

Probabilistic Prompt Learning for Dense Prediction

Hyeongjun Kwon¹ Taeyong Song² Somi Jeong³ Jin Kim¹ Jinhyun Jang¹ Kwanghoon Sohn^{1,4*}

¹Yonsei University, ²Hyundai Motor Company R&D Division,

³NAVER LABS, ⁴Korea Institute of Science and Technology (KIST)

{kwonjunn01, kimjin928, jr000192, khsohn}@yonsei.ac.kr,

taeyongsong@hyundai.com, somi.jeong@naverlabs.com

Abstract

Recent progress in deterministic prompt learning has become a promising alternative to various downstream vision tasks, enabling models to learn powerful visual representations with the help of pre-trained vision-language models. However, this approach results in limited performance for dense prediction tasks that require handling more complex and diverse objects, since a single and deterministic description cannot sufficiently represent the entire image. In this paper, we present a novel probabilistic prompt learning to fully exploit the vision-language knowledge in dense prediction tasks. First, we introduce learnable class-agnostic attribute prompts to describe universal attributes across the object class. The attributes are combined with class information and visual-context knowledge to define the class-specific textual distribution. Text representations are sampled and used to guide the dense prediction task using the probabilistic pixel-text matching loss, enhancing the stability and generalization capability of the proposed method. Extensive experiments on different dense prediction tasks and ablation studies demonstrate the effectiveness of our proposed method.

1. Introduction

Dense predictions, *e.g.*, semantic segmentation [3, 31], instance segmentation [30], and object detection [11, 42], are fundamental computer vision problems, which aim to produce pixel-level predictions to thoroughly understand the scene. Due to the expensive cost of collecting dense annotations, most approaches [10, 38] employ a “pre-training + fine-tuning” paradigm. Based on existing pre-trained networks [8, 15] trained on large-scale datasets such as ImageNet [6], semi- [36, 60] or self-supervised learning [29, 53] has been extensively researched to fine-tune additional

*Corresponding author.

This research was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (NRF2021R1A2C2006703).

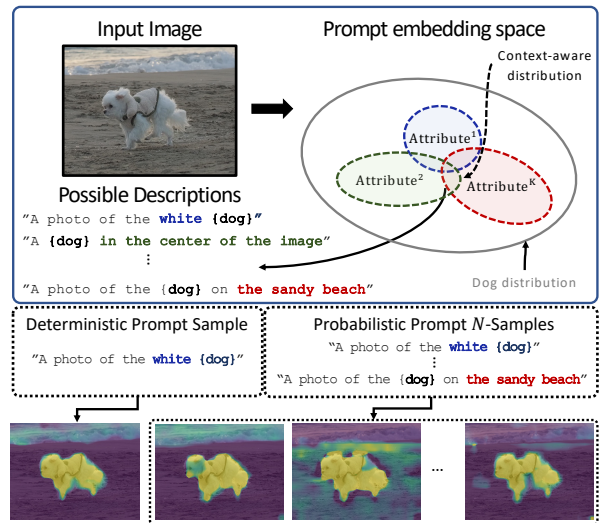


Figure 1. **Probabilistic Prompt Learning.** The proposed PPL exploits N multiple prompts sampled from probabilistic text embeddings, which can leverage granular textual representations, enabling a better understanding of the details of the image.

modules for dense prediction. However, due to the biased visual representations, they suffer from a lack of scalability when there exists a large semantic gap between pre-trained and target tasks w.r.t. dataset and objective, such as transferring ImageNet classification network to COCO object detection [13, 61].

Recently, applying vision-language pre-trained models (VLM) to the downstream tasks has demonstrated remarkable success, including zero-shot classification [58, 59], referring expression segmentation [45], and object detection [41]. VLM, such as CLIP [40] and ALIGN [21], is trained on large-scale web noisy image-text pair datasets via contrastive learning to align representations between text and image in a joint embedding space. In this way, VLM learns robust visual representations by exploiting the semantic relationship between text and image representations, which is beneficial to transfer knowledge to various

vision tasks. To efficiently leverage the language knowledge, there have been many pioneering attempts [16, 22, 46] to deploy VLM to the downstream tasks via *prompting*. For example, based on a prompt template "a photo of a {class}", it measures the confidence score by calculating image-text similarity and classifies the image into {class}. In practice, however, this hand-crafted prompt may not be the optimal description for a particular task, and furthermore, manually designing a task-specific prompt is laborious.

To tackle this issue, several methods [41, 58, 59] have introduced learning-based prompting techniques inspired by early works in NLP [18, 28]. The goal is to automatically construct the prompts according to the task by optimizing continuous prompt variables based on VLM. This simple and intuitive setup, referred to as *deterministic prompt learning*, is the most popular approach in the current literature [9, 46]. It has shown performance improvement on classification tasks [40, 52], where a single deterministic embedding is sufficient for representing a class. However, this approach is not fully compositional in dense prediction tasks due to a semantic ambiguity problem. Firstly, while the dense prediction tasks require granular information to generate precise pixel-wise results, not only complex and multiple objects within an input image but also their various attributes (*e.g.*, color, location, etc.) cannot be comprehensively represented in a single textual representation [12, 49]. Thus, a single prompt fails to comprehend the object of all classes in detail. Second, the visual representation has high randomness [5, 7] due to various contexts with external objects and object representations, and it results in high uncertainty in representing in language. For example, as shown in Fig. 1, the image can be described as "A photo of the dog on the sandy beach", but it can also be expressed differently such as "A photo of the dog near the ocean". Therefore, it is not appropriate to exploit a deterministically visual representation when transferring VLM in dense prediction tasks.

In this paper, we propose a Probabilistic Prompt Learning (PPL) that explores learning appropriate textual descriptions using visual cues in a probabilistic embedding space. We present a set of prompts that express class-agnostic attributes such as position, size, and color to represent objects without semantic ambiguity. To effectively learn the probabilistic distribution to describe diverse and informative attributes for the entire class, we model each attribute distribution as a distinct normal distribution. To this end, we set its variance as contextual relations between text and visual features to explicitly utilize visual-text information. With these attribute distributions, class-specific attribute distribution is approximated by a Mixture of Gaussian (MoG). Furthermore, we introduce a novel probabilistic pixel-text matching loss to attenuate the instability of

prediction probability caused by high uncertainty.

In summary, our contributions are as follows: (1) We propose a novel PPL to effectively represent class-agnostic attributes of objects in probabilistic text embedding space. To the best of our knowledge, this is the first attempt to leverage context-aware probabilistic prompt learning. (2) We introduce a novel probabilistic pixel-text matching loss to alleviate the adverse impact of uncertainty. (3) We demonstrate the effectiveness of the proposed probabilistic approach through extensive experiments on dense prediction tasks, including semantic segmentation, instance segmentation, and object detection.

2. Related Works

Vision-Language pre-trained Model. Vision-language pre-trained models (VLM) [17, 20, 23, 25] have been widely researched on various downstream vision tasks, including visual question answering [1], image captioning [47], text-to-image retrieval [39], and so on. Conventional VLM learns the connections between image content and language via extra modules. However, due to the relatively small dataset, most approaches had difficulty aligning the connection between image and language. Recently, applying contrastive learning on noisy web-scale image-text paired datasets has shown promising results in VLM such as CLIP [40] and ALIGN [21]. Leveraging this language supervision from VLM, there exists remarkable improvement in various downstream vision tasks under unannotated or restricted data regimes.

Prompt Learning. Prompt learning, inspired by the concurrent work in NLP [28], is widely researched for VLM, which aims to learn to generate optimal descriptions that enhance the visual-text representations. CoOp [59] is a pioneer work that applied prompt learning to vision tasks, and it leverages learnable continuous prompts trained on the freeze CLIP encoder. Lu *et al.* [32] proposed ProDA to optimize multiple sets of prompts by learning the distribution of prompts, which discovers task-related content with less bias than manual design. With the recent progress of prompt learning, these approaches have demonstrated impressive improvements in high-level vision tasks, including video recognition [22, 33], point cloud understanding [54], and dense prediction [41]. Especially, DenseCLIP [41] is proposed to apply that prompt learning to dense prediction tasks, where pixel-text matching loss is used as a guide for the task loss. Despite the progress of prompting leveraging the visual-context, they still do not consider the randomness of the visual-context.

Probabilistic Embedding. Learning probabilistic representations is a traditional approach in the word embedding approach [44]. Since they fully exploit the inherent hierarchical structure in language embeddings, it has been widely studied for structuring different distributions with word rep-

representations. Recently, probabilistic embedding approaches have received more attention in the field of computer vision. Hedged Instance Embedding (HIB) [34] is proposed to handle one-to-many mappings via the variational information bottleneck (VIB) principle. Based on the HIB concept, probabilistic embeddings for cross-modal retrieval [5] have been studied to learn joint embeddings to capture one-to-many associations.

Uncertainty in Computer Vision. Uncertainty is one of the main problems that degrade reliable performance in CNN-based methods. Therefore, various methods of handling this uncertainty to improve robustness and performance have been extensively studied in various applications such as image retrieval [5], face recognition [2], video retrieval [37], and semantic segmentation [24]. In general, uncertainty can be classified into two types depending on its cause: (1) Epistemic uncertainty, called model uncertainty, is derived from model parameters. (2) Aleatoric uncertainty, called data uncertainty, is originated from the inherent noise of data. Epistemic uncertainty can be reduced by providing sufficient data, while aleatoric uncertainty is irreducible with supplementary data. Although some works [5, 37, 43, 50, 51] have attempted to define uncertainty through the variance from their dataset and handle it, there is no direct applicable approach to VLM in conventional computer vision tasks due to the lack of linguistic datasets. In this work, we explore the aleatoric uncertainty of language with only the image-modality dataset.

3. Background

Contrastive Language-Image Pre-training (CLIP) [40] is a powerful vision-language pre-trained model that learns to align image-text representations via contrastive learning [35]. It considers the relevant image and class text description pairs $\{\mathbf{x}, \mathbf{t}_c\}$ as positive samples and the non-relevant pairs as negative samples. The contrastive objective is to maximize the similarity of the positive samples while minimizing the similarity of negative samples. To be specific, it consists of an image encoder \mathcal{F} and a text encoder \mathcal{G} . Given \mathbf{x} , it measures the similarity between image and class-embedded text representations $\mathbf{w}_{1:C} = \{w_1, \dots, w_C\}$, where $w_c = \mathcal{G}(\mathbf{t}_c)$ and C denotes the number of classes. Then, the prediction probability that \mathbf{x} belongs to class c is computed as follows:

$$p(y_c|\mathbf{x}, \mathbf{w}_{1:C}) = \frac{\exp(\mathcal{F}(\mathbf{x}) \cdot w_c^\top / \tau)}{\sum_{i=1}^C \exp(\mathcal{F}(\mathbf{x}) \cdot w_i^\top / \tau)}, \quad (1)$$

where τ is a hyper-parameter.

Context Optimization (CoOp) [59] introduced a learning-based prompt method that learns task-relevant prompts for better transferability in downstream recognition tasks. The prompt \mathbf{p} is revised as learnable continuous

variables, which is updated by back-propagation with the pre-trained CLIP text encoder. In particular, the class description with \mathbf{p} is obtained by concatenating their text token as $\mathbf{t}_c(\mathbf{p})$. It can optimize \mathbf{p} using M training samples $\{\mathbf{x}_i, y_i\}_{i=1}^M$, by minimizing the following objective:

$$\mathcal{L}(\mathbf{p}) = \mathbb{E}_{\mathbf{x}_i, y_i} [-\log p(y_i|\mathbf{x}_i, \mathbf{w}_{1:C}(\mathbf{p}))], \quad (2)$$

where $\mathbf{w}_{1:C}(\mathbf{p}) = \{w_1(\mathbf{p}), \dots, w_C(\mathbf{p})\}$ and $w_c(\mathbf{p}) = \mathcal{G}(\mathbf{t}_c(\mathbf{p}))$.

Prompt Distribution Learning (ProDA) [32] aims to learn the distribution of diverse prompts to handle various visual representations. They assume that the prompt distribution $p(\mathbf{p})$ can be modeled as a Gaussian distribution. Specifically, they define a set of K learnable prompts as $\mathbf{P} = \{\mathbf{p}^1, \dots, \mathbf{p}^K\}$, and model the distribution of $\mathbf{w}_{1:C}(\mathbf{P}) \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{w}_{1:C}(\mathbf{P})}, \boldsymbol{\sigma}_{\mathbf{w}_{1:C}(\mathbf{P})} \mathbf{I})$ with a mean and diagonal covariance of $\mathbf{w}_{1:C}(\mathbf{P}) = \{\mathbf{w}_{1:C}(\mathbf{p}^1), \dots, \mathbf{w}_{1:C}(\mathbf{p}^K)\}$. To learn an optimal prompt distribution, \mathbf{P} is optimized by minimizing the classification loss as:

$$\mathcal{L}(\mathbf{P}) = \mathbb{E}_{\mathbf{x}_i, y_i} \left[-\log \mathbb{E}_{\mathbf{w}_{1:C}(\mathbf{P})} [p(y_i|\mathbf{x}_i, \mathbf{w}_{1:C}(\mathbf{P}))] \right]. \quad (3)$$

While the ProDA has achieved state-of-the-art results in a large variety of downstream tasks, it is vulnerable to generating context-aware text representations.

DenseCLIP [41] leverages the pixel-text matching loss \mathcal{L}_{pixel} as a guidance of task objective. In particular, DenseCLIP suggests a post-model prompting approach, which directly adds visual-context bias $u = \text{TransDecoder}(\mathbf{w}_{1:C}(\mathbf{p}), \mathcal{F}(\mathbf{x}))$ to $\mathbf{w}_{1:C}$. It obtains the context-aware text representations as $\hat{\mathbf{w}}_{1:C} \leftarrow \mathbf{w}_{1:C} + \gamma u$, where γ is a learnable scale factor. Then, they utilize pixel-text contrastive loss as an auxiliary loss for task-specific loss, which is formulated follow as:

$$\mathcal{L}_{pixel} = \mathbb{E}_{\mathbf{x}_i, j, y_c} [-\log p(y_c|\mathbf{x}_{i,j}, \hat{\mathbf{w}}_{1:C})], \quad (4)$$

where i, j are each pixel location of input image \mathbf{x} . Although they utilize visual cues as context bias via context-aware prompting, a deterministic prompt is still not enough to deal with the randomness of visual representations.

4. Method

4.1. Overview

In this section, we present our proposed Probabilistic Prompt Learning (PPL) for dense predictions, illustrated in Fig. 2. Given an image \mathbf{x} , our goal is to predict plausible pixel-wise results y by taking advantage of general knowledge learned from VLM. It consists of an image encoder \mathcal{F} , a text encoder \mathcal{G} , and an image decoder \mathcal{D} for generating the

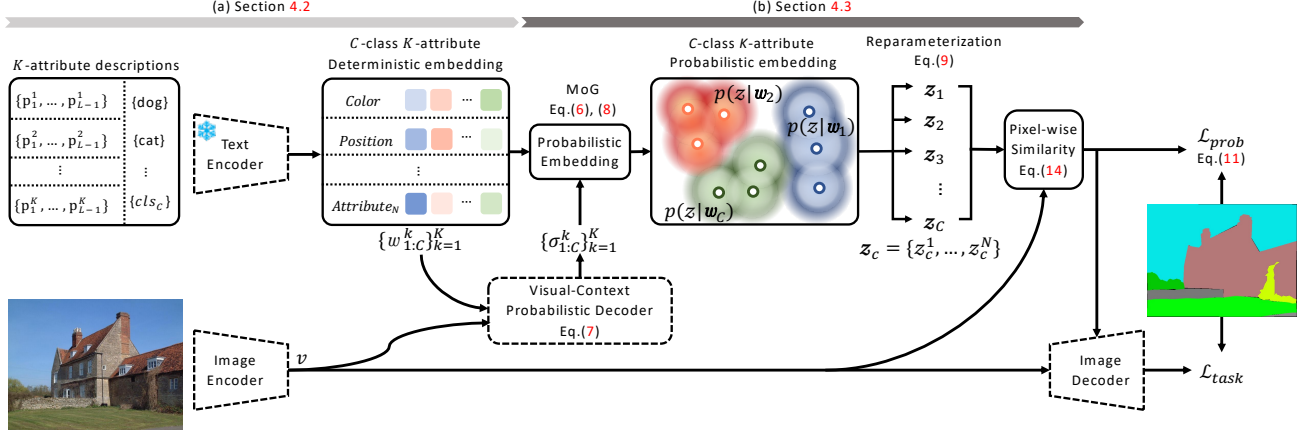


Figure 2. **Overall framework of PPL.** (a) The text encoder takes K -attribute prompts to generate the deterministic text embeddings w_c for each class c . The standard deviation σ_c^k of class-attribute embedding w_c^k is then computed via the visual-context probabilistic decoder. (b) Each class distribution is modeled as a Mixture of Gaussian (MoG), from which the text embeddings z_c are randomly sampled.

dense prediction results. We utilize multiple text representations, which are combinations of class-agnostic attribute prompts and object classes (Fig. 2 (a)). The visual and text features are fed into the visual-context probabilistic decoder to define a probabilistic embedding space (Fig. 2 (b)). The attributes’ distribution for each object class is represented as a Mixture of Gaussian (MoG), from which we sample text representations to boost the dense prediction task through pixel-text similarity loss.

4.2. Class-Agnostic Attribute Prompt

We first introduce class-agnostic attribute prompts to understand objects with diverse perspectives. It aims not only to learn diverse prompts from data but also to automatically assign efficient attributes that are universally available in all object classes. Specifically, it consists of a set of K learnable prompt templates $\mathbf{P} = \{\mathbf{p}^1, \dots, \mathbf{p}^K\}$, to describe various attributes across the object class. We define an association of prompt \mathbf{p}^k and class \mathbf{t}_c as $\mathbf{t}_c(\mathbf{p}^k)$, and set k -th attribute representations as $w_{1:C}^k = \mathcal{G}(\mathbf{t}_{1:C}(\mathbf{p}^k))$.

We further propose diversity loss to regularize each class-text representation to become different from others and prevent the multiple learned text representations from being identical. The diversity loss \mathcal{L}_{div} is formulated as:

$$\mathcal{L}_{div} = \frac{1}{C} \sum_{c=1}^C \max(\|w_c w_c^\top - \mathbf{I}\|_F^2 - b, 0), \quad (5)$$

where $w_c = \{w_c^1, \dots, w_c^K\}$ is the set of attribute representations of class c and $\|\cdot\|_F$ denotes Frobenius norm of a matrix. Note that we control the extent of overlap between attribute representations with a learnable vector $b \in \mathbb{R}^C$. When b is close to 0, each attribute representation is trained to be orthogonal. The representation set w_c is used to define a distribution, from which the probabilistic prompts for the class c are sampled.

4.3. Probabilistic Prompt Learning

We propose probabilistic prompt learning to infer the distribution of the class-attribute representations w_c^k that describe various visual-contexts for the target objects. From this perspective, for each w_c^k , we define a probability distribution $p(z|w_c^k)$ as a factorized Gaussian with its center vector μ_c^k and diagonal covariance matrix σ_c^k :

$$p(z|w_c^k) \sim \mathcal{N}(\mu_c^k, \sigma_c^k \mathbf{I}). \quad (6)$$

We set the center vector $\mu_c^k = w_c^k$, then exploit visual-context knowledge to calculate the covariance matrix σ_c^k .

Specifically, we design a visual-context probabilistic decoder, illustrated in Fig. 3. We feed the class-attribute representation w_c^k and visual embedding $v = \mathcal{F}(\mathbf{x})$ into a transformer module [8], whose output is fed into the multi-head self-attention layer to generate a new query q_c^k . Here, the key and value vectors are both obtained from the visual embedding v . The covariance matrix of each class-attribute representation is computed as:

$$\sigma_c^k = \text{MLP}(\text{LN}(q_c^k)) + \text{MLP}(\text{MHA}(q_c^k, v_k, v_v)), \quad (7)$$

where LN and MHA are layer norm and multi-head attention. Finally, we formulate the prompt distributions for the class as a Mixture of Gaussian (MoG) model such that

$$p(z|w_c) \sim \sum_{k=1}^K \mathcal{N}(\mu_c^k, \sigma_c^k \mathbf{I}). \quad (8)$$

It can be interpreted as a distribution of possible class-attribute representations which reflects visual-context knowledge from the input image.

From $p(z|w_c)$, we randomly sample N prompt representations for class c as $z_c = \{z_c^1, \dots, z_c^N\} \stackrel{i.i.d.}{\sim} p(z|w_c)$, and we treat z_c as “self-augmented” context descriptions. This

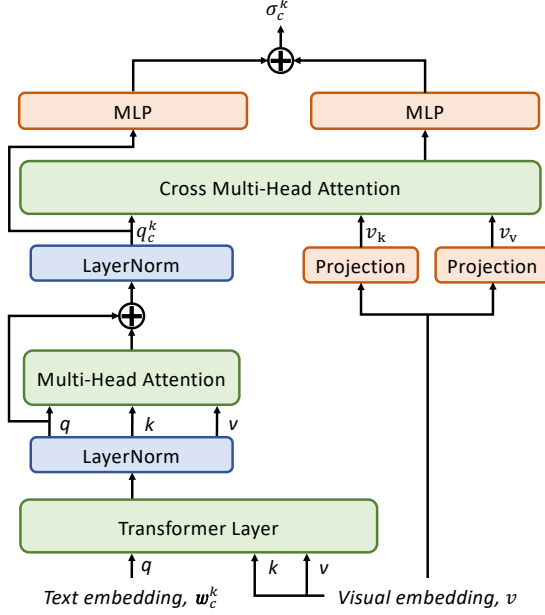


Figure 3. **Architecture of the visual-context probabilistic decoder.** It takes text embedding w_c^k and visual embedding v as inputs, and generates a standard deviation σ_c^k of attribute distributions $p(z|w_c^k) \sim \mathcal{N}(\mu_c^k, \sigma_c^k \mathbf{I})$.

sampling process is based on reparameterization trick [26]

$$z_c^n = \mu(w_c) + \epsilon^n \sigma(w_c), \quad (9)$$

where $\mu(w_c)$, $\sigma(w_c)$ are mean and standard deviation of $p(z|w_c)$, and $\epsilon^n \sim \mathcal{N}(0, I)$. Finally, we transfer the text knowledge to the dense prediction task by pixel-text matching loss \mathcal{L}_{pixel} :

$$\mathcal{L}_{pixel} = \mathbb{E}_{\mathbf{x}_{i,j}, y_c} \left[-\log \mathbb{E}_{\mathbf{w}_{1:C}} [p(y_c | \mathbf{x}_{i,j}, \mathbf{z}_{1:C})] \right], \quad (10)$$

where i and j denote the pixel location in the image.

The assumption that text representations can be modeled as Gaussian distribution seems similar to that in ProDA [32]. However, in [32], the distribution is simply estimated using the statistics of multiple prompts such as $\mathbf{w}_{1:C}(\mathbf{P}) \sim \mathcal{N}(\mu_{\mathbf{w}_{1:C}(\mathbf{P})}, \sigma_{\mathbf{w}_{1:C}(\mathbf{P})} \mathbf{I})$. In contrast, we model independent class-attribute distributions from multiple prompts considering the contextual information from image and text, and represent the class-specific prompt distribution as a mixture of distributions. Consequently, our method not only captures the diversity of visual representations but also alleviates the bias problem caused by aleatoric uncertainty, providing better generalization to the downstream tasks.

4.4. Training

Handling cross-modal uncertainty. Although *a priori* distribution $p(\mathbf{w}_{1:C})$ of class representations can provide granular text representations, high uncertainty resulting

from complex and diverse visual representation attenuates the reliability of predictive probabilities. To mitigate the negative impact of uncertainty, we replace Eq. (10) with probabilistic pixel-text matching loss \mathcal{L}_{prob} , defined as:

$$\mathcal{L}_{prob} = \frac{1}{(\prod_c \sigma(w_c)^2)^{1/C}} \mathcal{L}_{pixel} + \frac{1}{2C} \sum_c \log \sigma(w_c)^2, \quad (11)$$

where the denominator $(\prod_c \sigma(w_c)^2)^{1/C}$ in the first term is to alleviate the penalty originating from uncertainty, and the second term reduces the high uncertainties.

Total objectives. We employ additional KL divergence loss \mathcal{L}_{KL} between each class distribution and Gaussian prior distribution $\mathcal{N}(0, I)$ to prevent the learned variance from collapsing to zero, inspired by [5, 37]:

$$\mathcal{L}_{KL} = \frac{1}{C} \sum_c \text{KL}(p(z|w_c) \parallel \mathcal{N}(0, I)). \quad (12)$$

Therefore, with the task-specific loss, \mathcal{L}_{task} for dense prediction task, the overall objective of the proposed PPL is a weighted summation of all loss functions defined as:

$$\mathcal{L}_{PPL} = \mathcal{L}_{task} + \mathcal{L}_{prob} + \alpha \mathcal{L}_{div} + \beta \mathcal{L}_{KL}, \quad (13)$$

where α and β are hyper-parameters. Note that \mathcal{L}_{KL} and uncertainty regularization term in \mathcal{L}_{prob} have opposite objectives: \mathcal{L}_{KL} prevents $\sigma(w_c)$ from collapsing to zero, while the uncertainty regularization term aims to reduce $\sigma(w_c)$. The balance of these terms is controlled by β .

Inference. Given the learned prompts set $\{\mathbf{p}^1, \dots, \mathbf{p}^K\}$, class representation z_c follows $\mathcal{N}(\mu(w_c), \sigma(w_c))$. The prediction probability of the input image \mathbf{x} is formulated by $\mathbb{E}_{\mathbf{w}_{1:C}} [p(y|\mathbf{x}_{i,j}, \mathbf{z}_{1:C})]$. While computing the prediction probability is intractable over $p(\mathbf{w}_{1:C})$, it can be factorized via Monte-Carlo estimation, defined as:

$$\mathbb{E}_{\mathbf{w}_{1:C}} [p(y|\mathbf{x}_{i,j}, \mathbf{z}_{1:C})] \approx \frac{1}{N} \sum_n p(y|\mathbf{x}_{i,j}, \mathbf{z}_{1:C}^n). \quad (14)$$

5. Experiments

We present the experimental results to demonstrate the effectiveness of the proposed PPL. We conduct comparisons with the state-of-the-art methods on different dense prediction tasks: semantic segmentation, object detection, and instance segmentation. Then, we provide the results of extensive ablation studies.

5.1. Experimental Settings

For prompt learning, we set context length $L = 8$ and initialize the context vectors using Gaussian noise. We freeze the text encoder to conserve the pre-trained language knowledge. The visual-context probabilistic decoder

Table 1. **Quantitative results of semantic segmentation on ADE20k dataset** under different pre-trained dataset and backbone network settings. (*: Results are directly taken from [41].)

Backbone	Method	Pre-train	mIoU-SS	mIoU-MS	GFLOPs	Params(M)
ResNet-50	FCN [31]	ImageNet	36.1	38.1	793.6	49.6
	PSPNet [55]	ImageNet	41.1	41.9	716.2	49.1
	DeepLab V3+ [4]	ImageNet	42.7	43.8	711.5	43.7
	UperNet [48]	ImageNet	42.1	42.8	953.2	66.5
	Semantic FPN* [27]	ImageNet	38.6	40.6	227.1	31
	CLIP+Semantic FPN* [40]	CLIP	39.6	41.6	248.8	31
	DenseCLIP+Semantic FPN* [41]	CLIP	43.5	44.7	269.2	50.3
	PPL + Semantic FPN	CLIP	44.7	45.8	421.4	51.8
ResNet-101	FCN [31]	ImageNet	39.9	41.4	1104.4	68.6
	PSPNet [55]	ImageNet	43.6	44.4	1027.4	68.1
	DeepLab V3+ [4]	ImageNet	44.6	46.1	1022.7	62.7
	UperNet [48]	ImageNet	43.8	44.8	1031	85.5
	Semantic FPN* [27]	ImageNet	40.4	42.3	304.9	50
	CLIP+Semantic FPN* [40]	CLIP	42.7	44.3	236.6	50
	DenseCLIP+Semantic FPN* [41]	CLIP	45.1	46.5	346.23	67.9
	PPL + Semantic FPN	CLIP	46.4	47.8	496.9	69.4
ViT-B	SETR-MLA-DeiT [56]	ImageNet	46.2	47.7	-	-
	Semantic FPN* [27]	ImageNet	48.3	50.9	937.4	100.8
	Semantic FPN* [27]	ImageNet-21K	49.1	50.9	937.4	100.8
	CLIP+Semantic FPN* [40]	CLIP	49.4	50.4	937.4	100.8
	DenseCLIP+Semantic FPN* [41]	CLIP	50.6	50.3	933.1	105.3
	PPL + Semantic FPN	CLIP	51.6	51.8	1072.1	106.9

Table 2. **Semantic segmentation performance on ADE20k of ProDA [32] and the proposed method.**

Model	mIoU-SS	GFLOPs	Params(M)
ProDA+Semantic FPN [32]	42.6	1379	46.4
PPL+Semantic FPN	44.7	421	51.8

is composed of a transformer module with 5 layers and 1 variance prediction module. We train our network with AdamW optimizer. In the comparison experiments, we set the number of attributes $K = 3$, and sampled the embeddings $N = 15$ fairly and evenly from each of the K distributions, and the KL-divergence hyperparameter $\beta = 10^{-5}$ throughout the experiments. We also include FLOPs and the number of parameters for fair comparisons. Additional settings and implementation details for each experiment are presented in each subsection.

5.2. Semantic Segmentation

Settings. We evaluate the semantic segmentation performance of PPL on ADE20k [57], which contains 20k training and 2k validation images with 150 categories. Following the common protocol in [19, 48], we evaluate the performance on the validation set, using mIoU scores measured in single scale (mIoU-SS) and multiple scales (mIoU-MS). We adopt the Semantic FPN [27] framework, with different encoders: ResNet-50 (RN50), ResNet-101 (RN101) [15], and ViT-B [8]. The network is trained for 127 epochs with a batch size of 32 and a learning rate of 10^{-4} .

Results. Table 1 shows the semantic segmentation performance of different methods on ADE20k dataset [57]. The

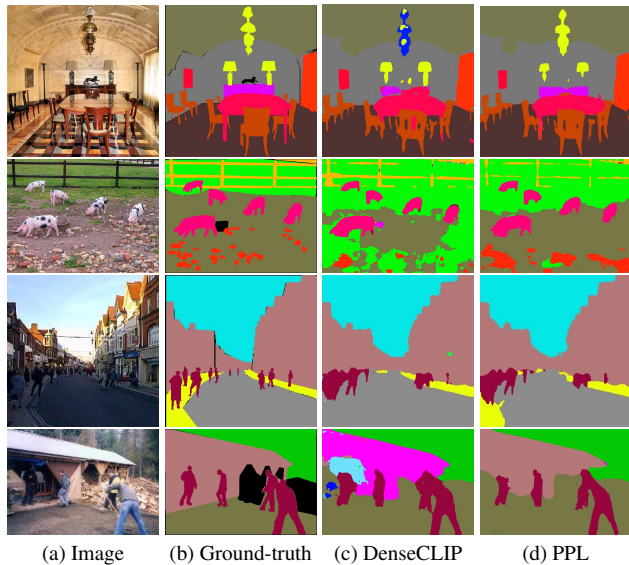


Figure 4. **Qualitative results** for semantic segmentation on ADE20k based on DenseCLIP and our proposed PPL.

comparison between DenseCLIP and our method shows that our probabilistic approach outperforms the deterministic approach by 1.2% for RN50, 1.3% for RN101, and 1% for ViT-B backbones, in terms of mIoU-SS. Furthermore, compared to the vanilla CLIP-based segmentation networks [41], our approach results in higher mIoU-MS by 5.1%, 3.7%, and 2.2% for RN50, RN101, and ViT-B encoders, respectively. We provide qualitative results of DenseCLIP [41] and the proposed method in Fig. 4. We observe that our method tends to capture fine-detailed objects

Table 3. **Quantitative results of object detection and instance segmentation on COCO** with different backbone networks.(*: Results are directly taken from [41].)

Model	GFLOPs	Params(M)	Object Detection						Instance Segmentation					
			AP^b	AP_{50}^b	AP_{75}^b	AP_S^b	AP_M^b	AP_L^b	AP^m	AP_{50}^m	AP_{75}^m	AP_S^m	AP_M^m	AP_L^m
RN50-IN1K [15]	275	44	38.2	58.8	41.4	21.9	40.9	49.5	34.7	55.7	37.2	18.3	37.4	47.2
RN50-CLIP* [40]	301	44	39.3	61.3	42.7	24.6	42.6	50.1	36.8	58.5	39.2	18.6	39.9	51.8
RN50-DenseCLIP* [41]	327	67	40.2	63.2	43.9	26.3	44.2	51	37.6	60.2	39.8	20.8	40.7	53.7
RN50-PPL	368	70	41.0	64.5	45.2	27.0	45.1	51.8	38.3	60.9	41.1	21.1	41.3	53.9
RN101-IN1K [15]	351	63	40.0	60.5	44.0	22.6	44	52.6	36.1	57.5	38.6	18.8	39.7	49.5
RN101-CLIP* [40]	377	63	42.2	64.2	46.5	26.4	46.1	54.0	38.9	61.4	41.8	20.5	42.3	55.1
RN101-DenseCLIP* [41]	399	84	42.6	65.1	46.5	27.7	46.5	54.2	39.6	62.4	42.4	21.4	43.0	56.2
RN101-PPL	444	87	43.2	66.0	47.0	28.1	47.1	54.8	40.4	63.3	43.2	21.8	43.7	57.0

Table 4. **Ablation studies** on component of objective functions.

\mathcal{L}_{div}	\mathcal{L}_{KL}	\mathcal{L}_{prob}	mIoU	Δ
			42.0	0
✓			42.4	+0.4
	✓		43.6	+1.6
✓	✓		44.3	+2.3
✓	✓	✓	44.7	+2.7

and segment the correct labels compared to DenseCLIP. Furthermore, since our method leverage various visual-context knowledge, we can reduce error in the ambiguous region to classify, and this is shown through Fig. 4.

Comparison with ProDA We further compare our approach with the existing probabilistic approach ProDA [32]. For a fair comparison, both methods use RN-50 backbone, the number of sampled prompts $N_{ProDA} = 16$, $N_{PPL} = 15$. The quantitative result is presented in Table 2. We observe higher mIoU-SS of the proposed method. Since the proposed PPL estimates the class-specific attribute distribution by considering the visual-text relationship, it is a more suitable segmentation. On the other hand, ProDA uses only textual information to model the distribution, so it tends to fail to capture the diverse and complex components.

5.3. Object Detection and Instance Segmentation

Settings. We evaluate the proposed PPL in object detection and instance segmentation on COCO dataset [30], which consists of 118k training samples and 5k validation images. We conduct experiments on our proposed method using Mask-RCNN architecture [14]. We report results using standard Average Precision metric measured using bounding box (AP^b) and segmentation mask (AP^m) with IoU= 0.5/0.75, and object sizes. We use RN50 and RN101 as backbone networks. The network is trained for 12 epochs with a batch size of 16 and 2×10^{-4} .

Results. We summarized the experimental results in Table 3. We observe the VLM methods outperform the conventional ImageNet-1K (IN1K) pre-trained model. The proposed PPL exploits multiple text representations to provide diverse visual-language knowledge to generate plausible results, and outperforms both Vanilla CLIP [59] and DenseCLIP [41] in both object detection and instance segmen-

Table 5. **Ablation studies** on hyperparameters. The number of attributes K , KL-divergence hyperparameter β , and the number of sampled representations N .

Number of attributes ($N = 15, \beta = 10^{-5}$)				
Parameter K	1	3	5	7
mIoU	43.8	44.7	44.2	44.0
KL-divergence hyperparameter ($N = 15, K = 3$)				
Parameter β	10^{-7}	10^{-6}	10^{-5}	10^{-4}
mIoU	42.4	42.8	44.7	44.1
Number of sampled representations ($\beta = 10^{-5}, K = 3$)				
Parameter N	5	10	15	20
mIoU	43.9	44.4	44.7	44.8

tation. Showing consistent improvements across the sub-measures, we confirm that the diverse expressions of visual-context from the PPL help the network to generate more accurate predictions across different scales with precise localization.

5.4. Ablation Study and Analysis

To further analyze and validate the components of our method, we conduct ablation study experiments on semantic segmentation. In addition, to show the impact of multiple attributes prompts, we visualize the activation maps and analyze them in detail.

Contribution of objectives. We conduct experiments to observe the effect of each loss function in (13). To this end, we train networks with different objective functions and present the results in Table 4. We first observe 42.0 mIoU of a baseline network trained using only \mathcal{L}_{task} and \mathcal{L}_{pixel} in (10). Applying each of the diversity loss \mathcal{L}_{div} and KL divergence loss \mathcal{L}_{KL} encourages the diversity in probabilistic distribution, resulting in improved performance. The performance is further improved when they are applied simultaneously. Finally, replacing \mathcal{L}_{pixel} with \mathcal{L}_{prob} regularizes high uncertainties and results in further improved performance, achieving 44.7 mIoU. The analysis of \mathcal{L}_{prob} is described in the last paragraph of Sec. 5.4.

Number of attribute prompts. We compare the semantic segmentation performance with respect to the number of attributes prompts K . As shown in the first block of Table 5, $K = 3$ notably outperforms $K = 1$, which justifies the introduction of multiple attribute representations. However,

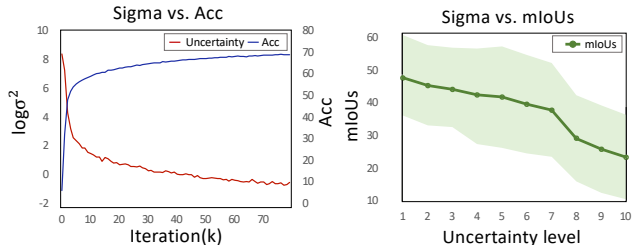


Figure 5. **Illustration of uncertainty analysis** on semantic segmentation. (Left) uncertainty versus accuracy during training, (Right) performance versus uncertainty level for the test set.

we find that the performance is rather degraded for $K > 3$ since many attributes are likely to provide redundant information, which degrades performance.

KL-divergence hyperparameters. To investigate the effect of the KL-divergence hyperparameter β in Eq. (13), we include the experimental result in the second block of Table 5. In general, the variance of a mixed distribution follows the unit variance as β increases, reducing the discriminability of distributions. Conversely, if β is too small (e.g. 10^{-6}), variance converges to zero and \mathcal{L}_{prob} diverges. In summary, β controls the range of use of the visual-context, and also adjusts \mathcal{L}_{KL} and \mathcal{L}_{prob} , which operate opposite to each other, learn in a balanced way. We find that our model achieves the best performance with $\beta = 10^{-5}$.

Number of sampled representations. To analyze the effect of the number of sampled representations N , we conduct experiments with 5, 10, 15, and 20 samples. The results are summarized in the third block of Table 5. We can observe that the performance increases as the number of sampled representations increases. For example, there was a 0.8% improvement in performance when $N = 15$ compared to when $N = 5$. However, due to the increase in computational cost as the number of sampled representations increases, we fixed $N = 15$ to balance between the computational cost and performance.

Uncertainty vs. Performance. To analyze the correlation between uncertainty and the prediction probability (14), we measured the uncertainty of text representations and report the performance in terms of segmentation accuracy, according to the training iterations. We define the uncertainty as the geometric mean of the variance $\sigma(w_{1:C})$ given the input image. As shown in the left plot in Fig. 5, the uncertainty is minimized and performance increases as the learning progress. Concretely, the network learns to focus on useful visual-contexts and neglect the redundant ones.

In addition, we divide the uncertainty value into 10 bins and measure the performance according to its level. We observed a negative correlation between uncertainty and performance. As shown in the right plot in Fig. 5, we observe that high uncertainty results in unreliable predictions, thus we suppress the high uncertainties via (14) to generate the improved results.

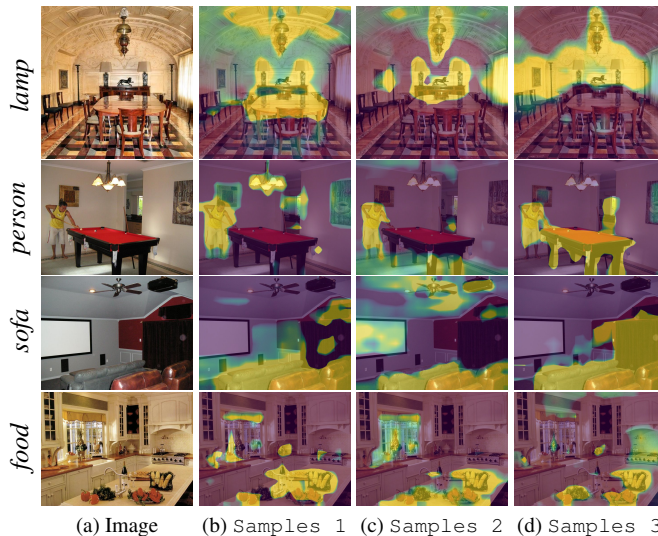


Figure 6. **Visualization of activation maps.** We report the activation maps of each attribute distribution (mean) of different classes indicated on the left side, with $K = 3$ on the ADE20k dataset. Different samples from MoG attend to different attributes of the object.

Visualization To better understand the advantage of probabilistic embedding, we demonstrate some visualization examples of the activation maps derived from each sample text representation in Fig. 6. We observed that the attended region by sampled representations associated with different contexts changes according to the given context. This means that these multiple prompts can efficiently represent objects of various shapes, sizes, and colors, and consequently, our method is suitable for dense prediction tasks.

6. Conclusion

In this paper, we presented a probabilistic prompt learning (PPL) for dense prediction. It aims to extract diverse text representations to fully exploit the knowledge from VLM. Leveraging the visual-context information, class-specific probabilistic text distribution is defined, from where diverse text representations are sampled to guide the dense prediction tasks. In addition, we learn the optimal distribution by suppressing the high uncertainty from the complex visual-context via the probabilistic pixel-text matching loss. The experimental results show that the proposed method achieves significantly improved performance compared to the previous method in semantic segmentation, object detection, and instance segmentation, demonstrating the potential extension of our method to comprehensive multi-modal scene understanding tasks.

Acknowledgement. This research was supported by the Yonsei Signature Research Cluster Program of 2022 (2022-22-0002) and the KIST Institutional Program (Project No.2E31051-21-203).

References

- [1] Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. Fusion of detected objects in text for visual question answering. In *EMNLP*, 2019. 2
- [2] Jie Chang, Zhonghao Lan, Changmao Cheng, and Yichen Wei. Data uncertainty learning in face recognition. In *CVPR*, 2020. 3
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4):834–848, 2017. 1
- [4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 6
- [5] Sanghyuk Chun, Seong Joon Oh, Rafael Sampaio De Rezende, Yannis Kalantidis, and Diane Larlus. Probabilistic embeddings for cross-modal retrieval. In *CVPR*, 2021. 2, 3, 5
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1
- [7] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *ICCV*, 2021. 2
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1, 4, 6
- [9] Chengjian Feng, Yujie Zhong, Zequn Jie, Xiangxiang Chu, Haibing Ren, Xiaolin Wei, Weidi Xie, and Lin Ma. Promptdet: Expand your detector vocabulary with uncurated images. In *ECCV*, 2022. 2
- [10] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Dropblock: A regularization method for convolutional networks. In *NeurIPS*, 2018. 1
- [11] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 1
- [12] Fusheng Hao, Fengxiang He, Jun Cheng, Lei Wang, Jianzhong Cao, and Dacheng Tao. Collect and select: Semantic alignment metric learning for few-shot learning. In *ICCV*, 2019. 2
- [13] Kaiming He, Ross Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *ICCV*, 2019. 1
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 7
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 6, 7
- [16] Tao He, Lianli Gao, Jingkuan Song, and Yuan-Fang Li. Towards open-vocabulary scene graph generation with prompt-based finetuning. In *ECCV*, 2022. 2
- [17] Yicong Hong, Qi Wu, Yuankai Qi, Cristian Rodriguez-Opazo, and Stephen Gould. Vln bert: A recurrent vision-and-language bert for navigation. In *CVPR*, 2021. 2
- [18] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *ICML*, 2019. 2
- [19] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *ICCV*, 2019. 6
- [20] Zhicheng Huang, Zhaoyang Zeng, Yupan Huang, Bei Liu, Dongmei Fu, and Jianlong Fu. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *CVPR*, 2021. 2
- [21] 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 *ICML*, 2021. 1, 2
- [22] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *ECCV*, 2022. 2
- [23] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *ICCV*, 2021. 2
- [24] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, 2017. 3
- [25] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. 2
- [26] Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *NeurIPS*, 2015. 5
- [27] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *CVPR*, 2019. 6
- [28] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 2
- [29] Yandong Li, Di Huang, Danfeng Qin, Liqiang Wang, and Boqing Gong. Improving object detection with selective self-supervised self-training. In *ECCV*, 2020. 1
- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 1, 7
- [31] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1, 6
- [32] Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. Prompt distribution learning. In *CVPR*, 2022. 2, 3, 5, 6, 7
- [33] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding language-image pretrained models for general video recognition. In *ECCV*, 2022. 2

- [34] Seong Joon Oh, Kevin Murphy, Jiyan Pan, Joseph Roth, Florian Schroff, and Andrew Gallagher. Modeling uncertainty with hedged instance embedding. In *ICLR*, 2019. 3
- [35] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [36] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *CVPR*, 2020. 1
- [37] Jungin Park, Jiyoun Lee, Ig-Jae Kim, and Kwanghoon Sohn. Probabilistic representations for video contrastive learning. In *CVPR*, 2022. 3, 5
- [38] Rudra PK Poudel, Stephan Liwicki, and Roberto Cipolla. Fast-scnn: Fast semantic segmentation network. In *BMVC*, 2019. 1
- [39] Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data. *arXiv preprint arXiv:2001.07966*, 2020. 2
- [40] 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 *ICML*, 2021. 1, 2, 3, 6, 7
- [41] 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 *CVPR*, 2022. 1, 2, 3, 6, 7
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 1
- [43] Luke Vilnis and Andrew McCallum. Word representations via gaussian embedding. In *ICLR*, 2015. 3
- [44] Jinpeng Wang, Yuting Gao, Ke Li, Jianguo Hu, Xinyang Jiang, Xiaowei Guo, Rongrong Ji, and Xing Sun. Enhancing unsupervised video representation learning by decoupling the scene and the motion. In *AAAI*, 2021. 2
- [45] Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. Cris: Clip-driven referring image segmentation. In *CVPR*, 2022. 1
- [46] 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 *ECCV*, 2022. 2
- [47] Qiaolin Xia, Haoyang Huang, Nan Duan, Dongdong Zhang, Lei Ji, Zhifang Sui, Edward Cui, Taroon Bharti, and Ming Zhou. Xgpt: Cross-modal generative pre-training for image captioning. In *CCF*, 2021. 2
- [48] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018. 6
- [49] Boyu Yang, Chang Liu, Bohao Li, Jianbin Jiao, and Qixiang Ye. Prototype mixture models for few-shot semantic segmentation. In *ECCV*, 2020. 2
- [50] Gengcong Yang, Jingyi Zhang, Yong Zhang, Baoyuan Wu, and Yujiu Yang. Probabilistic modeling of semantic ambiguity for scene graph generation. In *CVPR*, 2021. 3
- [51] Tianyuan Yu, Da Li, Yongxin Yang, Timothy M Hospedales, and Tao Xiang. Robust person re-identification by modelling feature uncertainty. In *ICCV*, 2019. 3
- [52] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Bayer. Lit: Zero-shot transfer with locked-image text tuning. In *CVPR*, 2022. 2
- [53] Xiaohang Zhan, Ziwei Liu, Ping Luo, Xiaoou Tang, and Chen Loy. Mix-and-match tuning for self-supervised semantic segmentation. In *AAAI*, 2018. 1
- [54] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *CVPR*, 2022. 2
- [55] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. 6
- [56] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*, 2021. 6
- [57] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 127(3):302–321, 2019. 6
- [58] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, 2022. 1, 2
- [59] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 130(9):2337–2348, 2022. 1, 2, 3, 7
- [60] Yanzhao Zhou, Xin Wang, Jianbin Jiao, Trevor Darrell, and Fisher Yu. Learning saliency propagation for semi-supervised instance segmentation. In *CVPR*, 2020. 1
- [61] Barret Zoph, Golnaz Ghiasi, Tsung-Yi Lin, Yin Cui, Hanxiao Liu, Ekin Dogus Cubuk, and Quoc Le. Rethinking pre-training and self-training. In *NeurIPS*, 2020. 1