

Effective SAM Combination for Open-Vocabulary Semantic Segmentation

Minhyeok Lee¹ Suhwan Cho¹ Jungho Lee¹ Sunghun Yang¹
 Heeseung Choi² Ig-Jae Kim² Sangyoun Lee¹

¹Yonsei University

²Korea Institute of Science and Technology (KIST)

{hydragon516, chosuhwan, 2015142131, sunghun98, syleee}@yonsei.ac.kr

{hschoi, drjay}@kist.re.kr

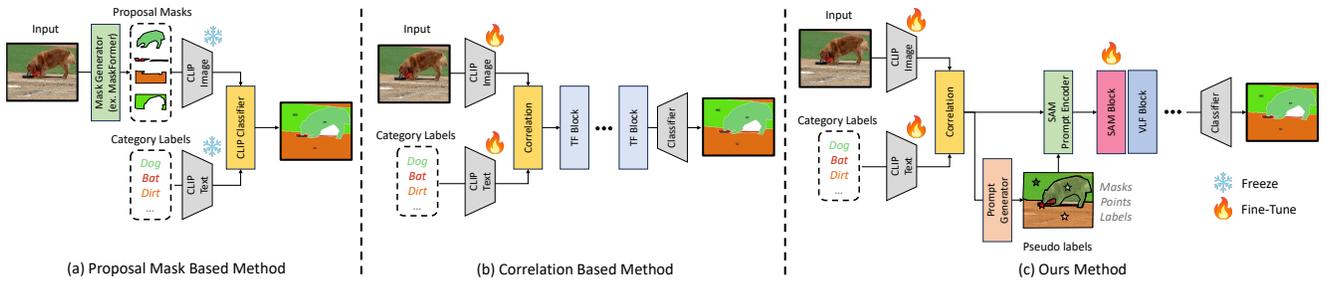


Figure 1. (a) A model structure that generates proposal masks using a mask generation model. (b) A model structure that refines the correlation between image and text. (c) The structure of the proposed ESC-Net. Our ESC-Net efficiently models the relationship between images and text by combining a pre-trained SAM block with pseudo prompts instead of an inefficient mask generation model. This approach enables much denser mask prediction compared to conventional correlation-based methods.

Abstract

Open-vocabulary semantic segmentation aims to assign pixel-level labels to images across an unlimited range of classes. Traditional methods address this by sequentially connecting a powerful mask proposal generator, such as the Segment Anything Model (SAM), with a pre-trained vision-language model like CLIP. But these two-stage approaches often suffer from high computational costs, memory inefficiencies. In this paper, we propose ESC-Net, a novel one-stage open-vocabulary segmentation model that leverages the SAM decoder blocks for class-agnostic segmentation within an efficient inference framework. By embedding pseudo prompts generated from image-text correlations into SAM’s promptable segmentation framework, ESC-Net achieves refined spatial aggregation for accurate mask predictions. Additionally, a Vision-Language Fusion (VLF) module enhances the final mask prediction through image and text guidance. ESC-Net achieves superior performance on standard benchmarks, including ADE20K, PASCAL-VOC, and PASCAL-Context, outperforming prior methods in both efficiency and accuracy. Comprehensive

ablation studies further demonstrate its robustness across challenging conditions.

1. Introduction

Open-vocabulary semantic segmentation aims to segment an image at the pixel level using an unlimited range of class labels. To achieve this, many approaches have gained attention by leveraging pre-trained vision-language models such as CLIP [28] and ALIGN [14], which associate images with various textual descriptions. However, since these models are trained at the image level using paired image-text datasets, they tend to perform poorly on pixel-level tasks.

To address this issue, many methods [7, 13, 21, 31, 35] adopt a two-stage approach that utilizes a mask proposal generator. As shown in Figure 1 (a), these methods first use a powerful class-agnostic mask proposal generator with strong pre-segmentation capabilities, such as MaskFormer [4] or Segment Anything Model (SAM) [19], to pre-segment the image. Next, each image region is fed into a pre-trained CLIP classifier to determine its class. However, these methods have two major issues. First, the two-stage

pipeline incurs significant computational costs and inefficient memory usage. Second, because these methods remove the background outside the segmented regions, a domain gap exists between the pre-trained CLIP model and the cropped regions. Additionally, as cropped regions typically include parts of other objects, classification accuracy often suffers.

To overcome these challenges, recent works [5, 16, 33] address open-vocabulary semantic segmentation by transferring the modeling of image-text correlation to pixel-level predictions. As shown in Figure 1 (b), these methods first establish multi-modal correlations from the embedded CLIP image and text features. Then, by refining this information through the model, they generate the final prediction masks in a manner similar to traditional semantic segmentation. Therefore, unlike the previous two-stage methods, these approaches can perform inference much more efficiently. Additionally, they generally enable direct fine-tuning of CLIP, further optimizing the cross-modality modeling between vision and language. However, CLIP’s vision-language contrastive learning aligns only the global representations of context, causing it to focus solely on the text-described information. As a result, these methods perform correlation at a low-resolution, highly embedded feature level to maximize the encoder’s localization capability. In other words, while previous two-stage methods can generate high-resolution boundary masks using a mask proposal generator, correlation-based models produce less accurate masks because of CLIP’s limited ability to reconstruct detailed local information.

In this paper, we propose a novel one-stage open-vocabulary semantic segmentation model called the Effective SAM Combination (ESC-Net) to address these challenges. Our ESC-Net leverages SAM’s powerful class-agnostic segmentation capability while maintaining an efficient inference pipeline. To achieve this, we focus on SAM’s ability to effectively transfer to downstream segmentation tasks through promptable segmentation. Figure 1 (c) illustrates the proposed ESC-Net’s process. First, the proposed ESC-Net computes the correlation between CLIP image features and text features, generating pseudo coordinate points and object masks that are highly correlated with the text class. This pre-generated information is then embedded as prompts via SAM’s prompt encoder, and used as pseudo features that incorporate approximate location information. Next, the generated prompt features and CLIP encoder features are used as inputs to the pre-trained SAM transformer blocks, enabling cross-model interaction between the prompts and the image. Therefore, due to SAM’s region-wise spatial aggregation capability, the image features are effectively represented to facilitate modeling the image-language relationships. In other words, this is made possible by SAM’s ability to generate valid segmen-

tation masks even when the prompts are ambiguous and can reference multiple objects. Finally, we design an Vision-Language Fusion (VLF) module to guide the model in generating the final class prediction masks based on image and text guidance.

The ESC-Net structure has two major advantages. First, the combination of correlation-based pseudo prompts and SAM blocks provides more concentrated spatial context information for image-text modeling. This allows for more accurate and dense prediction masks compared to previous one-stage methods that relied on CLIP’s weak spatial localization signals for images and text. Second, since ESC-Net does not use the SAM image encoder, it offers much more efficient inference compared to the existing two-stage methods that utilize the entire SAM pipeline. Notably, the proposed method achieves better performance compared to the previous correlation-based one-stage method, CAT-Seg [5], despite having similar computational costs.

The proposed ESC-Net is evaluated on standard open-vocabulary benchmarks, including the ADE20K [39], PASCAL-VOC [9], and PASCAL-Context [25] datasets, achieving state-of-the-art performance compared to previous methods. Additionally, we demonstrate through various ablation studies that our model maintains robust performance across a range of challenging scenarios.

Our main contributions can be summarized as follows:

- We propose ESC-Net, a novel one-stage open-vocabulary semantic segmentation model. It effectively combines CLIP and SAM to leverage SAM’s powerful class-agnostic segmentation capabilities while maintaining efficient inference.
- ESC-Net introduces correlation-based pseudo prompts derived from CLIP image-text features. These prompts are embedded through SAM’s prompt encoder and used to guide the SAM transformer blocks. This enhances image-text modeling and results in more accurate, dense prediction masks.
- Our method achieves state-of-the-art performance on standard open-vocabulary benchmarks such as ADE20K, PASCAL-VOC, and PASCAL-Context. It outperforms previous one-stage and two-stage methods with similar computational costs. Extensive ablation studies demonstrate its robust performance across various challenging scenarios.

2. Related Work

Open-vocabulary semantic segmentation. Open vocabulary semantic segmentation evolves as a specialized task within computer vision, aiming to enable models to segment arbitrary, unseen categories beyond predefined labels. Early methods [1, 32, 38] in this area focus on aligning visual features with pre-trained text embeddings by learning

a feature mapping that associates visual and text spaces effectively.

With the advent of vision-language models like CLIP, many works leverage its extensive vocabulary and visual-text alignment capabilities to address the vocabulary limitations of previous approaches. For instance, two-stage frameworks like OpenSeg [10], ZegFormer [6], and ZSseg [35] use class-agnostic mask proposals that are later matched with CLIP-based text embeddings for classification. OVSeg [21] improves this two-stage approach by fine-tuning CLIP with region-text pairs to enhance segmentation accuracy on diverse categories, while ODISE [34] employs a diffusion model to generate high-quality masks, thus bypassing the reliance on pre-trained mask generators. However, these methods often depend on region generators, which are typically trained on limited annotated data, thus constraining generalization to unseen categories.

Single-stage frameworks emerge as an alternative to bypass the need for region proposals. Models such as LSeg [20] and MaskCLIP [7] adopt pixel-level learning directly from CLIP’s text embeddings, aiming to predict segmentation maps in one forward pass. Methods like SAN [36] and ZegCLIP [40] introduce side adapter networks and learnable tokens to assist CLIP’s segmentation performance. CAT-Seg [5] adapts CLIP by aggregating cosine similarity between image and text embeddings, fine-tuning the encoders to enhance segmentation on unseen classes. Meanwhile, MAFT+ [16] proposes a collaborative optimization approach, where content-dependent transfer adaptively refines text embeddings based on image context, while a Representation Compensation strategy preserves CLIP’s zero-shot capability.

Segment Anything Model. The Segment Anything Model (SAM)[19], recently introduced by Meta AI Research, marks a pivotal advancement in the field of foundation models for computer vision, particularly image segmentation. It builds on the Vision Transformer[8] architecture and trains on the vast SA-1B dataset, which contains over a billion segmentation masks across 11 million images. SAM positions itself as a transformative tool in zero-shot image segmentation. Its primary innovation lies in its “promptable” design, which enables it to generalize across various image domains without fine-tuning, allowing effective segmentation in new environments through prompt engineering. This promptability aligns SAM with trends in natural language processing, where large language models, also pre-trained on web-scale data, demonstrate robust zero-shot and few-shot capabilities.

SAM’s zero-shot proficiency draws comparisons to the advancements in natural language processing seen in models like ChatGPT, marking it as a milestone toward artificial general intelligence in visual tasks. Prior to SAM, foundation models in computer vision, such as CLIP and

ALIGN, focus on aligning multimodal data—specifically text and image encoders—through contrastive learning for cross-modal zero-shot generalization. However, these earlier models primarily excel at tasks reliant on cross-modal alignment, such as image captioning and generation, rather than segmentation. SAM advances beyond these models by specifically targeting image segmentation as a foundational task, making it a versatile tool for addressing complex segmentation challenges.

3. Proposed Approach

3.1. Overall Architecture

Figure 2 shows the overall structure of the proposed ESC-Net. Our model consists of a CLIP vision encoder, language encoder, and a single decoder for mask prediction, similar to conventional one-stage methods, without the need for an additional mask proposal model. In this task, we provide the input image I and a set of text descriptions T for all candidate classes to the model, where T^n denotes the text description of the n -th class. The total number of candidate classes is N_c , where $n = 1, \dots, N_c$. Therefore, ESC-Net has the same input configuration as conventional one-stage methods. First, ESC-Net generates a vision-language correlation map $C_{v\&l}$ from the CLIP vision feature $F_v \in \mathbb{R}^{C \times H \times W}$ and language feature $F_l \in \mathbb{R}^{C \times N_c}$. $C_{v\&l} \in \mathbb{R}^{N_c \times H \times W}$ represents the image similarity relationships for each text class. Next, we design a Pseudo Prompt Generator (PPG) to generate pseudo point coordinates and mask regions for each class candidate based on $C_{v\&l}$. These pseudo prompts embed approximate location information about the object regions and are passed to the pre-trained SAM block. A detailed explanation of the PPG is provided in Section 3.3. Next, F_v passes through the SAM transformer block along with the pseudo prompts, refining it into $F'_v \in \mathbb{R}^{C \times H \times W}$ with enhanced class-agnostic spatial context information. Finally, we propose the Vision-Language Fusion (VLF) module to model the interaction between the refined vision feature and the text feature. In this module, $C_{v\&l}$ is aggregated with F'_v and F_l to produce a reconstructed correlation map $C'_{v\&l} \in \mathbb{R}^{N_c \times C \times H \times W}$. Details on the VLF module are discussed in Section 3.5. As shown in Figure 2, we progressively refine $C_{v\&l}$ using N ESC blocks from the above process and generate the final prediction mask M through the mask decoder.

3.2. Vision-Language Correlation

Before generating the pseudo prompt, we first create a correlation map using the correlation between F_v and F_l . This process is performed based on simple cosine similarity. Let $C_{v\&l}^n(i)$ be the image correlation map for the n -th class information, and let i represent the 2D spatial coordinates of the image embedding. Then, $C_{v\&l}^n(i)$ is expressed as fol-

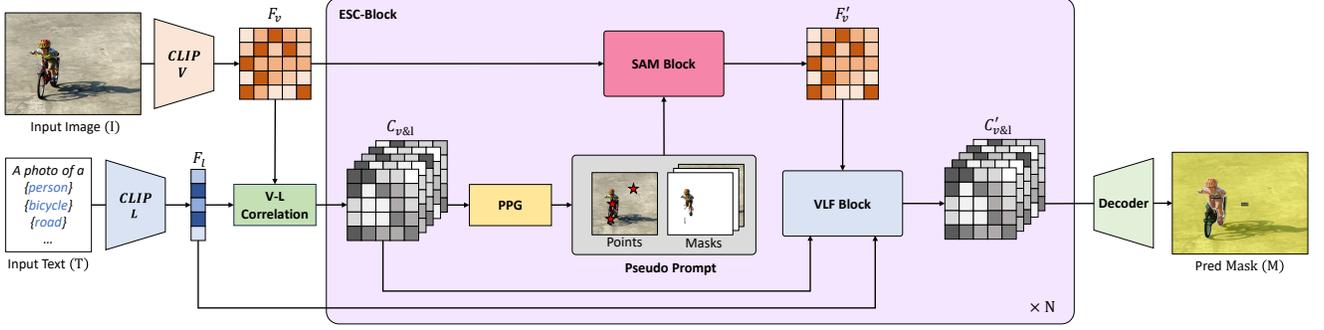


Figure 2. The proposed ESC-Net consists of the CLIP vision and language encoders, N consecutive ESCBlocks, and a decoder. Each ESC-Block generates a pseudo prompt from the image-text correlation map and uses it as input to the SAM block. The SAM block aggregates the CLIP image features. The VLF block models the image-text correlation using image features and text features, refining the correlation map through this process.

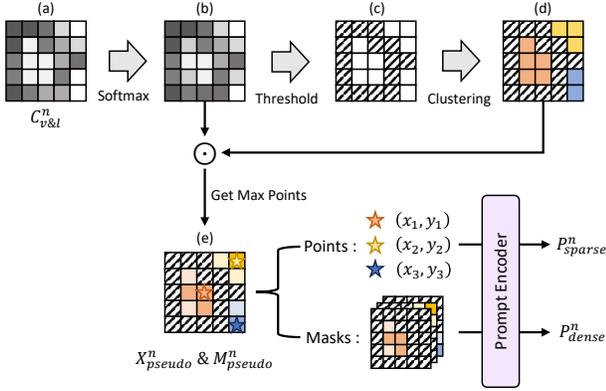


Figure 3. The process of the proposed Pseudo Prompt Generator (PPG). PPG aims to generate class-specific pseudo prompts from image-text correlation maps. For efficiency, all processes are computed in batch-wise parallelization over all classes.

lows:

$$C_{v\&l}^n(i) = \frac{F_v(i) \cdot F_l^n}{\|F_v(i)\| \|F_l^n\|}. \quad (1)$$

3.3. Pseudo Prompt Generator

The purpose PPG aim to generate location information for candidate objects from $C_{v\&l}^n$ and convert it into the SAM prompt format. Traditional SAM [19] can take some or all of points, bounding boxes, masks, and text as input prompts. Among these, we use only points and masks as prompts. Figure 3 shows the proposed PPG process.

First, PPG applies a softmax operation to the correlation map $C_{v\&l}^n \in \mathbb{R}^{1 \times H \times W}$ for each n -th class to generate class-specific object probability masks (b). Then, it binarizes this probability mask using a predefined threshold value α . As a result, this binary mask (c) represents the approximate location of candidate objects corresponding to class n . However, since multiple objects may exist for a single class, we cluster the mask by region, as shown in Figure 3 (d). To

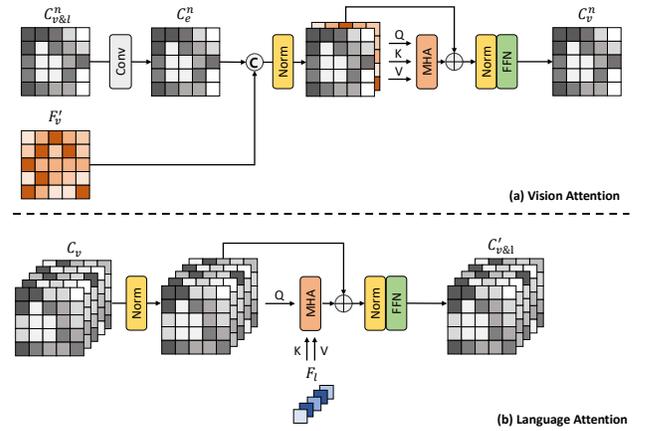


Figure 4. The structure of the proposed Vision-Language Fusion (VLF) block. VLF sequentially applies image and text guidance to the correlation map to refine it.

achieve this, we apply a k-means clustering algorithm [12] based on pixel locations to divide the binary mask into N_o mask regions. Next, we multiply the probabilistic mask (b) by the clustered mask region map (d) to generate a probability map (e) where only the object regions are filtered. Finally, we generate pseudo points and pseudo masks from (e). As shown in Figure 3, the pseudo points are determined as the highest probability areas for each object region in (e), and the pseudo masks correspond to the masks of each region. Consequently, a total of N_o pseudo points and N_o pseudo masks are generated.

The generated pseudo points and pseudo masks are used as inputs to the prompt encoder, as proposed by the traditional SAM. According to SAM’s approach, the points and masks are embedded as sparse prompt features P_{sparse}^n and dense prompt features P_{dense}^n , respectively.

3.4. SAM Block

The proposed ESC-Net utilizes a pre-trained SAM transformer block, as shown in Figure 2, to aggregate the spa-

tial context information of image features using generated pseudo-prompt information. Specifically, we leverage the transformer decoder block of the SAM mask decoder, which effectively maps images and prompts. The SAM block first applies a prompt self-attention layer $SA(\cdot)$, followed by bidirectional cross-attention $BCA(\cdot, \cdot)$ for image-to-prompt and prompt-to-image, updating the image embedding feature F_v to $F'_v \in \mathbb{R}^{C \times H \times W}$. However, unlike the original SAM, we use N_c sets of prompts, so the operations for each prompt are parallelized in batches. Therefore, the refined image feature $(F'_v)^n \in \mathbb{R}^{C \times H \times W}$ for the n -th prompt set is expressed as follows:

$$\begin{aligned} (F'_v)^n &= BCA(SA(F_l^n), F_v), \\ F'_v &= \text{Conv}([F_v^0; F_v^1; \dots; F_v^{N_c}]), \end{aligned} \quad (2)$$

where $\text{Conv}(\cdot)$ denotes a 1×1 convolution, and $[\cdot; \cdot]$ represents the concatenation operation. In this paper, we apply a total of N ESC blocks, thus utilizing the parameters of N consecutive pre-trained SAM blocks.

3.5. Vision-Language Fusion Module

The VLF module aims to refine the correlation map $C_{v\&l}$ by using the image embedding feature F'_v , which has enhanced spatial aggregation through SAM blocks, and the text feature F_l . First, as shown in Figure 4 (a), the correlation map $C_{v\&l}^n \in \mathbb{R}^{1 \times H \times W}$ for the n -th class is embedded using a 1×1 convolution layer. The embedded class correlation map $C_e^n \in \mathbb{R}^{C \times H \times W}$ is then concatenated with F'_v and used as input to the transformer block. To improve the model’s efficiency, this process is performed with batch-level parallel operations for all classes, similar to the SAM block. Furthermore, similar to CAT-Seg [5], we structure the transformer block using the Swin Transformer [23] architecture to gain additional memory advantages. As a result of the above process, the vision correlation map $C_v^n \in \mathbb{R}^{C \times H \times W}$ for the n -th class is generated, as shown in Figure 4 (a). Next, as shown in Figure 4 (b), we introduce an attention mechanism with F_l as the key and value to explicitly capture relationships between different class categories. For open-vocabulary segmentation, the number of input classes N_c can vary, and it needs to be order-agnostic. Therefore, we adopt a linear transformer [18] structure without positional embedding. As a result, the VLF module generates a multimodally refined embedding correlation map $C'_{v\&l}$.

3.6. Mask Prediction Decoder

We use the simple upsampling layers of a U-Net [29] architecture, composed of bilinear interpolation and 3×3 convolution layers, to generate the final predicted segmentation map. This approach is similar to existing one-stage methods and employs CLIP image encoder features as skip connections.

4. Experiments

4.1. Datasets

In this paper, we train the proposed ESC-Net on the COCO-Stuff [2] dataset, following established methods. The COCO-Stuff training set includes approximately 118k images annotated with masks across 171 semantic categories. To ensure a fair comparison with previous methods, we evaluate our model on the ADE20K [39], PASCAL-VOC [9], and PASCAL-Context [25] datasets. Details for each dataset are as follows:

ADE20K [39] is a large-scale semantic segmentation dataset featuring 20,000 training images and 2,000 validation images. In open-vocabulary semantic segmentation, it includes two distinct test sets: A-150, containing 150 commonly seen categories, and A-847, which includes a total of 847 categories.

PASCAL-VOC [9], one of the foundational datasets for object detection and segmentation, has approximately 1,500 training images and 1,500 validation images, covering 20 object categories. For the open-vocabulary segmentation task, this dataset is referred to as PAS-20. Additionally, following the approaches in [5, 10], we include PAS-20^b in the evaluation dataset. Unlike PC-59, PAS-20^b is defined without including the “background” class.

PASCAL-Context [25], an extension of PASCAL VOC for semantic segmentation, also features two test sets for open-vocabulary segmentation: PC-59 with 59 categories and PC-459 with 459 categories.

4.2. Evaluation Metric

To quantitatively evaluate the performance, we follow existing traditional open-vocabulary semantic segmentation. Semantic segmentation results are evaluated with mean Intersection over Union (mIoU) [9].

4.3. Implementation Details

In this study, we set the class-specific object count N_0 in PPG to 5 and the number of ESC-Blocks N to 4. In other words, PPG generates 5 points and masks per class, and the correlation map is refined through 4 ESC-Blocks after the CLIP encoder. Therefore, we take 4 SAM mask decoder transformer blocks from the pretrained SAM model. The input image is resized to 336×336 , and the AdamW [24] optimizer is used for training. For optimal training, the CLIP encoder and SAM blocks are trained with a learning rate of 2×10^{-6} , while the rest of the model part is trained with a learning rate of 2×10^{-4} . All experiments are conducted using the PyTorch [26] framework, and training is performed on four NVIDIA RTX A6000 GPUs.

Model	Publication	VLM	Additional Backbone	Training Dataset	Additional Dataset	A-847	PC-459	A-150	PC-59	PAS-20	PAS-20 ^b
SPNet [32]	CVPR'19	-	ResNet-101	PASCAL VOC	✗	-	-	-	24.3	18.3	-
ZS3Net [1]	NeurIPS'19	-	ResNet-101	PASCAL VOC	✗	-	-	-	19.4	38.3	-
LSeg [20]	ICLR'22	CLIP ViT-B/32	ResNet-101	PASCAL VOC-15	✗	-	-	-	-	47.4	-
LSeg+ [10]	ECCV'22	ALIGN	ResNet-101	COCO-Stuff	✓	2.5	5.2	13.0	36.0	-	59.0
ZegFormer [6]	CVPR'22	CLIP ViT-B/16	ResNet-101	COCO-Stuff-156	✗	4.9	9.1	16.9	42.8	86.2	62.7
ZSseg [35]	ECCV'22	CLIP ViT-B/16	ResNet-101	COCO-Stuff	✓	7.0	-	20.5	47.7	88.4	-
OpenSeg [10]	ECCV'22	ALIGN	ResNet-101	COCO Panoptic	✓	4.4	7.9	17.5	40.1	-	63.8
OVSeg [21]	CVPR'23	CLIP ViT-B/16	ResNet-101c	COCO-Stuff	✓	7.1	11.0	24.8	53.3	92.6	-
ZegCLIP [40]	CVPR'23	CLIP ViT-B/16	-	COCO-Stuff-156	✗	-	-	-	41.2	93.6	-
SAN [36]	CVPR'23	CLIP ViT-B/16	-	COCO-Stuff	✗	10.1	12.6	27.5	53.8	94.0	-
DeOP [11]	ICCV'23	CLIP ViT-B/16	ResNet-101c	COCO-Stuff-156	✗	7.1	9.4	22.9	48.8	91.7	-
SCAN [22]	CVPR'24	CLIP ViT-B/16	Swin-B	COCO-Stuff	✗	10.8	13.2	30.8	<u>58.4</u>	<u>97.0</u>	-
EBSeg [30]	CVPR'24	CLIP ViT-B/16	SAM ViT-B	COCO-Stuff	✗	11.1	17.3	30.0	56.7	94.6	-
SED [33]	CVPR'24	ConvNeXt-B	-	COCO-Stuff	✗	11.4	18.6	31.6	57.3	94.4	-
CAT-Seg [5]	CVPR'24	CLIP ViT-B/16	-	COCO-Stuff	✗	<u>12.0</u>	<u>19.0</u>	<u>31.8</u>	57.5	94.6	<u>77.3</u>
ESC-Net (ours)	-	CLIP ViT-B/16	-	COCO-Stuff	✗	13.3	21.1	35.6	59.0	97.3	80.1
						(+1.3)	(+2.1)	(+3.8)	(+0.6)	(+0.3)	(+2.8)
LSeg [20]	ICLR'22	CLIP ViT-B/32	ViT-L/16	PASCAL VOC-15	✗	-	-	-	-	52.3	-
OpenSeg [10]	ECCV'22	ALIGN	EfficientNet-B7	COCO Panoptic	✓	8.1	11.5	26.4	44.8	-	70.2
OVSeg [21]	CVPR'23	CLIP ViT-L/14	Swin-B	COCO-Stuff	✓	9.0	12.4	29.6	55.7	94.5	-
SAN [36]	CVPR'23	CLIP ViT-L/14	-	COCO-Stuff	✗	12.4	15.7	32.1	57.7	94.6	-
ODISE [34]	CVPR'23	CLIP ViT-L/14	Stable Diffusion	COCO-Stuff	✗	11.1	14.5	29.9	57.3	-	-
FC-CLIP [37]	NeurIPS'23	ConvNeXt-L	-	COCO Panoptic	✗	11.2	12.7	26.6	42.4	89.5	-
MAFT [15]	NeurIPS'23	CLIP ViT-L/14	-	COCO-Stuff	✗	12.7	16.2	33.0	59.0	92.1	-
USE [31]	CVPR'24	CLIP ViT-L/14	DINOv2, SAM	COCO-Stuff	✓	13.4	15.0	37.1	58.0	-	-
SCAN [22]	CVPR'24	CLIP ViT-L/14	Swin-B	COCO-Stuff	✗	14.0	16.7	33.5	59.3	97.0	-
EBSeg [30]	CVPR'24	CLIP ViT-L/14	SAM ViT-B	COCO-Stuff	✗	13.7	21.0	32.8	60.2	<u>97.2</u>	-
SED [33]	CVPR'24	ConvNeXt-L	-	COCO-Stuff	✗	13.9	22.6	35.2	60.6	96.1	-
CAT-Seg [5]	CVPR'24	CLIP ViT-L/14	-	COCO-Stuff	✗	<u>16.0</u>	<u>23.8</u>	<u>37.9</u>	<u>63.3</u>	97.0	<u>82.5</u>
MAFT+ [17]	ECCV'24	ConvNeXt-L	-	COCO-Stuff	✗	15.1	21.6	36.1	59.4	96.5	-
ESC-Net (ours)	-	CLIP ViT-L/14	-	COCO-Stuff	✗	18.1	27.0	41.8	65.6	98.3	86.3
						(+2.1)	(+3.2)	(+3.9)	(+2.3)	(+1.1)	(+3.8)

Table 1. Quantitative evaluation on open-vocabulary segmentation benchmarks. The best-performing results are presented in bold, while the second-best results are underlined. Improvements over the second-best are highlighted in red.

4.4. Results

Quantitative results. Table 1 presents the quantitative results of the proposed ESC-Net compared to previous state-of-the-art methods on standard open-vocabulary semantic segmentation benchmark datasets [9, 25, 39]. In open-vocabulary semantic segmentation tasks, the final performance generally varies depending on the type of vision-language model (VLM) used. Therefore, as shown in Table 1, we perform comparisons between models with similar-scale VLMs. Some approaches utilize auxiliary models such as SAM or mask formers for mask generation or additional information extraction, in addition to the VLM. Furthermore, some models enhance performance by leveraging additional training datasets [3, 27]. Hence, in Table 1, we provide information on the additional backbones and types of training datasets used for each method. As shown in Table 1, the proposed ESC-Net achieves state-of-the-art performance, surpassing existing methods with both CLIP ViT-B/16 and CLIP ViT-L/14 VLMs. Compared to previous state-of-the-art methods, MAFT+ [17] and CAT-Seg [5], there is a significant performance improvement across all evaluation datasets. Notably, ESC-Net shows an mIoU improvement of 2.1 on the A-847 dataset, which includes a large number of classes, and 3.2 on the PC-459

dataset. Furthermore, in Section 4.5, we demonstrate the effectiveness of the proposed ESC-Net over existing models through various ablation studies.

Qualitative results. Figure 5 presents the qualitative results comparing the proposed ESC-Net with the previous state-of-the-art single-stage model, CAT-Seg [5], across various benchmarks. As shown, our model demonstrates robustness to variations in object scale, lighting, and more, achieving denser segmentation compared to CAT-Seg. Furthermore, in scenarios with a high number of objects and complex backgrounds, the proposed model exhibits superior performance. This is due to the effective use of features from the SAM block, which learns class-agnostic segmentation across various scales, for mask generation.

4.5. Ablation Analysis.

This section includes various ablation experiments on the proposed model. All experiments are evaluated using the CLIP ViT-L/14 VLM with an image size of 336×336 .

Effect of SAM blocks. Table 2 demonstrates the effectiveness of the pre-trained SAM blocks in the proposed model. First, (a) represents the model without using the SAM blocks, meaning that F_v is not refined. Consequently, the proposed PPG is also removed in this exper-



Figure 5. Qualitative comparison of CAT-Seg and our ESC-Net across various datasets. Our model is capable of generating more accurate and robust masks compared to existing correlation-based state-of-the-art method.

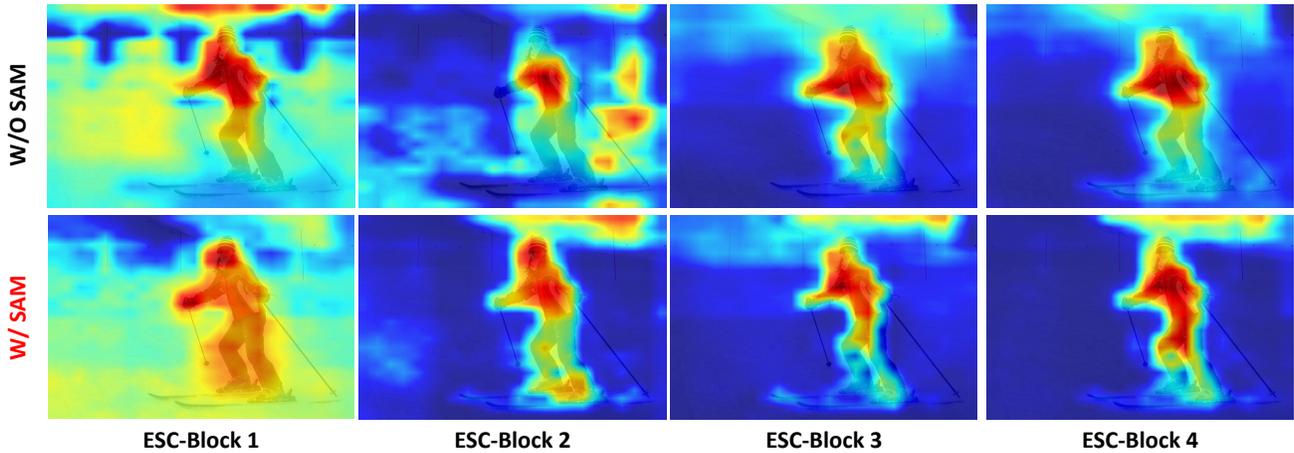


Figure 6. Visualization of image-text correlation maps with and without the SAM block. We visualize the model activation maps for the “Person” class for each ESC-Block. The proposed SAM-based method enables more accurate and dense object localization compared to the baseline.

iment. Next, (b) shows the model that utilizes the SAM structure but does not use pre-trained parameters. Lastly, (c) corresponds to the model that uses the pretrained SAM blocks, which is identical to the proposed ESC-Net. As shown in the table, instead of generating masks using a mask proposal network, significant performance improvements can be achieved by integrating pretrained SAM trans-

former blocks. In particular, when comparing (b) and (c), which have the same computational load, the approach using pretrained SAM blocks reaches much better segmentation performance. This demonstrates that the proposed ESC-Net effectively utilizes the spatial aggregation capabilities learned by SAM through the pseudo prompts.

Furthermore, to verify that the proposed SAM integra-

Index	Methods	A-847	PC-459	A-150	PC-59	PAS-20	PAS-20
(a)	W/O SAM	4.8	11.7	24.2	50.4	89.4	74.8
(b)	W/ SAM (<i>rand init</i>)	5.9	15.8	28.4	55.9	91.5	77.2
(c)	W/ SAM (<i>pretrained</i>)	18.1	27.0	41.8	65.6	98.3	86.3

Table 2. Quantitative performance comparison based on the use of SAM blocks. “*rand init*” refers to SAM blocks with random initialization.

tion structure aids in image-text modeling, we visualize $C'_{v&l}$ with and without the SAM block as class-specific activation maps, as shown in Figure 6. Since we include one SAM block in each ESC block, we generate a total of four activation maps. As illustrated in the figure, the proposed SAM integration method shows more accurate and denser mapping to the class target objects compared to the baseline. This demonstrates that the ESC-Net structure can achieve more precise image-text modeling compared to existing methods.

Effect of the Pseudo Prompt Generation Method. Table 3 presents the performance of the proposed model based on different combinations of pseudo prompt generation methods. Since the traditional SAM model takes points and bounding boxes as sparse prompt embeddings and masks as dense prompt embeddings, we explore various combinations of prompt embeddings. In this experiment, we utilize the clustered masks shown in Figure 3 (d) to generate bounding box pseudo prompts. For each mask, we find the smallest bounding box that covers it and use this as the pseudo bounding box. Similar to points and masks, 5 bounding boxes per class are sampled. As shown in Table 3 (d), the best performance was achieved when using both points and masks together. In contrast, bounding boxes, as in (a) and (b), were found to be less effective as pseudo prompt embeddings than points. This is likely because the accuracy of the generated pseudo bounding boxes is relatively low, and multiple objects can be overlapped within the box. Although the bounding boxes are generated from pseudo masks, in traditional SAM, mask prompts are trained as areas that approximately indicate the object’s location. Therefore, they function effectively as pseudo prompts that are less dependent on the precise position of the object. Consequently, using all three pseudo prompts, as in Table 3 (e), is inefficient; thus, we use only points and masks as pseudo prompts.

Model efficiency. In Table 4, we compare the model efficiency of the proposed ESC-Net with previous methods [5, 6, 21, 35]. To this end, we calculate the model’s parameters, inference time, and inference GFLOPs. All experiments were conducted on the same single RTX A6000 GPU, and except for the inference time, the rest of the data follows the results presented by [5]. As shown in Table 4, the proposed model exhibits lightweight and efficient inference performance compared to existing methods. Fur-

Index	Methods			A-847	PC-459	A-150	PC-59	PAS-20	PAS-20
	Point	B-box	Mask						
(a)	✓			15.2	23.6	37.8	59.4	96.7	81.3
(b)		✓		14.3	20.5	33.5	59.4	96.5	81.4
(c)	✓	✓		15.1	23.4	38.1	62.0	96.7	82.5
(d)	✓		✓	18.1	27.0	41.8	65.6	98.3	86.3
(e)		✓	✓	17.0	26.1	39.9	62.0	97.1	84.2
(f)	✓	✓	✓	<u>17.6</u>	<u>26.5</u>	41.9	65.8	<u>98.0</u>	<u>86.1</u>

Table 3. Quantitative performance comparison based on the combination of pseudo prompts used. “B-box” refers to a bounding box.

Model	# of Params. (M)	Inference time (s)	Inference GFLOPs
ZegFormer [6]	531.2	3.10	19,425.6
ZSSeg [35]	530.8	3.11	22,302.1
OVSeg [21]	532.6	2.98	19,345.6
CAT-Seg [5]	433.7	0.78	2,121.1
Ours	451.3	0.76	2,203.5

Table 4. Model efficiency comparison. All results are measured with a single RTX A6000 GPU.

thermore, it has comparable parameters and GFLOPs when compared to CAT-Seg, which is also a correlation-based one-stage model.

Limitations. While ESC-Net leverages SAM blocks for spatial aggregation, it operates on low-resolution feature levels, which limits the mask resolution and affects the precision of segmentation boundaries. This issue results in lower-quality masks compared to methods that can perform high-resolution aggregation. Additionally, similar to other correlation-based approaches, ESC-Net’s memory requirements increase with the number of classes due to the expanded text-image feature correlations. This growth in memory usage can reduce the efficiency of the model, particularly when handling a large variety of classes.

5. Conclusion

In this paper, we present ESC-Net, a one-stage open-vocabulary semantic segmentation model that efficiently segments images across diverse class labels. By combining SAM’s powerful class-agnostic segmentation with CLIP’s vision-language alignment, ESC-Net overcomes the limitations of prior two-stage and one-stage methods. Our model uses correlation-based pseudo prompts from CLIP’s image-text features to guide SAM, generating precise prediction masks without requiring a separate mask proposal stage, thus reducing computational costs. We evaluate ESC-Net on standard open-vocabulary benchmarks, including ADE20K, PASCAL-VOC, and PASCAL-Context, where it achieves state-of-the-art performance. Additionally, ablation studies confirm its robustness across various challenging scenarios, demonstrating ESC-Net as an efficient and accurate solution for open-vocabulary semantic segmentation.

References

- [1] Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 32, 2019. 2, 6
- [2] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Cocosuff: Thing and stuff classes in context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1209–1218, 2018. 5
- [3] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 6
- [4] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in neural information processing systems*, 34:17864–17875, 2021. 1
- [5] Seokju Cho, Heeseong Shin, Sunghwan Hong, Anurag Arnab, Paul Hongsuck Seo, and Seungryoung Kim. Catseg: Cost aggregation for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4113–4123, 2024. 2, 3, 5, 6, 8
- [6] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11583–11592, 2022. 3, 6, 8
- [7] Zheng Ding, Jieke Wang, and Zhuowen Tu. Open-vocabulary universal image segmentation with maskclip. *arXiv preprint arXiv:2208.08984*, 2022. 1, 3
- [8] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [9] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010. 2, 5, 6
- [10] 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*, pages 540–557. Springer, 2022. 3, 5, 6
- [11] Cong Han, Yujie Zhong, Dengjie Li, Kai Han, and Lin Ma. Open-vocabulary semantic segmentation with decoupled one-pass network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1086–1096, 2023. 6
- [12] John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108, 1979. 4
- [13] Dat Huynh, Jason Kuen, Zhe Lin, Jiuxiang Gu, and Ehsan Elhamifar. Open-vocabulary instance segmentation via robust cross-modal pseudo-labeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7020–7031, 2022. 1
- [14] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 1
- [15] Siyu Jiao, Yunchao Wei, Yaowei Wang, Yao Zhao, and Humphrey Shi. Learning mask-aware clip representations for zero-shot segmentation. *Advances in Neural Information Processing Systems*, 36:35631–35653, 2023. 6
- [16] Siyu Jiao, Hongguang Zhu, Jiannan Huang, Yao Zhao, Yunchao Wei, and Humphrey Shi. Collaborative vision-text representation optimizing for open-vocabulary segmentation. *arXiv preprint arXiv:2408.00744*, 2024. 2, 3
- [17] Siyu Jiao, Hongguang Zhu, Jiannan Huang, Yao Zhao, Yunchao Wei, and Humphrey Shi. Collaborative vision-text representation optimizing for open-vocabulary segmentation. In *European Conference on Computer Vision*, pages 399–416. Springer, 2025. 6
- [18] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR, 2020. 5
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 1, 3, 4
- [20] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022. 3, 6
- [21] Feng Liang, Bichen Wu, Xiaoliang Dai, Kungpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7061–7070, 2023. 1, 3, 6, 8
- [22] Yong Liu, Sule Bai, Guanbin Li, Yitong Wang, and Yansong Tang. Open-vocabulary segmentation with semantic-assisted calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3491–3500, 2024. 6
- [23] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 5
- [24] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [25] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 891–898, 2014. 2, 5, 6
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming

- Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5
- [27] Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. Connecting vision and language with localized narratives. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 647–664. Springer, 2020. 6
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1
- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015. 5
- [30] Xiangheng Shan, Dongyue Wu, Guilin Zhu, Yuanjie Shao, Nong Sang, and Changxin Gao. Open-vocabulary semantic segmentation with image embedding balancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28412–28421, 2024. 6
- [31] Xiaoqi Wang, Wenbin He, Xiwei Xuan, Clint Sebastian, Jorge Piazzentin Ono, Xin Li, Sima Behpour, Thang Doan, Liang Gou, Han-Wei Shen, et al. Use: Universal segment embeddings for open-vocabulary image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4187–4196, 2024. 1, 6
- [32] Yongqin Xian, Subhabrata Choudhury, Yang He, Bernt Schiele, and Zeynep Akata. Semantic projection network for zero-and few-label semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8256–8265, 2019. 2, 6
- [33] Bin Xie, Jiale Cao, Jin Xie, Fahad Shahbaz Khan, and Yanwei Pang. Sed: A simple encoder-decoder for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3426–3436, 2024. 2, 6
- [34] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. 3, 6
- [35] 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*, pages 736–753. Springer, 2022. 1, 3, 6, 8
- [36] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2945–2954, 2023. 3, 6
- [37] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. *Advances in Neural Information Processing Systems*, 36:32215–32234, 2023. 6
- [38] Hang Zhao, Xavier Puig, Bolei Zhou, Sanja Fidler, and Antonio Torralba. Open vocabulary scene parsing. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2002–2010, 2017. 2
- [39] 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. 2, 5, 6
- [40] Ziqin Zhou, Yinjie Lei, Bowen Zhang, Lingqiao Liu, and Yifan Liu. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11175–11185, 2023. 3, 6