CorrCLIP: Reconstructing Correlations in CLIP with Off-the-Shelf Foundation Models for Open-Vocabulary Semantic Segmentation

Dengke Zhang   Fagui Liu[1,2*]   Quan Tang[2*]

[1]South China University of Technology   [2]Pengcheng Laboratory

csdk@mail.scut.edu.cn, fgliu@scut.edu.cn, tangq@pcl.ac.cn

Abstract

Open-vocabulary semantic segmentation aims to assign semantic labels to each pixel without relying on a predefined set of categories. Contrastive Language-Image Pretraining (CLIP) demonstrates outstanding zero-shot classification capabilities but struggles with the pixel-wise segmentation task as the captured inter-patch correlations correspond to no specific visual concepts. Despite previous CLIP-based works improving inter-patch correlations by self-self attention, they still face the inherent limitation that image patches tend to have high similarity to outlier ones. In this work, we introduce CorrCLIP, a training-free approach for open-vocabulary semantic segmentation, which reconstructs significantly coherent inter-patch correlations utilizing foundation models. Specifically, it employs the Segment Anything Model (SAM) to define the scope of patch interactions, ensuring that patches interact only with semantically similar ones. Furthermore, CorrCLIP obtains an understanding of an image's semantic layout via self-supervised models to determine concrete similarity values between image patches, which addresses the similarity irregularity problem caused by the aforementioned restricted patch interaction regime. Finally, CorrCLIP reuses the region masks produced by SAM to update the segmentation map. As a training-free method, CorrCLIP achieves a notable improvement across eight challenging benchmarks regarding the averaged mean Intersection over Union, boosting it from 44.4% to 51.0%.

(a) Attention maps of selected patches (red-star symbol).

(b) Segmentation performance of CLIP.

Figure 1. Qualitative (a) and quantitative (b) improvements achieved by our proposed designs on COCO dataset with ViT-L. **Vanilla CLIP** denotes remaining query-key and **Baseline** denotes transforming query-key into query-query for the final layer's attention map of ViT in CLIP. **SR*** denotes scope reconstruction using labels in our preliminary experiment and **SR** denotes our scope reconstruction using SAM. **VR** denotes our value reconstruction.

1. Introduction

Open-vocabulary semantic segmentation (OVSS) [4, 46, 59] aims to partition an image into multiple segments and assign corresponding categories to each segment based on textual descriptions. This challenge requires models to demonstrate strong generalization abilities and effectively align visual representations with textual descriptions. Contrastive Language-Image Pre-training (CLIP) [35] models, trained on large-scale image-text pair datasets, have shown remarkable zero-shot classification capabilities, providing a viable solution for OVSS.

However, the goal of image-text alignment forces CLIP to focus on the image's overall visual representations, preventing the establishment of coherent inter-patch correla-

---

*Corresponding authors

1

tions required for localizing specific visual concepts. To this end, recent research modifies the computation of the attention map in the last layer of the Vision Transformer (ViT) [13] used in CLIP's image encoder, transforming vanilla query-key into query-query [23], key-key, value-value [27], or their combinations [3, 42]. Although these methods improve inter-patch correlations in CLIP to some extent, they still encounter an inherent limitation: *patches often exhibit high similarity to semantically unrelated patches*, known as outlier patches [11]. As illustrated in Fig. 1a (second column), the selected patches in CLIP display high similarity to certain semantically unrelated patches, leading to confusing inter-patch correlations.

To address the aforementioned challenge, we propose scope reconstruction to compel patches to interact only with semantically similar ones. To validate the effectiveness of this method, we conduct a preliminary experiment with image labels. Specifically, we set the similarity between all patches to be equal but restrict the patches to interact only with those belonging to the same semantic category. This adjustment leads to a significant improvement in segmentation performance, as demonstrated in Fig. 1b. Thus, we provide our key insight: *confining the scope of patch interactions to semantically similar regions can effectively improve inter-patch correlations and substantially enhance the segmentation capabilities of CLIP*.

It is challenging to accurately evaluate the semantic similarity between patches when relying solely on CLIP. To this end, we define semantically similar regions according to the impressive zero-shot class-agnostic segmentation capabilities of the Segment Anything Model (SAM) [22]. Specifically, We use SAM to divide an image into regions and limit patch interactions to occur only within these regions. However, it faces two challenges. First, the segmented regions tend to be fragmented. To tackle this issue, we apply a clustering algorithm to merge regions. Second, some regions remain unsegmented due to thresholding. We define patches in unsegmented regions with similarity greater than the threshold as having similar semantics. When combined with the above two strategies, scope reconstruction significantly improves the segmentation capabilities of CLIP as illustrated in Fig. 1b.

As indicated in Fig. 1a (third column), the similarity between patches in the region after scope reconstruction may be irregular because patches have high similarity to some patches but low similarity to one another. We further propose value reconstruction to address the above similarity irregularity problem and construct more coherent inter-patch correlations. Specifically, we use the semantic layout information included in the self-supervised model DINO [6, 11, 40] to determine similarity values. Inspired by previous work [38, 42], which combines query and key in CLIP, we employ the sum of DINO's query and key to acquire

more comprehensive similarity values. Figure 1 demonstrates the improvements in CLIP resulting from our proposed scope reconstruction and value reconstruction.

Besides, we propose two simple yet effective designs to correct the defects during CLIP's segmentation. We reuse region masks to correct spatial inconsistency and combine the category name with its plural form to correct concept ambiguity. By incorporating the four proposed designs, we present a training-free method called CorrCLIP that effectively reconstructs coherent inter-patch correlations and substantially enhances segmentation capacities of CLIP. Extensive experiments have confirmed the effectiveness of CorrCLIP, which surpasses state-of-the-art methods by a notable margin of 6.6% in averaged mIoU across eight benchmarks. CorrCLIP will continue to improve as the CLIP series, the SAM series, and self-supervised models develop, requiring only model replacements.

Our contributions are summarized as follows: (1) We reveal that limiting the scope of patch interactions to semantically similar regions can effectively improve inter-patch correlations and substantially enhance the segmentation capabilities of CLIP. (2) We propose scope reconstruction and value reconstruction to reconstruct coherent inter-patch correlations of CLIP. (3) We present segmentation map correction and class name correction to ensure spatial consistency and enhance concept understanding. (4) We introduce CorrCLIP for OVSS. Extensive experiments show that CorrCLIP outperforms state-of-the-art methods.

2. Related Work

Open-vocabulary semantic segmentation. The objective of OVSS is to segment an image based on categories described by texts. Recent advancements in OVSS are largely due to the development of large-scale vision-language models (VLMs) [8, 35, 51]. OVSS methods can be divided into two categories: training-based and training-free. Training-based methods rely on mask annotations [9, 28, 29, 43, 47, 52, 53], images [39, 45], or texts [7, 31, 44, 48, 57]. While training-based methods are typically more effective on specific datasets, they carry the potential risk of reducing the open-vocabulary capacity inherited from VLMs, as described in CaR [41]. In contrast, training-free methods [1, 24, 38, 41] do not need any training and fully use the open-vocabulary capabilities offered by VLMs. To enable CLIP with patch-text matching capability, some training-free OVSS methods [23, 27, 42, 62] modify the calculation of attention map in the last layer of ViT used in CLIP's image encoder, which does not consider the inherent limitation. In this paper, we use SAM to overcome this limitation.
Vision-language pre-training. Vision-language pre-training aims to enable models to learn cross-modal information through weakly supervised training on image-text pairs. This pre-training process allows models to under-

stand the associations between images and corresponding texts. The performance of early research [25, 26, 30] is constrained by the limited size of datasets. Recent studies [20, 35, 54], which leverage large-scale web data, have resulted in the development of more robust representations. Among these, CLIP [35] stands out as the most popular vision-language model. It employs contrastive learning to align images with corresponding captions, achieving impressive generalization capabilities on unseen data. Subsequent research [8, 16, 17, 32, 50, 51, 55] has further enhanced CLIP by optimizing training data and processes.

Vision foundation models. Vision foundation models (VFMs) undergo pre-training on large-scale datasets to capture general feature representations of the visual world. VFMs learn rich underlying visual patterns, which can be fine-tuned or directly applied to various visual tasks. One type of VFMs is self-supervised models [6, 18, 34, 64], which aim to learn general-purpose visual features solely from images. Among these, the DINO [6] algorithm has been shown to produce models that capture explicit details about the semantic layout of an image [11, 40]. Another type of VFMs is SAM [22] series, designed for general image segmentation. SAM has demonstrated impressive zero-shot, class-agnostic segmentation capabilities. Recent advancements have focused on improving the quality of the generated masks [19, 21] and enhancing efficiency to enable broader applications in real-world and mobile scenarios [36, 49, 56, 58, 60, 63]. By utilizing SAM's exceptional segmentation capabilities, we can precisely define the scope of patch interactions within CLIP.

3. Method

In this section, we first introduce the overall process of adapting CLIP to OVSS in Sec. 3.1. Then, we reconstruct the scope of patch interactions in Sec. 3.2 and the similarity values in Sec. 3.3. Finally, we reuse the region masks to ensure spatial consistency and use the class name's plural form to enhance concept understanding in Sec. 3.4. The overall framework is shown in Fig. 2.

3.1. Preliminary

An image is first divided into patches, which are then transformed into token embeddings using a linear layer. These tokens are flattened to form the token sequence represented as $X_C \in \mathbb{R}^{N \times d}$, where N is the sequence length and d is the dimension of tokens. Here, we exclude the class token used in the origin CLIP because it is not used in segmentation tasks. Next, positional encoding is added to provide position information. The token sequence is then input into the ViT of CLIP's image encoder. In each preceding layer before the final layer, the token sequence sequentially passes through a multi-head attention network and a feed-forward

network, with residual connections applied after each sub-layer.

In the last layer, the token sequence X_C is mapped to query, key, and value embeddings, denoted as Q_C, K_C, and $V_C \in \mathbb{R}^{N \times d}$, respectively. For easy description, we only show single-head attention. Next, the similarity matrix $S \in \mathbb{R}^{N \times N}$ is calculated. S represents the inter-patch correlations and its computation is the key to improving CLIP's localization ability. Previous studies [23, 27, 42] calculate S by the inner product of F_S, where F_S can be either Q_C, K_C, V_C, or their conbinations. They do not consider the inherent limitation of CLIP. In the following subsections, we will introduce how to eliminate this limitation and reconstruct coherent inter-patch correlations.

The attention map is derived from the similarity matrix by applying the softmax operation. Then, the attention map is applied to V_C to capture the correlations between different tokens in the input sequence. To reduce the detrimental impact of noise on the segmentation results as illustrated in previous work [23, 27], we remove the residual connections and the feed-forward neural network. Subsequently, the sequence of tokens is projected to image features $F_{img} \in \mathbb{R}^{N \times d}$:

$$Attn = \text{Softmax}(\frac{S}{\sqrt{d}}) \tag{1}$$

$$F_{img} = \text{Proj}(AttnV_C) \tag{2}$$

K class names are combined with standard ImageNet prompts [35] to construct textual descriptions, which are passed through CLIP's text encoder to obtain class embeddings $F_{text} \in \mathbb{R}^{K \times d}$. Image features F_{img} are projected to align with class embeddings. Finally, class embeddings are used as the parameters of the classifier to get the segmentation map:

$$pred = \arg \max_K (\text{Proj}(F_{img})F_{text}^{\text{T}}) \tag{3}$$

3.2. Scope Reconstruction

Applying CLIP to segmentation tasks as above encounters the limitation that patches tend to exhibit high similarity to outlier patches, which is detrimental to establishing coherent inter-patch correlations. To address this limitation, we reconstruct the scope of patch interactions, ensuring that patches only interact with semantically similar ones. SAM provides a feasible implementation based on its remarkable zero-shot segmentation capabilities.

SAM generates region masks for the entire image and then applies thresholding to these masks to discard those with lower confidence and stability. Consequently, we obtain Z non-overlapping region masks $M = \{m_1, m_2, ..., m_Z\} \in \mathbb{R}^{Z \times N}$, where $m_i \in \mathbb{R}^N$ represents
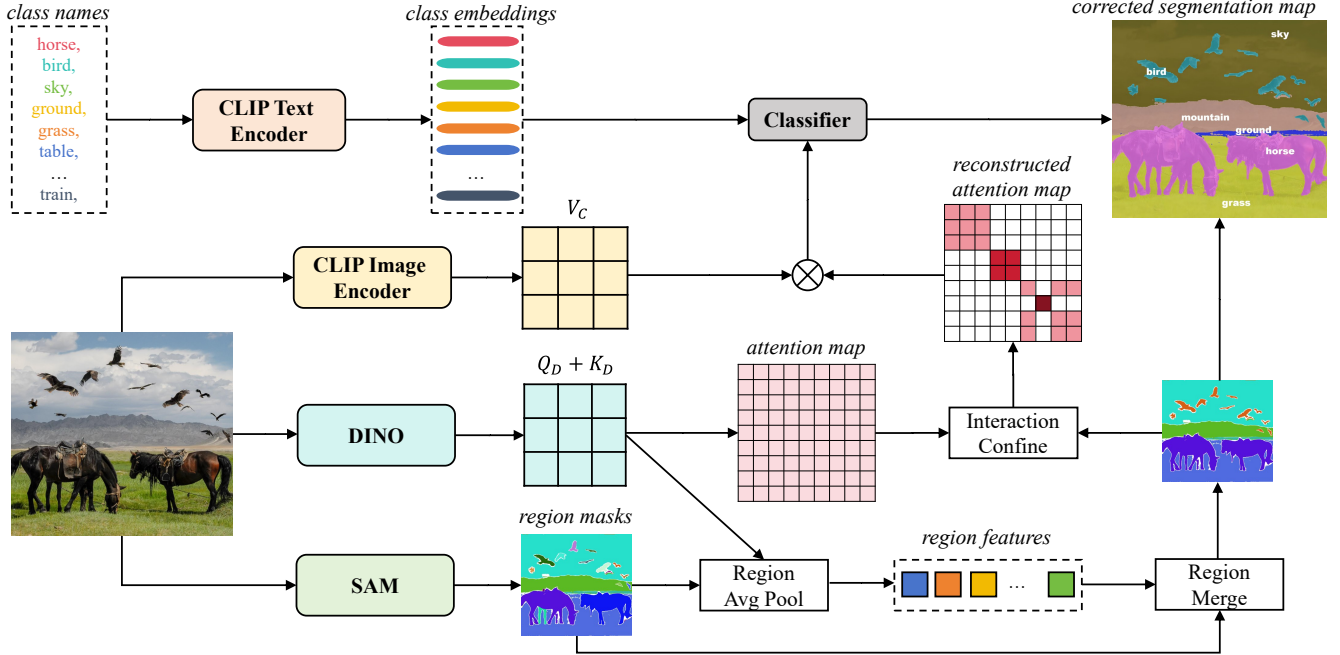
Figure 2. The overall framework of our method. An image passes through the CLIP image encoder to obtain the last layer's value V_C and the DINO to get the combination of the last layer's query and key $Q_D + K_D$. SAM segments this image to get region masks, which are merged according to the clustering algorithm using region features generated by region average pooling. Then, $Q_D + K_D$ is used to calculate the attention map, and the merged region masks are used to confine the scope of patch interactions and get the reconstructed attention map. The final attention map is applied to V_C to get image features F_{img}, passed through the classifier whose parameters are class embeddings to get the segmentation map. Finally, the region masks are reused to correct the segmentation map.

the ith region mask. We denote all unsegmented region masks as m_0. Additionally, the generated regions represent separate instances, and different instances that have similar semantics can be merged. We acquire the features of all segmented regions $F_{region} = \{f_1, f_2, ..., f_Z\} \in \mathbb{R}^{Z \times d}$ by performing mask average pooling on F_S. Subsequently, we use a clustering algorithm to merge semantically similar regions based on these region features:

$$f_i = \text{Mean}(m_i \odot F_S) \quad (4)$$
$$\hat{M} = \text{Cluster}(M, F_{region}) \quad (5)$$

where $\hat{M} = \{\hat{m}_1, \hat{m}_2, ..., \hat{m}_z\} \in \mathbb{R}^{z \times N}$ are merged region masks and z is the number of merged region masks. $\odot$ represents the element-wise product.

We then use these merged region masks to calculate the semantic matrix $E \in \mathbb{R}^{N \times N}$ where $E_{i,j} = 1$ denotes that ith patch and jth patch have similar semantics. Notably, patches belonging to m_0 cannot be regarded as semantically similar. However, they cannot be completely ignored due to their potentially large area. To tackle this issue, we define patches with a similarity greater than the average of the sim-

ilarity matrix as semantically similar:

$$E = (m_0 + m_0^{\mathrm{T}}) \odot (S > \text{Mean}(S)) + \sum_{i=1}^{z} \hat{m}_i \odot \hat{m}_i^{\mathrm{T}} \quad (6)$$

Notice that Eq. (6) employs the broadcasting mechanism. Finally, we confine the scope of patch interactions by ignoring the attention scores between patches that do not have similar semantics:

$$A_{i,j} = \begin{cases} 0, & E_{i,j} = 1 \\ -\infty, & E_{i,j} = 0 \end{cases} \quad (7)$$

$$Attn = \text{Softmax}(\frac{S + A}{\sqrt{d}}) \quad (8)$$

3.3. Value Reconstruction

Confining the scope of patch interactions leads to the similarity irregularity problem because patches display high similarity to some patches while showing low similarity to one another. To establish more coherent correlations between patches within regions, we utilize DINO's ability to recognize the semantic layout information to calculate the similarity matrix S. Inspired by previous work [27, 38, 42], which combines query and key in CLIP, we similarly em-

4

ploy this combination in DINO to obtain a more comprehensive similarity matrix.

Similar to CLIP, the token sequence in DINO's final layer is mapped to Q_D, K_D, and $V_D \in \mathbb{R}^{N \times d}$, respectively. In practice, the token sequence lengths of CLIP and DINO may be different due to the difference in patch size, which can be solved by interpolation. Then similarity matrix S is calculated by the combination of Q_D and K_D:

$$S = \frac{F_S F_S^{\mathrm{T}}}{\|F_S\|^2} = \frac{(Q_D + K_D)(Q_D + K_D)^{\mathrm{T}}}{\|Q_D + K_D\|^2} \quad (9)$$

$$Attn = \mathrm{Softmax}(\frac{S + A}{\tau}) \quad (10)$$

where $\tau < 1$ is the temperature coefficient used to sharpen attention scores so that patches have higher attention scores when patches are semantically similar. Notice that using DINO's features to calculate attention scores has an inevitable defect: multi-head attention needs to be replaced with single-head attention. The reason is that the visual concepts learned by each attention head in DINO are different from those learned in CLIP.

3.4. Segment Correction

In this subsection, we propose two strategies for correcting the defects in CLIP during image segmentation.

Segmentation map correction. Spatial consistency is essential in segmentation tasks, which refers to maintaining the continuity in space. Simply put, it means ensuring that the predictions of adjacent pixels are logically consistent, avoiding abrupt changes such as not having predictions of other objects within a coherent object. Fully supervised methods with mask annotations can learn this property, but CLIP lacks this property due to its weakly supervision. Although our method has somewhat improved this situation, it is still inferior to fully supervised training. We reuse the region masks generated above to ensure spatial consistency and correct the segmentation map. Specifically, we change the categories of all patches within a region to the category that is most common within this region:

$$pred[\hat{m}_i] = \mathrm{Mode}(pred[\hat{m}_i]), \quad i > 0 \quad (11)$$

Class name correction. The "background" class is too broad for CLIP to classify this concept accurately; therefore, the common practice [24, 38, 41, 42] is to replace "background" with more specific classes, such as "sky", "cloud", "wall", etc. Drawing inspiration from this approach, we propose a simple yet effective strategy to address the issue of conceptual ambiguity in CLIP. For instance, CLIP often confuses "hill" with "person", incorrectly categorizing images of a person as the hill. In this study, we combine the class "person" with its plural form "people", slightly improving classification accuracy. This

suggests that combining a class name with its plural form can provide CLIP with a more comprehensive understanding of the category concept. More details are provided in the supplementary.

4. Experiments

4.1. Dataset and Evaluation Metric

We evaluate our method on the validation splits of five datasets following [23, 24, 42, 45]. Specifically, **Pascal VOC** [15] comprises 1,449 images and is utilized for two benchmark configurations: VOC21 with 21 classes (including the background class) and VOC20 with 20 classes (excluding the background class). **Pascal Context** [33] contains 5,104 images and is also used for two benchmark configurations similar to Pascal VOC: PC60 with 60 classes and PC59 with 59 classes. **COCO Stuff** [5] (Stuff) includes 5,000 images divided into 171 classes, which encompass both stuff and object categories, excluding the background class. COCO Object (Object) is a derivative of COCO Stuff, which merges all stuff classes into the background class and has 81 classes. **ADE20k** [61] (ADE) has 2,000 images and 150 classes without the background class. **Cityscapes** [10] (City) has 500 images and 19 classes without the background class. The above five datasets generate a total of eight benchmarks. We use the mean Intersection over Union (mIoU) as the metric to compare our method with previous state-of-the-art methods.

4.2. Implementation Details

Our model has three different sizes based on CLIP: ViT-B/16 [51], ViT-L/14 [51], and ViT-H/14 [8]. All other configurations for the three variants are identical.

The backbone of DINO is ViT-B/8 [6]. We use SAM2 [36] with MAE [18] pre-trained Hiera-L [2, 37]. We collect 32×32 prompt points in a grid manner to generate region masks of an image. Mask thresholding has two parameters: "pred_iou_thresh" and "stability_score_thresh". We set both parameters to 0.7 for all datasets. The clustering algorithm used in region merging is DBSCAN [14]. DBSCAN has two parameters: the neighborhood radius (eps) and the minimum number of samples (samples) in the radius. We set "eps" to 0.2 and "samples" to 1 for all datasets. The temperature coefficient is set to 0.25. More details about implementation are provided in the supplementary.

We resize the images to meet the varying specifications of different datasets: a shorter side of 336 pixels for Pascal VOC, Pascal Context, and COCO, and 448 pixels for Cityscapes and ADE20K. We perform slide inference with a 336×336 window and 112×112 stride.

Method	Size	VOC21	VOC20	PC60	PC59	Object	Stuff	ADE	City	Avg
Training-based										
TCL[7]		55.0	83.2	30.4	33.9	31.6	22.4	17.1	24.0	37.2
CLIP-DINOiser[45]	ViT-B/16	62.1	80.9	32.4	35.9	34.8	24.6	20.0	31.7	40.3
CoDe[44]		57.7	-	30.5	-	32.3	23.9	17.7	28.9	-
Training-free										
MaskCLIP[62]		38.8	74.9	23.6	26.4	20.6	16.4	9.8	12.6	27.9
ClearCLIP[23]		51.8	80.9	32.6	35.9	33.0	23.9	16.7	30.0	38.1
SCLIP[42]		59.1	80.4	30.4	34.2	30.5	22.4	16.1	32.2	38.2
ProxyCLIP[24]	ViT-B/16	<u>61.3</u>	80.3	<u>35.3</u>	<u>39.1</u>	37.5	<u>26.5</u>	<u>20.2</u>	<u>38.1</u>	<u>42.3</u>
CLIPtrase[38]		53.0	<u>81.2</u>	30.8	34.9	**44.8**	24.1	17.0	-	-
CorrCLIP (Ours)		**72.5**	**88.7**	**42.0**	**46.2**	<u>43.7</u>	**30.6**	**25.3**	**48.3**	**49.7**
Training-free										
SCLIP[42]		43.5	69.1	22.3	25.2	25.0	17.6	10.9	18.6	29.0
ProxyCLIP[24]		60.6	83.2	34.5	37.7	<u>39.2</u>	25.6	22.6	<u>40.1</u>	43.0
FreeDA[1]	ViT-L/14	55.4	87.9	<u>38.3</u>	<u>43.5</u>	37.4	<u>28.8</u>	<u>23.2</u>	36.7	<u>43.9</u>
CaR[41]		<u>67.6</u>	**91.4**	30.5	39.5	36.6	-	17.7	-	-
CorrCLIP (Ours)		**73.2**	<u>90.6</u>	**41.0**	**47.5**	**46.0**	**32.0**	**29.1**	**49.0**	**51.0**
Training-free										
SCLIP[42]		43.8	67.5	23.5	25.6	24.6	16.8	11.3	19.5	29.1
ProxyCLIP[24]	ViT-H/14	<u>65.0</u>	<u>83.3</u>	<u>35.4</u>	<u>39.6</u>	<u>38.6</u>	<u>26.8</u>	<u>24.2</u>	<u>42.0</u>	<u>44.4</u>
CorrCLIP (Ours)		**74.1**	**91.6**	**40.3**	**45.5**	**43.6**	**30.6**	**27.0**	**47.7**	**50.0**

Table 1. Comparison with state-of-the-art OVSS methods on eight benchmarks in three different sizes of CLIP. Bold fonts indicate the optimal methods and underlined fonts indicate the suboptimal methods. "Avg" represents the averaged mIoU across eight benchmarks.

4.3. Comparison with State-of-the-Art Methods

We compare CorrCLIP with state-of-the-art training-free OVSS methods across eight benchmarks. Specifically, these methods include MaskCLIP [62], SCLIP [42], ClearCLIP [23], and CLIPtrase[38], which improve inter-patch correlations only using CLIP. Besides, CaR [41] progressively filters out irrelevant texts and enhances the mask quality of CLIP. ProxyCLIP [24] uses various VFMs' features as a form of proxy attention to augment CLIP. FreeDA [1] leverages diffusion models to visually localize generated concepts and local-global similarities to match class-agnostic regions with semantic classes.

The results of the comparison with these methods are summarized in Tab. 1. All three variants of CorrCLIP demonstrate superior performance across eight benchmarks, significantly outperforming other methods. In particular, our CorrCLIP achieves an increase of 7.4%, 7.1%, and 5.6% in averaged mIoU across three variants. We also compare our approach with several state-of-the-art training-based methods that are supervised by images or texts. These weakly supervised methods perform worse than some training-free methods, such as ProxyCLIP, FreeDA, and our CorrCLIP. We attribute it to CLIP's diminished open-vocabulary capability stemming from training on specific datasets.

The results presented above demonstrate the effectiveness of our method in reconstructing inter-patch correlations and enhancing segmentation capabilities of CLIP. In Fig. 3, we provide qualitative comparisons between our method and ProxyCLIP, which both utilize CLIP with ViT-H/14 [8]. Our method accurately identifies object categories and maintains spatial consistency.

4.4. Ablation Study

We perform a series of ablation studies to investigate the effects of each component in our method. Unless otherwise specified, we use CLIP with ViT-B/16 [35] and abbreviate VOC21 as VOC and PC60 as PC.

Impact of integrating different components. We investigate the impact of four components of CorrCLIP outlined in Tab. 2. These components are scope reconstruction (SR), value reconstruction (VR), segmentation map correction (MC), and class name correction (NC). We use the method that substitutes query-key with query-query as the baseline. Limiting the scope of patch interactions significantly enhances CLIP's segmentation performance, yielding an improvement of 6.6%, 3.8%, 4.4%, and 2.4% mIoU on VOC, PC, Object, and ADE benchmarks, respectively. When integrating value reconstruction, we observe a further
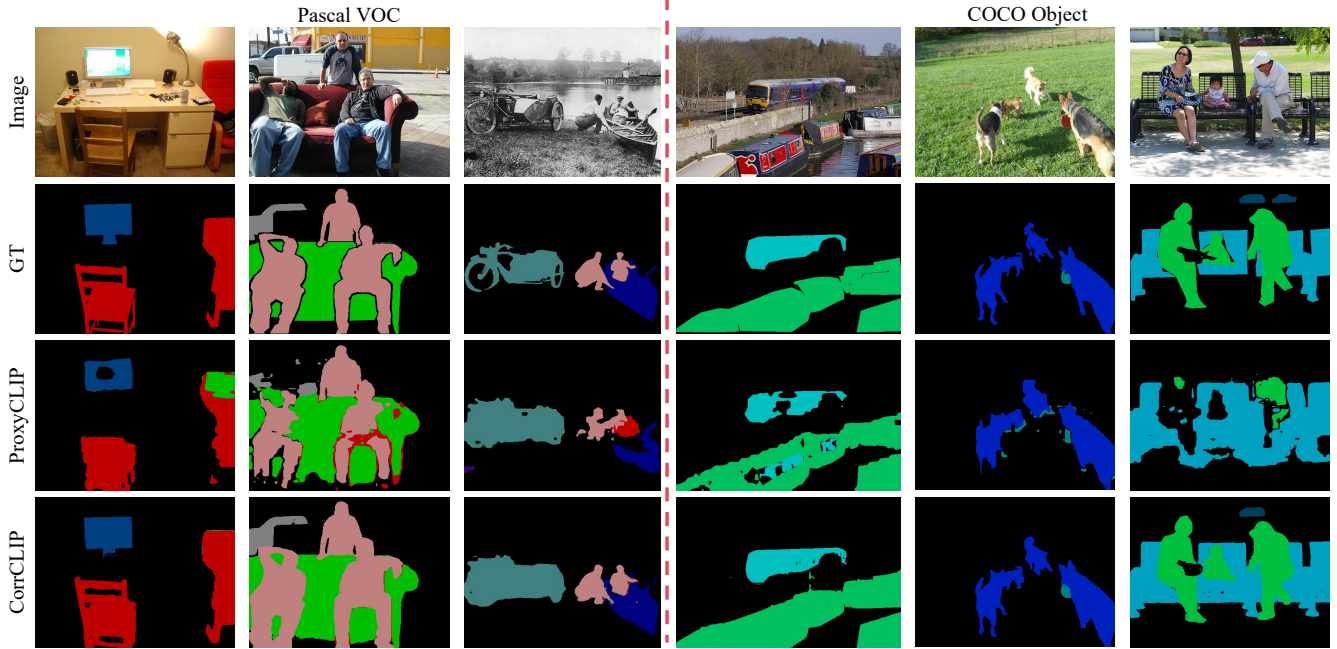
Figure 3. Qualitative comparison of segmentation maps. "GT" denotes ground truth.

SR	VR	MC	NC	VOC	PC	Object	ADE
				55.6	31.7	33.9	17.9
✓				62.2	35.5	38.3	20.3
✓	✓			66.0	37.2	40.7	21.2
✓	✓	✓		68.6	39.2	41.8	22.5
✓	✓	✓	✓	**69.2**	**39.5**	**41.8**	**22.8**

Table 2. Impact of integrating different components.

Thresh	0.4*	0.5*	0.6*	0.7	0.8	0.9
mIoU	44.6	44.5	44.4	44.6	44.2	43.4
Sampled Points	16×16		32×32		64×64	
mIoU	44.2		44.6		44.6	
w/o SAM			25.1			

Table 3. Impact of different parameters in SAM across all eight benchmarks. "*" indicates that SAM outputs one mask per prompt; otherwise, it outputs three masks.

improvement of 3.8%, 1.7%, 2.4%, and 0.9% mIoU, respectively. Incorporating segmentation map correction and class name correction also leads to a performance gain. These results demonstrate the effectiveness of each component.

Ablation study on SAM. Several hyperparameters of SAM affect the final region masks. "pred_iou_thresh" and "stability_score_thresh" are critical parameters that filter out masks with lower confidence and stability. Increasing the thresholds leads to more region masks being discarded. The number of sampled prompt points determines the number of generated masks. Increasing the sampled prompt points leads to more generated masks. SAM outputs three levels of masks (whole, part, and subpart) when the "multi-mask_output" parameter is specified; otherwise, it outputs only one mask.

We investigate the impact of these parameters on segmentation performance. Here, we use the sum of DINO's query and key to determine similarity values and set equal values for "pred_iou_thresh" and "stability_score_thresh".

As shown in Tab. 3, when SAM outputs three masks per prompt, the optimal threshold is 0.7; when SAM outputs one mask per prompt, the optimal threshold is 0.4. Increasing the threshold leads to decreased segmentation performance due to a rise in the size of unsegmented regions. The segmentation performance is not sensitive to the number of sampled points.

Notice that "w/o SAM" represents that our method only employs value reconstruction without scope reconstruction. But it yields poor results. The main reason is that the magnitude of similarity values between patches computed by DINO is unsuitable for CLIP. After confining the scope of patch interactions, it only needs to consider similarity values between semantically similar patches, alleviating this issue to some extent. The results highlight the need for scope reconstruction when using value reconstruction.

Effectiveness of scope reconstruction. To further validate the effectiveness of reconstructing the scope of patch inter-

Figure 4. Effectiveness of reconstructing the scope of patch interactions using SAM across eight benchmarks. "1" represents replacing the similarity matrix with an all-ones matrix.

Feature	Model	VOC	PC	Object	ADE
q-q	CLIP	68.4	38.9	40.7	22.4
x-x		68.2	39.2	41.3	22.7
q-q	DINO	69.1	38.9	**41.9**	22.7
k-k		67.8	39.1	41.0	22.6
qk-qk		**69.2**	**39.5**	41.8	**22.8**

Table 4. Impact of different features from CLIP and DINO used to calculate similarity matrix. "x" denotes the output features.

Type	Size	IN ZS Acc.	mIoU
CLIP		68.3	47.2
OpenCLIP	ViT-B/16	70.2	47.5
MetaCLIP		72.1	**49.7**
DFNCLIP		**76.2**	49.4
CLIP		75.5	48.3
OpenCLIP	ViT-L/14	75.3	46.7
MetaCLIP		79.2	**51.0**
DFNCLIP		**81.4**	49.4
OpenCLIP		78.0	**50.0**
MetaCLIP	ViT-H/14	80.5	47.8
DFNCLIP		**83.4**	**50.0**

Table 5. Effect of different types and sizes of CLIP models on CorrCLIP across eight benchmarks. "IN ZS ACC." is the zero-shot accuracy on ImageNet [12].

eps	0	0.1	0.15	0.2	0.25	0.3
mIoU	38.2	38.3	38.6	38.8	38.6	37.5

samples	0	1	2	3	4	5
mIoU	38.2	38.8	38.8	38.6	38.5	38.4

Table 6. Impact of different parameters in region merging on Cityscapes. Setting "eps" and "samples" to 0 indicates that region merging is not employed in scope reconstruction.

actions using SAM, we employ various features from CLIP to compute the attention map. As illustrated in Fig. 4, scope reconstruction significantly enhances segmentation performance across different model sizes and features used to compute the attention map. It's worth noting that substantial improvements can also be achieved without modifying the standard query-key calculation process.

Ablation study on value reconstruction. We investigate the impact of different features from DINO used to compute similarity matrix in value reconstruction. As shown in Tab. 4, DINO's features generally produced better results than CLIP's features in most cases, and the combination of query and key yielded the best outcomes. However, we observe that in some instances, the results using DINO's features are worse than those from CLIP. This is due to the removal of the last layer's multi-head attention in ViT.

Ablation study on CLIP. In Tab. 5, we investigate the effect of different types and sizes of CLIP models on CorrCLIP. We focus on four specific types of CLIP: CLIP [35], OpenCLIP [8], MetaCLIP [51], and DFNCLIP [16]. While these models share the same architecture, they differ from training data and processes. Our findings indicate that the zero-shot segmentation ability of CLIP models is positively correlated with their zero-shot classification ability, although this correlation is not particularly strong.

Generally, zero-shot classification accuracy improves as the model size increases. However, an increase in zero-shot classification accuracy does not always translate to better segmentation performance. For example, MetaCLIP with ViT-L/14 does not achieve the highest zero-shot classification accuracy but does exhibit the best segmentation performance. These results suggest that there is still room for improvement in leveraging CLIP's zero-shot classification capabilities for segmentation tasks.

Ablation study on region merging. We examine the impact of parameters in region merging, as shown in Tab. 6. The results indicate that merging similar regions helps establish better inter-patch correlations and enhances segmentation performance. The clustering method exhibits low sensitivity to "samples". When "eps" exceeds 0.25, the performance degrades as a larger neighborhood radius increases the likelihood of incorrectly merging regions that do not share similar semantics.

5. Conclusion

In this work, we propose a training-free method, CorrCLIP, which effectively adapts CLIP for OVSS by reconstructing coherent inter-patch correlations. We propose scope reconstruction to limit the scope of patch interactions to seman-

tically similar regions with the help of SAM. Furthermore, we propose value reconstruction to address the problem of similarity irregularity according to DINO's comprehension of an image's semantic layout. Moreover, we propose segmentation map correction and class name correction to correct the defects during CLIP's segmentation. Extensive experiments demonstrate the effectiveness of our method.

References

[1] Luca Barsellotti, Roberto Amoroso, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Training-free open-vocabulary segmentation with offline diffusion-augmented prototype generation. In *CVPR*, pages 3689–3698, 2024. 2, 6

[2] Daniel Bolya, Chaitanya Ryali, Judy Hoffman, and Christoph Feichtenhofer. Window attention is bugged: How not to interpolate position embeddings. In *ICLR*, 2023. 5

[3] Walid Bousselham, Felix Petersen, Vittorio Ferrari, and Hilde Kuehne. Grounding everything: Emerging localization properties in vision-language transformers. In *CVPR*, pages 3828–3837, 2024. 2

[4] Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. In *NeurIPS*, 2019. 1

[5] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, pages 1209–1218, 2018. 5

[6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 2, 3, 5

[7] Junbum Cha, Jonghwan Mun, and Byungseok Roh. Learning to generate text-grounded mask for open-world semantic segmentation from only image-text pairs. In *CVPR*, pages 11165–11174, 2023. 2, 6

[8] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *CVPR*, pages 2818–2829, 2023. 2, 3, 5, 6, 8

[9] Seokju Cho, Heeseong Shin, Sunghwan Hong, Anurag Arnab, Paul Hongsuck Seo, and Seungryong Kim. Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. In *CVPR*, pages 4113–4123, 2024. 2

[10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, pages 3213–3223, 2016. 5

[11] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In *ICLR*, 2024. 2, 3

[12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 8

[13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2

[14] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, pages 226–231, 1996. 5

[15] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88:303–338, 2010. 5

[16] Alex Fang, Albin Madappally Jose, Amit Jain, Ludwig Schmidt, Alexander Toshev, and Vaishaal Shankar. Data filtering networks. *arXiv preprint arXiv:2309.17425*, 2023. 3, 8

[17] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *CVPR*, pages 19358–19369, 2023. 3

[18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 3, 5

[19] You Huang, Wenbin Lai, Jiayi Ji, Liujuan Cao, Shengchuan Zhang, and Rongrong Ji. Hrsam: Efficiently segment anything in high-resolution images. *arXiv preprint arXiv:2407.02109*, 2024. 3

[20] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, pages 4904–4916, 2021. 3

[21] Lei Ke, Mingqiao Ye, Martin Danelljan, Yu-Wing Tai, Chi-Keung Tang, Fisher Yu, et al. Segment anything in high quality. In *NeurIPS*, 2024. 3

[22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, pages 4015–4026, 2023. 2, 3

[23] Mengcheng Lan, Chaofeng Chen, Yiping Ke, Xinjiang Wang, Litong Feng, and Wayne Zhang. Clearclip: Decomposing clip representations for dense vision-language inference. In *ECCV*, 2024. 2, 3, 5, 6

[24] Mengcheng Lan, Chaofeng Chen, Yiping Ke, Xinjiang Wang, Litong Feng, and Wayne Zhang. Proxyclip: Proxy attention improves clip for open-vocabulary segmentation. In *ECCV*, 2024. 2, 5, 6

[25] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *AAAI*, pages 11336–11344, 2020. 3

[26] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, pages 121–137, 2020. 3

[27] Yi Li, Hualiang Wang, Yiqun Duan, and Xiaomeng Li. Clip surgery for better explainability with enhancement in open-vocabulary tasks. *arXiv preprint arXiv:2304.05653*, 2023. 2, 3, 4

[28] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *CVPR*, pages 7061–7070, 2023. 2

[29] Yong Liu, Sule Bai, Guanbin Li, Yitong Wang, and Yansong Tang. Open-vocabulary segmentation with semantic-assisted calibration. In *CVPR*, pages 3491–3500, 2024. 2

[30] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019. 3

[31] Huaishao Luo, Junwei Bao, Youzheng Wu, Xiaodong He, and Tianrui Li. Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. In *ICML*, pages 23033–23044, 2023. 2

[32] Jiawei Ma, Po-Yao Huang, Saining Xie, Shang-Wen Li, Luke Zettlemoyer, Shih-Fu Chang, Wen-Tau Yih, and Hu Xu. Mode: Clip data experts via clustering. In *CVPR*, pages 26354–26363, 2024. 3

[33] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *CVPR*, pages 891–898, 2014. 5

[34] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3

[35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 1, 2, 3, 6, 8

[36] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 3, 5

[37] Chaitanya Ryali, Yuan-Ting Hu, Daniel Bolya, Chen Wei, Haoqi Fan, Po-Yao Huang, Vaibhav Aggarwal, Arkabandhu Chowdhury, Omid Poursaeed, Judy Hoffman, et al. Hiera: A hierarchical vision transformer without the bells-and-whistles. In *ICML*, pages 29441–29454, 2023. 5

[38] Tong Shao, Zhuotao Tian, Hang Zhao, and Jingyong Su. Explore the potential of clip for training-free open vocabulary semantic segmentation. In *ECCV*, 2024. 2, 4, 5, 6

[39] Gyungin Shin, Weidi Xie, and Samuel Albanie. Reco: Retrieve and co-segment for zero-shot transfer. In *NeurIPS*, pages 33754–33767, 2022. 2

[40] Oriane Siméoni, Gilles Puy, Huy V. Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. In *BMVC*, 2021. 2, 3

[41] Shuyang Sun, Runjia Li, Philip Torr, Xiuye Gu, and Siyang Li. Clip as rnn: Segment countless visual concepts without training endeavor. In *CVPR*, pages 13171–13182, 2024. 2, 5, 6

[42] Feng Wang, Jieru Mei, and Alan Yuille. Sclip: Rethinking self-attention for dense vision-language inference. In *ECCV*, 2024. 2, 3, 4, 5, 6

[43] Xiaoqi Wang, Wenbin He, Xiwei Xuan, Clint Sebastian, Jorge Piazentin Ono, Xin Li, Sima Behpour, Thang Doan, Liang Gou, Han-Wei Shen, et al. Use: Universal segment embeddings for open-vocabulary image segmentation. In *CVPR*, pages 4187–4196, 2024. 2

[44] Ji-Jia Wu, Andy Chia-Hao Chang, Chieh-Yu Chuang, Chun-Pei Chen, Yu-Lun Liu, Min-Hung Chen, Hou-Ning Hu, Yung-Yu Chuang, and Yen-Yu Lin. Image-text co-decomposition for text-supervised semantic segmentation. In *CVPR*, pages 26794–26803, 2024. 2, 6

[45] Monika Wysoczańska, Oriane Siméoni, Michaël Ramamonjisoa, Andrei Bursuc, Tomasz Trzciński, and Patrick Pérez. Clip-dinoiser: Teaching clip a few dino tricks. In *ECCV*, 2024. 2, 5, 6

[46] Yongqin Xian, Subhabrata Choudhury, Yang He, Bernt Schiele, and Zeynep Akata. Semantic projection network for zero-and few-label semantic segmentation. In *CVPR*, pages 8256–8265, 2019. 1

[47] Bin Xie, Jiale Cao, Jin Xie, Fahad Shahbaz Khan, and Yanwei Pang. Sed: A simple encoder-decoder for open-vocabulary semantic segmentation. In *CVPR*, pages 3426–3436, 2024. 2

[48] Yun Xing, Jian Kang, Aoran Xiao, Jiahao Nie, Ling Shao, and Shijian Lu. Rewrite caption semantics: Bridging semantic gaps for language-supervised semantic segmentation. In *NeurIPS*, 2023. 2

[49] Yunyang Xiong, Bala Varadarajan, Lemeng Wu, Xiaoyu Xiang, Fanyi Xiao, Chenchen Zhu, Xiaoliang Dai, Dilin Wang, Fei Sun, Forrest Iandola, et al. Efficientsam: Leveraged masked image pretraining for efficient segment anything. In *CVPR*, pages 16111–16121, 2024. 3

[50] Hu Xu, Po-Yao Huang, Xiaoqing Ellen Tan, Ching-Feng Yeh, Jacob Kahn, Christine Jou, Gargi Ghosh, Omer Levy, Luke Zettlemoyer, Wen-tau Yih, et al. Altogether: Image captioning via re-aligning alt-text. *arXiv preprint arXiv:2410.17251*, 2024. 3

[51] Hu Xu, Saining Xie, Xiaoqing Ellen Tan, Po-Yao Huang, Russell Howes, Vasu Sharma, Shang-Wen Li, Gargi Ghosh, Luke Zettlemoyer, and Christoph Feichtenhofer. Demystifying clip data. In *ICLR*, 2024. 2, 3, 5, 8

[52] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. In *CVPR*, pages 2945–2954, 2023. 2

[53] Qihang Yu, Ju He, Xueqing Deng, Xiaohui Shen, and Liang-Chieh Chen. Convolutions die hard: Open-vocabulary segmentation with single frozen convolutional clip. In *NeurIPS*, pages 32215–32234, 2023. 2

[54] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang,

Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. 3

[55] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *CVPR*, pages 11975–11986, 2023. 3

[56] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023. 3

[57] Fei Zhang, Tianfei Zhou, Boyang Li, Hao He, Chaofan Ma, Tianjiao Zhang, Jiangchao Yao, Ya Zhang, and Yanfeng Wang. Uncovering prototypical knowledge for weakly open-vocabulary semantic segmentation. In *NeurIPS*, 2023. 2

[58] Zhuoyang Zhang, Han Cai, and Song Han. Efficientvit-sam: Accelerated segment anything model without performance loss. *arXiv preprint arXiv:2402.05008*, 2024. 3

[59] Hang Zhao, Xavier Puig, Bolei Zhou, Sanja Fidler, and Antonio Torralba. Open vocabulary scene parsing. In *ICCV*, pages 2002–2010, 2017. 1

[60] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything. *arXiv preprint arXiv:2306.12156*, 2023. 3

[61] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, pages 633–641, 2017. 5

[62] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *ECCV*, pages 696–712, 2022. 2, 6

[63] Chong Zhou, Xiangtai Li, Chen Change Loy, and Bo Dai. Edgesam: Prompt-in-the-loop distillation for on-device deployment of sam. *arXiv preprint arXiv:2312.06660*, 2023. 3

[64] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. In *ICLR*, 2022. 3