

# Open-Vocabulary Semantic Segmentation with Decoupled One-Pass Network

Cong Han<sup>1\*</sup> Yujie Zhong<sup>1\*</sup> Dengjie Li<sup>1</sup> Kai Han<sup>2†</sup> Lin Ma<sup>1†</sup>  
<sup>1</sup>Meituan Inc. <sup>2</sup>The University of Hong Kong

hancong0911@163.com jaszhong@hotmail.com kaih anx@hku.hk

## Abstract

Recently, the open-vocabulary semantic segmentation problem has attracted increasing attention and the best performing methods are based on two-stream networks: one stream for proposal mask generation and the other for segment classification using a pre-trained visual-language model. However, existing two-stream methods require passing a great number of (up to a hundred) image crops into the visual-language model, which is highly inefficient. To address the problem, we propose a network that only needs a single pass through the visual-language model for each input image. Specifically, we first propose a novel network adaptation approach, termed *patch severance*, to restrict the harmful interference between the patch embeddings in the pre-trained visual encoder. We then propose *classification anchor learning* to encourage the network to spatially focus on more discriminative features for classification. Extensive experiments demonstrate that the proposed method achieves outstanding performance, surpassing state-of-the-art methods while being 4 to 7 times faster at inference. Code: <https://github.com/CongHan0808/DeOP.git>

## 1. Introduction

Semantic segmentation is a critical computer vision task that entails grouping image pixels into semantically significant regions and predicting their class labels. Previously, semantic segmentation networks have concentrated on a pre-defined set of semantic classes based on the dataset. Recently, more attention has been devoted to open-vocabulary (also known as zero-shot) semantic segmentation, thanks to the rise of large-scale pre-trained visual-language models (VLMs) like CLIP [32] and ALIGN [19].

OpenSeg [15] was among the first methods for the open-vocabulary semantic segmentation task, which proposes to adopt class-agnostic masks for possible semantic regions and then classify them using text embeddings extracted from a pre-trained VLM (ALIGN [19] in its case). It is ef-

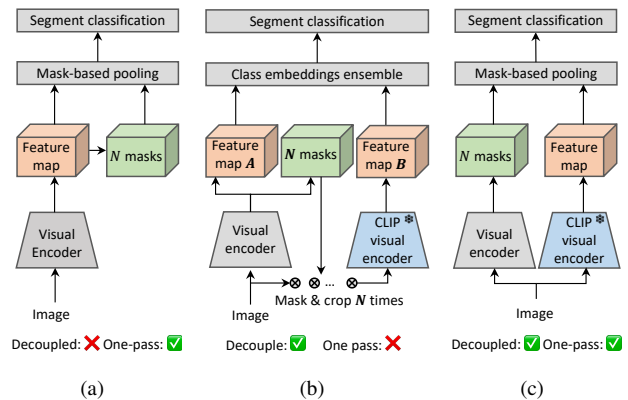


Figure 1: Comparisons between three macro-architectures for open-vocabulary semantic segmentation: (a) **coupled** network, e.g. OpenSeg [15]; (b) **decoupled multi-pass** network, e.g. SimBaseline [38]; (c) **decoupled one-pass** network (our baseline). The decoupled structure can maintain the generality of the pre-trained VLM, while the one-pass mechanism brings high computational efficiency.

ficient since it only requires passing the image through the visual encoder once for segmentation. **However, the single shared visual encoder that is updated during training inevitably breaks the original visual-language alignment of the pre-trained VLM.** Therefore, it requires training on numerous semantic concepts to ensure its zero-shot capability. For instance, training on Localized Narratives [31] is necessary. We refer to OpenSeg as the *coupled one-pass* method. Another line of work, including ZegFormer [10], SimBaseline [38] and OVSeg [26], decouples the two sub-tasks and designs a two-stream architecture. On this note, the class-agnostic mask proposal stream and the mask classification stream use two separate backbones, preserving the generality of the VLM’s visual encoder. **However, the downside of such approaches is their high computational overhead since image crops obtained by proposal masks are fed to the visual encoder individually.** We refer to methods like these as *decoupled multi-pass* methods.

In this work, we attempt to answer a question: can we develop a framework that can maintain the zero-shot abil-

\*Equal contribution.

†Corresponding authors.

ity of the VLM while being computationally efficient? To explore this possibility, we first design a baseline architecture based on two key design principles - *decoupled* and *one-pass*. As depicted in Figure 1(c), our baseline approach features two decoupled visual backbones, while the visual encoder of the pre-trained VLM (*i.e.* CLIP in this work) remains fixed to retain its generalization ability. Additionally, to achieve high efficiency, only one pass is required for the image in the classification stream. However, we empirically find that the performance of this baseline architecture is far from ideal. We further apply prompt learning to adapt both the CLIP visual encoder and text encoder to the segmentation task at hand, but the performance gap between the baseline model and previous decoupled methods is still significant.

After conducting a thorough analysis, we discover that the main reason for poor classification performance lies in the segment classification rather than the quality/recall of the masks. Specifically, we have identified two main problems that contribute to this issue: **(1) The patch embeddings that belong to different segments interact too much with each other in the original CLIP visual backbone.** In contrast, in multi-pass decoupled methods, patch embeddings in different masks have no interaction at all when passing through the VLM. Although patches belonging to different masks can provide some context information for classification, excessive interaction between different segment embeddings can actually harm the classification performance. **(2) The final segment embedding, which is obtained by mask-based pooling of patch embeddings, is not optimal for classification.** This is because the mask proposal network is trained in a class-agnostic manner, and the weightings in the mask only indicate how likely each patch embedding belongs to a particular segment, regardless of the category. Therefore, pooling the patch embeddings based on these masks is sub-optimal for classification. Therefore, by addressing these issues, we believe that it is possible to improve the classification performance of one-pass segmentation models while maintaining their computational efficiency.

To this end, we propose a two-stream framework, called Decoupled One-Pass network (DeOP), with the Generalized Patch Severance (GPS) and Classification Anchor Learning (CAL) to alleviate the above two problems, respectively. **Generalized Patch Severance** can be seen as a ‘severance’ operation on the patch tokens/embeddings in the CLIP visual encoder (*e.g.* ViT-Base). It aims to reduce the harmful interference between patch tokens in the encoder, while maintaining the embedding space of CLIP. **Classification Anchor Learning** is designed to find patches that are more suitable for segment classification, which we term classification anchors. It is achieved by appending a module at the end of the CLIP visual encoder, and the module learns

to generate a heatmap for each mask proposal, indicating which patch embeddings should be focused (*i.e.* the anchors) in the following spatial pooling for classification.

We extensively experiment on public benchmarks and show that DeOP consistently outperforms previous methods in both intra- and cross-dataset evaluation, while being significantly more efficient than other multi-pass methods (*e.g.*, SimBaseline), validating the effectiveness of our proposed GPS and CAL.

## 2. Related Work

**Vision-language pre-training.** Vision-language pre-training aims to connect image-text representations for tasks such as image-text retrieval, visual question answering, text-to-image generation, and dense prediction. One prominent pre-trained model, CLIP [32], learns image representations from a vast internet dataset of 400 million (image, text) pairs, delivering superior zero-shot image classification performance. To improve performance of pre-trained VLM for downstream tasks, text prompt learning has been proposed as a means to adapt models to new tasks by adding extra learnable tokens to the text prompt. Text prompt learning achieves strong generalization across various tasks in both few-shot and zero-shot settings, as demonstrated by studies such as [47, 21, 14].

**Semantic segmentation.** Semantic segmentation is a fundamental task in computer vision used to cluster parts of an image that belong to the same category. Fully Convolutional Networks (FCNs) [33] represent a seminal work in deep-net-based semantic segmentation, formulating it as a per-pixel classification task. Various variants, such as ASPP [5, 6], PPM [43], and OCNet [39], have followed FCNs. In recent years, MaskFormer [9, 8] has proposed predicting a set of binary masks, each associated with a single global class label prediction.

**Open-vocabulary semantic segmentation.** Open-vocabulary semantic segmentation learns pixel-wise classifiers for unseen object categories using either per-pixel-based (*e.g.*, ZS3Net [2], CaGNet [16], SPNet [36], Fusioner [29]) or region-based semantic segmentation (*e.g.* OpenSeg [15], ZegFormer [10], SimBaseline [38], OVSeg [26] and ODISE [37]). ZS3Net uses a generative model while CaGNet incorporates context-aware feature generation. SPNet utilizes a fixed word embedding projection matrix. ZegFormer, SimBaseline as well as two concurrent works, OVSeg and ODISE, generate per-segment embeddings and then perform class-agnostic grouping and per-segment open-vocabulary classification. This work adopts a decoupled architecture similar to them. Notably, the decoupled architecture can be easily combined with more advanced networks (such as SAM-series [22, 44, 40] or Mask DINO [24]) for better mask

proposal generation.

### 3. Baseline Method

**Problem definition.** We follow the settings of zero-shot segmentation in ZS3Net [2], which defines all classes as  $C = C_s \cup C_u$  and  $C_s \cap C_u = \emptyset$ , where  $C_s$  is the category set of seen classes and  $C_u$  is the set of unseen classes. In training stage, the model needs to project the images into different regions with seen classes. During inference, the model predicts the results for all seen and unseen classes.

As discussed, we aim to tackle the open-vocabulary semantic segmentation task based on two key principles: one is to keep the generalization ability of the pre-trained VLM for unseen classes, and the other one is to achieve high computational efficiency. In the following, we first propose a baseline network which possesses the two desired abilities (Section 3.1), and then present an improved baseline method (*i.e.* baseline+) which is obtained by applying prompt learning on the baseline (Section 3.2). Lastly, in Section 3.3, we discuss the limitations of the proposed baseline method.

#### 3.1. Baseline Network Architecture

Figure 1(c) shows the decoupled one-pass architecture of the baseline method. The network contains two streams: a class-agnostic mask proposal network and a mask classification network. The mask proposal network groups pixels to a series of class-agnostic masks. With this proposal network, we can convert the pixel-level classification to the region-level classification. The mask classification branch consists of a frozen CLIP visual encoder and classifiers obtained by the CLIP text encoder.

Specifically, we feed an image into the mask proposal network and CLIP visual encoder to generate  $N$  proposal masks  $M$  and the image feature map respectively. Then the  $N$  proposal masks are applied to the CLIP feature map to obtain  $N$  feature embeddings via mask-based pooling [15]. On the other hand, the class embeddings  $F_C$  (*i.e.* the classifiers) are obtained by computing the dot product between the visual features  $F_I$  with text the feature embedding  $F_T$  generated by CLIP text encoder which encodes the text categories. Finally, we get the final prediction  $O$  by combining the  $N$  proposal masks with their class predictions using matrix multiplication. We can formulate these as  $F_C = F_T * F_I$  and  $O = F_C * M$ .

**Mask proposal network based on MaskFormer.** We build our mask proposal network by adapting MaskFormer [9] into a class-agnostic variant. Unlike per-pixel classification segmentation models, MaskFormer is a mask classification model that predicts a set of masks, each associated with a single global class label. To make MaskFormer class-agnostic, we simply remove the class embed-

ding component while retaining the mask-related components. (Please refer to supplementary for a detailed description of the modified MaskFormer’s structure).

**Training of baseline.** In our baseline method, only the mask proposal network needs training, whereas the classification branch requires no training. During training of the proposal network, only the mask loss is employed, and classification loss is removed. Additionally, we learn a text prompt in an offline manner, same as [38].

#### 3.2. Improving Baseline with Prompt Learning

Notably, since the CLIP visual encoder is frozen in order to retain its zero-shot ability, a straightforward approach to enhance the classification performance of the baseline is prompt learning [47]. Prompt learning can slightly modify the embedding space of CLIP such that it can better adapt to the task at hand without breaking its visual-language alignment. In this work, we adopt both visual prompt tuning and text prompt learning to enhance our baseline, and the resultant method is denoted by baseline+.

**Visual prompt tuning.** Visual Prompt Tuning [20] (VPT) is an efficient and effective alternative to full fine-tuning for large-scale Transformer models in vision. We follow VPT to finetune the pre-trained CLIP visual encoder model. In specific, we add a set of learnable prompts to the sequence of the image patches embeddings in an element-wise manner, which can slightly steer the CLIP visual embedding space after training, so that it is more suitable for the segmentation task.

**Online text prompt learning.** In SimBaseline, the text prompts are learned offline via a classification task. To further improve the segmentation performance, we propose to train them with the whole segmentation network in an end-to-end manner. Specifically, the learnable text prompts are first initialized offline as [38], and then trained together with the visual prompts for the task of semantic segmentation.

**Training of baseline+.** The baseline+ requires end-to-end training to learn the visual and text prompts for the CLIP encoders. As it is trained for segmentation, we adopt a combination of dice loss [25] and focal loss [27], which is a common practice for semantic segmentation.

#### 3.3. Limitations

Experimentally, the performance gap between the baseline+ approach and previous decoupled approaches (*i.e.* ZegFormer and SimBaseline) remains substantial, as shown in Table 3. It can be deduced that the gap is mainly caused by the segment classification, since their mask proposal networks are similar. In this work, we consider two problems that lead to the inferior classification performance: (1) The original CLIP visual encoder may not

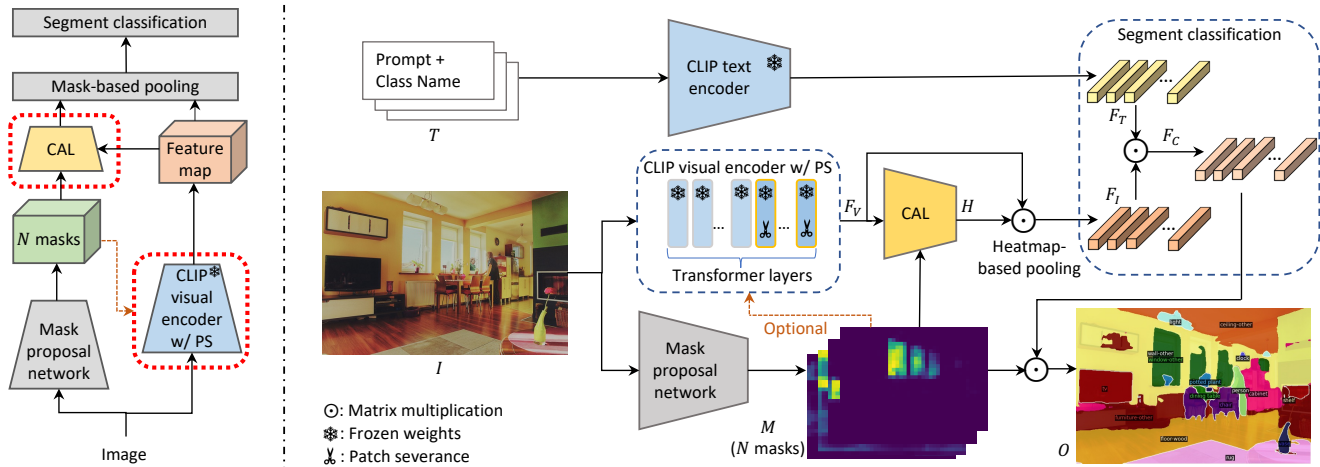


Figure 2: **Left:** the overview of DeOP. Comparing to our baseline network, the additional patch severance (PS) and classification learning (CAL) module are indicated by red dotted line. **Right:** the detailed pipeline. We feed an image to mask proposal network to generate  $N$  proposal masks  $M$ , and the image also goes into the CLIP visual encoder (with PS) to obtain visual features  $F_V$ .  $M$  and  $F_V$  are passed to the CAL module to produce classification anchor heatmaps  $H$  for  $M$ .  $H$  is used in the heatmap-based pooling to obtain  $F_I$ . The text embeddings  $F_T$  are generated by CLIP text encoder and used as classifiers for to get the classification scores for each mask. Finally, we get the segment prediction by combining  $F_C$  and  $M$ .

be aware of the correspondence between patch embeddings and semantic segments, and the patch embeddings belonging to different segments may overly/aimlessly interact in the encoder. Notice that this problem does not exist in previous decoupled methods since the patch embeddings belonging to different segments are not interacted, owing to the separate pass through the visual encoder of masked images. (2) The final segment embeddings produced by mask-based pooling on the patch embeddings are not ideal for classification. The mask proposal network is trained in a class-agnostic manner; therefore, the output masks may not be suitable for mask-based pooling for segment classification.

## 4. Our Method – DeOP

To mitigate the two problems discussed in Section 3.3, we propose a ViT-adapting approach coined patch severance to reduce the harmful interaction between patch embeddings in the CLIP visual encoder, and the Classification anchor learning (CAL) to produce better representations in the mask-based pooling for classification.

**Overall pipeline.** An overview of our framework, DeOP, which follows the two-stream design of our baseline method, is shown in Figure 2 (left). The two additional components, patch severance and CAL, are highlighted with red dotted lines. The proposed patch severance is applied to one or more layers in the CLIP visual encoder, while CAL is achieved by appending a query-based heatmap decoder after the CLIP visual encoder. The rest

of the pipeline (demonstrated in Figure 2 right), including the mask proposal network, classifiers generated from text and segment classification are the same as the baseline network. In terms of the training of DeOP, the same loss as baseline+ is adopted. Notably, at inference, DeOP does not need a heuristic ensemble of the classification scores from the classification stream and the mask proposal stream as required in [10, 38].

In the following, we elaborate on the proposed patch severance in Section 4.1 and classification anchor learning in Section 4.2.

### 4.1. Patch Severance

Generally, the feature maps extracted by the ViT-based CLIP visual encoder contain a large amount of global information of the whole image, due to the self-attention layer in the transformer layers. However, such property may be overwhelming for the semantic segmentation task.

The classification branch aims to classify each region. Although a certain extent of information from other regions may provide some context and benefit the classification, too much information may have negative effects and distract the classification.

To this end, we propose a ViT adaptation approach, termed **patch severance (PS)**, to restrict the interference between patch embeddings in the CLIP visual encoder. In the following, we compare two types of patch severance: mask-guided patch severance and generalized patch severance. They can be applied flexibly in one or more layers in the ViT-based visual encoder.

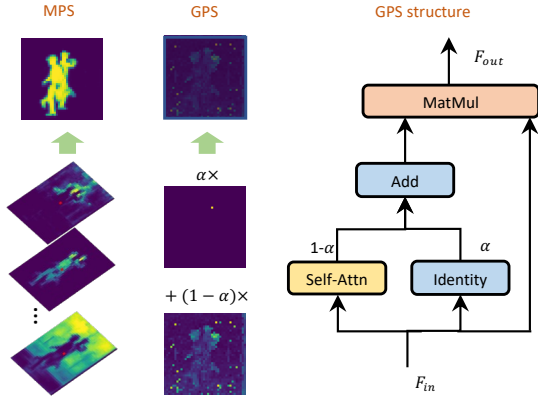


Figure 3: **Left:** the comparison between MPS and GPS in terms of how to obtain the final attention heatmap (for a particular patch embedding) in the self-attention layer. **Right:** the structure of GPS, which is used to replace the original self-attention layer in the transformer layer in ViT.

**Mask-guided patch severance.** Mask-guided Patch Severance (MPS) is an operation to guide the interaction between patch embeddings in the transformer layer. As Figure 3 (left) illustrates, MPS replaces the multiplication between key and query in the self-attention operation by proposal masks (generated by the mask proposal network), such that each patch embedding mainly interacts with those within the same predicted segment, while having little interaction between those belonging to other predicted segments. The mask-guided attention for each patch embedding is derived from the  $N$  proposal masks. Concretely, the mask-guided attention for the  $i$ th patch embedding is computed by Equation 1:

$$M_i = \sum_{j=1}^N M_{j,i} * M_j, \quad (1)$$

where  $M_i$  is the mask attention map for the  $i$ th patch embedding,  $M_j$  is the  $j$ th mask and  $M_{j,i}$  is the value of mask  $M_j$  for  $i$ th patch embedding.

**Generalized patch severance.** Different from MPS which utilizes proposal masks as a strong prior to guide the patch interaction, a Generalized Patch Severance (GPS) is proposed and it requires no prior (Figure 3 middle).

GPS replaces the query-key multiplication in the self-attention with a weighted combination of the query-key multiplication and an identity matrix. Mathematically, GPS can be expressed by Equation 2, where  $E$  is an identity matrix,  $Attn$  is the self-attention [34] matrix. In one extreme, GPS degrades to the normal self-attention when  $\alpha = 0$ . In the other extreme, when  $\alpha = 1$ , GPS is equivalent to complete patch severance (*i.e.* no interaction happens in the self-attention layer). The hyper-parameter  $\alpha$  enables a tran-

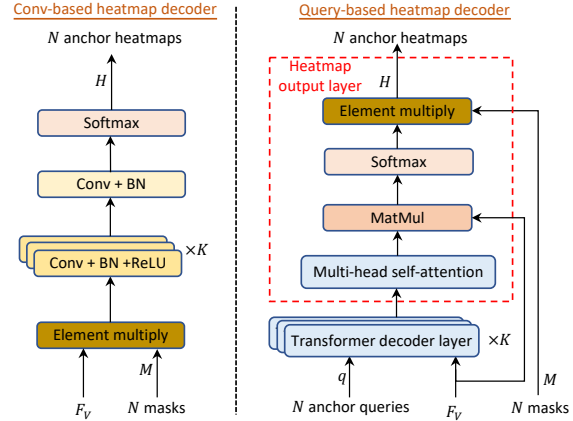


Figure 4: Comparison between the conv-based heatmap decoder and the novel query-based heatmap decoder for CAL.

sition between the two extreme cases.

$$F_{out} = (\alpha * E + (1 - \alpha) * Attn) \times F_{in}. \quad (2)$$

As shown in Figure 2, the patch severance can be applied multiple times to different transformer layers, while in practice we found that simply applying it to the last transformer layer is an effective choice.

#### 4.1.1 Comparison to Existing Adaptation Methods

**The embedding space is preserved.** Previous adaptation methods like ViT-adapter [7] usually require an updating of the backbone network, which could easily break the original embedding space. However, similar to VPT [20], the proposed patch severance can preserve the original visual embedding space of the pre-trained VLM (*e.g.* CLIP), such that the visual-language alignment and the zero-shot generality can be maintained. This property is in line with our first design principle.

**Training-free.** Unlike existing methods (such as ViT-adapter [7] and VPT [20]) which introduce additional parameters to the backbone network and thus requires further finetuning, patch severance brings no additional parameters and is training-free. Moreover, patch severance can work collaboratively with previous adapting methods (*e.g.* VPT) to achieve better adaptation for the task at hand.

#### 4.2. Query-Based Classification Anchor Learning

Classification Anchor Learning (CAL) is proposed to overcome the second limitation, namely, focusing on the more discriminative patches for classification (*i.e.* classification anchors). For example, to classify a human, the network may focus more on the head compared to the rest of the body. Essentially, CAL shares a similar spirit with the anchor selecting/learning for training object detectors [42, 41, 13], while it aims to output an anchor heatmap

Method	Decoupled	Passes	COCO-Stuff				VOC			
			hIoU	Seen	Unseen	FPS	hIoU	Seen	Unseen	FPS
ZS3Net [2]	✗	1	15.0	34.7	9.5	-	28.7	77.3	17.7	-
CaGNet [16]	✗	1	18.2	35.5	12.2	-	39.7	78.4	25.6	-
ZegFormer [10]	✓	$N'$	34.8	36.6	33.2	-	73.3	86.4	63.6	-
SimBaseline [38]	✓	$N'$	37.8	<b>39.3</b>	36.3	1.11	77.5	83.5	72.5	2.66
DeOP (ours)	✓	1	<b>38.2</b>	38.0	<b>38.4</b>	<b>4.37</b>	<b>80.8</b>	<b>88.2</b>	<b>74.6</b>	<b>6.55</b>

Table 1: Comparison with other methods on COCO-Stuff and Pascal VOC in the open-vocabulary setting. *passes* means the number of passing of the input image through the CLIP visual encoder. Note that  $N'$  is dependent on the input image since ZegFormer and SimBaseline only pass the image crops with scores above a threshold into CLIP, and  $N' \leq N$ .

for each proposal mask. To achieve this, we propose a query-based heatmap decoder module. As Figure 2 shows, the heatmap decoder (indicated by CAL) is appended to the CLIP visual encoder. It can generate attention heatmaps indicating which patches are more appropriate to represent the segments in terms of classification. The module is trained end-to-end with the whole segmentation network. We elaborate on it in the following.

#### 4.2.1 Anchor Heatmap Decoder

We first describe the proposed query-based heatmap decoder, and then introduce a simple counterpart based on convolution. Both of them can achieve CAL and they are compared experimentally in Section 5.3.

**Query-based heatmap decoder.** Overall, the decoder receives  $N$  learnable anchor queries (denoted by  $q$ ), feature maps output by CLIP visual encoder (denoted by  $F_V$ ) and  $N$  proposal masks (denoted by  $M$ ). The output is  $N$  classification anchor heatmaps  $H$ . The structure of the query-based heatmap decoder is displayed in Figure 4(right). The decoder is composed of  $K$  standard transformer decoder layers and a novel heatmap output layer. The  $K$  transformer layers are designed to update the anchor queries  $q$  by attending to the feature map of the input image.

The key component in the heatmap decoder is the **heatmap output layer**. Unlike the standard transformer decoder layer in which the output is the updated queries, the proposed heatmap decoder layer is designed to produce anchor heatmaps  $H$ . In specific, the anchor queries first go through a multi-head self-attention layer, and then are multiplied with the feature map followed by Softmax to obtain an attention map for each anchor query. Lastly, the attention maps are multiplied with the  $N$  proposal masks, enabling each anchor query to be aware of its corresponding mask/segment. Concretely, the heatmap output layer can be formulated as Equation 3:

$$H = \text{softmax}\left(\frac{q \times F_V}{\sqrt{d}}\right) * M, \quad (3)$$

where  $q \in \mathbb{R}^{B \times N \times C}$  is the  $N$  anchor queries,  $F_V \in \mathbb{R}^{B \times C \times H \times W}$  is the feature map generated by CLIP visual encoder,  $d$  is the dimension of  $F_V$ ,  $M \in \mathbb{R}^{B \times N \times H \times W}$  is  $N$  masks and  $*$  denotes the element-wise multiplication.

**Heatmap decoder based on convolution.** The proposed query-based heatmap decoder is not the only way to achieve CAL. In this work, we also compare the query-based heatmap decoder with a convolutional counterpart (shown in Figure 4 left). This conv-based anchor heatmap decoder first generates  $N$  weighted feature maps by applying each proposal mask to the feature map. Each weighted feature map goes into a stack of  $K$  conv-BN [18]-ReLU [1] layers, where conv layers have kernel size of 3 and channel size 512). Lastly, an extra conv layer maps the channel dimension to 1, followed by a Softmax layer to generate the  $N$  anchor heatmaps. Mathematically, the conv-based anchor heatmap decoder can be expressed as follows:

$$H = \sigma(\text{conv}(CBR_K(\dots CBR_1(F * M))), \quad (4)$$

where  $CBR$  refers to the conv-BN-ReLU layers, and  $\sigma$  denotes Softmax operation.

#### 4.2.2 Classification Based on Anchor Heatmap

Here we describe how to make use of the classification anchor heatmap for the segment classification. We simply follow the mask-based pooling described in [15], while the only difference is that the learned classification anchor heatmap is used for the pooling, instead of the mask proposal. We term this process heatmap-based pooling. It can be formulated as  $F_I = H \odot F_V$ , where  $F_I$  is the classification embeddings for the proposal masks/segments.

#### 4.2.3 Discussion on CAL

The proposed CAL can keep the original visual-alignment of CLIP because it only learns a spatial heatmap which down-weights or up-weights certain patch embeddings output by CLIP. Furthermore, although CAL is trained in seen

Method	Backbone	Training dataset	De.	Passes	VOC-20	PC-59	A-150	PC-459	A-847
ZS3Net [2]	R101	Pascal VOC 2012	✗	1	38.3	19.4	-	-	-
LSeg [23]	R101	Pascal VOC 2012	✗	1	47.4	-	-	-	-
OpenSeg [15]	R101	COCO	✗	1	60.0	36.9	15.3	6.5	4.0
OpenSeg [15]	R101	COCO + Loc. Narr.	✗	1	63.8	40.1	17.5	7.9	4.4
SimBase [38]	R101c	COCO-Stuff-156	✓	$N'$	88.4	47.7	20.5	-	7.0
DeOP (ours)	R101c	COCO-Stuff-156	✓	1	<b>91.7</b>	<b>48.8</b>	<b>22.9</b>	<b>9.4</b>	<b>7.1</b>

Table 2: Comparison with other methods on Pascal VOC, Pascal Context and ADE20K in the cross-dataset setting. The number after each dataset indicates the number of categories. *De.* means decoupled network. *Loc. Narr.* means Localized Narratives [31].

Method	COCO-Stuff				VOC			
	pAcc	hIoU	Seen	Unseen	pAcc	hIoU	Seen	Unseen
Baseline	15.6	7.0	6.2	8.0	46.6	30.7	37.5	26.0
Baseline+	27.6	8.3	8.0	8.7	64.3	33.4	53.2	24.4
Baseline+ w/ PS	53.9	25.8	25.9	25.4	72.7	47.6	64.0	37.9
Baseline+ w/ CAL	54.2	23.2	28.8	21.0	88.1	70.5	83.3	61.2
DeOP (ours)	<b>62.2</b>	<b>38.2</b>	<b>38.0</b>	<b>38.4</b>	<b>92.5</b>	<b>80.8</b>	<b>88.2</b>	<b>74.6</b>

Table 3: Ablation on the effectiveness of PS (*i.e.* patch severance) and CAL (*i.e.* classification anchor learning).

classes, it can be well-transferred to unseen classes. This impressive property is verified quantitatively in Section 5.3.

## 5. Experiments

**Datasets.** We study DeOP using four widely used image segmentation datasets: COCO-Stuff [3, 28], PASCAL VOC 2012 [12], PASCAL Conext [30] and ADE20K [45, 46]. In terms of data splitting and evaluation metrics, we simply follow [38]. The details regarding dataset statistics, data splitting and evaluation metrics are described in the supplementary material.

**Implementation details.** We conduct experiments based on Detectron2 [35]. We select ResNet-101c [4] which replaces the first  $7 \times 7$  convolution layer of ResNet-101 [17] with 3 consecutive  $3 \times 3$  convolutions as the backbone for MaskFormer, and set the number of queries  $N$  (*i.e.* proposal masks) in the transformer decoder to 100 by default. For CLIP, we use the text encoder and image encoder of the ViT-B/16 [11] model. The experiments are conducted on 8 Tesla V100 32G GPUs, and the batch size is set to 32. The training iterations are 20k for PASCAL VOC 2012 and 60k for COCO-Stuff. For the speed test, we test on COCO-Stuff dataset with batch size 1, using 1 Tesla V100 32G GPU.

### 5.1. Results in Open-Vocabulary Setting

We evaluate our method in a open-vocabulary segment setting by training the model on the COCO-Stuff dataset and the PASCAL VOC 2012 dataset. During training, we only use images belonging to seen classes in the training

dataset, while including all images in the validation dataset. Table 1 shows the comparison with other methods.

As shown in Table 1, DeOP significantly outperforms previous works. Impressively, we achieve 38.2 hIoU for the COCO-Stuff dataset and 80.8 hIoU for the PASCAL VOC 2012 dataset. The performance boost over ZegFormer and SimBaseline is mainly due to the proposed patch severance and learned anchor heatmaps for classification. Apart from that, we posit that those multi-pass methods get rid of the context (*i.e.* information around the segment) completely for each proposal segment (due to the mask operation in the pre-processing) which may hinder the performance; whereas DeOP utilizes context information naturally in the one-pass architecture.

**Inference speed.** To evaluate the efficiency of our method, we compare the inference speed between our method and SimBaseline in Table 1. Impressively, our method attains around 3 to 4 $\times$  increase in the inference speed (FPS). This demonstrates the benefits and efficiency of adopting the one-pass mechanism, which is our second design principle.

### 5.2. Results in Cross-Dataset Setting

We evaluate our method in a cross-dataset setting by training the model on the COCO-Stuff dataset and evaluating it on other datasets without fine-tuning. Table 2 shows that our approach generalizes well on other datasets due to the decoupled design (*i.e.* the first design principle) and frozen visual CLIP encoder, and outperforms all previous



Figure 5: The visualization of the heatmaps. From left to right are: the original image, the self-attention heatmap for a particular patch  $i$  in the last transformer layer in the CLIP visual encoder, the self-attention heatmap for another patch  $j$ , the proposal mask for the segment containing the patch  $i$  and the learned classification anchor heatmap for the patch  $i$ . The categories in the three rows are: human (seen), grass (**unseen**) and frisbee (**unseen**).

methods. Specifically, our method achieves 91.7 mIoU on the Pascal VOC dataset, 48.8 mIoU on the Pascal Context dataset, and 22.9 mIoU on the ADE20K dataset. It is worth noting that as the number of categories increases from left to right, the corresponding performance metrics decrease, demonstrating the importance of improving classification.

**Effectiveness of prompt learning, PS and CAL.** In Table 3, we validate the effectiveness of each of our proposed components over the plain *baseline*. Prompt learning is benefiting the pAcc metric, which increases from 15.6 to 27.6 on COCO-Stuff dataset and from 46.6 to 64.3 on VOC dataset, as shown on *baseline+*. Both PS and CAL bring improvements in all metrics. Not surprisingly, DeOP(containing the prompt learning, PS and CAL) achieves a significant improvement over all others, demonstrating that PS and CAL can work collaboratively and generalize well to unseen categories. In the supplementary material, we provide a detailed ablation on different combinations of text and visual prompt learning.

**Analysis on PS.** As discussed in Section 4.1, we compare two types of patch severance: MPS and GPS. The results are

Method	$\alpha_{-2}$	$\alpha_{-1}$	hIoU	Seen	<b>Unseen</b>
Baseline	0	0	7.0	6.2	8.0
MPS	-	-	12.5	11.0	14.4
GPS	0	0.5	20.0	17.1	24.0
GPS	0	1.0	<b>23.8</b>	<b>21.5</b>	<b>26.6</b>
GPS	0.5	1.0	16.2	15.2	17.4
GPS	0.5	0.5	13.3	12.1	14.8

Table 4: Ablation study on patch severance.  $\alpha_{-1}$  and  $\alpha_{-2}$  denote the hyper-parameter  $\alpha$  for the last layer and second last layer in ViT, respectively. Experiments are based on our baseline method (without prompt learning).

shown in Table 4. The first conclusion to draw is that both types of PS perform better than the original self-attention. In general, GPS achieves better performance than MPS. We think that it is because the prior (*i.e.* weighted proposal masks) is too strong and may not be optimal. Interestingly, the best combination is that the last layer has a complete patch severance. Notably, we also attempt to apply PS to more transformer layers in ViT. However, we find that PS



Type	$K$	pAcc	hIoU	Seen	Unseen
DeOP w/o CAL	0	53.9	25.7	25.9	25.4
Conv-based	3	58.8	31.4	31.4	31.4
Query-based	1	<b>62.2</b>	<b>38.2</b>	<b>38.0</b>	<b>38.4</b>
Query-based	3	62.2	37.7	37.9	37.6
Query-based	5	62.2	38.0	38.0	38.0

Table 5: Ablation study on the classification anchor learning.  $K$  is the number of layers in the heatmap decoder.

Method	pAcc	hIoU	Seen	Unseen	FPS
SimBaseline	62.1	38.6	<b>40.7</b>	37.0	0.27
DeOP (ours)	<b>62.6</b>	<b>40.0</b>	40.4	<b>39.6</b>	<b>1.85</b>

Table 6: Results with CLIP visual encoder based on ViT-L.

may break the embedding space if it is applied in shallower layers (*i.e.* the first layer) and results in poor performance which is shown in Table S3.

**Analysis on CAL.** We compare the query-based and conv-based heatmap decoder for CAL introduced in Section 4.2. For the query-based one, we also investigate the effect of the depth ( $K$ ) of the decoder. The results in Table 5 illustrate that query-based heatmap decoder is better than the conv-based counterpart, and the performance is not sensitive to the depth of the heatmap decoder.

**Results on larger CLIP visual encoder.** We also study the impact of the more powerful CLIP visual encoder on the outcomes. Specifically, we train SimBaseline and DeOP models with a CLIP visual encoder based on ViT-L/14-336 on the COCO-Stuff dataset. As displayed in Table 6, our approach consistently outperforms SimBaseline on unseen classes. Moreover, DeOP achieves a more impressive (almost  $7\times$ ) speed-up at inference.

### 5.3. Visualization

**Comparison in attention masks.** We visualize the attention maps, mask proposals, and anchor heatmaps in Figure 5. It is clear that the self-attention map for a patch is disorganized, and cannot provide useful information for the segmentation task. The attention map generated by CAL focuses on the local information surrounding the patch (*e.g.* attending to the tree to help for classifying the grass), and CAL also allows the model to pay more attention to the locations that are more discriminative for classification, *e.g.* the head being more important than the body of the human. More heatmap visualizations and segmentation results are displayed in the supplementary material.

## 6. Conclusion

In this work, we propose an efficient framework called DeOP for open-vocabulary semantic segmentation. Un-

like previous decoupled methods, DeOP only requires a single pass for the input image in the visual-language model. Specifically, DeOP is equipped with the proposed patch severance, which restricts redundant interference between patch embeddings in the visual encoder of the visual-language model, and classification anchor learning to identify visually discriminative regions for better segment classification. We conducted extensive experiments to validate the method and found that DeOP outperforms previous methods in both intra- and cross-dataset evaluation while being significantly faster than other multi-pass methods during inference.

**Acknowledgements** This work is supported by National Key R&D Program of China (No. 2022ZD0118700), Hong Kong Research Grant Council - Early Career Scheme (Grant No. 27208022), and HKU Seed Fund for Basic Research.

## References

- [1] Abien Fred Agarap. Deep learning using rectified linear units (relu). *ArXiv*, abs/1803.08375, 2018.
- [2] Max Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. *ArXiv*, abs/1906.00817, 2019.
- [3] Holger Caesar, Jasper R. R. Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1209–1218, 2016.
- [4] 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 transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin P. Murphy, and Alan Loddon Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:834–848, 2016.
- [6] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *ArXiv*, abs/1706.05587, 2017.
- [7] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Y. Qiao. Vision transformer adapter for dense predictions. *ArXiv*, abs/2205.08534, 2022.
- [8] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1280–1289, 2021.
- [9] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *Neural Information Processing Systems*, 2021.
- [10] Jian Ding, Nan Xue, Guisong Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. *2022 IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11573–11582, 2021.
- [11] 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. *ArXiv*, abs/2010.11929, 2020.
- [12] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.
- [13] Chengjian Feng, Yujie Zhong, Yu Gao, Matthew R Scott, and Weilin Huang. TOOD: Task-aligned one-stage object detection. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3490–3499. IEEE Computer Society, 2021.
- [14] 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.
- [15] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *European Conference on Computer Vision*, 2021.
- [16] Zhangxuan Gu, Siyuan Zhou, Li Niu, Zihan Zhao, and Liqing Zhang. Context-aware feature generation for zero-shot semantic segmentation. *Proceedings of the 28th ACM International Conference on Multimedia*, 2020.
- [17] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2015.
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ArXiv*, abs/1502.03167, 2015.
- [19] 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.
- [20] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge J. Belongie, Bharath Hariharan, and Ser Nam Lim. Visual prompt tuning. *ArXiv*, abs/2203.12119, 2022.
- [21] 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*, 2021.
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- [23] Boyi Li, Kilian Q. Weinberger, Serge J. Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *ArXiv*, abs/2201.03546, 2022.
- [24] Feng Li, Hao Zhang, Huaizhe Xu, Shilong Liu, Lei Zhang, Lionel M Ni, and Heung-Yeung Shum. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3041–3050, 2023.
- [25] Xiaoya Li, Xiaofei Sun, Yuxian Meng, Junjun Liang, Fei Wu, and Jiwei Li. Dice loss for data-imbalanced nlp tasks. In *Annual Meeting of the Association for Computational Linguistics*, 2019.
- [26] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Péter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. *ArXiv*, abs/2210.04150, 2022.
- [27] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:318–327, 2017.
- [28] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 2014.
- [29] Chao Ma, Yu-Hao Yang, Yanfeng Wang, Ya Zhang, and Weidi Xie. Open-vocabulary semantic segmentation with frozen vision-language models. In *British Machine Vision Conference*, 2022.
- [30] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [31] Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. Connecting vision and language with localized narratives. In *ECCV*, 2020.
- [32] 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.
- [33] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2014.
- [34] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *ArXiv*, abs/1706.03762, 2017.
- [35] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [36] Yongqin Xian, Subhabrata Choudhury, Yang He, Bernt Schiele, and Zeynep Akata. Semantic projection network for zero- and few-label semantic segmentation. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8248–8257, 2019.
- [37] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiulong Wang, and Shalini De Mello. Open-Vocabulary Panoptic Segmentation with Text-to-Image Diffusion Models. *arXiv preprint arXiv:2303.04803*, 2023.

- [38] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *European Conference on Computer Vision*, 2021.
- [39] Yuhui Yuan, Lang Huang, Jianyuan Guo, Chao Zhang, Xilin Chen, and Jingdong Wang. Ocnet: Object context for semantic segmentation. *International Journal of Computer Vision*, 129:2375 – 2398, 2021.
- [40] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023.
- [41] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9759–9768, 2020.
- [42] Xiaosong Zhang, Fang Wan, Chang Liu, Rongrong Ji, and Qixiang Ye. Freeanchor: Learning to match anchors for visual object detection. *Advances in neural information processing systems*, 32, 2019.
- [43] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6230–6239, 2016.
- [44] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything, 2023.
- [45] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017.
- [46] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127:302–321, 2019.
- [47] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130:2337 – 2348, 2021.