

SSR: Semantic and Spatial Rectification for CLIP-based Weakly Supervised Segmentation

Xiuli Bi¹, Die Xiao¹, Junchao Fan^{1*}, Bin Xiao^{1, 2},

¹Chongqing Key Laboratory of Image Cognition,
Chongqing University of Posts and Telecommunications, Chongqing, China

²Jinan Inspur Data Technology Co., Ltd., Jinan, China

bixl@cqupt.edu.cn, s230201127@stu.cqupt.edu.cn, fanjc@cqupt.edu.cn, xiaobin@cqupt.edu.cn

Abstract

In recent years, Contrastive Language-Image Pretraining (CLIP) has been widely applied to Weakly Supervised Semantic Segmentation (WSSS) tasks due to its powerful cross-modal semantic understanding capabilities. This paper proposes a novel Semantic and Spatial Rectification (SSR) method to address the limitations of existing CLIP-based weakly supervised semantic segmentation approaches: over-activation in non-target foreground regions and background areas. Specifically, at the semantic level, the Cross-Modal Prototype Alignment (CMPA) establishes a contrastive learning mechanism to enforce feature space alignment across modalities, reducing inter-class overlap while enhancing semantic correlations, to rectify over-activation in non-target foreground regions effectively; at the spatial level, the Superpixel-Guided Correction (SGC) leverages superpixel-based spatial priors to precisely filter out interference from non-target regions during affinity propagation, significantly rectifying background over-activation. Extensive experiments on the PASCAL VOC and MS COCO datasets demonstrate that our method outperforms all single-stage approaches, as well as more complex multi-stage approaches, achieving mIoU scores of 79.5% and 50.6%, respectively.

Introduction

The goal of WSSS is to generate high-quality pseudo labels by using annotations such as points (Bearman et al. 2016), scribbles (Vernaza and Chandraker 2017), bounding boxes (Lee et al. 2021), or image-level labels (Pinheiro and Collobert 2015; Lee et al. 2022), thereby addressing the issues of high creation cost and time consumption associated with pixel-level annotations in fully supervised semantic segmentation (Long, Shelhamer, and Darrell 2015). Among these annotations, image-level labels are the most popular yet most challenging form of annotation. In this work, we study WSSS with image-level labels.

Current WSSS methods typically follow a three-stage pipeline: 1) training a classification network to generate initial Class Activation Maps (CAM) (Zhou et al. 2016), 2) refining these CAM (Ahn, Cho, and Kwak 2019), and 3) producing pseudo-labels for segmentation model (Xie et al.

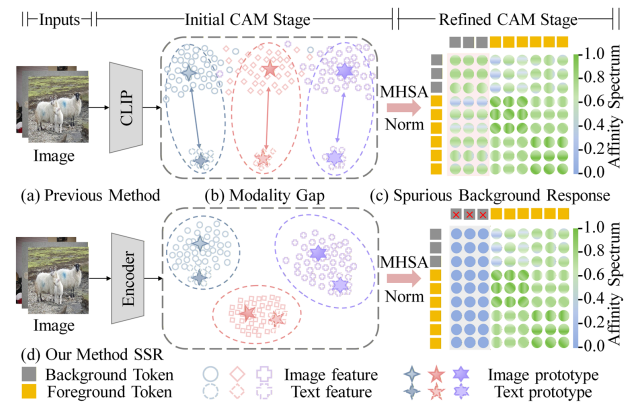


Figure 1: Our motivation. (a) Previous methods exhibit inherent limitations. (b) The initial CAM stage suffers from a modality gap, where the visual feature space faces dual challenges of intra-class dispersion and inter-class overlap. (c) The refined CAM stage is plagued by spurious background responses, as affinity estimation is corrupted by background noise. (d) To address these issues, we propose SSR.

2021) training. While conventional CNN-based CAM methods (Guo et al. 2022; He et al. 2024) suffer from under-activation (focusing only on the most discriminative object parts) due to limited receptive fields, recent Vision Transformer (ViT)-based approaches (Qin et al. 2022) demonstrate superior performance by capturing global contexts through Multi-Head Self-Attention (MHSA). However, ViT architectures introduce a new challenge of over-activation (erroneously activating background regions), which current research (Yoon et al. 2024b,a) addresses through improved positional encoding (Wu et al. 2021) or CNN-ViT hybrid architectures.

Recently, researchers have begun to explore the integration of CLIP (Radford et al. 2021) into WSSS (Yang and Gong 2024; Yang et al. 2025a). They employ gradient-based methods (Selvaraju et al. 2017) to generate initial CAM, which exhibit significantly superior performance compared to traditional CNN and ViT solutions, leading to their widespread adoption in WSSS applications.

As a multimodal model, CLIP projects images and text

*Corresponding author

into a shared representation space where matched image-text embeddings are pulled closer while mismatched ones are pushed apart. However, in downstream applications, CLIP still faces the issue of over activation in non-target foreground regions due to insufficient semantic alignment between image and text modalities, which stems from its inherent modality gap (Liang et al. 2022)—the fundamental differences in the underlying feature distributions of the two modalities, as shown in Figure 1. Specifically, visual features primarily focus on low-level patterns, such as colors and shapes, whereas textual features emphasize higher-level, abstract semantics. To address this semantic misalignment issue, existing improvement methods primarily focus on optimizing the quality of text prompts. However, these methods (Yang and Gong 2024; Yang et al. 2025a) merely modify the text-side representations without fundamentally bridging the cross-modal representational gap between visual and textual modalities. Consequently, they fail to resolve the activation errors caused by semantic misalignment completely.

During the feature refinement process, the abnormally high affinity values between background regions and target regions often lead to erroneous activation of background areas, consequently triggering the spurious background response issue (as illustrated in Figure 1). To address this problem, existing research primarily follows two solution approaches. The first employs a multi-stage iterative optimization strategy (Zhang et al. 2021) to eliminate background noise through progressively alternating training mechanisms. The second implements affinity matrix constraints (Ru et al. 2022) by applying threshold truncation or introducing constrained loss functions to enhance the reliability of supervision signals. However, these methods still suffer from interference by low-level features and the problem of global context confusion.

In this work, we propose a Semantic and Spatial Rectification to tackle the semantic misalignment caused by modality gaps and the affinity noise from spurious background response. At the semantic level, our approach aims to: 1) preserve cross-modal semantic consistency in features, and 2) ensure pixels of the same category exhibit similar representations in feature space. To achieve this, we design Cross-Modal Prototype Alignment, which enhances inter-modal alignment and feature discriminability, providing a more reliable basis for pseudo-label generation. To further enhance the quality of initial CAM for more precise pseudo-labels at the spatial level, we propose Superpixel Guided Correction. SGC addresses potential erroneous associations between target and background patch tokens by incorporating superpixel guidance, which effectively strengthens foreground semantic consistency while suppressing interference from background noise.

The main contributions of our work are listed as follows:

- We propose a Semantic and Spatial Rectification method to address the over-activation issues concerning non-target foreground regions and background areas in CLIP-based weakly supervised semantic segmentation.
- The Cross-Modal Prototype Alignment establishes a contrastive learning mechanism at the semantic level to align

the feature representations across different modalities in the feature space.

- Superpixel Guided Correction operates at the spatial level to rectify ViT’s affinity matrix by leveraging superpixel segmentation, enabling local affinity optimization.
- Extensive experiments on PASCAL VOC and MS COCO demonstrate that our method significantly outperforms recent state-of-the-art approaches.

Related Work

Weakly Supervised Semantic Segmentation

To address the limitation that CAM highlight only the most discriminative object regions, researchers have proposed various innovative solutions. Erasing (Kweon et al. 2021), cross-image mining (Sun et al. 2020; Li et al. 2021), self-supervised learning (Wang et al. 2020; Chen et al. 2022), and adversarial attack (Lee, Kim, and Yoon 2021) approaches have enhanced CAM coverage from different perspectives. Some works adopt prototype learning frameworks: PSDPM (Zhao et al. 2024) activates more secondary discriminative pixels using class prototypes, while FPR (Chen et al. 2023) constructs category-specific positive and negative prototypes. The emergence of ViTs has demonstrated promising localization capabilities. A2GNN (Zhang et al. 2021) builds semantic graph structures through affinity convolution, while AFA (Ru et al. 2022) leverages the MHSA mechanism to explore pixel-level semantic relationships. Other approaches focus on position encoding: CTI (Yoon et al. 2024b) injects category-specific tokens from images, and MCTformer (Xu et al. 2022) optimizes class-specific attention maps through inter-block pairwise affinity computation.

Contrastive Language-Image Pre-training

In the field of WSSS, CLIP has attracted significant attention due to its exceptional performance. CLIMS (Xie et al. 2022) employs contrastive learning in both foreground and background regions to perform cross-lingual image matching. CLIP-ES (Lin et al. 2023) employs carefully designed manual text prompts and utilizes the softmax function to generate GradCAM. WeakCLIP (Zhu et al. 2025) transforms the WSSS task into a continuous text-to-image matching problem, effectively leveraging vision-language pre-training knowledge. WeCLIP (Zhang et al. 2024) directly adopts the CLIP visual encoder to accomplish segmentation tasks. However, these methods all face the inherent modality gap (Liang et al. 2022) issue in multimodal models, leading to inconsistent semantic alignment between visual and textual features. To address this challenge, researchers have proposed various innovative solutions. The FMA (Yang and Gong 2024) method optimizes textual feature representation by designing learnable prompts separately for classification and segmentation tasks. ExCEL (Yang et al. 2025a) enriches text prompt information by generating fine-grained category descriptions with the help of large language models. VPL (Xu et al. 2025) employs a gradient descent approach to learn category-specific visual prototypes in the visual space, replacing traditional text prototypes to more accurately capture the features of semantic target regions.

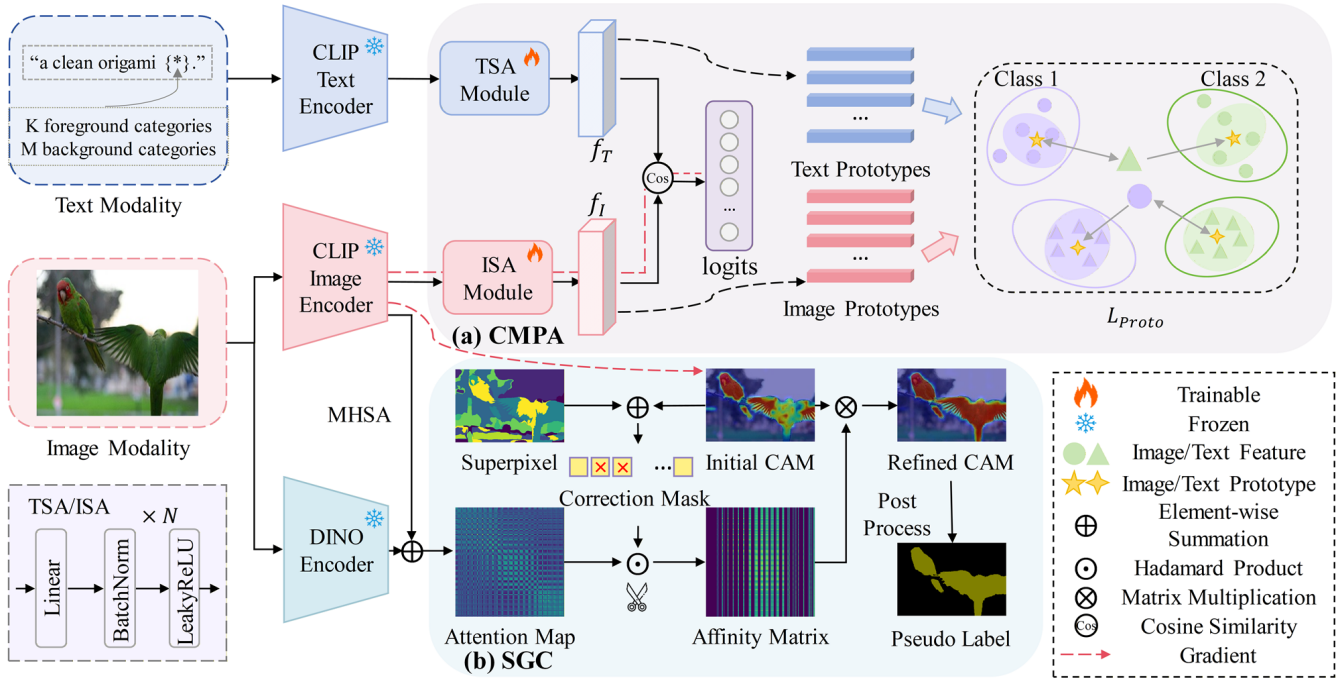


Figure 2: Overview of our SSR. We propose two novel components to address the key challenges of modality gap and erroneous activation: the CMPA and SGC. (a) The CMPA utilizes cross-modal prototype contrastive learning to establish precise matching relationships between visual features and textual prototypes in a shared embedding space, thereby effectively alleviating class confusion. (b) The SGC utilizes local spatial consistency priors derived from superpixel clustering to selectively filter the feature affinity matrix, eliminating erroneous cross-region propagation and guiding the feature refinement process toward semantically consistent directions, thereby significantly suppressing background over-activation phenomena.

Methodology

Framework Overview

Due to the significant discrepancy between textual and visual modalities, conventional text-feature-optimized alignment methods often struggle to establish precise pixel-level semantic correspondences, leading to erroneous activation in non-target foreground regions. During the refinement of CAM, background over-activation frequently occurs due to noise in affinity matrix. To systematically address these challenges, the proposed SSR Framework innovatively conducts collaborative modeling and joint optimization from dual dimensions of semantic understanding and spatial relationships, thereby achieving more accurate pixel-level semantic alignment across modalities.

Figure 2 illustrates the overall framework of SSR. The cross-modal input consists of an image modality I and a text modality T , where T comprises K foreground categories (class labels $Y=1,2,\dots,C$) and M background categories derived from CLIP-ES (Lin et al. 2023). At the semantic level, Cross-Modal Prototype Alignment significantly reduces the modality gap through contrastive learning between image and text prototypes, achieving compact cross-modal feature alignment. At the spatial level, Superpixel-Guided Correction incorporates a noise filtering mechanism that effectively suppresses erroneous information propagation in affinity modeling.

Cross-Modal Prototype Alignment

Multimodal Prototype Generation To address the inherent semantic gap between vision-language modalities, we propose a dual-branch alignment solution: For a batch of N image-text pairs $\{(I_i, T_i)_{i=1}^N\}$, the structurally identical yet parameter-independent Image Semantic Alignment (ISA) and Text Semantic Alignment (TSA) modules refine the visual features $v'_i \in \mathbb{R}^{1 \times d_1}$ and textual features $t'_i \in \mathbb{R}^{1 \times d_1}$ extracted by CLIP, respectively. With customized loss function constraints, this solution significantly improves fine-grained semantic alignment of cross-modal features. Higher semantically aligned image and text representations are obtained as:

$$v'_i = \text{ISA}(v_i), \quad t'_i = \text{TSA}(t_i), \quad (1)$$

where $v'_i \in \mathbb{R}^{1 \times d_2}$ and $t'_i \in \mathbb{R}^{c_f \times d_2}$ denote the projected image features and text features obtained through the ISA module and TSA module, respectively, where c_f represents the number of categories for the current image.

For each image, we generate CAM_{ij}^c using Grad-CAM (Selvaraju et al. 2017) and compute image/text-specific prototypes. These prototypes are constructed in the projected space of ISA and TSA for two key reasons: (1) to preserve CLIP’s inherent instance discrimination capability by shifting prototype discrimination to the projected space, and (2) to significantly reduce prototype construction costs through dimensionality reduction via the projection heads in

these modules. Specifically, the image features f_{image} and text features f_{text} containing foreground target information are computed by:

$$f_{image} = MAP(CAM^c \odot v'_i), \quad f_{text} = t'_i[index], \quad (2)$$

where $MAP(\cdot)$ represents masked average pooling. Since the text feature t'_i has dimensions $R^{c_f \times d_2}$, the text feature for the target category can be directly retrieved using the index.

We then collect the foreground-aware image features f_{image} and text features f_{text} from all image-text pairs in the dataset and perform K-means clustering to obtain the image prototypes $P^I \in \mathbb{R}^{K \times d_2} = [P_1^I, P_2^I, \dots, P_K^I]$ and text prototypes $P^T \in \mathbb{R}^{K \times d_2} = [P_1^T, P_2^T, \dots, P_K^T]$. After K-means clustering, cluster pseudo-labels can be generated based on the proximity between each sample's representation and its corresponding prototype.

Prototype Contrastive Learning This study proposes a prototype contrastive learning, which achieves fine-grained semantic alignment through triple constraints: 1) visual features are matched with text prototypes of the same category; 2) text prototypes are aggregated with visual prototypes of the same category; and 3) cross-modal prototypes of different categories are separated. As shown in Figure 2 (a), this design guides image features toward semantic content while enabling text features to focus on visually alignable attributes, directly narrowing the modality gap and enhancing inter-class discriminability in a shared embedding space, thereby effectively addressing the issue of category confusion. The core innovation lies in constructing cross-modal positive and negative sample pairs, which synchronously optimizes modality alignment and classification boundaries through contrastive learning. To establish contrastive learning, the positive and negative sample pairs constructed using visual features v'_i and text prototypes P^T are defined as:

$$p_i^I = \frac{v'_i \cdot P^T}{\tau_{proto}}, \quad (3)$$

$$\mathcal{S}_{pos}^i = \{p_{i,j} \mid j = pos_{idx}\}, \quad (4)$$

$$\{\mathcal{S}_{neg_1}^i, \dots, \mathcal{S}_{neg_k}^i\} = \{p_{i,j} \mid j \neq pos_{idx}\}, \quad (5)$$

where the temperature hyperparameter τ_{proto} is set as a learnable parameter to optimize model performance. For constructing contrastive learning samples, pos_{idx} denotes the cluster-generated pseudo-labels that guide the formation of positive and negative sample pairs, where \mathcal{S}_{pos}^i represents the positive sample pair and $\mathcal{S}_{neg_k}^i$ indicates the k -th negative sample pair. Based on this, we employ the cross-entropy loss function to compute the average loss of all in-batch samples, and the prototype contrastive loss \mathcal{L}_{proto} is implemented by:

$$\mathcal{L}_{proto} = - \left(\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\mathcal{S}_{pos}^i)}{\exp(\mathcal{S}_{pos}^i) + \sum_{j=1}^k \exp(\mathcal{S}_{neg_j}^i)} \right). \quad (6)$$

Superpixel-Guided Correction

Superpixel Clustering To address the issue of erroneous affinity propagation in attention mechanisms leading to the misactivation of background regions, we propose a Superpixel-Guided Correction. As shown in the Figure 2(b), this module leverages superpixel structural information to construct binary masks, selectively masking column vectors in the affinity matrix associated with non-target regions, thereby effectively suppressing the propagation of erroneous semantics in background regions. By constraining the attention propagation range with structured priors, only semantic correlations within target regions are preserved. Specifically, we define a binary mask matrix $Mask$, whose elements are defined as:

$$Mask = \begin{cases} a_{ij} = 1 & \text{if } j \in \text{target regions,} \\ a_{ij} = 0 & \text{if } j \notin \text{target regions,} \end{cases} \quad (7)$$

where i and j represent the row and column indices of the affinity matrix respectively.

To more accurately extract the target regions of the input image I_i , we employ the SLIC(Achanta et al. 2012) algorithm for superpixel segmentation. This method clusters pixels based on feature similarity, efficiently representing the image with a reduced number of superpixels while preserving object boundary integrity. Subsequently, the superpixel regions are clustered based on color space information to obtain the target region C as:

$$C = \text{K-means}(\text{SLIC}(I_i)). \quad (8)$$

We calculate the ratio of the sum of high-confidence pixel activations to the total activation value within each clustering region. Only regions with a ratio above a predefined threshold are classified as *target regions*. Additionally, the lightweight design of SLIC makes it more suitable than complex models like SAM for the role of spatial prior in SGC.

Affinity Matrix Correction CLIP's MHSA excels at extracting global semantic features, but its insufficient spatial detail capture leads to blurred boundaries in CAM. In contrast, DINO's MHSA, enhanced through self-supervised training, strengthens local-to-global consistency. We use DINO's local structures to improve CAM spatial priors. To address this, we integrate their MHSA features and normalize them to obtain the affinity matrix A : CLIP provides high-level semantic guidance, while DINO supplements fine-grained spatial relationships, thereby generating CAM that maintain category discriminability while achieving precise spatial localization, with the fused affinity matrix computed as:

$$A = \text{ConCat}(MHSA_{CLIP}, MHSA_{DINO}), \quad (9)$$

where $MHSA_{CLIP}$ represents the MHSA from CLIP and $MHSA_{DINO}$ denotes the MHSA from DINO, the affinity matrix A is obtained by normalizing their *ConCat* operation, effectively combining both attention mechanisms for cross-modal feature integration.

Subsequently, we refine the affinity matrix A using the obtained $Mask$ to derive an updated matrix A^* , where redundant non-target column elements are eliminated. This re-

Method	VOC		COCO
	Val	Test	Val
Multi-stage.			
SIPE (Chen et al. 2022)	68.8	69.7	43.6
CLIMS (Xie et al. 2022)	70.4	70.0	-
WeakTr (Zhu et al. 2023)	78.4	79.0	50.3
CLIP-ES (Lin et al. 2023)	73.8	73.9	45.4
PSDPM (Zhao et al. 2024)	74.1	74.9	47.2
CPAL (Tang et al. 2024)	74.5	74.7	46.8
CTI (Yoon et al. 2024b)	74.1	73.2	45.4
WeakCLIP (Zhu et al. 2025)	74.0	73.8	47.4
VPL (Xu et al. 2025)	79.3	79.0	49.8
Single-stage.			
DIAL (Jang et al. 2024)	74.5	74.9	44.4
DuPL (Wu et al. 2024)	73.3	72.8	44.6
WeCLIP (Zhang et al. 2024)	76.4	77.2	47.1
MoRe (Yang et al. 2025b)	76.4	75.0	47.4
ExCEL (Yang et al. 2025a)	78.4	78.5	50.3
Ours w/o CRF	78.2	78.1	49.2
Ours	79.5	79.6	50.6

Table 1: Segmentation performance comparison (mIoU%) on PASCAL VOC and COCO datasets, showing results on VOC val/test sets and COCO val set.

finer affinity matrix A^* is then employed to enhance the initial CAM through spatial propagation. The rectified affinity matrix A^* and the final CAM are computed as:

$$A^* = A \odot Mask, \quad (10)$$

$$CAM_{refine}^c = A^* \otimes CAM^c, \quad (11)$$

where CAM_{refine}^c denotes the refined CAM corresponding to foreground target class c .

Training Objectives

The overall loss of SSR consists of two components: a prototype contrastive loss \mathcal{L}_{proto} and a segmentation loss \mathcal{L}_{seg} . \mathcal{L}_{proto} encourages samples to move closer to their corresponding cross-modal prototypes of the same class while pushing them away from prototypes of different classes. \mathcal{L}_{seg} employs online-generated pseudo masks and adopts a cross-entropy formulation for end-to-end training to accomplish the segmentation task. The objective function of our approach can be expressed as:

$$\mathcal{L}_{SSR} = \mathcal{L}_{proto} + \gamma \mathcal{L}_{seg}. \quad (12)$$

Experiments

Experimental Settings

Datasets and Metrics The proposed method SSR is evaluated on both PASCAL VOC 2012 (Everingham et al. 2015) and MS COCO 2014 (Lin et al. 2014) datasets, where VOC comprises 21 categories (including 1 background class) and follows the established protocol (Li et al. 2021; Du et al. 2022) using augmented sets of 10,582 training images, 1,449 validation images, and 1,456 test images. COCO contains

Method	Sup.	N.	VOC
			Train
MCTformer CVPR'2022	$\mathcal{I} + \mathcal{D}$	R	61.7
CLIMS (Xie et al. 2022)	$\mathcal{I} + \mathcal{L}$	R	56.6
WeakTr (Zhu et al. 2023)	\mathcal{I}	D	66.2
POLE WACV'2023	$\mathcal{I} + \mathcal{L}$	R	59.0
CLIP-ES (Lin et al. 2023)	$\mathcal{I} + \mathcal{L}$	V	70.8
ToCo (Ru et al. 2023)	\mathcal{I}	V	71.6
CPAL (Tang et al. 2024)	$\mathcal{I} + \mathcal{L}$	R	71.9
CTI (Yoon et al. 2024b)	\mathcal{I}	D	69.5
SeCo (Yang et al. 2024)	\mathcal{I}	V	74.8
DuPL (Wu et al. 2024)	\mathcal{I}	V	75.0
WeCLIP (Zhang et al. 2024)	$\mathcal{I} + \mathcal{L}$	V	75.4
DIAL (Jang et al. 2024)	$\mathcal{I} + \mathcal{L}$	V	75.2
WeakCLIP (Zhu et al. 2025)	$\mathcal{I} + \mathcal{L}$	V	61.7
VPL (Xu et al. 2025)	$\mathcal{I} + \mathcal{L}$	V	77.8
MoRe (Yang et al. 2025b)	\mathcal{I}	V	77.0
ExCEL (Yang et al. 2025a)	$\mathcal{I} + \mathcal{L}$	V	78.0
Ours	$\mathcal{I} + \mathcal{L} + \mathcal{D}$	V	78.7

Table 2: CAM seed comparisons on VOC train set. D: DINO. N: network. R: ResNet. D:Deit. V:Vit-B.

81 categories with 82,081 training images and 40,137 validation images. We employ mIoU as the primary evaluation metric, complemented by secondary metrics including confusion ratio, precision (P) and recall (R) for comprehensive performance assessment.

Implementation Details CLIP adopts the ViT-B/16, while the DINO utilizes the ViT-S/16 architecture from DINOv1. We used the AdamW optimizer with a learning rate of $1e-5$ and weight decay of $2e-3$. For PASCAL VOC 2012, we set the batch size to 128 with a maximum of 30,000 iterations, whereas MS COCO-2014 uses a larger batch size of 256 and 80,000 iterations. The loss weight is set to 0.1 with a prototype temperature coefficient $\tau_{proto}=0.05$. Both ISA and TSA consist of stacked linear layers with batch normalization and ReLU. The prototypes are updated every 5,000 iterations. In SGC, the CLIP:DINO weighting ratio is 0.4:0.6.

Comparisons with State-of-the-art Methods

Performance of Semantic Segmentation Table 1 presents a comparison of segmentation performance between SSR and recent methods on the VOC and COCO datasets. SSR achieves new SOTA performance on both VOC and COCO, surpassing multi-stage methods by up to 0.6% and CLIP-based ExCEL by 0.3%. As shown in Figure 3, qualitative results demonstrate superior segmentation quality through: (1) enhanced cross-modal contrastive learning for accurate category prediction; (2) improved region integrity and boundary clarity; and (3) stronger inter-class discriminability in multi-category scenes.

Evaluation of CAM Seeds Table 2 reports the quality of CAM seeds on the VOC training set. Compared to recent methods, SSR further improves the CAM quality to 78.7%, surpassing the SOTA by at least 0.7%. As shown in Fig-

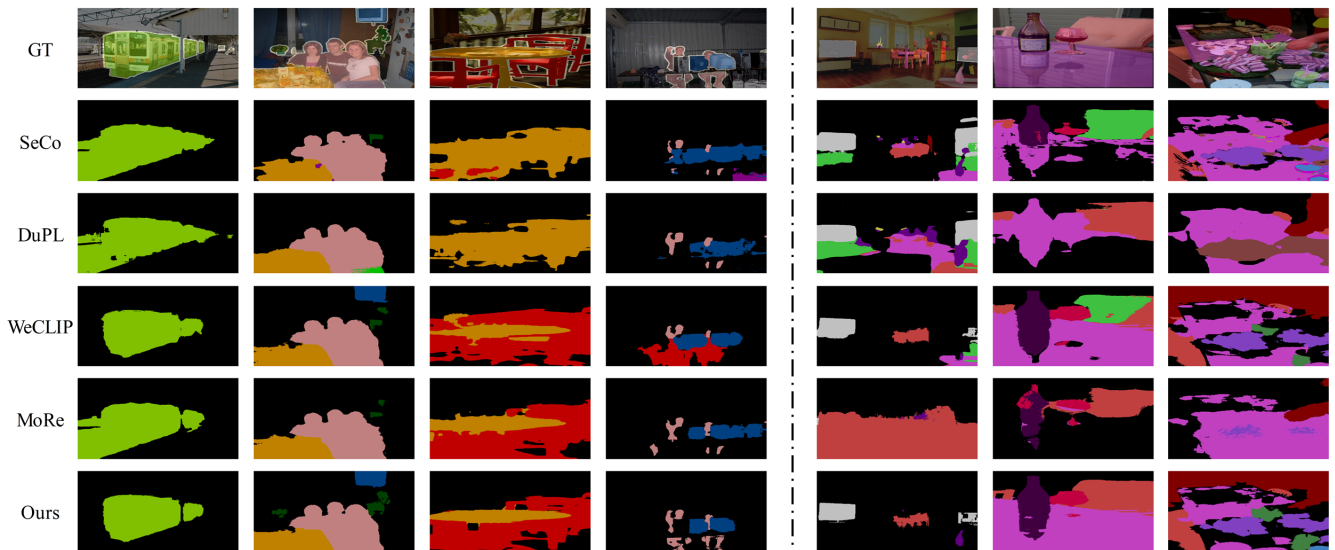


Figure 3: Segmentation visualizations of SeCo, DUPL, WeCLIP, MoRe, and Ours on VOC and COCO. Columns 1-4: Results on PASCAL VOC dataset. Columns 5-7: Results on the MS COCO dataset. SSR segments objects more precisely.




Figure 4: CAM visualizations on VOC val set. We conduct a comparative analysis between the initial CAMs generated by CLIP and those produced by our CMPA, followed by evaluating the performance gap between WeCLIP’s results and our final optimized outputs.

Figure 4, the CAM visualization results demonstrate that the initial CAMs generated by CMPA show significant advantages over CLIP, exhibiting superior completeness. Taking the “sofa” category as an example, the CLIP method erroneously activates adjacent “chair” regions, and this error gets amplified during subsequent processing. SSR not only achieves more precise focus on target regions but also shows stronger background noise suppression capability.

Ablation Study

Effectiveness of the SGC Table 3 presents the component-wise ablation study of SGC on PASCAL VOC: The CPMA-generated initial CAM achieves 63.3% mIoU, which improves to 74.6% after incorporating CLIP’s multi-head attention. Further integration with DINO’s

CPCM	At_{clip}	At_{dino}	SGC	P	R	mIoU
✓				72.8	84.6	63.3
✓	✓			85.2	88.9	74.6
✓	✓	✓		84.3	86.2	76.3
✓	✓	✓	✓	87.9	89.1	78.7

Table 3: Ablation study of SGC on the PASCAL VOC train set. At denotes attention map.

Conditions	$L_{feature}$	L_{in_modal}	L_{cross_modal}	mIoU
CLIP				58.6
w $L_{feature}$	✓			53.5
w L_{in_modal}		✓		57.8
w L_{cross_modal}			✓	63.3

Table 4: Performance evaluation of SSR using three different loss functions

attention mechanism yields 76.3% mIoU, demonstrating its complementary role in enhancing inter-patch semantic relations. Our complete SGC module ultimately achieves 78.7% mIoU by effectively rectifying fusion errors in attention maps during CAM refinement.

Effectiveness of the training loss We evaluated three different loss functions on the PASCAL VOC dataset. As shown in Table 4, baseline CLIP achieves 58.6% mIoU, while direct fine-tuning ($L_{feature}$) degrades performance by 5.1%; Intra-modal contrastive learning (L_{in_modal}) causes only a marginal 0.8% drop, indicating limited efficacy of single-modal prototype comparison; In contrast, our cross-modal contrastive loss (L_{cross_modal}) significantly improves mIoU by 4.7% by aligning cross-modal representations of

Method	mIoU	Precision	Recall	Confusion
SeCo	0.740	0.84	0.849	0.232
WeCLIP	0.764	0.844	0.861	0.237
MoRe	0.764	0.837	0.847	0.239
Ours	0.795	0.879	0.891	0.198

Table 5: Comparison with recent methods on the PASCAL VOC val set in terms of four metrics.

Methods	Sup.	Net.	Val	Ratio
DeepLabV2	F	RN101	77.7	-
DeepLabV2	F	ViT-B	82.3	-
CLIMS CVPR'2022	$\mathcal{I} + \mathcal{L}$	RN101	70.4	90.6%
CLIP-ES CVPR'2023	$\mathcal{I} + \mathcal{L}$	RN101	72.2	92.9%
CPAL CVPR'2024	$\mathcal{I} + \mathcal{L}$	RN101	74.5	95.9%
ToCo CVPR'2024	\mathcal{I}	ViT-B	71.1	86.4%
DuPL CVPR'2024	\mathcal{I}	ViT-B	73.3	89.1%
SeCo CVPR'2024	\mathcal{I}	ViT-B	74.0	89.9%
DIAL ECCV'2024	$\mathcal{I} + \mathcal{L}$	ViT-B	74.5	90.5%
WeCLIP CVPR'2024	$\mathcal{I} + \mathcal{L}$	ViT-B*	76.4	93.6%
ExCEL CVPR'2025	$\mathcal{I} + \mathcal{L}$	ViT-B*	78.4	96.1%
Ours	$\mathcal{I} + \mathcal{L}$	ViT-B*	79.5	97.4%

Table 6: Comparisons with the fully-supervised methods on VOC val set. F:fully-supervised. ViT-B*: pretrained from CLIP.

the same class while separating different classes, mitigating the modal semantic gap.

Model Performance Across Key Metrics Table 5 compares our method with SeCo, MoRe and WeCLIP on PASCAL VOC val set across four key metrics. Our method achieves superior performance on all indicators: surpassing the best baseline MoRe by 3.1% in mIoU and 3.5% in Precision, demonstrating enhanced detection accuracy. The 3% Recall improvement indicates more complete segmentation with fewer false activations, while the 3.4% reduction in Confusion ratio confirms better inter-class discrimination.

Fully-supervised Counterparts Table 6 systematically compares SSR with fully-supervised approaches. The smaller performance gap demonstrates our method’s effectiveness. SSR achieves 79.5% mIoU on VOC 2012 val set, reaching 97.4% of fully-supervised performance. The results demonstrate significant advantages over existing WSSS methods.

Visualization of CAM refinement Figure 5 illustrates the optimization effect of the SGC on the initial CAM. After superpixel refinement, the erroneous activations in background regions are significantly suppressed. This improvement primarily benefits from the carefully designed affinity matrix, which effectively eliminates interference from affinity relationships in non-target regions, thereby achieving more precise localization of activated areas.




Figure 5: Visualization of CAM refinement: (c) Initial CAM with background artifacts; (d) Refined CAM after SGC processing, showing cleaner background suppression and sharper target focus.




Figure 6: The t-SNE(Van der Maaten and Hinton 2008) visualization of feature embeddings on Pascal VOC 2012 validation images shows: (a) the original CLIP feature space distribution, and (b) the optimized feature distribution.

Effectiveness of feature alignment To validate CPMA’s effectiveness, Figure 6 compares CLIP features with CPMA-enhanced features on PASCAL VOC 2012. While CLIP features show good transferability but sparse distribution with class overlap, CPMA produces more compact intra-class clusters and clearer inter-class boundaries (e.g., distinct diningtable/bottle separation).

Conclusion

We propose SSR to address the issues of modality gap and spurious background responses in CLIP-based WSSS, effectively suppressing erroneous activations in non-target foreground and background regions. Specifically, our proposed CPMA establishes a contrastive relationship between features and cross-modal prototypes, achieving intra-class aggregation and inter-class separation. Meanwhile, SGC dynamically adjusts the propagation direction of feature affinity to effectively suppress redundant correlations in non-target regions within the affinity matrix. Extensive experiments demonstrate that our method exhibits significant advantages in segmentation accuracy and error suppression, fully validating the effectiveness of the proposed solution.

Acknowledgements

This work was supported in part by the National Science Foundation of China Joint Key Project under Grants U24B20173, and by the National Natural Science Foundation of China under Grant 62536002, Grant 62561160098, Grant 62576055 and Grant 32350009.

References

- Achanta, R.; Shaji, A.; Smith, K.; Lucchi, A.; Fua, P.; and Süsstrunk, S. 2012. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence*, 34(11): 2274–2282.
- Ahn, J.; Cho, S.; and Kwak, S. 2019. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2209–2218.
- Bearman, A.; Russakovsky, O.; Ferrari, V.; and Fei-Fei, L. 2016. What’s the point: Semantic segmentation with point supervision. In *European conference on computer vision*, 549–565. Springer.
- Chen, L.; Lei, C.; Li, R.; Li, S.; Zhang, Z.; and Zhang, L. 2023. Fpr: False positive rectification for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1108–1118.
- Chen, Q.; Yang, L.; Lai, J.-H.; and Xie, X. 2022. Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4288–4298.
- Du, Y.; Fu, Z.; Liu, Q.; and Wang, Y. 2022. Weakly supervised semantic segmentation by pixel-to-prototype contrast. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4320–4329.
- Everingham, M.; Eslami, S. A.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2015. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111: 98–136.
- Guo, J.; Han, K.; Wu, H.; Tang, Y.; Chen, X.; Wang, Y.; and Xu, C. 2022. Cmt: Convolutional neural networks meet vision transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12175–12185.
- He, J.; Cheng, L.; Fang, C.; Feng, Z.; Mu, T.; and Song, M. 2024. Progressive feature self-reinforcement for weakly supervised semantic segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 2085–2093.
- Jang, S.; Yun, J.; Kwon, J.; Lee, E.; and Kim, Y. 2024. Dial: Dense image-text alignment for weakly supervised semantic segmentation. In *European Conference on Computer Vision*, 248–266. Springer.
- Kweon, H.; Yoon, S.-H.; Kim, H.; Park, D.; and Yoon, K.-J. 2021. Unlocking the potential of ordinary classifier: Class-specific adversarial erasing framework for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6994–7003.
- Lee, J.; Kim, E.; and Yoon, S. 2021. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4071–4080.
- Lee, J.; Oh, S. J.; Yun, S.; Choe, J.; Kim, E.; and Yoon, S. 2022. Weakly supervised semantic segmentation using out-of-distribution data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16897–16906.
- Lee, J.; Yi, J.; Shin, C.; and Yoon, S. 2021. Bbam: Bounding box attribution map for weakly supervised semantic and instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2643–2652.
- Li, X.; Zhou, T.; Li, J.; Zhou, Y.; and Zhang, Z. 2021. Group-wise semantic mining for weakly supervised semantic segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 1984–1992.
- Liang, V. W.; Zhang, Y.; Kwon, Y.; Yeung, S.; and Zou, J. Y. 2022. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Advances in Neural Information Processing Systems*, 35: 17612–17625.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part v 13*, 740–755. Springer.
- Lin, Y.; Chen, M.; Wang, W.; Wu, B.; Li, K.; Lin, B.; Liu, H.; and He, X. 2023. 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*, 15305–15314.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440.
- Pinheiro, P. O.; and Collobert, R. 2015. From image-level to pixel-level labeling with convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1713–1721.
- Qin, J.; Wu, J.; Xiao, X.; Li, L.; and Wang, X. 2022. Activation modulation and recalibration scheme for weakly supervised semantic segmentation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 2117–2125.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Ru, L.; Zhan, Y.; Yu, B.; and Du, B. 2022. Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16846–16855.

- Ru, L.; Zheng, H.; Zhan, Y.; and Du, B. 2023. Token contrast for weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3093–3102.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.
- Sun, G.; Wang, W.; Dai, J.; and Van Gool, L. 2020. Mining cross-image semantics for weakly supervised semantic segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, 347–365. Springer.
- Tang, F.; Xu, Z.; Qu, Z.; Feng, W.; Jiang, X.; and Ge, Z. 2024. Hunting attributes: Context prototype-aware learning for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3324–3334.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Vernaza, P.; and Chandraker, M. 2017. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7158–7166.
- Wang, Y.; Zhang, J.; Kan, M.; Shan, S.; and Chen, X. 2020. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12275–12284.
- Wu, K.; Peng, H.; Chen, M.; Fu, J.; and Chao, H. 2021. Rethinking and improving relative position encoding for vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10033–10041.
- Wu, Y.; Ye, X.; Yang, K.; Li, J.; and Li, X. 2024. Dupl: Dual student with trustworthy progressive learning for robust weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3534–3543.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34: 12077–12090.
- Xie, J.; Hou, X.; Ye, K.; and Shen, L. 2022. Clims: Cross language image matching for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4483–4492.
- Xu, L.; Ouyang, W.; Bennamoun, M.; Boussaid, F.; and Xu, D. 2022. Multi-class token transformer for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4310–4319.
- Xu, Z.; Tang, F.; Chen, Z.; Su, Y.; Zhao, Z.; Zhang, G.; Su, J.; and Ge, Z. 2025. Toward modality gap: Vision prototype learning for weakly-supervised semantic segmentation with clip. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 9023–9031.
- Yang, X.; and Gong, X. 2024. Foundation model assisted weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 523–532.
- Yang, Z.; Fu, K.; Duan, M.; Qu, L.; Wang, S.; and Song, Z. 2024. Separate and conquer: Decoupling co-occurrence via decomposition and representation for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3606–3615.
- Yang, Z.; Meng, Y.; Fu, K.; Tang, F.; Wang, S.; and Song, Z. 2025a. Exploring CLIP’s Dense Knowledge for Weakly Supervised Semantic Segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 20223–20232.
- Yang, Z.; Meng, Y.; Fu, K.; Wang, S.; and Song, Z. 2025b. More: Class patch attention needs regularization for weakly supervised semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 9400–9408.
- Yoon, S.-H.; Kwon, H.; Jeong, J.; Park, D.; and Yoon, K.-J. 2024a. Diffusion-guided weakly supervised semantic segmentation. In *European Conference on Computer Vision*, 393–411. Springer.
- Yoon, S.-H.; Kwon, H.; Kim, H.; and Yoon, K.-J. 2024b. Class tokens infusion for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3595–3605.
- Zhang, B.; Xiao, J.; Jiao, J.; Wei, Y.; and Zhao, Y. 2021. Affinity attention graph neural network for weakly supervised semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 8082–8096.
- Zhang, B.; Yu, S.; Wei, Y.; Zhao, Y.; and Xiao, J. 2024. Frozen clip: A strong backbone for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3796–3806.
- Zhao, X.; Yang, Z.; Dai, T.; Zhang, B.; and Xiao, J. 2024. PSDPM: Prototype-based secondary discriminative pixels mining for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3437–3446.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.
- Zhu, L.; Li, Y.; Fang, J.; Liu, Y.; Xin, H.; Liu, W.; and Wang, X. 2023. Weaktr: Exploring plain vision transformer for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2304.01184*.
- Zhu, L.; Wang, X.; Feng, J.; Cheng, T.; Li, Y.; Jiang, B.; Zhang, D.; and Han, J. 2025. Weakclip: Adapting clip for weakly-supervised semantic segmentation. *International Journal of Computer Vision*, 133(3): 1085–1105.