1. Architecture of proposed PromptSync method for zero-shot generalization in CLIP. During test time, we updates the learnable prompts using discriminative and alignment of class prototypes. For a single test example, we obtain multiple augmented views and obtain the mean class probabilities after parameter updates with discriminating loss. Mean class probabilities act as weights in the class-prototype alignment with filtered augmented views. Gradient are accumulated over multiple iterations before final update to the learnable prompts.

when $m = v$. Resulting prompt update $\mathbf{p} \rightarrow \hat{\mathbf{p}}$ (learnable prompt tokens) is obtained after applying gradients for the discriminating loss. Since the proxy dataset will remain same for all the test instances, the updated prompt can be saved and restored each time for an incoming test sample. We presented the study on performance and latency with and without saving these updated prompts in Appendix 10. For the rest of the paper we generalize our method without requiring to save these updated prompts.

3.4. Prototype Alignment Loss

Loss \mathcal{L}_D can effectively separate different classes, it is not able to tune the prompt for the test sample, which comes from a different distribution than the source distribution. Hence, we propose the prototype alignment of the test sample (and its augmented views) with the class prototype obtained from the source distribution. We propose to weigh the prototype alignment by the probability of the test sample lying in the particular class. Lets denote the probability $\tilde{p}_{\hat{\mathbf{p}}}[c]$ (as mentioned in Eq.1) as the mean of probabilities (for class c) produced with the updated prompt $\hat{\mathbf{p}}$ across filtered augmented views (preserved after the confidence selection filter) including test sample (F). The amplitude and angle alignment of sample x_i with the class prototypes for both text and visual branches is calculated as follows:

$$\mathcal{L}'_{amp}(x_i) = \sum_{c=1}^C \tilde{p}_{\hat{\mathbf{p}}}[c] \|p_{x_i}^m - p_c^m\|^2 \quad (9)$$

$$\mathcal{L}'_{ang}(x_i) = \sum_{c=1}^C \tilde{p}_{\hat{\mathbf{p}}}[c] \text{sim}(p_{x_i}^m, p_c^m) \quad (10)$$

where $m \in \{t, v\}$. However, there is an issue with MSE loss since it gives an equal penalty (e.g. $\mathcal{L}'_{amp} = 0.1$) for an increase from 1.2 to 1.3 and 1.7 to 1.8. But we wanted to penalise more for 1.2 to 1.3 since increase in MSE in the smaller range should be penalised more to preserve the base class performance. Hence we penalise with logarithm i.e. we use $\mathcal{L}_{amp} = \log \mathcal{L}'_{amp}$. Similarly, the penalty should be applied to $1/\mathcal{L}'_{ang}$ for angle alignment. The updated amplitude and angle alignment loss is $\mathcal{L}_{amp}(x_i) = \log \mathcal{L}'_{amp}$ and $\mathcal{L}_{ang}(x_i) = \log(1/\mathcal{L}'_{ang}) = -\log \mathcal{L}'_{ang}$ respectively. We combined the amplitude and angle loss with equal importance and hence the prototype alignment loss \mathcal{L}_A is given as:

$$\begin{aligned} \mathcal{L}_A &= \frac{1}{|F|} \sum_{x_i \in F} (\mathcal{L}_{amp}(x_i) + \mathcal{L}_{ang}(x_i)) \\ &= -\frac{1}{|F|} \sum_{x_i \in F} \log \frac{\mathcal{L}'_{ang}(x_i)}{\mathcal{L}'_{amp}(x_i)} \end{aligned} \quad (11)$$

3.5. Algorithm Details

In order to compute the prototype discriminating loss on the source dataset, we require the pre-training dataset of the CLIP model. However, it was trained on over 400 million image-text pairs, which are not publicly available. Nevertheless, in previous works[2, 33], CLIP has been heavily tuned on the ImageNet[11] dataset to achieve excellent zero-shot performance. Hence, we use ImageNet as the proxy for the source dataset to compute prototypes for each class. These prototypes are computed offline for both the sample and its augmented views, and they are used directly during test-time adaptation. During each iteration of test-time adaptation, the

meta-train stage is entered first. The model starts training using the prototype discriminating objective $\mathop{\text{argmin}}_p \mathcal{L}_D$, and gradients are calculated, resulting in the prompt update $\mathbf{p} \rightarrow \hat{\mathbf{p}}$ (i^{th} iteration). Subsequently, the meta-test stage is executed. Here, the augmented views are first filtered using a confidence threshold over predicted probabilities using the updated prompts $\hat{\mathbf{p}}$. The mean probabilities $\tilde{p}_{\hat{\mathbf{p}}}$ are computed over F and used as weights in \mathcal{L}_A . The model is trained on F , and the gradient of prototype alignment loss $\nabla_{\mathbf{p}} \mathcal{L}_A$ is calculated. We average out the gradients over all samples in F . Finally, the prompts \mathbf{p} is updated using combined objective: $\mathcal{L}_{ent} + \mathop{\text{argmin}}_p \mathcal{L}_A$. For $n > 1$ we accumulate the averaged gradients before final prompt update.

4. Experiments

We have evaluated PromptSync on different benchmark settings (Appendix 9) with different datasets described below: **Datasets:** For domain generalisation setting, we follow PromptAlign [33] and evaluated our method on four out-of-distribution (OOD) variants of ImageNet [11]: ImageNetV2 [32], ImageNet-Sketch [40], ImageNet-A [18] and ImageNet-R [17]. We also consider the evaluation on a recent and challenging benchmark, namely, Photorealistic Unreal Graphics (PUG) dataset [4], comprised of different textures, sizes, orientations, and backgrounds. For cross-dataset transfer setting, we follow TPT [35] and evaluate the performance on 10 diverse image classification datasets with varying complexities for visual recognition tasks. This includes Caltech 101 [12] for generic objects. Five fine-grained datasets (spanning images of animals, flowers and transportation) are StanfordCars [24], Food101 [5], Flowers102 [27], FGVC-Aircraft [25], OxfordPets [29]. Moreover, four datasets, namely, SUN397 [37], DTD [10], UCF101 [36], and EUROSAT [15], comprise scenes, textures, human actions, and satellite imagery, respectively. For base-to-novel generalisation, we follow [21] and evaluate our method on ImageNet and the 10 image classification datasets.

Baselines: We compared PromptSync with existing few-shot prompt learning methods for CLIP adaptation; these are CoOp [47], CoCoOp [46], TPT [35], and PromptAlign [33]. MaPLe [21] is a multi-modal prompt learning baseline that adapts CLIP by learning prompts on both text and visual branches. TPT [35] and PromptAlign [33] are the test-time prompt tuning methods that tune the prompt for each incoming test sample, achieving state-of-the-art performance in prompt learning.

Implementation Details: We ran all experiments on a single NVIDIA A100 40GB GPU. Following [21], we trained on ImageNet with 16-shot training data selected at random for each class using 2 prompt tokens for a depth of 3 layers (on CLIP ViT-B/16 backbone architecture). We optimized the prompts on both the text and visual branches using a single

test image. We augmented each test image with 127 different views using random resized crops, background substitution, horizontal flip augmentations, and visual corruption. For text augmentation, we used hyponyms, synonyms, and meronyms from WordNet[26]. Moreover, we generated various text prompts from pre-trained LLMs [6]. Additionally, we randomly masked one of the learnable tokens for 15% of augmented views. We computed the gradients of alignment loss for a batch size of 128 images, including the original image. During the meta-train stage, we updated the original parameters (using a single iteration) and then optimized the prompts in the meta-test stage by calculating the gradients of alignment loss w.r.t. the updated parameters accumulated for a single ($n = 1$) iteration to facilitate the one-to-one comparison with baselines. We obtained the top 10% confident predictions of augmented views based on the lowest entropy. We used the AdamW optimizer and a learning rate β of $5e^{-4}$ for the fine-grained datasets and 0.04 for the rest of the datasets.

4.1. Domain Generalization

We demonstrate that all test-time adaptation methods exhibit better performance (Table 1) compared to the pre-trained CLIP model, highlighting the advantage of tuning V-L models at test time. PromptSync achieves the highest Top-1 accuracy averaged across all the domains of ImageNet variants. Furthermore, we evaluated the ImageNet-trained model on various out-of-distribution (OOD) datasets and observed consistent improvement in performance compared to existing state-of-the-art (SOTA) approaches. The detailed results for each domain dataset are presented in Tables 1 and 2. This confirms that alignment and discriminative training with augmented views on both the text and visual branches enhance the generalization performance of V-L models like CLIP.

4.2. Base to Novel Generalization

Table 3 presents the detailed performance report of PromptSync on base and novel classes across 11 recognition datasets. On average, our strategy outperforms the model performance by 1.29% on base classes and nearly 1% on novel classes. We observe that PromptAlign, based on a distribution alignment strategy, outperforms for novel classes in most cases, with an average improvement of 0.79% compared to the best-performing model. However, the margin of improvement is very low. In contrast, with TPT, the performance drops in some instances, such as for OxfordPets, Eurosat, and UCF101. This demonstrates that: 1) test-source alignment is crucial for prompt tuning. 2) Prompt tuning alone in the text branch is not sufficient for zero-shot generalization. Since distribution alignment does not promote discriminative learning and the entropy loss on the test dataset is noisy, PromptSync outperforms with class-aware prototype discrimination and alignment across different augmented

	Imagenet V2	Imagenet Sketch	Imagenet A	Imagenet R	OOD Avg
CLIP	60.86	46.09	47.87	73.98	57.20
CLIP+TPT	64.35	47.94	54.77	77.06	60.81
CoOp	64.20	47.99	49.71	75.21	59.28
CoOp+TPT	66.83	49.29	57.95	77.27	62.84
Co-CoOp	64.07	48.75	50.63	76.18	59.91
Co-CoOp+TPT	64.85	48.27	58.47	78.65	62.61
MaPLe	64.07	49.15	50.90	76.98	60.28
MaPLe+TPT	64.87	48.16	58.08	78.12	62.31
PromptAlign	65.29	50.23	59.37	79.33	63.55
PromptSync	67.54	53.42	61.92	80.64	65.88

Table 1. **Comparison on the domain generalization setting.** Prompt tuning methods are trained on ImageNet and evaluated on datasets with domain shifts

	Camera (Yaw/ Pitch/ Roll)	Pose (Yaw/ Pitch/ Roll)	Scale	Texture	Lighting	Worlds
MaPLe	48.73/ 39.93/ 32.13	48.10/ 28.40/ 27.80	46.90	37.90	15.50	32.13
MaPLe+TPT	57.04/ 45.99/ 39.23	56.26/ 35.64/ 33.26	54.87	43.73	22.52	42.00
PromptAlign	58.14/ 46.93/ 40.45	57.43/ 36.31/ 34.32	56.18	44.97	23.06	43.24
PromptSync	59.84/ 48.54/ 41.92	59.72/ 38.84/ 36.64	58.12	45.98	25.02	44.84

Table 2. **Comparison on the domain generalization setting for distribution alignment.** MaPLe is trained on ImageNet and evaluated on OOD dataset i.e., PUG

views. Averaging the gradients further motivates domain-agnostic prompt tuning on both the text and visual branches. This enhances the zero-shot generalization of the V-L model compared to other state-of-the-art approaches. Moreover, our strategy for prompt tuning does not lose information for base classes.

4.3. Cross-Dataset Transfer

In Table 4, we compared the transfer performance of PromptSync with existing state-of-the-art methods using prompt learning. We evaluated methods for transfer performance across diverse cross-datasets. PromptSync consistently outperforms the previous best method, i.e., PromptAlign [33], across all cross-datasets, providing an average improvement of 2.84%. Compared to PromptAlign, which outperforms the previous method MaPLe + TPT by a very small margin, i.e., 0.42%, our method shows a significant average improvement of 3.26% over MaPLe + TPT. Other methods, CoOp and CoCoOp, on average, perform worse than zero-shot CLIP + TPT (except ProDA [45]). This affirms that both text-visual alignment and domain-agnostic parameter updates result in better transfer generalization across cross-datasets in V-L models. As opposed to our method, the previous approaches were not consistent in performance across all datasets, which further affirms the advantage of a domain-agnostic training strategy.

5. Ablation

Class-Aware Prototype Alignment: Table 5 summarizes the comparison between two alignment strategies: distribution alignment of the test sample with the class-agnostic source distribution. All results are on the ImageNet-A dataset. PromptAlign adopted distribution alignment along with averaged cross-entropy for prompt tuning. However, we perform domain-agnostic parameter updates with class-aware prototype alignment for the test sample. As shown in Table 5, PromptAlign[†] without entropy loss is as good as vanilla MaPLe. This is due to the fact that distributional alignment does not promote any discriminative learning in the absence of entropy loss. However, because entropy loss is noisy due to the poor performance of the vanilla zero-shot V-L model, we propose the stronger discriminative loss of class prototype alignment for prompt tuning with source and test samples with augmented views. PromptSync[†] without entropy loss outperforms the corresponding counterpart PromptAlign[†]. This is because the class-aware prototype alignment has both alignment and discriminative properties, thus improving test-time adaptation on its own. With additional signals from predicted probabilities for each class, the class-aware prototype alignment acts as a geometric regularizer, mitigating class collapse in prompt representation.

Loss variants: We conducted an ablation study on amplitude and angle loss for the class-aware prototype alignment

Datasets	Sets	CoOp (IJCV22)	CoCoOp (CVPR22)	ProDA (CVPR22)	MaPLe (CVPR23)	MaPLe + TPT (CVPR23)	PromptAlign (NIPS23)	PromptSync
Average	Base	82.38	80.47	81.56	82.24	82.16	83.19	84.17
	Novel	67.96	71.69	72.30	75.09	74.95	75.88	77.17
ImageNet	Base	76.46	75.98	75.40	76.67	77.73	78.26	79.23
	Novel	66.31	70.43	70.23	70.54	72.24	72.59	73.84
Caltech101	Base	97.80	97.96	98.27	98.00	98.54	98.60	98.62
	Novel	93.27	93.81	93.23	94.27	94.29	94.50	94.67
OxfordPets	Base	94.47	95.20	95.43	95.43	95.23	95.38	95.44
	Novel	96.00	97.69	97.83	97.80	97.37	97.56	97.83
Stanford Cars	Base	75.67	70.49	74.70	72.90	74.00	75.02	76.42
	Novel	67.53	73.59	71.20	73.97	75.20	75.71	77.21
Flowers102	Base	97.27	94.87	97.70	95.93	96.24	96.61	97.73
	Novel	67.13	71.75	68.68	72.40	72.10	72.34	73.78
Food101	Base	89.37	90.70	90.30	90.70	91.13	91.63	92.39
	Novel	88.77	91.29	88.57	92.07	92.03	92.68	92.95
FGVC Aircraft	Base	39.67	33.41	36.90	37.27	34.31	37.21	40.91
	Novel	31.23	23.71	34.13	35.53	35.81	37.27	39.31
SUN397	Base	80.85	79.74	78.67	80.80	81.15	81.57	84.28
	Novel	68.34	76.86	76.93	78.70	79.18	79.48	83.01
DTD	Base	79.97	77.01	80.67	80.30	82.20	82.60	83.49
	Novel	48.60	56.00	56.48	59.23	59.91	60.55	62.03
EuroSAT	Base	90.10	87.49	83.90	93.63	91.02	94.10	94.63
	Novel	53.00	60.04	66.00	72.87	68.96	72.71	73.19
UCF101	Base	84.53	82.33	85.23	82.97	82.23	84.11	85.75
	Novel	67.37	73.45	71.97	78.57	77.34	79.30	81.29

Table 3. **Comparison on Base-to-novel generalization setting.** PromptSync shows consistent improvement on both base and novel classes over previous methods

	Caltech	Pets	Cars	Flowers	Food101	Aircraft	SUN397	DTD	EuroSAT	UCF101	Average
CLIP	93.35	88.25	65.48	67.44	83.65	23.67	62.59	44.27	42.01	65.13	63.58
CLIP+TPT	94.16	87.79	66.87	68.98	84.67	24.78	65.50	47.75	42.44	68.04	65.10
CoOp	93.70	89.14	64.51	68.71	85.30	18.47	64.15	41.92	46.39	66.55	63.88
CoOp + TPT	93.15	89.48	66.77	68.48	86.48	20.51	66.06	43.32	37.73	68.91	64.08
CoCoOp	93.79	90.46	64.90	70.85	83.97	22.29	66.89	45.45	39.23	68.44	64.63
CoCoOp + TPT	88.57	85.33	59.68	55.31	80.64	16.89	60.24	38.93	48.55	63.35	59.75
ProDA	86.70	88.20	60.10	77.50	80.80	22.20	-	50.90	58.50	-	65.62
MaPLe	93.53	90.49	65.57	72.23	86.20	24.74	67.01	46.49	48.06	68.69	66.30
MaPLe+TPT	93.59	90.72	66.50	72.37	86.64	24.70	67.54	45.87	47.80	69.19	66.50
PromptAlign	94.01	90.76	68.50	72.39	86.65	24.80	67.54	47.24	47.86	69.47	66.92
PromptSync	95.78	91.89	69.24	77.68	87.72	25.91	67.98	50.99	59.36	71.04	69.76

Table 4. **Comparison on cross-dataset transfer setting.** Prompt tuning methods are trained on ImageNet and evaluated on cross-datasets

objective. Table 6 compares three loss choices: 1) amplitude loss, 2) angle loss, and 3) amplitude + angle loss. Clearly, the combination of amplitude and angle performs better than

other choices. The formulation for the combination of amplitude and angle loss is the same as in equation 11. We further investigated other variants, i.e., combining two of them with

Method	Entropy Loss	Alignment Loss	Discriminative Loss	Top-1 Acc.
MaPLe	✗	✗	✗	50.90
MaPLe+TPT	✓	✗	✗	58.08
PromptAlign [†]	✗	✓	✗	50.85
PromptAlign	✓	✓	✗	59.37
PromptSync [†]	✗	✗	✓	56.67
PromptSync	✓	✓	✓	61.92

Table 5. **Ablation Study.** Analysis of Alignment, Discriminative and Entropy minimization loss. The average of Top-1 accuracy(%) across three seeds is reported

	\mathcal{L}_{amp}	\mathcal{L}_{ang}	sub	sum_exp	\mathcal{L}_A
PromptSync	59.84	58.81	57.86	59.83	61.92

Table 6. **Ablation analysis to alignment loss variants.** All results are on ImageNet-A dataset

out taking the log: 1) subtraction between amplitude and angle $\mathcal{L}'_{amp} - \mathcal{L}'_{ang}(sub)$ 2) the summation of exponential of both losses $\exp(\mathcal{L}'_{amp}) + \exp(1/\mathcal{L}'_{ang})(sum_exp)$. Clearly, the formulation in equation 11 (\mathcal{L}_A) performs best among other variants. Ablation on the proxy dataset is given in Appendix 12, and ablation on performance and latency with and without saving updated prompts is provided in Appendix 10. We also compared the number of augmented views and prompt updates in Appendix 11.

6. Performance and Latency

The experiments presented in the Table 7 (Appendix) involve a comparison of different methods, namely MaPLe + TPT, PromptAlign, PromptSync*, and PromptSync. In these experiments, we evaluated the top-1 average accuracy (%) and latency (in hours for a single prompt update) of each method. Specifically, we investigated PromptSync with and without saving the updated prompt obtained after prototype discrimination, with the variant denoted as PromptSync* indicating the adaptation of prompt tokens for test samples after restoring saved prompt tokens.

The results, as shown in Table 7, include latency measurements represented in hours for a single prompt update, and all evaluations are conducted on the ImageNet-A dataset. Notably, the PromptSync* variant demonstrates a faster processing time compared to the full PromptSync method, with only a marginal drop in performance. This outcome underscores the achieved generalization through prototype alignment. Furthermore, in comparison to previous methods such as MaPLe + TPT and PromptAlign, the PromptSync* variant exhibits only a slight increase in latency (0.03 hours) while still improving overall performance.

7. Sensitivity Comparison

We further performed the sensitivity comparison of our method as compared to other state-of-the-art baselines. In Appendix, Figure 2(a) shows the comparison of performance during test time adaptation as the number of views increases. All the results are on ImageNet-A dataset. In comparison to PromptAlign and MaPLe + TPT, their performance almost plateaus around 64 views with insignificant improvement further, while PromptSync shows a consistent improvement with the increase in views and insignificant improvement beyond 128. This proves the generalizability achieved by our method since it optimises base CLIP over a larger number of possible shifts in the dataset, resulting in better performance. Figure 2(b) shows the performance comparison as the number of prompt update steps increases. All the methods increase their performance with an increase in the number of steps; however, our method shows better adaptation to the test sample with more steps in comparison to PromptAlign and MaPLe + TPT. For apples-to-apples comparison we perform a single-step update (128 views) following TPT [35].

8. LAION400M Proxy Dataset Analysis

Given CLIP’s impressive zero-shot performance on ImageNet, we opted for ImageNet as a viable proxy source dataset, aligning with prior research [33]. We worked with a subset of LAION400M, comprising 2.5 million images (2 times the size of ImageNet). Furthermore, we carried out an ablation study on the alignment strategy using LAION400M as the source dataset, a dataset known to mirror CLIP’s training dataset [9]. The results for this ablation study is shown in Table 8 (Appendix). Notably, the performance impact remains consistent when utilizing this subset of LAION400M alongside ImageNet. Source class prototypes are computed on the proxy source data to derive the distribution for alignment during test time. As this proxy dataset aligns with the model’s training set, this offline computation remains unchanged despite environmental shifts and only necessitates computation once.

Conclusion

In summary, PromptSync significantly improves zero-shot generalization in vision-language models. Our approach, addressing class dominance and variance, outperforms existing methods by 2.33% overall, with a 1% boost in base-to-novel generalization and 2.84% in cross-dataset transfer on a domain generalization benchmark. This underscores PromptSync’s effectiveness in enhancing the robustness of vision-language models.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. [1](#)
- [2] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring visual prompts for adapting large-scale models. *arXiv preprint arXiv:2203.17274*, 2022. [4](#)
- [3] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Visual prompting: Modifying pixel space to adapt pre-trained models. *arXiv preprint arXiv:2203.17274*, 3:11–12, 2022. [2](#)
- [4] Florian Bordes, Shashank Shekhar, Mark Ibrahim, Diane Bouchacourt, Pascal Vincent, and Ari S Morcos. Pug: Photorealistic and semantically controllable synthetic data for representation learning. *arXiv preprint arXiv:2308.03977*, 2023. [5](#)
- [5] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 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*, pages 446–461. Springer, 2014. [5](#)
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [5](#)
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [3](#)
- [8] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [3](#)
- [9] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2818–2829, 2023. [8](#), [1](#)
- [10] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014. [5](#)
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [4](#), [5](#), [1](#)
- [12] Li Fei-Fei, Rob Fergus, and Pietro Perona. 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*, pages 178–178. IEEE, 2004. [5](#)
- [13] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*, 2021. [1](#)
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [3](#)
- [15] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. [5](#)
- [16] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019. [1](#)
- [17] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. [5](#)
- [18] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021. [5](#)
- [19] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. [1](#)
- [20] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. [1](#), [2](#)
- [21] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multimodal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19113–19122, 2023. [1](#), [2](#), [5](#)
- [22] Muhammad Uzair Khattak, Syed Talal Wasim, Muzammal Naseer, Salman Khan, Ming-Hsuan Yang, and Fahad Shahbaz Khan. Self-regulating prompts: Foundational model adaptation without forgetting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15190–15200, 2023. [1](#)
- [23] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. [3](#)
- [24] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013. [5](#)
- [25] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual clas-

- sification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. [5](#)
- [26] George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995. [5](#)
- [27] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pages 722–729. IEEE, 2008. [5](#)
- [28] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1717–1724, 2014. [1](#)
- [29] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012. [5](#)
- [30] Joaquin Quinonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. Mit Press, 2008. [1](#)
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#), [2](#)
- [32] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*, pages 5389–5400. PMLR, 2019. [1](#), [5](#)
- [33] Jameel Hassan Abdul Samadh, Hanan Gani, Noor Hazim Hussein, Muhammad Uzair Khattak, Muzammal Naseer, Fahad Khan, and Salman Khan. Align your prompts: Test-time prompting with distribution alignment for zero-shot generalization. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [34] Zeyang Sha, Xinlei He, Ning Yu, Michael Backes, and Yang Zhang. Can’t steal? cont-steal! contrastive stealing attacks against image encoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16373–16383, 2023. [3](#)
- [35] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. *Advances in Neural Information Processing Systems*, 35:14274–14289, 2022. [1](#), [2](#), [3](#), [5](#), [8](#)
- [36] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. A dataset of 101 human action classes from videos in the wild. *Center for Research in Computer Vision*, 2(11), 2012. [5](#)
- [37] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pages 9229–9248. PMLR, 2020. [5](#)
- [38] Yue Tan, Guodong Long, Lu Liu, Tianyi Zhou, Qinghua Lu, Jing Jiang, and Chengqi Zhang. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8432–8440, 2022. [3](#)
- [39] Shengbang Tong, Yubei Chen, Yi Ma, and Yann Lecun. Empssl: Towards self-supervised learning in one training epoch. *arXiv preprint arXiv:2304.03977*, 2023. [3](#)
- [40] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32, 2019. [5](#)
- [41] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021. [1](#)
- [42] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. [1](#)
- [43] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. *arXiv preprint arXiv:2210.07225*, 2022. [1](#), [2](#)
- [44] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18123–18133, 2022. [1](#)
- [45] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12414–12424, 2021. [6](#)
- [46] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. [1](#), [2](#), [5](#)
- [47] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [1](#), [2](#), [5](#)