

Active Learning for Vision-Language Models

Bardia Safaei
 Johns Hopkins University, USA
 bsafaei1@jhu.edu

Vishal M. Patel
 Johns Hopkins University, USA
 vpatel136@jhu.edu

Abstract

Pre-trained vision-language models (VLMs) like CLIP have demonstrated impressive zero-shot performance on a wide range of downstream computer vision tasks. However, there still exists a considerable performance gap between these models and a supervised deep model trained on a downstream dataset. To bridge this gap, we propose a novel active learning (AL) framework that enhances the zero-shot classification performance of VLMs by selecting only a few informative samples from the unlabeled data for annotation during training. To achieve this, our approach first calibrates the predicted entropy of VLMs and then utilizes a combination of self-uncertainty and neighbor-aware uncertainty to calculate a reliable uncertainty measure for active sample selection. Our extensive experiments show that the proposed approach outperforms existing AL approaches on several image classification datasets, and significantly enhances the zero-shot performance of VLMs.

1. Introduction

The rise of foundational vision-language models (VLM) has enabled impressive progress in various tasks in the field of computer vision [30, 34, 39, 46, 50, 55]. These models are pretrained on a large collection of image-text pairs and are typically trained using contrastive learning objectives. For example, both CLIP [39] and ALIGN [50] formulate their learning objectives as contrastive losses, which learn to bring images and their corresponding textual descriptions closer in the feature space while pushing unmatched pairs apart. By pretraining at a large scale, these models learn a broad understanding of visual concepts enabling them to effectively transfer to numerous downstream tasks. Furthermore, their pretraining enables generalization and zero-shot learning capabilities that surpass those of models trained under supervision on more limited datasets.

Although pretrained VLMs have demonstrated their effectiveness in zero-shot settings, their performance on domain-specific datasets is comparatively low when compared to the supervised models specifically trained on those




Figure 1. VLMs such as CLIP can effectively transfer to various downstream vision tasks. However, their performance on a novel target dataset still falls behind a supervised model specifically trained on the target dataset. Active learning approaches aim to reduce this performance gap by querying only a few beneficial samples from the unlabeled data, acquiring their labels, and efficiently utilizing them for improving the VLM performance.

datasets. Therefore, fine-tuning these foundational models on specialized datasets to enhance their performance is a promising direction. Prompt tuning (PT) [21, 26, 63, 64] has become an important line of research in recent years for efficiently adapting pre-trained VLMs to new tasks. Context Optimization (CoOp) [63] is a representative prompt tuning approach that replaces context tokens within a prompt into a set of learnable vectors while keeping the weights of CLIP’s image and text encoders frozen. Thus, CoOp requires only a minimal set of labeled images for training and has shown significant improvements compared to manually tuned prompts across various image recognition datasets.

Despite these advancements, PT methods typically employ the few-shot setting for fine-tuning the prompts. However, this setting assumes access to a few labeled samples from each class, which may not hold in real-world scenarios. For instance, in a time-sensitive task where a stream of unlabeled data becomes available for immediate training, labeling data until all classes are covered can be costly. Also, the randomly selected few-shot samples may not provide useful and complete information about some datasets. In contrast, active learning (AL) [40, 43] aims to select the most informative samples from the unlabeled data without requiring access to samples from each class. As a result, AL is a promising approach for more efficient adaptation of VLMs by focusing on the most beneficial samples, making it both cost-effective and scalable (see Fig. 1).

However, it has been shown that large pre-trained mod-

els like CLIP often produce uncalibrated outputs [28, 61], which can lead to imbalanced predictions, with some classes being predicted less frequently than others. This issue negatively impacts conventional uncertainty-based AL methods, as they rely on miscalibrated outputs to calculate uncertainty, resulting in suboptimal performance. To address this, we propose **Calibrated Entropy-weighted Clustering (CEC)**, a novel AL framework designed to enhance the selection of informative samples from the unlabeled data which after adding to the labeled training data, can maximally boost the performance of prompt tuning approaches in VLMs. Our method utilizes a calibrated uncertainty score, followed by a novel clustering approach that enables sampling from more vulnerable regions in the feature space. Specifically, we calibrate the predictive entropy of the CLIP model to reduce the bias towards the more frequently encountered categories. This entropy score quantifies the uncertainty for a single sample. Additionally, we leverage CLIP’s rich representations by proposing a neighbor uncertainty measure for each sample that enhances the reliability of uncertainty estimates. Finally, to ensure diversity, we adaptively cluster and perform an uncertainty-weighted sampling on each cluster according to the AL budget. Our experiments show that our method outperforms existing state-of-the-art AL methods on several image classification datasets and significantly improves the zero-shot performance of VLMs. The contributions of this paper can be summarized as follows.

- In this paper, we systematically analyze and show the advantages of utilizing AL strategies for PT of VLMs. In particular, we propose an AL method for VLMs that selects beneficial samples for PT and significantly improves the zero-shot performance of VLMs.
- We introduce an AL framework that uses both self- and neighbor-aware uncertainty of unlabeled samples to select informative samples during AL rounds.
- Our experimental results show that the proposed method outperforms existing state-of-the-art methods on several image classification datasets.

2. Related Work

Vision-Language Models (VLM). Vision-Language Pre-training (VLP) has emerged as a promising method for developing transferable and versatile recognition models by establishing connections between visual content and language descriptors. This approach has been explored in various studies [8, 14, 23, 32, 55], where the primary challenge is the limited size of the training datasets, such as Flickr [22] and COCO Captions [11]. However, recent advancements in VLP, exemplified by models like CLIP [39] and ALIGN [20] models, have achieved impressive outcomes. These models utilize web-scale noisy image-text

pairs and employ a contrastive objective that aligns matching image-text pairs while distancing non-matching ones. Unlike ALIGN and CLIP, which mostly focus on specific visual tasks, FLAVA [46] is a universal VLM targeting vision, language, and multi-modal tasks. BLIP [30] is a multi-modal mixture of encoder-decoder, which is pretrained on a bootstrapped dataset to obtain improved performance. BLIP-2 [29] enhances the cost-efficiency of BLIP by keeping image and language encoders frozen and pretraining a lightweight transformer to reduce the modality gap. SLIP [36] incorporates self-supervision into the multi-modal pre-training objective. By harnessing the power of natural language supervision, these VLMs not only obtain robust visual representations but also show remarkable adaptability to a wide range of downstream applications.

Prompt Tuning (PT). Recent advancements in natural language processing (NLP) have given rise to a novel paradigm known as prompt learning/engineering [6, 15, 31, 45, 57, 60], which aims to optimize learnable prompts instead of end-to-end fine-tuning of the model. has gradually supplanted the traditional fine-tuning approach in NLP. Recently, there have been enormous research efforts to develop prompt learning approaches for fine-tuning VLMs [21, 26, 33, 59, 63]. For example, CoOp [63] replaces context tokens within a prompt into a set of learnable vectors while keeping the weights of CLIP’s image and text encoders frozen. VPT [21] learns prompts on the visual side. It introduces learnable parameters to the input sequence of every Transformer layer, and these learnable parameters are jointly optimized with the linear classification layer during fine-tuning. Unlike CoOp and VPT, MaPLe learns the prompts in a multi-modal fashion, utilizing both the vision and language branches. ProDA [33] is another PT method that attempts to learn the category-wise distribution of prompts and utilize those distributions to generate a diverse set of prompts.

Active Learning (AL). AL focuses on enhancing the performance of a deep model on an unlabeled dataset by selecting a few informative samples for supervised training. AL algorithms are particularly important when labeling the entire unlabeled data is impractical and the annotation budget is limited. AL methods can be divided into three main groups, namely uncertainty-based [2, 3, 37], diversity-based [1, 2, 42], and ensemble-based [13] approaches. Uncertainty-based methods are based on selecting the most confusing samples for the model by relying on various uncertainty scoring functions, such as entropy [41, 49], softmax margin [3], and mutual information [27]. Diversity-based approaches [42, 56] attempt to choose samples from diverse regions of feature space to better adapt to the distribution of unlabeled data. Ensemble-based methods [13, 17] use a measure of disagreement among an en-

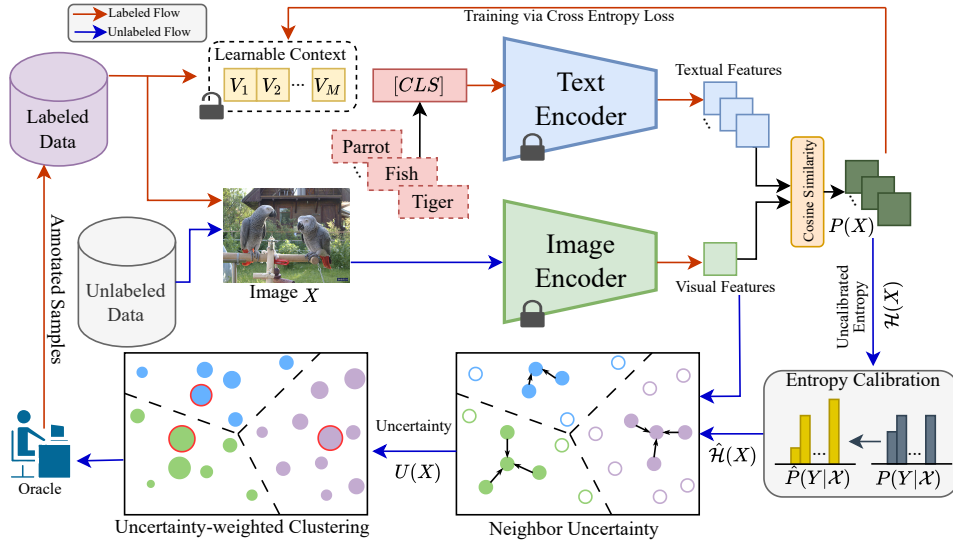


Figure 2. **An illustration of our CEC framework.** For a given unlabeled dataset, the visual features are extracted, and the prediction probabilities are calculated using textual features. Next, we calibrate the entropy score and utilize it to precisely quantify the uncertainty of an unlabeled sample, considering its similarities to textual embeddings. Moreover, we regularize this entropy by incorporating the uncertainty of a sample’s neighbors. Ultimately, this calibrated entropy is integrated into an uncertainty-weighted clustering approach to ensure diverse sample selection. The selected samples are then annotated and used for prompt tuning of CLIP.

semble of models with different characteristics (e.g. model initialization) for selection. Lately, there have been methods that improve AL performance by using a hybrid of multiple sample selection criteria [2, 37, 38, 51–53]. For example in [2], the authors combine diversity and uncertainty for improved performance by clustering gradient representations of unlabeled samples. In [53, 54], the authors optimize a parametric model to find the most informative samples for a pretraining-finetuning task. [37] utilizes a mixup technique to generate novel features for unlabeled data. Samples with confusing novel features are then clustered, and queries are made from different clusters. Recently, Bang et al. [4] proposed a class-balancing strategy added to conventional AL methods to improve their PT performance.

3. Background

In this section, we first introduce the problem of AL for prompt tuning of VLMs. Following that, we provide a brief overview of the training process for the CLIP model, followed by an explanation of the prompt tuning approach we employed for fine-tuning the CLIP model.

Problem Formulation. In this paper, we consider active learning for a K -way image classification task using the CLIP model, where K is the number of classes. In this setting, we begin with a pool of unlabeled data $\mathcal{D}_U = \{x_i^U\}_{i=1}^{N_U}$ and an initially empty labeled dataset, denoted as $\mathcal{D}_L = \emptyset$. At each AL round, the query function selects a batch of b samples X^q from the unlabeled data, and their

labels are queried from an oracle. We then add these annotated samples to the labeled dataset as $\mathcal{D}_L = \mathcal{D}_L \cup X^q$. The updated \mathcal{D}_L is employed to finetune the CLIP model by tuning CoOp’s learnable prompts.

3.1. Contrastive Language-Image Pretraining

In this section, we briefly review the well-known Contrastive Language-Image Pretraining (CLIP) model. The CLIP’s architecture consists of an image encoder and a text encoder. CLIP utilizes a contrastive loss [9] to align the feature vectors generated by these encoders. In particular, for a given batch of images and their corresponding text descriptions, CLIP initially extracts the visual and textual embeddings via encoders. Next, these embeddings undergo normalization, and the model maximizes the cosine similarity between a textual embedding and its corresponding visual embedding. Simultaneously, CLIP minimizes the cosine similarities between unmatched textual and visual embeddings. This approach enables CLIP to learn multimodal information from massive datasets of image-text pairs, contributing to its impressive zero-shot classification performance. During inference, for a given image x belonging to a dataset with K different classes and class names $\{\text{CLASS}_i\}_{i=1}^K$, K class-wise prompts are generated as ‘a photo of a $\{\text{CLASS}_i\}$ ’. By passing these prompts to the text encoder, we obtain class-wise textual embeddings t_i . The probability of x belonging to the class i is then calcu-

lated as follows:

$$P(y = i|x) = \frac{\exp(\text{sim}(f(x), t_i)/T)}{\sum_{j=1}^K \exp(\text{sim}(f(x), t_j)/T)}, \quad (1)$$

where $f(\cdot)$ denotes the image encoder, $\text{sim}(\cdot, \cdot)$ is the cosine similarity, and T is the temperature scaling value. The class with the highest probability is considered the prediction of x . The value of T is set to 0.01 in the original CLIP paper, and we follow the same setting in this paper.

3.2. Prompt Tuning for Training

One of the challenges associated with large multimodal pretrained models, such as CLIP, is fine-tuning them on new training samples while preserving their intrinsic multimodal representations. Typically, two common approaches are considered: end-to-end fine-tuning of the CLIP encoders and linear probing. The end-to-end fine-tuning approach updates all parameters of the model, whereas the linear probing method trains a linear classification layer on top of CLIP’s backbone. However, these approaches not only risk diminishing the CLIP’s adaptability to new domains but also show instability when employed in scenarios with limited training data.

To mitigate such limitations, we utilize a prompt tuning approach called Context Optimization (CoOp) [63] in this paper. CoOp replaces the fixed context tokens in prompts with learnable vectors.

In particular, it uses M learnable context vectors $\{V_1, V_2, \dots, V_M\}$ along with the class token C_i for the i -th class name as the prompt. Thus, the textual embedding for class i is calculated as $t_i = h([V_1, V_2, \dots, V_M, C_i])$ where h is the text encoder¹. CoOp trains learnable context vectors via cross-entropy loss while the parameters of the CLIP backbone are kept frozen.

4. Methodology

In this section, we provide a detailed elaboration on our proposed approach, Calibrated Entropy-weighted Clustering (CEC), which is tailored to the problem of AL for prompt tuning of pre-trained VLMs, such as the CLIP model. In CEC, we first utilize a calibrated entropy score to effectively quantify the uncertainty of an unlabeled sample based on its similarities to the textual embeddings. Furthermore, we regularize this entropy by considering the uncertainty of a sample’s neighbors to mitigate the problem of outlier selection. Finally, we incorporate our calibrated entropy into a clustering approach to ensure diverse sample selection (see Algorithm 1 and Fig. 2 for more details).

Notation. We consider a CLIP model with an image encoder $f(\cdot)$ and a text encoder $h(\cdot)$ for a K -way classification

¹ [63] has explored other variations such as positioning the class token in the middle of context tokens or using class-specific context tokens. In our experiments, we use the default setup as mentioned.

problem. We denote the textual embedding corresponding to the i -th class as t_i .

4.1. Entropy Calibration

To measure the uncertainty of a given unlabeled image x , we utilize CLIP’s predicted entropy as follows:

$$\mathcal{H}(x) = - \sum_{i=1}^K P(y = i|x) \cdot \log(P(y = i|x)), \quad (2)$$

where $P(y = i|x)$ was defined in Eq. 1. However, this entropy can be unreliable given the miscalibration issue inherent in large pretrained models including CLIP [44, 58, 61]. To mitigate this issue, we estimate the *contextualized prior* [18] for each class to calibrate the predicted probabilities. For the i -th class, we select the first N samples with the highest $P(y = i|x)$ values as follows:

$$S_i = \{x' | P(y = i|x') \in \text{Top}_N(P(y = i|x))\}, \quad (3)$$

where $x, x' \in \mathcal{D}_U$. The contextualized prior for class i is calculated by

$$Q(i) \approx \frac{1}{N} \sum_{x \in S_i} P(y = i|x). \quad (4)$$

We then calibrate the probabilities as

$$\hat{P}(y = i|x) = \left(\frac{P(y = i|x)}{Q(i)} \right) / \left(\sum_{j=1}^K \frac{P(y = j|x)}{Q(j)} \right). \quad (5)$$

Finally, the calibrated entropy $\hat{\mathcal{H}}$ is calculated by substituting Eq. 5 to Eq. 2.

4.2. Neighbor Uncertainty

While we have calibrated the entropy of each unlabeled sample, relying only upon $\hat{\mathcal{H}}$ might lead to selecting highly uncertain but outlier samples. To address this problem, for a given sample, we take the uncertainty of its nearest neighbors into consideration. Thus, an unlabeled sample is recognized as uncertain if it has both high self-uncertainty and neighbor uncertainty. Labeling such samples not only resolves their self-uncertainties but also helps reduce the model’s confusion in a local neighborhood around those samples. Formally, we denote the k -nearest neighbors of a given sample x in feature space as $k\text{NN}(x)$. The neighbor uncertainty is defined as:

$$\mathcal{H}_{\text{NN}}(x) = \frac{1}{k} \cdot \sum_{x_i \in k\text{NN}(x)} \exp(-\alpha \|z - z_i\|_2^2) \cdot \hat{\mathcal{H}}(x_i), \quad (6)$$

where $z = \frac{f(x)}{\|f(x)\|}$ is the normalized visual representation of x . We formulate the uncertainty of a sample x as

$$U(x) = \hat{\mathcal{H}}(x) + \mathcal{H}_{\text{NN}}(x). \quad (7)$$

4.3. Uncertainty-weighted Clustering

Similar to several existing AL methods [2, 37, 38], we propose to perform clustering on the visual features to ensure sampling from diverse regions of the feature space. Specifically, we assign weights to visual features based on

Algorithm 1 Our proposed *CEC* algorithm for AL

```
1: Input:
2:   unlabeled data  $\mathcal{D}_U$ , number of AL
   cycles  $R$ , known categories  $K$ , per-cycle budget  $b$ ,
   image encoder  $f$ , text encoder  $h$ ,
3: Process:
4:   for  $c = 0, 1, \dots, R - 1$  do
5:      $\forall x \in \mathcal{D}_U, U(x) \leftarrow 0$       # Initialization
6:     calibrate the predicted entropy of  $x$  via Eq. 5.
7:     compute neighbor uncertainty.  $\mathcal{H}_{NN}$  via 6.
8:     compute  $U(x)$  via 7.
9:     # Clustering
10:    perform weighted K-Means clustering on features.

11:    for the  $i$ -th cluster, find the sample  $x_i$  closest to
12:    the cluster center.
13:    # All queries for the current cycle:
14:     $X^q \leftarrow x_1 \cup x_2 \cup \dots \cup x_b$ 
15:     $\mathcal{D}_L \leftarrow \mathcal{D}_L \cup X^q$ ,
16:     $\mathcal{D}_U \leftarrow \mathcal{D}_U / X^q$       # Update datasets
17:    # Prompt tuning
18:    update CoOp’s learnable context tokens  $V_i$ 
19:    via cross-entropy on  $\mathcal{D}_L$ 
20:  Return updated prompts
```

their corresponding uncertainty (Eq. 7). Intuitively, this weighting mechanism increases the number of instances from a particular sample proportional to its uncertainty, enhancing the likelihood of its selection from the cluster. Our clustering approach strikes a balance between uncertainty and diversity resulting in significant performance improvement. We have empirically demonstrated that uncertainty-weighted clustering achieves superior performance compared to the alternative approach that clusters the features and selects the most uncertain samples within each cluster (see Fig. 3 (right)). Our clustering approach can be implemented using a weighted K-Means [19] algorithm.

4.4. Query Strategy and Training

As described in previous sections, at each AL cycle, we first calculate the uncertainty of samples via Eq. 7 and then perform uncertainty-weighted clustering to partition samples into b clusters. We query the label for the closest sample to each cluster center. After updating \mathcal{D}_L , we finetune CLIP on the labeled samples using the CoOp approach.

5. Experimental Setup

Datasets. Following prompt tuning works [26, 62], we select 6 different image classification datasets for our experiments, namely Describable Textures [10], Caltech-101 [5], EuroSAT [16], FGVC-Aircraft [35], Flowers-102 [24], and

UCF-101 [47]. These datasets cover a range of specialized computer vision tasks that are suitable for evaluating a large pre-trained model like CLIP that has zero-shot capabilities. More information about these datasets can be found in the supplementary material.

Implementation Details. In all experiments, we use a pretrained vision transformer ViT-B16 [12] model as the CLIP’s backbone. For CoOp prompt tuning, we initialize prompts with ‘a photo of a { }’. Also, we do not use class-specific context tokens, and the position of the class token is set to ‘end’. In each AL round, we load the zero-shot CLIP model and fine-tune it on the labeled data. We train for 200 epochs using the SGD optimizer [25] with an initial learning rate of 0.002, a momentum of 0.9, a weight decay of 0.005, and a cosine annealing scheduler. The batch size is set to 32 for all experiments. We fixed the value of N to 10. We utilize an NVIDIA A5000 GPU to run each experiment.

6. Experiments

We compare our method against a suite of state-of-the-art AL approaches, namely, **ALFA-Mix** [37], **BADGE** [2], **GCNAL** [7], **CoreSet** [42], **Entropy** [48], and **Random**. More details about these methods can be found in the supplementary material.

In our experiments, each method is applied within the same experimental setup to ensure a fair comparison. All datasets are split into train, validation, and test sets, and the train set is used as the unlabeled dataset. The effectiveness of each approach is measured by the model’s performance on the held-out test set after each round of AL.

Active Learning Setting. In all experiments, we perform 6 rounds of AL, where at each round we select 1% of the unlabeled data for annotation. For AL methods that require initial labeled data we perform random sampling in the first round. For a fair comparison, we conduct each experiment across three different seeds and report the mean and standard deviations of the results.

6.1. Main Results

CoOp results. Table 1 and Fig. 5 show the classification results corresponding to the CoOp prompt tuning method on six image classification datasets. In the table, we compare our method with uncertainty-based, diversity-based, hybrid, random, and zero-shot approaches.

The results are presented for 1%, 2%, and 5% annotation budgets. It can be seen that the proposed approach outperforms other baselines in most of the annotation ratios and datasets. In particular, on the Textures dataset, our method outperforms the SOTA method ALFA-Mix by margins of 3.2% and 2.6% for budgets 2% and 5%, respectively. Furthermore, it outperforms random sampling by large margins of 9.5% 8.6% and 4.1% for different annotation budgets on

Table 1. **Results for CoOp [63] Prompt Tuning.** Classification accuracy comparisons on 6 datasets. B is the per-cycle annotation budget. All experiments are conducted across three seeds, and the average results are reported. The second-best results are underlined.

Dataset	B (%)	Random	Entropy [48]	CoreSet [42]	BADGE [2]	ALFA-Mix [37]	GCNAL [2]	Ours	Zero-shot
Textures	1	38.4 ± 0.2	35.2 ± 0.8	-	40.2 ± 5.0	-	38.8 ± 0.9	47.9 ± 1.2	44.3
	2	44.2 ± 2.9	40.9 ± 2.0	44.9 ± 0.9	46.9 ± 1.4	49.6 ± 0.4	44.8 ± 0.4	52.8 ± 1.0	
	5	54.1 ± 2.9	49.3 ± 1.4	52.9 ± 2.3	56.2 ± 1.1	55.6 ± 1.9	55.0 ± 1.4	58.2 ± 2.0	
Caltech101	1	88.2 ± 3.4	86.1 ± 4.6	-	88.2 ± 1.7	-	88.4 ± 3.3	88.7 ± 1.5	91.3
	2	88.4 ± 3.2	89.3 ± 1.4	91.1 ± 2.0	89.8 ± 2.7	89.7 ± 0.8	89.8 ± 2.0	89.0 ± 1.3	
	5	91.1 ± 1.1	89.4 ± 0.8	91.3 ± 0.3	92.2 ± 0.1	92.3 ± 0.3	92.4 ± 0.9	92.8 ± 0.5	
EuroSAT	1	82.2 ± 1.0	70.5 ± 2.0	-	80.6 ± 0.7	-	82.1 ± 1.4	82.8 ± 1.6	42.0
	2	86.1 ± 1.0	78.1 ± 4.3	84.5 ± 1.5	85.6 ± 0.7	86.1 ± 0.3	84.0 ± 0.9	86.2 ± 0.6	
	5	87.8 ± 0.6	84.8 ± 2.1	87.9 ± 1.4	87.5 ± 0.5	88.3 ± 0.3	88.2 ± 0.6	88.0 ± 0.9	
FGVC-Aircraft	1	18.4 ± 0.6	19.7 ± 1.1	-	17.8 ± 1.7	-	18.4 ± 0.6	20.3 ± 1.1	24.9
	2	21.2 ± 1.4	22.0 ± 2.1	19.7 ± 1.4	20.7 ± 1.1	20.2 ± 1.5	21.2 ± 0.4	22.3 ± 1.0	
	5	26.0 ± 1.0	24.7 ± 0.6	23.0 ± 0.1	25.7 ± 1.0	28.7 ± 0.4	26.3 ± 0.7	27.1 ± 0.3	
Flowers102	1	60.2 ± 2.2	55.2 ± 4.7	-	53.5 ± 5.3	-	60.2 ± 2.3	64.1 ± 2.4	67.3
	2	66.3 ± 2.2	65.4 ± 3.5	62.2 ± 0.7	68.2 ± 1.7	74.0 ± 0.8	64.0 ± 4.0	75.6 ± 2.5	
	5	82.6 ± 2.2	80.3 ± 2.8	76.2 ± 2.6	86.2 ± 1.6	88.7 ± 1.2	80.0 ± 2.9	88.2 ± 1.6	
UCF101	1	55.4 ± 2.7	53.1 ± 3.9	-	50.7 ± 3.0	-	55.3 ± 3.7	57.6 ± 1.8	64.3
	2	66.6 ± 1.2	62.1 ± 1.9	65.4 ± 1.4	63.7 ± 1.5	66.9 ± 1.5	63.9 ± 1.9	67.0 ± 0.8	
	5	73.8 ± 0.4	72.8 ± 0.7	74.1 ± 3.0	75.3 ± 1.2	75.4 ± 1.9	73.1 ± 1.4	76.2 ± 0.6	
Average	1	57.1	53.3	-	55.2	-	57.2	60.2 (+3.0)	55.7
	2	62.1	59.6	61.3	62.5	64.4	61.3	65.5 (+1.1)	
	5	69.2	66.9	67.6	70.5	71.5	69.2	71.8 (+0.3)	




Figure 3. Effect of each component within our framework.




Figure 4. Hyperparameter analysis.

the Textures dataset. It can be observed that uncertainty-based approaches such as Entropy and GCNAL underperform random sampling in some budgets. This can be due to uncalibrated uncertainty utilization. However, our method consistently outperforms random sampling which shows the effectiveness of our approach. Moreover, Fig. 5 shows the performance of AL approaches over different AL cycles for Textures, Flowers102, and UCF101 datasets. Our method maintains a superior performance compared to other baselines in almost all rounds.

Results for other prompt tuning methods. To verify the generalization ability of our proposed method, we show the classification results for two additional prompt tuning algorithms, namely VPT [21] and MaPLe [26]. From the results in Table 2 and Fig. 7, it is evident that our approach achieves superior performance regardless of the prompt tun-

ing approach. Furthermore, the proposed method significantly improves the zero-shot performance of CLIP. For example in the MaPLe setting, our approach outperforms zero-shot CLIP by 0.6%, 7.7%, and 15% when sampling 1%, 2%, and 5% of the unlabeled samples, respectively.

Results for other architectures. Previous results are achieved using ViT-B/16 architecture as the CLIP image encoder. Here, we conduct experiments utilizing ResNet-50 and ViT-L/14 as the CLIP image encoder, with the prompt tuning method set to CoOp. Table 3 and Fig. 6 present the classification results for this setting. These results show the superiority of our approach compared to other AL methods and verify the effectiveness of our method regardless of the backbone architecture used.

6.2. Ablation Study

Effect of different components. In Fig. 3, we conduct an extensive ablation study on the Textures dataset to verify the effectiveness of each component within our framework. In Fig. 3 (left), we evaluate the effect of our *Entropy Calibration* and *Neighbor Uncertainty* components. We can observe that eliminating either of them leads to a performance drop, which highlights the importance of these elements. In Fig. 3 (right), we study the impact of our proposed *Entropy-Weighted Clustering* method. K-Means selects the nearest sample to the center from each cluster. Greedy Clustering selects the samples with the highest entropy from each cluster. We observe that our proposed entropy-weighted clustering method performs significantly better than such methods. This is because our method reduces the possibility of outlier selection. Even if an outlier shows high uncertainty, the increased number of inlier samples with relatively lower

Table 2. Results for VPT [21] and MaPLe [26] Prompt Tuning Methods.

Method	Dataset	$B(\%)$	Random	Entropy [48]	CoreSet [42]	BADGE [2]	Ours	Zero-shot	
VPT	Textures	1	38.9 ± 2.9	39.3 ± 1.5	-	42.5 ± 2.7	45.9 ± 1.0	44.3	
		2	49.3 ± 3.1	44.4 ± 1.3	47.0 ± 1.0	50.0 ± 0.8	52.8 ± 0.9		
		5	58.2 ± 2.5	57.0 ± 1.7	56.6 ± 0.7	59.7 ± 1.9	61.4 ± 0.4		
	Flowers102	1	61.0 ± 3.7	55.3 ± 4.2	-	56.7 ± 4.3	64.6 ± 5.2	67.3	
		2	70.6 ± 1.4	66.8 ± 3.8	65.9 ± 1.4	70.8 ± 2.1	75.9 ± 1.3		
		5	80.5 ± 1.5	81.6 ± 1.2	79.6 ± 2.1	84.8 ± 0.7	85.7 ± 0.9		
	UCF101	1	58.8 ± 2.7	60.0 ± 4.8	-	61.3 ± 3.8	63.7 ± 1.2	64.3	
		2	66.5 ± 0.8	67.4 ± 1.5	67.6 ± 1.9	68.6 ± 0.7	70.8 ± 0.9		
		5	75.0 ± 0.7	74.0 ± 1.5	75.2 ± 1.2	77.9 ± 0.4	78.0 ± 0.9		
	Average	1	52.9	51.5	-	53.5	58.1 (+4.6)	58.6	
		2	62.1	59.5	60.2	63.1	66.5 (+3.4)		
		5	71.2	70.9	70.5	74.1	75.0 (+0.9)		
	MaPLe	Textures	1	37.1 ± 3.5	36.0 ± 7.8	-	38.3 ± 0.7	45.8 ± 2.2	44.3
			2	46.0 ± 1.8	37.8 ± 6.5	40.7 ± 0.4	45.4 ± 1.7	50.8 ± 2.0	
			5	55.7 ± 2.0	54.9 ± 2.2	52.0 ± 1.2	57.7 ± 1.3	56.7 ± 0.9	
Flowers102		1	62.2 ± 1.4	61.2 ± 2.5	-	67.0 ± 3.1	66.0 ± 2.3	67.3	
		2	72.0 ± 2.7	64.0 ± 4.2	70.0 ± 4.4	70.1 ± 4.1	77.2 ± 1.9		
		5	82.5 ± 2.5	83.0 ± 2.1	79.0 ± 2.6	86.2 ± 0.2	86.5 ± 0.4		
UCF101		1	64.4 ± 3.9	59.7 ± 3.2	-	62.0 ± 0.7	65.8 ± 1.6	64.3	
		2	69.6 ± 1.3	66.2 ± 3.1	69.8 ± 0.1	68.8 ± 3.1	70.9 ± 2.4		
		5	76.9 ± 0.8	73.8 ± 2.4	76.3 ± 0.5	76.6 ± 0.8	77.6 ± 0.2		
Average		1	54.6	52.3	-	55.8	59.2 (+3.4)	58.6	
		2	62.5	56.0	60.2	61.4	66.3 (+3.8)		
		5	71.7	70.6	69.1	73.5	73.6 (+0.1)		

Table 3. Results for ViT-L/14 and ResNet-50 Architectures.

Model	Dataset	$B(\%)$	Random	Entropy [48]	CoreSet [42]	BADGE [2]	Ours	Zero-shot	
ViT-L/14	Textures	1	45.7 ± 0.9	40.0 ± 2.8	-	45.1 ± 2.5	55.1 ± 2.6	53.0	
		2	50.0 ± 1.9	43.8 ± 2.0	51.8 ± 1.9	52.1 ± 1.1	56.4 ± 0.8		
		5	60.4 ± 0.7	56.9 ± 2.0	59.7 ± 1.9	61.7 ± 2.1	63.7 ± 1.7		
	Flowers102	1	75.0 ± 1.3	76.2 ± 0.9	-	72.6 ± 3.8	80.0 ± 1.4	79.3	
		2	77.1 ± 1.7	73.2 ± 2.1	74.1 ± 2.1	79.1 ± 4.0	86.4 ± 2.0		
		5	87.1 ± 2.4	87.8 ± 2.0	84.2 ± 2.2	92.8 ± 0.2	93.6 ± 0.7		
	UCF101	1	73.3 ± 1.6	72.4 ± 1.0	-	73.7 ± 1.2	74.9 ± 1.3	74.2	
		2	76.4 ± 0.7	72.3 ± 1.4	74.1 ± 1.5	75.1 ± 2.6	77.5 ± 0.7		
		5	82.4 ± 1.7	80.0 ± 0.8	80.5 ± 0.3	82.0 ± 0.9	82.5 ± 1.5		
	Average	1	64.7	62.9	-	63.8	70.0 (+6.2)	68.8	
		2	67.8	63.1	66.7	68.8	73.4 (+4.6)		
		5	76.6	74.9	74.8	78.8	79.9 (+1.1)		
	ResNet-50	Textures	1	30.5 ± 2.6	26.8 ± 6.7	-	33.2 ± 2.5	34.1 ± 0.5	40.4
			2	38.9 ± 1.1	34.9 ± 1.7	38.9 ± 2.6	37.5 ± 1.1	43.0 ± 2.4	
			5	48.0 ± 1.6	46.2 ± 1.2	44.9 ± 1.0	48.4 ± 2.1	49.2 ± 2.3	
Flowers102		1	48.5 ± 3.8	41.8 ± 3.8	-	45.6 ± 2.3	53.6 ± 2.0	62.1	
		2	62.2 ± 2.9	55.7 ± 2.9	55.6 ± 2.1	60.1 ± 1.4	64.2 ± 2.9		
		5	72.8 ± 2.3	71.4 ± 1.7	64.8 ± 0.4	76.0 ± 1.4	77.7 ± 2.6		
UCF101		1	41.8 ± 2.9	42.7 ± 2.5	-	44.5 ± 0.6	51.0 ± 3.8	58.2	
		2	55.6 ± 1.8	51.7 ± 1.2	54.9 ± 1.6	56.2 ± 1.0	58.3 ± 1.2		
		5	66.2 ± 2.8	63.3 ± 2.3	65.2 ± 1.6	66.5 ± 0.9	67.0 ± 1.0		
Average		1	40.3	37.1	-	41.1	46.2 (+5.1)	53.6	
		2	52.2	47.4	49.8	51.3	55.2 (+3.0)		
		5	62.3	60.3	58.3	63.6	64.6 (+1.0)		

uncertainty will bias the cluster towards the high-density region of feature space, reducing the selection of outliers.

Hyper-parameter analysis. Our method introduces three hyperparameters: N (in entropy calibration), k (number of nearest neighbors), and α (in neighbor uncertainty). In Fig. 4, we study the effect of these parameters for three different sampling ratios on the Textures dataset. From the left figure, we can observe that our entropy calibration benefits performance compared to the no calibration setting. From

the middle and right figures, we see that our approach is robust to k and α , especially in higher sampling ratios. In our experiments, we fix the same value of N for all datasets. For other hyper-parameters, we perform a grid search to find the best values for each dataset. Specifically, we select k from $\{10, 20, 50\}$ and α from $\{0.05, 0.1, 1\}$. We run experiments for three seeds and report the average results.

Discussion on worse than zero-shot performance in certain cases. From the results we can observe that the per-




Figure 5. CoOp accuracy results over different 6 AL cycles. From left to right: Textures, Flowers102, and UCF101 datasets.




Figure 6. ResNet-50 accuracy results over 6 AL cycles. From left to right: Textures, Flowers102, and UCF101 datasets.




Figure 7. VPT accuracy results over 6 AL cycles. From left to right: Textures, Flowers102, and UCF101 datasets.

formance of AL methods, including ours, can be lower than zero-shot CLIP at $B = 1\%$ due to limited labeled samples and overfitting from training the prompts for 50-200 epochs. However, starting from the second round, our approach significantly surpasses zero-shot performance in almost all datasets, with an average gain of 10% at $B = 2\%$ and 16% at $B = 5\%$ in CoOp experiments.

7. Conclusion and Future Work

In this paper, we propose a novel framework for addressing the problem of AL for prompt tuning of VLMs. Specifically, our approach includes an entropy calibration component that reduces the prediction bias towards frequently observed categories. We further combine the self-uncertainty of samples with their neighbor’s uncertainty to provide a reliable uncertainty score. We leverage the uncertainty score to perform weighted clustering on the unlabeled data, leading to a diverse sampling of our approach. Moreover, we

conduct extensive experiments on a wide range of specialized datasets where we significantly outperform zero-shot CLIP by only sampling a small number of unlabeled data for annotation. Our paper shows the importance of active learning in adapting VLMs to new datasets and suggests a promising solution for prompt tuning of such models. In future work, we plan to investigate the integration of our AL framework with other fine-tuning approaches.

Acknowledgements. Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0008. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- [1] Sharat Agarwal, Himanshu Arora, Saket Anand, and Chetan Arora. Contextual diversity for active learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 137–153. Springer, 2020. [2](#)
- [2] Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. *arXiv preprint arXiv:1906.03671*, 2019. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [3] Maria-Florina Balcan, Andrei Broder, and Tong Zhang. Margin based active learning. In *International Conference on Computational Learning Theory*, pages 35–50. Springer, 2007. [2](#)
- [4] Jihwan Bang, Sumyeong Ahn, and Jae-Gil Lee. Active prompt learning in vision language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27004–27014, 2024. [3](#)
- [5] Monika Bansal, Munish Kumar, Monika Sachdeva, and Ajay Mittal. Transfer learning for image classification using vgg19: Caltech-101 image data set. *Journal of ambient intelligence and humanized computing*, pages 1–12, 2021. [5](#)
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [2](#)
- [7] Razvan Caramalau, Binod Bhattarai, and Tae-Kyun Kim. Sequential graph convolutional network for active learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9583–9592, 2021. [5](#)
- [8] Fei-Long Chen, Du-Zhen Zhang, Ming-Lun Han, Xiu-Yi Chen, Jing Shi, Shuang Xu, and Bo Xu. Vlp: A survey on vision-language pre-training. *Machine Intelligence Research*, 20(1):38–56, 2023. [2](#)
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [3](#)
- [10] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014. [5](#)
- [11] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11162–11173, 2021. [2](#)
- [12] 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. *arXiv preprint arXiv:2010.11929*, 2010. [5](#)
- [13] Yoav Freund, H Sebastian Seung, Eli Shamir, and Naftali Tishby. Selective sampling using the query by committee algorithm. *Machine learning*, 28:133–168, 1997. [2](#)
- [14] Zhe Gan, Linjie Li, Chunyuan Li, Lijuan Wang, Zicheng Liu, Jianfeng Gao, et al. Vision-language pre-training: Basics, recent advances, and future trends. *Foundations and Trends® in Computer Graphics and Vision*, 14(3–4):163–352, 2022. [2](#)
- [15] Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020. [2](#)
- [16] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. [5](#)
- [17] Hideitsu Hino and Shinto Eguchi. Active learning by query by committee with robust divergences. *Information Geometry*, 6(1):81–106, 2023. [2](#)
- [18] Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification. *arXiv preprint arXiv:2108.02035*, 2021. [4](#)
- [19] Joshua Zhexue Huang, Michael K Ng, Hongqiang Rong, and Zichen Li. Automated variable weighting in k-means type clustering. *IEEE transactions on pattern analysis and machine intelligence*, 27(5):657–668, 2005. [5](#)
- [20] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. [2](#)
- [21] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision*, pages 709–727. Springer, 2022. [1](#), [2](#), [6](#), [7](#)
- [22] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*, 2016. [2](#)
- [23] Armand Joulin, Laurens Van Der Maaten, Allan Jabri, and Nicolas Vasilache. Learning visual features from large weakly supervised data. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pages 67–84. Springer, 2016. [2](#)
- [24] Christopher Kanan and Garrison Cottrell. Robust classification of objects, faces, and flowers using natural image statistics. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2472–2479. IEEE, 2010. [5](#)
- [25] Nitish Shirish Keskar and Richard Socher. Improving generalization performance by switching from adam to sgd. *arXiv preprint arXiv:1712.07628*, 2017. [5](#)
- [26] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19113–19122, 2023. [1](#), [2](#), [5](#), [6](#), [7](#)

- [27] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical review E*, 69(6):066138, 2004. [2](#)
- [28] Will LeVine, Benjamin Pikus, Pranav Raja, and Fernando Amat Gil. Enabling calibration in the zero-shot inference of large vision-language models. *arXiv preprint arXiv:2303.12748*, 2023. [2](#)
- [29] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. [2](#)
- [30] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. [1](#), [2](#)
- [31] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. [2](#)
- [32] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021. [2](#)
- [33] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5206–5215, 2022. [2](#)
- [34] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7086–7096, 2022. [1](#)
- [35] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. [5](#)
- [36] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *European Conference on Computer Vision*, pages 529–544. Springer, 2022. [2](#)
- [37] Amin Parvaneh, Ehsan Abbasnejad, Damien Teney, Gholamreza Reza Haffari, Anton Van Den Hengel, and Javen Qin-feng Shi. Active learning by feature mixing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12237–12246, 2022. [2](#), [3](#), [4](#), [5](#), [6](#)
- [38] Viraj Prabhu, Arjun Chandrasekaran, Kate Saenko, and Judy Hoffman. Active domain adaptation via clustering uncertainty-weighted embeddings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8505–8514, 2021. [3](#), [4](#)
- [39] 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](#)
- [40] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40, 2021. [1](#)
- [41] Bardia Safaei, VS Vibashan, Celso M de Melo, and Vishal M Patel. Entropic open-set active learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 4686–4694, 2024. [2](#)
- [42] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*, 2017. [2](#), [5](#), [6](#), [7](#)
- [43] Burr Settles. Active learning literature survey. 2009. [1](#)
- [44] Jie-Jing Shao, Jiang-Xin Shi, Xiao-Wen Yang, Lan-Zhe Guo, and Yu-Feng Li. Investigating the limitation of clip models: The worst-performing categories. *arXiv preprint arXiv:2310.03324*, 2023. [4](#)
- [45] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020. [2](#)
- [46] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15638–15650, 2022. [1](#), [2](#)
- [47] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [5](#)
- [48] Dan Wang and Yi Shang. A new active labeling method for deep learning. In *2014 International joint conference on neural networks (IJCNN)*, pages 112–119. IEEE, 2014. [5](#), [6](#), [7](#)
- [49] Alfred Wehrl. General properties of entropy. *Reviews of Modern Physics*, 50(2):221, 1978. [2](#)
- [50] Fangyun Wei, Yue Gao, Zhirong Wu, Han Hu, and Stephen Lin. Aligning pretraining for detection via object-level contrastive learning. *Advances in Neural Information Processing Systems*, 34:22682–22694, 2021. [1](#)
- [51] Kai Wei, Rishabh Iyer, and Jeff Bilmes. Submodularity in data subset selection and active learning. In *International conference on machine learning*, pages 1954–1963. PMLR, 2015. [3](#)
- [52] Yichen Xie, Mingyu Ding, Masayoshi Tomizuka, and Wei Zhan. Towards free data selection with general-purpose models. *arXiv preprint arXiv:2309.17342*, 2023. [3](#)
- [53] Yichen Xie, Han Lu, Junchi Yan, Xiaokang Yang, Masayoshi Tomizuka, and Wei Zhan. Active finetuning: Exploiting annotation budget in the pretraining-finetuning paradigm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23715–23724, 2023. [3](#)
- [54] Wenshuai Xu, Zhenghui Hu, Yu Lu, Jinzhou Meng, Qingjie Liu, and Yunhong Wang. Activedc: Distribution calibration for active finetuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16996–17005, June 2024. [3](#)
- [55] Kaicheng Yang, Jiankang Deng, Xiang An, Jiawei Li, Ziyong Feng, Jia Guo, Jing Yang, and Tongliang Liu. Alip: Adaptive language-image pre-training with synthetic caption. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2922–2931, 2023. [1](#), [2](#)

- [56] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113:113–127, 2015. [2](#)
- [57] Hongbin Ye, Ningyu Zhang, Shumin Deng, Xiang Chen, Hui Chen, Feiyu Xiong, Xi Chen, and Huajun Chen. Ontology-enhanced prompt-tuning for few-shot learning. In *Proceedings of the ACM Web Conference 2022*, pages 778–787, 2022. [2](#)
- [58] Yue Yu, Rongzhi Zhang, Ran Xu, Jieyu Zhang, Jiaming Shen, and Chao Zhang. Cold-start data selection for few-shot language model fine-tuning: A prompt-based uncertainty propagation approach. *arXiv preprint arXiv:2209.06995*, 2022. [4](#)
- [59] Jiaming Zhang, Xingjun Ma, Xin Wang, Lingyu Qiu, Jiaqi Wang, Yu-Gang Jiang, and Jitao Sang. Adversarial prompt tuning for vision-language models. *arXiv preprint arXiv:2311.11261*, 2023. [2](#)
- [60] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5579–5588, 2021. [2](#)
- [61] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR, 2021. [2](#), [4](#)
- [62] 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. [5](#)
- [63] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [1](#), [2](#), [4](#), [6](#)
- [64] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15659–15669, 2023. [1](#)