

ProxyCLIP: Proxy Attention Improves CLIP for Open-Vocabulary Segmentation

Mengcheng Lan¹, Chaofeng Chen¹, Yiping Ke², Xinjiang Wang³,
Litong Feng³, and Wayne Zhang^{3,4*}

¹ S-Lab, Nanyang Technological University

² CCDS, Nanyang Technological University ³ SenseTime Research

⁴ Guangdong Provincial Key Laboratory of Digital Grid Technology
lanm0002@e.ntu.edu.sg {chaofeng.chen, ypke}@ntu.edu.sg
{wangxinjiang, fenglitong, wayne.zhang}@sensetime.com
<https://github.com/mc-lan/ProxyCLIP>

Abstract. Open-vocabulary semantic segmentation requires models to effectively integrate visual representations with open-vocabulary semantic labels. While Contrastive Language-Image Pre-training (CLIP) models shine in recognizing visual concepts from text, they often struggle with segment coherence due to their limited localization ability. In contrast, Vision Foundation Models (VFMs) excel at acquiring spatially consistent local visual representations, yet they fall short in semantic understanding. This paper introduces ProxyCLIP, an innovative framework designed to harmonize the strengths of both CLIP and VFMs, facilitating enhanced open-vocabulary semantic segmentation. ProxyCLIP leverages the spatial feature correspondence from VFMs as a form of proxy attention to augment CLIP, thereby inheriting the VFMs’ robust local consistency and maintaining CLIP’s exceptional zero-shot transfer capacity. We propose an adaptive normalization and masking strategy to get the proxy attention from VFMs, allowing for adaptation across different VFMs. Remarkably, as a *training-free* approach, ProxyCLIP significantly improves the average mean Intersection over Union (mIoU) across eight benchmarks **from 40.3 to 44.4**, showcasing its exceptional efficacy in bridging the gap between spatial precision and semantic richness for the open-vocabulary segmentation task.

Keywords: Semantic segmentation · Vision language model · Vision foundation model · Open vocabulary

1 Introduction

Open-vocabulary semantic segmentation [51, 63] aims to partition an image into coherent segments and assign each segment a semantic label from arbitrary visual concepts. The success of this task hinges on a model’s ability to generate

* Corresponding author.

semantically coherent visual representations that align well with textual descriptions. The emergence of Contrastive Language-Image Pre-training (CLIP) models [7, 16, 31, 50], trained on large-scale image-text pair datasets, has showcased remarkable capabilities in semantic understanding, such as image recognition [31] and image-text retrieval [8, 20]. However, while excelling in recognizing concepts within images, CLIP models often struggle with producing coherent image segments for dense prediction tasks due to limitations in localization ability. In contrast, Vision Foundation Models (VFMs) like self-supervised methods [4, 15, 29] and the Segment Anything Model (SAM) [17] have demonstrated proficiency in learning spatially consistent visual representations. These models have found extensive applications in unsupervised semantic segmentation [14, 19] and object detection [36, 37, 44] to generate *class-agnostic* image segments. However, they often lack semantic understanding, requiring fine-tuning on downstream tasks.

Significant efforts have been dedicated to enhance the open-vocabulary semantic segmentation capabilities of CLIP models. A key research area focuses on exploiting the inherent localization properties of CLIP. This involves modifying its attention mechanism during the inference stage to rearrange local spatial information [2, 22, 40, 63], or fine-tuning the CLIP model to improve its local consistency via the self-distillation technique [6, 27, 28, 46]. Despite their efficacy, these methods are constrained by the intrinsic performance of CLIP models, as the contrastive learning objective does not explicitly encourage learning spatially consistent local features [4]. Another research path emphasizes improving CLIP representations with semantically coherent properties through knowledge distillation [41, 47, 56] from VFM features [4, 17]. While these fine-tuning approaches are usually more effective, they carry the potential risk of sacrificing the strength of the original models during fine-tuning, which may yield sub-optimal results.

In this study, our aim is to integrate the expertise of CLIP models and VFMs in a *training-free manner* without compromising their individual strengths for the open-vocabulary semantic segmentation. The key principle is to borrow the robust feature correspondence properties from VFMs and maintain the zero-shot transfer ability of CLIP. We first verify the feasibility of our idea by comparing the semantic coherence between CLIP’s attention scores and VFMs’ feature correspondence maps. Our investigation reveals that VFMs’ feature correspondence significantly surpasses various types of attention scores in CLIP, highlighting the significant potential of their integration. Consequently, we propose ProxyCLIP, a general framework designed to harness VFMs’ feature correspondence as a proxy attention for CLIP.

Designing a universal proxy attention from various VFMs is challenging due to their diverse inductive biases in visual representations, which leads to the difficulties in ensuring the consistency and separateness of the proxy attention. Our research explores different configurations, leading to two crucial designs for improved performance in the open-vocabulary segmentation:

- We introduce the adaptive normalization and masking strategy to the VFMs’ feature correspondence, ensuring our proxy attention remains effective and consistent across different VFMs.

- We find that VFM architectures with smaller patch size produce clearer boundaries, which contributes to superior segmentation results.

To our knowledge, ProxyCLIP is the first training-free method to successfully achieve better performance with larger CLIP models, while previous approaches such as MaskCLIP [63] fail to benefit from stronger CLIP backbones. Extensive experiments conducted across 8 segmentation benchmarks establish ProxyCLIP as a new state-of-the-art solution for the open-vocabulary segmentation.

2 Related Work

Vision-language pre-training models. Vision-Language Models (VLMs) constitute a significant stride in AI, designed to imbibe generic multimodal representations through extensive training on large-scale image-text pair datasets. Recent advancements in contrastive vision-language pre-training, facilitated by dual encoders, have significantly enhanced the generalization capabilities of image recognition systems [7, 16, 21, 31, 50, 54, 57, 58]. Among these advancements, the CLIP series [7, 31, 50] stands out. Pre-trained on billion-scale image-text pairs, CLIP models exhibit exceptional performance in zero-shot image classification and image-text retrieval tasks. However, despite their success, the reliance on image-level supervision often poses challenges in aligning local image regions with corresponding text, leading to subpar performance in dense prediction tasks [52, 61, 63]. Towards this challenge, various weakly-supervised approaches have emerged, aiming to indirectly align region and language representations using solely image-text pairs [27, 46]. Additionally, other works [24, 32, 60] have sought to perform region-text alignment leveraging visual grounding datasets [18, 30].

Vision foundation models. The field of computer vision has witnessed the rise of potent models built upon Vision Transformers (ViT) [12]. Recent advancements in self-supervised learning (SSL) [4, 11, 15, 29, 64] have yielded features with strong localization properties, opening avenues for novel applications across various downstream tasks. These features have been effectively harnessed in diverse areas, including unsupervised semantic segmentation [14, 19, 55], unsupervised object detection [36, 37, 44], unsupervised instance segmentation [42, 43], and zero-shot semantic segmentation [1, 45, 47, 48]. DINO [4] and DINOv2 [29] have contributed significantly to the advancement of SSL models, demonstrating remarkable capabilities in producing spatially coherent visual representations. More recently, the Segment Anything Model (SAM) [17] has emerged as proficient in the object segmentation, adeptly producing precise masks within images, albeit limited to class-agnostic masks. Our work introduces a generic framework designed to harness the spatial coherence properties of VFMs’ features to enhance the dense prediction performance of CLIP models.

Open-vocabulary semantic segmentation. The emergence of vision-language pre-training models has spurred recent advancements in open-vocabulary semantic segmentation methods. Many studies [2, 22, 23, 38, 40, 63] aim to leverage

CLIP’s inherent localization properties for the dense prediction tasks. For instance, MaskCLIP [63] demonstrates that the *value* embeddings of the last layer of CLIP exhibit superior localization capabilities compared to token embeddings, resulting in impressive performance in open-vocabulary dense prediction tasks. Additionally, some researchers [2, 22, 40] have identified noisy attention maps as a bottleneck in CLIP’s performance on dense prediction. To address this, they propose to revise the vanilla *query-key* attention to *self-self* attention, such as *value-value* [22], *query-query*, *key-key* attention, or their combination [2, 40], leading to substantial performance improvements. Despite these efforts, these methods are constrained by the intrinsic performance of CLIP models. Besides these training-free methods, various weakly-supervised retraining methods [5, 25, 27, 33, 51, 53, 59] have been proposed to address the limitations of CLIP. These methods incorporate group tokens [25, 51, 53] or prototype tokens [59] into the image encoder to enhance object discovery and image segmentation. However, this enhancement may come at the cost of compromising some of CLIP’s open-vocabulary properties. Related approaches include recent CLIP-DINOiser [47] and SAM-CLIP [41], which integrate the favorable spatial consistent properties of DINO [4] and SAM [17] with CLIP through knowledge distillation techniques, aiming to enhance the local patch representations of images for the dense prediction tasks. Besides, ReCo [35] proposes to dynamically curate training sets by retrieving unlabeled images using CLIP, and to co-segment all entities among these collections based on the visual representations of DINO [4]. In this work, we attempt to seamlessly integrate the strengths of CLIP and VFMs in a training-free manner by designing a novel proxy attention module, without compromising their unique strengths, ultimately striving to achieve state-of-the-art results.

3 Method

3.1 Preliminaries and Motivation

CLIP’s representations. The CLIP ViT architecture primarily consists of residual attention blocks. Let $\mathbf{x} = [x_0, x_1, \dots, x_L]^T$ denotes the input to the last block, where x_0 represents the global class embedding, and $\{x_i | i = 1, 2, \dots, L\}$ denotes local patch embeddings. The forward process of a residual attention block is as follows:

$$\mathbf{q} = \text{Emb}_q(\mathbf{x}), \mathbf{k} = \text{Emb}_k(\mathbf{x}), \mathbf{v} = \text{Emb}_v(\mathbf{x}), \quad (1)$$

$$\mathbf{y} = \mathbf{x} + \text{Proj}(\text{Attn}_{qk} \cdot \mathbf{v}), \quad (2)$$

$$\mathbf{z} = \mathbf{y} + \text{FFN}(\text{LN}(\mathbf{y})), \quad (3)$$

where \mathbf{q} , \mathbf{k} and \mathbf{v} represent the *query*, *key* and *value* embeddings, respectively. Emb comprises a layer norm (LN) and a projection layer, and FFN is a feed-forward network. $\text{Attn}_{qk} = \text{SoftMax}(\frac{\mathbf{q}\mathbf{k}^T}{\tau})$ denotes the attention score with a constant scaling factor τ . Typically, $z_{\text{image}} = \mathbf{z}[0] \in \mathbb{R}^{1 \times d}$ is selected as

the class representation with dimension d , representing the entire image, while $\mathbf{z}_{\text{dense}} = \mathbf{z}[1 : L] \in \mathbb{R}^{L \times d}$ denotes the feature representations. For brevity, we omit the final output layer following the last attention block. Additionally, the class representation is removed from the visual representations of both CLIP models and VFMs in the subsequent analysis.

Discussion and motivation of ProxyCLIP. The image-level contrastive learning paradigm adopted by CLIP often encounters challenges in aligning local image patches with corresponding textual representations [27, 63]. This difficulty arises from its global attention mechanism’s indiscriminate propagation of information across all image patches, while effective at capturing the global context, tends to weaken the specificity of local patch representations. Recent studies [2, 22, 40, 63] suggest that modifying the attention mechanism to better organize the spatial information of local patches can benefit the open-vocabulary segmentation. We hypothesize that an improved local patch representation should selectively incorporate contextually relevant patches sharing the same semantics, thus enhancing the precision of patch representations for dense prediction tasks. However, CLIP’s local feature correspondence is not ideal, which cannot ensure spatially coherent segments using CLIP alone. Meanwhile, recent vision foundation models like DINO [4], DINOv2 [11, 29], and SAM [17] have demonstrated the ability to learn semantically coherent representations with strong localization capabilities. This motivates us to improve CLIP with these VFMs.

We begin with a comparative analysis of the semantic coherence within attention scores [2, 22, 40] of CLIP and feature correspondences [14] of VFMs, using data from the COCO stuff [3] dataset. Specifically, we collect all different attention scores at the last attention layer of CLIP, and extract all patch representations of VFMs for each image in COCO Stuff. For CLIP, let $\text{Attn}_{qk} \in \mathbb{R}^{L \times L}$, as well as the self-self attention Attn_{qq} , Attn_{kk} and Attn_{vv} , signify the similarities between image patches. For VFMs, we extract the patch representations $F \in \mathbb{R}^{L_v \times D_v}$ using their frozen backbones, where L_v and D_v denote the number and dimensionality of image patches, respectively. We then compute their cosine similarity $S_{ij} = \frac{F_i \cdot F_j}{|F_i| |F_j|}$. To evaluate the semantic coherence, we use the attention scores of CLIP and similarity scores of VFMs as a binary classifier to predict whether two patches share the same semantic label. We determine the ground true label for each patch through majority voting using the segmentation maps. Given a pair of semantic labels l_i and l_j , we assign a target value of 1 for the classifier if the labels $l_i = l_j$; otherwise, the target value is 0.

Fig. 1 presents the precision-recall curves for different attention scores of CLIP and the feature




Fig. 1: Precision recall curves of different classifiers. Higher average precision (AP) indicates better semantic correspondence.

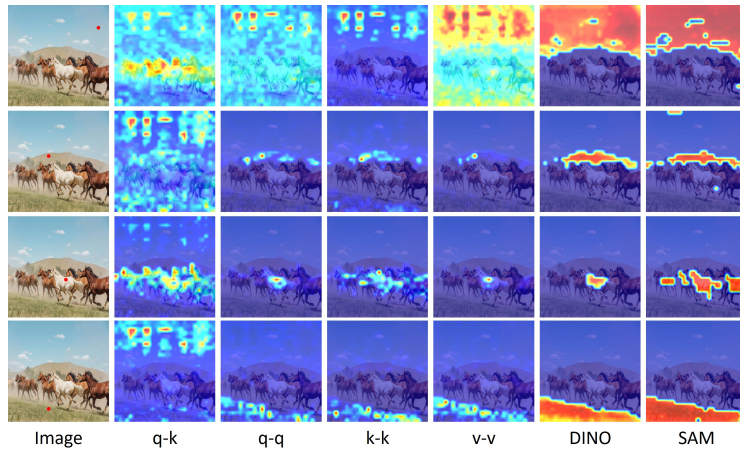


Fig. 2: Attention scores (maps) between CLIP, DINO and SAM using different seeds (in red). For CLIP’s attention maps, we display only the first head of multi-head self-attention maps.

correspondence of VFMs using DINO [4], SAM [17], and MAE [15] backbones. Notably, the self-self attention scores exhibit significantly higher average precision than the vanilla attention scores, reflecting their better semantic coherence, which focuses on contextually relevant information. This finding aligns with earlier studies [2, 22, 40] that reported improved outcomes over the vanilla CLIP. Remarkably, the feature correspondence of VFMs, especially with DINO and SAM backbones, far surpasses the attention scores of CLIP. We also showcase qualitative results in Fig. 2, illustrating that attention maps generated by DINO and SAM effectively highlight semantically relevant patches in different contexts, such as sky, mountain, horses, and grass. This is highly encouraging, as it suggests that the quality of CLIP’s dense patch representations could be enhanced by employing VFMs’ feature representations as a proxy attention for CLIP.

3.2 ProxyCLIP

In this section, we detail our method for improving CLIP’s dense representations by utilizing pre-trained Vision Foundation Models (VFMs). Our objective is to seamlessly combine the advanced spatial coherence of VFMs with CLIP’s semantic understanding capabilities within a training-free framework, thereby enhancing dense vision-language inference.

Architecture. The overall framework, illustrated in Fig. 3, includes two frozen image encoders, *i.e.*, CLIP image encoder and VFM backbone, and a novel proxy attention module (PAM). In a word, we expect CLIP’s image encoder to provide the semantic base (*value* embeddings), while the VFMs’ image encoder contributes attention scores (*query* and *key* embeddings) to reorganize this semantic




Fig. 3: Overview of the ProxyCLIP architecture. ProxyCLIP consists of two frozen image encoders and a novel proxy attention module (PAM). On the right, the flow of the proxy attention mechanism with an adaptive normalization and masking strategy is illustrated, corresponding to Eqs. (6) to (8).

base. The goal is to produce semantically coherent local patch representations that align well with language cues.

Proxy attention module. The Proxy Attention Module (PAM) serves as a pivotal component in our approach. PAM is designed to integrate visual representations $\mathbf{x} \in \mathbb{R}^{L_x \times D_x}$ from the VFM backbones and *value* embeddings $\mathbf{v} \in \mathbb{R}^{n \times L_v \times D_v}$ from CLIP’s last attention layer. Here, L_x and L_v denote the length of the visual sequences, D_x and D_v represent their respective dimensions, and n is the number of heads for \mathbf{v} . The primary objective of PAM is to generate dense representations of images for subsequent vision-language inference tasks. Let ℓ_2 -normalized visual representations \mathbf{x} as the *query* and *key* embeddings, the process of the proxy attention module is formulated as follows:

$$\text{Attn}_p = \text{SoftMax}(\mathbf{x}\mathbf{x}^T), \quad (4)$$

$$\mathbf{z} = \text{Proj}(\text{Attn}_p \cdot \mathbf{v}), \quad (5)$$

where $\text{Attn}_p \in \mathbb{R}^{L_x \times L_x}$ represents the proxy attention score, which is expanded to encompass n heads in a new axis. Proj denotes a projection layer responsible for fusing the multiple self-attention outputs into the final representations, and its parameters are inherited from the output projection layer of CLIP’s last self-attention layer. It’s noteworthy that our PAM differs from the cross-attention mechanism in existing works [17, 24, 39], which typically utilize a pair of keys and values from one encoder while employing the queries from another.




Fig. 4: The statistics of similarity matrix before (left) and after (right) normalization.

Normalization and masking. In practice, we observed that the attention score calculated by Eq. (4) may not always ensure good consistency

and separateness across various VFMs due to their diverse inductive biases in visual representations. For example, as illustrated in Fig. 4, when analyzing a subset of the COCO Stuff dataset, we found that the median cosine similarity among patch representations learned from DINO with ViT-B/16 architecture is 0.22, whereas it is 0.57 for SAM with the same architecture. To this end, we propose an adaptive normalization and masking approach to the similarity matrix:

$$A = \gamma(\mathbf{x}\mathbf{x}^T - \frac{\beta}{L_v^2} \sum_{i,j} [\mathbf{x}\mathbf{x}^T]_{ij}), \quad (6)$$

$$\mathcal{M}_{ij} = \begin{cases} 0, & A_{ij} \geq 0 \\ -\infty, & A_{ij} < 0 \end{cases} \quad (7)$$

$$\text{Attn}_p = \text{SoftMax}(A + \mathcal{M}). \quad (8)$$

Here, β and γ denote the shifting and scaling factors, respectively. We illustrate the distribution of similarity matrices of different VFMs after normalization in Fig. 4, showing their aligned statistics with similar median cosine similarity. Based on the normalized similarity matrix, we introduce the masking matrix \mathcal{M} to suppress negative similarities in A . This ensures that each patch representation positively integrates contextually relevant information, leading to improved discriminativeness in dense patch representations. To ensure that the attention scores concentrate exclusively on contextually relevant patches, which constitute a minor fraction of the total image patches, we empirically set these two parameters to $\beta = 1.2$ and $\gamma = 3$ as default values. Moreover, through experiments on parameter sensitivity, we find that these settings are robust.

Different resolution. Given the flexibility of our framework, ProxyCLIP can utilize various VFMs and CLIP models. We observed that proxy attention with higher spatial resolution enhances the performance. Specifically, we propose to interpolate the spatial resolution of \mathbf{v} to align with the size of \mathbf{x} , resulting in $L_v = L_x$. This modification allows ProxyCLIP to benefit from VFMs with smaller patch sizes, such as the DINO ViT-B/8 architecture. Consequently, ProxyCLIP is able to produce more detailed segmentation maps, as demonstrated in Fig. 7, where the segmentation maps produced by ProxyCLIP with DINO ViT-B/8 as the proxy attention outperform that with DINO ViT-B/16.

Open-vocabulary segmentation. To perform dense inference with CLIP, \mathbf{z} is first projected to the aligned vision-language space, denoted as \mathbf{z}_v . Then, we obtain the textual representations $\mathbf{z}_t \in \mathbb{R}^{C \times d}$ by feeding the text prompt “a photo of a {label names}.” with different label names to CLIP’s text encoder, where C denotes the number of classes, and d represents the dimension of the shared latent vision-language space. Next, we conduct patch-level classification by computing the cosine similarity between \mathbf{z}_v and \mathbf{z}_t , and determine the label for each patch based on the most similar prompt. Finally, we obtain the segmentation maps $\mathcal{S} \in \mathbb{R}^{h \times w \times 1}$ for the image by reshaping and up-sampling.

4 Experiments

4.1 Experiment Setups

Benchmark Datasets. We conduct a comprehensive evaluation on eight well-established datasets following [5, 48]. These datasets are divided into two main categories. (i) With background class: PASCAL VOC [13] (VOC), PASCAL Context [26] (Context) and COCO Object [3] (Object). (ii) Without background class: PASCAL VOC20 [13] (VOC20), PASCAL Context59 [26] (Context59), COCO Stuff [3] (Stuff), Cityscapes [10] (City) and ADE20K [62] (ADE).

Implementation & metric. Our implementations are based on MMSegmentation [9]. For image pre-processing, we resize the images to accommodate varying dataset specifications: a shorter side of 336 pixels for PASCAL and COCO datasets and 448 pixels for Cityscapes and ADE20K datasets. We adopt a sliding window strategy with a 336×336 window and 112×112 stride to benefit both CLIP models and VFMs. For all datasets, we construct textual descriptions using the standard ImageNet prompts [31] along with their corresponding class names. Our method undergoes direct evaluation on the validation set of all datasets without the need of retraining or fine-tuning. We adopt the mean Intersection over Union (mIoU) metric to evaluate the semantic segmentation performance and report the results without any post-processing step.

Comparison Methods. We compare our method against several baselines and state-of-the-arts. The baselines include CLIP [31] with ViT-B/16 and ViT-L/14 architectures, and OpenCLIP [7] with the ViT-H/14 architecture. Additionally, We compare our method with two categories of state-of-the-art approaches.

- **Training-free** methods: MaskCLIP [63], CLIPSurgery [22], GEM [2], ReCo [35] and SCLIP [40];
- **Weakly-supervised training-based** methods: GroupViT [51], SegCLIP [25], ViewCo [33], OVSegmentor [53], CoCu [49], TCL [5], SAM-CLIP [41] and CLIP-DINOiser [48].

For baseline methods and methods based on ViT-L/14 and ViT-H/14 architectures, we report their results based on our implementation. For state-of-the-art methods, if not otherwise specified, all reported results are directly sourced from the respective papers.

4.2 Main results

Quantitative results. We first conduct experiments to evaluate ProxyCLIP against various open-vocabulary semantic segmentation models. The VFM adopted in our ProxyCLIP is DINO with ViT-B/8 architecture for its good performance. Tab. 1 summarises the comparison results of all comparison methods on eight datasets, where our ProxyCLIP consistently achieves the best performance. Our

Table 1: Open-vocabulary semantic segmentation comparison on 8 datasets. [†] Results are directly cited from TCL [5].

Methods	VOC	Context	Object	VOC20	Context59	Stuff	City	ADE	Avg.
Training-based									
SegCLIP [25]	52.6	24.7	26.5	-	-	-	-	-	-
ViewCo [33]	52.4	23.0	23.5	-	-	-	-	-	-
OVSegmentor [53]	53.8	20.4	25.1	-	-	-	-	5.6	-
CoCu [49]	40.9	21.2	20.3	-	-	13.6	15.0	11.1	-
SAM-CLIP [41]	60.6	29.2	-	-	-	-	-	17.1	-
GroupViT [†] [51]	50.4	18.7	27.5	79.7	23.4	15.3	11.1	9.2	27.7
TCL [5]	51.2	24.3	30.4	77.5	30.3	19.6	23.1	14.9	33.1
CLIP-DINOiser [48]	62.2	32.4	35.0	80.2	35.9	24.6	31.7	20.0	40.3
CLIP-ViT-B/16									
CLIPSurgery [22]	-	29.3	-	-	-	21.9	31.4	-	-
GEM [2]	46.2	32.6	-	-	-	-	-	15.7	-
ReCo [†] [35]	25.1	19.9	15.7	57.7	22.3	14.8	21.1	11.2	23.5
CLIP [31]	16.4	8.4	5.6	41.9	9.2	4.4	5.0	2.9	11.7
MaskCLIP [†] [63]	38.8	23.6	20.6	74.9	26.4	16.4	12.6	9.8	27.9
SCLIP [40]	59.1	30.4	30.5	80.4	34.2	22.4	32.2	16.1	38.2
ProxyCLIP	61.3	35.3	37.5	80.3	39.1	26.5	38.1	20.2	42.3
CLIP-ViT-L/14									
CLIP [31]	8.2	4.1	2.7	15.6	4.4	2.4	2.5	1.7	5.2
MaskCLIP [63]	23.3	11.7	7.2	29.4	12.4	8.8	11.5	7.2	13.9
SCLIP [40]	43.5	22.3	25.0	69.1	25.2	17.6	18.6	10.9	29.0
ProxyCLIP	60.6	34.5	39.2	83.2	37.7	25.6	40.1	22.6	43.0
OpenCLIP-ViT-H/14									
OpenCLIP [7]	8.8	5.0	5.3	21.7	5.5	3.2	4.3	2.8	7.1
MaskCLIP [63]	31.4	13.3	16.2	41.7	15.8	8.4	17.7	10.4	19.3
SCLIP [40]	43.8	23.5	24.6	67.5	25.6	16.8	19.5	11.3	29.1
ProxyCLIP	65.0	35.4	38.6	83.3	39.6	26.8	42.0	24.2	44.4

investigation yields interesting findings: 1) Among trained-based methods, CLIP-DINOiser [48] emerges as a standout performer, with the average mIoU (40.3) significantly higher than TCL [5] (33.1). CLIP-DINOiser’s success can be attributed to the improved CLIP features distilled from DINO features, benefiting from the good spatial consistency of DINO. This highlights the importance and benefits of integrating VFMs with CLIP over approaches like self-distillation used in TCL. Meanwhile, the proposed ProxyCLIP outperforms CLIP-DINOiser by a significant margin, with average improvements of **+2.0 mIoU**, **+2.7 mIoU**, **and +4.1 mIoU** under CLIP-ViT-B/16, CLIP-ViT-L/14, and OpenCLIP-ViT-H/14 architectures, respectively. This underscores ProxyCLIP’s superior efficiency in merging the strengths of VFMs and CLIP. 2) Notably, all training-based methods fall short when compared to ProxyCLIP, suggesting potential compromises to CLIP’s original zero-shot transfer properties during fine-tuning.

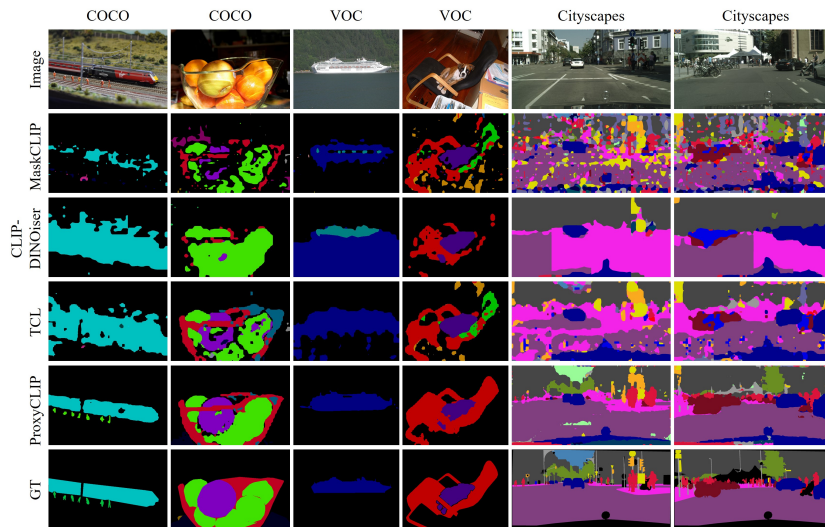


Fig. 5: Qualitative comparison of semantic segmentation results.

In contrast, as a training-free approach, ProxyCLIP seamlessly combines the spatial consistency of VFMs with the visual language alignment strengths of CLIP, without sacrificing their advantages. 3) Among training-free methods, ProxyCLIP significantly outperforms all competitors, achieving an average mIoU improvement of 4.1 with CLIP-ViT-B/16 architectures. This improvement expands to 14.0 mIoU and 15.3 mIoU with CLIP-ViT-L/14 and OpenCLIP-ViT-H/14, respectively. Notably, ProxyCLIP is the first training-free method to effectively utilize larger CLIP models to boost performance. In conclusion, these results clearly affirm ProxyCLIP’s efficacy in open-vocabulary segmentation tasks.

Qualitative results. Figure 5 demonstrates the qualitative comparison results of MaskCLIP [63], CLIP-DINOiser [48], TCL [5] and our ProxyCLIP. For a fair comparison, we do not use any post-refinement technique. We observe that our proposed ProxyCLIP consistently produce higher quality and more precise segmentation maps compared to the other methods. By leveraging the higher resolution and spatially coherent features of VFMs, ProxyCLIP successfully localizes smaller objects, such as the person in the first and the last two columns. Moreover, ProxyCLIP produces image segments that accurately represent the scene layout (*e.g.*, streetscape in the last two columns) and the boundaries of objects (*e.g.*, the train in the first column and the ship in the third column), whereas other approaches tend to yield coarser results.

4.3 Ablation study

Impact of the backbones. We then investigate the impact of various VFMs and CLIP backbones on the performance of ProxyCLIP. In Tab. 2, we observe

Table 2: Comparison of open-vocabulary semantic segmentation performance under different models and architectures. “Arch” denotes the architecture, where “B/16” represents ViT-Base architecture with a patch size of 16.

VFMs	Arch	VOC	Context	Object	VOC20	Context59	Stuff	City	ADE	Avg.
CLIP-ViT-B/16										
MAE	B/16	52.2	30.4	30.8	76.3	33.5	23.1	30.1	17.1	36.7
SAM	B/16	59.3	33.6	35.4	80.4	37.0	25.0	37.0	19.1	40.8
SD	UNet	60.3	34.7	36.5	82.0	38.2	25.9	37.7	19.6	41.8
DINOv2	B/14	58.6	33.8	37.4	83.0	37.2	25.4	33.9	19.7	41.1
DINO	B/16	59.3	34.4	36.2	79.7	38.1	25.7	36.0	19.4	41.1
OpenCLIP-ViT-H/14										
MAE	B/16	54.7	29.8	32.2	80.6	32.9	21.8	34.9	19.4	38.3
SAM	B/16	63.5	34.1	36.7	84.0	37.9	25.0	41.1	22.0	43.1
SD	UNet	62.8	34.6	34.7	84.9	38.5	26.2	41.5	22.9	43.3
DINOv2	B/14	61.5	34.0	37.3	86.1	37.8	26.2	37.8	23.4	43.0
DINO	B/16	62.9	34.7	37.8	83.7	38.7	26.4	38.6	23.5	43.3





Fig. 6: Qualitative results on PASCAL VOC. (3)(4): Different configurations of ProxyCLIP. It is clear that proxy attention from DINO largely improves the spatial consistency while better CLIP backbone improves the semantics.

that the results of ProxyCLIP under the CLIP-ViT-B/16 using DINO [4], DINOv2 [11, 29], and SAM [17] are very close, averaging 41.0 mIoU, as well as under the OpenCLIP-ViT-H/14. This indicates that these VFMs provide a comparable quality of proxy attention. Interestingly, the self-supervised DINO performs as well as the supervised SAM. We attribute this to the low resolution of the input images (336×336), which does not benefit SAM, as it relies on high-resolution inputs. We also observe that, under CLIP-ViT-B/16 model, utilizing SD [34] as VFM yields better performance compared to using VFMs with ViT-B/16 or ViT-B/14 architecture. However, SD features still lag behind DINO features with ViT-B/8 architecture (42.3 mIoU in Tab. 3). On the other hand, we observe an improvement of around 2.0 mIoU for each variant when using OpenCLIP-ViT-H/14 models. This suggests that ProxyCLIP benefits from the enhanced semantic features of CLIP, confirming its effectiveness and generality.

Table 3: Comparison of open-vocabulary semantic segmentation performance under DINO architectures with different patch size.

VFMs Arch	VOC	Context	Object	VOC20	Context59	Stuff	City	ADE	Avg.
CLIP-ViT-B/16									
DINO S/16	59.1	34.2	35.8	79.8	37.9	25.6	35.8	19.2	40.9
DINO S/8	60.0	34.7	36.9	79.6	38.5	26.2	37.7	19.9	41.7
DINO B/16	59.3	34.4	36.2	79.7	38.1	25.7	36.0	19.4	41.1
DINO B/8	61.3	35.3	37.5	80.3	39.1	26.5	38.1	20.2	42.3
OpenCLIP-ViT-H/14									
DINO S/16	62.8	34.5	37.5	83.8	38.4	26.1	38.6	23.1	43.1
DINO S/8	63.5	34.8	38.0	82.2	38.9	26.3	41.7	23.9	43.7
DINO B/16	62.9	34.7	37.8	83.7	38.7	26.4	38.6	23.5	43.3
DINO B/8	65.0	35.4	38.6	83.3	39.6	26.8	42.0	24.2	44.4

**Fig. 7:** Qualitative comparison of different patch size. VFMs with smaller patch size of 8 helps ProxyCLIP to produce sharper boundaries.

We further conduct a qualitative comparison between different configurations of ProxyCLIP and MaskCLIP. To show the role of proxy attention, we further make qualitative comparison between different configurations of ProxyCLIP and MaskCLIP. Fig. 6(2)→(3) clearly shows that incorporating proxy attention from DINO-ViT-B/8 significantly enhances the local consistency of the segmentation maps. Fig. 6(3)→(4) highlights that utilizing a larger OpenCLIP backbone increases classification accuracy.

Impact of the patch size. The resolution of model’s feature maps plays a crucial role in determining the quality of the final segmentation results. Smaller patch sizes in the ViT architecture usually lead to higher-resolution feature maps, which in turn produces more detailed segmentation maps. In Tab. 3, our analysis of ProxyCLIP’s performance with DINO at different patch sizes reveals that DINO with a smaller patch size (8) enables ProxyCLIP to achieve superior outcomes. Specifically, for the same model size, ProxyCLIP’s performance with a patch size of 8 exceeds that of 16 by approximately 1.0 mIoU on average. This finding highlights ProxyCLIP’s advantage when using VFMs with smaller patch

Table 4: Comparison of open-vocabulary semantic segmentation performance using different attention scores.

Attention		VOC	Context	Object	VOC20	Context59	Stuff	City	ADE	Avg.
Norm	Mask									
✗	✗	32.5	18.8	17.0	70.0	22.1	14.0	15.8	10.3	25.0
✓	✗	39.8	23.4	21.4	73.1	27.0	17.5	20.0	13.5	29.5
✓	✓	61.3	35.3	37.5	80.3	39.1	26.5	38.1	20.2	42.3
Attn _{DINO-B/8}		28.4	16.2	18.8	75.7	18.4	12.2	13.0	10.2	24.1
Attn _{DINOv2-B/14}		48.4	26.5	33.7	83.0	30.2	21.4	25.4	17.0	35.7
Attn _{SAM-B/16}		56.3	32.4	37.0	83.0	35.6	24.3	32.3	18.7	40.0

sizes, yielding more precise and higher-quality segmentation maps. As depicted in Fig. 7, ProxyCLIP with DINO-ViT-B/8 produce sharper boundaries, such as the swings of the bird, and the tail of the giraffe.

Impact of the normalization and masking. In Tab. 4, we evaluate the effect of our adaptive normalization and masking strategy on performance by using ProxyCLIP with CLIP-ViT-B/16 and DINO-ViT-B/8 architectures. We observe that the combination of normalization and masking leads to the best results, increasing from an average of 25.0 mIoU to 42.3 mIoU. These results demonstrate that our proposed proxy attention mechanism with an adaptive normalization and masking strategy contributes to the success of incorporating VFMs and CLIP for open-vocabulary semantic segmentation. Moreover, we present findings utilizing the attention scores from the last self-attention layer of each VFM. It is observed that employing the attention scores from DINO-ViT-B/8 and DINOv2-ViT-B/16 results in significantly poorer outcomes, while SAM-ViT-B/16 shows marginally lower performance (40.0 mIoU) compared to our proxy attention formula in Eq. (8) (40.8 mIoU). This suggests that our proposed proxy attention method consistently delivers favorable outcomes across various VFMs.

5 Conclusion

In this paper, we introduce ProxyCLIP, a training-free method using spatial feature correspondence from VFMs as proxy attention to enhance CLIP for open-vocabulary segmentation. Our comparative analysis of semantic coherence within CLIP’s attention scores and VFMs’ feature correspondences demonstrates that VFMs offer much better local correspondence than CLIP’s inherent features, highlighting the benefits of integrating VFMs with CLIP for segmentation tasks. We explore various configurations of VFMs and CLIP and propose two main designs for the proxy attention: an adaptive normalization and masking strategy, and the use of small patch sizes. The experimental results indicate that ProxyCLIP not only significantly outperforms other training-free methods but is also superior to weakly-supervised approaches.

Acknowledgments. This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

1. Barsellotti, L., Amoroso, R., Cornia, M., Baraldi, L., Cucchiara, R.: Training-free open-vocabulary segmentation with offline diffusion-augmented prototype generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3689–3698 (2024)
2. Bousselham, W., Petersen, F., Ferrari, V., Kuehne, H.: Grounding everything: Emerging localization properties in vision-language transformers. arXiv preprint arXiv:2312.00878 (2023)
3. Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1209–1218 (2018)
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021)
5. Cha, J., Mun, J., Roh, B.: Learning to generate text-grounded mask for open-world semantic segmentation from only image-text pairs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11165–11174 (2023)
6. Chen, J., Zhu, D., Qian, G., Ghanem, B., Yan, Z., Zhu, C., Xiao, F., Culatana, S.C., Elhoseiny, M.: Exploring open-vocabulary semantic segmentation from clip vision encoder distillation only. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 699–710 (2023)
7. Cherti, M., Beaumont, R., Wightman, R., Wortsman, M., Ilharco, G., Gordon, C., Schuhmann, C., Schmidt, L., Jitsev, J.: Reproducible scaling laws for contrastive language-image learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2818–2829 (2023)
8. Cho, J., Lei, J., Tan, H., Bansal, M.: Unifying vision-and-language tasks via text generation. In: International Conference on Machine Learning. pp. 1931–1942. PMLR (2021)
9. Contributors, M.: Mmsegmentation: Openmmlab semantic segmentation toolbox and benchmark (2020)
10. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
11. Darcet, T., Oquab, M., Mairal, J., Bojanowski, P.: Vision transformers need registers. arXiv preprint arXiv:2309.16588 (2023)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
13. Everingham, M., Winn, J.: The pascal visual object classes challenge 2012 (voc2012) development kit. Pattern Anal. Stat. Model. Comput. Learn., Tech. Rep **2007**(1-45), 5 (2012)

14. Hamilton, M., Zhang, Z., Hariharan, B., Snavely, N., Freeman, W.T.: Unsupervised semantic segmentation by distilling feature correspondences. arXiv preprint arXiv:2203.08414 (2022)
15. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
16. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International conference on machine learning. pp. 4904–4916. PMLR (2021)
17. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
18. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision* **123**, 32–73 (2017)
19. Lan, M., Wang, X., Ke, Y., Xu, J., Feng, L., Zhang, W.: Smooseg: smoothness prior for unsupervised semantic segmentation. *Advances in Neural Information Processing Systems* **36** (2024)
20. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: *International Conference on Machine Learning*. pp. 12888–12900. PMLR (2022)
21. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems* **34**, 9694–9705 (2021)
22. Li, Y., Wang, H., Duan, Y., Li, X.: Clip surgery for better explainability with enhancement in open-vocabulary tasks. arXiv preprint arXiv:2304.05653 (2023)
23. Li, Y., Li, Z., Zeng, Q., Hou, Q., Cheng, M.M.: Cascade-clip: Cascaded vision-language embeddings alignment for zero-shot semantic segmentation. arXiv preprint arXiv:2406.00670 (2024)
24. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023)
25. Luo, H., Bao, J., Wu, Y., He, X., Li, T.: Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. In: *International Conference on Machine Learning*. pp. 23033–23044. PMLR (2023)
26. Mottaghi, R., Chen, X., Liu, X., Cho, N.G., Lee, S.W., Fidler, S., Urtasun, R., Yuille, A.: 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*. pp. 891–898 (2014)
27. Mukhoti, J., Lin, T.Y., Poursaeed, O., Wang, R., Shah, A., Torr, P.H., Lim, S.N.: Open vocabulary semantic segmentation with patch aligned contrastive learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19413–19423 (2023)
28. Naeem, M.F., Xian, Y., Zhai, X., Hoyer, L., Van Gool, L., Tombari, F.: Silc: Improving vision language pretraining with self-distillation. arXiv preprint arXiv:2310.13355 (2023)
29. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)

30. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: Proceedings of the IEEE international conference on computer vision. pp. 2641–2649 (2015)
31. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
32. Rasheed, H., Maaz, M., Shaji, S., Shaker, A., Khan, S., Cholakkal, H., Anwer, R.M., Xing, E., Yang, M.H., Khan, F.S.: Glamm: Pixel grounding large multimodal model. arXiv preprint arXiv:2311.03356 (2023)
33. Ren, P., Li, C., Xu, H., Zhu, Y., Wang, G., Liu, J., Chang, X., Liang, X.: Viewco: Discovering text-supervised segmentation masks via multi-view semantic consistency. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=2XLRBjY4606>
34. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
35. Shin, G., Xie, W., Albanie, S.: Reco: Retrieve and co-segment for zero-shot transfer. *Advances in Neural Information Processing Systems* **35**, 33754–33767 (2022)
36. Siméoni, O., Puy, G., Vo, H.V., Roburin, S., Gidaris, S., Bursuc, A., Pérez, P., Marlet, R., Ponce, J.: Localizing objects with self-supervised transformers and no labels. arXiv preprint arXiv:2109.14279 (2021)
37. Siméoni, O., Sekkat, C., Puy, G., Vobecký, A., Zablocki, É., Pérez, P.: Unsupervised object localization: Observing the background to discover objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3176–3186 (2023)
38. Sun, S., Li, R., Torr, P., Gu, X., Li, S.: Clip as rnn: Segment countless visual concepts without training endeavor. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13171–13182 (2024)
39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
40. Wang, F., Mei, J., Yuille, A.: Sclip: Rethinking self-attention for dense vision-language inference. arXiv preprint arXiv:2312.01597 (2023)
41. Wang, H., Vasu, P.K.A., Faghri, F., Vemulapalli, R., Farajtabar, M., Mehta, S., Rastegari, M., Tuzel, O., Pouransari, H.: SAM-CLIP: Merging vision foundation models towards semantic and spatial understanding. In: UniReps: the First Workshop on Unifying Representations in Neural Models (2023), <https://openreview.net/forum?id=md0mU6rN2u>
42. Wang, X., Yu, Z., De Mello, S., Kautz, J., Anandkumar, A., Shen, C., Alvarez, J.M.: Freesolo: Learning to segment objects without annotations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14176–14186 (2022)
43. Wang, X., Girdhar, R., Yu, S.X., Misra, I.: Cut and learn for unsupervised object detection and instance segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3124–3134 (2023)
44. Wang, Y., Shen, X., Yuan, Y., Du, Y., Li, M., Hu, S.X., Crowley, J.L., Vaufreydaz, D.: Tokencut: Segmenting objects in images and videos with self-supervised transformer and normalized cut. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)

45. Wang, Y., Sun, R., Luo, N., Pan, Y., Zhang, T.: Image-to-image matching via foundation models: A new perspective for open-vocabulary semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3952–3963 (2024)
46. Wu, S., Zhang, W., Xu, L., Jin, S., Li, X., Liu, W., Loy, C.C.: Clipself: Vision transformer distills itself for open-vocabulary dense prediction. arXiv preprint arXiv:2310.01403 (2023)
47. Wysoczańska, M., Ramamonjisoa, M., Trzciński, T., Siméoni, O.: Clip-diy: Clip dense inference yields open-vocabulary semantic segmentation for-free. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1403–1413 (2024)
48. Wysoczańska, M., Siméoni, O., Ramamonjisoa, M., Bursuc, A., Trzciński, T., Pérez, P.: Clip-dinoiser: Teaching clip a few dino tricks. arXiv preprint arXiv:2312.12359 (2023)
49. Xing, Y., Kang, J., Xiao, A., Nie, J., Shao, L., Lu, S.: Rewrite caption semantics: Bridging semantic gaps for language-supervised semantic segmentation. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=9iafshF7s3>
50. Xu, H., Xie, S., Tan, X.E., Huang, P.Y., Howes, R., Sharma, V., Li, S.W., Ghosh, G., Zettlemoyer, L., Feichtenhofer, C.: Demystifying clip data. arXiv preprint arXiv:2309.16671 (2023)
51. Xu, J., De Mello, S., Liu, S., Byeon, W., Breuel, T., Kautz, J., Wang, X.: Groupvit: Semantic segmentation emerges from text supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18134–18144 (2022)
52. Xu, J., Liu, S., Vahdat, A., Byeon, W., Wang, X., De Mello, S.: Open-vocabulary panoptic segmentation with text-to-image diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2955–2966 (2023)
53. Xu, J., Hou, J., Zhang, Y., Feng, R., Wang, Y., Qiao, Y., Xie, W.: Learning open-vocabulary semantic segmentation models from natural language supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2935–2944 (2023)
54. Yao, L., Huang, R., Hou, L., Lu, G., Niu, M., Xu, H., Liang, X., Li, Z., Jiang, X., Xu, C.: Filip: Fine-grained interactive language-image pre-training. arXiv preprint arXiv:2111.07783 (2021)
55. Yin, Z., Wang, P., Wang, F., Xu, X., Zhang, H., Li, H., Jin, R.: Transfgu: a top-down approach to fine-grained unsupervised semantic segmentation. In: European conference on computer vision. pp. 73–89. Springer (2022)
56. Yuan, H., Li, X., Zhou, C., Li, Y., Chen, K., Loy, C.C.: Open-vocabulary sam: Segment and recognize twenty-thousand classes interactively. arXiv preprint arXiv:2401.02955 (2024)
57. Yuan, L., Chen, D., Chen, Y.L., Codella, N., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., et al.: Florence: A new foundation model for computer vision. arXiv preprint arXiv:2111.11432 (2021)
58. Zhai, X., Wang, X., Mustafa, B., Steiner, A., Keysers, D., Kolesnikov, A., Beyer, L.: Lit: Zero-shot transfer with locked-image text tuning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18123–18133 (2022)

59. Zhang, F., Zhou, T., Li, B., He, H., Ma, C., Zhang, T., Yao, J., Zhang, Y., Wang, Y.: Uncovering prototypical knowledge for weakly open-vocabulary semantic segmentation. arXiv preprint arXiv:2310.19001 (2023)
60. Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: Glipv2: Unifying localization and vision-language understanding. *Advances in Neural Information Processing Systems* **35**, 36067–36080 (2022)
61. Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., et al.: Regionclip: Region-based language-image pretraining. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16793–16803 (2022)
62. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision* **127**, 302–321 (2019)
63. Zhou, C., Loy, C.C., Dai, B.: Extract free dense labels from clip. In: *European Conference on Computer Vision*. pp. 696–712. Springer (2022)
64. Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T.: ibot: Image bert pre-training with online tokenizer. arXiv preprint arXiv:2111.07832 (2021)