

# PREFIX CONDITIONING UNIFIES LANGUAGE AND LABEL SUPERVISION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Pretraining visual models on web-scale image-caption datasets has recently emerged as a powerful alternative to traditional pretraining using image classification data. Image-caption datasets are more “open-domain”, containing broader scene types and vocabulary words, and result in models that have strong performance in few- and zero-shot recognition tasks. In this work, we study a pretraining strategy that uses both classification and caption datasets to unite their complementary benefits. First, we show that naively unifying the two datasets results in sub-optimal performance in downstream zero-shot recognition tasks, as the model is affected by dataset bias. While large-scale classification datasets can provide fine-grained categories with a balanced label distribution, the coverage of image domains and vocabulary words are biased compared to caption data. We address the problem with novel Prefix Conditioning, a simple yet effective method that helps disentangle dataset biases from visual concepts. This is done by introducing prefix tokens that inform the language encoder of the input data type (e.g., classification or caption) at training time. Our approach allows the language encoder to learn from both datasets while also tailoring feature extraction to each dataset. Our method is generic and can be easily integrated into existing VL pretraining objectives, such as CLIP or UniCL. In experiments, we show that prefix conditioning improves the performance in zero-shot image recognition and the robustness to the image-level distribution shift.

## 1 INTRODUCTION

Supervised classification datasets (e.g., ImageNet (Deng et al., 2009)) have traditionally been used to pretrain image representations for use in downstream tasks. However, web-scale image-caption datasets have recently emerged as a powerful pretraining alternative (Radford et al., 2021; Jia et al., 2021; Li et al., 2021a). Such datasets are more “open-domain”, containing a wider variety of scene types and vocabularies than traditional classification datasets, which are biased towards specific categories in their fixed label sets. Consequently, models trained on web-scale image-caption datasets have shown stronger generalization in novel tasks (Radford et al., 2021; Chan et al., 2022) and demonstrated remarkable performance on few and zero-shot image classification tasks (Radford et al., 2021). Nevertheless, classification datasets are still useful for pre-training as they have a more balanced coverage of categories, including rare and fine-grained categories, and a better focus on the labeled objects in each image.

Recent works (Yang et al., 2022a; Yu et al., 2022) therefore propose to combine caption and classification datasets for pre-training. Yang et al. (2022a) convert classification labels to “label-prompts” by inserting the label into a template sentence, e.g., “a photo of a <label>.”<sup>1</sup> Although training on the caption and label-prompt data achieves promising results, it does not fully resolve distribution differences between the open-domain caption data and the classification data. In particular, it produces a language embedding entangled with the classification dataset “bias”. We note that classification datasets tend to be biased in at least two ways: 1) the images mostly contain single objects from restricted domains, and 2) the vocabulary is limited and lacks the linguistic flexibility required for zero-shot learning. Therefore, the class embedding of “a photo of a dog” optimized for ImageNet may really mean *a photo of a dog from ImageNet* instead, which involves the bias in ImageNet and does not generalize well to other datasets. We empirically show that such dataset

<sup>1</sup>We use the term *prompt* to indicate a template sentence filled with a class name.

biases negatively affect unified pretraining by reducing the generalization of learned representations and thus jeopardizing zero-shot performance.

To recognize diverse concepts in the open domain, the language model needs to disentangle the dataset bias from the visual concepts and extract language embeddings generalizable to the open domain, e.g., the language embedding representing *a photo of a dog from an open-domain dataset, such as image-caption dataset*, instead of *a photo of a dog from ImageNet*. Given this intuition, we propose to learn dataset-specific language embeddings, while sharing knowledge from both datasets during training. We achieve this by a simple yet effective approach called *Prefix Conditioning*. The idea is to train a dataset-specific text token for each dataset so that the bias of the dataset can be absorbed into this token, and in return the remaining text tokens can focus on learning visual concepts. Specifically, we prepend a token for each dataset (e.g., image classification or image caption dataset) to the text input token sequence during pre-training. We note that the idea is in part inspired by the prefix or prompt tuning (Li & Liang, 2021; Lester et al., 2021; Zhou et al., 2021), where they have shown that learnable tokens prepended to the input token sequences of the pre-trained language models are able to learn task-specific knowledge and thus can be used to solve downstream tasks by combining the knowledge of pre-trained large language models and task-specific prefix tokens. In contrast to this, the proposed prefix conditioning learns prefix tokens and the vision and language encoders jointly for VL pre-training. Our method effectively unifies the supervision of the two datasets by allowing prefix tokens to learn the dataset bias and the remaining tokens, shared across datasets, to learn visual concepts.

In our experiments, the proposed simple technique achieves superior performance on zero-shot evaluation if we use the prefix of the caption dataset to extract the language embedding at test time. Meanwhile, inserting the prefix of the classification dataset leads to better performance on classification data. We also observe a drastic performance improvement when combining our prefix conditioning with the UniCL (Yang et al., 2022a) objective because of their complementarity. Our contributions are summarized as follows:

- We propose novel Prefix Conditioning *at pre-training time* to unify image-label and image-caption supervision. It is the first mechanism to use prefixes to condition the source of the dataset during vision language contrastive pre-training, rather than post pre-training.
- This simple approach improves zero-shot recognition performance by more than 6% on average in experiments on ImageNet21K (Deng et al., 2009) and CC12M (Changpinyo et al., 2021).
- Our comprehensive ablation study shows that prefix conditioning enables the model to switch its approach to extracting language features, e.g., attend to different words.

## 2 RELATED WORK

**Vision-Language Contrastive Learning.** Zero-shot recognition is conventionally solved by learning the relationship between visual representations and word embeddings of the class names (Frome et al., 2013; Akata et al., 2015; Xian et al., 2017; 2016; Wang et al., 2018; Mensink et al., 2014; Jayaraman & Grauman, 2014). Vision-language contrastive learning models, such as CLIP (Radford et al., 2021), pre-train a model with a large-scale image-caption data (400M) and achieve a remarkable improvement in zero-shot recognition. ALIGN (Jia et al., 2021) demonstrated the effect of scaling up the size of image-caption data. Various techniques have been proposed to improve the data efficiency given a relatively small amount of image-caption data (order of 10M). ALBEF (Li et al., 2021a) employs model distillation and masked language modeling. DeCLIP (Li et al., 2021b), SLIP (Mu et al., 2021) and TCL (Yang et al., 2022b) harness self-supervised contrastive learning. FILIP (Yao et al., 2021) uses token-to-token contrastive learning rather than the global contrastive learning used in CLIP. BLIP (Li et al., 2022) generates pseudo captions to diversify the language modality for each image. Unlike these works that handle only caption-style supervision, we focus on making good use of label supervision in vision-language pre-training together with captions. Our approach brings orthogonal improvement to the aforementioned works as they seek to improve training on image-caption data.

UniCL (Yang et al., 2022a) and K-Lite (Shen et al., 2022) unite the image-caption and image-label supervision by converting labels into text with pre-defined template sentences. UniCL leverages a supervised contrastive loss (Khosla et al., 2020) for image-label pairs. K-Lite (Shen et al., 2022) utilizes external knowledge from WordNet (Miller, 1995) and Wikitionary (Meyer & Gurevych, 2012).

The input noun is augmented with the class hierarchy and definition to enrich the supervision. Our method is complementary to these approaches as well since both UniCL and K-Lite do not consider the distribution shift between datasets. In experiments, we observe a significant performance boost when UniCL is combined with the proposed prefix conditioning.

**Learning with Prompts.** Prompt tuning is a popular technique to adapt a large language model to a specific task with few training data and low computational cost (Li & Liang, 2021; Gao et al., 2020; Lester et al., 2021; Liu et al., 2021a; Qin & Eisner, 2021). To avoid tuning all parameters of the model and using hand-crafted prompts, prefix embeddings are added to the training input and are the only parameters optimized during fine-tuning. The prefix embedding can be viewed as the knowledge of the downstream task. In this paper, since the target task is the zero-shot classification, the bias of the language embedding needs to be from the dataset covering a wide range of domains rather than a specific domain. Therefore, we choose to use the prefix embedding learned for image-caption dataset during test time. This technique is also effective in adapting a pre-trained vision-language model (Zhou et al., 2021; 2022) to few-shot classification by tuning the prompts of the language encoder to adapt to a downstream task. Additionally, prompt-tuning is effective in adapting a pre-trained vision model to a target task (Jia et al., 2022). While these works aim to tailor a large pre-trained model to a specific downstream task with a small amount of data or low computational cost, our goal here is to condition a model with the prefix during the pre-training stage by distinguishing between the image label and image caption data. We show that this allows a model to effectively share the knowledge obtained from two different types of data sources.

**Dataset bias in image recognition.** A large-scale image recognition dataset such as ImageNet (Deng et al., 2009) is known to be biased towards a specific image domain. Therefore, a model trained on such a dataset shows vulnerability to the distribution shift, e.g., shift in object pose (Barbu et al., 2019) and style of the images (Wang et al., 2019). Nevertheless, Wortsman et al. (2022); Kumar et al. (2022) show that adapting only a linear layer on the pre-trained models can improve performance on the downstream tasks with distribution shifts. This indicates the importance of having a good classifier on top of image encoders, such as linear classifiers generated by language encoders with preconditioning in our work. Dubey et al. (2021) propose a method for domain generalization. They condition image recognition models with the domain embedding, which discriminates the input image domains, and demonstrate the importance of the domain-specific image classifier. Our prefix conditioning can be seen as an attempt to de-bias the linear classifier to obtain a domain-specific classifier and adapt it from the classification to the captioning domain. Also, Lee et al. (2021); Bahng et al. (2020) approach the dataset bias in image classification by de-biasing image representations. By contrast, we tackle the problem in the framework of vision-language learning, disentangle the dataset bias in the language embedding and utilize the classifier obtained by the caption domain. We note that while captioning datasets can also have data biases, they tend to be more open-domain than existing classification datasets.

### 3 METHOD

In this section, we introduce the Prefix Conditioning technique for pretraining a deep learning model on both image-caption and image-label (classification) data. In Sec. 3.1, we discuss our problem setting and the background of contrastive learning with image-caption data. In Sec. 3.2, we explain the details of our training approach, and in Sec. 3.3 our inference procedure.

#### 3.1 PRELIMINARIES

**Setup.** Suppose we have access to two datasets: (i) an image label dataset  $\mathcal{S}_L = \{(\mathbf{x}_n, \mathbf{t}_n^P, y_n)\}_{n=1}^{N_L}$ , where  $\mathbf{x} \in \mathcal{X}$  is the image and  $\mathbf{t}^P \in \mathcal{P}$  is a prompt-style language description based on its class label  $y \in \mathcal{Y}$ , and (ii) a dataset of image-caption pairs  $\mathcal{S}_C = \{(\mathbf{x}_n, \mathbf{t}_n^C)\}_{n=1}^{N_C}$ , where  $\mathbf{t}^C \in \mathcal{T}$  is a caption. We assume that  $\mathbf{t}$  is the tokenized language description. For each image  $\mathbf{x}$ , an image encoder model  $f_\theta$  parameterized by  $\theta$  extracts a visual representation  $\tilde{\mathbf{v}} \in \mathbb{R}^{d \times 1}$ :  $\tilde{\mathbf{v}} = f_\theta(\mathbf{x})$ . For each caption or prompt  $\mathbf{t} \in \mathcal{T}$ , a text encoder  $f_\phi$  parameterized by  $\phi$  extracts a language representation  $\tilde{\mathbf{u}} \in \mathbb{R}^{d \times 1}$ :  $\tilde{\mathbf{u}} = f_\phi(\mathbf{t})$ .

**Contrastive Loss.** CLIP (Radford et al., 2021) is designed to find representations that match an image to its paired caption while separating unpaired ones. For  $i$ -th image  $\mathbf{x}_i$  and  $j$ -th language description  $\mathbf{t}_j$  in a batch  $\mathcal{B}$ , their features are normalized using  $\mathbf{v}_i = \frac{\tilde{\mathbf{v}}_i}{\|\tilde{\mathbf{v}}_i\|}$  and  $\mathbf{u}_j = \frac{\tilde{\mathbf{u}}_j}{\|\tilde{\mathbf{u}}_j\|}$ . Finally,

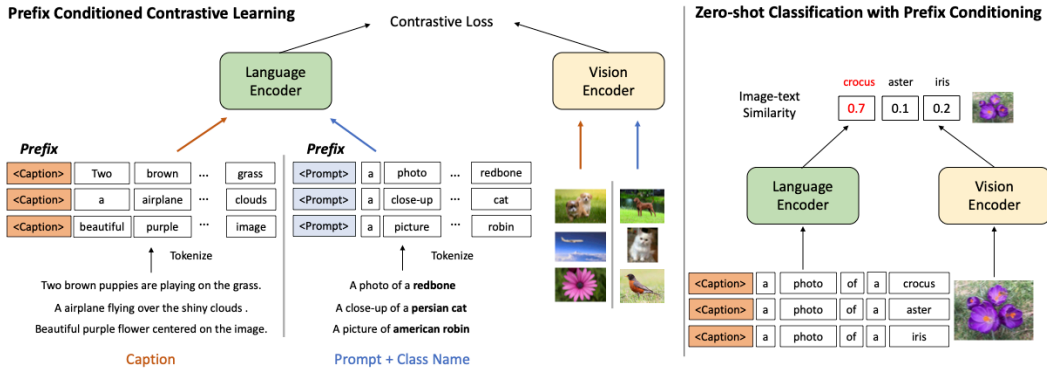


Figure 1: **Left:** Prefix conditioning at training time. Dataset-specific token is added to the input tokens with a contrastive learning objective applied. **Right:** Prefix conditioning at test time. Given a class name, we construct a class prompt with pre-defined templates and add a token used to condition real caption during training considering that image-caption dataset covers much wider range of image domains and vocabulary words than image classification dataset.

CLIP optimizes the symmetric multi-class N-pair loss (Sohn, 2016):

$$\min_{\{\theta, \phi\}} \mathcal{L}_{con} = \mathcal{L}_{t2i} + \mathcal{L}_{i2t}, \quad (1)$$

which includes two contrastive terms (a temperature hyper-parameter  $\tau$  controls the strength of penalties on hard negative samples):

$$\mathcal{L}_{t2i} = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \mathbf{u}_i^T \mathbf{v}_i)}{\sum_{j \in \mathcal{B}} \exp(\tau \mathbf{u}_i^T \mathbf{v}_j)}, \quad \mathcal{L}_{i2t} = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \mathbf{v}_i^T \mathbf{u}_i)}{\sum_{j \in \mathcal{B}} \exp(\tau \mathbf{v}_i^T \mathbf{u}_j)}. \quad (2)$$

UniCL (Yang et al., 2022a) composes each mini-batch with samples from both  $\mathcal{S}_L$  and  $\mathcal{S}_C$ . Then, for pairs from  $\mathcal{S}_L$ , they regard all samples from the same class as positive pairs while a sample from  $\mathcal{S}_C$  has a unique pair. Except for the number of positive pairs, no special treatment is given to differentiate between the image-caption and image-label data.

### 3.2 PREFIX CONDITIONED CONTRASTIVE LEARNING

Fig. 1 describes the overview of our approach. The goal is to enable the language encoder to learn embedding strategies conditioned on the type of input dataset. The conditioning can then be used to manipulate the bias at inference time.

Prefix-tuning (Li & Liang, 2021; Gao et al., 2020; Lester et al., 2021; Liu et al., 2021a; Qin & Eisner, 2021) shares the intuition that the prefix tokens are responsible for switching the context of a language model from the pre-trained task to the downstream task. These approaches leverage the prefix to tailor a model to a single task during fine-tuning and construct different prefixes for different natural language tasks (Lester et al., 2021). In our problem setting, there is no task distinction between the image-caption and image-prompt matching since both are formulated as contrastive learning. However, we focus on the fact that the two datasets have different biases in the image distributions and vocabulary words. The label-prompt sentences are embedded closer to the image classification data, even though we may want to use them to match a new label to an image from the open-domain image distribution during zero-shot classification.

To solve this problem, we propose to inform the model of the type of dataset at the input level to switch the feature extraction. Specifically, to make the model aware of the dataset type, prefix-conditioning prepends a prefix token to an input sentence to obtain  $\bar{\mathbf{t}}^P = [\text{PREFIX}^P; \mathbf{t}^P]$ ,  $\bar{\mathbf{t}}^C = [\text{PREFIX}^C; \mathbf{t}^C]$ . The brackets indicate the concatenation of two lists of discrete tokens;  $\text{PREFIX}^P$  and  $\text{PREFIX}^C$  denote a prompt-style and caption-style token respectively. In this way, we prepend the token to learn the dataset-specific bias, which enables us to disentangle the bias in language representations and utilize the embedding learned on the image-caption dataset at test time *even without an input caption*.

The number of prefix tokens can affect the performance of the model (Lester et al., 2021; Zhou et al., 2021; Li & Liang, 2021). To avoid significantly increasing the training cost, we set the number of prefixes to one in all experiments. Then, the language representations for each data source are extracted as  $\tilde{\mathbf{u}}^P, \tilde{\mathbf{u}}^C = f_\phi(\bar{\mathbf{t}}^P), f_\phi(\bar{\mathbf{t}}^C)$ . This input design is independent from the training objectives, and therefore we can easily apply the technique to optimize Eq. 1 or UniCL’s loss.

**Data Sampling.** Cui et al. (2021) argue that the data sampling matters when learning from multiple data sources in a contrastive learning framework, as the model may learn to distinguish the samples by exploiting the dataset bias. As such, we need to take data sampling into consideration in our problem setting as we learn from two different data sources. One option is a debiased sampling (Cui et al., 2021), which constructs each mini batch to contain samples from a single data source. Alternatively, as done in UniCL (Yang et al., 2022a), we can compose each mini-batch with samples from both data sources (image-caption and image-label) with equal probability. In experiments, we choose the debiased sampling, but empirically find that the choice of sampling does not significantly affect the performance.

### 3.3 INFERENCE WITH PREFIX CONDITIONING

During inference (the right side of Fig. 1), an input image is classified as one of  $K$  classes by embedding the corresponding label-prompts and choosing the one most similar to the image embedding. Following (Radford et al., 2021), we obtain class prompts by filling the default prompt templates with class names, and add a prefix. Considering the wider coverage of domains in the image-caption dataset, the caption-style prefix conditioning may work better to classify novel downstream data. In our experiments, we empirically find that the caption-style prefix indeed outperforms the prompt-style prefix with a large margin in zero-shot recognition while prompt prefix performs better on the image classification dataset used to train the model. We provide a detailed analysis of different conditioning in Section 4.3.

## 4 EXPERIMENTS

The goal of experiments is twofold: comparing our approach with baselines in zero-shot recognition, and analyzing the behavior of prefix conditioning. We describe the experimental setup in Sec. 4.1, show the main results in Sec. 4.2, and analyze the properties of prefix-conditioning in Sec. 4.3.

### 4.1 SETUP

**Training Datasets.** We conduct experiments on the setting where we have a large source of image-caption and image-label datasets. Following UniCL (Yang et al., 2022a), we utilize CC3M (Sharma et al., 2018) and CC12M (Changpinyo et al., 2021) as image-caption data. For the image classification dataset, we utilize ImageNet21K and ImageNet1K (Deng et al., 2009). While ImageNet1k contains 1,000 classes, ImageNet21K has more than 20,000 categories that include fine-grained and general object concepts. To observe the behavior in diverse image classification data, we also run experiments on ImageNet21K while excluding the classes of ImageNet1K. Details are explained in each section.

**Training.** We use the same prompt strategy and 80 prompt templates as used in CLIP (Radford et al., 2021). During training, we randomly sample one prompt template and fill it with the class names, followed by a tokenization step before feeding into the text encoder. We average language embeddings extracted from all 80 templates in validation. We use the same language encoder as CLIP (Radford et al., 2021) and Swin-Tiny transformer (Liu et al., 2021b) as the vision encoder following UniCL (Yang et al., 2022a). All models are optimized with AdamW (Loshchilov & Hutter, 2019) where the learning rate is set to 0.001, and weight decay to 0.1. All models are trained with a batch size of 1024. Considering the amount of training data, we train the models for 15 and 50 epochs in the experiments on ImageNet21K and ImageNet1K respectively.<sup>2</sup> For all training, we used a cosine learning rate schedule with a warm-up of 10,000 iterations.

**Baselines.** We train CLIP (Radford et al., 2021) and UniCL (Yang et al., 2022a) as our baselines. For comparison, we present results on CLIP trained only on image-caption or image classification data, as well as CLIP and UniCL trained on both image-caption and IN21K data. Unless otherwise

<sup>2</sup>When training a model on two different datasets, e.g., IN21K and CC12M, we count the epochs based on how many samples are used from the image classification dataset. For instance, in UniCL, each mini-batch consists of approximately the same number of samples from IN21K and CC12M. Then, to train a model for 15 epochs, we train for  $N/1024 \times 2 \times 15$  iterations, where  $N$  indicates the number of samples in IN21K.



Table 1: Performance comparison among different training datasets and training objectives. The proposed prefix conditioning shows improved zero-shot recognition accuracy across models trained with different combinations of image-classification and image-caption datasets and training objectives.

Training Data		Objective	Prefix Conditioning	Metric		
Classification	Caption			Size	IN-1K	Zero-shot 11 datasets
–	CC-3M	3M	CLIP		18.1	28.7
–	CC-12M	12M	CLIP		33.4	41.2
ImageNet-1K	–	1M	CLIP		72.1	20.2
ImageNet-21K	–	12M	CLIP		47.1	39.6
ImageNet-1K	CC-12M	13M	CLIP		68.7	43.3
ImageNet-1K	CC-12M	13M	CLIP	✓	<b>71.5</b>	<b>45.5</b>
ImageNet-1K	CC-12M	13M	UniCL		68.8	43.1
ImageNet-1K	CC-12M	13M	UniCL	✓	<b>71.7</b>	<b>44.5</b>
ImageNet-21K	CC-12M	25M	CLIP		56.8	49.5
ImageNet-21K	CC-12M	25M	CLIP	✓	<b>67.3</b>	<b>57.8</b>
ImageNet-21K	CC-12M	25M	UniCL		58.2	51.7
ImageNet-21K	CC-12M	25M	UniCL	✓	<b>66.5</b>	<b>58.4</b>
ImageNet-21K w/o IN-1K	CC-12M	24M	CLIP		29.1	46.9
ImageNet-21K w/o IN-1K	CC-12M	24M	CLIP	✓	<b>47.8</b>	<b>56.4</b>

Table 2: Ablation study for sampling in IN21K + CC12M. Equal sampling (ES) composes a mini-batch with roughly equal number of samples from two datasets. Debiased sampling (DS) samples a mini-batch of either IN21K or CC12M with equal probability.

Train Prefix	Sampling	IN-1K	Cal	CF100	CF10	ESTAT	Food	Flower	Pets	Patch	R45	VOC	DTD	AVG
✓	ES	56.8	70.2	55.0	79.4	21.1	46.0	60.3	57.2	51.2	24.8	57.7	21.4	49.5
	ES	<b>65.4</b>	<b>81.2</b>	<b>62.6</b>	<b>88.9</b>	<b>30.4</b>	<b>51.7</b>	<b>61.8</b>	<b>71.9</b>	50.0	<b>28.2</b>	<b>78.1</b>	<b>27.7</b>	<b>57.5</b>
✓	DS	58.7	65.9	55.0	85.7	22.8	40.8	55.7	60.2	50.0	20.6	45.2	23.8	47.8
	DS	<b>67.3</b>	<b>79.7</b>	<b>63.8</b>	<b>87.9</b>	<b>31.5</b>	<b>53.4</b>	<b>58.8</b>	<b>69.6</b>	<b>50.6</b>	<b>31.5</b>	<b>80.5</b>	<b>28.4</b>	<b>57.8</b>

stated, CLIP and UniCL are trained with equal sampling (ES) strategy as in Yang et al. (2022a), while our prefix conditioning model is trained with debiased sampling (DS) (Cui et al., 2021). We provide an analysis of the sampling in Sec. 4.2 and find that DS itself does not have a noticeable advantage over ES.

**Evaluation.** We evaluate the learned representations on supervised and zero-shot image classification on ImageNet1K<sup>3</sup> and on 11 datasets chosen from the ones used in CLIP (Radford et al., 2021) including object classification (e.g., CIFAR10, CIFAR100), fine-grained classification (e.g., Oxford-IIIT Pets, Oxford Flowers 102, and Food-101), and aerial images (e.g., EuroSAT and Resisc45). Although our main focus is at the zero-shot generalization, we also provide an analysis of a linear-probe evaluation of the image encoder.

## 4.2 MAIN RESULTS

We describe our main results in Table 1, followed by the analysis of prefix conditioning in Sec. 4.3.

There are three observations. First, the improvements upon a model trained only with image-caption or image-label data are obvious in almost all cases. As the previous work indicates (Yang et al., 2022a), the effectiveness of combining two types of supervision is clear from these results.

Second, in all cases, our prefix conditioning significantly improves performance on both ImageNet-1K (supervised recognition) and 11 zero-shot recognition tasks. When training on ImageNet-21K,

<sup>3</sup>While we follow the same zero-shot evaluation protocol when evaluating on ImageNet1K, we note that it is zero-shot only where we explicitly exclude ImageNet1K from the training, last two rows of Table 1

Table 3: Linear evaluation accuracy on models trained with and without prefix conditioning. Prefix conditioning slightly improves the performance upon a model without it (second row vs. last row).

Train Data	Prefix Conditioning	IN-1K	Cifar10	Cifar100	Caltech	Food	Pet	Patch	VOC	DTD
ImageNet-21K		71.5	94.3	79.1	83.5	79.1	86.3	82.3	88.9	61.3
ImageNet-21K + CC12M		69.2	93.0	76.4	82.4	78.4	82.2	81.4	88.7	61.4
ImageNet-21K + CC12M	✓	69.4	93.5	77.3	83.2	78.8	83.6	82.0	88.8	62.5

Table 4: Ablation study for test-time prefix conditioning. Note that the difference between two results come from the prefix used in test time and we use the same model for this evaluation.

Data	Test-time Prefix	IN-1K	Cal	C100	C10	ESTAT	Food	Flower	Pets	Patch	R45	VOC	DTD	AVG
IN-1K + CC12M	Prompt	<b>75.4</b>	71.7	35.5	63.9	24.2	20.0	8.1	72.2	50.4	24.2	61.1	15.3	40.6
	Caption	71.5	<b>75.1</b>	<b>39.4</b>	<b>70.5</b>	<b>26.7</b>	<b>33.9</b>	<b>13.9</b>	<b>72.3</b>	<b>50.5</b>	<b>25.8</b>	<b>67.8</b>	<b>25.4</b>	<b>45.5</b>
IN-21K + CC12M	Prompt	<b>71.4</b>	76.5	59.0	86.0	20.1	45.7	<b>62.3</b>	69.1	<b>52.4</b>	26.3	76.8	21.4	54.1
	Caption	67.3	<b>79.7</b>	<b>63.8</b>	<b>87.9</b>	<b>31.5</b>	<b>53.4</b>	58.8	<b>69.6</b>	50.6	<b>31.5</b>	<b>80.5</b>	<b>28.4</b>	<b>57.8</b>
IN-21K w/o 1K + CC12M	Prompt	40.8	74.9	61.0	84.6	31.2	48.1	58.7	45.2	<b>51.2</b>	23.5	67.5	21.4	51.6
	Caption	<b>47.8</b>	<b>81.9</b>	<b>63.3</b>	<b>87.3</b>	<b>32.4</b>	<b>52.9</b>	<b>62.8</b>	<b>57.0</b>	50.6	<b>25.6</b>	<b>80.1</b>	<b>26.2</b>	<b>56.4</b>

the conditioning improves the baseline by more than 8% in ImageNet-1K and more than 6% in zero-shot recognition on average. In training with ImageNet-1K, the margin from the baseline is smaller than training with ImageNet-21K, probably because the size of ImageNet-1K (1M) is much smaller than that of ImageNet-21K (12M). Also, prefix conditioning is effective in both UniCL and CLIP objectives. Due to its simplicity, our approach can be easily integrated with various objectives.

Finally, our method is less affected by ablating a part of categories. The classes of ImageNet-1K are excluded from ImageNet-21K in the last two rows of Table 1. Therefore, both approaches significantly drop performance on ImageNet-1K, whose task now becomes true zero-shot recognition, compared to other settings. Even in this setting, prefix conditioning maintains high accuracy on zero-shot recognition and outperforms a CLIP baseline model by a large margin.

**Sampling Method.** We analyze the data sampling scheme to construct a mini-batch in Table 2. We apply debiased sampling (DS) in our method, namely, sampling one data source with equal probability and getting a mini-batch of it. The other option is mixing two data sources with equal probability (ES). The table indicates that prefix conditioning works well with ES sampling and the sampling strategy itself is not advantageous. Ablating prefix conditioning during training clearly drops the performance in both sampling strategies, and the performance is worse than ES on average in zero-shot results (49.5 vs. 47.8). ES sampling should allow the model to differentiate sentences by using the prepended prefix. Interestingly, this result implies that differentiating sentences by prefix information does not much degrade the performance. The distinguished sentences enable the model to associate images from different datasets. Since images of two datasets are different with respect to the categories and the locations of objects in images, distinguishing the two kinds of images may not harm generalizability of the representations.

**Linear-probe Evaluation.** We evaluate the linear-probe performance in Table 3 to see the quality of learned image representations. Although the accuracy is better than the model trained without prefix conditioning (second line), the improvements are not substantial. This result indicates that the zero-shot performance gain obtained by our method is not due to the image representations. We investigate the learned language and image features in the next subsection.

#### 4.3 ANALYSIS OF PREFIX-CONDITIONING

We present a detailed analysis of prefix conditioning. We first study how different prefixes impact the zero-shot recognition performance and analyze their behaviors by looking into the attention weights of the language transformer encoder. We also demonstrate improved robustness with respect to the image-level domain shift. Unless otherwise stated, we employ a model trained with CLIP objective on ImageNet-21K and CC12M in this analysis. Finally, this section concludes that prefix conditioning enables the language encoder to switch its role during training, which eases learning from different types of datasets, e.g., image classification and image caption dataset.




Figure 2: An example of attention weights for an end token. Best viewed in color. The sentence shown here is one of class prompts in the VOC 2007 dataset. Different rows show the weights of different transformer layers. With a prompt prefix (leftmost), the model focuses on a class name (*airplane*) while caption prefix (middle) allows a model to pay attention to another noun, *sculpture*. By prefix conditioning, the attention of the model changes as intended.




Figure 3: T-SNE (Van der Maaten & Hinton, 2008) visualization of the class-prompt features of 20 classes of VOC 2007 with different prefix conditions. (a): Language embeddings with prompt (red) and caption (blue) prefixes, respectively. (b)(c)(d): Different colors indicate language embeddings of different classes.

**Test Time Prefix.** We analyze the role of the prefix token in Table 4, where the table describes the comparison in the choice of test time prefix conditioning. As explained in Sec. 3, the choice of prefix during test time should change the behavior of the model since the prefix should tailor the language encoder for classification-style or caption-style feature extraction. Except for the IN-1K results of a model trained with the entire IN21K or IN-1K, conditioning with the caption prefix shows much better results. The superiority of the caption prefix is noticeable in several datasets. This means caption prefix works better if the target comes from outside the image classification data, indicating that the class-prompt prefix conditioning makes the model tailored for the image classification dataset. Class-prompts prefix works better to categorize IN-1K data because the prefix is trained to specialize in classifying it. Note that caption-style prefix performs better than prompt-style prefix in IN-1K for a model trained with IN21K excluding IN1k classes. This indicates that the caption-style prefix works better when the vocabulary of the class name comes from outside the image classification data since the caption data covers much more diverse words.

**Prefix controls attention.** Fig. 2 visualizes the attention weights for an end token in different prefix conditions and models. The input sentence, *a sculpture of an airplane*, is one of the class-prompts. When a prompt prefix (leftmost) is employed, the language model pays attention to the class name at the first layer, it does not focus on the noun in other layers. The only noun the encoder focuses on is *airplane*. On the other hand, the model attends to both *sculpture* and *airplane* in the case of the caption prefix and unconditional model. Note that this behavior does not mean that the prompt-prefix performs better in zero-shot recognition as shown in experiments due to the effect of the bias in image classification dataset.

While we visualize only one example in the main text due to the space limit and defer more examples to the appendix, this highlights a general trend that the prompt prefix guides the language encoder to focus on a single word (e.g., class name), whereas the caption prefix makes the model attend to multiple words in the input sentence. In other words, prefix conditioning allows the language encoder to “switch gears” to represent sentences from different datasets (i.e., image-classification vs image-caption). On the other hand, the baseline model without prefix conditioning attends to multiple words (e.g., Fig. 2 rightmost) even though the input sentence is a class prompt. This indicates that it is hard to switch the gears without explicitly informing of the type of dataset.



Table 5: Evaluation on the robustness to the image-level domain shift. Prefix conditioned training achieves better robustness, and caption-prefix outperforms prompt-prefix in the images distinct from those used in training (IN-R and IN-S).

Train Data	Prefix Conditioning	Test-Time Prefix	IN	IN-V2	IN-R	IN-S
ImageNet-1K		N/A	72.1	59.3	19.9	17.8
ImageNet-1K + CC12M		N/A	68.7	57.4	27.7	27.8
ImageNet-1K + CC12M	✓	Caption	71.5	60.2	<b>31.8</b>	<b>30.7</b>
ImageNet-1K + CC12M	✓	Prompt	<b>75.4</b>	<b>63.3</b>	29.2	27.9
ImageNet-21K		N/A	47.1	41.1	20.1	16.1
ImageNet-21K + CC12M		N/A	56.8	48.6	29.4	30.6
ImageNet-21K + CC12M	✓	Caption	67.3	57.5	<b>35.2</b>	<b>34.6</b>
ImageNet-21K + CC12M	✓	Prompt	<b>71.4</b>	<b>61.1</b>	32.1	32.2

**Language Feature Visualization.** Fig. 3 visualizes extracted language features conditioned with different prefixes. As seen in Fig. 3a, language features extracted with caption-prefix (blue) and prompt-prefix (red) are clearly separated. In addition, prompt-prefix (Fig. 3b) has lower intra-class and higher inter-class variance, whereas caption-prefix (Fig. 3c) shows higher intra-class variance across prompts. Interestingly, results in Table 4 suggest that the caption-prefix conditioned language features result in a better zero-shot recognition performance than those conditioned on the prompt-prefix. Although the prompt-prefix mode extracts discriminative language embeddings, the embeddings do not perform well on the zero-shot recognition because the embeddings contain significant bias from image-classification dataset.

**Robustness in image domain shift.** Test samples can be unseen with respect to image classification data in two ways (or combinations of two): 1) The image is similar to the training distribution, but the class name is different from the seen image classification labels. 2) Although the class label is the same, the image data comes from a different distribution. Datasets evaluated in the zero-shot recognition include both two cases since the vocabularies and image are from different domains. To understand these two factors, we analyze the test-time prefix by using ImageNet-1K and evaluate the performance on image-level domain shift using variants of ImageNet, i.e., ImageNet-V2 (Recht et al., 2019), ImageNet-R (Hendrycks et al., 2021), and ImageNet-S (Wang et al., 2019). Table 5 describes the results of ablating prefix-conditioned training and changing the test-time prefix. The prefix-conditioned training outperforms all baselines. This reveals that the prefix-conditioned training achieves class embeddings that are generalizable across image domains. The prompt-style prefix performs the best in IN, IN-V2, both of which have image styles similar to ImageNet. By contrast, the caption-style prefix performs the best in IN-R and IN-S, which has art-style and sketch-style images respectively. Thus, the caption-style prefix generates more generalizable class embeddings for the domain dissimilar from the ImageNet training data. This observation is consistent with the results in the paragraph *Test time Prefix*.

## 5 CONCLUSION

In this paper, we explore a simple yet effective mechanism for unified pre-training on image-caption and image classification data. We propose to learn prefix tokens at training time to condition the language encoder to switch the input source. By learning a dataset-specific token for each dataset, the bias of the dataset can be absorbed into one token. Moreover, specifying the prefix allows the model to switch the manner of feature extraction and can control which visual domain the embedding is projected to. This approach boosts the performance of zero-shot recognition accuracy of the contrastive learning models. Our analysis suggests that the trained language encoder provides robustness to the image-level domain shift. Although we limit our scope to unifying image-caption and image-label supervision, incorporating other supervision types such as object detection or semantic segmentation is an interesting next step.

**Limitations and Negative Societal Impacts.** While our approach shows promising results on various zero-shot recognition benchmarks, the performance is poor on the dataset dissimilar to both image-caption and image classification pretraining data, PatchCamelyon (denoted as Patch in tables). In addition, training a model on webly-crawled data may have unwanted social implications since it may contain unintended private information, harmful images or texts.

## REFERENCES

- Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2927–2936, 2015. [2](#)
- Hyojin Bahng, Sanghyuk Chun, Sangdoo Yun, Jaegul Choo, and Seong Joon Oh. Learning de-biased representations with biased representations. In *International Conference on Machine Learning*, pp. 528–539. PMLR, 2020. [3](#)
- Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Advances in neural information processing systems*, 32, 2019. [3](#)
- Stephanie CY Chan, Adam Santoro, Andrew K Lampinen, Jane X Wang, Aaditya Singh, Pierre H Richemond, Jay McClelland, and Felix Hill. Data distributional properties drive emergent few-shot learning in transformers. *arXiv preprint arXiv:2205.05055*, 2022. [1](#)
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12M: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *CVPR*, 2021. [2](#), [5](#)
- Quan Cui, Boyan Zhou, Yu Guo, Weidong Yin, Hao Wu, and Osamu Yoshie. Zerovl: A strong baseline for aligning vision-language representations with limited resources. *arXiv preprint arXiv:2112.09331*, 2021. [5](#), [6](#)
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009. [1](#), [2](#), [3](#), [5](#)
- Abhimanyu Dubey, Vignesh Ramanathan, Alex Pentland, and Dhruv Mahajan. Adaptive methods for real-world domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14340–14349, 2021. [3](#)
- Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013. [2](#)
- Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*, 2020. [3](#), [4](#)
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8340–8349, 2021. [9](#)
- Dinesh Jayaraman and Kristen Grauman. Zero-shot recognition with unreliable attributes. *Advances in neural information processing systems*, 27, 2014. [2](#)
- 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*, pp. 4904–4916. PMLR, 2021. [1](#), [2](#)
- Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. *arXiv preprint arXiv:2203.12119*, 2022. [3](#)
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020. [2](#)
- Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. *arXiv preprint arXiv:2202.10054*, 2022. [3](#)

- Jungsoo Lee, Eungyeup Kim, Juyoung Lee, Jihyeon Lee, and Jaegul Choo. Learning debiased representation via disentangled feature augmentation. *Advances in Neural Information Processing Systems*, 34:25123–25133, 2021. 3
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 2, 3, 4, 5
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34, 2021a. 1, 2
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*, 2022. 2
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 2, 3, 4, 5
- 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*, 2021b. 2
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021a. 3, 4
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021b. 5
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- Thomas Mensink, Efstratios Gavves, and Cees GM Snoek. Costa: Co-occurrence statistics for zero-shot classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2441–2448, 2014. 2
- Christian M Meyer and Iryna Gurevych. *Wiktionary: A new rival for expert-built lexicons? Exploring the possibilities of collaborative lexicography*. na, 2012. 2
- George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11): 39–41, 1995. 2, 15
- Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. *arXiv preprint arXiv:2112.12750*, 2021. 2
- Guanghui Qin and Jason Eisner. Learning how to ask: Querying lms with mixtures of soft prompts. *arXiv preprint arXiv:2104.06599*, 2021. 3, 4
- 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*, pp. 8748–8763. PMLR, 2021. 1, 2, 3, 5, 6
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pp. 5389–5400. PMLR, 2019. 9
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2556–2565, 2018. 5

- Sheng Shen, Chunyuan Li, Xiaowei Hu, Yujia Xie, Jianwei Yang, Pengchuan Zhang, Anna Rohrbach, Zhe Gan, Lijuan Wang, Lu Yuan, et al. K-lite: Learning transferable visual models with external knowledge. *arXiv preprint arXiv:2204.09222*, 2022. 2, 15
- Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29, 2016. 4
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 8, 16, 17
- Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *Advances in Neural Information Processing Systems*, pp. 10506–10518, 2019. 3, 9
- Xiaolong Wang, Yufei Ye, and Abhinav Gupta. Zero-shot recognition via semantic embeddings and knowledge graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6857–6866, 2018. 2
- Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7959–7971, 2022. 3
- Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 69–77, 2016. 2
- Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning-the good, the bad and the ugly. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4582–4591, 2017. 2
- Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Bin Xiao, Ce Liu, Lu Yuan, and Jianfeng Gao. Unified contrastive learning in image-text-label space. *arXiv preprint arXiv:2204.03610*, 2022a. 1, 2, 4, 5, 6, 12, 13, 15
- Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. *arXiv preprint arXiv:2202.10401*, 2022b. 2
- Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021. 2
- Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022. 1
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *arXiv preprint arXiv:2109.01134*, 2021. 2, 3, 5
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. *arXiv preprint arXiv:2203.05557*, 2022. 3

## A APPENDIX

## B EXPERIMENTAL DETAILS

**Dataset.** Table 6 describes the statistics of dataset used for evaluation. We pick the test datasets based on UniCL (Yang et al., 2022a) and availability in [Tensorflow dataset](#). We use the test set to evaluate zero-shot recognition and linear probe while the train set is used to train a linear classifier. Note that since EuroSAT and Resisc45 utilize the training split for evaluation, we exclude the

Table 6: Statistics of datasets used in zero-shot and linear probe.

Abbreviation	Dataset	#Concepts	Train size	Test size	Source link
Food	Food-101	102	75,750	25,250	<a href="#">Tensorflow</a>
CF10	CIFAR-10	10	50,000	10,000	<a href="#">Tensorflow</a>
CF100	CIFAR-100	100	50,000	10,000	<a href="#">Tensorflow</a>
VOC	VOC2007 classification	20	5,011	4,952	<a href="#">Tensorflow</a>
DTD	Describable Textures	47	3,760	1,880	<a href="#">Tensorflow</a>
Pets	Oxford-IIIT Pets	37	3,680	3,669	<a href="#">Tensorflow</a>
Cal	Caltech-101	102	3,060	6084	<a href="#">Tensorflow</a>
Flower	Oxford Flowers 102	102	1,020	6,149	<a href="#">Tensorflow</a>
Patch	PatchCamelyon	2	294,912	32,768	<a href="#">Tensorflow</a>
ESTAT	EuroSAT	10	N/A	27,000	<a href="#">Tensorflow</a>
R45	Resisc45	45	N/A	31,500	<a href="#">Tensorflow</a>

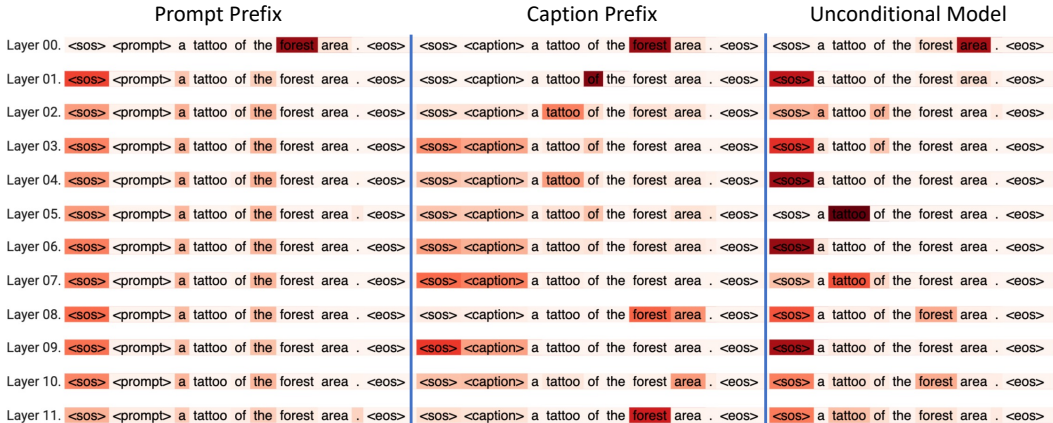


Figure 4: Attention visualization for a *class prompt*. Note that the attention weights are for and end token. Best viewed in color. The class name shown here is one of class prompts in the EUROSAT dataset. Different rows show the weights of different transformer layers. With a prompt prefix (leftmost), the model focuses on a class name (*forest area*) while caption prefix (middle) allows a model to pay attention to another noun, *tattoo*. By prefix conditioning, the attention of the model changes as intended.

two datasets from linear probe evaluation. Also, since Oxford Flowers do not have many training samples (10 samples per class), we exclude the dataset from the evaluation too.

**Data Augmentation.** Following UniCL (Yang et al., 2022a), only random cropping is applied to train all models for a fair comparison.

**Computation.** We use 32 Nvidia Tesla V100 GPUs to train all models. 4 nodes, where each node has 8 GPUs, are used to run experiments.

## C ADDITIONAL RESULTS

**Attention Visualization.** Fig. 4 visualizes attention weights for the class *forest area*, where a prompt template, *a tattoo of*, is employed. The model focuses on a word, *forest* when prompt prefix is employed. In other two cases, the model also pays much attention to *tattoo* probably because the word should provide useful information to distinguish a sentence from others for image-caption contrastive learning. Fig. 5 represents attention for a real caption from CC3M. While the model conditioned with caption prefix and unconditional model attend to several words through many layers, the model conditioned with prompt prefix shows clear attention only in the first layer. Since





Figure 5: Attention visualization for a *real caption*. Note that the attention weights are for and end token. Best viewed in color. The sentence shown here is from CC3M. Different rows show the weights of different transformer layers. Caption prefix conditioning helps to attend to many words while prompt conditioning fails to do that.

the prompt-conditioned model has never seen the real caption during training, it fails in attending to discriminative words.

Table 7: Evaluation on the robustness to the class name shift using ImageNet-1K. *Original* refers to the subset of ImageNet-1K classes while *synonym* refers to their synonyms taken from Wordnet. The last two rows indicate the models trained with the synonyms, thus showing superior performance on *synonym* whereas degrading performance on *Original*.

Train Data	Train on Synonym	Prefix Training	Test-Time Prefix	Original		Synonym	
				top-1	top-5	top-1	top-5
IN1K + CC12M			N/A	69.3	89.3	31.2	49.5
IN1K + CC12M		✓	Prompt	<b>75.0</b>	<b>92.9</b>	<b>38.3</b>	54.8
IN1K + CC12M		✓	Caption	71.4	91.6	36.6	<b>56.7</b>
IN21K + CC12M			N/A	54.5	83.2	23.1	43.9
IN21K + CC12M		✓	Prompt	<b>69.9</b>	<b>92.4</b>	32.1	53.7
IN21K + CC12M		✓	Caption	65.3	90.6	<b>33.5</b>	<b>56.9</b>
IN21K + CC12M	✓	✓	Prompt	54.4	78.6	<b>70.8</b>	<b>92.8</b>
IN21K + CC12M	✓	✓	Caption	<b>54.5</b>	<b>82.6</b>	59.0	86.1

Table 8: Image-text retrieval results on CC3M and COCO. The performance is evaluated on the subset of CC3M and validation set of COCO. All models are trained on CC12M and ImageNet-21K. Caption conditioning (last row) slightly improves retrieval performance compared to the unconditional model (first row). Since prompt conditioning (middle) tailors a model for class-prompt, it fails to extract discriminative information from real captions.

Prefix Training	Test-time Prefix	CC3M				COCO			
		I2T@1	I2T@5	T2I@1	T2I@5	I2T@1	I2T@5	T2I@1	T2I@5
	N/A	21.8	47.4	21.0	45.7	23.9	49.5	18.7	43.2
✓	Prompt	13.1	31.3	8.1	21.8	17.2	38.1	16.8	37.7
✓	Caption	<b>22.6</b>	<b>47.5</b>	<b>21.6</b>	<b>46.1</b>	<b>24.7</b>	<b>49.7</b>	<b>19.7</b>	<b>43.9</b>

**Class Name Shift.** Test samples can be unseen with respect to image classification data in two ways (or combinations of two): 1) The image is similar to training distribution, but the class name used for testing is different from the image classification label. 2) Although the class label is the same, the image data comes from the different distributions. Datasets evaluated in the zero-shot recognition include both two cases since class names and images are from different domains. 2) is analyzed in Subsection 4.3 of the main paper, *Robustness in image domain shift*. We analyze 1) by evaluating the recognition performance of ImageNet-1K by changing its class name from the

one used during training. We find a synonym for each class with WordNet (Miller, 1995), where we exclude synonyms substantially similar to the original class name and obtain synonyms for 525 classes. Then, we use the synonym to classify images during evaluation. Since the input image distribution does not vary, we can evaluate the performance on the class name shift. If the model is robust to the change in the class name, the degrade in the performance should be small.

The first 6 rows of Table 7 describe the models trained with the original class names and evaluated on both original ones and synonyms, and the last two rows represent a model trained with synonyms, where the original class names are replaced with synonyms. Prompt prefix outperforms caption prefix with a large margin in testing with class names used in training time. Generally, caption prefix performs better when tested with the class names different from the ones used during training. Prompt prefix is tailored to handle class names employed during training time while caption prefix enables the language encoder to extract more general representations.

Interestingly, the choice of class names seems to significantly change the generalization as shown in the comparison between a model trained with synonyms and original class names. The original model decreases the accuracy more than 30% by changing the class name while the model trained with synonym decreases less than 20%.

**Image-Caption Retrieval.** In Table 8, we evaluate the performance of image-caption retrieval using the subset of CC3M (12288 pairs of image and caption) and COCO validation set (5000 pairs of image and caption), where all models are trained with CC12M and ImageNet-21K. First, our model (last row) slightly performs better than the model without conditioning (first row). Second, prompt prefix conditioning (second row) significantly performs worse than caption prefix conditioning (last row). Since the prompt prefix conditioning specializes a model for the class name prompts of ImageNet21K, the conditioning does not generalize well to real captions.

**Larger Batch-size and Training Epochs.** We examine the effect of increasing batch-size and training epochs in Table 9. In CLIP, increasing the batch-size and training epochs improves the performance of both ImageNet-1K and zero-shot recognition. On the other hand, the zero-shot performance of UniCL is not benefited from training with longer epochs (compare last and second to last row). UniCL attempts to ensure the invariance of images from the same classes by supervised contrastive loss while CLIP does not consider it. However, such invariance is not necessarily required in zero-shot recognition, which leads to the degraded performance.

**Comparison to Reported UniCL’s Results.** In the main paper, we provide our reproduced results of UniCL, which is based on our implementation, since the authors have not released the code and did not report the numerical accuracy of each zero-shot recognition. In this paragraph, we compare our approach and the reported performance of UniCL (Yang et al., 2022a) and K-Lite (Shen et al., 2022) by aligning several hyper-parameters, e.g., batch-size and training epochs, using ImageNet-1K. When using ImageNet-22K and CC-15M for training, our method (batch-size:4096, training epochs: 30) shows 73.9 while UniCL (batch-size:4096, training epochs 32) reports 71.5. When using ImageNet-21K excluding ImageNet-1K and CC-15M, our method (batch-size:1024, training epochs 30) shows 49.7 whereas UniCL (batch-size: 4096, training epochs: 32) and K-Lite (batch-size: 4096, training epochs: 32) perform 46.6 and 48.7 respectively according to K-Lite results (See last two rows of Table 3 in (Shen et al., 2022)). These results suggest that our method performs better than the reported numbers of UniCL and K-Lite in ImageNet-1K. Also, the knowledge augmentation technique proposed by K-Lite can be complementary to our approach, thus combining two approaches is an interesting research direction.

**T-SNE visualization for language features.** Fig. 6 visualizes extracted language features (ImageNet-1K) conditioned with different prefixes. The prompt-prefix (left) has lower intra-class and higher inter-class variance, whereas caption-prefix (right) shows higher intra-class variance across prompts.

**T-SNE visualization for image features.** Fig. 7 visualizes image features from ImageNet-1K (blue) and CC3M (red). Since ImageNet-1K is object-centered while CC3M covers more diverse scenes, the distributions are separated. This is consistent across baseline (w/o conditioning) and our method (with conditioning).

**Comparison between unconditioned and conditioned model by language features.** Fig. 8 visualizes language features of ImageNet-1K class prompts (Blue) and CC3M captions (Red) for

Table 9: Performance comparison among different batch-size and training epochs. ImageNet-22K denotes the combination of ImageNet-21K and ImageNet-1K, CC-15M indicates that of CC-12M and CC-3M.

Training Data		Objective	Batch-size	Epochs	Metric	
Classification	Caption				IN-1K	Zero-shot 11 datasets
ImageNet-21K	CC-12M	CLIP	1024	15	67.3	57.8
ImageNet-21K	CC-12M	CLIP	1024	30	<b>69.1</b>	<b>58.3</b>
ImageNet-22K	CC-15M	CLIP	1024	15	69.3	58.5
ImageNet-22K	CC-15M	CLIP	4096	15	71.1	59.5
ImageNet-22K	CC-15M	CLIP	4096	30	<b>72.2</b>	<b>59.8</b>
ImageNet-22K	CC-15M	UniCL	1024	15	69.7	58.5
ImageNet-22K	CC-15M	UniCL	4096	15	70.3	<b>60.4</b>
ImageNet-22K	CC-15M	UniCL	4096	30	<b>73.9</b>	58.9




Figure 6: T-SNE [Van der Maaten & Hinton \(2008\)](#) visualization of the class-prompt features of ImageNet-1K with different prefix conditions. Different colors indicate language embeddings of different classes. Prompt conditioning extracts more class discriminative representations than caption conditioning.

unconditioned (left) and conditioned (right) respectively. Note that the conditioned model utilizes prompt prefix for class prompts and caption prefix for real captions respectively. As seen from the visualization, unconditioned model cannot distinguish some prompts from captions of CC3M. This is probably because some captions are similar to class prompts of ImageNet. By contrast, the conditioned model differentiate class prompts from captions better than unconditioned model due to the prefix conditioning.

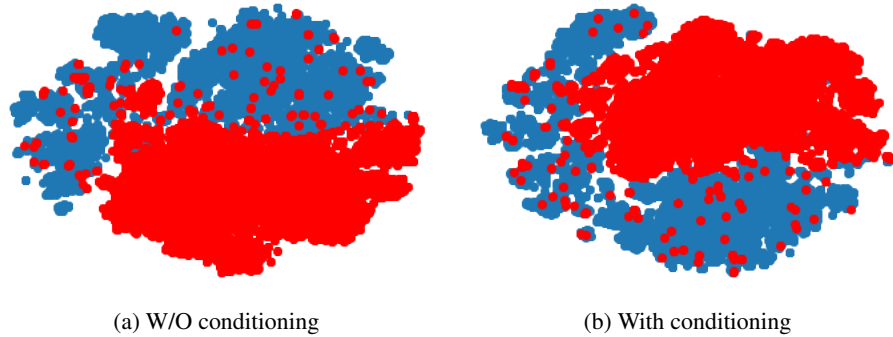


Figure 7: T-SNE [Van der Maaten & Hinton \(2008\)](#) visualization of the image features of ImageNet-1K (blue) and CC3M (red). Since ImageNet-1K is object-centered while CC3M covers more diverse scenes, the distributions are separated. This is consistent across baseline (w/o conditioning) and our method (with conditioning).

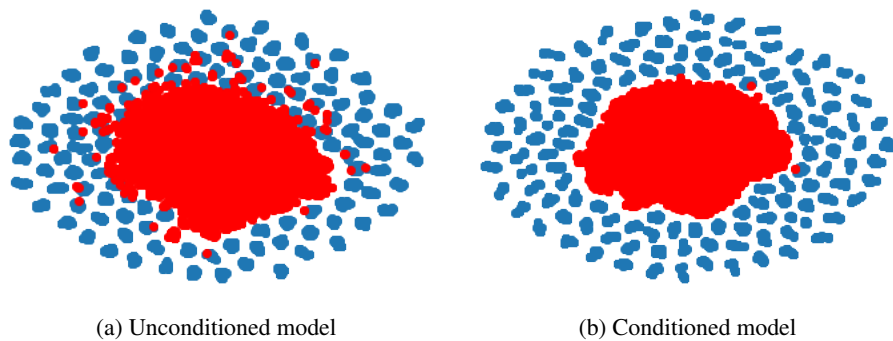


Figure 8: T-SNE [Van der Maaten & Hinton \(2008\)](#) visualization of language features of ImageNet-1K class prompts (Blue) and CC3M captions (Red) for unconditioned (left) and conditioned (right) respectively. Our proposed condition better differentiates prompts from real captions.