

DECOOP: Robust Prompt Tuning with Out-of-Distribution Detection

Zhi Zhou¹ Ming Yang^{1,2} Jiang-Xin Shi^{1,2} Lan-Zhe Guo^{1,3} Yu-Feng Li^{1,2}

Abstract

Vision-language models (VLMs), such as CLIP, have demonstrated impressive zero-shot capabilities for various downstream tasks. Their performance can be further enhanced through few-shot prompt tuning methods. However, current studies evaluate the performance of learned prompts separately on base and new classes. This evaluation lacks practicality for real-world applications since downstream tasks cannot determine whether the data belongs to base or new classes in advance. In this paper, we explore a problem setting called *Open-world Prompt Tuning* (OPT), which involves tuning prompts on base classes and evaluating on a combination of base and new classes. By introducing *Decomposed Prompt Tuning* framework (DEPT), we theoretically demonstrate that OPT can be solved by incorporating out-of-distribution detection into prompt tuning, thereby enhancing the base-to-new discriminability. Based on DEPT, we present a novel prompt tuning approach, namely, *Decomposed Context Optimization* (DECOOP), which introduces new-class detectors and sub-classifiers to further enhance the base-class and new-class discriminability. Experimental results on 11 benchmark datasets validate the effectiveness of DEPT and demonstrate that DECOOP outperforms state-of-the-art methods, providing a significant 2% average accuracy improvement.

1. Introduction

Vision-language models (VLMs), such as CLIP (Radford et al., 2021), have been developed to align images and language, demonstrating impressive zero-shot capabilities for

¹National Key Laboratory for Novel Software Technology, Nanjing University, China ²School of Artificial Intelligence, Nanjing University, China ³School of Intelligence Science and Technology, Nanjing University, China. Correspondence to: Lan-Zhe Guo <guolz@nju.edu.cn>, Yu-Feng Li <liyf@nju.edu.cn>.

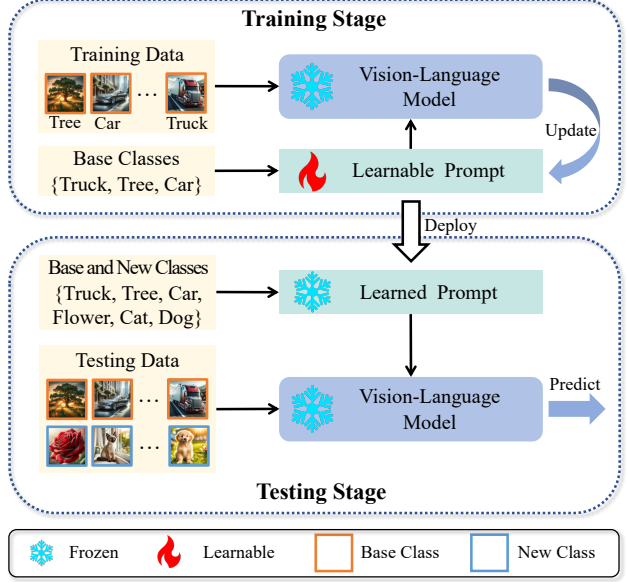


Figure 1. An illustration of the OPT evaluation paradigm. During the training, we finetune the model with data from base classes. During the testing, we evaluate the model on a mix of base and new classes.

a variety of downstream tasks (Deng et al., 2009; Maji et al., 2013; Krause et al., 2013), using only class names. The classification prediction is determined by calculating the cosine similarity between the image embedding, generated by the image encoder, and the text embedding, generated by the text encoder, using prompting techniques (Liu et al., 2023). For example, by inputting “a photo of class”, the text encoder generates the corresponding text embedding, where “class” represents the class name.

In addition, it is possible to improve the performance of CLIP, particularly when dealing with downstream tasks that have limited labeled data. Few-shot prompt tuning methods (Lu et al., 2022; Zhou et al., 2022b; Shu et al., 2022b) utilize a small amount of labeled data from downstream datasets to fine-tune learnable prompts while keeping the other parameters unchanged. These approaches can yield substantial performance improvement compared to the zero-shot VLMs in downstream classification tasks. However, previous studies (Zhou et al., 2022a; Wang et al., 2023b)

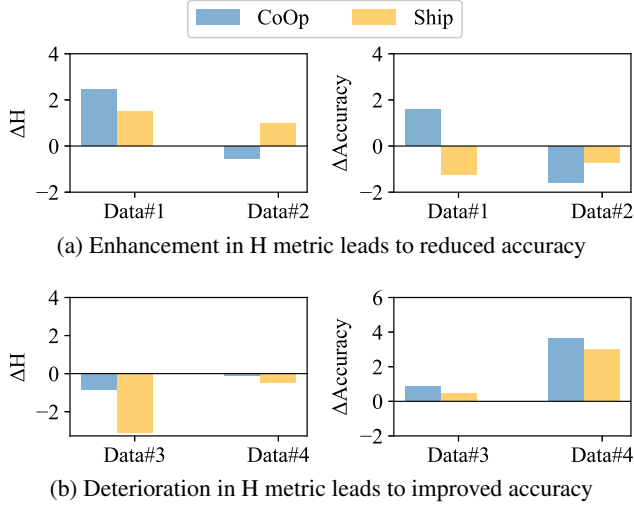


Figure 2. Delta performance of CoOp and SHIP method compared to zero-shot baseline CLIP method. Subfigures (a) and (b) show that the changes in the H metric are not necessary indicators of performance improvements or degradation of accuracy, highlighting the significance of addressing the OPT problem.

have identified a limitation in which the learned prompts only operate effectively with labeled data from base classes. This limitation leads to a decrease in zero-shot performance for new classes which are unseen in the training set. To address this, the researchers propose an evaluation paradigm that assesses the performance of both base and new classes separately, as well as their harmonic mean, i.e., H metric.

Although this evaluation paradigm can comprehensively evaluate the performance of both base and new classes, it lacks practicality for real-world applications, which necessitate prior knowledge of whether the data belongs to base or new classes in the downstream task. For instance, in the context of biological underpinnings (Hayes et al., 2021; Kudithipudi et al., 2022) and visual classification (Lange et al., 2022; Mai et al., 2022), both base classes and new classes that arise during testing will be evaluated together. Therefore, we introduce a realistic problem setting, namely, *Open-world Prompt Tuning* (OPT), which evaluates the performance of the model on a mix of base and new classes while training model with base classes. An illustration of the OPT problem is shown in Figure 1. The results in Figure 2 show that the changes in the H metric are not necessary indicators of performance improvement or degradation when evaluating the combination of base and new classes, which highlights the significance of the OPT problem.

To address the OPT problem, we first analyze the original problem, which consists of three parts: base-to-new discriminability, base-class discriminability, and new-class discriminability. We observe that existing methods and settings fail to adequately consider the base-to-new discriminability. Motivated by this analysis, we propose the

DEPT framework, which incorporates out-of-distribution (OOD) detection into prompt tuning to enhance the base-to-new discriminability and thereby prevents performance degradation on new classes. We theoretically prove that the DEPT framework can improve performance compared to the zero-shot baseline and prompt tuning methods. Building upon the DEPT framework, we introduce a novel prompt tuning approach called *Decomposed Context Optimization* (DECoOP). This approach incorporates new-class detectors and sub-classifiers to further enhance the base-class and new-class discriminability, respectively. Empirical results validate the effectiveness of the DEPT framework and demonstrate that DECoOP approach outperforms current state-of-the-art (SOTA) methods by a significant margin.

The contributions of this paper are summarized as follows:

- (1) We explore a practical OPT problem and break down the problem into two sub-problems: OOD detection and prompt tuning. Through decomposition, we uncover that base-to-new discriminability is crucial to address OPT, overlooked in existing methods and settings.
- (2) We propose a novel DEPT framework, which introduces OOD detection into prompt tuning. Both our theoretical analysis and experimental results demonstrate the effectiveness of DEPT framework for OPT.
- (3) Based on DEPT framework, we propose a novel prompt tuning approach DECoOP, which additionally enhances the base-class and new-class discriminability by introducing new-class detectors and sub-classifiers.
- (4) We conduct comprehensive experiments on DECoOP using 11 benchmark datasets. The results show that our proposed scheme outperforms current SOTA comparison methods, delivering a significant 2% average improvement in accuracy.

2. Problem and Analysis

In this section, we first describe the notions and problem formulation for the OPT setting. Subsequently, we conduct an empirical analysis using a real-world dataset (Krause et al., 2013), wherein we identify two primary challenges to address: base-to-new discriminability and new-class discriminability. Finally, we decompose the original problem to demonstrate that the incorporation of the OOD detection technique can effectively resolve these two challenges.

2.1. Problem Formulation

We focus on the prompt tuning setting for multi-class classification problems that involve an input space \mathcal{X} , a class space $\mathcal{Y} = \mathcal{Y}_b \cup \mathcal{Y}_n = [C]$, and the text space \mathcal{T} , where C represents the number of classes. Here, \mathcal{Y}_b denotes the set of

base classes, and \mathcal{Y}_n represents the set of new classes. The name of the i -th class is denoted as $t_i \in \mathcal{T}$. Furthermore, $\mathbf{x} \in \mathcal{X}$ represents the data. $f(\mathbf{x}) \in \mathcal{Y}$ and $g(\mathbf{x}) \in \{\mathbf{b}, \mathbf{n}\}$ denote the label of \mathbf{x} and the specific class space to which it belongs, where f and g are the mapping functions of the ground truth of the labels and the class space.

In OPT problem, we are given a pre-trained vision-language model $\mathcal{F} = \{E_V, E_T\}$, which consists of a visual encoder $E_V : \mathcal{X} \mapsto \mathbb{R}^d$ and a textual encoder $E_T : \mathcal{T} \mapsto \mathbb{R}^d$, where d represents the dimension of model \mathcal{F} . During the training stage, we learn the prompt vector \mathbf{p} on a few-shot dataset \mathcal{D} containing data derived from \mathcal{Y}_b . To simplify the notation, we define $t_i(\mathbf{p})$ as the concatenation of the tokens of the class name t_i and the learned prompt \mathbf{p} . Consequently, weight vectors $\{\mathbf{w}_i(\mathbf{p})\}_{i=1}^C$ are generated for each class as textual embeddings, where $\mathbf{w}_i(\mathbf{p}) = E_T(t_i(\mathbf{p})) / \|E_T(t_i(\mathbf{p}))\|$. In the testing stage, given the test data \mathbf{x} drawn from \mathcal{Y} , we initially obtain its visual embedding $\mathbf{z} = E_V(\mathbf{x}) / \|E_V(\mathbf{x})\|$. Subsequently, we calculate the prediction probabilities as follows:

$$P(y|\mathbf{x}) = \frac{\exp(\mathbf{z}^T \mathbf{w}_y / \tau)}{\sum_{i=1}^C \exp(\mathbf{z}^T \mathbf{w}_i / \tau)} \quad (1)$$

where τ represents the temperature determined by VLMs. For convenience, we will also use $P(\mathbf{x})$ to represent $P(y|\mathbf{x})$ in the subsequent paper. The prediction for \mathbf{x} is given by $\arg \max_{y \in \mathcal{Y}} P(y|\mathbf{x})$. The objective of OPT is to train a model that can make robust predictions on \mathcal{Y} , which includes both base and new classes, without experiencing overall performance degradation due to the presence of new classes. In our following analyses and experiments, we perform a comparison between the zero-shot baseline method (referred to as ZS) and the prompt tuning method (referred to as PT) on OPT problem.

2.2. Problem Analysis

To tackle the OPT problem, we investigate a real-world dataset (Krause et al., 2013) to conduct detailed analyses of the challenges inherent in OPT. Our observation demonstrates that while prompt tuning methods can improve base-class discriminability, they compromise both base-to-new discriminability and new-class discriminability. To illustrate this observation, we present a comparison between the ZS methods and PT methods, where we employ CLIP as ZS method and COOP as PT method, in Figures 3 and 4.

Figure 3 indicates that the prompt tuning method results in a decreased base-to-new discriminability compared to the zero-shot baseline. Specifically, the AUROC for detecting new classes using the MSP technique (Hendrycks & Gimpel, 2016) decreases, and more false positive predictions are introduced for base classes. Moreover, Figure 4 illustrates




Figure 3. Performance of ZS and PT methods to distinguish data from base classes and new classes (base-to-new discriminability).




Figure 4. Performance of ZS and PT methods to distinguish data within new classes (new-class discriminability).

that the prompt tuning method also exhibits reduced new-class discriminability compared to the zero-shot baseline.

We emphasize that the existing H metric is incapable of measuring base-to-new discriminability, making it unsuitable for comprehensive practical applications. In OPT problem, the accuracy evaluated in the entire class space can effectively address this limitation.

2.3. Problem Decomposition

The above analysis reveals that the zero-shot baseline surpasses the prompt tuning method in terms of both new-class discriminability and base-to-new discriminability. This observation motivates us to incorporate OOD detection technique to combine ZS method and PT method. This approach aims to preserve the new-class discriminability using ZS while enhancing the base-class discriminability using PT. Therefore, we decompose the original classification problem into separate OOD detection and two classification problems:

$$\begin{aligned} P(y|\mathbf{x}) &= \sum_{i \in \{\mathbf{b}, \mathbf{n}\}} P(y|y \in \mathcal{Y}_i, \mathbf{x}) \cdot P(y \in \mathcal{Y}_i|\mathbf{x}) \\ &= P(y|y \in \mathcal{Y}_k, \mathbf{x}) \cdot P(y \in \mathcal{Y}_k|\mathbf{x}) \end{aligned} \quad (2)$$

where k always equals $g(\mathbf{x})$ for the sake of simplicity, representing the ground-truth label space of \mathbf{x} . The second term is an OOD detector to determine whether \mathbf{x} belongs to the base or new class space. The first term is a classifier for the corresponding class space.

Equation 2 motivates us to propose a novel *Decomposed*

Prompt Tuning framework (DEPT), which synergistically leverages the advantages of both the zero-shot baseline ZS and the prompt tuning method PT. The prediction probability $P_{\text{DEPT}}(y|\mathbf{x})$ of DEPT framework is:

$$\begin{cases} P_{\text{PT}}(y|\mathbf{x}), & P_{\text{OOD}}(y \in \mathcal{Y}_b|\mathbf{x}) \geq P_{\text{OOD}}(y \in \mathcal{Y}_n|\mathbf{x}), \\ P_{\text{ZS}}(y|\mathbf{x}), & P_{\text{OOD}}(y \in \mathcal{Y}_b|\mathbf{x}) < P_{\text{OOD}}(y \in \mathcal{Y}_n|\mathbf{x}). \end{cases} \quad (3)$$

where $P_{\text{OOD}}(y \in \mathcal{Y}_b|\mathbf{x})$ is the OOD detector to determine whether \mathbf{x} belongs to the base or new class space. $P_{\text{ZS}}(y|\mathbf{x})$ and $P_{\text{PT}}(y|\mathbf{x})$ are classifiers of ZS and PT. In following theoretical analysis and empirical experiment, we adopt the ZS method using MSP method as the OOD detector, i.e., $P_{\text{OOD}}(y \in \mathcal{Y}_i|\mathbf{x}) = \max_{j \in \mathcal{Y}_i} P_{\text{ZS}}(y = j|\mathbf{x})$ for $i \in \{b, n\}$.

Then, we adopt the cross-entropy metric of two probability distributions \mathbf{p} and \mathbf{q} , i.e., $H(\mathbf{p}, \mathbf{q}) = -\sum_{i=1}^C p_i \log q_i$, to evaluate the performance of $P_{\text{ZS}}(y|\mathbf{x})$ and our DEPT framework $P_{\text{DEPT}}(y|\mathbf{x})$. We denote distributions $\mathbf{k} = \{\mathbb{I}[k=b], \mathbb{I}[k=n]\}$ and $\tilde{\mathbf{y}} = \{\mathbb{I}[f(\mathbf{x})=i]\}_{i=1}^C$ for \mathbf{x} . Finally, we denote the following cross-entropy values for zero-shot baseline, prompt tuning method, and DEPT framework:

$$\begin{aligned} H_{\text{ZS}}^{\text{OOD}}(\mathbf{x}) &= H(\tilde{\mathbf{k}}, \{P_{\text{ZS}}(y \in \mathcal{Y}_i|\mathbf{x})\}_{i=\{b,n\}}), \\ H_{\text{ZS}}^{\text{CLS}}(\mathbf{x}) &= H(\tilde{\mathbf{y}}, \{P_{\text{ZS}}(y = j|y \in \mathcal{Y}_k, \mathbf{x})\}_{j=1}^C), \\ H_{\text{PT}}^{\text{CLS}}(\mathbf{x}) &= H(\tilde{\mathbf{y}}, \{P_{\text{PT}}(y = j|y \in \mathcal{Y}_k, \mathbf{x})\}_{j=1}^C), \\ H_{\text{ZS}}(\mathbf{x}) &= H(\tilde{\mathbf{y}}, \{P_{\text{ZS}}(y = j|\mathbf{x})\}_{j=1}^C), \\ H_{\text{DEPT}}(\mathbf{x}) &= H(\tilde{\mathbf{y}}, \{P_{\text{DEPT}}(y = j|\mathbf{x})\}_{j=1}^C). \end{aligned} \quad (4)$$

Theorem 2.1. If $\mathbb{E}_{\mathbf{x}}[H_{\text{ZS}}^{\text{CLS}}(\mathbf{x})] \leq \delta$ for \mathbf{x} belonging to both base and new classes, $\mathbb{E}_{\mathbf{x}}[H_{\text{PT}}^{\text{CLS}}(\mathbf{x})] \leq \delta - \Delta$ for \mathbf{x} belonging to base classes, and $\mathbb{E}_{\mathbf{x}}[H_{\text{ZS}}^{\text{OOD}}(\mathbf{x})] \leq \epsilon$, given a uniform mixing ratio ($\alpha : 1 - \alpha$) of base classes and new classes in the testing data, we can determine that:

$$\begin{cases} \mathbb{E}_{\mathbf{x}}[H_{\text{ZS}}(\mathbf{x})] & \leq \epsilon + \delta, \\ \mathbb{E}_{\mathbf{x}}[H_{\text{DEPT}}(\mathbf{x})] & \leq \epsilon + \delta - \alpha \cdot \Delta. \end{cases} \quad (5)$$

Remark 2.2. Theorem 2.1 demonstrates that decomposing the zero-shot baseline into an OOD detector and classifiers, and incorporating prompt tuning methods to aid in classifying base classes, can effectively decrease the upper bound of classification error. Moreover, enhancing the reliability of the OOD detector helps reduce the error term ϵ and ensures that the performance on new classes remains uncompromised compared to the baseline method. Consequently, this framework preserves base-to-new discriminability and new-class discriminability of ZS method. Additionally, refining the PT method increases Δ , further enhancing base-class discriminability and reducing the upper bound of error.

The proof is presented in Appendix A. Theorem 2.1 motivates us to design a robust prompt tuning method based on Equation 3 using OOD detection techniques to solve OPT.

3. DECoOP Approach

We propose a novel prompt tuning framework, called DEPT, to address the OPT problem. The DEPT framework effectively maintains the discriminability between base classes and new classes, thus preventing degradation of discriminability when prompt tuning is applied. Our theoretical analysis, as presented in Theorem 2.1, demonstrates the superiority of DEPT when combining the zero-shot baseline and prompt tuning method. However, there are still two challenges that need to be addressed in order to further enhance the performance in complex real-world applications: (1) How can we train reliable OOD detectors to identify new-class data using limited labeled data from base classes? (2) With reliable OOD detectors, how to separately improve the base-class and new-class discriminability?

To tackle the challenges above, we present a novel prompt tuning approach named **Decomposed Context Optimization** (DECoOP) based on our DEPT framework, containing K new-classes detectors $\{\mathcal{M}_D^i\}_{i=1}^K$ and sub-classifiers $\{\mathcal{M}_C^i\}_{i=1}^K$. The introduction of new-class detectors aids in the improved detection of data from new classes in OPT problem, where the names of new classes are known and can be utilized. This differs from the traditional OOD detection problems and presents an opportunity for further performance enhancement. The sub-classifiers are designed to better classify the data from base classes and reduce the potential risks for new classes, which aims to enhance the base-class and new-class discriminability with a reliable base-to-new discriminability. The overall illustration of DECoOP approach is shown in Figure 5 and each component is described thoroughly in the following subsections.

3.1. New-class Detector \mathcal{M}_D

In the OPT problem, the model is trained with \mathcal{Y}_b but has knowledge of the entire class space \mathcal{Y} during testing. Therefore, the main challenge for new class detectors is to train the model to effectively utilize the knowledge of the new class \mathcal{Y}_n , which is only known during testing.

Specifically, Our proposed solution incorporates a leave-out strategy which divides the base class space \mathcal{Y}_b into two distinct subsets during training stage: simulated base classes $\hat{\mathcal{Y}}_b$ and simulated new classes $\hat{\mathcal{Y}}_n$, where $\hat{\mathcal{Y}}_b \cup \hat{\mathcal{Y}}_n = \mathcal{Y}_b$. Respectively, we split the original training set \mathcal{D} into $\mathcal{D}_b = \{(\mathbf{x}, y) | (\mathbf{x}, y) \sim \mathcal{D} \wedge y \in \hat{\mathcal{Y}}_b\}$ and $\mathcal{D}_n = \{(\mathbf{x}, y) | (\mathbf{x}, y) \sim \mathcal{D} \wedge y \in \hat{\mathcal{Y}}_n\}$. Then, our optimization objective function for the new-class detector is defined as:

$$\begin{aligned} \ell_{\text{OOD}} &= \frac{1}{|\mathcal{D}_b|} \sum_{(\mathbf{x}, y) \sim \mathcal{D}_b} \ell_{CE}(\mathbf{x}, y) \\ &\quad + \max\{0, \gamma + \ell_E^b - \ell_E^n\} \end{aligned} \quad (6)$$

where $\ell_{CE}(\mathbf{x}, y) = -\log P(\mathbf{x})_y$ represents the cross-




Figure 5. The overall illustration of DECoOP approach.

entropy loss, $\ell_E(\mathbf{x}) = -\sum_{i=1}^C P(\mathbf{x})_i \log P(\mathbf{x})_i$ represents the entropy loss, $\ell_E^b = \frac{1}{|\mathcal{D}_b|} \sum_{(\mathbf{x}, y) \sim \mathcal{D}_b} \ell_E(\mathbf{x})$ represents the average entropy on the simulated base classes, and $\ell_E^n = \frac{1}{|\mathcal{D}_n|} \sum_{(\mathbf{x}, y) \sim \mathcal{D}_n} \ell_E(\mathbf{x})$ represents the average entropy on the simulated new classes. Additionally, γ is a hyperparameter that controls the margin between ℓ_E^b and ℓ_E^n to ensure stable optimization. The objective function in Equation 6 encourages the model to make low-entropy predictions on simulated base classes and high-entropy predictions on simulated new classes, thereby enhancing base-to-new discriminability. However, partitioning the base class space causes the model’s cognition to be limited to a subset of base classes, leading to the failure to distinguish between other base classes and new classes during testing. To address this issue, we propose the adoption of an ensemble of K new-class detectors $\{\mathcal{M}_D^i\}_{i=1}^K$ that cover the entire base class space during training. Each new-class detector is trained with Equation 6 with a different class partition. Our class partitions of K new-class detectors ensure each base class is considered as a simulated new class for at least one new-class detector. We denote $\mathcal{M}_D^i(\mathbf{x})$ as the new-class score computed for \mathbf{x} . Lower scores indicate a higher likelihood that \mathbf{x} belongs to new classes.

In addition, a threshold remains crucial for the detection of new classes, even when well-trained new-class detectors are provided. Leveraging the benefits of our partition and ensemble strategy, we can directly estimate the threshold for each new-class detector during training using the Otsu algorithm (Otsu, 1979; Liu & Yu, 2009) and training data. This is possible due to the presence of naturally simulated base classes and new classes in the training data for each new-class detector. Subsequently, these estimated thresholds can

be averaged to determine the threshold value, denoted as τ , for testing.

3.2. Sub-Classifier \mathcal{M}_C

After training reliable new-class detectors, we proceed to train a sub-classifier for each detector, as each detector focuses on a specific subset of the base class space. Each of the K sub-classifiers, denoted as $\{\mathcal{M}_C^i\}_{i=1}^K$, is designed to specialize in a particular base class space, thereby achieving better discriminability for the corresponding subset class space. For the i -th sub-classifier \mathcal{M}_C^i , we first utilize the trained new-class detector \mathcal{M}_D^i partition the training data into two subsets: \mathcal{D}_b^i and \mathcal{D}_n^i . Here, $\mathcal{D}_b^i = \{(\mathbf{x}, y) | (\mathbf{x}, y) \sim \mathcal{D} \wedge \mathcal{M}_D^i(\mathbf{x}) \geq \tau\}$ and $\mathcal{D}_n^i = \{(\mathbf{x}, y) | (\mathbf{x}, y) \sim \mathcal{D} \wedge \mathcal{M}_D^i(\mathbf{x}) < \tau\}$. Subsequently, we optimize the following objective function:

$$\ell_{CLS} = \sum_{(\mathbf{x}, y) \sim \mathcal{D}_b^i} \ell_{CE}(\mathbf{x}, y) + \sum_{(\mathbf{x}, y) \sim \mathcal{D}_n^i} \ell_{KL}(P(\mathbf{x}), P_{ZS}(\mathbf{x})) \quad (7)$$

Here, ℓ_{KL} denotes KL-divergence loss, and $P(\mathbf{x})$ and $P_{ZS}(\mathbf{x})$ represent the prediction probabilities of DECoOP approach and zero-shot CLIP baseline. We denote $\mathcal{M}_C^i(\mathbf{x})$ as the prediction probabilities computed for \mathbf{x} .

3.3. Inference

During testing, we evaluate an ensemble of K new-class detectors $\{\mathcal{M}_D^i\}_{i=1}^K$ to determine whether each testing data should be predicted by one of the learned sub-classifiers \mathcal{M}_C^i or the zero-shot CLIP baseline. Specifically, for a testing instance \mathbf{x} , we first compute the scores of the new-

Table 1. Comparison of average performance across 11 datasets was conducted among three approaches: ZS, PT, and our DEPT framework, utilizing ViT-B/16 and ViT-B/32 architectures. These results are consistent with our theoretical analysis.

METHOD	ViT-B/16		ViT-B/32	
	NEW ACC.	ACCURACY	NEW ACC.	ACCURACY
ZS	65.49	63.92	63.95	60.36
PT	57.73	65.57	53.01	61.03
DEPT	68.15	68.03	65.45	62.92

class detectors, $\{\mathcal{M}_D^i(\mathbf{x})\}_{i=1}^K$, and then make the prediction according to our DECoOP approach, defined as:

$$P_{\text{DECoOP}}(\mathbf{x}) = \begin{cases} P_{\text{ZS}}(\mathbf{x}), & \text{if } \max_{i \in \{1, \dots, K\}} \mathcal{M}_D^i(\mathbf{x}) < \tau, \\ \mathcal{M}_C^{i^*}(\mathbf{x}), & \text{if } \max_{i \in \{1, \dots, K\}} \mathcal{M}_D^i(\mathbf{x}) \geq \tau, \end{cases} \quad (8)$$

where $i^* = \arg \max_{i \in \{1, \dots, K\}} \mathcal{M}_D^i(\mathbf{x})$. DECoOP approach selects single sub-classifier to predict each testing data instead of aggregating the results from all sub-classifiers. As a result, our approach requires K times computation for the new-class detectors compared to the zero-shot CLIP baseline. In our experiments, we set K to 3, which does not impose a heavy computational burden. We conduct experiments about evaluation time in Appendix B.7, demonstrating that DECoOP is relatively efficient.

4. Experiments

In this section, we conduct experiments to answer the following three research questions:

RQ1: Can the empirical results of the DEPT framework on real-world datasets conform to our theoretical analysis?

RQ2: Can the DECoOP method surpass existing baseline and SOTA methods, thereby demonstrating its robustness?

RQ3: Does the DECoOP successfully improve the base-to-new discriminability, as designed?

4.1. Experimental Setup

Evaluation Protocol. We adopt the few-shot prompt tuning setting as previously explored in studies such as (Radford et al., 2021; Zhou et al., 2022a; Wang et al., 2023b). This setting involves partitioning the class space of each dataset equally, with 50% of the classes designated as base classes and the remaining 50% as new classes. Consequently, for each dataset, prompts are learned for downstream tasks using 16 labeled samples per base class, drawn from the training set. The efficacy of these learned prompts

Table 2. The average performance across 11 datasets using ViT-B/16 and ViT-B/32 architectures. The best performance is in bold.

METHOD	ViT-B/16		ViT-B/32	
	H	ACCURACY	H	ACCURACY
CLIP	70.84	63.92	67.13	60.36
PROMPT ENS.	71.65	65.39	67.76	60.73
CoOp	72.14	65.57	67.86	61.03
CoCoOp	74.72	67.67	70.77	62.96
SHIP	72.26	64.51	69.25	59.91
DECoOP(OURS)	76.13	69.69	72.51	65.75

is subsequently evaluated on the entire testing set, encompassing both base and new classes. In DECoOP method, we report the Accuracy as well as previously reported H metric. As per the definition in CoCoOp (Zhou et al., 2022a), H metric separately evaluates the accuracy on base classes and new classes, denoted as Acc_{base} and Acc_{new} . Then, H metric is computed using their harmonic mean, defined as $H = \frac{2 \times \text{Acc}_{\text{base}} \times \text{Acc}_{\text{new}}}{\text{Acc}_{\text{base}} + \text{Acc}_{\text{new}}}$. The metric H evaluates the overall performance of classifying both base and new classes separately, which we refer to as base-class discriminability and new-class discriminability. We evaluate the accuracy of the entire class space, which includes a mix of base and new classes, denoted as Accuracy. This metric evaluates the overall performance of classifying both base and new classes, while additionally measuring base-to-new discriminability compared to the H metric.

Datasets. Following the CoOp framework (Zhou et al., 2022b), we conducted evaluations of our proposed DECoOP framework along with comparison methods on various image classification tasks. These tasks included general object recognition using ImageNet (Deng et al., 2009) and Caltech-101 (Fei-Fei et al., 2007) datasets, fine-grained object recognition involving datasets such as Oxford Pets (Krause et al., 2013), Food-101 (Bossard et al., 2014), Stanford Cars (Krause et al., 2013), Oxford Flowers 102 (Nilsback & Zisserman, 2008), and FGVC Aircraft (Maji et al., 2013). Additionally, we performed a remote sensing recognition task using the EuroSAT (Helber et al., 2019) dataset, a texture recognition task using the DTD (Cimpoi et al., 2014) dataset, an action recognition task using UCF101 (Soomro et al., 2012) dataset and a large-scale scene understanding task using SUN397 (Xiao et al., 2010) dataset. For each dataset, we developed a few-shot training set for prompt tuning and employed the full testing set to evaluate the effectiveness of the learned prompts.

Compared Methods. We compare our approach with five existing prompt-based methods. CLIP (Radford et al., 2021) uses a hand-crafted prompt to generate the target classifier on the downstream task. Furthermore, we compare the

Table 3. Performance comparison on 11 datasets using ViT-B/16 architecture. The best performance is in bold.

	AVERAGE		IMAGENET		CALTECH101		OXFORDPETS	
	H	ACC.	H	ACC.	H	ACC.	H	ACC.
CLIP	70.84	63.92	70.20 \pm 0.00	66.73 \pm 0.00	95.41 \pm 0.00	92.90 \pm 0.00	92.93 \pm 0.00	88.03 \pm 0.00
PROMPT ENS.	71.65	65.39	72.00 \pm 0.00	68.48 \pm 0.00	96.20 \pm 0.00	94.08 \pm 0.00	92.42 \pm 0.00	86.37 \pm 0.00
CoOp	72.14	65.57	64.95 \pm 1.11	61.79 \pm 1.09	95.96 \pm 0.39	93.24 \pm 0.68	95.38 \pm 0.33	89.61 \pm 0.34
CoCoOp	74.72	67.67	72.71 \pm 0.33	69.41 \pm 0.36	95.55 \pm 0.24	93.43 \pm 0.37	95.71 \pm 0.76	90.24 \pm 1.32
SHIP	72.26	64.51	67.29 \pm 0.38	63.65 \pm 0.32	95.83 \pm 0.23	92.93 \pm 0.37	94.44 \pm 0.54	86.78 \pm 1.32
DECoOP(OURS)	76.13	69.69	72.98 \pm 0.04	69.62 \pm 0.08	96.52 \pm 0.09	94.50 \pm 0.22	95.27 \pm 0.08	88.87 \pm 0.28
	STANDFORDCARS		FLOWERS102		FOOD101		FGVCAIRCRAFT	
	H	ACC.	H	ACC.	H	ACC.	H	ACC.
CLIP	68.75 \pm 0.00	65.39 \pm 0.00	72.74 \pm 0.00	67.28 \pm 0.00	90.18 \pm 0.00	85.40 \pm 0.00	30.25 \pm 0.00	23.94 \pm 0.00
PROMPT ENS.	69.36 \pm 0.00	65.95 \pm 0.00	72.14 \pm 0.00	67.03 \pm 0.00	90.32 \pm 0.00	85.54 \pm 0.00	29.42 \pm 0.00	23.31 \pm 0.00
CoOp	68.22 \pm 0.49	63.81 \pm 0.44	78.33 \pm 2.26	72.11 \pm 2.36	86.65 \pm 1.38	80.84 \pm 1.50	29.38 \pm 1.78	24.80 \pm 1.23
CoCoOp	71.49 \pm 0.62	67.75 \pm 0.68	80.04 \pm 1.46	71.95 \pm 1.24	90.41 \pm 0.24	85.61 \pm 0.43	27.87 \pm 11.36	21.46 \pm 7.42
SHIP	69.71 \pm 0.43	64.67 \pm 0.55	76.85 \pm 2.18	70.40 \pm 2.01	86.84 \pm 1.49	77.39 \pm 2.19	27.13 \pm 1.10	24.44 \pm 0.96
DECoOP(OURS)	73.24 \pm 0.15	69.64 \pm 0.19	84.16 \pm 0.27	78.61 \pm 0.59	90.68 \pm 0.09	85.83 \pm 0.07	31.44 \pm 0.39	25.15 \pm 0.31
	SUN397		DTD		EUROSAT		UCF101	
	H	ACC.	H	ACC.	H	ACC.	H	ACC.
CLIP	72.26 \pm 0.00	62.57 \pm 0.00	57.32 \pm 0.00	44.56 \pm 0.00	58.16 \pm 0.00	41.40 \pm 0.00	71.00 \pm 0.00	64.97 \pm 0.00
PROMPT ENS.	75.04 \pm 0.00	65.97 \pm 0.00	59.63 \pm 0.00	46.28 \pm 0.00	58.45 \pm 0.00	48.91 \pm 0.00	73.17 \pm 0.00	67.33 \pm 0.00
CoOp	71.37 \pm 1.21	61.82 \pm 1.11	57.22 \pm 2.37	48.18 \pm 1.78	74.33 \pm 4.35	59.65 \pm 5.07	71.68 \pm 2.84	65.41 \pm 2.18
CoCoOp	77.17 \pm 0.27	68.17 \pm 0.33	60.59 \pm 1.51	47.90 \pm 1.43	73.77 \pm 3.58	58.08 \pm 1.49	76.59 \pm 0.79	70.39 \pm 1.25
SHIP	72.57 \pm 0.38	60.42 \pm 0.48	56.82 \pm 2.18	47.58 \pm 1.62	73.29 \pm 2.67	54.11 \pm 1.73	74.09 \pm 2.09	67.24 \pm 1.94
DECoOP(OURS)	78.11 \pm 0.09	69.33 \pm 0.05	62.72 \pm 1.23	51.44 \pm 1.04	74.61 \pm 3.82	61.90 \pm 3.72	77.67 \pm 0.50	71.71 \pm 0.79

PROMPT ENS. method, an ensemble technique that utilizes multiple classifiers to enhance the performance of CLIP, adhering to the guidelines set by CLIP. COOP (Zhou et al., 2022b) learns a soft prompt by minimizing the classification loss, and CoCoOp (Zhou et al., 2022b) extends COOP by further learning a lightweight neural network to generate for each image an input-conditional token. SHIP (Wang et al., 2023b) follows variational autoencoders to introduce a generator that reconstructs the visual features by inputting the synthesized prompts and the corresponding class names to the textual encoder of CLIP.

Implementation Details. The number of tokens in each prompt is set to 16 for DECoOP approach and comparison methods. We train the prompts of new-class detectors for 50 epochs using the SGD optimizer and subsequently train the prompts for sub-classifiers for 100 epochs, also using the SGD optimizer. The learning rate lr is set to 0.002, and it follows a cosine decay schedule. The margin γ is set to 0.4 for all datasets. We use the PROMPT ENS. method as our zero-shot baseline within the DECoOP approach. The batch size for images is 32 across all datasets. All experiments were conducted on Linux servers equipped with NVIDIA A800 GPUs. We report the average results over 5 runs with different random seed $\{1, 2, 3, 4, 5\}$.

4.2. Empirical Results

RQ1: Can the empirical results of the DEPT framework on real-world datasets conform to our theoretical analysis?

To verify Theorem 2.1, we conducted experiments on 11 datasets using ViT-B/16 and ViT-B/32 architectures. We employed CLIP as the zero-shot baseline Zs and COOP as the prompt tuning method PT. Subsequently, we constructed our DEPT framework by integrating these two methods, as presented in Equation 3. Here, the OOD detector used in our DEPT framework directly derives from CLIP using MSP method (Hendrycks & Gimpel, 2016). Each method is evaluated on the entire class space \mathcal{Y} , and the average performance across all datasets is reported. The results include New Acc. and Accuracy, indicating the average performance of new classes and all classes, respectively. The results presented in Table 1 consistently demonstrate that our DEPT framework outperforms both Zs and PT methods when evaluated using the New Acc. and Accuracy metrics. This observation suggests that the DEPT framework effectively mitigates performance degradation on new classes through the utilization of the OOD detector, which aligns well with our theoretical analysis.

RQ2: Can the DECoOP method surpass existing baseline and SOTA methods, thereby demonstrating its robustness?

To assess the effectiveness of the DECoOP approach, we conducted experiments on 11 datasets using ViT-B/16 and ViT-B/32 architectures. The average performance across all datasets, as well as the detailed performance on each dataset measured by two metrics, i.e., H and Accuracy, is reported. The results obtained using the ViT-B/16 architecture are presented in Table 3. Our DECoOP approach demonstrates superior average performance on both the H metric and

Table 4. The base-to-new discriminability of each method evaluated using MSP method (Hendrycks & Gimpel, 2016) and AUROC metrics. The best performance is in bold.

DATASET	CLIP	CoCoOp	SHIP	DECoOP(OURS)
IMAGENET	88.34	88.05	84.71	97.48
CALTECH101	97.03	95.71	96.94	99.58
OXFORDPETS	92.66	91.15	93.30	98.12
STANFORDCARS	86.24	83.00	87.23	97.63
FLOWERS102	84.92	79.63	84.84	95.75
FOOD101	89.88	88.19	89.92	97.59
FGVCAIRCRAFT	75.08	69.00	75.78	84.06
SUN397	72.46	73.75	74.78	90.21
DTD	62.29	60.65	60.66	75.47
EUROSAT	56.40	57.74	59.32	77.78
UCF101	82.03	79.03	80.35	93.56
AVERAGE	80.67	78.72	80.71	91.57




Figure 6. The ROC curve for detecting new classes of each method on Flowers102 and StandfordCars datasets.

Accuracy, showcasing its robustness. Regarding the detailed performance on each dataset, our approach outperforms the comparison methods on 10 out of 11 datasets, while achieving comparable performance on the remaining dataset. The detailed results using the ViT-B/32 architecture are provided in Appendix B.1, which yield similar conclusions.

Furthermore, these results reveal a positive correlation between the H metric and Accuracy in most cases. However, specific datasets such as FGVCAircraft (Maji et al., 2013) show that higher H metric values do not necessarily lead to improved Accuracy. This observation suggests that the H metric is inadequate for measuring base-to-new discriminability, emphasizing the significance of OPT problem.

RQ3: Does the DECoOP successfully improve the base-to-new discriminability, as designed?

The DECoOP approach introduces new-class detectors with the aim of improving base-to-new discriminability while simultaneously enhancing the discriminability of new classes. We evaluate the base-to-new discriminability of our approach and selected methods using the MSP (Hendrycks & Gimpel, 2016) method with the ViT-B/16 architecture. Specifically, for each method, we use the maximum probability on base classes as the scores and report the AUROC (Bradley, 1997) in Table 4. The results clearly indicate

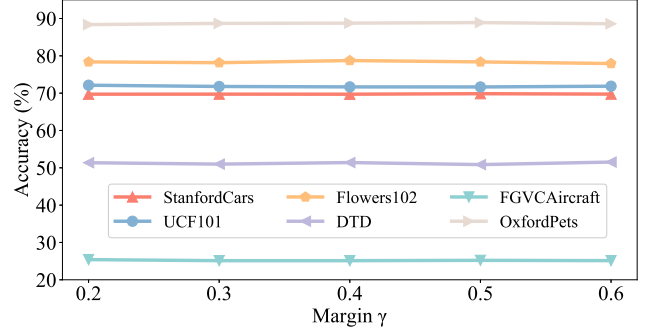


Figure 7. Performance using different values of margin γ .

that our DECoOP approach significantly improves base-to-new discriminability, which accounts for its SOTA performance. We have omitted some methods and standard deviations due to space limitations. Please refer to Appendix B.2 for full results. Additionally, we present the ROC curves for two representative datasets in Figure 6, which demonstrates similar findings. Due to space limitations, the ROC curves for all datasets are provided in Appendix B.3. Furthermore, we explore the correlation between the performance of new-class detectors and the model in Appendix B.4.

Hyperparameter. The margin γ serves as a hyperparameter for learning new-class detectors in our DECoOP approach. It controls the margin in the optimization process of the detectors, which may affect their performance. To answer the robustness question of γ , we conduct experiments on six datasets. Figure 7 demonstrates the robustness of the DECoOP approach to changes in γ .

Comparison with Ensembling of CoOp. In Appendix B.6, we conduct an experiment to determine if directly combining multiple CoOp prompts can lead to performance improvement. The results demonstrate that combining 2, 4, or 6 CoOp prompts does not effectively enhance performance and, at times, even deteriorates the performance. This indicates that our performance gains cannot be attributed to simple prompt ensembling.

Running Time. In Appendix B.7, we conduct an experiment to compare the CoOp, CoCoOp, and DECoOP methods as shown in Table 9. On the EuroSAT dataset, the runtime of DECoOP increased only slightly compared to CoOp (14.1s vs. 34.1s), but it is significantly more efficient than CoCoOp (62.0s), demonstrating the efficiency of the DECoOP algorithm.

5. Related Work

Few-shot Prompt Tuning. Prompt learning aims to formalize various NLP tasks to mask language modeling problems, which is similar to the pre-training of language mod-

els (Devlin et al., 2018; Radford et al., 2019; 2021) by adopting different prompt templates. The previous works (Petroni et al., 2019; Radford et al., 2019; Brown et al., 2020) elaborately design human-crafted prompts, which is known as prompt engineering. Despite considerable progress in NLP, prompt learning remains underexplored in computer vision. Pretrained VLMs (Jia et al., 2021; Radford et al., 2021) introduce hand-crafted prompts to perform zero-shot inference on the downstream tasks. However, designing specific prompts for various downstream tasks is inefficient and costly and several studies (Shi et al., 2024b;c) performs parameter-efficient fine-tuning to address this problem. CoOp (Zhou et al., 2022b) makes prompt learnable via minimizing the classification loss on the target task, adopting the prompt tuning approach of NLP. However, CoOp decreases the zero-shot capability of VLMs. To fix the problem, CoCoOp (Zhou et al., 2022a) introduces meta net to conditionally fine-tune prompts. LFA (Ouali et al., 2023) adopts a simple linear approach for vision-language alignment. VPT (Derakhshani et al., 2022) attempts to learn a collection of continuous prompts to capture the variational visual representation. SHIP (Wang et al., 2023b) follows the paradigm of variational autoencoders to generate visual features according to the prompts via the generative method. ProDA (Lu et al., 2022) proposes to learn the distribution of instance-specific prompts via variational inference. Ding et al. (2024) explores the integration of OOD detection methods for VLMs and present meaningful observations. However, these studies do not consider the OPT evaluation setting. Recent studies (Zhang et al., 2024; Shu et al., 2022a) also make the attempts to perform prompt tuning on changing datastreams in a test-time adaptation manner (Zhou et al., 2023; 2024; Zhao et al., 2024). These studies can be explored to address OPT problem in the future studies.

Out-of-distribution Detection. Out-of-distribution detection refers to training the model on in-distribution (ID) dataset to classify OOD and ID samples. MSP (Hendrycks & Gimpel, 2016) takes the maximum softmax probability over ID categories as the score. RotPred (Hendrycks et al., 2019) includes an extra head to predict the rotation angle of rotated inputs in a self-supervised manner, and the rotation head together with the classification head is used for OOD detection. MCD (Yu & Aizawa, 2019) considers an ensemble of multiple classification heads and promotes the disagreement between each head’s prediction on OOD samples. StyleAugment (Geirhos et al., 2018) applies style transfer to clean images to emphasize the shape bias over the texture bias. STEP (Zhou et al., 2021) focuses on exploring out-of-distribution detection within a semi-supervised setting (Tong et al., 2022; Lan-Zhe & Yu-Feng, 2024; Shi et al., 2024a; Jia et al., 2024). CIDER (Ming et al., 2022) regularizes the model’s hyperspherical space by increasing inter-class separability and intra-class compactness. MixOE (Zhang et al.,

2023) performs pixel-level mixing operations between ID and OOD samples and regularizes the model such that the prediction confidence smoothly decays as the input transitions from ID to OOD. RegMixup (Pinto et al., 2022) trains the model with both clean images and mixed images obtained from the convex combination. Recent studies (Ming & Li, 2024; Sun et al., 2024), such as Clipn (Wang et al., 2023a), LoCoOp (Miyai et al., 2023), attempt to explore the capability of zero-shot and few-shot ood detection via VLMs respectively. However, while these studies primarily focus on OOD detection tasks, our research utilizes OOD detection to enhance the generalization of VLMs.

6. Conclusion

In this paper, we explore the OPT problem in detail and uncover that the base-to-new discriminability is crucial but often overlooked by existing methods and settings. We first introduce the DEPT framework and demonstrate, through theoretical analysis, that incorporating OOD detection into prompt tuning can enhance the base-to-new discriminability and prevent degradation of new-class discriminability. Building upon DEPT, we propose a novel prompt tuning approach called DECoOP that introduces new-class detectors and sub-classifiers to further enhance the discriminability of both the base and new classes. Experimental results validate our analysis of DEPT and demonstrate the effectiveness of our DECoOP approach.

One limitation of our work is that we only take the initial step in combining OOD detection and prompt tuning. We believe there is potential for future researchers to explore. The other limitations are included in Appendix C.

Code Availability Statement

The implementation code for this work is available at <https://github.com/WNJXYK/DeCoOp>.

Acknowledgements

This research was supported by National Science and Technology Major Project (2022ZD0114803) and the National Science Foundation of China (62306133, 62176118).

Impact Statement

This paper aims to advance prompt tuning for vision-language models. The work carried out in this study has various potential societal implications. We firmly believe that the majority of these impacts are positive and it is unnecessary to explicitly emphasize any specific ones in this paper. Additionally, we anticipate that the responsible utilization of technology will foster discourse concerning the best practices and regulations for implementing methods.

References

- Bossard, L., Guillaumin, M., and Gool, L. V. Food-101 - mining discriminative components with random forests. In *Proceedings of the 13th European Conference on Computer Vision*, pp. 446–461, 2014.
- Bradley, A. P. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, pp. 1145–1159, 1997.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. In *Advances in neural information processing systems*, pp. 1877–1901, 2020.
- Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3606–3613, 2014.
- Deng, J., Dong, W., Socher, R., Li, L., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- Derakhshani, M. M., Sanchez, E., Bulat, A., da Costa, V. G. T., Snoek, C. G., Tzimiropoulos, G., and Martinez, B. Variational prompt tuning improves generalization of vision-language models. *arXiv preprint arXiv:2210.02390*, 2022.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Ding, K., Zhang, H., Yu, Q., Wang, Y., Xiang, S., and Pan, C. Weak distribution detectors lead to stronger generalizability of vision-language prompt tuning. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, pp. 1528–1536, 2024.
- Fei-Fei, L., Fergus, R., and Perona, P. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Computer Vision and Image Understanding*, pp. 59–70, 2007.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., and Brendel, W. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*, 2018.
- Hayes, T. L., Krishnan, G. P., Bazhenov, M., Siegelmann, H. T., Sejnowski, T. J., and Kanan, C. Replay in deep learning: Current approaches and missing biological elements. *Neural Computation*, pp. 2908–2950, 2021.
- Helber, P., Bischke, B., Dengel, A., and Borth, D. 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*, pp. 2217–2226, 2019.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- Hendrycks, D., Mazeika, M., Kadavath, S., and Song, D. Using self-supervised learning can improve model robustness and uncertainty. In *Advances in neural information processing systems*, volume 32, 2019.
- Ilharco, G., Wortsman, M., Gadre, S. Y., Song, S., Hajishirzi, H., Kornblith, S., Farhadi, A., and Schmidt, L. Patching open-vocabulary models by interpolating weights. In *Advances in Neural Information Processing Systems*, 2022.
- Jia, C., Yang, Y., Xia, Y., Chen, Y.-T., Parekh, Z., Pham, H., Le, Q., Sung, Y.-H., Li, Z., and Duerig, T. Scaling up visual and vision-language representation learning with noisy text supervision. In *Proceedings of International conference on machine learning*, pp. 4904–4916. PMLR, 2021.
- Jia, L.-H., Guo, L.-Z., Zhou, Z., and Li, Y.-F. Lamda-ssl: a comprehensive semi-supervised learning toolkit. *Science China Information Sciences*, 67(117101), 2024.
- Krause, J., Stark, M., Deng, J., and Fei-Fei, L. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 554–561, 2013.
- Kudithipudi, D., Aguilar-Simon, M., Babb, J., Bazhenov, M., Blackiston, D., Bongard, J. C., Brna, A. P., Raja, S. C., Cheney, N., Clune, J., Daram, A. R., Fusi, S., Helfer, P., Kay, L., Ketz, N., Kira, Z., Kolouri, S., Krichmar, J. L., Kriegman, S., Levin, M., Madireddy, S., Manicka, S., Marjaninejad, A., McNaughton, B., Miikkulainen, R., Navratilova, Z., Pandit, T., Parker, A., Pilly, P. K., Risi, S., Sejnowski, T. J., Soltoggio, A., Soures, N., Tolia, A. S., Urbina-Meléndez, D., Cuevas, F. J. V., van de Ven, G. M., Vogelstein, J. T., Wang, F., Weiss, R., Yanguas-Gil, A., Zou, X., and Siegelmann, H. T. Biological underpinnings for lifelong learning machines. *Nature Machine Intelligence*, pp. 196–210, 2022.
- Lan-Zhe, G. and Yu-Feng, L. Robust pseudo-label selection for holistic semi-supervised learning. *Science China Information Sciences*, 53(3):623–637, 2024.

- Lange, M. D., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G. G., and Tuytelaars, T. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 3366–3385, 2022.
- Liu, D. and Yu, J. Otsu method and k-means. In *In Proceedings of the 9th International Conference on Hybrid Intelligent Systems*, pp. 344–349, 2009.
- Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., and Neubig, G. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, pp. 1–35, 2023.
- Lu, Y., Liu, J., Zhang, Y., Liu, Y., and Tian, X. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5206–5215, 2022.
- Mai, Z., Li, R., Jeong, J., Quispe, D., Kim, H., and Sanner, S. Online continual learning in image classification: An empirical survey. *Neurocomputing*, 469:28–51, 2022.
- Maji, S., Rahtu, E., Kannala, J., Blaschko, M., and Vedaldi, A. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- Ming, Y. and Li, Y. How does fine-tuning impact out-of-distribution detection for vision-language models? *International Journal of Computer Vision*, 132(2):596–609, 2024.
- Ming, Y., Sun, Y., Dia, O., and Li, Y. How to exploit hyperspherical embeddings for out-of-distribution detection? *arXiv preprint arXiv:2203.04450*, 2022.
- Miyai, A., Yu, Q., Irie, G., and Aizawa, K. Locoop: Few-shot out-of-distribution detection via prompt learning. *arXiv preprint arXiv:2306.01293*, 2023.
- Nilsback, M. and Zisserman, A. Automated flower classification over a large number of classes. In *Proceedings of the 6th Indian Conference on Computer Vision*, pp. 722–729, 2008.
- Otsu, N. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979.
- Ouali, Y., Bulat, A., Martínez, B., and Tzimiropoulos, G. Black box few-shot adaptation for vision-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15488–15500, 2023.
- Petroni, F., Rocktäschel, T., Lewis, P., Bakhtin, A., Wu, Y., Miller, A. H., and Riedel, S. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*, 2019.
- Pinto, F., Yang, H., Lim, S. N., Torr, P., and Dokania, P. Using mixup as a regularizer can surprisingly improve accuracy & out-of-distribution robustness. In *Advances in Neural Information Processing Systems*, pp. 14608–14622, 2022.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, pp. 9, 2019.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763, 2021.
- Shi, J.-X., Wei, T., and Li, Y.-F. Residual diverse ensemble for long-tailed multi-label text classification. *Science CHINA Information Science*, 2024a.
- Shi, J.-X., Wei, T., Zhou, Z., Shao, J.-J., Han, X.-Y., and Li, Y.-F. Long-tail learning with foundation model: Heavy fine-tuning hurts. In *Proceedings of the 41st International Conference on Machine Learning*, 2024b.
- Shi, J.-X., Zhang, C., Wei, T., and Li, Y.-F. Efficient and long-tailed generalization for pre-trained vision-language model. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024c.
- Shu, M., Nie, W., Huang, D., Yu, Z., Goldstein, T., Anandkumar, A., and Xiao, C. Test-time prompt tuning for zero-shot generalization in vision-language models. In *Advances in Neural Information Processing Systems*, 2022a.
- Shu, M., Nie, W., Huang, D.-A., Yu, Z., Goldstein, T., Anandkumar, A., and Xiao, C. Test-time prompt tuning for zero-shot generalization in vision-language models. In *Advances in Neural Information Processing Systems*, pp. 14274–14289, 2022b.
- Soomro, K., Zamir, A. R., and Shah, M. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012.
- Sun, H., He, R., Han, Z., Lin, Z., Gong, Y., and Yin, Y. Clip-driven outliers synthesis for few-shot OOD detection. *CoRR*, abs/2404.00323, 2024.
- Tong, W., Hai, W., Weiwei, T., and Yufeng, L. Robust model selection for positive and unlabeled learning with constraints. *Science China Information Sciences*, 65(212101), 2022.
- Wang, H., Li, Y., Yao, H., and Li, X. Clipn for zero-shot ood detection: Teaching clip to say no. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1802–1812, 2023a.

- Wang, Z., Liang, J., He, R., Xu, N., Wang, Z., and Tan, T. Improving zero-shot generalization for clip with synthesized prompts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3032–3042, 2023b.
- Wortsman, M., Ilharco, G., Kim, J. W., Li, M., Kornblith, S., Roelofs, R., Lopes, R. G., Hajishirzi, H., Farhadi, A., Namkoong, H., and Schmidt, L. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7949–7961, 2022.
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. SUN database: Large-scale scene recognition from abbey to zoo. In *Proceedings of the 2010 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3485–3492, 2010.
- Yang, X., Shao, J., Tu, W., Li, Y., Dai, W., and Zhou, Z. Safe abductive learning in the presence of inaccurate rules. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, pp. 16361–16369, 2024a.
- Yang, X., Wei, W., Shao, J., Li, Y., and Zhou, Z. Analysis for abductive learning and neural-symbolic reasoning shortcuts. In *Proceedings of the 41st International Conference on Machine Learning*, 2024b.
- Yu, Q. and Aizawa, K. Unsupervised out-of-distribution detection by maximum classifier discrepancy. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9518–9526, 2019.
- Zhang, D., Zhou, Z., and Li, Y. Robust test-time adaptation for zero-shot prompt tuning. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, pp. 16714–16722, 2024.
- Zhang, J., Inkawhich, N., Linderman, R., Chen, Y., and Li, H. Mixture outlier exposure: Towards out-of-distribution detection in fine-grained environments. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 5531–5540, 2023.
- Zhao, P., Zhang, Y.-J., Zhang, L., and Zhou, Z.-H. Adaptivity and non-stationarity: Problem-dependent dynamic regret for online convex optimization. *Journal of Machine Learning Research*, 25(98):1 – 52, 2024.
- Zhou, K., Yang, J., Loy, C. C., and Liu, Z. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16795–16804, 2022a.
- Zhou, K., Yang, J., Loy, C. C., and Liu, Z. Learning to prompt for vision-language models. *International Journal of Computer Vision*, pp. 2337–2348, 2022b.
- Zhou, Z., Guo, L., Cheng, Z., Li, Y., and Pu, S. STEP: out-of-distribution detection in the presence of limited in-distribution labeled data. In *Advances in Neural Information Processing Systems*, pp. 29168–29180, 2021.
- Zhou, Z., Guo, L., Jia, L., Zhang, D., and Li, Y. ODS: test-time adaptation in the presence of open-world data shift. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 42574–42588, 2023.
- Zhou, Z., Zhang, D.-C., Li, Y.-F., and Zhang, M.-L. Towards robust test-time adaptation method for open-set recognition. *Journal of Software*, 35(4):1667–1681, 2024.

A. Proof of Theorem 2.1

Proof. We first compute $H_{ZS}^{CLS}(\mathbf{x})$ and $H_{ZS}^{OOD}(\mathbf{x})$ for one specific instance \mathbf{x} . Recall that for an instance \mathbf{x} , we denote its ground-truth label space as k (which always equals to $g(\mathbf{x})$) and its ground-truth label as $f(\mathbf{x})$. To facilitate the proof, we define additional label spaces:

$$\mathcal{Y}_{i,j} = \begin{cases} \{j\}, & j \in \mathcal{Y}_i, \\ \emptyset, & \text{otherwise}, \end{cases} \quad (9)$$

and additional class vectors for \mathbf{x} :

$$\tilde{y}_{i,j} = \begin{cases} 1, & f(\mathbf{x}) = j \wedge f(\mathbf{x}) \in \mathcal{Y}_i, \\ 0, & \text{otherwise}. \end{cases} \quad (10)$$

Our computational results are presented as follows:

$$\begin{aligned} H_{ZS}^{CLS}(\mathbf{x}) &= H(\tilde{\mathbf{y}}, \{P_{ZS}(y = j | y \in \mathcal{Y}_k, \mathbf{x})\}_{j=1}^C) \\ &= H(\tilde{\mathbf{y}}, \{P_{ZS}(y \in \mathcal{Y}_{k,j} | y \in \mathcal{Y}_k, \mathbf{x})\}_{j=1}^C) \\ &= - \sum_{j=1}^C \tilde{y}_j \log P_{ZS}(y = j | y \in \mathcal{Y}_k, \mathbf{x}) \\ &= - \log P_{ZS}(y = f(\mathbf{x}) | y \in \mathcal{Y}_k, \mathbf{x}), \end{aligned} \quad (11)$$

and

$$\begin{aligned} H_{ZS}^{OOD}(\mathbf{x}) &= H(\tilde{\mathbf{k}}, \{P_{ZS}(y \in \mathcal{Y}_i | \mathbf{x})\}_{i=\{b,n\}}) \\ &= - \sum_{i \in \{b,n\}} \tilde{k}_i \log P_{ZS}(y \in \mathcal{Y}_i | \mathbf{x}) \\ &= - \log P_{ZS}(y \in \mathcal{Y}_k | \mathbf{x}). \end{aligned} \quad (12)$$

Then, we can bound $\mathbb{E}_{\mathbf{x}} [H_{ZS}(\mathbf{x})]$ as follows:

$$\begin{aligned} \mathbb{E}_{\mathbf{x}} [H_{ZS}(\mathbf{x})] &= \mathbb{E}_{\mathbf{x}} [H(\tilde{\mathbf{y}}, \{P_{ZS}(y = j | \mathbf{x})\}_{j=1}^C)] \\ &= \mathbb{E}_{\mathbf{x}} [H(\tilde{\mathbf{y}}, \{P_{ZS}(y \in \mathcal{Y}_{k,j} | \mathbf{x})\}_{j=1}^C)] \\ &= \mathbb{E}_{\mathbf{x}} [-\log P_{ZS}(y \in \mathcal{Y}_{k,f(\mathbf{x})} | \mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x}} [-\log P_{ZS}(y \in \mathcal{Y}_{k,f(\mathbf{x})} | y \in \mathcal{Y}_k, \mathbf{x}) - \log P_{ZS}(y \in \mathcal{Y}_k | \mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x}} [H_{ZS}^{CLS}(\mathbf{x}) + H_{ZS}^{OOD}(\mathbf{x})] \\ &= \mathbb{E}_{\mathbf{x}} [H_{ZS}^{CLS}(\mathbf{x})] + \mathbb{E}_{\mathbf{x}} [H_{ZS}^{OOD}(\mathbf{x})] \\ &\leq \delta + \epsilon. \end{aligned} \quad (13)$$

Further, we can similarly compute $H_{PT}^{CLS}(\mathbf{x})$ as follows:

$$\begin{aligned} H_{PT}^{CLS}(\mathbf{x}) &= H(\tilde{\mathbf{y}}, \{P_{PT}(y = j | y \in \mathcal{Y}_k, \mathbf{x})\}_{j=1}^C) \\ &= H(\tilde{\mathbf{y}}, \{P_{PT}(y \in \mathcal{Y}_{k,j} | y \in \mathcal{Y}_k, \mathbf{x})\}_{j=1}^C) \\ &= - \sum_{j=1}^C \tilde{y}_{k,j} \log P_{PT}(y \in \mathcal{Y}_{k,j} | y \in \mathcal{Y}_k, \mathbf{x}) \\ &= - \log P_{PT}(y \in \mathcal{Y}_{k,f(\mathbf{x})} | y \in \mathcal{Y}_k, \mathbf{x}). \end{aligned} \quad (14)$$

Table 5. Performance comparison between our proposed DECoOP method and comparison methods on 11 datasets using ViT-B/32 architecture. The best performance is in bold.

	AVERAGE		IMAGENET		CALTECH101		OXFORDPETS	
	H	Acc.	H	Acc.	H	Acc.	H	Acc.
CLIP	67.13	60.36	65.69 \pm 0.00	62.05 \pm 0.00	93.78 \pm 0.00	91.08 \pm 0.00	91.30 \pm 0.00	85.01 \pm 0.00
PROMPT ENS.	67.76	60.73	66.91 \pm 0.00	63.22 \pm 0.00	94.06 \pm 0.00	91.20 \pm 0.00	89.73 \pm 0.00	83.18 \pm 0.00
CoOp	67.86	61.03	60.99 \pm 0.09	57.61 \pm 0.12	93.55 \pm 0.76	91.09 \pm 0.45	92.17 \pm 0.77	85.21 \pm 0.65
CoCoOp	70.77	62.96	67.74 \pm 1.23	64.06 \pm 1.39	93.78 \pm 0.92	91.01 \pm 0.87	94.05 \pm 0.56	87.84 \pm 0.89
SHIP	69.25	59.91	61.72 \pm 0.61	56.93 \pm 1.26	93.35 \pm 0.93	89.80 \pm 0.83	92.19 \pm 1.47	81.22 \pm 1.03
DECoOP(OURS)	72.51	65.75	68.07 \pm 0.06	64.49 \pm 0.04	95.56 \pm 0.22	93.36 \pm 0.48	93.13 \pm 0.50	86.25 \pm 0.96
	STANDFORDCARS		FLOWERS102		FOOD101		FGVCAIRCRAFT	
	H	Acc.	H	Acc.	H	Acc.	H	Acc.
CLIP	65.14 \pm 0.00	60.39 \pm 0.00	70.50 \pm 0.00	64.27 \pm 0.00	85.10 \pm 0.00	79.16 \pm 0.00	23.62 \pm 0.00	18.30 \pm 0.00
PROMPT ENS.	64.67 \pm 0.00	59.82 \pm 0.00	68.60 \pm 0.00	63.30 \pm 0.00	85.55 \pm 0.00	79.59 \pm 0.00	23.45 \pm 0.00	18.30 \pm 0.00
CoOp	62.33 \pm 1.21	56.95 \pm 1.37	71.13 \pm 1.95	65.25 \pm 1.43	81.55 \pm 0.91	74.32 \pm 1.17	23.15 \pm 1.71	18.88 \pm 0.85
CoCoOp	65.48 \pm 0.66	60.27 \pm 0.84	74.46 \pm 1.10	65.86 \pm 1.53	86.11 \pm 0.29	80.09 \pm 0.40	21.68 \pm 5.89	15.28 \pm 4.87
SHIP	64.38 \pm 0.81	56.22 \pm 1.00	70.41 \pm 1.72	62.41 \pm 1.88	81.76 \pm 0.90	72.14 \pm 1.43	19.34 \pm 2.64	19.00 \pm 0.98
DECoOP(OURS)	67.45 \pm 0.15	62.55 \pm 0.23	79.06 \pm 0.43	72.84 \pm 0.77	86.04 \pm 0.10	79.98 \pm 0.11	25.58 \pm 0.33	20.03 \pm 0.16
	SUN397		DTD		EUROSAT		UCF101	
	H	Acc.	H	Acc.	H	Acc.	H	Acc.
CLIP	71.35 \pm 0.00	61.99 \pm 0.00	53.60 \pm 0.00	42.85 \pm 0.00	50.81 \pm 0.00	38.17 \pm 0.00	67.56 \pm 0.00	60.67 \pm 0.00
PROMPT ENS.	73.27 \pm 0.00	63.74 \pm 0.00	53.81 \pm 0.00	43.44 \pm 0.00	56.90 \pm 0.00	40.75 \pm 0.00	68.39 \pm 0.00	61.49 \pm 0.00
CoOp	69.48 \pm 1.01	59.89 \pm 0.85	57.52 \pm 1.82	48.90 \pm 1.23	67.46 \pm 7.70	51.07 \pm 8.05	67.11 \pm 3.56	62.12 \pm 2.48
CoCoOp	75.51 \pm 0.37	65.96 \pm 0.45	59.57 \pm 2.21	47.08 \pm 1.30	66.98 \pm 8.67	49.19 \pm 5.78	73.17 \pm 1.24	65.98 \pm 1.06
SHIP	70.33 \pm 0.63	58.86 \pm 0.71	57.22 \pm 3.14	45.91 \pm 1.07	77.74 \pm 3.74	50.23 \pm 1.92	73.27 \pm 1.21	66.31 \pm 0.72
DECoOP(OURS)	75.87 \pm 0.14	66.59 \pm 0.19	60.61 \pm 0.48	50.39 \pm 0.40	72.35 \pm 2.42	58.93 \pm 2.62	73.87 \pm 0.36	67.83 \pm 0.81

Finally, we can bound $\mathbb{E}_{\mathbf{x}} [H_{\text{DEPT}}(\mathbf{x})]$ as follows:

$$\begin{aligned}
 \mathbb{E}_{\mathbf{x}} [H_{\text{DEPT}}(\mathbf{x})] &= \mathbb{E}_{\mathbf{x}} [H(\tilde{\mathbf{y}}, \{P_{\text{DEPT}}(y = j|\mathbf{x})\}_{j=1}^C)] \\
 &= \mathbb{E}_{\mathbf{x}} [H(\tilde{\mathbf{y}}, \{P_{\text{DEPT}}(y \in \mathcal{Y}_{k,i}|\mathbf{x})\}_{i=1}^C)] \\
 &= \mathbb{E}_{\mathbf{x}} [-\log P_{\text{DEPT}}(y \in \mathcal{Y}_{k,i}|\mathbf{x})] \\
 &= \mathbb{E}_{\mathbf{x} \wedge k=b} [-\log P_{\text{DEPT}}(y \in \mathcal{Y}_{k,i}|\mathbf{x})] + \mathbb{E}_{\mathbf{x} \wedge k=n} [-\log P_{\text{DEPT}}(y \in \mathcal{Y}_{k,i}|\mathbf{x})] \\
 &= \mathbb{E}_{\mathbf{x} \wedge k=b} [-\log P_{\text{PT}}(y \in \mathcal{Y}_{k,f(\mathbf{x})}|y \in \mathcal{Y}_k, \mathbf{x}) - \log P_{\text{ZS}}(y \in \mathcal{Y}_k|\mathbf{x})] \\
 &\quad + \mathbb{E}_{\mathbf{x} \wedge k=n} [-\log P_{\text{ZS}}(y \in \mathcal{Y}_{k,f(\mathbf{x})}|y \in \mathcal{Y}_k, \mathbf{x}) - \log P_{\text{ZS}}(y \in \mathcal{Y}_k|\mathbf{x})] \\
 &= \mathbb{E}_{\mathbf{x} \wedge k=b} [H_{\text{PT}}^{\text{CLS}}(\mathbf{x}) + H_{\text{ZS}}^{\text{OOD}}(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \wedge k=n} [H_{\text{ZS}}^{\text{CLS}}(\mathbf{x}) + H_{\text{ZS}}^{\text{OOD}}(\mathbf{x})] \\
 &\leq \alpha \cdot (\delta - \Delta + \epsilon) + (1 - \alpha) \cdot (\delta + \epsilon) \\
 &\leq \delta + \epsilon - \alpha \cdot \Delta.
 \end{aligned} \tag{15}$$

□

B. Additional Experimental Results

B.1. Detailed Results on ViT-B/32 Architecture

To address the consistent performance of our proposal, we conduct experiments and compare our proposed DECOOP method, baseline methods, and SOTA prompting tuning methods on 11 datasets using ViT-B/32 architectures. Each dataset is trained with random seeds from 1 to 5. In terms of detailed performance on each dataset, our proposed method outperforms the comparison methods on 9 out of 11 datasets, while achieving comparable performance on the remaining 2 datasets, showcasing its robustness to different pre-trained architectures.

B.2. Detailed AUROC Results

The full experimental results of Table 4 are presented in Table 6. Our DECOOP approach achieves the best base-to-new discriminability among all comparison methods.

Table 6. AUROC performance is compared with CLIP, Prompt Ensemble, CoOP, CoCoOP, SHIP and our proposed DECoOP. The results demonstrate that our proposal enhances base-to-new discriminability.

DATASET	CLIP	PROMPT ENS.	CoOP	CoCoOP	SHIP	DECoOP(OURS)
IMAGENET	88.34 \pm 0.00	89.79 \pm 0.00	77.14 \pm 1.62	88.05 \pm 1.22	84.71 \pm 1.62	97.48 \pm 0.03
CALTECH101	97.03 \pm 0.00	97.09 \pm 0.00	94.53 \pm 0.87	95.71 \pm 0.50	96.94 \pm 0.79	99.58 \pm 0.03
OXFORDPETS	92.66 \pm 0.00	92.21 \pm 0.00	91.06 \pm 1.00	91.15 \pm 0.95	93.30 \pm 1.23	98.12 \pm 0.24
STANFORDCARS	86.24 \pm 0.00	87.46 \pm 0.00	78.25 \pm 2.00	83.00 \pm 2.24	87.23 \pm 1.16	97.63 \pm 0.02
FLOWERS102	84.92 \pm 0.00	87.78 \pm 0.00	78.06 \pm 1.82	79.63 \pm 2.20	84.84 \pm 1.41	95.75 \pm 0.18
FOOD101	89.88 \pm 0.00	90.26 \pm 0.00	87.53 \pm 1.20	88.19 \pm 1.07	89.92 \pm 1.00	97.59 \pm 0.04
FGVCAIRCRAFT	75.08 \pm 0.00	75.86 \pm 0.00	75.25 \pm 1.36	69.00 \pm 7.91	75.78 \pm 1.65	84.06 \pm 0.26
SUN397	72.46 \pm 0.00	75.29 \pm 0.00	70.29 \pm 1.47	73.75 \pm 1.11	74.78 \pm 1.14	90.21 \pm 0.10
DTD	62.29 \pm 0.00	61.10 \pm 0.00	56.78 \pm 1.93	60.65 \pm 0.94	60.66 \pm 1.22	75.47 \pm 1.02
EUROSAT	56.40 \pm 0.00	57.74 \pm 0.00	52.26 \pm 8.68	57.74 \pm 2.49	59.32 \pm 6.31	77.78 \pm 3.85
UCF101	82.03 \pm 0.00	83.56 \pm 0.00	72.72 \pm 2.21	79.03 \pm 1.52	80.35 \pm 1.99	93.56 \pm 0.62
AVERAGE	80.67	81.65	75.81	78.72	80.71	91.57

Table 7. Ablation study. We report average performance across 11 datasets was conducted among baselines, CoOP, DEPT, DECoOP approaches, utilizing ViT-B/16 and ViT-B/32 architectures. The best performance is in bold. The second-best performance is underlined.

METHOD	ViT-B/16		ViT-B/32	
	H	ACCURACY	H	ACCURACY
CLIP	70.84	63.92	67.13	60.36
PROMPT ENS.	71.65	65.39	67.76	60.73
CoOP	72.14	65.57	67.86	61.03
DEPT	74.82	<u>68.03</u>	69.96	<u>62.92</u>
DECoOP	76.13	69.69	72.51	65.75

B.3. Detailed ROC Curves

To evaluate whether our proposal can improve the performance for detecting, we conduct experiments on 11 datasets using ViT-B/16 architecture. Each curve is drawn using our experiment results with random seeds to 1. For each method, we adopt the maximum softmax probability over new classes as the detecting score for drawing the curve. The results in Figure 8 show that our proposal can achieve the best detection performance.

B.4. Correlation between \mathcal{M}_O and Performance

The objective of the DECoOP approach is to enhance the base-to-new discriminability through the \mathcal{M}_O , leading to improved performance. Hence, a key question arises: does a better \mathcal{M}_O result in enhanced performance? To investigate this, we employ different new-class detectors with varying AUROC values for training and evaluate the performance as shown in Figure 9. This figure illustrates the correlation between the AUROC of the new-class detector and the performance metric. The results indicate a positive correlation between these two variables, validating our claim and aligning with our research objective.

B.5. Ablation Study

We conduct ablation studies to validate the effectiveness of each component of our proposed DECoOP approach in Table 7. In this paper, we first propose a novel prompt tuning framework DEPT to introduce OOD detection into prompt tuning. Then, two advanced modules are integrated into DEPT framework to form our DECoOP approach. As the our two modules cannot be separated to perform classification, we compare baseline methods, DEPT framework, and our proposed DECoOP approach. The results show that DEPT framework enhances the base-to-new discriminability and prevents performance degradation of new classes, thereby outperforming CLIP, PROMPT ENS., and CoOP methods. Further, our proposed DECoOP approach achieves the best performance among all methods, demonstrating it additionally enhances the base-class and new-class discriminability.

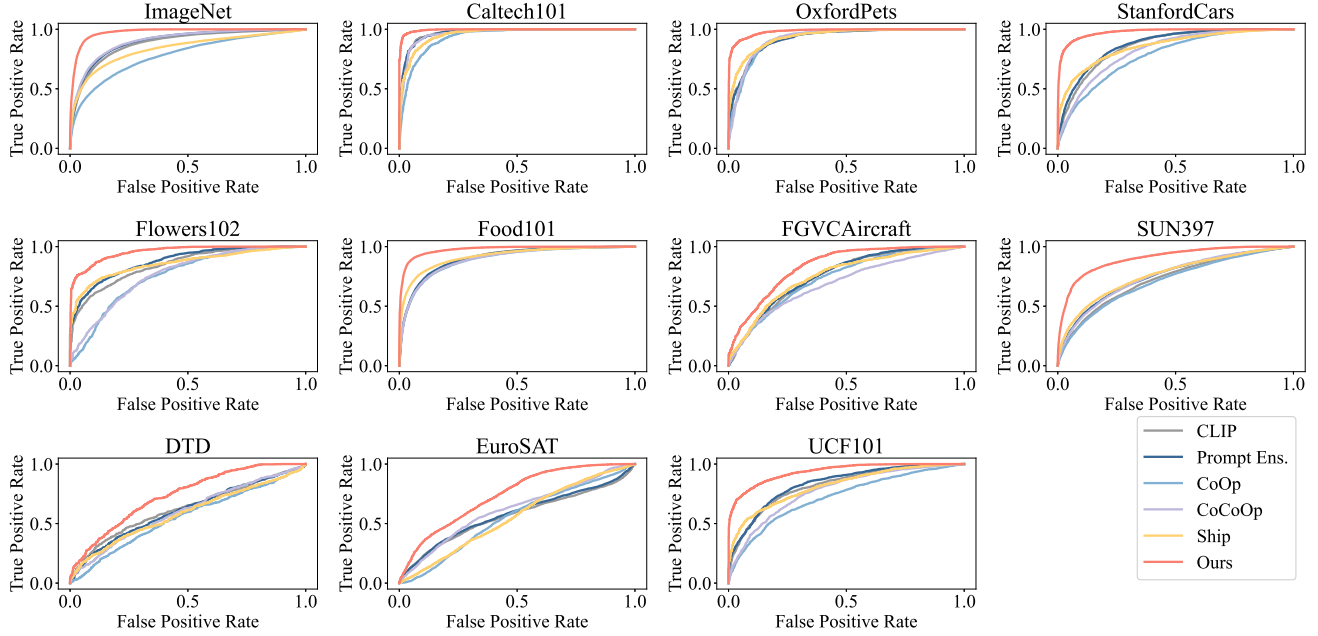


Figure 8. The roc curve for detecting new classes of each method on 11 datasets.

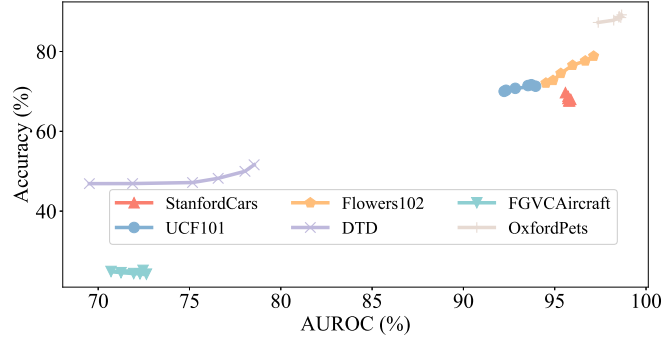


Figure 9. Correlation between performance of \mathcal{M}_O and accuracy.

B.6. Simple Ensembling of CoOP Method

We also conduct experiments to evaluate whether directly ensemble multiple CoOP learners can achieve similar performance. The results, shown in Table 8, indicate that the ensemble of multiple CoOP prompts does not yield significantly better performance compared to the CoOP method. These results prove that the performance gain does not derive from simple prompt ensembling.

B.7. Evaluation Time

Our DECoOP approach adopts multiple prompts to detect OOD, so it may take more time. We compared the running time taken by CoOP, CoCoOP, and DECoOP methods when evaluating the testing set of two datasets in Table 9. The results show that the running time of the DECoOP is not significantly longer than the CoOP method since the computation can be performed in parallel. However, our DECoOP approach runs in two stages (i.e., OOD detection and classification stages), therefore, the running time will be approximately double compared to the CoOP method. However, the running time of CoCoOP rises significantly as the number of categories increases, where demonstrates our DECoOP is efficient.

Table 8. Performance comparison between our proposed method and the ensemble of multiple CoOP prompts is conducted. The results demonstrate that directly combining multiple CoOP learners does not yield significantly better performance compared to the CoOP method. Moreover, our proposed algorithm outperforms other methods.

METHOD	FLOWERS102	DTD	CALTECH101	STANDFORDCARS
CoOP	72.11	48.18	93.24	63.81
CoOP×2	71.62	50.08	93.31	64.20
CoOP×4	73.12	49.99	92.89	64.32
CoOP×6	71.89	49.69	92.87	65.03
DECoOP	78.61	51.44	94.50	69.64

Table 9. Evaluation running time of CoOP, CoCoOP, and DE-CoOP methods.

DATASETS	#CLASSES	CoOP	CoCoOP	DECoOP
EUROSAT	10	14.1S	62.0S	34.1S
FOOD101	101	50.5S	711.7S	131.5S

Table 10. Comparison with weight interpolating methods.

	NEW ACC.	ACCURACY
CLIP	65.48	63.92
CoOP	57.75	65.58
RFT (WORTSMAN ET AL., 2022)	65.34	69.26
DECoOP	66.54	69.69

B.8. Comparison with Weight Interpolating Methods

Existing studies (Wortsman et al., 2022; Ilharco et al., 2022) observe that interpolating weights for tuned and original vision-language models can improve the generalization capacity. In the context of prompt tuning, we can interpolate weights of the tuned prompt and original prompt. We report the average results on all datasets using ViT-B/16 architecture in Table 10. The results show that interpolating weights can give better performance compared to both the original model and the tuned model, which aligns with the conclusion of existing studies. Our DECoOP outperforms other methods, demonstrating its effectiveness. Note that Weight Interpolating Methods and DECoOP have studied the different stages in fine-tuning, therefore, the combination of both to further enhance performance can be a direction for future research.

C. Limitation and Future Work

Our paper proposes the integration of OOD detection into prompt tuning to prevent performance degradation on new classes. In addition to the content discussed at the end of Section 6, one limitation of our approach is the potentially increased time consumption due to the adoption of a two-stage classification process. Integrating knowledge (Yang et al., 2024b;a) into the prompt tuning to achieve the better generalization is also a future direction to explore. Experiments detailed in Appendix B.7 demonstrate that our method’s running time is shorter than some existing methods (e.g., CoCoOP), proving that the running time of our proposal is within an acceptable range. One possible solution is to integrate the two-stage classification into prompt training through the utilization of advanced training strategies, which can be explored as potential research directions in the future.