



UvA-DARE (Digital Academic Repository)

Bayesian Prompt Learning for Image-Language Model Generalization

Derakhshani, M.M.; Sanchez, E.; Bulat, A.; Turrisi da Costa, V.G.; Snoek, C.G.M.; Tzimiropoulos, G.; Martinez, B.

DOI

[10.48550/arXiv.2210.02390](https://doi.org/10.48550/arXiv.2210.02390)

[10.1109/ICCV51070.2023.01398](https://doi.org/10.1109/ICCV51070.2023.01398)

Publication date

2023

Document Version

Final published version

Published in

2023 IEEE/CVF International Conference on Computer Vision

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Derakhshani, M. M., Sanchez, E., Bulat, A., Turrisi da Costa, V. G., Snoek, C. G. M., Tzimiropoulos, G., & Martinez, B. (2023). Bayesian Prompt Learning for Image-Language Model Generalization. In *2023 IEEE/CVF International Conference on Computer Vision: ICCV 2023 : Paris, France, 2-6 October 2023 : proceedings* (pp. 15191-15200). IEEE Computer Society. <https://doi.org/10.48550/arXiv.2210.02390>, <https://doi.org/10.1109/ICCV51070.2023.01398>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Bayesian Prompt Learning for Image-Language Model Generalization

Mohammad Mahdi Derakhshani^{1*}, Enrique Sanchez⁴, Adrian Bulat⁴

Victor Guilherme Turrise da Costa^{2*}, Cees G. M. Snoek^{1†}, Georgios Tzimiropoulos^{3,4†}, Brais Martinez^{4†}

¹University of Amsterdam ²University of Trento ³Queen Mary University of London

⁴Samsung AI Cambridge

Abstract

Foundational image-language models have generated considerable interest due to their efficient adaptation to downstream tasks by prompt learning. Prompt learning treats part of the language model input as trainable while freezing the rest, and optimizes an Empirical Risk Minimization objective. However, Empirical Risk Minimization is known to suffer from distributional shifts which hurt generalizability to prompts unseen during training. By leveraging the regularization ability of Bayesian methods, we frame prompt learning from the Bayesian perspective and formulate it as a variational inference problem. Our approach regularizes the prompt space, reduces overfitting to the seen prompts and improves the prompt generalization on unseen prompts. Our framework is implemented by modeling the input prompt space in a probabilistic manner, as an a priori distribution which makes our proposal compatible with prompt learning approaches that are unconditional or conditional on the image. We demonstrate empirically on 15 benchmarks that Bayesian prompt learning provides an appropriate coverage of the prompt space, prevents learning spurious features, and exploits transferable invariant features. This results in better generalization of unseen prompts, even across different datasets and domains.

Code available at: <https://github.com/saic-fi/Bayesian-Prompt-Learning>

1. Introduction

In the continuous quest for better pre-training strategies, models based on image and language supervision have set impressive milestones, with CLIP [38], ALIGN [22] and Flamingo [1] being leading examples. Contrastively trained image-language models consist of image and text encoders that align semantically-related concepts in a joint embedding space. Such models offer impressive zero-shot image classification by using the text encoder to generate classifier weights from arbitrarily newly defined category classes

*Most work done during an internship at Samsung AI Cambridge.

†Equal advising

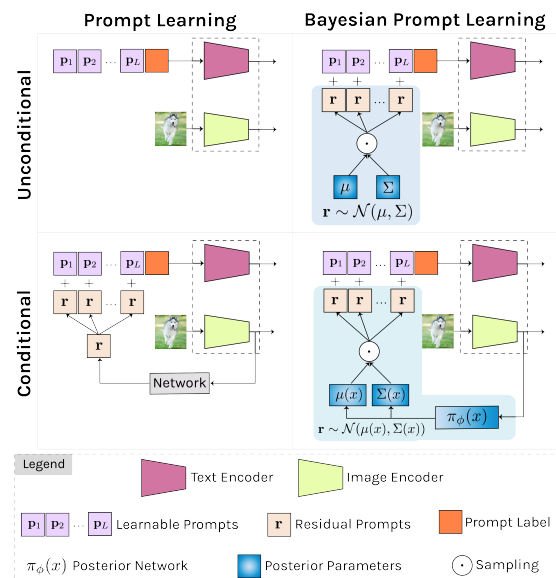


Figure 1: We present a Bayesian perspective on prompt learning by formulating it as a variational inference problem (right column). Our framework models the prompt space as an a priori distribution which makes our proposal compatible with common prompt learning approaches that are unconditional (top) or conditional on the image (bottom).

without relying on any visual data. In particular, the class name is used within a handcrafted *prompt* template and then tokenized and encoded into the shared embedding space to generate new classifier weights. Rather than manually defining prompts, both Lester *et al.* [28] and Zhou *et al.* [55] demonstrated prompts can instead be optimized in a data-driven manner through backpropagation. However, as prompt learning typically has access to only a few training examples per prompt, overfitting to the seen prompts in lieu of the unseen prompts is common [55]. In this paper, we strive to mitigate the overfitting behavior of prompt learning so as to improve generalization for unseen prompts.

Others before us have considered the generalization problem in prompt learning as well, *e.g.*, [55, 54], be it they all seek to optimize a so-called Empirical Risk Minimization

tion. It is, however, well known that Empirical Risk Minimization based models degrade drastically when training and testing distributions are different [37, 2]. To relax the i.i.d. assumption, Peters *et al.* [37] suggest exploiting the “invariance principle” for better generalization. Unfortunately, Invariant Risk Minimization methods for deep neural networks as of yet fail to deliver competitive results, as observed in [15, 31, 30]. To alleviate this limitation, Lin *et al.* [30] propose a Bayesian treatment of Invariant Risk Minimization that alleviates overfitting with deep models by defining a regularization term over the posterior distribution of classifiers, minimizing this term, and pushing the model’s backbone to learn invariant features. We take inspiration from this Bayesian Invariant Risk Minimization [30] and propose the first *Bayesian prompt learning* approach.

We make three contributions. First, we frame prompt learning from the Bayesian perspective and formulate it as a variational inference problem (see Figure 1). This formulation provides several benefits. First, it naturally injects noise during prompt learning and induces a regularization term that encourages the model to learn informative prompts, scattered in the prompt space, for each downstream task. As a direct result, we regularize the prompt space, reduce overfitting to seen prompts, and improve generalization on unseen prompts. Second, our framework models the input prompt space in a probabilistic manner, as an a priori distribution which makes our proposal compatible with prompt learning approaches that are unconditional [55] or conditional on the image [54]. Third, we empirically demonstrate on 15 benchmarks that Bayesian prompt learning provides an appropriate coverage of the prompt space, prevents learning spurious features, and exploits transferable invariant features, leading to a better generalization of unseen prompts, even across different datasets and domains.

2. Related Work

Prompt learning in language. Prompt learning was originally proposed within natural language processing (NLP), by models such as GPT-3 [4]. Early methods constructed prompts by combining words in the language space such that the model would perform better on downstream evaluation [46, 24]. Li and Liang [29] prepend a set of learnable prompts to the different layers of a frozen model and optimize through back-propagation. In parallel, Lester *et al.* [28] demonstrate that with no intermediate layer prefixes or task-specific output layers, adding prefixes alone to the input of the frozen model is enough to compete with fine-tuning. Alternatively, [16] uses a HyperNetwork to conditionally generate task-specific and layer-specific prompts pre-pended to the values and keys inside the self-attention layers of a frozen model. Inspired by progress from NLP, we propose a prompt learning method intended for image and language models.

Prompt learning in image and language. Zhou *et al.* [55] propose Context Optimization (CoOp), a prompt learner for CLIP, which optimizes prompts in the continuous space through back-propagation. The work demonstrates the benefit of prompt learning over prompt engineering. While CoOp obtains good accuracy on prompts seen during training, it has difficulty generalizing to unseen prompts. It motivated Zhou *et al.* to introduce Conditional Context Optimization (CoCoOp) [54]. It generates instance-specific prompt residuals through a conditioning mechanism dependent on the image data, which generalizes better. ProGrad by Zhu *et al.* [56] also strives to bridge the generalization gap by matching the gradient of the prompt to the general knowledge of the CLIP model to prevent forgetting. Alternative directions consist of test-time prompt learning [47], where consistency across multiple views is the supervisory signal, and unsupervised prompt learning [21], where a pseudo-labeling strategy drives the prompt learning. More similar to ours is ProDA by Lu *et al.* [32], who propose an ensemble of a fixed number of hand-crafted prompts and model the distribution exclusively within the language embedding. Conversely, we prefer to model the input prompt space rather than relying on a fixed number of templates, as it provides us with a mechanism to cover the prompt space by sampling. Moreover, our approach is not limited to unconditional prompt learning from the language embedding, but like Zhou *et al.* [54], also allows for prompt learning conditioned on an image.

Prompt learning in vision and language. While beyond our current scope, it is worth noting that prompt learning has been applied to a wider range of vision problems and scenarios, which highlights its power and flexibility. Among them are important topics such as unsupervised domain adaptation [13], multi-label classification [49], video classification [25], object detection [9, 12] and pixel-level labelling [40]. Finally, prompt learning has also been applied to vision only models [23, 44] providing an efficient and flexible means to adapt pre-trained models.

Variational inference in computer vision. Variational inference and, more specifically, variational autoencoder variants have been extensively applied to computer vision tasks as diverse as image generation [39, 43, 42], action recognition [34], instance segmentation [20], few-shot learning [53, 45], domain generalization [10], and continual learning [7]. For example, Zhang *et al.* [53] focus on the reduction of noise vulnerability and estimation bias in few-shot learning through variational inference. In the same vein, Du *et al.* [10] propose a variational information bottleneck to better manage prediction uncertainty and unknown domains. Our proposed method also shares the advantages of variational inference in avoiding overfitting in low-shot settings, improving generalization, and encourages the prompt space to be resilient against these chal-

lenges. To the best of our knowledge we are the first to introduce variational inference in prompt learning.

3. Method

3.1. Background

Contrastive Language-Image Pretraining (CLIP) [38] consists of an *image encoder* $f(\mathbf{x})$ and *text encoder* $g(\mathbf{t})$, each producing a d -dimensional (L_2 normalized) embedding from an arbitrary image $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$, and word embeddings $\mathbf{t} \in \mathbb{R}^{L \times e}$, with L representing the text length and e the embedding dimension¹. Both encoders are trained together using a contrastive loss from a large-scale dataset composed of paired images and captions. Once trained, CLIP enables zero-shot C -class image classification by generating each of the c classifier weights \mathbf{w}_c as the d -dimensional text encoding $g(\mathbf{t}_c)$. Here \mathbf{t}_c results from adding the class-specific word embedding \mathbf{e}_c to a pre-defined prompt $\mathbf{p} \in \mathbb{R}^{L-1 \times e}$, i.e., $\mathbf{w}_c = g(\mathbf{t}_c)$ with $\mathbf{t}_c = \{\mathbf{p}, \mathbf{e}_c\}$. The prompt \mathbf{p} is manually crafted to capture the semantic meaning of the downstream task, e.g., $\mathbf{t}_c = \text{“An image of a \{class\}”}$. The probability of image \mathbf{x} being classified as $y \in \{1 \dots C\}$ is thus defined as

$$p(y|\mathbf{x}) = \frac{e^{f(\mathbf{x})^T \mathbf{w}_y}}{\sum_c e^{f(\mathbf{x})^T \mathbf{w}_c}}.$$

Context Optimization (CoOp) [55] provides a learned alternative to manually defining prompts. CoOp learns a fixed prompt from a few annotated samples. The prompt is designed as a learnable embedding matrix $\mathbf{p} \in \mathbb{R}^{L \times e}$ which is updated via back-propagating the classification error through the frozen CLIP model. Specifically, for a set of N annotated meta-training samples $\{\mathbf{x}_i, y_i\}_{i=1}^N$, the prompt \mathbf{p} is obtained by minimizing the cross-entropy loss, as:

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \mathbb{E}_{\mathbf{x}_i, y_i} [-\log p(y_i | \mathbf{x}_i, \mathbf{p})]. \quad (1)$$

Note that this approach, while resembling that of common meta-learning approaches, can still be deployed in a zero-shot scenario provided that for new classes the classification weights will be given by the text encoder. Although this approach generalizes to new tasks with few training iterations, learning a fixed prompt is sensitive to domain shifts between the annotated samples and the unseen prompts.

Conditional Prompt Learning (CoCoOp) [54] attempts to overcome domain shifts by learning an instance-specific continuous prompt that is conditioned on the input image. To ease the training of a conditional prompt generator, CoCoOp defines each conditional token in a residual way, with a task-specific, learnable set of tokens \mathbf{p} and a residual prompt that is conditioned on

¹In CLIP the word embedding is learned together with the text encoder. A *tokenizer* is used to convert the text into one-hot vectors, or tokens, that can be directly mapped into the word embeddings. For the sake of clarity we refer indistinctly to words and word embeddings.

the input image. Assuming \mathbf{p} to be composed of L learnable tokens $\mathbf{p} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_L]$, the residual prompt $\mathbf{r}(\mathbf{x}) = \pi_\phi(f(\mathbf{x})) \in \mathbb{R}^e$ is produced by a small neural network π_ϕ with as input the image features $f(\mathbf{x})$. The new prompt is then computed as $\mathbf{p}(\mathbf{x}) = [\mathbf{p}_1 + \mathbf{r}(\mathbf{x}), \mathbf{p}_2 + \mathbf{r}(\mathbf{x}), \dots, \mathbf{p}_L + \mathbf{r}(\mathbf{x})]$. The training now comprises learning the task-specific prompt \mathbf{p} and the parameters ϕ of the neural network π_ϕ . Defining the context-specific text embedding $\mathbf{t}_c(\mathbf{x}) = \{\mathbf{p}(\mathbf{x}), \mathbf{e}_c\}$, and $p(y|\mathbf{x})$ as:

$$p(y|\mathbf{x}) = \frac{e^{f(\mathbf{x})^T g(\mathbf{t}_c(\mathbf{x}))}}{\sum_c e^{f(\mathbf{x})^T g(\mathbf{t}_c(\mathbf{x}))}}, \quad (2)$$

the learning is formulated as:

$$\mathbf{p}^*, \phi^* = \arg \min_{\mathbf{p}, \phi} \mathbb{E}_{\mathbf{x}_i, y_i} [-\log p(y_i | \mathbf{x}_i, \mathbf{p}, \phi)]. \quad (3)$$

While CoCoOp achieves good results in many downstream tasks, it is still prone to the domain shift problem, considering that π_ϕ provides a deterministic residual prompt from the domain-specific image features $f(\mathbf{x})$.

Prompt Distribution Learning (ProDA) [32] learns a distribution of prompts that generalize to a broader set of tasks. It learns a collection of prompts $\mathbf{P} = \{\mathbf{p}_k\}_{k=1}^K$ that subsequently generate an *a posteriori* distribution of the classifier weights for each of the target classes. For a given mini-batch of K sampled prompts $\mathbf{p}_k \sim \mathbf{P}$, the classifier weights \mathbf{w}_c are sampled from the posterior distribution $\mathbf{q} = \mathcal{N}(\mu_{\mathbf{w}_{1:C}}, \Sigma_{\mathbf{w}_{1:C}})$, with mean $\mu_{\mathbf{w}_{1:C}}$ and covariance $\Sigma_{\mathbf{w}_{1:C}}$ computed from the collection $\{\mathbf{w}_{k,c} = g(\mathbf{t}_{k,c})\}_{c=1:C, k=1:K}$, with $\mathbf{t}_{k,c} = \{\mathbf{p}_k, \mathbf{e}_c\}$. The objective is formulated as:

$$\mathbf{P}^* = \arg \min_{\mathbf{P}} \mathbb{E}_{\mathbf{x}_i, y_i} [-\log \mathbb{E}_{\mathbf{w}_l \sim \mathbf{q}} p(y_i | \mathbf{x}_i, \mathbf{w}_l)]. \quad (4)$$

Computing $\mathbb{E}_{\mathbf{w}_l} p(y_i | \mathbf{x}_i, \mathbf{w}_l)$ is intractable and an upper bound to Eq. 4 is derived. During inference, the classifier weights are set to those given by the predictive mean $\mathbf{w}_c = \mu_{\mathbf{w}_{1:C}}$, computed across the set of learned prompts \mathbf{P} .

3.2. Conditional Bayesian Prompt Learning

We propose to model the input prompt space in a probabilistic manner, as an *a priori*, conditional distribution. We define a distribution p_γ over the prompts \mathbf{p} that is conditional on the image, i.e., $\mathbf{p} \sim p_\gamma(\mathbf{x})$. To this end, we assume that \mathbf{p} can be split into a fixed set of prompts \mathbf{p}_i and an conditional residual prompt \mathbf{r} that act as a latent variable over \mathbf{p} . The conditional prompt is then defined as:

$$\mathbf{p}_\gamma(\mathbf{x}) = [\mathbf{p}_1 + \mathbf{r}_\gamma, \mathbf{p}_2 + \mathbf{r}_\gamma, \dots, \mathbf{p}_L + \mathbf{r}_\gamma], \quad \mathbf{r}_\gamma \sim p_\gamma(\mathbf{x}), \quad (5)$$

where $p_\gamma(\mathbf{x})$ refers to the real posterior distribution over \mathbf{r} conditioned on the observed features \mathbf{x} . Denoting the class-

specific input as $\mathbf{t}_{c,\gamma}(\mathbf{x}) = \{\mathbf{p}_\gamma(\mathbf{x}), \mathbf{e}_c\}$, the marginal likelihood $p(y|\mathbf{x})$ is:

$$p(y|\mathbf{x}) = \int_{\gamma} \frac{e^{f(\mathbf{x})^T g(\mathbf{t}_{c,\gamma}(\mathbf{x}))}}{\sum_{c'} e^{f(\mathbf{x})^T g(\mathbf{t}_{c',\gamma}(\mathbf{x}))}} p(\mathbf{p}_\gamma(\mathbf{x})) d\gamma. \quad (6)$$

Solving Eq. 3 with the marginal likelihood from Eq. 6 is intractable, as it requires computing $p_\gamma(\mathbf{r}|\mathbf{x})p_\gamma(\mathbf{x})$. Instead, we resort to deriving a lower bound by introducing a variational posterior distribution $\pi_\phi(\mathbf{x})$ from which the residual \mathbf{r}_γ can be sampled. The variational bound is defined as:

$$\log p(y|\mathbf{x}) \geq \mathbb{E}_{\pi_\phi(\mathbf{r}|\mathbf{x})}[\log p(y|\mathbf{x}, \mathbf{r})] - D_{\text{KL}}[\pi_\phi(\mathbf{r}|\mathbf{x}) \| p_\gamma(\mathbf{r})], \quad (7)$$

with $p(y|\mathbf{x}, \mathbf{r}) \propto e^{f(\mathbf{x})^T g(\mathbf{t}_{c,\gamma}(\mathbf{x}))}$, where the dependency on \mathbf{r} comes through the definition of $\mathbf{t}_{c,\gamma}$. Following standard variational optimization practices [26, 14], we define π_ϕ as a Gaussian distribution conditioned on the input image features \mathbf{x} , as $\mathbf{r}(\mathbf{x}) \sim \mathcal{N}(\mu(\mathbf{x}), \Sigma(\mathbf{x}))$, with μ and Σ parameterized by two linear layers followed by two linear heads on top to estimate the μ and Σ of the residual distribution. The prior $p_\gamma(\mathbf{r})$ is defined as $\mathcal{N}(\mathbf{0}, \mathbf{I})$, and we use the reparameterization trick to generate Monte-Carlo samples from π_ϕ to maximize the right side of Eq. 7. The optimization of Eq. 7 comprises learning the prompt embeddings $\{\mathbf{p}_i\}_{i=1}^L$ as well as the parameters of the posterior network π_ϕ and the linear layers parameterizing μ and Σ . Note that this adds little complexity as it requires learning \mathbf{p} and π_ϕ , given that μ and Σ are defined as two linear layers on top of π_ϕ .

Inference. At test time, K residuals are sampled from the conditional distribution $\pi_\phi(\mathbf{x})$, which are used to generate K different prompts per class $\mathbf{p}_k = [\mathbf{p}_1 + \mathbf{r}_k, \mathbf{p}_2 + \mathbf{r}_k, \dots, \mathbf{p}_L + \mathbf{r}_k]$. Each prompt is prepended to the class-specific embedding to generate a series of K separate classifier weights $\mathbf{w}_{k,c}$. We then compute $p(y=c|\mathbf{x}) = (1/K) \sum_{k=1}^K p(y=c|\mathbf{x}, \mathbf{w}_{k,c})$ and select $\hat{c} = \arg \max_c p(y=c|\mathbf{x})$ as the predicted class. It is worth noting that because the posterior distribution is generated by the text encoder, it is not expected that for $K \rightarrow \infty$, $(1/K) \sum_k p(y=c|\mathbf{x}, \mathbf{w}_{k,c}) \rightarrow p(y=c|\mathbf{x}, g(\{\mu(\mathbf{x}), \mathbf{e}_c\}))$, meaning that sampling at inference time remains relevant.

3.3. Unconditional Bayesian Prompt Learning

Notably, our framework can also be reformulated as an unconditional case by simply removing the dependency of the input image from the latent distribution. In such a scenario, we keep a fixed set of prompt embeddings and learn a global latent distribution p_γ over the residual prompts \mathbf{r} , as $\mathbf{r} \sim \mathcal{N}(\mu, \Sigma)$, where μ and Σ are parameterized by two learnable vectors. To this end, we assume that \mathbf{p} can be split into a fixed set of prompts \mathbf{p}_i and a residual prompts \mathbf{r} that act as a latent variable over \mathbf{p} . For the unconditional case, the prompt is defined as:

$$\mathbf{p}_\gamma = [\mathbf{p}_1 + \mathbf{r}_\gamma, \mathbf{p}_2 + \mathbf{r}_\gamma, \dots, \mathbf{p}_L + \mathbf{r}_\gamma], \quad \mathbf{r}_\gamma \sim p_\gamma. \quad (8)$$

In this case, p_γ is a general distribution learned during training with no dependency on the input sample \mathbf{x} .

Having defined both *unconditional Bayesian prompt learning* and *conditional Bayesian prompt learning* we are now ready to evaluate their generalization ability.

4. Experiments and Results

4.1. Experimental Setup

Three tasks and fifteen datasets. We consider three tasks: *unseen prompts generalization*, *cross-dataset prompts generalization*, and *cross-domain prompts generalization*. For the first two tasks, we rely on the same 11 image recognition datasets as Zhou *et al.* [55, 54]. These include image classification (ImageNet [6] and Caltech101 [11]), fine-grained classification (OxfordPets [36], StanfordCars [27], Flowers102 [35], Food101 [3] and FGVCAircraft [33]), scene recognition (SUN397 [52]), action recognition (UCF101 [48]), texture classification (DTD [5]), and satellite imagery recognition (EuroSAT [17]). For cross-domain prompts generalization, we train on ImageNet and report on ImageNetV2 [41], ImageNet-Sketch [51], ImageNet-A [19], and ImageNet-R [18].

Evaluation metrics. For all three tasks we report average accuracy and standard deviation.

Implementation details. Our conditional variational prompt learning contains three sub-networks: an image encoder $f(\mathbf{x})$, a text encoder $g(\mathbf{t})$, and a posterior network π_ϕ . The image encoder $f(\mathbf{x})$ is a ViT-B/16 [8] and the text encoder $g(\mathbf{t})$ a transformer [50], which are both initialized with CLIP's pre-trained weights and kept frozen during training, as in [55, 54]. The posterior network π_ϕ consists of two linear layers followed by an ELU activation function as trunk and two linear heads on top to estimate the μ and Σ of the residual distribution. For each task and dataset, we optimize the number of samples K and epochs. Other hyper-parameters as well as the training pipeline in terms of few-shot task definitions are identical to [55, 54] (see Table 9 and 10 in the appendix).

4.2. Comparisons

We first compare against CoOp [55], CoCoOp [54], and ProDA [32] in terms of the generalization of learned prompts on unseen classes, datasets, or domains. For CoOp and CoCoOp, all results are adopted from [54], and we report results for ProDA using our re-implementation.

Task I: unseen prompts generalization We report the unseen prompts generalization of our method on 11 datasets for three different random seeds. Each dataset is divided into two disjoint subsets: seen classes and unseen classes. We train our method on the seen classes and evaluate it on the unseen classes. For a fair comparison, we follow [55, 54] in terms of dataset split and number of shots. From

Table 1: **Task I: unseen prompts generalization** comparison between conditional Bayesian prompt learning and alternatives. Our model provides better generalization on unseen prompts compared to CoOp, CoCoOp and ProDA.

	CoOp [55]	CoCoOp [54]	ProDA [32]	Ours
Caltech101	89.81	93.81	93.23	94.93 ± 0.1
DTD	41.18	56.00	56.48	60.80 ± 0.5
EuroSAT	54.74	60.04	66.00	75.30 ± 0.7
FGVCAircraft	22.30	23.71	34.13	35.00 ± 0.5
Flowers102	59.67	71.75	68.68	70.40 ± 1.8
Food101	82.26	91.29	88.57	92.13 ± 0.1
ImageNet	67.88	70.43	70.23	70.93 ± 0.1
OxfordPets	95.29	97.69	97.83	98.00 ± 0.1
StanfordCars	60.40	73.59	71.20	73.23 ± 0.2
SUN397	65.89	76.86	76.93	77.87 ± 0.5
UCF101	56.05	73.45	71.97	75.77 ± 0.1
<i>Average</i>	63.22	71.69	72.30	74.94 ± 0.2

Table 1, it can be seen that our best-performing method, conditional Bayesian prompt learning, outperforms CoOp and CoCoOp in terms of unseen prompts generalization by 11.72% and 3.25%, respectively. Our proposal also demonstrates minimal average variance across all datasets. This is achieved by regularizing the optimization by virtue of the variational formulation. Moreover, our model also performs better than ProDA [32], its probabilistic counterpart, by 2.64%. This mainly happens because our proposed method learns the prompt distribution directly in the prompt space and allows us to sample more informative prompts for the downstream tasks.

Task II: cross-dataset prompts generalization For the next task, cross-dataset prompts generalization, the model is trained on a source dataset (ImageNet) and then assessed on 10 distinct target datasets. This experiment tries to determine how effectively our method generalizes beyond the scope of a single dataset. As reported in Table 2, our conditional Bayesian prompt learning outperforms CoOp and CoCoOp on the target dataset by 2.07% and 0.36%. This highlights that our method encourages the model to exploit transferable invariant features beneficial for datasets with a non-overlapping label space. Furthermore, our method improves accuracy for 7 out of 10 target datasets. Unlike CoCoOp, which performs better in ImageNet-like datasets such as Caltech101 and StanfordCars, our Bayesian method exhibits improvement on dissimilar datasets (*e.g.*, FGVC-Aircraft, DTD, and EuroSAT), demonstrating its capacity to capture the unique characteristics of each dataset.

Task III: cross-domain prompts generalization Lastly, we examine our conditional Bayesian prompt learning through the lens of distribution shift and robustness. We train our model on the source dataset ImageNet for three different random seeds, and assess it on ImageNetV2, ImageNet-Sketch, ImageNet-A, and ImageNet-R. Prior works such as CoOp [55] and CoCoOp [54] demonstrate empirically that learning a soft-prompt improves the

Table 2: **Task II: cross-dataset prompts generalization.** Our proposed model is evaluated on 11 datasets with different label spaces. As shown, conditional Bayesian prompt learning performs better than non-Bayesian alternatives on 7 out of 10 datasets. **Task III: cross-domain prompts generalization.** Our model is evaluated on four datasets sharing the same label space as the training data. Our method outperforms alternatives on 3 out of 4 datasets.

	CoOp [55]	CoCoOp [54]	Ours
Task II: cross-dataset prompts generalization			
Caltech101	93.70	94.43	93.67 ± 0.2
DTD	41.92	45.73	46.10 ± 0.1
EuroSAT	46.39	45.37	45.87 ± 0.7
FGVCAircraft	18.47	22.94	24.93 ± 0.2
Flowers102	68.71	70.20	70.90 ± 0.1
Food101	85.30	86.06	86.30 ± 0.1
OxfordPets	89.14	90.14	90.63 ± 0.1
StanfordCars	64.51	65.50	65.00 ± 0.1
SUN397	64.15	67.36	67.47 ± 0.1
UCF101	66.55	68.21	68.67 ± 0.2
<i>Average</i>	63.88	65.59	65.95 ± 0.2
Task III: cross-domain prompts generalization			
ImageNetV2	64.20	64.07	64.23 ± 0.1
ImageNet-Sketch	47.99	48.75	49.20 ± 0.0
ImageNet-A	49.71	50.63	51.33 ± 0.1
ImageNet-R	75.21	76.18	77.00 ± 0.1
<i>Average</i>	59.27	59.88	60.44 ± 0.1

Table 3: **In-domain performance.** CoOp provides the best in-domain performance, but suffers from distribution shifts. Our proposal provides the best trade-off.

	CoOp [55]	CoCoOp [54]	ProDA [32]	Ours
<i>Task I</i>	82.66	80.47	81.56	80.10 ± 0.1
<i>Task II (III)</i>	71.51	71.02	-	70.70 ± 0.2

model’s resilience against distribution shift and adversarial attack. Following their experiments, we are also interested in determining if treating prompts in a Bayesian manner maintains or improves performance. As reported in Table 2 compared with CoOp, our method improves the performance on ImageNet-Sketch, ImageNet-A, and ImageNet-R by 1.21%, 1.62%, and 1.79%. Compared with CoCoOp, our proposed method consistently enhances the model accuracy on ImageNetV2, ImageNet-Sketch, ImageNet-A, and ImageNet-R by 0.16%, 0.45%, 0.7%, and 0.82%. This is because our proposed method prevents learning spurious features, *e.g.*, high-frequency features, and instead learns invariant ones by virtue of the Bayesian formulation.

In-domain performance Different from other prompt works, we focus on prompt generalization when we have distribution shift in the input space, the label space, or both.

Table 4: **Effect of variational formulation.** Formulating prompt learning as variational inference improves model generalization on unseen prompts compared to a non-Bayesian baseline [55], for both the unconditional and conditional setting.

	DTD	EuroSAT	FGVC	Flowers102	UCF101
Baseline	41.18	54.74	22.30	59.67	56.05
Unconditional	58.70	71.63	33.80	75.90	74.63
Conditional	60.80	75.30	35.00	70.40	75.77

Table 5: **Benefit of the posterior distribution.** The conditional posterior distribution $\mathcal{N}(\mu(\mathbf{x}), \Sigma(\mathbf{x}))$ outperforms the two distributions $\mathcal{U}(0, 1)$ and $\mathcal{N}(0, \mathbf{I})$ by a large margin for all datasets, indicating that the conditional variant captures more informative knowledge regarding the underlying distribution of prompts in the prompt space.

	DTD	EuroSAT	FGVC	Flowers102	UCF101
$\mathcal{U}(0, 1)$	33.20	54.20	10.50	45.30	55.70
$\mathcal{N}(0, \mathbf{I})$	26.60	50.00	07.70	36.10	48.80
$\mathcal{N}(\mu(\mathbf{x}), \Sigma(\mathbf{x}))$	56.40	64.50	33.00	72.30	75.60
$\mathcal{N}(\mu(\mathbf{x}), 0)$	59.80	59.90	34.10	73.50	76.50

For reference, we provide the in-domain average performance in Table 3. As expected, our generalization ability comes with reduced overfitting on the in-domain setting, leading to CoOp exhibiting better in-domain performance. However, when comparing average results in Tables 1, 2 and 3, our proposal brings the best trade-off and we maintain performance on the in-domain setting while delivering a consistent performance gain on the out-of-domain setting.

4.3. Ablations

Effect of variational formulation. When prompt learning is formulated as a variational inference problem, it improves model generalization by regularizing the prompt space. To demonstrate this, we first report the model performances on unseen prompts for DTD, EuroSAT, FGVC-Aircraft, Flowers102, and UCF101 in Table 4. Regardless of whether a variational formulation is unconditional or conditional, it consistently increases unseen prompt generalization compared with the CoOp baseline [55]. This is achieved by sampling more informative prompts and better coverage of the prompt space (which we detail in a following ablation). Moreover, conditional variational prompt learning enhances the model performance on 4 out of 5 datasets compared to its unconditional counterpart, indicating its effectiveness in capturing the prompt distribution. This happens as the conditional formulation provides us with the potential to successfully capture fine-grained prompts by conditioning on the input image.

Benefit of the posterior distribution q_ϕ . We ablate the effectiveness of the variational posterior distribution. To

Table 6: **Influence of Monte Carlo sampling** on unseen prompts accuracy. As demonstrated, increasing the number of Monte Carlo samples boosts performance initially but reaches a plateau after 20 samples for all datasets.

	DTD	EuroSAT	FGVC	Flowers102	UCF101
1	56.40	64.50	33.00	72.30	75.60
2	60.00	67.40	33.90	73.90	76.20
5	62.20	71.00	34.20	74.00	76.60
10	61.60	73.60	34.40	73.50	77.00
20	62.90	74.80	35.00	74.00	77.10
40	62.60	76.10	35.50	73.80	77.15
80	63.50	76.20	34.70	74.20	77.20

Table 7: **Influence of Monte Carlo sampling** on EuroSAT. Increasing samples increases the time complexity at inference, with a good trade-off for 20 samples (see Table 6).

	Monte Carlo Samples						
	1	5	10	20	40	60	80
Time (s)	0.02	0.02	0.03	0.04	0.07	0.10	0.13

do so, we consider sampling one residual prompt from the uniform distribution $\mathcal{U}(0, 1)$, normal distribution $\mathcal{N}(0, \mathbf{I})$, normal distribution $\mathcal{N}(\mu(\mathbf{x}), \Sigma(\mathbf{x}))$, and normal distribution $\mathcal{N}(\mu(\mathbf{x}), 0)$ and report the unseen prompts accuracy for one random seed in Table 5. We find that drawing one sample from $\mathcal{N}(\mu(\mathbf{x}), \Sigma(\mathbf{x}))$ yields superior results compared to drawing one sample from uniform distribution $\mathcal{U}(0, 1)$ and normal distribution $\mathcal{N}(0, \mathbf{I})$, demonstrating the informativeness of the prompt samples generated by our conditional setting. In addition, except for the EuroSAT dataset, a sample from the normal distribution $\mathcal{N}(\mu(\mathbf{x}), 0)$ obtains the best performance in comparison with alternatives, indicating the mean of the normal distribution $\mu(\mathbf{x})$ is the most effective sample.

Influence of Monte Carlo sampling. When approximating the log-likelihood of input data, the number of Monte Carlo samples is an important hyperparameter. Generally, a large number of samples should lead to a better approximation and unseen prompt generalization. We ablate this hyperparameter on unseen prompts by varying the number of samples from the normal distribution $\mathcal{N}(\mu(\mathbf{x}), \Sigma(\mathbf{x}))$ to understand the informativeness of the learned variational distribution at inference time in Table 6. Increasing the Monte Carlo samples from 1 to 20 consistently improves the unseen prompts accuracy, afterwards, the accuracy saturates. Table 7 further shows the impact of the time complexity as the number of Monte Carlo samples increases on one NVIDIA RTX 2080ti GPU. As expected, it shows the increase in sample size is accompanied by an increase in model inference time. We obtain a good trade-off in terms of accuracy and inference cost for 20 samples. For bet-

Table 8: **Effect of prompt initialization.** Initializing the context tokens with an appropriate prompt “An image of a `{class}`” (denoted as Guided) improves the performance compared to random initialization.

	DTD	EuroSAT	FGVC	Flowers102	UCF101
Random	74.10	75.10	34.30	67.23	81.30
Guided	74.80	77.90	35.50	70.05	82.50

ter accuracy we recommend evaluating on a larger number of Monte Carlo samples as they capture the prompt space more appropriately. We optimize this hyper-parameter using cross-validation for all datasets and use them in the remaining experiments (see Table 10 in the appendix.)

Effect of prompt initialization. Next we ablate the effectiveness of the prompt initialization on the unseen prompts accuracy for one random seed. We consider two variants. In the first we initialize the context tokens randomly using a normal distribution and in the second we initialize the context tokens with “An image of a `{class}`”. Comparing the two variants in Table 8 demonstrates that an appropriately initialized prompt consistently outperforms a randomly initialized prompt. This is mainly because the prompt samples distributed in the vicinity of a guided prompt, e.g., “An image of a `{class}`”, are more informative than a random prompt. It can be viewed as informative prior knowledge about the prompt distribution, which facilitates the optimization towards a better prompt sub-space.

Benefit of variational ensemble. Variational prompt learning can be considered as an efficient ensemble approach as it generates several samples from the prompt space and uses them in conjunction. To demonstrate the value of our proposal over a naïve ensemble of prompts we compare against a modified version of CoCoOp [54]. We implement this version by initializing several learnable prompts per class using M hand-crafted templates provided in the CLIP codebase and fine-tune them together. The final text feature representation is computed as a simple average of the ensemble of prompts for each class (as in Section 3.1.4 of [38]). For a fair comparison, the number of ensembles, M , equals the number of Monte Carlo samples K per dataset (see Table 10 in the appendix). The results in Table 9 demonstrate that adding a naïve ensemble of prompts on top of CoCoOp does not necessarily lead to an improvement. For instance on DTD and UCF101, CoCoOp performs better by itself, while on Flowers102, EuroSAT, and FGVC Aircraft, the naïve ensemble obtains a higher unseen prompts accuracy. When we compare the naïve ensemble with our proposal, we note our method performs better in terms of unseen prompts accuracy for 4 out of 5 datasets. The reason is that, in naïve ensemble, these M learnable prompts likely converge into contradictory prompts, that deteriorate the final performance, due to the lack of a proper regularization term in its objective function. Moreover, the

Table 9: **Benefit of variational ensemble.** Our conditional Bayesian prompt learning outperforms CoCoOp and its naïve ensemble for 4 out of 5 datasets regarding unseen prompts accuracy, demonstrating that our method is effective in constructing informative prompt samples for ensemble learning in a data-driven manner.

	DTD	EuroSAT	FGVC	Flowers102	UCF101
CoCoOp [54]	56.00	60.04	23.71	71.75	73.45
Naïve ensemble	52.17	66.10	32.63	72.17	73.90
<i>This paper</i>	60.80	75.30	35.00	70.40	75.77

KL divergence in Eq. 7 serves to regulate and expand the norm of the prompt space. This allows our proposal to freely explore the prompt space, leading to complementary prompt samples. We believe the prompt samples are not contradictory as they are jointly optimized to cover inherent variability in the model predictions and maximize model performance.

Variational distribution interpretation. Next, we provide an intuition on how prompt samples are distributed within the variational distribution. For an input image \mathbf{x} and its corresponding target y , K residuals are sampled from the conditional distribution $\pi_\phi(\mathbf{x})$, which are used to generate K different prompts per class $\mathbf{p}_k = [\mathbf{p}_1 + \mathbf{r}_k, \mathbf{p}_2 + \mathbf{r}_k, \dots, \mathbf{p}_L + \mathbf{r}_k]$. We also construct a new prompt based on the mean of the variational distribution $\mu(\mathbf{x})$ dubbed as the mean prompt $\mathbf{p}_{\mu(\mathbf{x})} = [\mathbf{p}_1 + \mu(\mathbf{x}), \mathbf{p}_2 + \mu(\mathbf{x}), \dots, \mathbf{p}_L + \mu(\mathbf{x})]$. These K prompts are sorted based on their distance to the mean of the variational distribution $\mu(\mathbf{x})$, where for any i and j ($i \leq j$), $d(\mu(\mathbf{x}), \mathbf{r}_i) \leq d(\mu(\mathbf{x}), \mathbf{r}_j)$. All these K prompts and the mean prompt are prepended to the class-specific embedding to generate a series of K separate text encodings $\mathbf{w}_{\mu(\mathbf{x}),c}, \mathbf{w}_{1,c}, \mathbf{w}_{2,c}, \dots, \mathbf{w}_{K,c}$. For each class y , we compute the cosine similarity of the image encoding $f(\mathbf{x})$ and $\mathbf{w}_{\mu(\mathbf{x}),y}, \mathbf{w}_{1,y}, \mathbf{w}_{2,y}, \dots, \mathbf{w}_{K,y}$, and visualize them for $K=20$ prompts in Figure 2. The mean prompt $\mathbf{p}_{\mu(\mathbf{x})}$ has the largest cosine similarity between its corresponding text and image encoding among other prompts, indicating it is the most informative sample. As we move further away from the mean prompt, the cosine similarity score decreases. Thus, the further a prompt is from the mean prompt, the less informative it becomes.

Next we demonstrate that prompts distributed within the variational distribution positively complement each other and maximize cosine similarity. Given the text encodings $\mathbf{w}_{\mu(\mathbf{x}),y}, \mathbf{w}_{1,y}, \mathbf{w}_{2,y}, \dots, \mathbf{w}_{M,y}$ we further define a new weight as $\mathbf{w}_y^J = \mathbf{w}_{\mu(\mathbf{x}),y} + \sum_{j=1}^J \mathbf{w}_{j,y}$, where \mathbf{w}_y^J is the cumulative sum of text encodings regarding prompt \mathbf{p}_1 to \mathbf{p}_J and mean prompt $\mathbf{p}_{\mu(\mathbf{x})}$. The cosine similarity between the image encoding $f(\mathbf{x})$ and all $\{\mathbf{w}_y^j\}_{j=1}^M$ are computed and visualized in Figure 2. As shown, summing up, or ensembling, text encodings together increases the cosine similarity, where the maximum similarity is obtained when all clas-

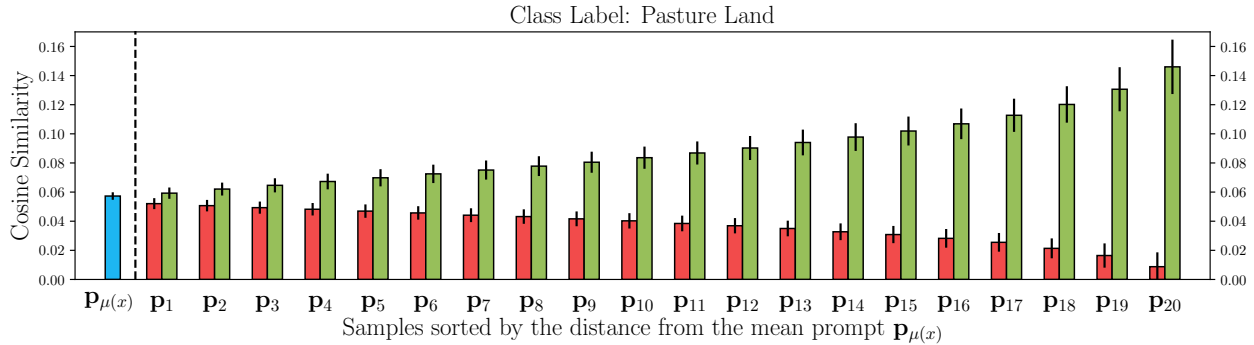


Figure 2: **Variational distribution interpretation** on the EuroSAT dataset. The text encoding of the mean prompt $\mathbf{p}_{\mu(\mathbf{x})}$ (■) is the most similar to the image encoding. As we move further away from the mean prompt, the cosine similarity scores between the text encoding and image encoding decrease further (■). When we ensemble the text encoding of different prompts the cosine similarity increases (■), where the maximum similarity is obtained when all text encodings are combined.

sifier weights are combined. Hence, our knowledge from the prompt space can be adjusted by gradually increasing the number of Monte Carlo samples.

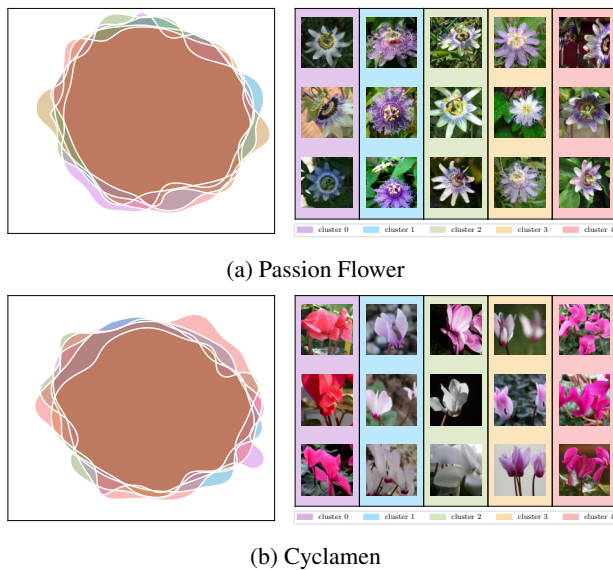


Figure 3: **Factor of variation analysis** on Flowers102 for two different classes. Left: we plot prompt samples across five clusters. Right: we show the top-3 most representative samples within each cluster (e.g., 3 closest images to the centroids). There is a region where the contours for five different clusters intersect, indicating shared knowledge related to its corresponding class, while they diverge slightly, indicating a particular variation factor.

Factor of variation analysis Here we analyze whether prompts that are sampled at the distributional modes correlate with features that characterize any known factors of variation (e.g., sub-domains) in a dataset. We perform this experiment on our conditional Bayesian prompt learning model. First, we randomly select two classes c_1 and c_2 from

the Flowers102 dataset (Passion Flower, and Cyclamen, for reference). Then, for each class, we compute image features by forwarding the respective input images \mathbf{x} into the image encoder of CLIP. Next, we apply k -means to group the image features into $Q=5$ cluster centroids. Note that each cluster centroid is assumed to capture some factors of visual variation within each class. Afterward, we treat the cluster assignments as pseudo-labels of those images. Note that this is done separately for each class. Additionally, for an input image \mathbf{x} , we generate $M=21$ different prompts. These prompts are fed into the clip’s text encoder, which generates M different weights $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_M$ per image instance. We now have a set of M weights for Q pseudo-labels across two classes, which are visualized in each row of Figure 3. From this figure, it can be seen that there is a high intersection between the distribution of all sub-classes, representing shared knowledge. However, some particular visual differences are expressed in regions without intersections in the distribution. Hence, we believe there is a correlation between the modes of the variational distribution and visual particularities of subsets of the same class.

5. Conclusion

In this paper, we formulate prompt learning for image-language models as a variational inference problem. The Bayesian perspective allows regularizing the prompt space, reducing overfitting to seen prompts, and improving unseen prompt generalization. We model the input prompt space as an a priori distribution, making our framework compatible with unconditional or conditional prompt learning approaches. Using 15 benchmarks, we show that Bayesian prompt learning provides an appropriate coverage of the prompt space, prevents learning spurious features, and exploits transferable invariant features, leading to better generalization of unseen prompts, even across datasets and domains. We believe prompt distributions can be estimated

also using other probabilistic methods, like normalizing flows, energy-based models and diffusion models, which we will leave open for future investigation.

References

- [1] Jean-Baptiste Alayrac et al. Flamingo: a visual language model for few-shot learning. In *Advances on Neural Information Processing Systems*, 2022.
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- [3] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014.
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Advances on Neural Information Processing Systems*, 2020.
- [5] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [7] Mohammad Mahdi Derakhshani, Xiantong Zhen, Ling Shao, and Cees Snoek. Kernel continual learning. In *International Conference on Machine Learning*, 2021.
- [8] 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.
- [9] Yu Du, Fangyun Wei, Zihe Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [10] Yingjun Du, Jun Xu, Huan Xiong, Qiang Qiu, Xiantong Zhen, Cees GM Snoek, and Ling Shao. Learning to learn with variational information bottleneck for domain generalization. In *European Conference on Computer Vision*, 2020.
- [11] 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 *IEEE Conference on Computer Vision and Pattern Recognition - Workshops*, 2004.
- [12] Chengjian Feng, Yujie Zhong, Zequn Jie, Xiangxiang Chu, Haibing Ren, Xiaolin Wei, Weidi Xie, and Lin Ma. PromptDet: Towards open-vocabulary detection using uncurated images. In *European Conference on Computer Vision*, 2022.
- [13] Chunjiang Ge, Rui Huang, Mixue Xie, Zihang Lai, Shiji Song, Shuang Li, and Gao Huang. Domain adaptation via prompt learning. *arXiv preprint arXiv:2202.06687*, 2022.
- [14] Jonathan Gordon, John Bronskill, Matthias Bauer, Sebastian Nowozin, and Richard Turner. Meta-learning probabilistic inference for prediction. In *International Conference on Learning Representations*, 2019.
- [15] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.
- [16] Yun He, Steven Zheng, Yi Tay, Jai Gupta, Yu Du, Vamsi Aribandi, Zhe Zhao, YaGuang Li, Zhao Chen, Donald Metzler, Heng-Tze Cheng, and Ed H. Chi. HyperPrompt: Prompt-based task-conditioning of transformers. In *International Conference on Machine Learning*, 2022.
- [17] 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*, 2019.
- [18] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [19] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [20] Namdar Homayounfar, Yuwen Xiong, Justin Liang, Wei-Chiu Ma, and Raquel Urtasun. Levelset r-cnn: A deep variational method for instance segmentation. In *European Conference on Computer Vision*, 2020.
- [21] Tony Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language models. *arXiv preprint arXiv:2204.03649*, 2022.
- [22] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. 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*, 2021.
- [23] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision*, 2022.
- [24] Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 2020.
- [25] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, 2022.
- [26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- [27] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3D object representations for fine-grained categorization. In *IEEE Conference on Computer Vision and Pattern Recognition - Workshops*, 2013.

- [28] 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.
- [29] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *Annual Meeting of the Association for Computational Linguistics*, 2021.
- [30] Yong Lin, Hanze Dong, Hao Wang, and Tong Zhang. Bayesian invariant risk minimization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [31] Yong Lin, Qing Lian, and Tong Zhang. An empirical study of invariant risk minimization on deep models. In *International Conference on Machine Learning - Workshops*, 2021.
- [32] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [33] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- [34] Nazanin Mehrasa, Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. A variational auto-encoder model for stochastic point processes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [35] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics & Image Processing*, 2008.
- [36] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [37] Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 2016.
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021.
- [39] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, 2021.
- [40] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [41] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, 2019.
- [42] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [43] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.
- [44] Mark Sandler, Andrey Zhmoginov, Max Vladymyrov, and Andrew Jackson. Fine-tuning image transformers using learnable memory. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [45] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [46] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020.
- [47] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. In *Advances on Neural Information Processing Systems*, 2022.
- [48] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. A dataset of 101 human action classes from videos in the wild. *Center for Research in Computer Vision*, 2012.
- [49] Ximeng Sun, Ping Hu, and Kate Saenko. DualCoOp: Fast adaptation to multi-label recognition with limited annotations. *arXiv preprint arXiv:2206.09541*, 2022.
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances on Neural Information Processing Systems*, 2017.
- [51] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. 2019.
- [52] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [53] Jian Zhang, Chenglong Zhao, Bingbing Ni, Minghao Xu, and Xiaokang Yang. Variational few-shot learning. In *IEEE International Conference on Computer Vision*, 2019.
- [54] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [55] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal on Computer Vision*, 2022.
- [56] Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. Prompt-aligned gradient for prompt tuning. *arXiv preprint arXiv:2205.14865*, 2022.