

Structure-Aware Feature Rectification with Region Adjacency Graphs for Training-Free Open-Vocabulary Semantic Segmentation

Qiming Huang, Hao Ai, Jianbo Jiao
The MIx Group, School of Computer Science
University of Birmingham

{qxh366, hxa456}@student.bham.ac.uk, j.jiao@bham.ac.uk

Abstract

Benefiting from the inductive biases learned from large-scale datasets, open-vocabulary semantic segmentation (OVSS) leverages the power of vision-language models, such as CLIP, to achieve remarkable progress without requiring task-specific training. However, due to CLIP’s pre-training nature on image-text pairs, it tends to focus on global semantic alignment, resulting in suboptimal performance when associating fine-grained visual regions with text. This leads to noisy and inconsistent predictions, particularly in local areas. We attribute this to a dispersed bias stemming from its contrastive training paradigm, which is difficult to alleviate using CLIP features alone. To address this, we propose a structure-aware feature rectification approach that incorporates instance-specific priors derived directly from the image. Specifically, we construct a region adjacency graph (RAG) based on low-level features (e.g. colour and texture) to capture local structural relationships and use it to refine CLIP features by enhancing local discrimination. Extensive experiments show that our method effectively suppresses segmentation noise, improves region-level consistency, and achieves strong performance on multiple open-vocabulary segmentation benchmarks. Project page: <https://qiming-huang.github.io/RAG-OVS/>.

1. Introduction

Pretrained vision-language models such as CLIP [27] have demonstrated remarkable performance in zero-shot and open-vocabulary recognition tasks. Despite its effectiveness in capturing global image-text alignment, CLIP suffers from notable inductive biases at the local image level, limiting its applicability to fine-grained visual understanding. Specifically, CLIP was trained on image-text pairs without explicit supervision, enforcing implicit alignment between visual regions and textual descriptions. Consequently, it captures coarse semantic correspondences rather than fine-

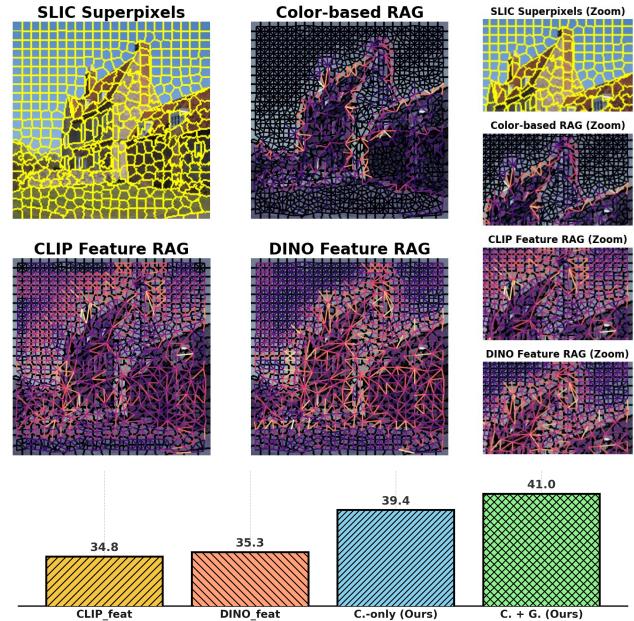


Figure 1. **Illustration of the main idea and performance.** High-level feature region adjacency graphs (RAGs) introduce local noise, while low-level colour-based RAGs maintain clean structure. The RAGs built on CLIP [27] and DINO [4] pretrained features exhibit noisy and inconsistent connectivity in local regions (see zoomed-in areas), when compared to the low-level based one. This highlights the potential of low-level cues for tasks requiring fine-grained local modelling, *e.g.* image segmentation. **Bottom:** Comparison of average performance across multiple datasets using different features for RAG construction. C.-only: colour-based features, and C. + G.: colour and texture features.

grained regional details, making it less effective for tasks requiring high spatial granularity, such as training-free open-vocabulary semantic segmentation (OVSS). This limitation manifests as noisy and inconsistent region-level predictions in training-free OVSS. As shown in Fig. 1, we observe that features extracted from CLIP [27] and DINO [4] lack clear discrimination across local (superpixel) regions, with high

levels of noise and blurred boundaries. In contrast, simple cues such as average colour differences are able to reflect local structural differences more clearly. This motivates us to ask: *Can low-level region-adjacency information be leveraged to guide CLIP toward more localised attention?*

These observations align with the contrastive training paradigm of CLIP, which encourages semantic alignment based on paired image-text data. Since high-resolution images naturally contain more discriminative details, they improve the model’s ability to form stable and structured similarity matrices. However, the inductive bias introduced by CLIP’s global training procedure cannot be easily mitigated using its own representations. As shown in Fig. 1, CLIP and DINO features are scattered and unstructured in local regions. Fortunately, the image itself inherently provides instance-specific priors that are largely immune to such global alignment biases. Specifically, the Region Adjacency Graph (RAG), constructed purely from low-level cues such as colour and texture, effectively captures spatial relationships between regions without being affected by CLIP’s global feature behaviour. Motivated by this, we propose a structure-aware feature rectification approach that incorporates RAG-based guidance into attention mechanisms. By constructing RAGs from low-level features (*e.g.* colour, GLCM texture statistics), we introduce local structure-aware biases to guide patch-level attention and similarity computation.

Extensive experiments validate that the proposed method enhances training-free open-vocabulary semantic segmentation performance. It improves regional consistency, reduces noise in segmentation outputs, and better preserves fine-grained structures, which are clearly visible in qualitative results (see Fig. 5).

2. Related Work

2.1. Contrastive Language-Image Pre-training

Contrastive Language-Image Pre-training (CLIP) [27] is a large-scale multi-modal foundation model that leverages contrastive learning to align visual and textual features, enhancing generalisation on unseen samples. Due to its strong zero-shot capabilities, CLIP has been widely adopted in Few-Shot/Zero-Shot Learning (FSL/ZSL) [17, 20, 24, 45, 46], Prompt Learning [17, 20, 45, 46], and Out-of-Distribution (OoD) detection tasks [33].

More recently, researchers have extended CLIP to dense prediction tasks [30, 38, 41, 42], such as semantic segmentation [24, 34]. However, a major challenge in utilising CLIP is the inherent noise in its features. Li *et al.* [21] analyse this issue from an explainability perspective and propose self-attention improvements to enhance CLIP’s performance in open-vocabulary tasks.

Unlike conventional pipelines that fine-tune pre-trained

models on additional datasets, CLIP’s encoder is typically kept frozen to maintain its alignment with the text feature space [44]. As a result, researchers tend to use CLIP directly as an encoder to extract preliminary features while focusing on designing sophisticated decoders [7, 9, 13, 28, 39] to refine image-level representations for dense prediction tasks.

2.2. Open-Vocabulary Semantic Segmentation

Open-vocabulary semantic segmentation (OVSS) extends segmentation [26, 35, 36] and refers to segmenting semantic regions via textual names or descriptions for the open world without any mask annotations. Early works [44] verify the importance of modal alignment in CLIP, and common downstream fine-tuning may destroy its generalisation ability. MaskCLIP [44] attempts to improve the Vision Transformer (ViT) [10] structure of CLIP to allow the model to obtain coarse feature localisation, and combines transductive learning to improve performance. CLIP-Surgery [22] analyses the difficulty of the current semantic segmentation task introduced by CLIP from the perspective of image-text noise, and makes certain improvements to the model using the idea of self-attention. SCLIP [37] inherits the idea of self-attention from MaskCLIP and directly adapts the improved CLIP structure to the semantic segmentation task.

Both CLIP-Surgery and SCLIP utilise the idea of self-attention to improve CLIP, while only CLIP-Surgery mentions the noise problem caused by the open category of text. None of them explores and analyses why CLIP lacks the semantic correlation between patches. Our work complements this point that it is the global patch formed during the attention interaction between [CLS] token and patches that leads to this.

Beyond architectural improvements, recent work scrutinizes OVS evaluation protocols regarding task ambiguity. Huang *et al.* [15] argue that rigid pixel-wise metrics contradict the open-world premise by penalizing plausible synonyms (*e.g.*, ‘sofa’ vs. ‘couch’). They propose a mask-wise evaluation protocol, demonstrating that mitigating category ambiguity significantly enhances model capabilities and suggesting a need for evolved benchmarks.

Regarding methodology, recent approaches leverage CLIP as an encoder within a “mask generation and classification” pipeline, inspired by MaskFormer [5] and Mask2Former [6]. These methods utilise pixel and query decoders to refine features and generate masks via query embeddings. By calculating the similarity between these embeddings and text prompts, the model weights query masks to produce final object boundaries and categories.

2.3. Training-free OVSS

Trident [31] proposes a training-free framework that addresses CLIP’s resolution limitation in semantic segmentation through a splice-then-segment approach. Trident first splices features extracted by CLIP and DINO from sub-images, then leverages the Segment Anything Model (SAM) for global aggregation, expanding the receptive field and improving segmentation performance. The kNN-CLIP [12] proposes a training-free approach for open-vocabulary continual segmentation that mitigates catastrophic forgetting. Instead of traditional continual training, kNN-CLIP augments the model with a database of instance embeddings, enabling segmentation methods to adapt to growing vocabularies without retraining or high memory costs. These methods primarily modify the internal attention structure of CLIP-like models to better capture relationships between image regions and textual descriptions.

In contrast, our approach takes a different direction by directly modifying the visual patch embeddings instead of adjusting attention maps. Specifically, we improve the accuracy of the visual patch-text embedding similarity matrix, ensuring a more precise alignment between visual and textual representations. By refining the embedding space at the patch level, our method enhances feature interaction and boosts segmentation performance, complementing and surpassing attention-based optimisation strategies.

3. Preliminaries of Training-free OVSS

Training-free OVSS aims to segment an image into meaningful regions by assigning semantic labels given arbitrary vocabulary, without requiring extra training. Instead of learning a segmentation model with annotated data, this approach leverages large pretrained vision-language models, such as CLIP, to directly match visual features with text embeddings through similarity computations. The visual patches embedding $\{\mathbf{v}_i\}_{i=1}^N$, where each patch \mathbf{v}_i is represented by a feature embedding of dimension $1 \times \mathbb{R}^D$, extracted from the vision encoder of CLIP. The text embedding $\{\mathbf{t}_j\}_{j=1}^M$ is obtained from the text encoder, where each \mathbf{t}_j corresponds to a text and is also represented as a feature embedding of dimension $1 \times \mathbb{R}^D$. The core idea is to compute the cosine similarity between each visual patch feature and all text embeddings:

$$s_{i,j} = \frac{\langle \mathbf{v}_i, \mathbf{t}_j \rangle}{\|\mathbf{v}_i\| \|\mathbf{t}_j\|}, \quad (1)$$

where $s_{i,j}$ is the similarity score between visual patch \mathbf{v}_i and text embedding \mathbf{t}_j . The semantic label for each visual patch is assigned based on the highest similarity score:

$$\hat{y}_i = \arg \max_j s_{i,j}, \quad (2)$$

where \hat{y}_i denotes the predicted semantic label for patch \mathbf{v}_i .

4. Structure-Aware Feature Rectification

Due to CLIP’s global training paradigm on image-text pairs, it lacks the capability for fine-grained local alignment [2, 14, 19, 29], resulting in structural inconsistency and noisy predictions when directly applied to segmentation. This issue is especially pronounced in training-free open-vocabulary semantic segmentation, where no additional data is available for model adaptation. To mitigate this, our method leverages a region adjacency graph (RAG) constructed from image low-level features to enhance structural awareness. It comprises two key modules: RAG-guided Attention, which introduces a structure-aware bias into CLIP’s attention mechanism to encourage local semantic consistency; and Similarity Fusion, which refines cross-modal similarity computation to suppress noisy matches.

4.1. RAG-guided Attention

Region adjacency graph (RAG) is a graph-based representation that captures the spatial relationships between image regions. Formally, a RAG is defined as an undirected graph $G = (V, E)$, where each node $v_i \in V$ corresponds to a low-level region R_i (e.g. a superpixel), and an edge $e_{ij} \in E$ exists if regions R_i and R_j are spatially adjacent in the image. By encoding both the appearance and structural proximity of regions, the RAG provides a compact yet informative structure that reflects the local layout of the image.

RAG construction. However, traditional RAGs are typically constructed based on the average colour differences between adjacent superpixel regions. The weight of each edge e_{ij} is defined as:

$$w_{ij}^{\text{color}} = \|\mu_i - \mu_j\|_2, \quad (3)$$

where μ_i and μ_j denote the mean RGB colour vectors of regions R_i and R_j , respectively. While this formulation provides clear local structural cues—as illustrated in Fig. 1—colour differences alone are insufficient for robust region discrimination. In real-world scenarios, colour ambiguity often arises, such as between a “white toilet” and a “white wall”. To build a more robust RAG, we incorporate not only colour differences but also texture information. Specifically, for each region R_i , we compute the *Grey-Level Co-occurrence Matrix* (GLCM) P_i and extract several statistical features from it. The edge weight is redefined as a combination of colour and texture similarities:

$$w_{ij} = w_{ij}^{\text{color}} + w_{ij}^{\text{texture}}, \quad (4)$$

where w_{ij}^{texture} is computed using the GLCM-based feature difference:

$$w_{ij}^{\text{texture}} = \sum_k \left| f_i^{(k)} - f_j^{(k)} \right|, \quad (5)$$

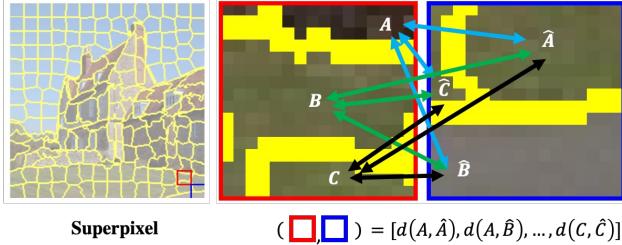


Figure 2. **Illustration of superpixel-to-patch encoding.** The distance between two patches is first represented as a list of all pairwise superpixel regions (\square, \square), then the patch distance is computed from this list using Eq. 8.

in which $f_i^{(k)}$ represents the k -th texture feature (e.g. contrast, homogeneity, energy, correlation) extracted from region R_i 's GLCM. These features are defined as follows: contrast: $\text{Contrast} = \sum_{m,n} (m - n)^2 P_i(m, n)$, homogeneity: $\text{Homogeneity} = \sum_{m,n} \frac{P_i(m, n)}{1 + |m - n|}$, energy: $\text{Energy} = \sum_{m,n} P_i(m, n)^2$, and correlation: $\text{Correlation} = \sum_{m,n} \frac{(m - \mu_m)(n - \mu_n) P_i(m, n)}{\sigma_m \sigma_n}$, where $P_i(m, n)$ denotes the normalized co-occurrence probability at position (m, n) , and $\mu_m, \mu_n, \sigma_m, \sigma_n$ are the means and standard deviations of the marginal distributions of P_i .

Superpixel-aligned patch encoding. Another challenge lies in the mismatch between the superpixel-based RAG and the patch-based tokenisation used in transformers, where inputs are typically divided into fixed-size square patches. To address this, we design a mechanism that preserves the structural advantages of superpixels—such as their ability to align flexibly with object boundaries—while enabling compatibility with patch-wise representations required by standard transformer attention. As illustrated in Fig. 2, for two adjacent patches, denoted as \square and \square , the computation of their edge weight is based on all pairwise distances between the superpixel regions contained within each patch. Specifically, let patch i contain superpixels $\{s_1^i, s_2^i, \dots, s_m^i\}$ and patch j contain superpixels $\{s_1^j, s_2^j, \dots, s_n^j\}$. We compute the pairwise distances:

$$\mathcal{D}_{ij} = \{d(s_p^i, s_q^j) \mid s_p^i \in i, s_q^j \in j\}, \quad (6)$$

where $d(\cdot, \cdot)$ is the distance function defined in Eq. 5. Therefore the computed edge weight w_{ij} is a list rather than a scalar, *i.e.*

$$w_{ij} = [d(s_1^i, s_1^j), d(s_1^i, s_2^j), \dots, d(s_m^i, s_n^j)]. \quad (7)$$

To preserve the structural variations within each patch, we compute the mean and variance of \mathcal{D}_{ij} , and use them

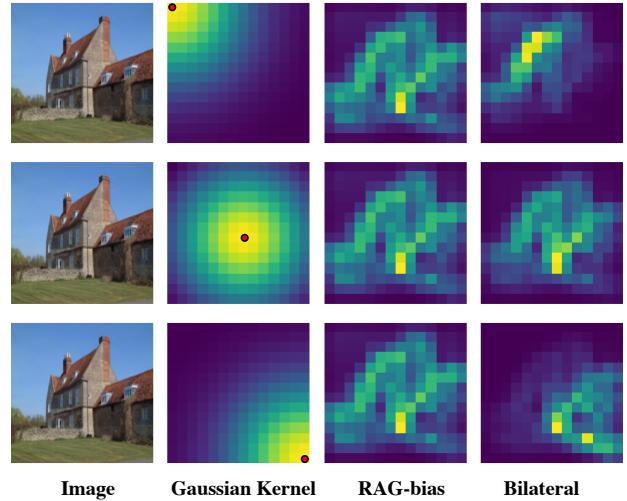


Figure 3. **Illustration of different attention bias mechanisms.** The first column shows the input images. The second column visualises the traditional Gaussian kernel, which models spatial proximity in a local window. The third column shows the RAG-bias computed from the Region Adjacency Graph (RAG), capturing structural relationships between neighbouring regions. The fourth column combines both the Gaussian kernel and the RAG-bias to form a bilateral attention bias, which accounts for both spatial distance and local structure.

as the final edge weight representation between patch i and patch j :

$$\mu_{ij} = \frac{1}{|\mathcal{D}_{ij}|} \sum_{d \in \mathcal{D}_{ij}} d, \quad \sigma_{ij}^2 = \frac{1}{|\mathcal{D}_{ij}|} \sum_{d \in \mathcal{D}_{ij}} (d - \mu_{ij})^2. \quad (8)$$

We then define the final edge weight list as (we use the standard deviation $\sigma_{i,j}$):

$$w_{ij}^{\text{final}} = [\mu_{ij}, \sigma_{ij}]. \quad (9)$$

RAG-guided attention via RAG bias. We leverage the constructed Region Adjacency Graph (RAG) to compute a structure-aware prior, referred to as the *RAG bias*, which serves as a local structural constraint in the attention mechanism. As illustrated in Fig. 3 (third column), the RAG bias is calculated for each token (patch) based on the topology of its local neighbourhood in the RAG.

Specifically, for a node (patch) i , we consider its adjacent neighbours $\mathcal{N}(i)$ —defined either as 4-connected (cross-shaped) or 8-connected neighbours. For each neighbour $j \in \mathcal{N}(i)$, we use the final edge weight $w_{ij}^{\text{final}} = [\mu_{ij}, \sigma_{ij}^2]$. The RAG bias b_{ij} is then computed by averaging the structural affinities from node i 's neighbourhood:

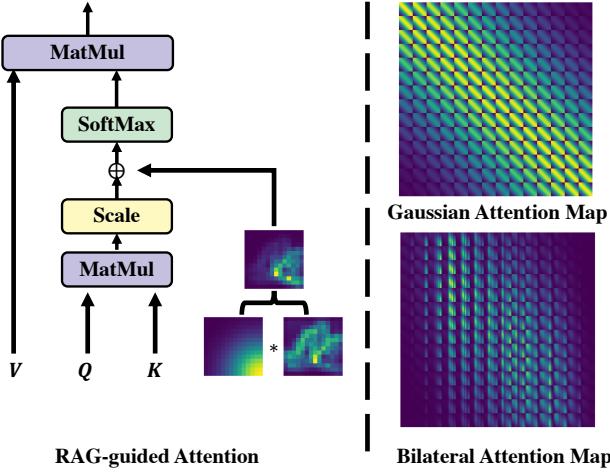


Figure 4. **Overview of the proposed RAG-guided attention mechanism.** The bilateral attention bias is computed by combining a spatial Gaussian kernel with a structure-aware RAG-bias. This combined bias is integrated into the attention weights to enhance structural sensitivity. **Right:** visualisation of the Gaussian and bilateral attention maps.

$$b_{ij} = \frac{1}{|\mathcal{N}(i)|} \sum_{k \in \mathcal{N}(i)} (\mu_{ik} + \sigma_{ik}). \quad (10)$$

Although the RAG bias $b_{i,j}$ effectively encodes structural context by aggregating information from a node’s local neighbourhood, it is fixed across all positions within the same image. This static nature makes it insufficient for capturing the pairwise relationships required in self-attention, where different attention weights are computed between every pair of tokens. To address this limitation, we draw inspiration from bilateral filtering and introduce a more flexible bias mechanism, Bilateral Bias, that combines both spatial proximity and structural similarity. Specifically, we compute a spatial Gaussian kernel $g(i, j)$ between any two positions i and j :

$$g(i, j) = \exp \left(-\frac{\|p_i - p_j\|^2}{2\sigma^2} \right), \quad (11)$$

where p_i and p_j denote the 2D coordinates of patches i and j , and σ controls the spatial range.

We then define the bilateral bias B_{ij} as the product of the spatial kernel and the structural RAG bias:

$$B_{ij} = g(i, j) \cdot \exp(b_{i,j}) \quad (12)$$

as shown in Fig. 3 (fourth column). Then the final biased attention is computed as:

$$\text{atten}_{ij}^{\text{biased}} = \text{softmax}_j \left(\frac{\mathbf{q}_i \cdot \mathbf{k}_j}{\sqrt{d}} + B_{ij} \right). \quad (13)$$

The overall of this RAG-guided attention process is shown in Fig. 4.

4.2. Similarity Fusion Module

Although introducing a bias that emphasises local structures helps the model attend to fine-grained region boundaries and maintain local consistency, it may also cause the model to respond to irrelevant local noise, such as background textures, illumination variations, or boundary artefacts. To mitigate this, we propose similarity fusion.

Specifically, given the original visual-textual similarity matrix $S_{i,j}$ defined in Eq. 1, we compute a refined similarity matrix $\tilde{S}_{i,j}$ by first applying a Gaussian kernel to smooth the visual features. Let $\hat{\mathbf{v}}_i$ denote the smoothed visual feature at position i ; we can then compute the cosine similarity between the smoothed visual features and the text feature \mathbf{t} :

$$\tilde{S}_{i,j} = \frac{\hat{\mathbf{v}}_i \cdot \mathbf{t}_j}{\|\hat{\mathbf{v}}_i\| \|\mathbf{t}_j\|}. \quad (14)$$

Finally, we fuse the original and smoothed similarities using geometric mean fusion:

$$S_{i,j}^{\text{fused}} = \left(\tilde{S}_{i,j} \right)^\alpha \cdot (S_{i,j})^{1-\alpha}. \quad (15)$$

5. Experiment

5.1. Implementation Details

Datasets. We evaluate our method on the following segmentation benchmarks, whose names are abbreviated (in parentheses) to conserve table space: PASCAL VOC 2012 (V21) [11], ADE20K-150 (ADE) [43], PASCAL Context (PC60) [25], COCO-Stuff (C-Stf) [3], Cityscapes (City) [8], COCO-Object (C-Obj) [23]. Additionally, alongside the original benchmarks on these datasets, we follow [37] and evaluate on variants of PASCAL VOC 2012 (V20) and PASCAL Context (PC59) in which the background class is removed from the evaluation.

Baselines. We compare our method to a set of relevant works in OVSS, including: MaskCLIP [44], ReCo [32], GroupVit [40], SCLIP [37], OVDiff [16], CLIPtrace [29], NACLIP [14], and ProxyCLIP [19]. It is worth noting that, for the sake of fair comparison, none of the methods, including our baselines, involve any post-processing during evaluation. This includes commonly used techniques such as Conditional Random Fields (CRF), multi-scale testing, mask refinement, or other enhancement strategies. All methods are evaluated based on their raw model outputs to ensure a fair and consistent comparison. Specifically, we adopt SCLIP, CLIPtrace, NACLIP, and ProxyCLIP as our baselines and integrate our proposed module into their original frameworks. All other settings strictly follow those described in the respective original papers.

Table 1. **Quantitative results on various OVSS benchmarks.** Our method consistently improves different CLIP-based baselines across all datasets, showing its generality and effectiveness. Best performance in **bold**.

Method	Venue	V21	PC60	C-Obj	V20	PC59	Stuff	City	ADE	Avg
CLIP [27]	ICML'21	16.4	8.4	5.6	41.9	9.2	4.4	5.0	2.9	11.7
GroupViT [40]	CVPR'22	52.3	18.7	27.5	79.7	18.5	23.4	10.4	15.3	30.7
MaskCLIP [44]	ECCV'22	43.4	23.2	20.6	74.9	26.4	16.7	24.9	11.9	30.3
Reco [32]	NeurIPS'22	25.1	19.9	15.7	57.7	21.6	22.3	11.2	14.8	23.5
OVDiff [16]	ECCV'24	66.3	29.7	34.6	80.9	32.9	20.3	23.4	14.1	37.8
CLIP-Surgery [21]	Pattern Recognition	59.0	30.1	30.2	80.1	33.9	22.1	31.8	15.8	37.9
↪ + Ours	-	61.1	32.2	31.8	81.3	34.8	23.6	33.5	17.4	39.4
SCLIP [37]	ECCV'24	59.1	30.4	30.5	80.4	34.2	22.4	32.2	16.1	38.2
↪ + Ours	-	61.9	32.9	32.3	81.8	35.0	24.1	33.9	18.1	39.9
CLIPtrace [29]	ECCV'24	53.0	30.8	33.8	81.2	35.0	24.1	35.0	17.0	38.7
↪ + Ours	-	56.2	32.1	35.2	83.8	36.4	25.1	36.6	18.5	40.5
NACLIP [14]	WACV'25	58.9	32.2	33.2	79.7	35.2	23.3	35.5	17.4	39.4
↪ + Ours	-	60.3	33.5	34.6	81.2	36.0	25.7	36.8	19.1	40.9
ProxyCLIP [19]	ECCV'24	61.3	35.3	37.5	80.3	39.1	26.5	38.1	20.2	42.3
↪ + Ours	-	62.9	36.6	38.9	82.1	39.8	27.7	40.1	21.1	43.6

Implementation details. All experiments are conducted on a single NVIDIA RTX 4090 GPU. We adopt mean Intersection-over-Union (mIoU) as the evaluation metric across all experiments. For the Similarity Fusion Module, we set the weighting parameter α to 0.6, selected based on the best performance on cocostuff171-val. The Gaussian kernel used in the module has a kernel size of 3 and a standard deviation σ of 3. SLIC [1] is used as our default super-pixel method, with n_segments=300 and compactness=10. More details about the hyperparameter sensitivity analysis can be found in Supplementary Material Section S1.

5.2. Results

Table 1 presents quantitative results on various OVSS benchmarks, comparing our method with several state-of-the-art baselines. As shown, integrating our proposed module into different CLIP-based models consistently improves performance across all datasets. Notably, our approach yields gains on challenging datasets such as ADE20K, Cityscapes, and PC60, regardless of the baseline model. These improvements validate the effectiveness and generality of our method in enhancing open-vocabulary semantic segmentation performance. Specifically, our method boosts the average mIoU by +1.8 on SCLIP (38.2 → 40.0), +1.8 on CLIPtrace (38.7 → 40.5), +1.5 on NACLIP (39.4 → 40.9), and +1.4 on ProxyCLIP (42.3 → 43.7). These results highlight the generality and effectiveness of our approach in enhancing training-free CLIP-based open-vocabulary segmentation models.

Table 2. Component ablation results based on the NACLIP model.

Method	V21	PC60	C-Obj	V20	PC59	Stuff	City	ADE
w/o	58.9	32.2	33.2	79.7	35.2	23.3	35.5	17.4
SimFusion	RAG-bias							
✓	✗	59.4	32.6	33.5	80.5	35.9	24.2	36.0
✗	✓	60.0	33.1	34.2	81.0	36.1	25.5	36.7
✓	✓	60.2	33.4	34.4	81.0	36.2	25.8	36.9
								19.2

Table 3. **Performance comparison using different feature types to construct RAG edges.** The top half shows results under standard input conditions, while the bottom half (marked with \diamond) represents experiments with colour perturbations to evaluate robustness. *C.-only* denotes colour-only input, and *C. + G.* indicates colour with additional GLCM texture statistics. **Bold** numbers highlight the best performance per column and per setting.

RAG edge	V21	PC60	C-Obj	V20	PC59	Stuff	City	ADE
CLIP_feat	55.5	28.4	23.9	74.3	33.2	19.2	28.9	14.9
DINO_feat	55.2	27.9	25.2	75.3	32.9	20.2	29.2	15.4
C.-only	58.6	32.1	32.2	80.0	35.8	24.0	35.2	17.2
C. + G.	60.0	33.1	34.2	81.0	36.1	25.5	36.7	19.0
CLIP_feat \diamond	53.0	26.7	22.4	72.0	30.2	16.4	25.3	12.8
DINO_feat \diamond	53.5	25.4	21.9	71.8	31.7	18.9	26.0	13.0
C.-only \diamond	50.4	27.6	25.3	74.2	29.9	18.0	30.4	10.2
C. + G. \diamond	58.5	32.0	32.9	79.9	35.1	23.9	35.8	18.2

5.3. Ablation Study

In this section, we conduct extensive ablation studies to analyse our model. Unless otherwise specified, the baseline used in these ablation studies is NACLIP [14].

Ablation on proposed components. Table 2 presents a component-level ablation study based on the NACLIP model, evaluating the individual and combined contribu-

tions of the Similarity Fusion module (SimFusion) and the RAG-bias mechanism. The baseline model without either component achieves an average mIoU of 39.4. Introducing SimFusion or RAG-bias individually improves the performance to 40.0 (+0.6) and 40.9 (+1.5), respectively. When both components are enabled, the model reaches the highest performance with an average mIoU of 41.2, showing consistent gains across all datasets.

These results suggest that the two components are complementary and jointly contribute to performance improvements. SimFusion enhances cross-region similarity integration, while RAG-bias introduces semantically meaningful structural bias to the attention. Notably, RAG-bias yields a greater standalone improvement than SimFusion, highlighting its stronger impact on the model’s effectiveness.

Ablation on RAG edge. Table 3 presents an ablation study comparing different feature types for constructing RAG edges. We evaluate four configurations: CLIP features, DINO features, colour-only input (*C.-only*), and colour with additional GLCM texture statistics (*C. + G.*). The top half reports results under standard conditions, while the bottom half (grey-shaded rows, marked with \heartsuit) includes colour perturbations to assess robustness. Since our RAG is primarily constructed using low-level features, it may be susceptible to common image perturbations. To analyse model behaviour under such conditions and evaluate robustness, we apply random colour jitter using the ColorJitter function with the following parameters: brightness=0.2, contrast=0.3, saturation=0.3, and hue=0.1. This augmentation introduces appearance shifts while preserving semantic content.

We find that combining colour with GLCM texture (*C.+G.*) consistently outperforms other settings under both clean and perturbed conditions. Under colour jitter, *C.-only* degrades noticeably, while *C.+G.* remains strong (e.g. 35.8 on City, 18.2 on ADE), surpassing even CLIP and DINO. This confirms the effectiveness of integrating texture features for robust RAG edge construction.

Ablation on the number of neighbours. When computing the RAG-bias (see $\mathcal{N}(i)$ in Eq. 10), we can aggregate information from a varying number of neighbouring nodes. Table 4 reports the results using 4 and 8 neighbours. While using 8 neighbours yields slightly better performance, the differences are marginal. This suggests that, once the RAG is constructed, our method is relatively insensitive to the number of aggregated neighbours.

Ablation on patch size and image size. Since our proposed superpixel-aligned patch encoding method computes representations based on the superpixel regions within each patch, both the patch size and input image resolution may

Table 4. Performance comparison with different neighbour configurations. “#neigh.” is the number of neighbours.

#neigh.	V21	PC60	C-Obj	V20	PC59	Stuff	City	ADE
4	60.1	33.0	34.2	81.0	35.8	25.6	36.5	19.0
8	60.2	33.4	34.4	81.0	36.2	25.8	36.9	19.2

Table 5. Effect of patch size and image resolution on performance. The results are reported on the NACLIP model.

Patch	Img	V21	PC60	C-Obj	V20	PC59	Stuff	City	ADE
B/16	336	60.2	33.4	34.4	81.0	36.2	25.8	36.9	19.2
B/32	336	57.2	30.0	31.2	78.7	34.3	22.7	32.5	16.4
B/16	224	58.5	31.2	30.9	79.2	35.0	23.6	34.0	16.1
B/32	224	55.2	28.0	29.6	75.0	31.7	20.0	30.6	13.6

Table 6. Comparison of different superpixel segmentation methods for RAG construction. Best performance in **bold**. A discussion on using masks generated by the Segment Anything Model (SAM) [18] is provided in Supplementary Material Section 5.

Method	V21	PC60	C-Obj	V20	PC59	Stuff	City	ADE
SLIC	60.2	33.4	34.4	81.0	36.2	25.8	36.9	19.2
Watershed	58.2	32.8	33.2	80.2	35.3	23.2	34.9	18.2
Felzenszwalb	55.2	30.1	28.9	76.9	33.0	22.5	32.9	14.5

affect the model performance. As shown in Table 5, using smaller patch sizes (e.g. B/16 vs. B/32) and higher image resolutions (e.g. 336 vs. 224) consistently leads to better performance across all benchmarks. In particular, the best results (an average mIoU of 40.9) are achieved with the B/16 patch size and 336 image resolution setting. These results suggest that finer spatial granularity in the patch-level representation helps better capture region-boundary alignment with superpixels, enhancing segmentation quality.

Ablation on different superpixel methods. In Table 6, we compare different superpixel segmentation methods for RAG construction, including SLIC, Watershed, and Felzenszwalb. Among them, SLIC achieves the best overall performance, while Felzenszwalb performs the worst across all benchmarks. We attribute the weaker performance of Felzenszwalb to its irregular and region-driven segmentation outputs, which often deviate from the grid-like patch structure used in our superpixel-to-patch encoding. In contrast, SLIC produces more compact and uniformly shaped superpixels that align better with patch boundaries, making it more compatible with our proposed encoding strategy.

Ablation on α in similarity fusion. Table 8 shows the performance of our model under different weighting parameter α in the Similarity Fusion module. We observe that the module with any $\alpha > 0$ consistently improves performance over the baseline (w/o), suggesting the effec-

Table 7. Efficiency analysis of our method. Comparison of inference speed (FPS) and computational cost (FLOPs) on different baseline models.

Method	FPS (Δ)	FLOPs (Δ)
CLIP / + Ours	72.5 \rightarrow 71.1 (-1.4)	41.7 \rightarrow \approx 41.7 (+0.0)
SCLIP / + Ours	68.8 \rightarrow 66.9 (-1.9)	44.2 \rightarrow \approx 44.2 (+0.0)
ProxyCLIP / + Ours	52.9 \rightarrow 51.5 (-1.4)	81.1 \rightarrow \approx 81.1 (+0.0)

Table 8. Performance comparison under different α values on SimFusion module based on the NACLIP model.

α	V21	PC60	C-Obj	V20	PC59	Stuff	City	ADE
w/o	58.9	32.2	33.2	79.7	35.2	23.3	35.5	17.4
0.1	58.9	32.2	33.2	79.7	35.2	23.3	35.5	17.6
0.2	59.0	32.5	33.4	79.9	35.3	23.5	35.6	17.7
0.5	59.2	32.5	33.6	80.3	35.8	24.5	35.9	18.0
0.6	59.4	32.6	33.5	80.5	35.9	24.2	36.0	18.1
0.7	58.9	32.0	33.1	80.3	35.8	24.1	35.9	18.0

tiveness of combining the original similarity matrix with its smoothed counterpart. This fusion helps suppress noisy or unreliable similarity signals, leading to more robust region aggregation. The best performance is achieved at $\alpha = 0.6$, which is used as the default setting in all experiments.

Generalisation analysis. We provide an extensive analysis of our model’s generalisation capabilities in Section S2 of the Supplementary Material, evaluating its performance on challenging transformations such as overexposure, underexposure, grayscale, style transfer, and texture destruction. In addition, we assess its zero-shot performance on a remote sensing dataset, with results reported in the Supplementary Material in Table S2 and Fig. S5.

Failure case analysis. Our method’s primary failure cases occur in underexposed conditions, where the loss of fine-grained local details causes our attention bias to be misweighted. For a detailed analysis of these failure cases, please refer to Section S3 in the Supplementary Material.

Qualitative results. In Fig. 5 we present qualitative comparisons between our method and CLIPtrace. As illustrated, our approach consistently yields more coherent and accurate segmentation results, especially in challenging regions such as object boundaries, fine-grained structures, and texture-rich areas. The highlighted areas (in red boxes) emphasise our model’s ability to preserve local consistency and object integrity, reducing the fragmented or noisy predictions that are often observed in the CLIPtrace outputs. These qualitative observations underscore the effectiveness of our proposed components in enhancing performance at finer levels of granularity.

Computational cost. As shown in Table 7, our method is computationally efficient. When integrated with various baselines, it introduces no additional FLOPs while only causing a negligible decrease in inference speed (FPS).

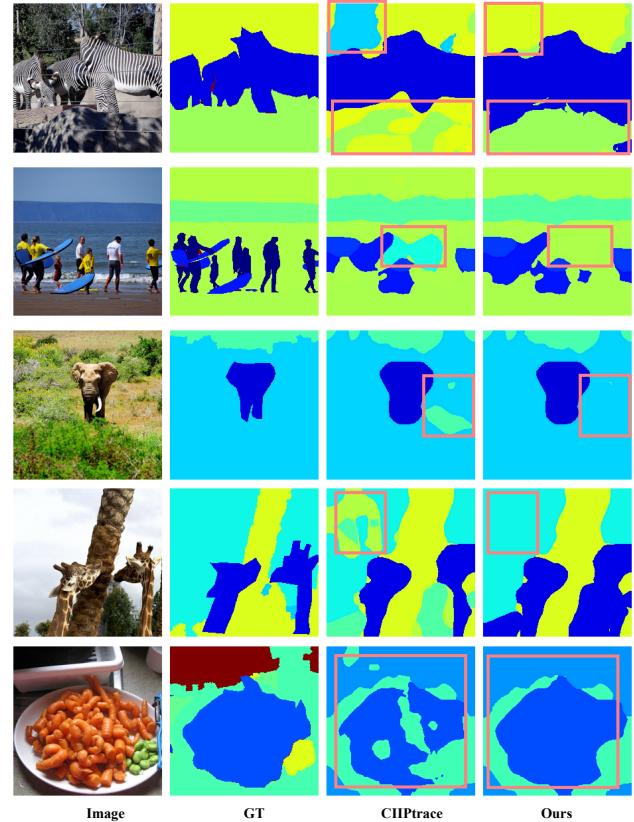


Figure 5. The qualitative results of our method. For more challenging cases, such as grayscale and stylised images (e.g. oil paintings), please refer to Figs. S4, S6, and S7 in the Supplementary Materials.

6. Conclusion

In this work, we proposed a new feature rectification approach for training-free open-vocabulary semantic segmentation. Our method leverages a Region Adjacency Graph (RAG) to refine visual patch embeddings and address local inconsistency in CLIP-based models. Specifically, we introduced a RAG-bias to guide attention toward semantically relevant regions, and a Similarity Fusion module to better align visual patches with textual categories. Extensive experimental analysis showed the effectiveness and generalisability of our approach, presenting consistent improvements across multiple datasets without additional training. Ablation studies further highlighted the importance of neighbourhood design and RAG construction, providing insights into utilising low-level priors for semantic refinement.

Acknowledgements

This project is partially supported by an Amazon Research Award. Qiming Huang is supported by the China Scholarship Council (Grant No. 202408060321). The computations in this research were performed using the Baskerville Tier 2 HPC service. Baskerville was funded by the EPSRC and UKRI through the World Class Labs scheme (EP\T022221\1) and the Digital Research Infrastructure programme (EP\W032244\1) and is operated by Advanced Research Computing at the University of Birmingham.

References

[1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2274–2282, 2012. 6

[2] Sule Bai, Yong Liu, Yifei Han, Haoji Zhang, and Yansong Tang. Self-calibrated clip for training-free open-vocabulary segmentation. *arXiv preprint arXiv:2411.15869*, 2024. 3

[3] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. 5

[4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 1

[5] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. 2

[6] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 2

[7] Seokju Cho, Heeseong Shin, Sunghwan Hong, Seungjun An, Seungjun Lee, Anurag Arnab, Paul Hongsuck Seo, and Seunghyong Kim. Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. *arXiv preprint arXiv:2303.11797*, 2023. 2

[8] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 5

[9] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. De-coupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11583–11592, 2022. 2

[10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2

[11] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 2015. 5

[12] Zhongrui Gui, Shuyang Sun, Runjia Li, Jianhao Yuan, Zhaochong An, Karsten Roth, Ameya Prabhu, and Philip Torr. knn-clip: Retrieval enables training-free segmentation on continually expanding large vocabularies. *arXiv preprint arXiv:2404.09447*, 2024. 3

[13] Jie Guo, Qimeng Wang, Yan Gao, Xiaolong Jiang, Shaohui Lin, and Baochang Zhang. Mvp-seg: Multi-view prompt learning for open-vocabulary semantic segmentation. In *Pattern Recognition and Computer Vision: 6th Chinese Conference*, page 158–171. Springer-Verlag, 2023. 2

[14] Sina Hajimiri, Ismail Ben Ayed, and Jose Dolz. Pay attention to your neighbours: Training-free open-vocabulary semantic segmentation. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5061–5071. IEEE, 2025. 3, 5, 6

[15] Qiming Huang, Han Hu, and Jianbo Jiao. Revisit the open nature of open vocabulary segmentation. In *Thirteenth International Conference on Learning Representations*, 2025. 2

[16] Laurynas Karazija, Iro Laina, Andrea Vedaldi, and Christian Rupprecht. Diffusion models for zero-shot open-vocabulary segmentation. *arXiv preprint arXiv:2306.09316*, 2023. 5, 6

[17] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19113–19122, 2023. 2

[18] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026, 2023. 7

[19] Mengcheng Lan, Chaofeng Chen, Yiping Ke, Xinjiang Wang, Litong Feng, and Wayne Zhang. Proxyclip: Proxy attention improves clip for open-vocabulary segmentation. In *European Conference on Computer Vision*, pages 70–88. Springer, 2024. 3, 5, 6

[20] Dongjun Lee, Seokwon Song, Jihee Suh, Joonmyeong Choi, Sanghyeon Lee, and Hyunwoo J Kim. Read-only prompt optimization for vision-language few-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1401–1411, 2023. 2

[21] 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, 6

[22] Yi Li, Hualiang Wang, Yiqun Duan, Jiheng Zhang, and Xiaomeng Li. A closer look at the explainability of contrastive language-image pre-training. *Pattern Recognition*, 162:111409, 2025. 2

[23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. [5](#)

[24] Yuqi Lin, Minghao Chen, Wenxiao Wang, Boxi Wu, Ke Li, Binbin Lin, Haifeng Liu, and Xiaofei He. Clip is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15305–15314, 2023. [2](#)

[25] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2014. [5](#)

[26] Bohao Peng, Zhuotao Tian, Xiaoyang Wu, Chengyao Wang, Shu Liu, Jingyong Su, and Jiaya Jia. Hierarchical dense correlation distillation for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23641–23651, 2023. [2](#)

[27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [1, 2, 6](#)

[28] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022. [2](#)

[29] Tong Shao, Zhuotao Tian, Hang Zhao, and Jingyong Su. Explore the potential of clip for training-free open vocabulary semantic segmentation. In *European Conference on Computer Vision*, pages 139–156. Springer, 2024. [3, 5, 6](#)

[30] Hengcan Shi, Munawar Hayat, Yicheng Wu, and Jianfei Cai. Proposalsclip: Unsupervised open-category object proposal generation via exploiting clip cues. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9611–9620, 2022. [2](#)

[31] Yuheng Shi, Minjing Dong, and Chang Xu. Harnessing vision foundation models for high-performance, training-free open vocabulary segmentation. *arXiv preprint arXiv:2411.09219*, 2024. [3](#)

[32] Gyungin Shin, Weidi Xie, and Samuel Albanie. Reco: Retrieve and co-segment for zero-shot transfer. *Advances in Neural Information Processing Systems*, 2022. [5, 6](#)

[33] Yang Shu, Xingzhuo Guo, Jialong Wu, Ximei Wang, Jianmin Wang, and Mingsheng Long. Clipood: Generalizing clip to out-of-distributions. *arXiv preprint arXiv:2302.00864*, 2023. [2](#)

[34] Jiajin Tang, Ge Zheng, Cheng Shi, and Sibei Yang. Contrastive grouping with transformer for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23570–23580, 2023. [2](#)

[35] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 44(2):1050–1065, 2020. [2](#)

[36] Zhuotao Tian, Xin Lai, Li Jiang, Shu Liu, Michelle Shu, Hengshuang Zhao, and Jiaya Jia. Generalized few-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11563–11572, 2022. [2](#)

[37] Feng Wang, Jieru Mei, and Alan Yuille. Sclip: Rethinking self-attention for dense vision-language inference. *European Conference on Computer Vision*, 2024. [2, 5, 6](#)

[38] Yixuan Wei, Yue Cao, Zheng Zhang, Zhuliang Yao, Zhenda Xie, Han Hu, and Baining Guo. icar: Bridging image classification and image-text alignment for visual recognition. *arXiv preprint arXiv:2204.10760*, 2022. [2](#)

[39] Letian Wu, Wenyao Zhang, Tengping Jiang, Wankou Yang, Xin Jin, and Wenjun Zeng. [cls] token is all you need for zero-shot semantic segmentation. *arXiv preprint arXiv:2304.06212*, 2023. [2](#)

[40] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. [5, 6](#)

[41] Hao Zhang, Feng Li, Xueyan Zou, Shilong Liu, Chunyuan Li, Jianwei Yang, and Lei Zhang. A simple framework for open-vocabulary segmentation and detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1020–1031, 2023. [2](#)

[42] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022. [2](#)

[43] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127:302–321, 2019. [5](#)

[44] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision*, pages 696–712. Springer, 2022. [2, 5, 6](#)

[45] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. [2](#)

[46] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [2](#)