This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# **Bayesian Test-Time Adaptation for Vision-Language Models**

Lihua Zhou<sup>1</sup>, Mao Ye<sup>2</sup>, Shuaifeng Li<sup>2</sup>, Nianxin Li<sup>2</sup>, Xiatian Zhu<sup>3</sup>, Lei Deng<sup>4</sup>, Hongbin Liu<sup>1,5</sup>, Zhen Lei<sup>1,5,6,7\*</sup>

<sup>1</sup> CAIR, HKSIS, CAS, <sup>2</sup> UESTC, <sup>3</sup> University of Surrey, <sup>4</sup> Shenzhen University, <sup>5</sup> MAIS, Institute of Automation, CAS, <sup>6</sup> SAI, UCAS, <sup>7</sup> M.U.S.T

#### Abstract

Test-time adaptation with pre-trained vision-language models, such as CLIP, aims to adapt the model to new, potentially out-of-distribution test data. Existing methods calculate the similarity between visual embedding and learnable class embeddings, which are initialized by text embeddings, for zero-shot image classification. In this work, we first analyze this process based on Bayes theorem, and observe that the core factors influencing the final prediction are the likelihood and the prior. However, existing methods essentially focus on adapting class embeddings to adapt likelihood, but they often ignore the importance of prior. To address this gap, we propose a novel approach, **B**ayesian *Class Adaptation (BCA), which in addition to continuously* updating class embeddings to adapt likelihood, also uses the posterior of incoming samples to continuously update the prior for each class embedding. This dual updating mechanism allows the model to better adapt to distribution shifts and achieve higher prediction accuracy. Our method not only surpasses existing approaches in terms of performance metrics but also maintains superior inference rates and memory usage, making it highly efficient and practical for real-world applications.

## 1. Introduction

Recent advances in visual tasks have seen a growing preference for using pre-trained vision-language models like CLIP [18, 22, 33, 46], which leverage large-scale paired image-text data enables it to learn rich, multi-modal representations that generalize well to a variety of downstream tasks [24]. However, when deploying CLIP in real-world environments, the distribution discrepancy between the pretrained data and the environment data can lead to performance degradation [23, 25, 45, 56, 57]. To address this issue in real-time, test-time adaptation (TTA), which aims adapt the pre-trained model to new, potentially out-of-distribution



Figure 1. Fixed Prior *vs.* Adaptive Prior: Comparison of Diagnosis Outcomes. In the fixed prior scenario, patients with fever are consistently diagnosed with the common cold, regardless of whether it is a normal period or a COVID-19 period. In contrast, the adaptive prior scenario adjusts the diagnosis based on the current context. During normal periods, patients with fever are diagnosed with the common cold, while during the COVID-19 period, they are more likely to be diagnosed with COVID-19. This demonstrates the importance of performing prior adaptation in different environments.

test data during the inference phase, have been applied to the CLIP [9, 11, 19, 38].

Existing TTA methods based on CLIP can be roughly divided into two categories. Early methods [9, 38] fine-tune the text prompts by minimizing the entropy of prediction to generate new class embeddings, which requires backpropagation and can be time-consuming, and then calculate the similarity between visual embedding and these new class embeddings for zero-shot classification. In fact, this approach is not in line with the TTA problem setting, which emphasizes real-time adaptation. Therefore, recent methods [11, 19] have begun to use memory to continuously store visual embeddings to increase the number of class embeddings or employ statistical methods to update class embeddings with incoming visual embedding, and then per-

<sup>\*</sup>Corresponding author

form classification with these updated class embeddings. These approaches can perform real-time test-time adaptation, aligning more closely with the TTA problem setting.

Essentially, current test-time adaptation methods based on CLIP generate initial class embeddings  $\boldsymbol{U} = \{\boldsymbol{\mu}_m\}_{m=1}^M$ by encoding text prompts and dynamically update these class embeddings U to update the logits (or probability)  $P(\boldsymbol{x}|\boldsymbol{U})$ , which is also known as the likelihood, to better adapt to the test environment distribution during the test phase, thereby making a better prediction P(Y|x), which is also called the posterior. In this work, we leverage the Bayes theorem to analyze and observe that the posterior  $P(Y|\mathbf{x})$  is influenced not only by the likelihood  $P(\mathbf{x}|\mathbf{U})$ but also by the prior P(Y|U), which reflects initial belief about the categories Y given the class embeddings U. However, existing methods [9, 11, 19, 38] all ignore this point and use the fixed prior from pre-trained CLIP, which inevitably leads to suboptimal solutions when the test data distribution significantly deviates from the pre-trained data distribution. We also give a related example in Figure 1. Please note that we focus on the conditional prior P(Y|U)rather than P(Y). Unless otherwise specified, the prior we refer to in this work is P(Y|U).

In this work, we propose Bayesian Class Adaptation (BCA), which not only considers only consider the likelihood adaptation as in previous methods, but also incorporate the prior adaptation. When a test image arrives, its visual embedding is computed using the visual encoder of CLIP, and the probability of belonging to each class embedding is calculated based on the current likelihood. For likelihood adaptation, the most similar class embedding is selected and updated using a statistical method by recalculating the mean of the selected class embedding with the current visual embedding. For prior adaptation, the probability of belonging to each class embedding are combined with the current prior to compute the posterior for each category. The category with the highest posterior probability is selected as the most likely category for the image. The prior of the selected class embedding is updated by recalculating the mean with the posterior. This dynamic update mechanism helps the model adapt to the test distribution, thereby enhancing the robustness and accuracy of zero-shot image classification in real-world applications. Despite its simplicity, the method ensures high inference rates and low memory usage, making it highly suitable for TTA scenarios. Our contributions are as follows:

 Comprehensive Analysis: We leverage the Bayes theorem to analyze the process of CLIP performing zero-shot image classification, observing that the key of making better prediction lies in considering both the likelihood adaptation and the prior adaptation. We find that existing methods often ignore the prior adaptation. By addressing this issue, we propose a more general method. 2. Effective and Fast Method: We propose a efficient method that not only achieves good performance on Out-of-Distribution (OOD) and Cross Domain benchmarks, but also maintains high inference speed and low memory usage, making it highly suitable for TTA scenarios.

## 2. Related Work

**Vision-language models.** Most visual recognition studies rely on training a separate deep neural network for each task, which is laborious [12]. Recently, Vision-Language Models (VLMs) have emerged, learning rich vision-language correlations from web-scale image-text pairs and enabling zero-shot predictions on various visual recognition tasks with a single model [3, 18, 33, 47]. CLIP [33] is the pioneering work that uses contrastive learning to align image and text embeddings. ALIGN [18] scales up the training of vision-language models by leveraging a dataset of over one billion noisy image alt-text pairs. FILIP [47] models the fine-grained semantic alignment through a novel cross-modal late interaction mechanism in the contrastive loss. PaLI [3] performs joint scaling on both the language and vision components for a wide range of parameters.

**Transfer learning with vision-language models.** Since pre-trained VLMs are generally trained to learn task-agnostic concepts, their performance can degrade when applied to downstream tasks due to the task-specific styles of images and text. To address this issue, transfer learning has been introduced to further adapt them to downstream tasks. From a methodological perspective, these methods can be roughly divided into two categories: prompt tuning [1, 35, 37, 50, 54, 55] and feature adapter [10, 31, 52].

For prompt tuning, it fits the downstream tasks by finding the optimal prompts. CoOp [55] propose Context Optimization to learn a specific set of context tokens for each class, while CoCoOp [54] further propose conditional Context Optimization to generate a specific prompt for each image. Unlike these two methods perform text prompt tuning, VP [1] uses a learnable perturbations on images to perform visual prompt tuning. RePrompt [35] proposes retrievalenhanced visual prompt learning to cache the knowledge of downstream tasks. Further, UPT [50] tunes both the visual and text prompt under the few-shot learning and domain generalization settings. MVLPT [37] incorporates crosstask knowledge to enable information sharing for visual and text prompt tuning. For feature adapter, it fine-tunes VLMs to align both text embedding and visual embedding with a light-weight adapter. CLIP-Adapter [10] adopts an additional bottleneck layer to learn new features and performs residual-style feature blending with the original pretrained features. Tip-Adapter [52] creates the weights by a keyvalue cache model constructed from the few-shot training set. SVL-Adapter [31] introduces a new encoder which is trained using self-supervision on the target dataset to combine the strengths of both vision-language pretraining and self-supervised learning.

**Test-time adaptation with vision-language models.** Testtime adaptation aims to adapt pre-trained model to test data, which may have the distribution discrepancy with the training data, during the test time. Early TTA methods often applied entropy minimization [27, 53], batch normalization calibration [29, 41], pseudo-labeling [16, 17] and consistency regularization [30, 43] to improve the performance.

Recently, with the increasing popularity of VLMs, many TTA works have also been applied to CLIP. TPT [38] optimizes the text prompt to encourage consistent predictions across augmented views by minimizing the marginal entropy. DiffTPT [9] leverages pre-trained diffusion models to generate diverse and informative augmented data to improve TPT. C-TPT [48] jointly optimizes the prompt to achieve better calibration by maximizing average text feature dispersion during test time. Unlike these methods that require backpropagation during the test time, TDA [19] adapts to test data gradually via progressive pseudo label refinement with a lightweight key-value cache. DOTA [11] continually estimates the distribution of test samples to adapt the model work well in test environment. MTA [49] manages augmented views by optimizing inlierness variables to generate final visual embeddings in a trainingfree manner. PromptAlign [36] explicitly aligns the test sample statistics with that of the source data distribution through token distribution alignment to optimizes the text prompt. HisTPT [51] introduces three types of knowledge banks to build up comprehensive memorization, mitigating the knowledge forgetting. Zero [7] sets the Softmax temperature to zero to converting probability distributions into one-hot encodings for test-time adaptation.

### 3. Method

**Problem Setting.** In this work, we address the challenge of test-time adaptation for a pre-trained visual-language model, CLIP [33]. Given a sequence of images  $\{x_i\}_{i=1}^n$  that arrive sequentially, the model needs to predict the label for each single image  $x_i$  immediately in an online manner. The goal is to improve the classification accuracy of the model on test data  $\{x_i\}_{i=1}^n$  while preserving inference efficiency.

For a K-class classification problem at test time, the CLIP model typically prepares K hand-crafted prompts, such as  $\{a \text{ photo of}[Class k]\}_{k=1}^{K}$ , and sends them to the text encoder  $E_t$  to obtain their normalized text embeddings  $F^t = \{f_k^t\}_{k=1}^{K}$ . The visual encoder  $E_v$  is then used to encode the test images  $x_i$  to obtain their normalized visual embeddings  $f_i^v$ . Finally, the model computes the cosine similarity between  $f_i^v$  and  $F^t$  to determine the prediction as follows:

$$P(Y|\boldsymbol{x}_i) = \frac{\exp(\cos(\boldsymbol{f}_i^v, \boldsymbol{F}^t))}{\sum_{j=1}^{K} \exp(\cos(\boldsymbol{f}_i^v, \boldsymbol{f}_j^t))},$$
(1)

where  $\exp(\cdot)$  denotes the exponential function and  $\cos(\cdot, \cdot)$  represents the cosine similarity function.

#### 3.1. Analysis

In this section, we use Bayes theorem to analyze the process of how the model generates predictions when a sample  $x_i$  arrives. We consider a more generalized scenario where we have M class embeddings  $\{\mu_m\}_{m=1}^M$  to perform K category zero-shot classification on the sample  $x_i$ . Here, Mcan be greater than or equal to K. Specifically,  $M \ge K$ because some methods [19] try to store visual embeddings during the test phase to increase class embeddings. In contrast, many other methods [9, 38] maintain M = K, updating each class embedding directly. The posterior can be formulated as follows:

$$P(Y|\boldsymbol{x}_{i}) = \sum_{m=1}^{M} P(Y, \boldsymbol{\mu}_{m} | \boldsymbol{x}_{i})$$
$$= \sum_{m=1}^{M} P(\boldsymbol{\mu}_{m} | \boldsymbol{x}_{i}) * P(Y|\boldsymbol{x}_{i}, \boldsymbol{\mu}_{m}) \qquad (2)$$
$$= \sum_{m=1}^{M} P(\boldsymbol{\mu}_{m} | \boldsymbol{x}_{i}) * P(Y|\boldsymbol{\mu}_{m}).$$

The first line is based on the law of total probability. The second line uses the definition of conditional probability. The key assumption in the third line is that once we know the class embeddings  $\mu_m$ , the information of  $x_i$  does not further change the probability of Y [21]. From the Eq. (2), we can divide the steps of predicting the label of  $x_i$  into two parts. First, we calculate the probability of  $x_i$  belonging to each class embedding  $\mu_m$  and then multiply it with the prior  $P(Y|\mu_m)$ . Based on Bayes theorem,  $P(\mu_m|x_i)$  can be further written as:

$$P(\boldsymbol{\mu}_{m}|\boldsymbol{x}_{i}) = \frac{P(\boldsymbol{x}_{i}|\boldsymbol{\mu}_{m}) * P(\boldsymbol{\mu}_{m})}{P(\boldsymbol{x}_{i})}$$
$$= \frac{P(\boldsymbol{x}_{i}|\boldsymbol{\mu}_{m}) * P(\boldsymbol{\mu}_{m})}{\sum_{j=1}^{M} P(\boldsymbol{x}_{i}, \boldsymbol{\mu}_{j})}$$
$$= \frac{P(\boldsymbol{x}_{i}|\boldsymbol{\mu}_{m}) * P(\boldsymbol{\mu}_{m})}{\sum_{j=1}^{M} P(\boldsymbol{x}_{i}|\boldsymbol{\mu}_{j}) * P(\boldsymbol{\mu}_{j})}$$
(3)

where  $P(\boldsymbol{x}_i | \boldsymbol{\mu}_m)$  represents likelihood, and  $P(\boldsymbol{\mu}_m)$  is the probability of each class embedding. Generally, it is assumed that the probability of all class embeddings  $P(\boldsymbol{\mu}_m)$  is the same, that is, assuming  $P(\boldsymbol{\mu}_m) = \frac{1}{M}$ . By substituting Eq. (3) and  $P(\boldsymbol{\mu}_m) = \frac{1}{M}$  back into Eq. (2), we can obtain:

$$P(Y|\boldsymbol{x}_i) = \sum_{m=1}^{M} \frac{P(\boldsymbol{x}_i|\boldsymbol{\mu}_m)}{\sum_{j=1}^{M} P(\boldsymbol{x}_i|\boldsymbol{\mu}_j)} * P(Y|\boldsymbol{\mu}_m). \quad (4)$$

**Remark: 1): Example:** For the original CLIP model, the number of class embeddings M = K, the class embeddings are set as  $\mu_k = f_k^t$ , and the likelihood is defined as



Figure 2. Overview of the proposed Bayesian Class Adaptation (BCA) method. When deploying CLIP to a test environment, M class embeddings are initialized based on hand-crafted prompts, and the prior for each class embedding is initialized as a one-hot vector with the corresponding class set to 1. (a) Embedding: when *i*-th image arrives, it is encoded into a visual embedding  $f_i^v$  using visual encoder. (b) Likelihood adaptation: the probability  $P(\boldsymbol{U}|\boldsymbol{x}_i)$  is calculated based on current likelihood to find class embedding  $\boldsymbol{\mu}_s$  with the highest probability. This  $\boldsymbol{\mu}_s$  is then updated using statistical method with  $f_i^v$  to adapt the likelihood. (c) Prior adaptation: the posterior  $P(Y|\boldsymbol{x}_i)$  is calculated by multiplying  $P(\boldsymbol{U}|\boldsymbol{x}_i)$  by the current prior. And the prior of *s*-th class embedding  $P(Y|\boldsymbol{\mu}_s)$  is adapted with this posterior.

 $P(\boldsymbol{x}_i | \boldsymbol{\mu}_k) \propto \exp(\cos(E_v(\boldsymbol{x}_i), \boldsymbol{\mu}_k)) = \exp(\cos(f_i^v, \boldsymbol{\mu}_k))).$ It is important to note that the prior of pre-trained CLIP  $P(Y | \boldsymbol{\mu}_k) \in R^K$  is a one-hot vector, where the k-th element is 1 and all other are 0. This means that when a sample belongs to the k-th class embedding, its probability of being in k-th category is 1. By substituting these into Eq. (4), we can obtain Eq. (1). **2): Core factors:** From the above derivation, it is evident that the core factors affecting posterior are the likelihood  $P(\boldsymbol{x}|\boldsymbol{\mu})$  and prior  $P(Y|\boldsymbol{\mu})$ . Existing methods focus on adapting the likelihood while keeping the prior fixed. In the following sections, we aim to adapt both the likelihood and the prior to improve the model's performance in new environments.

#### 3.2. Likelihood Adaptation

Here, we consider a setting where  $M \ge K$ , which is more general. From Eq. (4), it is clear that the accuracy of the final posterior is directly influenced by the likelihood  $P(\mathbf{x}_i | \boldsymbol{\mu}_m)$ . Recently, various methods have focused on updating this term to make class embedding more suitable for test distribution. For example, TPT [38], DiffTPT [9] and C-TPT [48] apply prompt tuning to update text embeddings to update  $\boldsymbol{\mu}_m$ , TDA [19] applies feature adapter to add the visual embedding as the new class embeddings, DOTA [11] uses Gaussian distribution  $P(\mathbf{x}_i | \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m) = N(\mathbf{x}_i | \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m)$ , and continuously update  $\mu_m$  and  $\Sigma_m$  based on statistical method. In this work, we follows original CLIP that assume the likelihood  $P(\boldsymbol{x}_i | \boldsymbol{\mu}_m) \propto \exp(\cos(E_v(\boldsymbol{x}_i), \boldsymbol{\mu}_m))$ , which only requires updating  $\boldsymbol{\mu}_m$  using statistical method.

Specifically, the text encoder is used to encode M handcrafted prompts to obtain M text embeddings as the initial class embeddings  $\{\mu_m\}_{m=1}^M$ , where the m-th text embedding is obtained by encoding the hand-crafted prompt of the m%K-th category. And as the test image  $x_i$  continues to arrive, the  $\mu_m$  is constantly updated based on  $x_i$ . Specifically, the visual encoder is used to map the image  $x_i$  to a visual embedding  $f_i^v$ . Then the probability that it belongs to each class embedding  $\mu_m$  is calculated based on Eq. (3):

$$P(\boldsymbol{\mu}_{m}|\boldsymbol{x}_{i}) = \frac{P(\boldsymbol{x}_{i}|\boldsymbol{\mu}_{m})}{\sum_{j=1}^{M} P(\boldsymbol{x}_{i}|\boldsymbol{\mu}_{j})} = \frac{\exp(\cos(\boldsymbol{f}_{i}^{v},\boldsymbol{\mu}_{m}))}{\sum_{j=1}^{M} \exp(\cos(\boldsymbol{f}_{i}^{v},\boldsymbol{\mu}_{j}))}$$
(5)

Then the class embedding  $\mu_s$  with highest probability is selected. If the probability  $P(\mu_s | x_i)$  exceeds the threshold  $\tau$ ,  $\mu_s$  is updated as follows:

$$\mu_{s} = Norm(\frac{C_{1}[s] * \mu_{s} + f_{i}^{v}}{C_{1}[s] + 1}),$$
  

$$C_{1}[s] = C_{1}[s] + 1,$$
(6)

where  $Norm(\cdot)$  denotes vector normalization, and  $C_1[s]$  is a counter initialized as  $n_1$  to count samples for the  $\mu_s$ .

#### **3.3. Prior Adaptation**

As mentioned earlier, the prior  $P(Y|\mu_m)$  in pre-trained CLIP model is a one-hot vector, where the m%K-th element is 1. However, the prior should be different in different environments. For example, consider a person who has a fever. During the COVID-19 period, we naturally think that he is more likely to have COVID-19, whereas in normal period, we might assume he is more likely to have common cold. This illustrates that the prior P(disease|fever) can vary significantly depending on the environment. Therefore, it is crucial to develop a more flexible and adaptive prior that can handle distribution shifts and improve the robustness of the model in real-world applications.

In our method, the prior  $P(Y|\boldsymbol{\mu}_m)$  is initialized as a *K*-dimensional one-hot vector, in which its m%K-th element is 1 from original CLIP. And the prior will be continuously updated as the samples arrive. When the *i*-th image  $\boldsymbol{x}_i$  arrives, the posterior  $P(Y|\boldsymbol{x}_i)$  is calculated based on the current prior  $P(Y|\boldsymbol{\mu}_m)$  and the probability  $P(\boldsymbol{x}_i|\boldsymbol{\mu}_m)$ , which is calculated in the likelihood adaptation, according to Eq. (2) as follows:

$$P(Y|\boldsymbol{x}_i) = \sum_{m=1}^{M} \frac{\exp(\cos(\boldsymbol{f}_i^v, \boldsymbol{\mu}_m))}{\sum_{j=1}^{M} \exp(\cos(\boldsymbol{f}_i^v, \boldsymbol{\mu}_j))} * P(Y|\boldsymbol{\mu}_m).$$
(7)

where the posterior  $P(Y|\mathbf{x}_i)$  is also the prediction for  $\mathbf{x}_i$ . Then this posterior is also used to update the prior of *s*-th class embedding  $P(Y|\boldsymbol{\mu}_s)$ , where  $\boldsymbol{\mu}_s$  is the one selected in the likelihood adaptation, as follows:

$$P(Y|\boldsymbol{\mu}_{s}) = \frac{\boldsymbol{C}_{2}[s] * P(Y|\boldsymbol{\mu}_{s}) + P(Y|\boldsymbol{x}_{i})}{\boldsymbol{C}_{2}[s] + 1},$$
  
$$\boldsymbol{C}_{2}[s] = \boldsymbol{C}_{2}[s] + 1,$$
(8)

where  $C_2[s]$  is a counter initialized as  $n_2$  to count samples for the s-th class embedding. It is worth noting that there are two different counters  $C_1$  and  $C_2$  in our method. The most intuitive understanding is that the learning rates used to update the prior  $P(Y|\mu_s)$  and class embedding  $\mu_s$  are different. By using the  $P(Y|\mathbf{x}_i)$  to update  $P(Y|\mu_s)$ , the model can continuously adapt to the evolving test distribution, leading to more accurate and robust predictions.

### 3.4. Our method

By integrating the above two points, we first initialize M class embeddings based on hand-crafted prompts, and initialize the prior of each class embedding as a one-hot vector, where the corresponding category is set to 1, that is defining  $U = [\mu_1^T; \mu_2^T; \cdots; \mu_M^T] \in \mathbb{R}^{d*M}$ , where d is the is the dimension of the embedding and T represents transpose, and  $V = [P(Y|\mu_1), P(Y|\mu_2), \cdots, P(Y|\mu_M)] \in \mathbb{R}^{M*K}$ . In addition,  $C_1$  and  $C_2$  are initialized to an M-dimensional vector with the values  $n_1$  and  $n_2$  respectively to update the

### Algorithm 1 Bayesian Class Adaptation

**Input:** A pre-trained CLIP model, unlabeled test data  $\{(x_i)\}_{i=1}^n$ , M hand-crafted prompts, hyperparameters  $\tau$ ,  $n_1$  and  $n_2$ 

## Procedure:

- 1: Initialize the *M* class embeddings  $U = [\mu_1^T; \dots; \mu_M^T]$ by encoding *M* hand-crafted prompts with text encoder of pre-trained CLIP;
- 2: Initialize the prior  $V = [P(Y|\mu_1), \dots, P(Y|\mu_M)]$ , where the *m*-th row is a one-hot vector with the m% Kth element being 1;
- Initialize C<sub>1</sub> and C<sub>2</sub> as an M-dimensional vector with the values n<sub>1</sub> and n<sub>2</sub> respectively;
- 4: **for** *i* = 1:*n* **do**
- 5: Map test image  $x_i$  to visual embedding  $f_i^v$ ;
- 6: Calculate the probability  $P(\boldsymbol{U}|\boldsymbol{x}_i)$  of belonging to each class embedding and posterior  $P(Y|\boldsymbol{x}_i)$  with current  $\boldsymbol{U}$  and  $\boldsymbol{V}$  based on Eq. (9);
- 7: Select the most likely class embedding index s based on  $P(U|x_i)$ ;
- 8: **if**  $P(\boldsymbol{U}[s]|\boldsymbol{x}_i) > \tau$  then
- 9: Adapt U[s] with  $f_i^v$ , and  $C_1[s]$  based on Eq. (6);
- 10: Adapt V[s] with  $P(Y|\boldsymbol{x}_i)$ , and  $C_2[s]$  based on Eq. (8);
- 11: end if
- 12: **return** posterior  $P(Y|\boldsymbol{x}_i)$ .

13: end for

class embedding and its prior. As the test data  $x_i$  continues to arrive, we first calculate the probability belongs to each class embedding based on Eq. (5) and posterior based on Eq. (4), which can be rewritten in matrix form as follows:

$$P(\boldsymbol{U}|\boldsymbol{x}_i) = \texttt{Softmax}(\boldsymbol{f}_i^v * \boldsymbol{U}),$$
  

$$P(Y|\boldsymbol{x}_i) = \texttt{Softmax}(\boldsymbol{f}_i^v * \boldsymbol{U}) * \boldsymbol{V},$$
(9)

where  $Softmax(\cdot)$  represents the softmax function.

Then the most likely class embedding  $\mu_s$  is selected based on  $P(\boldsymbol{U}|\boldsymbol{x}_i)$ , where  $s = \arg\max_m P(\boldsymbol{U}[m]|\boldsymbol{x}_i)$ . If  $P(\boldsymbol{U}[s]|\boldsymbol{x}_i) > \tau$ ,  $\boldsymbol{U}[s]$  and  $\boldsymbol{V}[s]$  are updated by the visual embedding  $\boldsymbol{f}_i^v$  and the posterior  $P(Y|\boldsymbol{x}_i)$ , which affect the likelihood and prior. These updates are performed using statistical method as shown in Eq. (6) and Eq. (8) respectively. Our method is summarized in Algorithm 1.

#### 4. Experiments

## 4.1. Experimental Setup

**Benchmarks.** In our experiments, we evaluated the effectiveness and robustness of our proposed method using two benchmarks: the Cross Domain benchmark and the Out-of-Distribution (OOD) benchmark. These benchmarks have been previously utilized in related research [38] to test the

Table 1. Comparison with the state-of-the-art methods on *OOD benchmark*. Metric: classification accuracy (%); Bp-free: backpropagation-free at test time; Average: mean accuracy across all datasets; OOD Average: mean accuracy across four OOD datasets excluding ImageNet.

Visual Backbone: ResNet-50										
Method	Venue	Bp-free	ImageNet	ImageNet-A	ImageNet-V2	ImageNet-R	ImageNet-S	Average	OOD Average	
TPT [38]	NIPS22	X	60.74	26.67	54.70	59.11	35.09	47.26	43.89	
DiffTPT [9]	ICCV23	X	60.80	31.06	55.80	58.80	37.10	48.71	45.69	
C-TPT [48]	ICLR24	×	61.2	25.6	54.8	59.7	35.7	47.4	44.0	
TDA [19]	CVPR24	1	61.35	30.29	55.54	62.58	38.12	49.58	46.63	
CLIP [33]	ICML21	1	59.81	23.24	52.91	60.72	35.48	46.43	43.09	
CLIP with LA	Ours	1	60.95	28.53	55.59	61.88	37.21	48.83	45.80	
CLIP with PA	Ours	1	61.15	28.66	55.41	61.97	37.41	48.92	45.86	
BCA	Ours	1	61.81	30.35	56.58	62.89	38.08	49.94	46.98	
Visual Backbone: ViT-B/16										
TPT [38]	NIPS22	X	68.98	54.77	63.45	77.06	47.94	62.44	60.81	
DiffTPT [9]	ICCV23	X	70.30	55.68	65.10	75.00	46.80	62.28	60.52	
C-TPT [48]	ICLR24	X	69.3	52.9	63.4	78.0	48.5	62.4	60.7	
TDA [19]	CVPR24	1	69.51	60.11	64.67	80.24	50.54	65.01	63.89	
MTA [49]	CVPR24	1	70.08	58.06	64.24	78.33	49.61	64.06	62.56	
PromptAlign [36]	NIPS24	×	-	59.37	65.29	79.33	50.23	-	63.55	
CLIP [33]	ICML21	1	68.34	49.89	61.88	77.65	48.24	61.20	59.42	
CLIP with LA	Ours	1	68.99	58.62	63.64	79.71	49.58	64.11	62.89	
CLIP with PA	Ours	1	69.16	59.39	64.21	79.68	49.88	64.46	63.29	
BCA	Ours	~	70.22	61.14	64.90	80.72	50.87	65.37	64.16	

adaptability of vision-language models during the inference phase. The Cross Domain benchmark, on the other hand, evaluates the model's performance across a wide range of image classification tasks from different domains, consisting of ten datasets: Aircraft [26], Caltech101 [8], Cars [20], DTD [4], EuroSAT [13], Flower102 [28], Food101 [2], Pets [32], SUN397 [44], and UCF101 [39]. Each dataset represents a unique domain with its own set of classes, allowing us to assess the model's adaptability and generalization capabilities in real-world scenarios where the class spaces may vary significantly. The OOD benchmark focuses on assessing the model's ability to handle data that differs significantly from the training set, using four specific datasets derived from ImageNet [5]: ImageNet-A [15], ImageNet-V2 [34], ImageNet-R [14], and ImageNet-S [42]. These datasets are designed to challenge the model's generalization to novel and unseen scenarios, providing insights into its robustness and reliability. The OOD benchmark focuses on assessing the model's ability to handle data that differs significantly from the training set, using four specific datasets derived from ImageNet [5]: ImageNet-A [15], ImageNet-V2 [34], ImageNet-R [14], and ImageNet-S [42]. These datasets are designed to challenge the model's generalization to novel and unseen scenarios, providing insights into its robustness and reliability.

**Implementation details.** Our experiment is conducted on the Pytorch platform. Following [9, 38], we adopt pre-trained CLIP models with ResNet-50 [12] and ViT-B/16 [6] backbones as the visual encoder and a Transformer [40] as

the text encoder in our experiments. The batch size is set to 1 to suit the application scenario of test time adaptation. For hyperparameters  $\tau/n_1/n_2$ , these set as 0.3/30000/10in OOD benchmark, and set as 0.35/50000/10 in Cross Domain benchmark. The top-1 accuaracy is used as evaluation metric. All the experiments are conducted on RTX 4070 Ti SUPER GPU.

#### 4.2. Comparisons with State-of-the-art

For fair comparisons with other method, we set M = K in this section, and follow [19] to use multiple context prompt templates for prompt ensembling and generate K initialized class embeddings for image classification. And we compare our method with the following two major categories of methods: (1) Backpropagation-based methods: these methods require backpropagation during the test time, such as TPT [38], DiffTPT [9], C-TPT [48], PromptAlign [36] and HisTPT[51]. (2) Backpropagation-free methods: these methods do not require backpropagation and are designed to be computationally efficient, such as TDA [19], MTA [49] and Zero [7]. In addition, we also report our method in detail, where CLIP [33] represents our baseline method; CLIP with LA represents the addition of Likelihood adaptation on the basis of CLIP; CLIP with PA represents the addition of prior adaptation on the basis of CLIP; BCA represents the complete algorithm.

**Results on the OOD benchmark.** We first compare BCA with state-of-the-art methods over the OOD benchmark. Table 1 presents the experimental results, highlighting the superior performance of the proposed BCA compared to other

Visual Backbone: ResNet-50											
Method	Aircraft	Caltech101	Cars	DTD	EuroSAT	Flower102	Food101	Pets	SUN397	UCF101	Average
TPT [38]	17.58	87.02	58.46	40.84	28.33	62.69	74.88	84.49	61.46	60.82	57.66
DiffTPT [9]	17.60	86.89	60.71	40.72	41.04	63.53	79.21	83.40	62.72	62.67	59.85
C-TPT [48]	17.5	87.4	57.3	43.1	29.4	65.3	76.0	84.0	62.1	60.7	58.3
TDA [19]	17.61	89.70	57.78	43.74	42.11	68.74	77.75	86.18	62.53	64.18	61.03
HisTPT[51]	18.1	87.2	61.3	41.3	42.5	67.6	81.3	84.9	63.5	64.1	61.2
CLIP [33]	16.11	87.26	55.89	40.37	25.79	62.77	74.82	82.97	60.85	59.48	56.63
CLIP with LA	17.85	88.96	57.11	43.55	35.50	65.57	76.94	84.24	62.54	60.71	59.30
CLIP with PA	18.10	89.08	57.56	46.06	38.53	65.02	76.92	85.28	62.20	60.45	59.92
BCA	19.89	89.70	58.13	48.58	42.12	66.30	77.19	85.58	63.38	63.51	61.44
Visual Backbone: ViT-B/16											
TPT [38]	24.78	94.16	66.87	47.75	42.44	68.98	84.67	87.79	65.50	68.04	65.10
DiffTPT [9]	25.60	92.49	67.01	47.00	43.13	70.10	87.23	88.22	65.74	62.67	65.47
C-TPT [48]	23.9	94.1	66.7	46.8	48.7	69.9	84.5	87.4	66.0	66.7	65.5
TDA [19]	23.91	94.24	67.28	47.40	58.00	71.42	86.14	88.63	67.62	70.66	67.53
MTA [49]	25.20	94.21	68.47	45.90	45.36	68.06	85.00	88.24	66.67	68.69	65.58
PromptAlign [36]	24.80	94.01	68.50	47.24	47.86	72.39	86.65	90.76	67.54	69.47	66.92
HisTPT[51] (NIPS24)	26.9	94.5	69.2	48.9	49.7	71.2	89.3	89.1	67.2	70.1	67.6
Zero [7] (NIPS24)	24.42	94.14	68.48	45.86	43.77	66.82	84.58	87.20	66.90	68.57	65.07
CLIP [33]	23.22	93.55	66.11	45.04	50.42	66.99	82.86	86.92	65.63	65.16	64.59
CLIP with LA	26.85	93.50	65.34	51.89	50.82	71.41	85.60	89.99	67.09	66.87	66.94
CLIP with PA	26.70	94.19	65.63	52.60	55.17	72.71	85.75	90.30	68.04	67.01	67.81
BCA	28.59	94.69	66.86	53.49	56.63	73.12	85.97	90.43	68.41	67.59	68.59

Table 2. Comparison with the state-of-the-art methods on Cross Domain benchmark. Metric: classification accuracy (%).

methods across various OOD datasets derived from ImageNet. Specifically, for the ResNet-50 backbone, BCA outperforms TDA by an average of 0.36 in overall accuracy and 0.35 in OOD accuracy. For instance, BCA achieves 61.81 accuracy on the ImageNet dataset, a 0.46 improvement over TDA's 61.35 and a 0.61 improvement over C-TPT's 61.2. On the ImageNet-V2 dataset, BCA reaches 56.58 accuracy, a 1.04 improvement over TDA's 55.54 and a 0.78 improvement over DiffTPT's 55.80. On the ImageNet-R dataset, BCA attains 62.89 accuracy, a 0.31 improvement over TDA's 62.58 and a 3.19 improvement over C-TPT's 59.7. For the ViT-B/16 backbone, BCA continues to exhibit strong performance. On the ImageNet-R dataset, BCA achieves 80.72 accuracy, a 0.48 improvement over TDA's 80.24 and a 1.39 improvement over PromptAlign's 79.33. On the ImageNet-S dataset, BCA reaches 50.87 accuracy, a 0.33 improvement over TDA's 50.54 and a 0.64 improvement over PromptAlign's 50.23. Overall, BCA achieves an average accuracy of 65.37 and an OOD average accuracy of 64.16, representing improvements of 0.36 and 0.27 over TDA's 65.01 and 63.89, respectively.

**Results on the Cross Domain benchmark.** We also compare BCA with state-of-the-art methods over the cross-domain benchmark. Table 2 presents the experimental results, highlighting the superior performance of the proposed BCA across various cross-domain datasets. Specifically, for the ResNet-50 backbone, BCA achieves 19.89 accuracy on the Aircraft dataset, a 2.28 improvement over TDA's 17.61 and a 1.79 improvement over HisTPT's 18.1. On the DTD dataset, BCA attains 48.58 accuracy, a 4.84 improvement over TDA's 43.74 and a 5.48 improvement over C-TPT's

43.1. Overall, BCA achieves an average accuracy of 61.44, representing a 0.41 improvement over TDA's 61.03 and a 0.24 improvement over HisTPT's 61.2. For the ViT-B/16 backbone, BCA continues to exhibit strong performance. On the Aircraft dataset, BCA achieves 28.59 accuracy, a 4.68 improvement over TDA's 23.91 and a 1.69 improvement over HisTPT's 26.9. On the DTD dataset, BCA attains 53.49 accuracy, a 6.09 improvement over TDA's 47.40 and a 4.59 improvement over HisTPT's 48.9. On the Flower102 dataset, BCA reaches 73.12 accuracy, a 1.70 improvement over TDA's 71.42 and a 0.72 improvement over HisTPT's 72.4. Overall, BCA achieves an average accuracy of 68.59, representing a 1.06 improvement over TDA's 67.53 and a 0.99 improvement over HisTPT's 67.6.

### 4.3. Ablation Studies

**Component analysis.** In this experiment, we validated the effectiveness of each component of the proposed method on two benchmarks as shown in Table 1 and Table 2. Based on the experimental results, we have summarized the following patterns: 1) Adapting likelihood and prior can both enable the CLIP model to better adapt to the test environment: By comparing the original CLIP model with the CLIP models that incorporate LA and PA, we found that CLIP with LA and CLIP with PA achieved significant improvements in classification accuracy across all datasets. This demonstrates that adapting the pre-trained CLIP model to the test environment can effectively improve the model's accuracy, thereby validating the correctness of Analysis in Section 3.1. 2) Adapting likelihood and prior can work synergistically: By comparing BCA with CLIP with LA and BCA



Figure 3. Sensitivity analysis with respect to  $\tau/n_1/n_2$  on ImageNet for OOD benchmark and Aircraft for Cross Domain benchmark using ViT-B/16 as the visual backbone.

with CLIP with PA, it is evident that BCA achieved even further performance enhancements, indicating that there is a complementary effect between LA and PA, allowing them to combine to enhance the model's generalization ability. 3) Prior adaptation is equally crucial: By comparing CLIP with PA to the original CLIP and BCA to CLIP with LA, we found that adding prior adaptation could significantly boost the model's accuracy, which is often ignored in previous model designs, highlighting the importance of domainspecific prior knowledge in improving the model's crossdomain adaptability.

Table 3. Comparison with the state-of-the-art methods on *Ima-geNet* in terms of efficiency (Testing Time, Memory usage) and effectiveness (Accuracy). Visual Backbone: ResNet-50; GPU: RTX 4070 Ti SUPER GPU.

Method	Testing time	Memory usage	Accuracy	Gain
CLIP	2.23min	753M	59.81	0
TPT	572.13min	21396M	60.74	+0.93
TDA	11.93min	1174M	61.35	+1.54
BCA	2.42min	757M	61.81	+2.00

Inference time comparison. In this experiment, we evaluated the efficiency and effectiveness of the proposed BCA method on the ImageNet dataset using ResNet-50 as the visual backbone on RTX 4070 Ti SUPER GPU. We compared our BCA with TPT and TDA, the experimental results are shown in Table 3. From the table, we can see that the BCA method outperforms other methods in several aspects: 1) Inference Time: The inference time of the BCA method is significantly reduced compared to both methods that require gradient backpropagation and those that do not. For example, TPT takes 572.13 minutes, while BCA only takes 2.42 minutes. Even compared to TDA, which does not require gradient backpropagation, BCA is still much faster, taking only 2.42 minutes compared to TDA's 11.93 minutes. 2) Memory Usage: The memory usage of the BCA method is also significantly lower than that of other methods, especially those that require gradient backpropagation. For example, TPT uses 21396M of memory, while BCA uses only 757M. Compared to TDA, which uses 1174M, BCA

still uses less memory, which is on par with the baseline CLIP. 3) Classification Accuracy: Despite its excellent performance in memory usage and inference time, BCA also achieves higher classification accuracy. For example, BCA achieves an accuracy of 61.81, which is a 2.00 improvement over the baseline CLIP. This is slightly better than TDA, which achieves an accuracy of 61.35.

**Parameter sensitivity analysis.** In this experiment, we investigated the settings of hyperparameters  $\tau/n_1/n_2$  to analyze the sensitivity of our method. We conducted experiments on the ImageNet for OOD benchmark and Aircraft for Cross Domain benchmark using ViT-B/16 as the visual backbone. The experimental results are shown in Figure 3. It can be found that the accuracy first increases and then decreases with the increase of the three hyperparameter values. And our method is also robust to the hyperparameters  $\tau/n_1/n_2$ . This suggests that our method maintains good performance under various hyperparameter settings, making it suitable for different tasks and datasets.

### 5. Conclusion

In this work, we propose Bayesian Class Adaptation (BCA). Unlike previous methods that focus solely on likelihood estimation, BCA incorporates the continuous updating of the prior distribution, which is crucial for handling distribution shifts in real-world applications. By employing a statistical approach to update both the likelihood and the prior, BCA enhances the robustness and accuracy of zero-shot image classification, particularly when the test data distribution differs significantly from the pre-trained data. Our experiments demonstrate that BCA not only outperforms existing methods on Out-of-Distribution (OOD) and Cross Domain tasks but also maintains high inference speed and efficiency, making it a effective solution for real-world scenarios.

#### Acknowledgement

This work was supported in part by Chinese National Natural Science Foundation Projects U23B2054, 62276254, 62306313, the Beijing Science and Technology Plan Project Z231100005923033, Beijing Natural Science Foundation L221013, and InnoHK program.

### References

- Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring visual prompts for adapting largescale models. *arXiv preprint arXiv:2203.17274*, 2022. 2
- [2] 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. 6
- [3] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointlyscaled multilingual language-image model. In *The Eleventh International Conference on Learning Representations*. 2
- [4] Mircea Cimpoi, Subhransu 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. 6
- [5] 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. 6
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 6
- [7] Matteo Farina, Gianni Franchi, Giovanni Iacca, Massimiliano Mancini, and Elisa Ricci. Frustratingly easy test-time adaptation of vision-language models. In Advances in Neural Information Processing Systems, 2024. 3, 6, 7
- [8] 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. 6
- [9] Chun-Mei Feng, Kai Yu, Yong Liu, Salman Khan, and Wangmeng Zuo. Diverse data augmentation with diffusions for effective test-time prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2704–2714, 2023. 1, 2, 3, 4, 6, 7
- [10] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *International Journal of Computer Vision*, 132(2): 581–595, 2024. 2
- [11] Zongbo Han, Jialong Yang, Junfan Li, Qinghua Hu, Qianli Xu, Mike Zheng Shou, and Changqing Zhang. Dota: Distributional test-time adaptation of vision-language models. arXiv preprint arXiv:2409.19375, 2024. 1, 2, 3, 4
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 6
- [13] 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. 6

- [14] 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. 6
- [15] 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. 6
- [16] Yusuke Iwasawa and Yutaka Matsuo. Test-time classifier adjustment module for model-agnostic domain generalization. Advances in Neural Information Processing Systems, 34:2427–2440, 2021. 3
- [17] Minguk Jang, Sae-Young Chung, and Hye Won Chung. Testtime adaptation via self-training with nearest neighbor information. In *The Eleventh International Conference on Learning Representations*. 3
- [18] 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
- [19] Adilbek Karmanov, Dayan Guan, Shijian Lu, Abdulmotaleb El Saddik, and Eric Xing. Efficient test-time adaptation of vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14162–14171, 2024. 1, 2, 3, 4, 6, 7
- [20] 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. 6
- [21] John K Kruschke. Bayesian data analysis. Wiley Interdisciplinary Reviews: Cognitive Science, 1(5):658–676, 2010.
   3
- [22] Samuel Lavoie, Polina Kirichenko, Mark Ibrahim, Mido Assran, Andrew Gordon Wilson, Aaron Courville, and Nicolas Ballas. Modeling caption diversity in contrastive visionlanguage pretraining. In *Forty-first International Conference* on Machine Learning. 1
- [23] Shuaifeng Li, Mao Ye, Xiatian Zhu, Lihua Zhou, and Lin Xiong. Source-free object detection by learning to over-look domain style. In *Proceedings of the IEEE/CVF con-ference on computer vision and pattern recognition*, pages 8014–8023, 2022. 1
- [24] Shuaifeng Li, Mao Ye, Lihua Zhou, Nianxin Li, Siying Xiao, Song Tang, and Xiatian Zhu. Cloud object detector adaptation by integrating different source knowledge. In Advances in Neural Information Processing Systems, pages 25251– 25283, 2024. 1
- [25] Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, pages 1–34, 2024. 1

- [26] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151, 2013.
   6
- [27] M Jehanzeb Mirza, Jakub Micorek, Horst Possegger, and Horst Bischof. The norm must go on: Dynamic unsupervised domain adaptation by normalization. In *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 14765–14775, 2022. 3
- [28] 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. 6
- [29] Shuaicheng Niu, Jiaxiang Wu, Yifan Zhang, Zhiquan Wen, Yaofo Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. In *The Eleventh International Conference on Learning Representations*. 3
- [30] Shuaicheng Niu, Jiaxiang Wu, Yifan Zhang, Yaofo Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. In *International conference on machine learning*, pages 16888–16905. PMLR, 2022. 3
- [31] Omiros Pantazis, Gabriel Brostow, Katherine Jones, and Oisin Mac Aodha. Svl-adapter: Self-supervised adapter for vision-language pretrained models. In *Proceedings of The* 33rd British Machine Vision Conference. The British Machine Vision Association (BMVA), 2022. 2
- [32] 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. 6
- [33] 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, 3, 6, 7
- [34] 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. 6
- [35] Jintao Rong, Hao Chen, Tianxiao Chen, Linlin Ou, Xinyi Yu, and Yifan Liu. Retrieval-enhanced visual prompt learning for few-shot classification. *arXiv preprint arXiv:2306.02243*, 2023. 2
- [36] 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 Advances in Neural Information Processing Systems, 2024. 3, 6, 7
- [37] Sheng Shen, Shijia Yang, Tianjun Zhang, Bohan Zhai, Joseph E Gonzalez, Kurt Keutzer, and Trevor Darrell. Multitask vision-language prompt tuning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 5656–5667, 2024. 2
- [38] 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 visionlanguage models. In *Advances in Neural Information Processing Systems*, 2022. 1, 2, 3, 4, 5, 6, 7

- [39] K Soomro. Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402, 2012. 6
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, pages 6000–6010, 2017. 6
- [41] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*. 3
- [42] 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. 6
- [43] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7201–7211, 2022. 3
- [44] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In 2010 IEEE computer society conference on computer vision and pattern recognition, pages 3485–3492. IEEE, 2010. 6
- [45] Yizhe Xiong, Hui Chen, Zijia Lin, Sicheng Zhao, and Guiguang Ding. Confidence-based visual dispersal for fewshot unsupervised domain adaptation. In *Proceedings of* the IEEE/CVF International Conference on Computer Vision (ICCV), pages 11621–11631, 2023. 1
- [46] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15671– 15680, 2022. 1
- [47] 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. In *International Conference on Learning Representations*. 2
- [48] Hee Suk Yoon, Eunseop Yoon, Joshua Tian Jin Tee, Mark A Hasegawa-Johnson, Yingzhen Li, and Chang D Yoo. C-tpt: Calibrated test-time prompt tuning for vision-language models via text feature dispersion. In *The Twelfth International Conference on Learning Representations*. 3, 4, 6, 7
- [49] Maxime Zanella and Ismail Ben Ayed. On the test-time zeroshot generalization of vision-language models: Do we really need prompt learning? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23783–23793, 2024. 3, 6, 7
- [50] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Unified vision and language prompt learning. arXiv preprint arXiv:2210.07225, 2022. 2

- [51] Jingyi Zhang, Jiaxing Huang, Xiaoqin Zhang, Ling Shao, and Shijian Lu. Historical test-time prompt tuning for vision foundation models. In Advances in Neural Information Processing Systems, 2024. 3, 6, 7
- [52] Renrui Zhang, Rongyao Fang, Wei Zhang, Peng Gao, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better visionlanguage modeling. arXiv preprint arXiv:2111.03930, 2021.
- [53] Bowen Zhao, Chen Chen, and Shu-Tao Xia. Delta: Degradation-free fully test-time adaptation. In *The Eleventh International Conference on Learning Representations*. 3
- [54] 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. 2
- [55] 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. 2
- [56] Lihua Zhou, Siying Xiao, Mao Ye, Xiatian Zhu, and Shuaifeng Li. Adaptive mutual learning for unsupervised domain adaptation. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(11):6622–6634, 2023. 1
- [57] Lihua Zhou, Mao Ye, Xiatian Zhu, Siying Xiao, Xu-Qian Fan, and Ferrante Neri. Homeomorphism alignment for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (*ICCV*), pages 18699–18710, 2023. 1