

MaPLe: Multi-modal Prompt Learning

Muhammad Uzair Khattak¹ Hanoona Rasheed¹ Muhammad Maaz¹
Salman Khan^{1,2} Fahad Shahbaz Khan^{1,3}

¹Mohamed bin Zayed University of AI ²Australian National University ³Linköping University

Abstract

Pre-trained vision-language (V-L) models such as CLIP have shown excellent generalization ability to downstream tasks. However, they are sensitive to the choice of input text prompts and require careful selection of prompt templates to perform well. Inspired by the Natural Language Processing (NLP) literature, recent CLIP adaptation approaches learn prompts as the textual inputs to fine-tune CLIP for downstream tasks. We note that using prompting to adapt representations in a single branch of CLIP (language or vision) is sub-optimal since it does not allow the flexibility to dynamically adjust both representation spaces on a downstream task. In this work, we propose Multi-modal Prompt Learning (MaPLe) for both vision and language branches to improve alignment between the vision and language representations. Our design promotes strong coupling between the vision-language prompts to ensure mutual synergy and discourages learning independent uni-modal solutions. Further, we learn separate prompts across different early stages to progressively model the stage-wise feature relationships to allow rich context learning. We evaluate the effectiveness of our approach on three representative tasks of generalization to novel classes, new target datasets and unseen domain shifts. Compared with the state-of-the-art method Co-CoOp, MaPLe exhibits favorable performance and achieves an absolute gain of 3.45% on novel classes and 2.72% on overall harmonic-mean, averaged over 11 diverse image recognition datasets. Our code and pre-trained models are available at <https://github.com/muzairkhattak/multimodal-prompt-learning>.

1. Introduction

Foundational vision-language (V-L) models such as CLIP (Contrastive Language-Image Pretraining) [32] have shown excellent generalization ability to downstream tasks. Such models are trained to align language and vision modalities on web-scale data *e.g.*, 400 million text-image pairs in CLIP. These models can reason about open-vocabulary visual concepts, thanks to the rich supervision provided by

natural language. During inference, hand-engineered text prompts are used *e.g.*, ‘a photo of a <category>’ as a query for text encoder. The output text embeddings are matched with the visual embeddings from an image encoder to predict the output class. Designing high quality contextual prompts have been proven to enhance the performance of CLIP and other V-L models [17, 42].

Despite the effectiveness of CLIP towards generalization to new concepts, its massive scale and scarcity of training data (*e.g.*, few-shot setting) makes it infeasible to fine-tune the full model for downstream tasks. Such fine-tuning can also forget the useful knowledge acquired in the large-scale pretraining phase and can pose a risk of overfitting to the downstream task. To address the above challenges, existing works propose language prompt learning to avoid manually adjusting the prompt templates and providing a mechanism to adapt the model while keeping the original weights frozen [14, 25, 29, 48, 49]. Inspired from Natural Language Processing (NLP), these approaches only explore prompt learning for the text encoder in CLIP (Fig. 1:a) while adaptation choices together with an equally important image encoder of CLIP remains an unexplored topic in the literature.

Our motivation derives from the multi-modal nature of CLIP, where a text and image encoder co-exist and *both* contribute towards properly aligning the V-L modalities. We argue that any prompting technique should adapt the model completely and therefore, learning prompts only for the text encoder in CLIP is not sufficient to model the adaptations needed for the image encoder. To this end, we set out to achieve completeness in the prompting approach and propose **Multi-modal Prompt Learning (MaPLe)** to adequately fine-tune the text and image encoder representations such that their optimal alignment can be achieved on the downstream tasks (Fig. 1:b). Our extensive experiments on three key representative settings including base-to-novel generalization, cross-dataset evaluation, and domain generalization demonstrate the strength of MaPLe. On base-to-novel generalization, our proposed MaPLe outperforms existing prompt learning approaches across 11 diverse image recognition datasets (Fig. 1:c) and achieves absolute average gain of 3.45% on novel classes and 2.72% on harmonic-mean

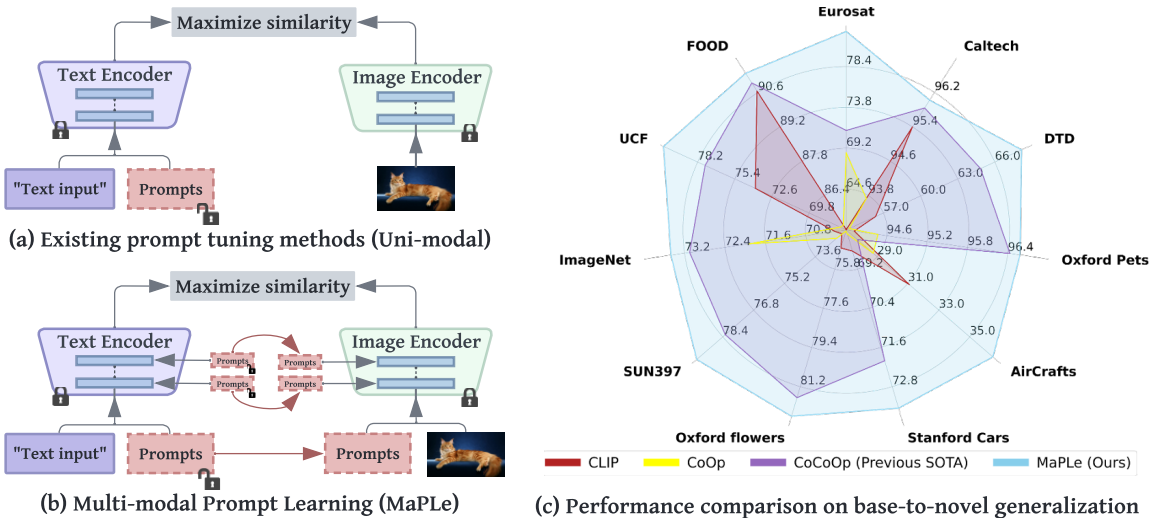


Figure 1. Comparison of MaPLE with standard prompt learning methods. **(a)** Existing methods adopt uni-modal prompting techniques to fine-tune CLIP representations as prompts are learned only in a single branch of CLIP (language or vision). **(b)** MaPLE introduces branch-aware hierarchical prompts that adapt both language and vision branches simultaneously for improved generalization. **(c)** MaPLE surpasses state-of-the-art methods on 11 diverse image recognition datasets for novel class generalization task.

over the state-of-the-art method Co-CoOp [48]. Further, MaPLE demonstrates favorable generalization ability and robustness in cross-dataset transfer and domain generalization settings, leading to consistent improvements compared to existing approaches. Owing to its streamlined architectural design, MaPLE exhibits improved efficiency during both training and inference without much overhead, as compared to Co-CoOp which lacks efficiency due to its image instance conditioned design. In summary, the main contributions of this work include:

- We propose *multi-modal* prompt learning in CLIP to favourably align its vision-language representations. To the best of our knowledge, this is the first multi-modal prompting approach for fine-tuning CLIP.
- To link prompts learned in text and image encoders, we propose a *coupling function* to explicitly condition vision prompts on their language counterparts. It acts as a bridge between the two modalities and allows mutual propagation of gradients to promote synergy.
- Our multi-modal prompts are learned across multiple transformer blocks in both vision and language branches to *progressively* learn the synergistic behaviour of both modalities. This deep prompting strategy allows modeling the contextual relationships independently, thus providing more flexibility to align the vision-language representations.

2. Related Work

Vision Language Models: The combined use of language supervision with natural images is found to be of great in-

terest in the computer vision community. In contrast to models learned with only image supervision, these vision-language (V-L) models encode rich multimodal representations. Recently, V-L models like CLIP [32], ALIGN [15], LiT [45], FILIP [41] and Florence [43] have demonstrated exceptional performance on a wide spectrum of tasks including few-shot and zero-shot visual recognition. These models learn joint image-language representations in a self-supervised manner using abundantly available data from the web. For example, CLIP and ALIGN respectively use $\sim 400\text{M}$ and $\sim 1\text{B}$ image-text pairs to train a multi-modal network. Although these pre-trained V-L models learn generalized representations, efficiently adapting them to downstream tasks is still a challenging problem. Many works have demonstrated better performance on downstream tasks by using tailored methods to adapt V-L models for few-shot image-recognition [9, 19, 46], object detection [8, 10, 27, 34, 44, 50], and segmentation [5, 22, 26, 33]. In this work, we propose a novel multi-modal prompt learning technique to effectively adapt CLIP for few-shot and zero-shot visual recognition tasks.

Prompt Learning: The instructions in the form of a sentence, known as text prompt, are usually given to the language branch of a V-L model, allowing it to better understand the task. Prompts can be handcrafted for a downstream task or learned automatically during fine-tuning stage. The latter is referred to as ‘Prompt Learning’ which was first used in NLP [21, 23, 24] followed by the adaptation in V-L [48, 49, 51] and vision-only [16, 38, 39, 47] models. Similar to [16] our design also uses deep ‘vision’ prompting. However, ours is the first multi-modal prompting de-

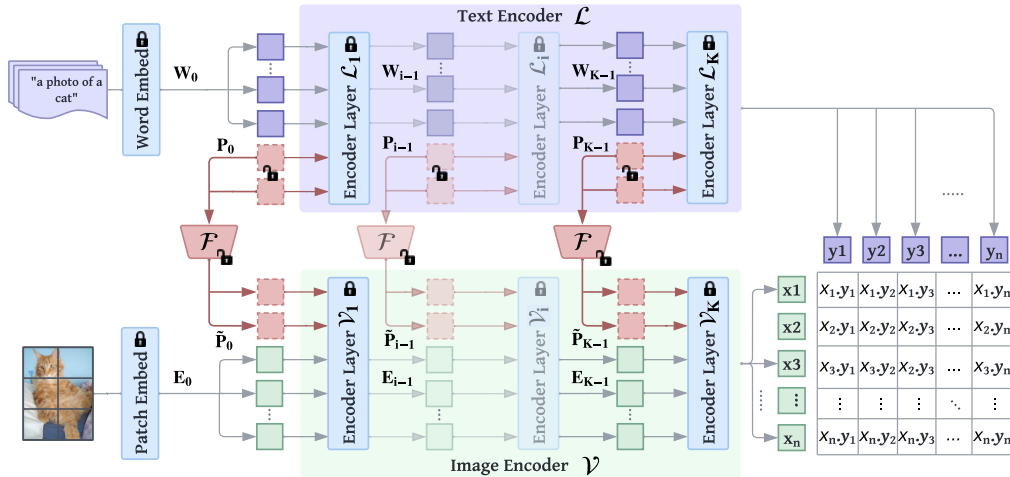


Figure 2. Overview of our proposed MaPLE (Multi-modal Prompt Learning) framework for prompt learning in V-L models. MaPLE tunes both vision and language branches where only the **context prompts** are learned, while the rest of the model is frozen. MaPLE conditions the vision prompts on language prompts via a V-L coupling function \mathcal{F} to induce mutual synergy between the two modalities. Our framework uses deep contextual prompting where separate context prompts are learned across multiple transformer blocks.

sign while [16] is uni-modal.

Prompt Learning in Vision Language models: Full fine-tuning and linear probing [9] are two typical approaches to adapt a V-L model (*i.e.* CLIP) to the downstream tasks. The complete fine-tuning results in degrading the previously learned joint V-L representation while linear probing limits the zero-shot capability of CLIP. To this end, inspired from prompt learning in NLP, many works have proposed to adapt V-L models by learning the prompt tokens in an end-to-end training. CoOp [49] fine-tunes CLIP for few-shot transfer by optimizing continuous set of prompt vectors at its language branch. Co-CoOp [48] highlights the inferior performance of CoOp on novel classes and solves the generalization issue by explicitly conditioning prompts on image instances. [25] proposes to optimize multiple set of prompts by learning the distribution of prompts. [18] adapt CLIP by learning prompts for video understanding tasks. [1] perform visual prompt tuning on CLIP by prompting on the vision branch. We note that the existing methods follow independent *uni-modal* solutions and learn prompts either in the language or in the vision branch of CLIP, thus adapting CLIP partially. In this paper, we explore an important question: given the multimodal nature of CLIP, is complete prompting (*i.e.*, in both language and vision branches) better suited to adapt CLIP? Our work is the first to answer this question by investigating the effectiveness of multi-modal prompt learning in order to improve alignment between vision and language representations.

3. Method

Our approach concerns with fine-tuning a pre-trained multi-modal CLIP for better generalization to downstream tasks

through context optimization via prompting. Fig. 2 shows the overall architecture of our proposed MaPLE (Multi-modal Prompt Learning) framework. Unlike previous approaches [48, 49] which learn context prompts only at the language branch, MaPLE proposes a joint prompting approach where the context prompts are learned in both vision and language branches. Specifically, we append learnable context tokens in the language branch and explicitly condition the vision prompts on the language prompts via a coupling function to establish interaction between them. To learn hierarchical contextual representations, we introduce deep prompting in both branches through separate learnable context prompts across different transformer blocks. During fine-tuning, only the context prompts along with their coupling function are learned while the rest of the model is frozen. Below, we first outline the pre-trained CLIP architecture and then present our proposed fine-tuning approach.

3.1. Revisiting CLIP

We build our approach on a pre-trained vision-language (V-L) model, CLIP, which consists of a text and vision encoder. Consistent with existing prompting methods [48, 49], we use a vision transformer (ViT) [6] based CLIP model. CLIP encodes an image $I \in \mathbb{R}^{H \times W \times 3}$ and a corresponding text description as explained below.

Encoding Image: Image encoder \mathcal{V} with K transformer layers $\{\mathcal{V}_i\}_{i=1}^K$, splits the image I into M fixed-size patches which are projected into patch embeddings $E_0 \in \mathbb{R}^{M \times d_v}$. Patch embeddings E_i are input to the $(i+1)^{\text{th}}$ transformer block (\mathcal{V}_{i+1}) along with a learnable class (CLS) token c_i




Figure 3. t-SNE plots of image embeddings in uni-modal prompting method Co-CoOp, and MaPLE on 3 diverse image recognition datasets. MaPLE shows better separability in both base and novel classes.

and sequentially processed through K transformer blocks,

$$[c_i, E_i] = \mathcal{V}_i([c_{i-1}, E_{i-1}]) \quad i = 1, 2, \dots, K.$$

To obtain the final image representation x , the class token c_K of last transformer layer (\mathcal{V}_K) is projected to a common V-L latent embedding space via ImageProj,

$$x = \text{ImageProj}(c_K) \quad x \in \mathbb{R}^{d_{vl}}.$$

Encoding Text: CLIP text encoder generates feature representations for text description by tokenizing the words and projecting them to word embeddings $W_0 = [w_0^1, w_0^2, \dots, w_0^N] \in \mathbb{R}^{N \times d_i}$. At each stage, W_i is input to the $(i + 1)^{\text{th}}$ transformer layer of text encoding branch (\mathcal{L}_{i+1}),

$$[W_i] = \mathcal{L}_i(W_{i-1}) \quad i = 1, 2, \dots, K.$$

The final text representation z is obtained by projecting the text embeddings corresponding to the last token of the last transformer block \mathcal{L}_K to a common V-L latent embedding space via TextProj,

$$z = \text{TextProj}(w_K^N) \quad z \in \mathbb{R}^{d_{vl}}.$$

Zero-shot Classification: For zero-shot classification, text prompts are hand-crafted with class labels $y \in \{1, 2, \dots, C\}$ (e.g., ‘a photo of a <category>’) having C classes. Prediction \hat{y} corresponding to the image I having the highest cosine similarity score ($\text{sim}(\cdot)$) is calculated with a temperature parameter τ ,

$$p(\hat{y}|x) = \frac{\exp(\text{sim}(x, z_{\hat{y}})/\tau)}{\sum_{i=1}^C \exp(\text{sim}(x, z_i))}.$$

3.2. MaPLE: Multi-modal Prompt Learning

To efficiently fine-tune CLIP for downstream image recognition tasks, we explore the potential of *multi-modal* prompt

tuning. We reason that prior works that have predominantly explored uni-modal approaches are less suitable as they do not offer the flexibility to dynamically adapt both language and vision representation spaces. Thus to achieve completeness in prompting, we underline the importance of multi-modal prompting approach. In Fig. 3, we visualize and compare the image embeddings of MaPLE with recent state-of-the-art work, Co-CoOp. Note that the image embeddings of CLIP, CoOp and Co-CoOp will be identical as they do not learn prompts in the vision branch. The visualization shows that image embeddings of MaPLE are more separable indicating that learning vision prompts in addition to language prompts leads to better adaptation of CLIP.

In addition to multi-modal prompting, we find that it is essential to learn prompts in the deeper transformer layers to progressively model stage-wise feature representations. To this end, we propose to introduce learnable tokens in the first J (where $J < K$) layers of both vision and language branches. These multi-modal hierarchical prompts utilize the knowledge embedded in CLIP model to effectively learn task relevant contextual representations (see Fig. 4).

3.2.1 Deep Language Prompting

To learn the language context prompts, we introduce b learnable tokens $\{P^i \in \mathbb{R}^{d_i}\}_{i=1}^b$, in the language branch of CLIP. The input embeddings now follow the form $[P^1, P^2, \dots, P^b, W_0]$, where $W_0 = [w^1, w^2, \dots, w^N]$ corresponds to fixed input tokens. New learnable tokens are further introduced in each transformer block of the language encoder (\mathcal{L}_i) up to a specific depth J ,

$$[_, W_i] = \mathcal{L}_i([P_{i-1}, W_{i-1}]) \quad i = 1, 2, \dots, J. \quad (1)$$

Here $[_, \cdot]$ refers to the concatenation operation. After J^{th} transformer layer, the subsequent layers process previous

layer prompts and final text representation z is computed,

$$[P_j, W_j] = \mathcal{L}_j([P_{j-1}, W_{j-1}]) \quad j = J + 1, \dots, K, \quad (2)$$

$$z = \text{TextProj}(w_K^N). \quad (3)$$

When $J = 1$, the learnable tokens P are only applied at the input of first transformer layer, and this deep language prompting technique degenerates to CoOp [49].

3.2.2 Deep Vision Prompting

Similar to deep language prompting, we introduce b learnable tokens $\{\tilde{P}^i \in \mathbb{R}^{d_v}\}_{i=1}^b$ in the vision branch of CLIP alongside the input image tokens. New learnable tokens are further introduced in deeper transformer layers of the image encoder (\mathcal{V}) up to depth J .

$$[c_i, E_i, _] = \mathcal{V}_i([c_{i-1}, E_{i-1}, \tilde{P}_{i-1}]) \quad i = 1, 2, \dots, J,$$

$$[c_j, E_j, \tilde{P}_j] = \mathcal{V}_j([c_{j-1}, E_{j-1}, \tilde{P}_{j-1}]) \quad j = J + 1, \dots, K,$$

$$x = \text{ImageProj}(c_K).$$

Our deep prompting provides the flexibility to learn prompts across different feature hierarchies within the ViT architecture. We find that sharing prompts across stages is better compared to independent prompts as features are more correlated due to successive transformer block processing. Thus, the later stages do not provide independently-learned complimentary prompts as compared to the early stages.

3.2.3 Vision Language Prompt Coupling

We reason that in prompt tuning it is essential to take a multi-modal approach and *simultaneously* adapt both the vision and language branch of CLIP in order to achieve completeness in context optimization. A simple approach would be to naively combine deep vision and language prompting, where both the language prompts P , and the vision prompts \tilde{P} , will be learned during the same training schedule. We name this design as ‘*Independent V-L Prompting*’. Although this approach satisfies the requirement of completeness in prompting, this design lacks synergy between vision and language branch as both branches do not interact while learning the task relevant context prompts.

To this end, we propose a branch-aware multi-modal prompting which tunes vision and language branch of CLIP together by sharing prompts across both modalities. Language prompt tokens are introduced in the language branch up to J^{th} transformer block similar to deep language prompting as illustrated in Eqs. 1-3. To ensure mutual synergy between V-L prompts, vision prompts \tilde{P} , are obtained by projecting language prompts P via vision-to-language projection which we refer to as *V-L coupling function* $\mathcal{F}(\cdot)$, such that $\tilde{P}_k = \mathcal{F}_k(P_k)$. The coupling function is implemented as a linear layer which maps d_l dimensional inputs

to d_v . This acts as a bridge between the two modalities, thus encouraging mutual propagation of gradients.

$$[c_i, E_i, _] = \mathcal{V}_i([c_{i-1}, E_{i-1}, \mathcal{F}_{i-1}(P_{i-1})]) \quad i = 1, \dots, J$$

$$[c_j, E_j, \tilde{P}_j] = \mathcal{V}_j([c_{j-1}, E_{j-1}, \tilde{P}_{j-1}]) \quad j = J + 1, \dots, K$$

$$x = \text{ImageProj}(c_K)$$

Unlike independent V-L prompting, explicit conditioning of \tilde{P} on P helps learn prompts in a shared embedding space between the two branches, thus improving mutual synergy.

4. Experiments

4.1. Benchmark setting

Generalization from Base-to-Novel Classes: We evaluate the generalizability of MaPLe, and follow a zero-shot setting where the datasets are split into base and novel classes. The model is trained only on the base classes in a few-shot setting and evaluated on base and novel categories.

Cross-dataset Evaluation: To validate the potential of our approach in cross-dataset transfer, we evaluate our ImageNet trained model directly on other datasets. Consistent with Co-CoOp, our model is trained on all 1000 ImageNet classes in a few-shot manner.

Domain Generalization: We evaluate the robustness of our method on out-of-distribution datasets. Similar to cross-dataset evaluation, we test our ImageNet trained model directly on four other ImageNet datasets that contain various types of domain shifts.

Datasets: For generalization from base-to-novel classes and cross-dataset evaluation, we follow [48, 49] and evaluate the performance of our method on 11 image classification datasets which covers a wide range of recognition tasks. This includes two generic-objects datasets, ImageNet [4] and Caltech101 [7]; five fine-grained datasets, OxfordPets [31], StanfordCars [20], Flowers102 [30], Food101 [2], and FGVC Aircraft [28]; a scene recognition dataset SUN397 [40]; an action recognition dataset UCF101 [36]; a texture dataset DTD [3] and a satellite-image dataset EuroSAT [11]. For domain generalization, we use ImageNet as source dataset and its four variants as target datasets including ImageNetV2 [35], ImageNet-Sketch [37], ImageNet-A [13] and ImageNet-R [12].

Implementation Details We use a few-shot training strategy in all experiments at 16 shots which are randomly sampled for each class. We apply prompt tuning on a pre-trained ViT-B/16 CLIP model where $d_l = 512$, $d_v = 768$ and $d_{vl} = 512$. For MaPLe, we set prompt depth J to 9 and the language and vision prompt lengths to 2. All models are trained for 5 epochs with a batch-size of 4 and a learning rate of 0.0035 via SGD optimizer on a single NVIDIA A100 GPU. We report base and novel class accuracies and their harmonic mean (HM) averaged over 3 runs. We initial-

Method	Base Acc.	Novel Acc.	HM	GFLOPS
1: MaPLE shallow ($J = 1$)	80.10	73.52	76.67	167.1
2: Deep vision prompting	80.24	73.43	76.68	18.0
3: Deep language prompting	81.72	73.81	77.56	166.8
4: Independent V-L prompting	82.15	74.07	77.90	167.0
5: MaPLE (Ours)	82.28	75.14	78.55	167.0

Table 1. Comparison of MaPLE with different prompting designs in base-to-novel generalization. Results are averaged over 11 datasets. HM refers to harmonic mean.

ize the language prompts of the first layer P_0 with the pre-trained CLIP word embeddings of the template ‘a photo of a <category>’, while for the subsequent layers they are randomly initialized from a normal distribution. For training MaPLE on all 1000 classes of ImageNet as a source model, prompt depth J is set to 3 and the model trained for 2 epochs with learning rate of 0.0026. Hyper-parameters for deep language prompting, deep vision prompting, and independent V-L prompting are detailed in Appendix A. The hyper-parameters are fixed across all datasets.

4.2. Prompting CLIP via Vision-Language Prompts

Prompting Variants: We first evaluate the performance of different possible prompting design choices as an ablation for our proposed branch-aware multi-modal prompting, MaPLE. These variants include shallow MaPLE, deep language prompting, deep vision prompting and independent V-L prompting. In Table 1, we present the results averaged over the 11 image recognition datasets. Shallow MaPLE (row-1) provides consistent improvements over CoOp and Co-CoOp in terms of generalization. Deep language prompting (row-3) shows improvements over deep vision prompting (row-2), indicating that prompts learned at the language branch provide better adaptation of CLIP. Although separately combining the above two approaches (row-4) further improves the performance, it struggles to achieve comprehensive benefits from the language and vision branches. We hypothesize that this is due to the lack of synergy between the learned vision and language prompts as they do not interact with each other during training. Meanwhile, MaPLE tied with deep prompting (row-4) combines the benefits of prompting in both branches by enforcing interactions through explicit conditioning of vision prompts on the language prompts. It provides improvements on novel and base class accuracies which leads to the best HM of 78.55%. We explore other possible design choices and present the ablations in Appendix B.

4.3. Base-to-Novel Generalization

Generalization to Unseen Classes: Table 3 presents the performance of MaPLE in base-to-novel generalization setting on 11 recognition datasets. We compare its performance with CLIP zero-shot, and recent prompt learning works including CoOp [49] and Co-CoOp [48]. In case of

CLIP, we use hand-crafted prompts that are specifically designed for each dataset.

In comparison with the state-of-the-art Co-CoOp, MaPLE shows improved performance on both base and novel categories on all 11 datasets with an exception of marginal reduction on only the base class performance of Caltech101. With mutual synergy from the branch-aware multi-modal prompting, MaPLE better generalizes to novel categories on all 11 datasets in comparison with Co-CoOp, and obtains an overall gain from 71.69% to 75.14%. When taking into account both the base and novel classes, MaPLE shows an absolute average gain of 2.72% over Co-CoOp.

In comparison with CLIP on novel classes, Co-CoOp improves only on 4/11 datasets dropping the average novel accuracy from 74.22% to 71.69%. MaPLE is a strong competitor which improves accuracy over CLIP on novel classes on 6/11 datasets, with an average gain from 74.22% to 75.14%.

Generalization and Performance on Base Classes: Co-CoOp solves the poor generalization problem in CoOp by conditioning prompts on image instances and shows significant gains in novel categories. However on base classes, it improves over CoOp only on 3/11 datasets with an average drop in performance from 82.69% to 80.47%. Meanwhile, the completeness in prompting helps MaPLE improve over CoOp on base classes in 6/11 datasets maintaining the average base accuracy to around 82.28%, in addition to its improvement in generalization to novel classes.

We find that the training strategies of Co-CoOp can be used to substantially boost the generalization performance of vanilla CoOp (6.8% gain in novel classes). We therefore compare our method with CoOp[†], which trains CoOp in Co-CoOp setting (refer to Appendix A for more details).

	Base	Novel	HM
CoOp	82.69	63.22	71.66
Co-CoOp	80.47	71.69	75.83
CoOp [†]	80.85	70.02	75.04
MaPLE	82.28	75.14	78.55

Table 2. Generalization comparison of MaPLE with CoOp[†].

Compare to CoOp[†], the vanilla CoOp model seems to overfit on base classes. When compared to CoOp[†] which attains an average base accuracy of 80.85%, MaPLE shows an improvement of 1.43% with the average base accuracy of 82.28% (Table 2).

4.4. Cross-Dataset Evaluation

We test the cross-dataset generalization ability of MaPLE by learning multi-modal prompts on all the 1000 ImageNet classes and then transferring it directly on the remaining 10 datasets. Table 4 shows the performance comparison between MaPLE, CoOp and Co-CoOp. On the ImageNet source dataset, MaPLE achieves performance comparable

(a) Average over 11 datasets				(b) ImageNet.				(c) Caltech101			
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CLIP	69.34	74.22	71.70	CLIP	72.43	68.14	70.22	CLIP	96.84	94.00	95.40
CoOp	82.69	63.22	71.66	CoOp	76.47	67.88	71.92	CoOp	98.00	89.81	93.73
Co-CoOp	80.47	71.69	75.83	Co-CoOp	75.98	70.43	73.10	Co-CoOp	97.96	93.81	95.84
MaPLe	82.28	75.14	78.55	MaPLe	76.66	70.54	73.47	MaPLe	97.74	94.36	96.02
	+1.81	+3.45	+2.72		+0.68	+0.11	+0.37		-0.22	+0.55	+0.18
(d) OxfordPets				(e) StanfordCars				(f) Flowers102			
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CLIP	91.17	97.26	94.12	CLIP	63.37	74.89	68.65	CLIP	72.08	77.80	74.83
CoOp	93.67	95.29	94.47	CoOp	78.12	60.40	68.13	CoOp	97.60	59.67	74.06
Co-CoOp	95.20	97.69	96.43	Co-CoOp	70.49	73.59	72.01	Co-CoOp	94.87	71.75	81.71
MaPLe	95.43	97.76	96.58	MaPLe	72.94	74.00	73.47	MaPLe	95.92	72.46	82.56
	+0.23	+0.07	+0.15		+2.45	+0.41	+1.46		+1.05	+0.71	+0.85
(g) Food101				(h) FGVC Aircraft				(i) SUN397			
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CLIP	90.10	91.22	90.66	CLIP	27.19	36.29	31.09	CLIP	69.36	75.35	72.23
CoOp	88.33	82.26	85.19	CoOp	40.44	22.30	28.75	CoOp	80.60	65.89	72.51
Co-CoOp	90.70	91.29	90.99	Co-CoOp	33.41	23.71	27.74	Co-CoOp	79.74	76.86	78.27
MaPLe	90.71	92.05	91.38	MaPLe	37.44	35.61	36.50	MaPLe	80.82	78.70	79.75
	+0.01	+0.76	+0.39		+4.03	+11.90	+8.76		+1.08	+1.84	+1.48
(j) DTD				(k) EuroSAT				(l) UCF101			
	Base	Novel	HM		Base	Novel	HM		Base	Novel	HM
CLIP	53.24	59.90	56.37	CLIP	56.48	64.05	60.03	CLIP	70.53	77.50	73.85
CoOp	79.44	41.18	54.24	CoOp	92.19	54.74	68.69	CoOp	84.69	56.05	67.46
Co-CoOp	77.01	56.00	64.85	Co-CoOp	87.49	60.04	71.21	Co-CoOp	82.33	73.45	77.64
MaPLe	80.36	59.18	68.16	MaPLe	94.07	73.23	82.35	MaPLe	83.00	78.66	80.77
	+3.35	+3.18	+3.31		+6.58	+13.19	+11.14		+0.67	+5.21	+3.13

Table 3. Comparison with state-of-the-art methods on base-to-novel generalization. MaPLe learns multi-modal prompts and demonstrates strong generalization results over existing methods on 11 recognition datasets. Absolute gains over Co-CoOp are indicated in blue.

	Source			Target								
	ImageNet	Caltech101	OxfordPets	StanfordCars	Flowers102	Food101	Aircraft	SUN397	DTD	EuroSAT	UCF101	Average
CoOp	71.51	93.70	89.14	64.51	68.71	85.30	18.47	64.15	41.92	46.39	66.55	63.88
Co-CoOp	71.02	94.43	90.14	65.32	71.88	86.06	22.94	67.36	45.73	45.37	68.21	65.74
MaPLe	70.72	93.53	90.49	65.57	72.23	86.20	24.74	67.01	46.49	48.06	68.69	66.30

Table 4. Comparison of MaPLe with existing approaches on cross-dataset evaluation. Overall, MaPLe achieves competitive performance providing highest average accuracy, indicating better generalization.

to competing approaches but demonstrates a much stronger generalization performance by surpassing CoOp in 9/10 and Co-CoOp in 8/10 datasets. Overall, MaPLe shows competitive performance leading to the highest averaged accuracy of 66.30%. This suggests that the use of branch-aware V-L prompting in MaPLe facilitates better generalization.

4.5. Domain Generalization

We show that MaPLe generalizes favourably on out-of-distribution datasets as compared to CoOp and Co-CoOp. We evaluate the direct transferability of ImageNet trained model to various out-of-domain datasets, and observe that it consistently improves against all the existing approaches

	Source		Target		
	ImageNet	ImageNetV2	ImageNet-S	ImageNet-A	ImageNet-R
CLIP	66.73	60.83	46.15	47.77	73.96
CoOp	71.51	64.20	47.99	49.71	75.21
Co-CoOp	71.02	64.07	48.75	50.63	76.18
MaPLe	70.72	64.07	49.15	50.90	76.98

Table 5. Comparison of MaPLe with existing approaches in domain generalization setting. MaPLe shows consistent improvements on all target datasets.

as indicated in Table 5. This indicates that utilizing multi-modal branch-aware prompting helps MaPLe in enhancing the generalization and robustness of V-L models like CLIP.

4.6. Ablation Experiments

Prompt Depth: In Fig. 4 (left), we illustrate the effect of prompt depth J for MaPLe and ablate on the depth of language and vision branch *individually*. In general, the performance improves as prompt depth increases. We note that performance sensitivity increases when randomly initialized prompts are inserted in the deeper layers of a frozen model where the model feature space is already mature. Similar trend is also reported by [16]. As earlier methods utilize shallow language prompting ($J = 1$), we compare our method with deep language prompting. Overall, MaPLe achieves better performance than deep language prompting and achieves maximum performance at a depth of 9.

Prompt Length: Fig. 4 (right) shows the effect of prompt length for MaPLe. As the prompt length increases, the performance on base classes is generally maintained, while the novel class accuracy decreases. This indicates over-fitting which inherently hurts the generalization to novel classes.

Effectiveness of Multi-modal Prompting: Fig. 5 shows the analysis of per class accuracy for selected datasets in the order of increasing domain shift. It indicates that the performance gains of MaPLe in comparison to Co-CoOp varies across different datasets. MaPLe provides significant gains over Co-CoOp for datasets that have large distribution shifts from the pretraining dataset of CLIP, and vision concepts that are usually rare and less generic. Further detailed analysis is provided in Appendix C.

Prompting complexity: Table 6 shows the computational complexity of MaPLe in comparison with other approaches.




Figure 4. Ablation on prompt depth (*left*) and prompt length (*right*) in MaPLe. We report average results on the held-out validation sets of all datasets.




Figure 5. Percentage classes where MaPLe shows improved performance over Co-CoOp, which increases as dataset domain shift from generic categories increases (\rightarrow).

Although MaPLe utilizes multi-modal prompts, its overall FLOPS (Floating Point Operations) exceeds only by 0.1% over CoOp and Co-CoOp. The independent V-L prompting also provides comparable FLOP count. In terms of inference speed, Co-CoOp is significantly slower and the FPS (Frames Per Second) remains constant as the batch size increases. In contrast, MaPLe has no such overhead and provides much better inference and training speeds. Further, MaPLe provides better convergence as it requires only half training epochs as compared to Co-CoOp (5 vs 10 epochs). MaPLe adds about 2.85% training parameters on top of CLIP. To study if the performance gain is mainly attributed to more parameters, we experiment with MaPLe \dagger , which uses a unified V-L coupling function for all layer prompts. MaPLe \dagger with about 9x lesser parameters than MaPLe also improves over existing methods. We also ablate by comparing MaPLe with heavier CoCoOp in Appendix D.

Method	Params	Params % CLIP	FPS (with BS)			HM
			1	4	100	
CoOp	2048	0.002	13.8	55.3	1353.0	71.66
CoCoOp	35360	0.03	64.6	114.7	15.1	75.83
Independent V-L	31488	0.02	62.5	239.4	1383.8	77.90
MaPLe	3.55 M	2.85	60.2	239.0	1365.1	78.55
MaPLe \dagger	0.41 M	0.33	60.2	238.0	1365.0	78.11

Table 6. Comparison of computational complexity among different prompting methods. MaPLe \dagger is a MaPLe version which utilizes a common V-L coupling function for all layers.

5. Conclusion

Adaptation of large-scale V-L models, *e.g.*, CLIP [32] to downstream tasks is a challenging problem due to the large number of tunable parameters and limited size of downstream datasets. Prompt learning is an efficient and scalable technique to tailor V-L models to novel downstream tasks. To this end, the current prompt learning approaches either consider only the vision or language side prompting. Our work shows that it is critical to perform prompting for both vision and language branches to appropriately adapt V-L models to downstream tasks. Further, we propose a strategy to ensure synergy between vision-language modalities by explicitly conditioning the vision prompts on textual prompt across different transformer stages. Our approach improves the generalization towards novel categories, cross-dataset transfer and datasets with domain shifts.

References

- [1] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Visual prompting: Modifying pixel space to adapt pre-trained models. *arXiv preprint arXiv:2203.17274*, 2022. [3](#), [11](#), [12](#)
- [2] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *The European Conference on Computer Vision*, pages 446–461. Springer, 2014. [5](#)
- [3] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3606–3613, 2014. [5](#)
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009. [5](#)
- [5] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11583–11592, 2022. [2](#)
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. [3](#)
- [7] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004. [5](#)
- [8] Chengjian Feng, Yujie Zhong, Zequn Jie, Xiangxiang Chu, Haibing Ren, Xiaolin Wei, Weidi Xie, and Lin Ma. Prompt-det: Towards open-vocabulary detection using uncurated images. In *The European Conference on Computer Vision*, 2022. [2](#)
- [9] Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*, 2021. [2](#), [3](#)
- [10] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. [2](#)
- [11] 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](#)
- [12] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. [5](#)
- [13] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021. [5](#)
- [14] Tony Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language models. *arXiv preprint arXiv:2204.03649*, 2022. [1](#)
- [15] 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](#)
- [16] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *The European Conference on Computer Vision*, 2022. [2](#), [3](#), [8](#)
- [17] Woojeong Jin, Yu Cheng, Yelong Shen, Weizhu Chen, and Xiang Ren. A good prompt is worth millions of parameters? low-resource prompt-based learning for vision-language models. *arXiv preprint arXiv:2110.08484*, 2021. [1](#)
- [18] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *The European Conference on Computer Vision*, 2021. [3](#)
- [19] Konwoo Kim, Michael Laskin, Igor Mordatch, and Deepak Pathak. How to adapt your large-scale vision-and-language model, 2022. [2](#)
- [20] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 554–561, 2013. [5](#)
- [21] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Conference on Empirical Methods in Natural Language Processing*, 2021. [2](#)
- [22] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *International Conference on Learning Representations*, 2022. [2](#)
- [23] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 2021. [2](#)
- [24] Xiao Liu, Kaixuan Ji, Yicheng Fu, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 2021. [2](#)
- [25] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *Proceed-*

- ings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5206–5215, 2022. 1, 3
- [26] 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. 2
- [27] Muhammad Maaz, Hanoona Rasheed, Salman Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Ming-Hsuan Yang. Class-agnostic object detection with multi-modal transformer. In *The European Conference on Computer Vision*. Springer, 2022. 2
- [28] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 5
- [29] Shu Manli, Nie Weili, Huang De-An, Yu Zhiding, Goldstein Tom, Anandkumar Anima, and Xiao Chaowei. Test-time prompt tuning for zero-shot generalization in vision-language models. 2022. 1
- [30] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 5
- [31] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012. 5
- [32] 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, 8
- [33] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022. 2
- [34] Hanoona Rasheed, Muhammad Maaz, Muhammad Uzair Khattak, Salman Khan, and Fahad Shahbaz Khan. Bridging the gap between object and image-level representations for open-vocabulary detection. In *Advances in Neural Information Processing Systems*, 2022. 2
- [35] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400. PMLR, 2019. 5
- [36] 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
- [37] 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*, volume 32, 2019. 5
- [38] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *The European Conference on Computer Vision*, 2022. 2
- [39] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. 2
- [40] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010. 5
- [41] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021. 2
- [42] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. Cpt: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797*, 2021. 1
- [43] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. 2
- [44] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Open-vocabulary detr with conditional matching. In *The European Conference on Computer Vision*, 2022. 2
- [45] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18123–18133, 2022. 2
- [46] Renrui Zhang, Rongyao Fang, Peng Gao, Wei Zhang, Kunchang Li, Jifeng Dai, Yu Qiao, and Hongsheng Li. Tip-adapter: Training-free clip-adapter for better vision-language modeling. In *The European Conference on Computer Vision*, 2022. 2
- [47] Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. Neural prompt search. *arXiv preprint arXiv:2206.04673*, 2022. 2
- [48] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. 1, 2, 3, 5, 6
- [49] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, pages 1–12, 2022. 1, 2, 3, 5, 6
- [50] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *The European Conference on Computer Vision*, 2022. 2
- [51] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. *arXiv preprint arXiv:2205.14865*, 2022. 2