

EOV-Seg: Efficient Open-Vocabulary Panoptic Segmentation

Hongwei Niu^{1,2*}, Jie Hu^{3*}, Jianghang Lin^{1*}, Guannan Jiang⁴, Shengchuan Zhang^{1†}

¹Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University, 361005, P.R. China

²Institute of Artificial Intelligence, Xiamen University, Fujian, China

³National University of Singapore, Singapore

⁴Contemporary Amperex Technology Co., Limited (CATL), Fujian, China
{niuhw649, hunterjlin007}@stu.xmu.edu.cn

hujie.cpp@gmail.com, jianggn@catl.com, zsc.2016@xmu.edu.cn

Abstract

Open-vocabulary panoptic segmentation aims to segment and classify everything in diverse scenes across an unbounded vocabulary. Existing methods typically employ two-stage or single-stage framework. The two-stage framework involves cropping the image multiple times using masks generated by a mask generator, followed by feature extraction, while the single-stage framework relies on a heavyweight mask decoder to make up for the lack of spatial position information through self-attention and cross-attention in multiple stacked Transformer blocks. Both methods incur substantial computational overhead, thereby hindering the efficiency of model inference. To fill the gap in efficiency, we propose EOV-Seg, a novel single-stage, shared, efficient, and spatial-aware framework designed for open-vocabulary panoptic segmentation. Specifically, EOV-Seg innovates in two aspects. First, a Vocabulary-Aware Selection (VAS) module is proposed to improve the semantic comprehension of visual aggregated features and alleviate the feature interaction burden on the mask decoder. Second, we introduce a Two-way Dynamic Embedding Experts (TDEE), which efficiently utilizes the spatial awareness capabilities of ViT-based CLIP backbone. To the best of our knowledge, EOV-Seg is the first open-vocabulary panoptic segmentation framework towards efficiency, which runs faster and achieves competitive performance compared with state-of-the-art methods. Specifically, with COCO training only, EOV-Seg achieves 24.5 PQ, 32.1 mIoU, and 11.6 FPS on the ADE20K dataset and the inference time of EOV-Seg is 4-19 times faster than state-of-the-art methods. Especially, equipped with ResNet50 backbone, EOV-Seg runs 23.8 FPS with only 71M parameters on a single RTX 3090 GPU.

Introduction

Panoptic segmentation aims to assign a semantic label and a unique instance identifier to each pixel in an image, actually covering both semantic segmentation and instance segmentation. Currently, several methods (Cheng et al. 2022; Hu et al. 2023; Cavagnero et al. 2024) have established a unified

*These authors contributed equally.

†Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

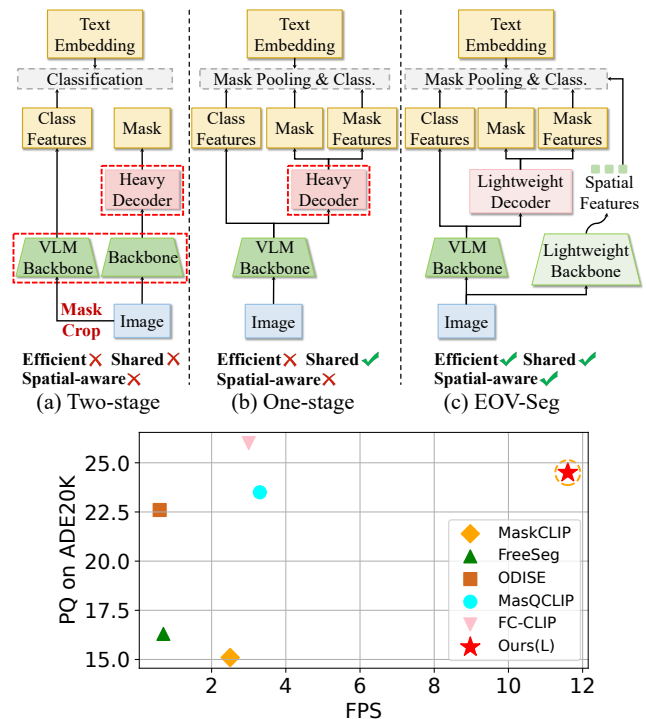



Figure 1: (Top) Comparisons between three macro-frameworks for open-vocabulary panoptic segmentation. (Bottom) FPS w.r.t. PQ on the ADE20K datasets for open-vocabulary panoptic segmentation.

framework capable of simultaneously addressing semantic, instance, and panoptic segmentation tasks. However, they are predominantly trained on specific datasets with small-scale predefined categories. This constraint significantly restricts their adaptability to novel scenarios that feature a diverse and extensive semantic vocabulary.

To overcome this limitation, there has been growing interest in more flexible open-vocabulary segmentation, as explored in works such as (Xu et al. 2022b; Liang et al. 2023; Qin et al. 2023; Zheng Ding 2023; Chen et al. 2023; Xu et al.



(a) Input Image (b) GT (c) Grad-CAM (d) K-means

Figure 2: Visualization of Grad-CAM (Gildenblat and contributors 2021) and K-means clustering of CNN-based CLIP (Radford et al. 2021) backbone features (ConvNeXt-L). (a) Input Image; (b) Ground Truth; (c) Grad-CAM shows the features focusing on local instances; (d) K-means clustering features produce semantically meaningful clusters.

2023a; Yu et al. 2023; VS et al. 2024). These approaches leverage pre-trained Visual-Language Models (VLMs) like CLIP (Radford et al. 2021) and ALIGN (Jia et al. 2021). Despite their innovative nature, these methods still encounter prevalent issues. Firstly, as shown in Fig. 1 (a), several methods (Xu et al. 2022b; Zheng Ding 2023; Liang et al. 2023) adopt a two-stage, non-shared, inefficient pipeline that first generates category-agnostic masks, followed by another Vision-Language Model (VLM) backbone that processes image crops obtained from these masks to extract features for individual classification, which results in high visual feature computational overhead and loss of contextual information. Secondly, alternatives (Yu et al. 2023; Li et al. 2024) employ a single-stage, shared, inefficient pipeline, as illustrated in Fig. 1 (b). They take a single shared frozen CNN-based CLIP visual encoder as backbone to extract multi-scale features, which is suitable for image segmentation tasks with high-resolution images. However, as shown in Fig. 2 (c) and (d), we experimentally observe that the CNN-based CLIP backbone only endows features with the capabilities to discriminate instances while lacking the spatial location capabilities. This results in that they often employ a heavyweight mask decoder consisting of many transformer layers, including self-attention and cross-attention mechanisms (self-attention captures contextual information between different queries, whereas cross-attention focuses on specific regions of the feature maps that are relevant to each query to provide awareness of spatial detail) to compensate for the lack of spatial aware capabilities, leading to unacceptable computational overhead and slower inference speed. In contrast, recent advances in Visual Foundation Models (VFM), such as DINOv2 (Oquab et al. 2023), SAM (Kirillov et al. 2023), and CLIP (Radford et al. 2021), those utilizing Vision Transformer (ViT) (Dosovitskiy et al. 2021) architectures as backbones, have demonstrated exceptional zero-shot generalization capabilities. Through qualitative analysis, as illustrated in Fig. 3, we observed nuanced distinctions in the feature representations across different blocks within these models, revealing insights into their discriminative and spatial localization capabilities. For instance, the clustering of features in the last block of SAM, depicted in Fig. 3 (c), demonstrates an ability for fine-grained spatial location. However, the process of se-

quentially calculating block features in obtaining these last block features at a significant computational cost. Conversely, as illustrated in Fig. 3 (d), the feature clustering of the first block in a ViT-based CLIP model reveals fine-grained object locations with precision comparable to SAM’s last block, yet achieves this with significantly less computational overhead. Therefore, it naturally fits the role of our spatial awareness extractor.

In response to these observations, we introduce EOVSeg: a novel single-stage, shared, efficient, and spatial-aware framework for open-vocabulary panoptic segmentation as illustrated in Fig. 1 (c). **1) Single-stage:** We integrate the mask generator and the text embeddings from CLIP text encoder as classifier into a single framework to implement end-to-end open-vocabulary panoptic segmentation. **2) Shared:** Instance features extracted by the VLM backbone are shared across classification and mask prediction while keeping image and text features aligned. **3) Efficient:** We propose a Vocabulary-Aware Selection (VAS) module to guide the visual aggregated features to select the features that are more relevant to the texts based on the semantic importance of the texts, thereby improving the semantic comprehension of visual aggregated features and alleviating the feature interaction burden on the mask decoder. This innovation allows for the deployment of a lightweight decoder as a mask generator, cutting down on computational demands and speeding up the inference process. **4) Spatial-aware:** In view of the observed advantages of ViT-based CLIP backbone, we introduce a Two-way Dynamic Embedding Experts (TDEE) to use weight allocation routers to evaluate the importance of embeddings experts and dynamically allocate expert weights to generate instance embeddings with semantic awareness and spatial awareness for mask recognition.

Extensive experiments show that EOVSeg runs faster and achieves competitive performance compared with state-of-the-art methods. In particular, with COCO training only, EOVSeg achieves 24.5 PQ, 32.1 mIoU, and 11.6 FPS on the ADE20K dataset. When taking ResNet50 as backbone, it runs 23.8 FPS with only 71M parameters on a single RTX 3090 GPU.

Related Work

Vision-Language Models. Vision-language pre-training aims to learn multimodal foundation models with improved performance on various vision-language tasks. Recently, large-scale Vision-Language Models (VLMs), such as CLIP (Radford et al. 2021) and ALIGN (Jia et al. 2021), which are contrastively pretrained on billion-scale, internet-sourced image-text pair datasets, have demonstrated remarkable zero-shot performance on image classification tasks. With the emergence of large Vision-Language Models (VLMs), they have been employed across various downstream vision tasks, including object detection (Gu et al. 2021; Zhong et al. 2022; Mi et al. 2022; Wu et al. 2023; Lin et al. 2024), image segmentation (Yue et al. 2024), video understanding (Weng et al. 2023; Rasheed et al. 2023), 3D scene understanding (Zhang et al. 2022; Zhu et al. 2023; Qu, Wang, and Qi 2023; Qu et al. 2024; Gong et al. 2024). In this paper, we are dedicated to exploring CLIP’s powerful



Figure 3: Visualization of K-means clustering of backbone features concerning different blocks across various VFMs.

instance discrimination and spatial awareness abilities to advance open-vocabulary panoptic segmentation.

Open-Vocabulary Panoptic Segmentation. Open-vocabulary panoptic segmentation is an emerging task in the field of image segmentation, aimed at evaluating the generalization ability to new visual categories that do not exist in the training set. Most open-vocabulary panoptic segmentation methods follow a two-stage pipeline. For example, MaskCLIP (Zheng Ding 2023) refines generated masks using transformer-based visual encoder and relative mask attention mechanism, combined with pre-trained CLIP models for panoptic segmentation. FreeSeg (Qin et al. 2023) proposes a unified, universal framework to seamlessly handle multiple segmentation tasks through one-time training. OSPNet (Chen et al. 2023) carefully designs embedding modulation modules and multiple fine-grained components for open-vocabulary segmentation. In contrast, some methods adopt a single-stage framework. ODISE (Xu et al. 2023a) uses a text-to-image diffusion model to generate mask proposals and perform classification. (Yu et al. 2023; Li et al. 2024) builds a single-stage framework by using a frozen CNN-based CLIP backbone, significantly surpassing previous two-stage methods. However, these methods suffer from high computational overhead, single-task focus, or lack of spatial location information.

Method

Problem Definition. Given an input image $I \in \mathbb{R}^{3 \times H \times W}$ and a set of candidate class labels $C \in C_{train} \cup C_{test}$. Open-vocabulary panoptic segmentation aims to segment image I into a set of masks $M \in \{0, 1\}^{N \times H \times W}$, each associated with a class label C , where N is the number of masks. The key challenge of this task arises from training exclusively with classes in C_{train} , while during inference, C_{test} may contain classes that are not encountered in C_{train} .

Single-Stage, Shared, Efficient Framework

Initially, we design a vanilla architecture of the proposed EOVSeg, which consists of a feature extractor (CNN-based

CLIP backbone), a lite aggregator, a lightweight decoder and a CLIP text encoder. In the following sections, we detail our elegant designs for elevating the vanilla baseline into the powerful EOVSeg framework.

Feature Extractor. Since the ViT-based backbone is more sensitive to the input resolution, we take CNN-based CLIP visual encoder as the backbone to extract multi-scale vision features C_i , $i \in \{2, 3, 4, 5\}$ from high-resolution images, with resolutions of $\frac{1}{4}$, $\frac{1}{8}$, $\frac{1}{16}$, and $\frac{1}{32}$, respectively.

Lite Aggregator. After obtaining multi-scale vision features C_i , $i \in \{2, 3, 4, 5\}$, we use a modulated deformable convolutional (Dai et al. 2017) feature pyramid structure to enhance and fuse features from different scales to obtain the intermediate FPN features P_i , $i \in \{2, 3, 4, 5\}$. To further reduce computational overhead and capture contextual information of different granularities, different scales features are aggregated to obtain the visual aggregated features $F_{agg} \in \mathbb{R}^{D \times \frac{H}{4} \times \frac{W}{4}}$ as follows:

$$F_{agg} = Fuse \left(\sum_{i=3}^N Up_i (Conv(P_i)) + Conv(P_2) \right), \quad (1)$$

where Up_i indicates an $2^{(i-2)} \times$ bilinear interpolation operation, N is number of feature layers, $Conv$ and $Fuse$ is 1×1 and 3×3 convolutions, respectively, used for adjusting feature dimensions and further fusion of the visual aggregated features F_{agg} .

Lightweight Decoder. As shown in the red part of Fig. 4, unlike the heavyweight mask decoder of the Mask2Former (Cheng et al. 2022), our decoder only consists of three modules at each layer in the following order: a initial-attention module, a dynamic depthwise attention module, and a late-attention module. Specifically, to further reduce the computational overhead of subsequent feature interactions, the initial-attention module is designed to extract the initial-attention features $F_{init} \in \mathbb{R}^{N \times D}$ by performing the dot product of the visual-semantic aggregated features $\hat{F}_{agg} \in \mathbb{R}^{D \times \frac{H}{4} \times \frac{W}{4}}$ from VAS module (which will be introduced in the following section) with the init masks or

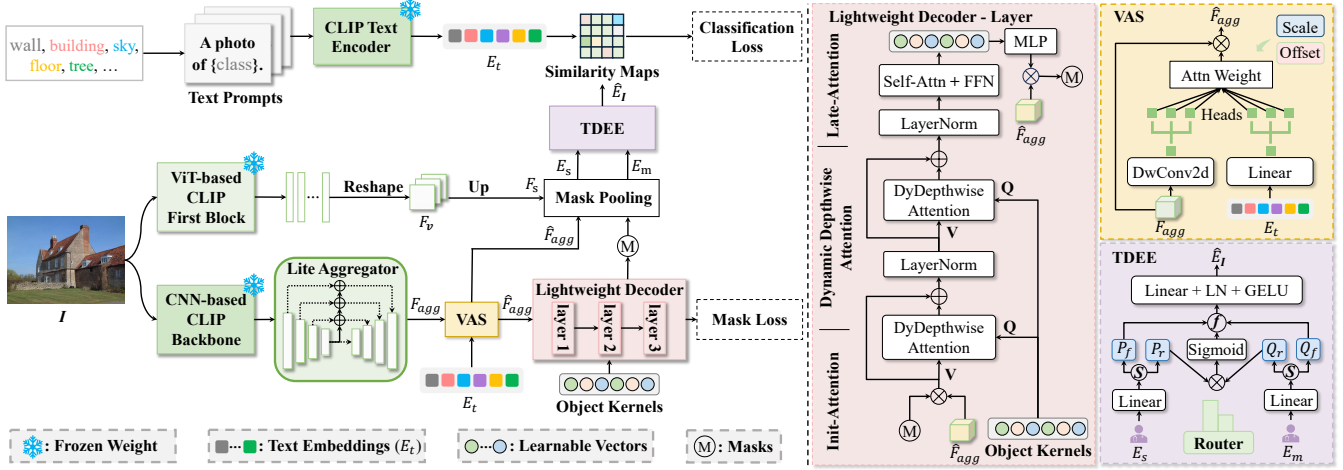


Figure 4: **Overview of EOVSeg.** First, the initial block of the ViT-based CLIP is used as a spatial awareness extractor to obtain spatial awareness features F_s . Then, the visual-semantic aggregated features \hat{F}_{agg} generated by VAS module, the masks generated by the light-weight decoder, and the spatial awareness features F_s are fed into mask pooling and TDEE sequentially to obtain instance embeddings \hat{E}_I , which will be used to calculate cosine similarity with E_t for classification.

masks $M \in \mathbb{R}^{N \times \frac{H}{4} \times \frac{W}{4}}$ from previous layer. Next, dynamic depthwise attention is used to conduct cross-dimensional feature interaction between a set of learnable object kernels $K \in \mathbb{R}^{N \times D}$ and the initial-attention features F_{init} , producing the refined object kernels $\hat{K} \in \mathbb{R}^{N \times D}$ as follows:

$$\begin{aligned} \hat{K} &= DyDepthwiseAttention(K, F_{init}) \\ &= F_{init} * r(KW_m) \end{aligned} \quad (2)$$

where $W_m \in \mathbb{R}^{D \times m}$ is a linear layer that projects K to generate kernels for depthwise convolution, $r(\cdot)$ is view operation that reshapes the input to $N \times 1 \times m$, $*$ denotes the Conv1d operation that takes $r(\cdot)$ as parameters and F_{init} as input. The refined object kernels \hat{K} further concentrate on the relation between different object kernels to enrich information by a multi-head self-attention and a feed-forward network (FFN). These refined object kernels \hat{K} function similarly to a Region Proposal Network (Ren et al. 2016), offering the potential to generate mask proposals. Then, we mapped the refined object kernels \hat{K} to be mask kernels through a three-layer MLP. Each binary mask $M_i \in \mathbb{R}^{H \times W}$ can be obtained by performing a dot product between the i^{th} mask kernel and the visual-semantic aggregated features \hat{F}_{agg} . Additionally, the mask embeddings $E_m \in \mathbb{R}^{N \times D}$ is obtained by mask pooling over the visual-semantic aggregated features \hat{F}_{agg} with predicted masks M .

Text Encoder. Given N class labels, we follow the precedent of prompt engineering (Gu et al. 2021) by providing M prompt templates for each class, such as ‘‘a photo of a {class}.’’. We use a frozen CLIP text encoder to encode all text prompts and then average them across the dimension of the number of templates to obtain the text embeddings $E_t \in \mathbb{R}^{N_{class} \times D}$.

Vocabulary-Aware Selection

To alleviate the feature interaction burden on the decoder and facilitate the deployment of a light-weight decoder, we propose Vocabulary-Aware Selection (VAS) as illustrated in the yellow part of Fig. 4. It can guide the selection of image features based on the semantic importance of the text, thereby aggregating features that are more relevant to the input text and enhancing the semantic understanding of visual features. Specifically, the visual aggregated features F_{agg} from the Lite Aggregator and the text embeddings E_t from the text encoder are mapped to the same feature dimension through linear layers, respective. Next, they are divided into multi-head visual aggregated features $\hat{F}_{agg} \in \mathbb{R}^{h \times \frac{D}{h} \times \frac{H}{4} \times \frac{W}{4}}$ and multi-head text embeddings $\hat{E}_t \in \mathbb{R}^{N_{class} \times h \times \frac{D}{h}}$ along the channel dimension, where h denotes the number of heads. This allows us to process complex semantic feature relationships in different subspaces. Then, they are processed through matrix operations to obtain multi-head attention $A \in \mathbb{R}^{h \times \frac{H}{4} \times \frac{W}{4} \times N_{class}}$, which is then subjected to a softmax function to produce smoothed logits. Following this, the maximum value is selected across the vocabulary dimension to acquire vocabulary-aware attention weights $\ddot{A} \in \mathbb{R}^{h \times \frac{H}{4} \times \frac{W}{4}}$:

$$\ddot{A} = Max^{(-1)}(Softmax^{(-1)}(\sum_d \hat{F}_{agg} \cdot \hat{E}_t)), \quad (3)$$

where (-1) denotes the vocabulary dimension. Additionally, two learnable parameters, namely scale factor γ and offset factor δ , are introduced to enable the network to adaptively adjust the vocabulary-aware attention weights \ddot{A} under various semantic distributions. Finally, the vocabulary-aware attention weights \ddot{A} are extended along the channel dimension, followed by element-wise multiplication with the multi-head aggregated features \hat{F}_{agg} . The weighted features

are then reshaped into the original aggregated features dimension to obtain the visual-semantic aggregated features $\hat{F}_{agg} \in \mathbb{R}^{D \times \frac{H}{4} \times \frac{W}{4}}$:

$$\hat{F}_{agg} = Reshape((\gamma \cdot \ddot{A} + \delta) \cdot \ddot{F}_{agg}), \quad (4)$$

Two-way Dynamic Embedding Experts

According to the analysis in the introduction section, spatial location information is essential for features that possess only instance-level discrimination. To efficiently extract spatial location information, we take the first block of the CLIP ViT-B visual encoder as spatial awareness extractor. Specially, the input image I is processed by the first block of ViT-B V_{b_0} to generate visual tokens (excluding the class token), which are then reshaped to two-dimensional spatial dimensions to obtain the visual-spatial features F_v . Subsequently, we employ two transposed convolutions for $4 \times$ upsampling to obtain the spatial awareness features $F_s \in \mathbb{R}^{D \times \frac{H}{4} \times \frac{W}{4}}$. Similar to mask embeddings E_m , mask pooling is used to extract regions of interest from spatial awareness features F_s as the spatial awareness embeddings $E_s \in \mathbb{R}^{N \times D}$.

Moreover, we propose Two-way Dynamic Embedding Experts (TDEE) as illustrated in the purple part of Fig. 4, which can adaptively incorporate spatial location information and instance-level discriminative information into mask embeddings to improve mask recognition capabilities under different data distributions and a wide range of semantic categories. Inspired by (Jacobs et al. 1991), we model the mask embeddings E_m and the spatial awareness embeddings E_s as experts with spatial awareness and instance-level discrimination capabilities, respectively. They should fulfill their respective roles while also synergistically collaborating. However, directly correlating the two using fixed weight coefficients prevents the network from perceiving the varying importance of embeddings, leading to either treating all embeddings equally or introducing bias. Therefore, a weight allocation router is introduced to estimate the significance of the current expert embeddings and adaptively adjust the correlation weight coefficients, thereby facilitating a better integration between the two. Specifically, the dynamic parameters $P \in \mathbb{R}^{N \times d}$ for the mask embeddings E_m and $Q \in \mathbb{R}^{N \times d}$ for the spatial awareness embeddings E_s are generated through linear layers, respectively. The dynamic parameters are evenly split along the channel dimension into router parameters $P_r \in \mathbb{R}^{N \times d/2}$ and $Q_r \in \mathbb{R}^{N \times d/2}$, and fusion parameters $P_f \in \mathbb{R}^{N \times d/2}$ and $Q_f \in \mathbb{R}^{N \times d/2}$. Next, total router parameters $P_t \in \mathbb{R}^{N \times d/2}$ are aggregated via an element-wise product between P_r and Q_r . Two router linear layers followed by sigmoid activation functions are then employed to provide the experts with adaptive weighting ability, as follows:

$$(P_f, P_r) = s(E_m W_m), \quad (5)$$

$$(Q_f, Q_r) = s(E_s W_s), \quad (6)$$

$$P_t = P_r \cdot Q_r, \quad (7)$$

$$\alpha_m = \sigma(LN(Linear_m(P_t))), \quad (8)$$

$$\alpha_s = \sigma(LN(Linear_s(P_t))), \quad (9)$$

where $W_m \in \mathbb{R}^{D \times d}$, $W_s \in \mathbb{R}^{D \times d}$ project E_m and E_s to generate dynamic parameters, respectively. $s(\cdot)$ is split operation along the channel dimension. σ is sigmoid activation function, LN indicates layer normalization. This enables the router to adaptively assign weights to the embeddings experts. Finally, a weighted summation is performed as follows:

$$\hat{E}_I = \alpha_m \cdot LN(P_f) + \alpha_s \cdot LN(Q_f), \quad (10)$$

To further aggregate, a linear layer, LayerNorm (LN), and GELU are employed to obtain the instance embeddings \hat{E}_I . In general, TDEE takes advantage of the weight allocation router to adaptively assign individual weights to each embeddings expert based on the significance of the embeddings. As a result, the spatial awareness embeddings E_s , which have strong spatial awareness capabilities, are integrated with the mask embeddings E_m , which contain only instance discrimination information. This integration improves the spatial and semantic understanding of the model in open scenes.

Experiments

Implementation Details

Architecture. We use a CNN-based CLIP (Radford et al. 2021) visual encoder as backbone, S, M, L denote ResNet50, ResNet50x4, ConvNeXt-Large, respectively. Additionally, we utilize the first block of the ViT-B/16 CLIP model as a spatial awareness extractor.

Training Strategy. Following prior works (Xu et al. 2023a; Yu et al. 2023), the training batch size is 16. We train our EOVSeg on 4 NVIDIA 3090 GPUs for a total of 200k iterations. We only use the COCO Panoptic (Lin et al. 2014) dataset for training, with a crop size of 1024×1024 .

Evaluation Protocols. We evaluated our EOVSeg for open-vocabulary semantic, instance, and panoptic segmentation on the ADE20K (Zhou et al. 2017) dataset, and for semantic segmentation on the ADE20K (Zhou et al. 2017), PASCAL Context (Mottaghi et al. 2014) and PASCAL VOC (Everingham et al. 2010) datasets. During inference, the shortest side of input images is resized to 640, while ensuring the longer side does not exceed 2560. See Appendix for more details.

Main Results

Open-Vocabulary Panoptic Segmentation. Tab. 1 reports the performance of different open vocabulary panoptic segmentation methods on the ADE20k dataset. Compared to existing state-of-the-art methods, EOVSeg (L) runs faster and achieves competitive performance. Specifically, although ODISE (Xu et al. 2023a) take powerful text-to-image diffusion (Rombach et al. 2022) model as backbone and is trained on the additional COCO Caption (Chen et al. 2015) dataset, EOVSeg (L) outperforms ODISE by +1.1 PQ with faster speed ($19 \times$), fewer parameters (-1297M), and lower computational complexity (-576GFLOPs). Compared to FC-CLIP (Yu et al. 2023), EOVSeg (L) is about 4 times faster than it and still shows competitive performance. It is worth mentioning that compared to other state-of-the-art methods like MasQCLIP (Xu et al. 2023c), OP-SNet (Chen et al. 2023), FreeSeg (Qin et al. 2023), and

Method	VLM	Backbone	Training datasets	PQ	SQ	RQ	AP	mIoU	FPS \uparrow	Params \downarrow	FLOPs \downarrow
OPNet (Chen et al. 2023)	ResNet50	Swin-L	COCO Panoptic+IM-21K	19.0	52.4	23.0	-	-	-	-	485G
ODISE (Xu et al. 2023a)	ViT-L/14	Stable Diffusion	COCO Panoptic	22.6	65.1	27.0	14.4	29.9	0.6	1522M	953G
ODISE (Xu et al. 2023a)	ViT-L/14	Stable Diffusion	COCO Panoptic+Caption	23.4	78.1	28.3	13.9	28.7	0.6	1522M	953G
FreeSeg (Qin et al. 2023)	ViT-B/16	ResNet101	COCO-Stuff-156	16.3	71.8	21.6	6.5	24.6	0.7	219M	-
MaskCLIP (Zheng Ding 2023)	ViT-L/14@336px	ResNet50	COCO Panoptic	15.1	70.5	19.2	6.0	23.7	2.5	367M	542G
FC-CLIP [†] (Yu et al. 2023)	ConvNeXt-L	shared	COCO Panoptic	26.0	70.6	31.3	16.4	33.6	3.0	221M	381G
MasQCLIP (Xu et al. 2023c)	ViT-L/14@336px	ResNet50	COCO Panoptic	23.5	69.6	29.2	12.8	30.3	3.3	373M	357G
EOV-Seg (L)	ConvNeXt-L	shared	COCO Panoptic	24.5	70.2	30.1	13.7	32.1	11.6	225M	377G
EOV-Seg (M)	ResNet50x4	shared	COCO Panoptic	18.7	63.5	23.2	8.5	25.5	18.4	127M	176G
EOV-Seg (S)	ResNet50	shared	COCO Panoptic	15.1	57.0	18.9	7.2	21.9	23.8	71M	133G

Table 1: Open-vocabulary panoptic segmentation performance for training on COCO and testing on ADE20K. “ \dagger ” means the results are from a re-implementation of FC-CLIP (Yu et al. 2023). “ \uparrow ” and “ \downarrow ” means the larger or smaller the value, the better. “IM-21K” denotes ImageNet-21K datasets.

Method	VLM	Backbone	Training datasets	A-847	PC-459	A-150	PC-59	PAS-20	FPS \uparrow	Params \downarrow	FLOPs \downarrow
GroupViT (Xu et al. 2022a)	ViT-S/16	ViT-S/16	GCC+YFCC	4.3	4.9	10.6	25.9	50.7	-	-	-
LSeg+ (Ghiasi et al. 2022)	ALIGN	ResNet101	COCO-Stuff	2.5	5.2	13.0	36.0	59.0	-	-	-
OpenSeg (Ghiasi et al. 2022)	ALIGN	ResNet101	COCO Panoptic+LN	4.4	7.9	17.5	40.1	-	-	-	-
OVSeg (Liang et al. 2023)	ViT-L/14	Swin-B	COCO-Stuff+Caption	9.0	12.4	29.6	55.7	94.5	1.1	531M	830G
CAT-Seg (Cho et al. 2024)	ViT-L/14	Swin-B	COCO-Stuff	10.8	20.4	31.5	62.0	96.6	2.2	490M	473G
DeOP (Han et al. 2023)	ViT-B/16	ResNet101c	COCO-Stuff-156	7.1	9.4	22.9	48.8	91.7	2.5	211M	480G
OVSeg (Liang et al. 2023)	ViT-B/16	ResNet101c	COCO-Stuff+Caption	7.1	11.0	24.8	53.3	92.6	3.4	211M	830G
ZegFormer (Ding et al. 2022)	ViT-B/16	ResNet101	COCO-Stuff-156	4.9	9.1	16.9	42.8	86.2	5.4	211M	456G
ZegFormer (Ding et al. 2022)	ViT-B/16	ResNet101	COCO-Stuff	5.6	10.4	18.0	45.5	89.5	5.4	211M	456G
SAN (Xu et al. 2023b)	ViT-L/14	ViT-liked	COCO-Stuff	12.4	15.7	32.1	57.7	94.6	8.0	437M	403G
EOV-Seg (L)	ConvNeXt-L	shared	COCO Panoptic	12.8	16.8	32.1	56.9	94.8	11.8	225M	377G
EOV-Seg (M)	ResNet50x4	shared	COCO Panoptic	7.8	12.2	25.5	51.8	91.2	18.9	127M	176G
EOV-Seg (S)	ResNet50	shared	COCO Panoptic	6.6	11.5	21.9	46.0	87.2	24.5	71M	133G

Table 2: Open-vocabulary semantic segmentation performance. We report the mIoU results on five widely used test sets for open-vocabulary semantic segmentation. “ \uparrow ” and “ \downarrow ” means the larger or smaller the value, the better. “LN” denotes Localized Narrative datasets.

MaskCLIP (Zheng Ding 2023), EOv-seg shows significant improvements, such as PQ increased by +1.0 to +9.4, SQ increased by +0.6 to +17.9, RQ increased by +0.9 to +10.9, mIoU increased by +1.8 to +8.4 and the inference speed is 3.5 to 19.3 times faster. These results show that EOv-Seg achieves the optimal trade-off between performance and speed.

Task-general EOv-Seg Is All You Need. In Tab. 2, EOv-Seg delivers performance comparable to current state-of-the-art models and it boasts the fastest inference speed among its peers. Specially, EOv-Seg (L) outperforms CAT-Seg (Cho et al. 2024) by +2.0 mIoU on the A-150 dataset and is 5.3 \times faster, has fewer parameters (-265M) and lower FLOPs than CAT-Seg. These results demonstrate that even though EOv-Seg is trained only on the COCO Panoptic dataset with fewer semantic classes, it successfully adapts to semantic segmentation tasks, generalizes efficiently to unknown categories, achieves accurate segmentation, and runs faster with fewer parameters.

Ablation Study

Spatial Awareness Extractor. In Tab. 3, we compared the performance and efficiency of using ViT-B versions of different pre-trained models as spatial-aware extractors. Specifically, when using different blocks of the DINOv2 backbone as spatial awareness extractor, it shows poor performance. As shown in Fig. 3 (b), it contains a lot of noise

Extractor	index	PQ	SQ	RQ	mIoU	FPS \uparrow	Params \downarrow
DINOv2	0	21.9	67.3	27.3	29.5	11.7	225M
	5	22.4	67.7	28.8	30.1	8.9	247M
	11	23.1	68.2	29.4	31.2	5.8	260M
SAM	0	22.4	68.3	27.8	30.8	11.6	225M
	5	23.7	70.4	28.8	31.6	8.7	247M
	11	24.9	71.2	29.5	32.5	5.2	260M
ViT-based CLIP	0	24.5	70.2	30.1	32.1	11.6	225M
	5	23.8	68.6	29.4	32.0	9.1	247M
	11	23.7	68.5	29.6	31.8	5.9	260M

Table 3: Ablation studies on different blocks of different pre-trained models as spatial awareness extractor. index: denotes block index of the pre-trained models.

and lacks fine-grained spatial information, we argue the difficulty in learning feature space patterns due to self-supervised pre-training in DINOv2, coupled with the use of unlabeled data, results in diminished instance discrimination and spatial awareness capabilities. Since SAM is pre-trained on the fine-grained SA-1B (Kirillov et al. 2023) dataset, it exhibits enhanced fine-grained spatial representation capabilities in its later blocks, progressively improving the segmentation quality. However, the increase in RQ is modest, likely constrained by its class-agnostic pre-training. One key finding is that the first ViT-based block demonstrates compa-

Fusion	VAS	PQ	SQ	RQ	mIoU	FPS \uparrow	Params \downarrow
None	\times	22.1	67.2	27.3	30.3	13.8	218M
	\checkmark	22.7	67.3	28.2	31.4	13.6	219M
+EAF	\times	23.0	68.2	28.3	30.8	12.2	223M
	\checkmark	22.3	66.9	27.6	30.4	11.8	224M
+SDI	\times	22.5	67.1	27.8	31.1	11.4	223M
	\checkmark	22.9	67.4	28.5	31.5	11.0	224M
+TDEE	\times	23.6	68.7	30.0	31.1	11.9	225M
	\checkmark	24.5	70.2	30.1	32.1	11.6	225M

Table 4: Ablation studies on key components for EOV-Seg. EAF: Early Aggregation Fusion; SDI: Separable Dynamic Interaction; TDEE: Two-way Dynamic Embedding Experts; VAS: Vocabulary-Aware Selection;

Feat. Inter.	Feat. Scale	PQ	mIoU	FPS \uparrow	Params \downarrow
CA	multi-scale	24.7	32.5	3.4	304M
CA	single-scale	23.8	30.4	6.6	282M
DDA	multi-scale	22.7	29.7	9.1	245M
DDA	single-scale	24.5	32.1	11.6	225M

Table 5: Ablation studies on different feature interaction methods in lightweight decoder with different feature scales. CA: Cross-Attention; DDA: Dynamic Depthwise Attention.

rable performance to SAM, with about $2.2\times$ faster inference speed and fewer parameters ($\sim 35M$). This is consistent with the observation in Fig. 3 (d). As the number of blocks increases, performance is saturated and has a downward trend. To achieve the best trade-off between performance and inference speed, the first block of the ViT-based CLIP is the best choice.

Our Two Innovative Components. In Tab. 4, we designed two additional variants of spatial awareness fusion, namely EAF and SDI, for more details in the appendix. All three variants of spatial awareness fusion significantly improved performance, demonstrating the importance of spatial location information for the lightweight decoder. When we introduce the VAS module for EAF, the performance tends to decline. This phenomenon is attributed to the early integration of features, which changes the spatial distribution of features and leads to the loss of key spatial information in important semantic features. However, VAS can cooperate with SDI or TDEE spatial awareness fusion module can bring additional performance improvements. In summary, We finally choose the TDEE and VAS modules, which bring performance improvements of $+2.4$ PQ, $+3.0$ SQ, $+2.8$ RQ and $+1.8$ mIoU without introducing additional computational burden.

Feature Interaction Methods. In Tab. 5, we explored the performance differences between Cross-Attention (CA) and Dynamic Depthwise Attention (DDA) two attention modules and their interaction with features of different scales. When using CA with multi-scale features only yields marginal improvements over our method, with a $+0.2$ increase in PQ and $+0.4$ in mIoU, but results in a significant slowdown, approximately $3.4\times$ slower in FPS. Since single scale avoids the fuzzy spatial distribution caused by spatial

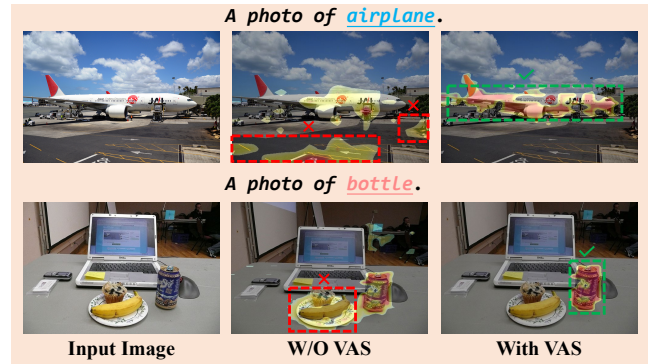


Figure 5: Visualization of similarity map in VAS.




Figure 6: Visualization of segmentation results.

dimension alignment, and the pre-perception of VAS can reduce the decoder interaction burden, using DDA with single-scale can bring a performance improvement of $+1.8$ PQ and $+2.4$ mIoU over DDA with multi-scale.

Qualitative Analysis

Alleviate Mask Decoder Burden. In Fig. 5, we visualize the similarity map between single text embedding and visual-semantic aggregated features. We found that without VAS, the features may focus on objects with similar colors or nearby positions to the prompt object, confounding feature discrimination. After VAS enhancement, the features will be more focused on the text-guided area.

Infusion of Spatial Awareness. We visualize the segmentation results with and without Two-way Dynamic Embedding Experts (TDEE), respectively. As shown in the red box in Fig. 6, using only the masks generated by the lightweight decoder will result in category confusion and incomplete segmentation. This demonstrates that spatial awareness capabilities is very critical for lightweight decoder.

Conclusion

We propose EOV-Seg, an efficient framework designed for open-vocabulary panoptic segmentation. It uses VAS and TDEE to improve instance discrimination capabilities for alleviating the interaction burden on mask decoder and endow mask recognition with spatial awareness capabilities. Extensive experiments show that our EOV-Seg is effective in terms of both accuracy and speed.

Acknowledgments

This work was supported by National Science and Technology Major Project (No.2022ZD0118201), the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No.U21B2037, No.U22B2051, No.U23A20383, No.U21A20472, No.62176222, No.62176223, No.62176226, No.62072386, No.62072387, No.62072389, No.62002305 and No.62272401), and the Natural Science Foundation of Fujian Province of China (No.2021J06003, No.2022J06001).

References

- Cavagnero, N.; Rosi, G.; Cuttano, C.; Pistilli, F.; Ciccone, M.; Averta, G.; and Cermelli, F. 2024. Pem: Prototype-based efficient maskformer for image segmentation. In *CVPR*.
- Chen, X.; Fang, H.; Lin, T.-Y.; Vedantam, R.; Gupta, S.; Dollár, P.; and Zitnick, C. L. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Chen, X.; Li, S.; Lim, S.; Torralba, A.; and Zhao, H. 2023. Open-vocabulary Panoptic Segmentation with Embedding Modulation. In *ICCV*.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girshick, R. 2022. Masked-attention mask transformer for universal image segmentation. In *CVPR*.
- Cho, S.; Shin, H.; Hong, S.; Arnab, A.; Seo, P. H.; and Kim, S. 2024. Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. In *CVPR*.
- Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; and Wei, Y. 2017. Deformable convolutional networks. In *ICCV*.
- Ding, J.; Xue, N.; Xia, G.-S.; and Dai, D. 2022. Decoupling zero-shot semantic segmentation. In *CVPR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshly, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2010. The Pascal Visual Object Classes (VOC) Challenge. *IJCV*.
- Ghiasi, G.; Gu, X.; Cui, Y.; and Lin, T.-Y. 2022. Scaling open-vocabulary image segmentation with image-level labels. In *ECCV*.
- Gildenblat, J.; and contributors. 2021. PyTorch library for CAM methods. <https://github.com/jacobgil/pytorch-grad-cam>.
- Gong, Y.; Zhong, Z.; Qu, Y.; Luo, Z.; Ji, R.; and Jiang, M. 2024. Cross-modality perturbation synergy attack for person re-identification. *arXiv preprint arXiv:2401.10090*.
- Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2021. Open-Vocabulary Detection via Vision and Language Knowledge Distillation. *arXiv preprint arXiv:2104.13921*.
- Han, C.; Zhong, Y.; Li, D.; Han, K.; and Ma, L. 2023. Open-vocabulary semantic segmentation with decoupled one-pass network. In *ICCV*.
- Hu, J.; Huang, L.; Ren, T.; Zhang, S.; Ji, R.; and Cao, L. 2023. You Only Segment Once: Towards Real-Time Panoptic Segmentation. In *CVPR*.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive mixtures of local experts. *Neural computation*.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollar, P.; and Girshick, R. 2023. Segment Anything. In *ICCV*.
- Li, X.; Yuan, H.; Li, W.; Ding, H.; Wu, S.; Zhang, W.; Li, Y.; Chen, K.; and Loy, C. C. 2024. OMG-Seg: Is one model good enough for all segmentation? In *CVPR*.
- Liang, F.; Wu, B.; Dai, X.; Li, K.; Zhao, Y.; Zhang, H.; Zhang, P.; Vajda, P.; and Marculescu, D. 2023. Open-vocabulary semantic segmentation with mask-adapted clip. In *CVPR*.
- Lin, J.; Shen, Y.; Wang, B.; Lin, S.; Li, K.; and Cao, L. 2024. Weakly supervised open-vocabulary object detection. In *AAAI*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*.
- Mi, P.; Lin, J.; Zhou, Y.; Shen, Y.; Luo, G.; Sun, X.; Cao, L.; Fu, R.; Xu, Q.; and Ji, R. 2022. Active teacher for semi-supervised object detection. In *CVPR*.
- Mottaghi, R.; Chen, X.; Liu, X.; Cho, N.-G.; Lee, S.-W.; Fidler, S.; Urtasun, R.; and Yuille, A. 2014. The Role of Context for Object Detection and Semantic Segmentation in the Wild. In *CVPR*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Qin, J.; Wu, J.; Yan, P.; Li, M.; Yuxi, R.; Xiao, X.; Wang, Y.; Wang, R.; Wen, S.; Pan, X.; et al. 2023. FreeSeg: Unified, Universal and Open-Vocabulary Image Segmentation. In *CVPR*.
- Qu, Y.; Dai, S.; Li, X.; Lin, J.; Cao, L.; Zhang, S.; and Ji, R. 2024. Goi: Find 3d gaussians of interest with an optimizable open-vocabulary semantic-space hyperplane. In *ACM MM*.
- Qu, Y.; Wang, Y.; and Qi, Y. 2023. Sg-nerf: Semantic-guided point-based neural radiance fields. In *ICME*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Rasheed, H.; Khattak, M. U.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Fine-tuned clip models are efficient video learners. In *CVPR*.

Ren, S.; He, K.; Girshick, R.; and Sun, J. 2016. Faster R-CNN: Towards real-time object detection with region proposal networks. *TPAMI*.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *CVPR*.

VS, V.; Borse, S.; Park, H.; Das, D.; Patel, V.; Hayat, M.; and Porikli, F. 2024. Possam: Panoptic open-vocabulary segment anything. *arXiv preprint arXiv:2403.09620*.

Weng, Z.; Yang, X.; Li, A.; Wu, Z.; and Jiang, Y.-G. 2023. Open-vclip: Transforming clip to an open-vocabulary video model via interpolated weight optimization. In *ICML*.

Wu, X.; Zhu, F.; Zhao, R.; and Li, H. 2023. Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. In *CVPR*.

Xu, J.; De Mello, S.; Liu, S.; Byeon, W.; Breuel, T.; Kautz, J.; and Wang, X. 2022a. Groupvit: Semantic segmentation emerges from text supervision. In *CVPR*.

Xu, J.; Liu, S.; Vahdat, A.; Byeon, W.; Wang, X.; and De Mello, S. 2023a. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *CVPR*.

Xu, M.; Zhang, Z.; Wei, F.; Hu, H.; and Bai, X. 2023b. Side adapter network for open-vocabulary semantic segmentation. In *CVPR*.

Xu, M.; Zhang, Z.; Wei, F.; Lin, Y.; Cao, Y.; Hu, H.; and Bai, X. 2022b. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *ECCV*.

Xu, X.; Xiong, T.; Ding, Z.; and Tu, Z. 2023c. MasQ-CLIP for Open-Vocabulary Universal Image Segmentation. In *ICCV*.

Yu, Q.; He, J.; Deng, X.; Shen, X.; and Chen, L.-C. 2023. Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. In *NeurIPS*.

Yue, P.; Lin, J.; Zhang, S.; Hu, J.; Lu, Y.; Niu, H.; Ding, H.; Zhang, Y.; Jiang, G.; Cao, L.; et al. 2024. Adaptive Selection based Referring Image Segmentation. In *ACM MM*.

Zhang, R.; Guo, Z.; Zhang, W.; Li, K.; Miao, X.; Cui, B.; Qiao, Y.; Gao, P.; and Li, H. 2022. Pointclip: Point cloud understanding by clip. In *CVPR*.

Zheng Ding, Z. T., Jieke Wang. 2023. Open-Vocabulary Universal Image Segmentation with MaskCLIP. In *ICML*.

Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; et al. 2022. Regionclip: Region-based language-image pretraining. In *CVPR*.

Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *CVPR*.

Zhu, X.; Zhang, R.; He, B.; Guo, Z.; Zeng, Z.; Qin, Z.; Zhang, S.; and Gao, P. 2023. Pointclip v2: Prompting clip and gpt for powerful 3d open-world learning. In *ICCV*.