

# Spatial Structure Constraints for Weakly Supervised Semantic Segmentation

Tao Chen<sup>1</sup>, Member, IEEE, Yazhou Yao<sup>1</sup>, Member, IEEE, Xingguo Huang<sup>1</sup>,  
Zechao Li<sup>1</sup>, Senior Member, IEEE, Liqiang Nie<sup>2</sup>, Senior Member, IEEE,  
and Jinhui Tang<sup>1</sup>, Senior Member, IEEE

**Abstract**—The image-level label has prevailed in weakly supervised semantic segmentation tasks due to its easy availability. Since image-level labels can only indicate the existence or absence of specific categories of objects, visualization-based techniques have been widely adopted to provide object location clues. Considering class activation maps (CAMs) can only locate the most discriminative part of objects, recent approaches usually adopt an expansion strategy to enlarge the activation area for more integral object localization. However, without proper constraints, the expanded activation will easily intrude into the background region. In this paper, we propose spatial structure constraints (SSC) for weakly supervised semantic segmentation to alleviate the unwanted object over-activation of attention expansion. Specifically, we propose a CAM-driven reconstruction module to directly reconstruct the input image from deep CAM features, which constrains the diffusion of last-layer object attention by preserving the coarse spatial structure of the image content. Moreover, we propose an activation self-modulation module to refine CAMs with finer spatial structure details by enhancing regional consistency. Without external saliency models to provide background clues, our approach achieves 72.7% and 47.0% mIoU on the PASCAL VOC 2012 and COCO datasets, respectively, demonstrating the superiority of our proposed approach. The source codes and models have been made available at <https://github.com/NUST-Machine-Intelligence-Laboratory/SSC>.

**Index Terms**—Semantic segmentation, weak supervision, image-level label, spatial structure constraints.

Manuscript received 24 July 2023; revised 25 December 2023; accepted 19 January 2024. Date of publication 1 February 2024; date of current version 6 February 2024. This work was supported in part by the National Natural Science Foundation of China under Grant 62102182, Grant 62202227, and Grant 62302217; in part by the Natural Science Foundation of Jiangsu Province under Grant BK20220934, Grant BK20220938, and Grant BK20220936; in part by the China Postdoctoral Science Foundation under Grant 2022M721626 and Grant 2022M711635; in part by the Jiangsu Funding Program for Excellent Postdoctoral Talent under Grant 2022ZB267; and in part by the Fundamental Research Funds for the Central Universities under Grant 30923010303. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Yulan Guo. (Corresponding author: Yazhou Yao.)

Tao Chen, Yazhou Yao, Zechao Li, and Jinhui Tang are with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing 210094, China (e-mail: taochen@njust.edu.cn; yazhou.yao@njust.edu.cn; zechao.li@njust.edu.cn; jinhuitang@njust.edu.cn).

Xingguo Huang is with the College of Instrumentation and Electrical Engineering, Jilin University, Changchun 130061, China (e-mail: xingguohuang@jlu.edu.cn).

Liqiang Nie is with the School of Computer Science and Technology, Harbin Institute of Technology (Shenzhen), Shenzhen 518055, China (e-mail: nieliqiang@gmail.com).

Digital Object Identifier 10.1109/TIP.2024.3359041

1941-0042 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.  
See <https://www.ieee.org/publications/rights/index.html> for more information.

## I. INTRODUCTION

WITH the development of deep learning algorithms, semantic segmentation has achieved remarkable progress and played an increasingly important role in many practical applications, such as autonomous driving and medical image analysis [1], [2], [3]. However, training a deep segmentation model requires a large number of high-cost pixel-wise accurate labels. Therefore, weakly supervised semantic segmentation (WSSS) has recently attracted the attention of researchers, which aims to alleviate the reliance on pixel-level annotations by resorting to weak supervision. Compared to bounding box [4], [5], [6], [7], point [8], and scribble [9], [10], image-level label [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22] is the most challenging annotation form which only indicates the presence or absence of certain classes of objects. It is also the most popular weak label due to its easy availability. In this paper, we concentrate on addressing the WSSS task based on image-level labels.

The typical pipeline of the image-level label based weakly supervised semantic segmentation is composed of three steps: 1) training a classification network with image-level annotations; 2) relying on visualization techniques to locate target objects and generating pixel-level pseudo-labels, *e.g.*, class activation maps (CAMs) [23]; 3) leveraging the obtained pseudo-labels to train a segmentation model in the fully-supervised setting. Though CAMs provide location clues for targets in images, they can only locate a small discriminative region of the target objects, which leads to unsatisfactory pseudo-labels with a high false negative rate. Therefore, recent efforts for WSSS tasks mainly focus on expanding the object activation to cover more integral object regions. For example, Kim et al. [24] proposed a discriminative region suppression approach to suppress the attention on discriminative regions and spread it to the adjacent non-discriminative area. They also leverage localization map refinement learning to further refine localization maps by recovering more missing parts. Though these methods can enlarge activated object regions, they will also inevitably result in over-activation, *i.e.*, the expanded object activation intrudes into the background area, illustrated in Fig. 1(c). Therefore, many approaches suffer from false positive activation and have to rely on additional saliency maps to provide background clues for post-processing or joint

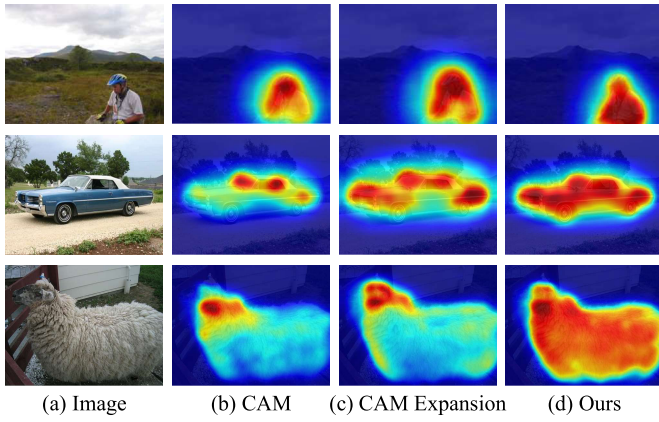


Fig. 1. Comparison between the traditional methods and ours. (a) Input image. (b) Localization maps produced by CAM [23] only identify the most discriminative part of the object, *e.g.*, the window of a car. (c) CAM expansion [24] results of traditional methods. They mainly focus on expanding the activation region and rely on saliency maps to provide background clues. Consequently, they will also inevitably result in over-activation, *i.e.*, the expanded object activation intrudes into the background area. (d) Our results. Our proposed spatial structure constraints can constrain the activation within the object area to alleviate object over-activation (first and second rows) and help activate more integral object regions to mitigate object under-activation (last row). Best viewed in color.

training. Few works focus on constraining the object attention during the expansion process.

In this paper, we propose to exploit spatial structure constraints (SSC) for weakly supervised semantic segmentation to alleviate the object over-activation problem mentioned above. Though CAMs can help locate regions of target objects, they can hardly provide any boundary clues due to the lack of shape information during training. Such amorphous localization accounts for the under-activation of naive CAMs and the over-activation of enlarged CAMs with attention expansion. Therefore, inspired by the self-supervised learning work of Autoencoder [25], we propose a CAM-driven reconstruction module to directly reconstruct the input image from its CAM features *i.e.*, the last-layer features used to generate CAMs. To endow CAMs with spatial structure information of the image content, we train the reconstruction network together with the classification backbone using a perceptual loss [26] that depends on high-level features from a pre-trained loss network, instead of using the per-pixel loss function (L1 or L2 loss) depending only on low-level pixel information. It is noteworthy that we reconstruct images directly from the class-specific CAM features. This is essential to guide the classification network to generate CAMs that maintain the spatial structure of image content by penalizing the reconstructed image that semantically and spatially deviates from the corresponding input image.

However, as we reconstruct images from deep semantic features with reduced resolution, only overall image content and coarse spatial structure are preserved in CAMs but exact shapes are not. Therefore, we propose an activation self-modulation module to further refine CAMs with higher resolution spatial structure details via enhancing regional consistency. Specifically, we resort to superpixels that group pixels that are similar in color and other low-level features.

We first refine the CAM features by averaging their values in each superpixel to obtain regional consistent features, which will also lead to CAMs with local smoothness and consistency. Then we align the CAMs to the obtained regional consistent CAM for encouraging the classification network to maintain regional consistency. However, the average operation will also significantly suppress the activation of the discriminative part if the high activation only occupies a small portion of the corresponding superpixel, which will lead to unwanted regional under-activation. Therefore, we also leverage a reliable activation selection strategy to maintain the high activation of the most discriminative region. Our proposed activation self-modulation module simultaneously enhances the regional consistency of CAM and keeps the reliable high activation, thus can significantly preserve the spatial structure details of the target objects. As can be seen in the first two rows of Fig. 1 (d), our proposed spatial structure constraints can constrain the activation within the object area to alleviate the object over-activation problem mentioned above. Besides, as shown in the last row of Fig. 1 (d), constraining the object attention with image content can also help activate more integral object regions to mitigate object under-activation not tackled by previous approaches.

Our proposed SSC is trained jointly with the classification network in a single round, which can be directly plugged into existing networks, as shown in Fig. 2. Without external saliency models to provide background clues, our approach achieves 72.7% and 47.0% mIoU on the PASCAL VOC 2012 and COCO datasets, respectively. These results demonstrate the advantage of alleviating the object over-activation problem by exploiting spatial structure constraints. Our contributions can be summarized as follows:

- We propose spatial structure constraints (SSC) for weakly supervised semantic segmentation to alleviate the object over-activation problem of CAM expansion.
- To preserve coarse spatial structure of the image content, we propose a CAM-driven reconstruction module with the perceptual loss that directly reconstructs the input image from its CAM features.
- We propose an activation self-modulation module with a reliable activation selection strategy to further refine CAMs with finer spatial structure details by enhancing regional consistency.

The rest of this paper is organized as follows: the related work is described in Section II and our approach is introduced in Section III; we then report our evaluations and ablation studies on two widely-used datasets in Section IV and finally conclude our work in Section V.

## II. RELATED WORK

### A. Weakly Supervised Semantic Segmentation

Semantic segmentation aims to identify the category of each image pixel. It has achieved significant progress with the development of deep learning. For example, while prevalent works like DeepLabV2 [27] typically learn parametric class representations, the work of [28] represents each class as a

set of non-learnable prototypes [29] to achieve dense prediction by nonparametric nearest prototype retrieving. Recently, LogicSeg [30] integrates neural inductive learning and logic reasoning into a structured parser to holistically interpret visual semantics. However, collecting dense annotations for training fully supervised segmentation models is labor-intensive and time-consuming.

Compared to pixel-level accurate labels, weakly supervised semantic segmentation resorts to weaker annotations for alleviating the annotation burden of the segmentation task, *e.g.*, image-level label [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], bounding box [4], [5], [7], point [8], and scribble [9], [10]. Among them, the image-level label is the most popular one and has attracted much attention of researchers. Visualization techniques like CAMs [23] are widely adopted to leverage image-level tags to locate target object regions for facilitating the generation of pixel-level pseudo labels. Considering object activation in CAMs is usually incomplete and lacks shape information, recent works mainly focus on: (1) expanding object attention to cover more integral target regions, and (2) refining CAMs for accurate boundaries.

An effective way for attention expansion is to transfer the surrounding discriminative information to non-discriminative object regions. For example, dilated [16] and deformable [19] convolutions are exploited to sample increasingly less discriminative object regions and Kim et al. [24] propose a discriminative region suppression approach. Adversarial erasing is proposed to hide highly activated regions and drive the network to find more object parts [31], [32]. Instead of treating each image independently, recent researchers also propose group-wise learning to mine cross-image semantic relations [33], [34]. While these methods can discover more object parts, they suffer from the object over-activation problem and rely on additional saliency maps to provide background clues [17], [24], [35], [36], [37]. The generation of saliency maps usually involves other datasets to train a salient object detection model. In contrast, in this paper, we propose spatial structure constraints for weakly supervised semantic segmentation to constrain the activation of CAM expansion.

For CAM refinement, Ahn and Kwak propose AffinityNet [15] to predict semantic affinity between a pair of adjacent image coordinates, which are then propagated via a random walk to recover delicate shapes of objects. To learn semantic affinities more effectively and efficiently, class boundary detection is exploited in IRN [38] to discover the entire instance areas with more accurate boundaries. The work of BES [39] proposes explicitly exploring object boundaries through learning with synthetic boundary labels derived from localization maps. In this paper, we propose CAM-driven reconstruction and activation self-modulation modules, which can refine CAMs with more accurate boundaries by exploiting both coarse and fine-level spatial structure information of the input image.

### B. Self-Supervised Learning

Deep supervised learning algorithms typically require large-scale annotated data to achieve satisfactory performance.

In contrast, self-supervised learning (SSL), a branch of unsupervised learning methods, aims to learn general image and video features from large numbers of unlabeled examples [40], [41], and [42]. A pre-defined pretext task is usually designed for SSL with pseudo labels automatically generated based on attributes of data, *e.g.*, relative position [43], jigsaw [44], and rotation [45]. Recently, through constructing the pretext task of instance discrimination, contrastive learning-based SSL methods, such as MoCo and SimCLR series [46], [47], [48], [49], have pushed the performance of self-supervised pretraining to a level comparable to that of supervised learning. With the advent of vision transformer, masked image modeling (MIM) methods like MAEs [50] and SimMIM [51] have also become very popular due to their massive potential in SSL. Following the idea of Autoencoder [25], the pioneering work about image generation-based SSL methods, we propose a CAM-driven reconstruction module to reconstruct the input image from features. However, different from the goal of Autoencoder to reduce the image dimension or extract features for downstream tasks, our reconstruction module aims to preserve the spatial structure of the image content. Therefore, we propose to directly reconstruct the input image from its CAM features and leverage perceptual loss to train the networks.

### C. Superpixel Guidance in Weakly Supervised Segmentation

Superpixel aims to group pixels that are similar in color and other low-level features. After obtaining visually significant entities, superpixels can lower the number of primitives required for future processing stages and thus have been widely applied in tasks like medical and satellite imaging [52]. Weakly supervised approaches also explore the contexts of superpixels to improve the segmentation performance [53], [54], [55], [56]. For example, Li et al. [54] apply a superpixel-CRF refinement model to rectify the mistakes of the initial pixel annotations by an iteration framework. SGWS [56] introduces the local and global consistency information of superpixel to optimize the coarse initial seeds that only cover the most discriminative parts of objects. Different from these methods that leverage superpixels at the post-processing stage to refine pseudo labels, we propose to exploit their regional consistency during the training of the classification network to derive more integral object activation. The most related work is ICD [55], which leverages superpixels to refine the ICD score for intra-class discrimination. However, focusing on separating the foreground and the background pixels, ICD [55] is mainly learned by single-class images to avoid mixing multi-class objects, which hampers its application in complex datasets like COCO [57]. In contrast, our proposed spatial structure constraints are conducted directly with the CAM feature and thus can be extended to multi-class images. Besides, our proposed activation self-modulation keeps the reliable high activation rather than simply learning from the regional consistency.

## III. THE PROPOSED APPROACH

In this paper, we propose spatial structure constraints (SSC) for weakly supervised semantic segmentation to alleviate the



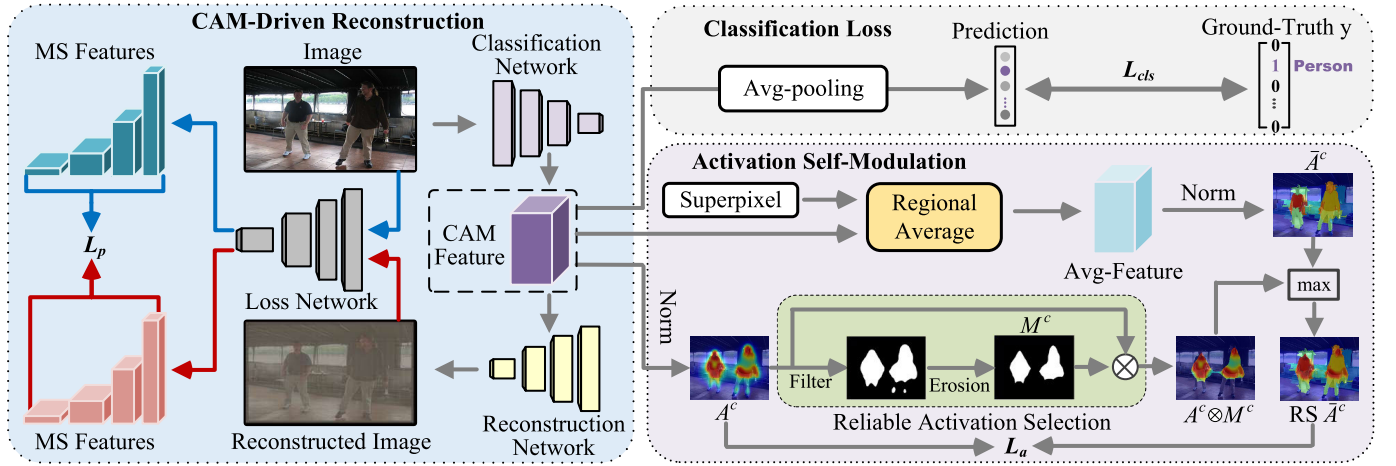


Fig. 2. The architecture of our proposed approach. While training the classification network with the image-level labels, we propose a CAM-driven reconstruction module to reconstruct the input image from its CAM-related features. Moreover, we propose an activation self-modulation module to further refine CAMs with finer spatial structure details through enhancing regional consistency. Our proposed modules help the classification network learn to preserve the spatial structure of the image content and constrain high activation within the object area.  $\otimes$  is the Hadamard product. Best viewed in color.

object over-activation problem induced by CAM expansion. Our framework is illustrated in Fig. 2. Given only image-level weak labels, a classification network is trained. Then CAMs are derived from the last-layer convolutional features to help locate the object regions and generate pixel-level pseudo labels for standard segmentation network training. Considering existing attention expansion methods suffer from the over-activation problem, we first introduce a CAM-driven reconstruction network to reconstruct the input image from its CAM-related features. This helps the classification network learn to preserve the coarse spatial structure of the image content and constrain high activation within the object area. In addition, we propose an activation self-modulation module to further refine CAMs with finer spatial structure details by enhancing regional consistency.

#### A. CAM Generation and Attention Expansion

Following the recent progress of WSSS approaches, we train a classification network and leverage CAMs to locate the target objects. To facilitate the implementation of our proposed CAM-driven reconstruction and activation self-modulation modules, we remove the final fully-connected layer in the classification network and set the output channel of the last convolutional layer to the number of classes  $C$ . We can thus directly generate object localization maps from the class-aware feature maps  $F$  of the last convolutional layer in the forward pass, which is proven by [58] identical to the attention generation process in the original CAMs [23]. In detail, for each target category  $c$ , we feed attention map  $F^c$  into a ReLU layer and then normalize it to range from 0 to 1:

$$A^c = \frac{\text{ReLU}(F^c)}{\max(F^c)}. \quad (1)$$

For the training of the classification network, we adopt the multi-label soft margin loss as follows:

$$L_{cls} = -\frac{1}{C} \sum_{c=1}^C y^c \log \sigma(q^c) + (1 - y^c) \log [1 - \sigma(q^c)]. \quad (2)$$

Here,  $\sigma(\cdot)$  is the sigmoid function.  $y^c$  is the image-level label for the  $c$ -th class. Its value is 1 if the class is present in the image; otherwise, its value is 0.

Since the classification network tends to only identify patterns from the most discriminative parts for recognition, the generated object activation is usually sparse and incomplete. Therefore, we follow the recent work of DRS [24] to expand the activated object regions, which is a simple yet effective method and requires few or no additional parameters. Given an intermediate feature map  $X \in \mathbb{R}^{K \times H \times W}$ , where  $H$ ,  $W$ , and  $K$  are the height, width, and number of channels of  $X$ , we first leverage global max pooling to extract the max-element of each channel  $X_{max} \in \mathbb{R}^{K \times 1 \times 1}$ . Then we suppress the discriminative regions by using a constant value of 0.55 (suggested in DRS [24]) to determine the upper bound of  $X$ .

$$X_{new} = \min(X, 0.55 \cdot X_{max}). \quad (3)$$

Such a parameter-free module suppresses the attention on the most discriminative regions and forces the classification network to focus on more non-discriminative regions.

#### B. CAM-Driven Reconstruction

While suppressing the highest activation can successfully spread the attention on discriminative regions to the surrounding non-discriminative object regions, it will also cause the over-activation problem. Without proper constraints, attention will easily intrude background area and parts of objects belonging to other categories. Therefore, many recent approaches resort to additional saliency maps to provide background clues for post-processing or joint training. However, the usage of saliency maps usually requires extra training of the salient object detection model with corresponding ground truth. In contrast, we propose a CAM-driven reconstruction module to help constrain the activation within the target object area. Specifically, we propose to leverage a reconstruction network to recover input images  $I$  directly from the class-specific CAM features  $F$  of the last convolutional layers:

$$\hat{I} = \text{Reconstruct}(F), \quad (4)$$

where  $\hat{I}$  is the reconstructed image. Note that all channels of the CAM features are used here to reconstruct the image (e.g., 20-channel features for the VOC dataset) rather than one or several specific feature channels related to the categories that exist in the image. For driving the classification network to preserve the spatial structure of image content, we penalize the reconstructed image that semantically and spatially deviates from the original image. To encourage the network to concentrate on the semantic content and spatial structure more than the color or texture of images, we propose to train the reconstruction network together with the backbone using a perceptual loss. This loss relies on high-level features from a pre-trained loss network rather than the per-pixel loss function depending only on low-level pixel information. Specifically, after obtaining the reconstructed image  $\hat{I}$ , we input it together with the original input image  $I$  into a pre-trained loss network  $\phi$ . Let  $\phi_j(I)$  and  $\phi_j(\hat{I})$  be the loss network's feature maps of  $j$ -th convolutional layer (stage) output with shape  $C_j \times H_j \times W_j$ , the feature reconstruction loss is the Mean Absolute Error between the two feature representations:

$$\ell_j^{\text{rec}}(\hat{I}, I) = \frac{1}{C_j H_j W_j} \|\phi_j(\hat{I}) - \phi_j(I)\|. \quad (5)$$

The perceptual loss is then defined as the sum of several reconstruction losses for multi-stage output features of the loss network:

$$L_p = \sum_{j=1}^J \frac{1}{2^{(J+1)-j}} \cdot \ell_j^{\text{rec}}(\hat{I}, I). \quad (6)$$

Here  $J$  denotes the number of stages for reconstruction loss calculation.  $\frac{1}{2^{(J+1)-j}}$  is the weight (designed to match the gradually reduced feature resolution) that controls the relative importance of reconstruction loss for different-stage features. While the higher level features contain rich semantic information, the lower level ones can provide rich clues of object edge and shape. Benefiting from the preserved spatial structure of image content, our CAM-driven reconstruction module can help constrain the activation of target objects (within their own region) from intruding into their surrounding background area. At the same time, the proposed CAM-driven reconstruction also promotes the activation to cover more integral object regions under the situation of object under-activation.

### C. Activation Self-Modulation

Benefiting from the reconstruction module with perceptual loss, the activation derived from the class-aware features tends to be more compact and consistent to align with the content of the image. However, as we reconstruct images from deep semantic features with reduced resolution, only overall image content and coarse spatial structure are preserved in CAMs but exact shapes are not. To endow CAM with more detailed spatial structure knowledge, we propose an activation self-modulation module to further refine CAMs via enhancing regional consistency. For regional consistent activation learning, we first refine the CAM features  $F$  by averaging their values in each superpixel  $S_k$  to obtain regional consistent

feature representations:

$$\bar{F}_{ij} = \frac{\sum_{m,n \in S_k(ij)} F_{mn}}{|S_k(ij)|}. \quad (7)$$

Here,  $S_k(ij)$  denotes the  $k$ -th superpixel that contains the pixel at position  $(i, j)$ .  $|\cdot|$  is the operation that counts the number of pixels in  $S_k(ij)$ . Then we apply the ReLU and maximum normalization illustrated in Equation 1 to generate the regional consistent CAM. After that, we propose to align the CAM  $A^c$  to the obtained regional consistent CAM  $\bar{A}^c$  with a mean-squared loss for driving the classification network to extract features maintaining the regional consistency:

$$L_a = \frac{1}{C' H' W'} \sum_{c=1}^{C'} \|A^c - \bar{A}^c\|^2, \quad (8)$$

where  $H'$ ,  $W'$ , and  $C'$  are the height, width of the CAM, and the number of classes that exist in the image. Note that we up-sample the CAM here to match the half resolution of the superpixel for taking advantage of finer spatial structure information. With our regional consistent activation learning, we can increase the attention of non-discriminative object areas, especially those within the superpixel that contain the most discriminative object part.

However, the average operation will also dilute the activation of the discriminative part if it only occupies a small portion of the corresponding superpixel, resulting in unwanted regional low-activation. In the scenario where several sparse high activation regions exist, the original high attention within a large superpixel might be significantly suppressed after the average dilation. This will lead to the false modulation towards the opposite direction, and the localization map might gradually lose the activation of the object region. Therefore, we propose a reliable activation selection strategy to maintain the high activation of the most discriminative regions. Specifically, we first leverage an object threshold  $T_{obj}$  to filter out the area with relatively higher attention values. Then we apply an erosion operation to further narrow the region mask  $M^c$  with high activation for reliable activation selection. Our erosion operation can help remove the unwanted high activation intruding into the background. Consequently, we obtain the CAM with reliable activation selection (RS  $\bar{A}^c$ ) for alignment as follows:

$$\bar{A}^c = \max(\bar{A}^c, A^c \otimes M^c), \quad (9)$$

where  $\otimes$  is the Hadamard product. Our proposed activation self-modulation module simultaneously enhances the regional consistency of CAM and keeps the reliable high activation, thus can significantly preserve the spatial structure details of target objects.

### D. Training Objective

With our proposed CAM-driven reconstruction and activation self-modulation modules, the overall training loss of the classification network is as follows:

$$L = L_{cls} + \beta_p L_p + \beta_a L_a. \quad (10)$$

Here,  $\beta_p$  and  $\beta_a$  are the hyperparameters that control the relative importance of perceptual loss and alignment loss. After leveraging the trained classification network to obtain CAMs to locate target objects, we follow the recent works [59], [60], [61], [62], [63] and adopt IRN [38] to further refine CAMs for pseudo label generation.

#### IV. EXPERIMENTS

##### A. Datasets and Evaluation Metrics

Following previous works, we evaluate our approach on the PASCAL VOC 2012 dataset [64] and COCO dataset [57]. As the most popular benchmark for WSSS, the PASCAL VOC 2012 dataset contains 21 classes (20 object categories and the background) for semantic segmentation. The official dataset split contains 1,464 images for training, 1,449 for validation and 1,456 for testing. Following the common protocol in semantic segmentation, we expand the training set to 10,582 images with additional data from SBD [65]. The COCO dataset is a more challenging benchmark with 80 semantic classes and the background. Following previous works [60], [66], [67], we use the default train/val splits (80k images for training and 40k for validation) in the experiment. For all the experiments, we only adopt the image-level classification labels for training. Mean intersection over union (mIoU) is adopted as the metric to evaluate the quality of our generated CAMs, pseudo labels and segmentation results. The results for the PASCAL VOC test set are obtained from the official evaluation server.

##### B. Implementation Details

For the classification network, we follow the work of IRN [38] and adopt the ResNet-50 [68] model as our backbone for a fair comparison, which is pre-trained on ImageNet [69]. A  $1 \times 1$  convolutional layer of  $C$  channels is adopted as the pixel-wise classifier to generate CAM features. Then, the CAM features are directly input into our CAM-driven reconstruction module, and coarse spatial structure constraint on features is added through minimizing the perceptual loss between the reconstructed and the original images. On the other hand, CAM features are also fed into our activation self-modulation module to further refine CAMs with finer spatial structure details by enhancing regional consistency with an alignment loss. The momentum and weight decay of the SGD [70] optimizer are 0.9 and  $1 \times 10^{-4}$ . The initial learning rate is set to 0.1. The input images are resized and cropped to  $512 \times 512$ , and the size of generated CAM feature is  $32 \times 32$  with stride = 16. We train the classification network for 10 epochs with batch size = 6.

The reconstruction network comprises one head ConvBlock, one ResBlock, four Up-sample blocks and one tail ConvBlock to recover the size of CAM features to the input image. Specifically, each Up-sample block contains a  $4 \times 4$  transpose convolution to up-sample the features and a  $3 \times 3$  convolution to further aggregate the representation. Its detailed architecture is demonstrated in Table I. For the loss network, we adopt the VGG-19 model [71] pre-trained on ImageNet [69]. For the perceptual loss, we set  $J = 5$  to choose features before each

TABLE I  
DETAILED ARCHITECTURE OF THE RECONSTRUCTION NETWORK

Num	Block	Specification
1	ConvBlock	$3 \times 3$ Conv, LayerNorm, Relu
1	ResBlock	$3 \times 3$ Conv, LayerNorm, Relu $3 \times 3$ Conv, LayerNorm
4	Up-sample	$4 \times 4$ ConvTranspose, InstanceNorm, Relu $3 \times 3$ Conv, LayerNorm, Relu
1	ConvBlock	$3 \times 3$ Conv, Tanh

TABLE II  
ACCURACY (mIoU) OF PSEUDO-MASKS EVALUATED ON PASCAL VOC 2012 TRAINING SET

Methods	seed	w/ IRN [38]
IRN [38] <sub>CVPR19</sub>	48.3	66.3
MBMNet [59] <sub>ACMMM20</sub>	50.2	66.8
CONTA [60] <sub>NIPS20</sub>	48.8	67.9
AdvCAM [61] <sub>CVPR21</sub>	55.6	69.9
RIB [62] <sub>NIPS21</sub>	56.6	70.6
ReCAM [63] <sub>CVPR22</sub>	56.6	70.5
ESOL [19] <sub>NIPS22</sub>	53.6	68.7
D2CAM [76] <sub>ICCV23</sub>	58.0	71.4
<b>SSC (Ours)</b>	<b>58.3</b>	<b>71.9</b>

pooling layer for reconstruction loss calculation. For the superpixels, we directly leverage the ones provided by the work of ICD [55] for fair comparison. Specifically, superpixels are first generated with the method of [72] and then hierarchically merged with selective search [73] so that each image contains at most 64 superpixels.

For the second stage training of WSSS, following the recent work of BECO [74], we adopt DeeplabV2 [27] as the segmentation network, which uses ResNet101 [68] as the backbone with an output stride of 16. The momentum and weight decay of the SGD optimizer are 0.9 and  $10^{-4}$ . The initial learning rate is set to  $10^{-2}$  and is decreased using polynomial decay. The segmentation model is trained for 80 epochs and 40 epochs on VOC and MS COCO datasets, respectively, with a common batch size of 16. We also follow the default setting of DeeplabV2 [27] and conduct experiments with VGG16 backbone for a more comprehensive comparison with previous approaches [15], [35], [36], [39], [55], [75]. Both ResNet101 and VGG16 backbones are pretrained on ImageNet dataset [69].

##### C. Comparisons to the State-of-the-Arts

1) *Accuracy of Pseudo-Masks*: For weakly supervised semantic segmentation, the quality of generated pseudo labels directly influences the performance of the trained segmentation network. Therefore, we first present the comparison of the quality of pseudo-masks derived from our approach and other state-of-the-arts. As shown in Table II, the segmentation seed generated by our approach can arrive at the mIoU of 58.3%, bringing a gain of 10% compared to the baseline reported by IRN [38]. Our method can obtain more accurate seeds than state-of-the-art methods like RIB [62] and ReCAM [63] by 1.7%. Our proposed SSC can also outperform the recent D2CAM [76] by 0.3% mIoU. Leveraging the random walk

TABLE III

QUANTITATIVE COMPARISONS TO PREVIOUS STATE-OF-THE-ART APPROACHES ON PASCAL VOC 2012 VAL AND TEST SETS WITH VGG BACKBONE. SUP: SUPERVISION, I: IMAGE-LEVEL LABEL, S: SALIENCY MAPS

Methods	Sup	Backbone	Val	Test
DRS [24] <sub>AAAI21</sub>	I+S	VGG16	63.6	64.4
GSM [67] <sub>AAAI21</sub>	I+S	VGG16	63.3	63.6
NSROM [77] <sub>CVPR21</sub>	I+S	VGG16	65.5	65.3
EPS [35] <sub>CVPR21</sub>	I+S	VGG16	67.0	67.3
L2G [36] <sub>CVPR22</sub>	I+S	VGG16	68.5	68.9
AffinityNet [15] <sub>CVPR18</sub>	I	VGG16	58.4	60.5
ICD [55] <sub>CVPR20</sub>	I	VGG16	61.2	60.9
BES [39] <sub>ECCV20</sub>	I	VGG16	60.1	61.1
ECS [75] <sub>ICCV21</sub>	I	VGG16	62.1	63.4
<b>SSC (Ours)</b>	I	VGG16	<b>66.8</b>	<b>67.4</b>

algorithm in IRN [38] to further refine CAMs for pseudo label generation, the mIoU of our pseudo-masks can reach 71.9%.

2) *Accuracy of Segmentation Maps on PASCAL VOC 2012*: We present our segmentation results on PASCAL VOC 2012 for the backbone of VGG and ResNet in Table III and Table IV, respectively. As can be seen, for the VGG backbone, our approach achieves better results than other state-of-the-art methods with only image-level labels. Specifically, our segmentation results can reach 67.4% on the test set, which outperforms the recent work of ECS [75] by 4%. Our performance is also competitive to many approaches that rely on saliency maps, *e.g.*, NSROM [77] and EPS [35]. For the ResNet backbone, we can get 72.8% on the validation set and 72.7% on the test set. Our proposed approach can outperform the recent work of ACR [78] and BECO [74] by 0.6% on the validation set and 0.9% on the test set. Though approaches of W-OoD [79] and CLIP-ES [80] leverage additional out-of-distribution images and language supervision to help distinguish the foreground from the background, our method can still outperform them by 1.4%. Besides, our proposed SSC can also achieve competitive performance with recent SOTA approaches like PPC [81] and RCA [82] that require additional saliency models. Some example prediction maps on the PASCAL VOC 2012 val set can be viewed in Fig. 3.

3) *Accuracy of Segmentation Maps on COCO*: For the more challenging COCO dataset, we provide performance comparisons with state-of-the-art WSSS methods for the backbone of VGG and ResNet in Table V and Table VI, respectively. As shown in Table V, our proposed SSC with VGG backbone can achieve the performance of 38.1% mIoU, much better than previous methods supervised with only image-level labels, *e.g.*, 14.4% mIoU higher than CONTA [60]. Besides, our approach can also obtain better results compared to previous SOTA methods with additional saliency guidance (*e.g.*, outperform RCA [82] and MDBA [1] by 1.3% and 0.3% mIoU on the validation set, respectively). Similarly, with the ResNet backbone, our proposed SSC reaches the best results of 47.0% mIoU compared to previous SOTA WSSS methods. Specifically, our approach outperforms ACR [78] and BECO [74] by 1.7% and 1.9% mIoU, respectively. Our SSC can also obtain 1.6% higher mIoU than CLIP-ES [80] than relies on language-image

TABLE IV

QUANTITATIVE COMPARISONS TO PREVIOUS STATE-OF-THE-ART APPROACHES ON PASCAL VOC 2012 VAL AND TEST SETS WITH RESNET BACKBONE. SUP: SUPERVISION, I: IMAGE-LEVEL LABEL, S: SALIENCY MAPS, O: OUT-OF-DISTRIBUTION DATA, L: LANGUAGE SUPERVISION

Methods	Sup	Backbone	Val	Test
DRS [24] <sub>AAAI21</sub>	I+S	R101	71.2	71.4
EPS [35] <sub>CVPR21</sub>	I+S	R101	71.0	71.8
AuxSegNet [83] <sub>ICCV21</sub>	I+S	WR38	69.0	68.6
PPC [81] <sub>CVPR22</sub>	I+S	R101	72.6	73.6
RCA [82] <sub>CVPR22</sub>	I+S	R101	72.2	72.8
L2G [36] <sub>CVPR22</sub>	I+S	R101	72.1	71.7
W-OoD [79] <sub>CVPR22</sub>	I+O	R101	69.8	69.9
CLIP-ES [80] <sub>CVPR23</sub>	I+L	R101	71.1	71.4
LPCAM [84] <sub>CVPR23</sub>	I+S	R101	71.8	72.1
IRN [38] <sub>CVPR19</sub>	I	R50	63.5	64.8
ICD [55] <sub>CVPR20</sub>	I	R101	64.1	64.3
CONTA [60] <sub>NIPS20</sub>	I	R101	66.1	66.7
AdvCAM [61] <sub>CVPR21</sub>	I	R101	68.1	68.0
CDA [85] <sub>ICCV21</sub>	I	WR38	66.1	66.8
ECS [75] <sub>ICCV21</sub>	I	WR38	66.6	67.6
RIB [62] <sub>NIPS21</sub>	I	R101	68.3	68.6
AMR [86] <sub>AAAI22</sub>	I	R101	68.8	69.1
MCTformer [87] <sub>CVPR22</sub>	I	WR38	71.9	71.6
AFA [88] <sub>CVPR22</sub>	I	MiT-B1	66.0	66.3
SIPE [89] <sub>CVPR22</sub>	I	R101	68.8	69.7
ReCAM [63] <sub>CVPR22</sub>	I	R101	68.5	68.4
PPC [81] <sub>CVPR22</sub>	I	R101	67.7	67.4
ViT-PCM [90] <sub>ECCV22</sub>	I	R101	70.3	70.9
Spatial-BCE [90] <sub>ECCV22</sub>	I	WR38	70.0	71.3
AEFT [91] <sub>ECCV22</sub>	I	WR38	70.9	71.7
ESOL [19] <sub>NIPS22</sub>	I	R101	69.9	69.3
TOCO [92] <sub>CVPR23</sub>	I	VIT-B	69.8	70.5
OCR [78] <sub>CVPR23</sub>	I	WR38	72.7	72.0
ACR [78] <sub>CVPR23</sub>	I	WR38	71.9	71.9
BECO [74] <sub>CVPR23</sub>	I	R101	72.1	71.8
D2CAM [76] <sub>ICCV23</sub>	I	R101	71.2	70.7
MARS [93] <sub>ICCV23</sub>	I	R101	70.3	71.2
<b>SSC (Ours)</b>	I	R101	<b>72.7</b>	<b>72.8</b>

TABLE V

QUANTITATIVE COMPARISONS TO PREVIOUS STATE-OF-THE-ART APPROACHES ON COCO VAL SET WITH VGG BACKBONE. SUP: SUPERVISION, I: IMAGE-LEVEL LABEL, S: SALIENCY MAPS

Methods	Sup	Backbone	Val
DSRG [14] <sub>CVPR18</sub>	I+S	VGG16	26.0
IAL [66] <sub>IJCV20</sub>	I+S	VGG16	27.7
GWSM [67] <sub>AAAI21</sub>	I+S	VGG16	28.4
EPS [35] <sub>CVPR21</sub>	I+S	VGG16	35.7
I2CRC [18] <sub>TMM22</sub>	I+S	VGG16	31.2
RCA [82] <sub>CVPR22</sub>	I+S	VGG16	36.8
MDBA [1] <sub>TIP23</sub>	I+S	VGG16	37.8
BFBP [94] <sub>ECCV16</sub>	I	VGG16	20.4
SEC [11] <sub>ECCV16</sub>	I	VGG16	22.4
CONTA [60] <sub>NIPS20</sub>	I	VGG16	23.7
<b>SSC (Ours)</b>	I	VGG16	<b>38.1</b>

pretraining models with super-large datasets. Some example prediction maps on the COCO val set can be viewed in Fig. 4.

#### D. Ablation Studies

1) *Element-Wise Component Analysis*: In this part, we demonstrate the contribution of each component proposed in our approach to improving the quality of pseudo-masks.






Fig. 3. Example prediction maps on the PASCAL VOC 2012 val set. For each (a) image, we show (b) the ground truth (GT), predictions of (c) baseline, (d) DRS [24], (e) DRS + CDR, and (f) DRS + CDR + ASM. Best viewed in color.

TABLE VI

QUANTITATIVE COMPARISONS TO PREVIOUS STATE-OF-THE-ART APPROACHES ON COCO VAL SET WITH RESNET BACKBONE. SUP: SUPERVISION, I: IMAGE-LEVEL LABEL, S: SALIENCY MAPS, L: LANGUAGE SUPERVISION

Methods	Sup	Backbone	Val
AuxSegNet [83] <sub>ICCV21</sub>	I+S	WR38	33.9
L2G [36] <sub>CVPR22</sub>	I+S	R101	44.2
CLIP-ES [80] <sub>CVPR23</sub>	I+L	R101	45.4
LPCAM [84] <sub>CVPR23</sub>	I+S	R101	42.1
IRN [38] <sub>CVPR19</sub>	I	R101	41.4
SEAM [95] <sub>CVPR20</sub>	I	WR38	31.9
CONTA [60] <sub>NIPS20</sub>	I	R50	33.4
PMM [96] <sub>ICCV21</sub>	I	WR38	36.7
OC-CSE [97] <sub>ICCV21</sub>	I	WR38	36.4
CDA [85] <sub>ICCV21</sub>	I	WR38	33.2
RIB [62] <sub>NIPS21</sub>	I	R101	43.8
MCTformer [87] <sub>CVPR22</sub>	I	WR38	42.0
SIPE [89] <sub>CVPR22</sub>	I	R101	40.6
ESOL [19] <sub>NIPS22</sub>	I	R101	42.6
TOCO [92] <sub>CVPR23</sub>	I	VIT-B	41.3
OCR [78] <sub>CVPR23</sub>	I	WR38	42.5
ACR [78] <sub>CVPR23</sub>	I	WR38	45.3
BECO [74] <sub>CVPR23</sub>	I	R101	45.1
D2CAM [76] <sub>ICCV23</sub>	I	R101	44.0
<b>SSC (Ours)</b>	I	R101	<b>47.0</b>

The experimental results are given in Table VII. While the accuracy of pseudo-labels from baseline is 48.3%, DRS [24] can improve the result to 50.2%. With our proposed CAM-driven reconstruction module, we can significantly improve the quality of pseudo-masks and the accuracy arrives at 54.5%. We can notice that if we reconstruct the input image from early features, the accuracy will drop to 52.3%. This highlights the importance of directly reconstructing the input image from its CAM features to preserve the spatial structure in attention maps. Some visualization of the image reconstruction

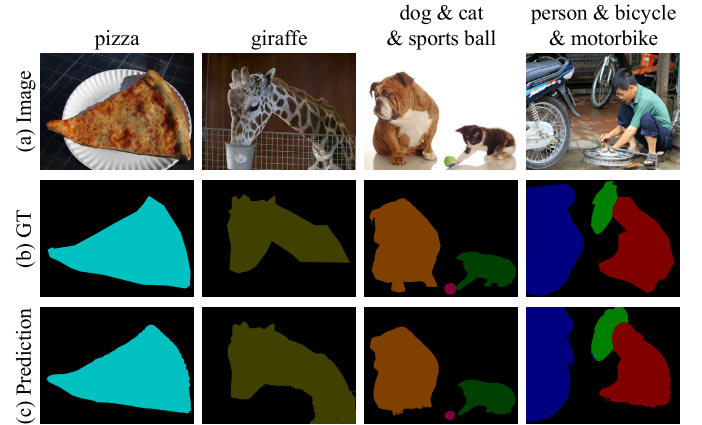


Fig. 4. Example prediction maps on the COCO val set. For each (a) image, we show (b) the ground truth (GT), and (c) prediction. Best viewed in color.

TABLE VII

ELEMENT-WISE COMPONENT ANALYSIS. THE ACCURACY (mIoU) OF PSEUDO-MASKS ON PASCAL VOC 2012 TRAINING SET IS REPORTED. DRS [24]: DISCRIMINATIVE REGION SUPPRESSION; CDR: CAM-DRIVEN RECONSTRUCTION; ASM: ACTIVATION SELF-MODULATION; RAS: RELIABLE ACTIVATION SELECTION

Methods	mIoU
baseline	48.3
baseline + DRS	50.2
baseline + DRS + CDR early	52.3
baseline + DRS + CDR	54.5
baseline + DRS + CDR + ASM w/o RAS	57.2
baseline + DRS + CDR + ASM	<b>58.3</b>

results can be observed in Fig. 6. As can be seen, the reconstructed results with the perceptual loss focus more on the content and spatial structure of images rather than low-level pixel information. With our activation self-modulation module,



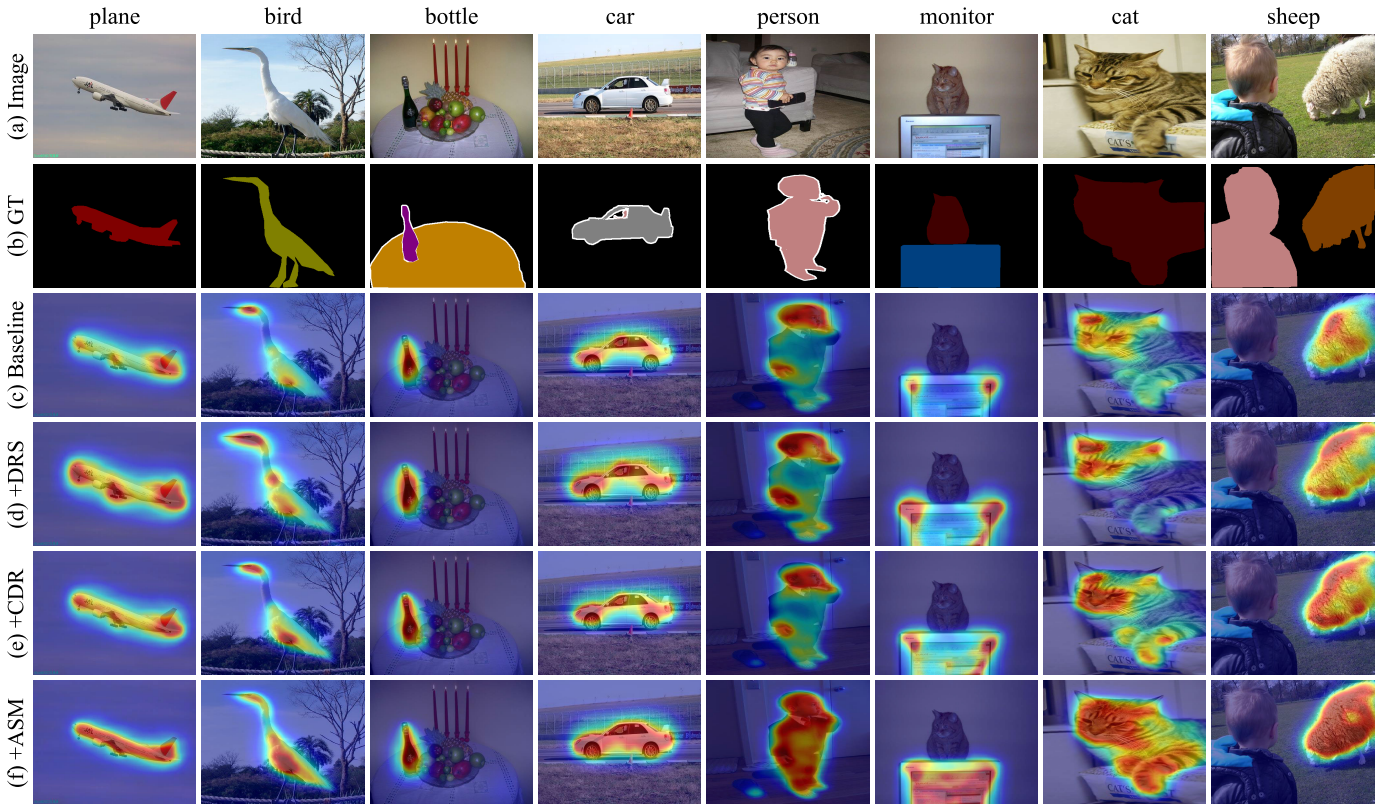


Fig. 5. Example localization maps on the PASCAL VOC 2012 training set. For each (a) image, we show (b) the ground truth (GT), localization maps produced by (c) baseline, (d) DRS [24], (e) DRS + CDR, and (f) DRS + CDR + ASM. Best viewed in color.

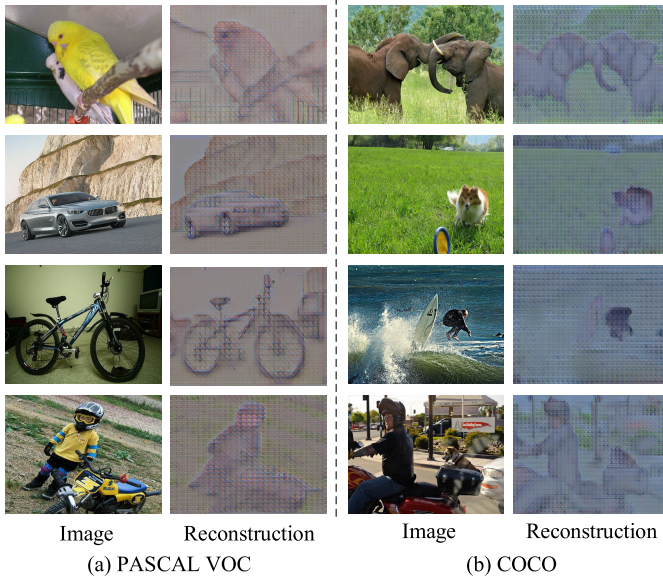


Fig. 6. Example image reconstruction results on the PASCAL VOC 2012 and COCO datasets. Best viewed in color.

we further constrain the localization maps with detailed spatial structure information and improve the accuracy of pseudo-masks to 58.3%. As can be seen, without our reliable activation selection strategy, the accuracy will decline to 57.2%. This highlights the importance of maintaining the high activation of the most discriminative regions to alleviate the false modulation in the opposite direction.

TABLE VIII

COMPARISON OF PERCEPTUAL LOSS WITH L1 AND L2 LOSSES FOR CAM-DRIVEN RECONSTRUCTION (CDR). THE ACCURACY (mIoU) OF PSEUDO-MASKS ON THE PASCAL VOC 2012 TRAINING SET IS REPORTED

Methods	mIoU
w/o CDR	50.2
w/ CDR + L1 loss	51.3
w/ CDR + L2 loss	50.9
w/ CDR + perceptual loss	<b>54.5</b>

Some example localization maps on the PASCAL VOC 2012 training set can be viewed in Fig. 5. As can be seen from the first two columns, though DRS [24] can help enlarge the activated region, it will also inevitably result in over-activation where the expanded object activation intrudes into the background area. With our proposed CAM-driven reconstruction module, we can help constrain the activation within the object region (e.g., the plane in the first column and the beak of bird in the second column). Our proposed activation self-modulation module further constrains the network attention on target objects by enhancing regional consistency, which leads to more compact object activation. In addition, benefiting from our spatial structure constraints, our proposed CAM-driven reconstruction and activation self-modulation modules can also help activate more integral object regions to alleviate the remaining under-activation problem in difficult images, illustrated in the last column. Some example prediction maps on the PASCAL VOC 2012 val set can be viewed in Fig. 3.




Fig. 7. The analysis of parameter sensitivity. The accuracy (mIoU) of pseudo-masks on the PASCAL VOC 2012 training set is reported.

2) *Discussion of Perceptual Loss*: For driving the classification network to preserve the spatial structure of image content, we propose a CAM-driven reconstruction module to recover the input image. Then the perceptual loss is adopted to penalize the reconstructed image that semantically and spatially deviates from the original image. The motivation for selecting perceptual loss rather than L1 or L2 losses which are also widely used in self-supervised learning is twofold. First, L1 and L2 losses highlight the alignment of low-level information. Such strict per-pixel alignment will encourage the network to focus more on color and texture than image content and spatial structure. Second, image reconstruction with L1 or L2 loss requires the network to preserve more subtle information than perceptual loss, which might exceed the capacity of CAM features and thus deteriorate its localization ability. The qualitative comparison of perceptual loss with L1 and L2 losses for CAM-driven reconstruction is demonstrated in Table VIII. As can be seen, only about 1% performance gain is achieved with L1 or L2 losses. In contrast, with the perceptual loss, our proposed CAM-driven reconstruction can improve the accuracy from 50.2% to 54.5%.

3) *Parameter Analysis*: For the CAM-driven reconstruction module, we conduct experiments to study the effect of perceptual loss. As shown in Fig. 7 (a), we vary the weight  $\beta_p$  over the range {0.1, 0.5, 1, 2, 5}. As we can see, we get better performance when the weight is between 1 and 2. A too-small  $\beta_p$  may not improve the results very much, and a large  $\beta_p$  deteriorates the performance seriously. We conjecture that increasing the importance of perceptual loss will make the network overemphasize reconstruction results and weaken the localization ability. Similarly, we vary  $\beta_a$  with the same range to study the effect of the activation self-modulation module. Though more stable performance can be achieved while  $\beta_a$  is between 0.1 and 2, the results drop significantly with a large weight as well. In our experiments, we empirically set  $\beta_p = 1$  and  $\beta_a = 1$ .

For the reliable activation selection strategy, the object threshold  $T_{obj}$  will affect the parameter of the erosion operation. Therefore, we first empirically fix  $T_{obj} = 0.3$  to locate the area with relatively higher attention values, which is a threshold adopted in OAA+ [17] and NSROM [77] as the foreground object threshold. Then we conduct experiments to study the effect of the erosion kernel size  $r$ . As shown in Fig. 7 (b), we vary the kernel size  $r$  over the range {2, 4, 6, 8, 10, 12}. As we can see, we get better performance when the kernel size is between 6 and 10. A too-large or small kernel size may not improve the results very much. We conjecture that a too-small kernel size keeps too much

TABLE IX  
GPU CONSUMPTION ANALYSIS. DRS [24]: DISCRIMINATIVE REGION SUPPRESSION; CDR: CAM-DRIVEN RECONSTRUCTION; ASM: ACTIVATION SELF-MODULATION

Methods	GPU
baseline	
baseline + DRS	3.2G
baseline + DRS + CDR	10.6G
baseline + DRS + CDR + ASM	14.2G




Fig. 8. Failure cases on PASCAL VOC 2012 val set. Best viewed in color.

high activation in the CAM and mistakes the background as foreground, which will deteriorate the CAM modulation. Meanwhile, a too-large kernel size keeps less object activation, which dilutes the effect of reliable activation selection. In our experiments, we empirically set  $T_{obj} = 0.3$  and  $r = 8$ .

#### E. Limitation and Failure Cases

1) *GPU Consumption*: As can be seen in Table IX, compared to the parameter-free DRS [24], the introduction of our CAM-driven reconstruction module will increase the GPU consumption from 3.2G to 10.6G. However, our full approach consumes only 14.2G GPU memory, which can be trained with a single 16G or 24G GPU.

2) *Failure Cases*: Though our proposed SSC has achieved great success in the WSSS task, it still faces difficulties in some challenging scenes. Some typical failure cases are demonstrated in Fig. 8. First, our approach may fail to tackle the co-occurrence problem. For example, in the 1st row of Fig. 8, our model fails to distinguish the train from the railway background since these two categories often co-occur in images. Similarly, our method will treat the keyboard as the part of the monitor in the 2nd row. In addition, our model sometimes cannot recognize objects of different categories with similar appearances, e.g., it mistakes the person in the last row for the motorcycle. In the future, we would like to explore the solutions to the above challenges.

#### V. CONCLUSION

In this work, we proposed spatial structure constraints for weakly supervised semantic segmentation to alleviate the object over-activation problem during attention expansion.



Specifically, we proposed a CAM-driven reconstruction module that directly reconstructs the input image from its CAM features. A perceptual loss was adopted to encourage the classification backbone to preserve the coarse spatial structure of the image content. Moreover, we proposed an activation self-modulation module to further refine CAMs with finer spatial structure details through enhancing regional consistency. Our proposed approach can help activate more integral target areas and constrain the activation within object regions. Extensive experiments on the PASCAL VOC 2012 and COCO datasets demonstrated the superiority of our proposed approach.

## REFERENCES

- [1] T. Chen, Y. Yao, and J. Tang, "Multi-granularity denoising and bidirectional alignment for weakly supervised semantic segmentation," *IEEE Trans. Image Process.*, vol. 32, pp. 2960–2971, 2023.
- [2] H. Ding, X. Jiang, B. Shuai, A. Q. Liu, and G. Wang, "Semantic segmentation with context encoding and multi-path decoding," *IEEE Trans. Image Process.*, vol. 29, pp. 3520–3533, 2020.
- [3] J.-Y. He, S.-H. Liang, X. Wu, B. Zhao, and L. Zhang, "MGSeg: Multiple granularity-based real-time semantic segmentation network," *IEEE Trans. Image Process.*, vol. 30, pp. 7200–7214, 2021.
- [4] J. Dai, K. He, and J. Sun, "BoxSup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1635–1643.
- [5] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele, "Simple does it: Weakly supervised instance and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2017, pp. 876–885.
- [6] L. Jing, Y. Chen, and Y. Tian, "Coarse-to-fine semantic segmentation from image-level labels," *IEEE Trans. Image Process.*, vol. 29, pp. 225–236, 2020.
- [7] C. Song, Y. Huang, W. Ouyang, and L. Wang, "Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation," in *Proc. CVPR*, Jun. 2019, pp. 3136–3145.
- [8] A. Bearman, O. Russakovsky, V. Ferrari, and L. Fei-Fei, "What's the point: Semantic segmentation with point supervision," in *Proc. ECCV*, 2016, pp. 549–565.
- [9] D. Lin, J. Dai, J. Jia, K. He, and J. Sun, "ScribbleSup: Scribble-supervised convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3159–3167.
- [10] P. Vernaza and M. Chandraker, "Learning random-walk label propagation for weakly-supervised semantic segmentation," in *Proc. CVPR*, Jul. 2017, pp. 7158–7166.
- [11] A. Kolesnikov and C. Lampert, "Seed, expand and constrain: Three principles for weakly-supervised image segmentation," in *Proc. ECCV*, 2016, pp. 695–711.
- [12] G. Pei, F. Shen, Y. Yao, G.-S. Xie, Z. Tang, and J. Tang, "Hierarchical feature alignment network for unsupervised video object segmentation," in *Proc. ECCV*, 2022, pp. 596–613.
- [13] T. Chen et al., "Semantically meaningful class prototype learning for one-shot image segmentation," *IEEE Trans. Multimedia*, vol. 24, pp. 968–980, 2022.
- [14] Z. Huang, X. Wang, J. Wang, W. Liu, and J. Wang, "Weakly-supervised semantic segmentation network with deep seeded region growing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7014–7023.
- [15] J. Ahn and S. Kwak, "Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4981–4990.
- [16] Y. Wei, H. Xiao, H. Shi, Z. Jie, J. Feng, and T. S. Huang, "Revisiting dilated convolution: A simple approach for Weakly- and semi-supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7268–7277.
- [17] P. Jiang, Q. Hou, Y. Cao, M. Cheng, Y. Wei, and H. Xiong, "Integral object mining via online attention accumulation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 2070–2079.
- [18] T. Chen, Y. Yao, L. Zhang, Q. Wang, G. Xie, and F. Shen, "Saliency guided inter- and intra-class relation constraints for weakly supervised semantic segmentation," *IEEE Trans. Multimedia*, vol. 25, pp. 1727–1737, 2022.
- [19] J. Li, Z. Jie, X. Wang, X. Wei, and L. Ma, "Expansion and shrinkage of localization for weakly-supervised semantic segmentation," in *Proc. NeurIPS*, vol. 35, 2022, pp. 16037–16051.
- [20] W. Xu, H. Huang, M. Cheng, L. Yu, Q. Wu, and J. Zhang, "Masked cross-image encoding for few-shot segmentation," in *Proc. IEEE Int. Conf. Multimedia Expo (ICME)*, Jul. 2023, pp. 744–749.
- [21] H. Huang et al., "PADDLES: Phase-amplitude spectrum disentangled early stopping for learning with noisy labels," in *Proc. ICCV*, Oct. 2023, pp. 16719–16730.
- [22] H. Huang, J. Zhang, J. Zhang, J. Xu, and Q. Wu, "Low-rank pairwise alignment bilinear network for few-shot fine-grained image classification," *IEEE Trans. Multimedia*, vol. 23, pp. 1666–1680, 2021.
- [23] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. CVPR*, Jun. 2016, pp. 2921–2929.
- [24] B. Kim, S. Han, and J. Kim, "Discriminative region suppression for weakly-supervised semantic segmentation," in *Proc. AAAI*, 2021, vol. 35, no. 2, pp. 1754–1761.
- [25] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, Jul. 2006.
- [26] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. ECCV*, 2016, pp. 694–711.
- [27] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, Apr. 2018.
- [28] T. Zhou, W. Wang, E. Konukoglu, and L. Van Gool, "Rethinking semantic segmentation: A prototype view," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 2582–2593.
- [29] W. Wang, T. Zhou, F. Yu, J. Dai, E. Konukoglu, and L. V. Gool, "Exploring cross-image pixel contrast for semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 7303–7313.
- [30] L. Li, W. Wang, and Y. Yang, "LogicSeg: Parsing visual semantics with neural logic learning and reasoning," in *Proc. ICCV*, Oct. 2023, pp. 4122–4133.
- [31] Y. Wei, J. Feng, X. Liang, M.-M. Cheng, Y. Zhao, and S. Yan, "Object region mining with adversarial erasing: A simple classification to semantic segmentation approach," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1568–1576.
- [32] Y. Tu et al., "Learning with noisy labels via self-supervised adversarial noisy masking," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 16186–16195.
- [33] T. Zhou, L. Li, X. Li, C.-M. Feng, J. Li, and L. Shao, "Group-wise learning for weakly supervised semantic segmentation," *IEEE Trans. Image Process.*, vol. 31, pp. 799–811, 2022.
- [34] W. Wang, G. Sun, and L. Van Gool, "Looking beyond single images for weakly supervised semantic segmentation learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022, doi: 10.1109/TPAMI.2022.3168530.
- [35] S. Lee, M. Lee, J. Lee, and H. Shim, "Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 5495–5505.
- [36] P.-T. Jiang, Y. Yang, Q. Hou, and Y. Wei, "L2G: A simple local-to-global knowledge transfer framework for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 16886–16896.
- [37] C. Zhao, Y. Tu, Z. Lai, F. Shen, H. T. Shen, and D. Miao, "Saliency-guided iterative asymmetric mutual hashing for fast person re-identification," *IEEE Trans. Image Process.*, vol. 30, pp. 7776–7789, 2021.
- [38] J. Ahn, S. Cho, and S. Kwak, "Weakly supervised learning of instance segmentation with inter-pixel relations," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 2209–2218.
- [39] L. Chen, W. Wu, C. Fu, X. Han, and Y. Zhang, "Weakly supervised semantic segmentation with boundary exploration," in *Proc. ECCV*, 2020, pp. 347–362.
- [40] L. Jing and Y. Tian, "Self-supervised visual feature learning with deep neural networks: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 11, pp. 4037–4058, Nov. 2021.



- [41] T. Chen, S.-H. Wang, Q. Wang, Z. Zhang, G.-S. Xie, and Z. Tang, "Enhanced feature alignment for unsupervised domain adaptation of semantic segmentation," *IEEE Trans. Multimedia*, vol. 24, pp. 1042–1054, 2022.
- [42] J. Gui et al., "A survey on self-supervised learning: Algorithms, applications, and future trends," 2023, *arXiv:2301.05712*.
- [43] C. Doersch, A. Gupta, and A. A. Efros, "Unsupervised visual representation learning by context prediction," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1422–1430.
- [44] M. Noroozi and P. Favaro, "Unsupervised learning of visual representations by solving jigsaw puzzles," in *Proc. ECCV*, 2016, pp. 69–84.
- [45] N. Komodakis and S. Gidaris, "Unsupervised representation learning by predicting image rotations," in *Proc. ICLR*, 2018, pp. 71–84.
- [46] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 9729–9738.
- [47] Y. Yao et al., "Jo-SRC: A contrastive approach for combating noisy labels," in *Proc. CVPR*, Jun. 2021, pp. 5192–5201.
- [48] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *Proc. ICML*, 2020, pp. 1597–1607.
- [49] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. E. Hinton, "Big self-supervised models are strong semi-supervised learners," in *Proc. NeurIPS*, vol. 33, 2020, pp. 22243–22255.
- [50] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 16000–16009.
- [51] Z. Xie et al., "SimMIM: A simple framework for masked image modeling," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 9653–9663.
- [52] B. V. Kumar, "An extensive survey on superpixel segmentation: A research perspective," *Arch. Comput. Methods Eng.*, vol. 30, pp. 3749–3767, Apr. 2023.
- [53] F. Z. Xing, E. Cambria, W.-B. Huang, and Y. Xu, "Weakly supervised semantic segmentation with superpixel embedding," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2016, pp. 1269–1273.
- [54] Y. Li, Y. Liu, G. Liu, and M. Guo, "Weakly supervised semantic segmentation by iterative superpixel-CRF refinement with initial clues guiding," *Neurocomputing*, vol. 391, pp. 25–41, May 2020.
- [55] J. Fan, Z. Zhang, C. Song, and T. Tan, "Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 4283–4292.
- [56] S. Yi, H. Ma, X. Wang, T. Hu, X. Li, and Y. Wang, "Weakly-supervised semantic segmentation with superpixel guided local and global consistency," *Pattern Recognit.*, vol. 124, Apr. 2022, Art. no. 108504.
- [57] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.
- [58] X. Zhang, Y. Wei, J. Feng, Y. Yang, and T. Huang, "Adversarial complementary learning for weakly supervised object localization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1325–1334.
- [59] W. Liu, C. Zhang, G. Lin, T.-Y. Hung, and C. Miao, "Weakly supervised segmentation with maximum bipartite graph matching," in *Proc. 28th ACM Int. Conf. Multimedia*, Oct. 2020, pp. 2085–2094.
- [60] D. Zhang, H. Zhang, J. Tang, X.-S. Hua, and Q. Sun, "Causal intervention for weakly-supervised semantic segmentation," in *Proc. NeurIPS*, vol. 33, 2020, pp. 1–12.
- [61] J. Lee, E. Kim, and S. Yoon, "Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 4071–4080.
- [62] J. Lee, J. Choi, J. Mok, and S. Yoon, "Reducing information bottleneck for weakly supervised semantic segmentation," in *Proc. NeurIPS*, vol. 34, 2021, pp. 27408–27421.
- [63] Z. Chen, T. Wang, X. Wu, X.-S. Hua, H. Zhang, and Q. Sun, "Class re-activation maps for weakly-supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 969–978.
- [64] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, Jun. 2010.
- [65] B. Hariharan, P. Arbeláez, L. D. Bourdev, S. Maji, and J. Malik, "Semantic contours from inverse detectors," in *Proc. ICCV*, Nov. 2011, pp. 991–998.
- [66] X. Wang, S. Liu, H. Ma, and M.-H. Yang, "Weakly-supervised semantic segmentation by iterative affinity learning," *Int. J. Comput. Vis.*, vol. 128, no. 6, pp. 1736–1749, Jun. 2020.
- [67] X. Li, T. Zhou, J. Li, Y. Zhou, and Z. Zhang, "Group-wise semantic mining for weakly supervised semantic segmentation," in *Proc. AAAI*, 2021, pp. 1984–1992.
- [68] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [69] J. Deng et al., "ImageNet: A large-scale hierarchical image database," in *Proc. CVPR*, Jun. 2009, pp. 248–255.
- [70] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proc. COMPSTAT*, 2010, pp. 177–186.
- [71] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015, pp. 1–14.
- [72] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *Int. J. Comput. Vis.*, vol. 59, no. 2, pp. 167–181, Sep. 2004.
- [73] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *Int. J. Comput. Vis.*, vol. 104, no. 2, pp. 154–171, 2013.
- [74] S. Rong, B. Tu, Z. Wang, and J. Li, "Boundary-enhanced co-training for weakly supervised semantic segmentation," in *Proc. CVPR*, Jun. 2023, pp. 19574–19584.
- [75] K. Sun, H. Shi, Z. Zhang, and Y. Huang, "ECS-Net: Improving weakly supervised semantic segmentation by using connections between class activation maps," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 7283–7292.
- [76] C. Wang, R. Xu, S. Xu, W. Meng, and X. Zhang, "Treating pseudo-labels generation as image matting for weakly supervised semantic segmentation," in *Proc. ICCV*, Oct. 2023, pp. 755–765.
- [77] Y. Yao et al., "Non-salient region object mining for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 2623–2632.
- [78] Z. Cheng et al., "Out-of-candidate rectification for weakly supervised semantic segmentation," in *Proc. CVPR*, Jun. 2023, pp. 23673–23684.
- [79] J. Lee, S. J. Oh, S. Yun, J. Choe, E. Kim, and S. Yoon, "Weakly supervised semantic segmentation using out-of-distribution data," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 16897–16906.
- [80] Y. Lin et al., "CLIP is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation," in *Proc. CVPR*, Jun. 2023, pp. 15305–15314.
- [81] Y. Du, Z. Fu, Q. Liu, and Y. Wang, "Weakly supervised semantic segmentation by pixel-to-prototype contrast," in *Proc. CVPR*, Jun. 2022, pp. 4320–4329.
- [82] T. Zhou, M. Zhang, F. Zhao, and J. Li, "Regional semantic contrast and aggregation for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 4299–4309.
- [83] L. Xu, W. Ouyang, M. Bennamoun, F. Boussaid, F. Sohel, and D. Xu, "Leveraging auxiliary tasks with affinity learning for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 6984–6993.
- [84] Z. Chen and Q. Sun, "Extracting class activation maps from non-discriminative features as well," in *Proc. CVPR*, Jun. 2023, pp. 3135–3144.
- [85] Y. Su, R. Sun, G. Lin, and Q. Wu, "Context decoupling augmentation for weakly supervised semantic segmentation," in *Proc. ICCV*, Oct. 2021, pp. 7004–7014.
- [86] J. Qin, J. Wu, X. Xiao, L. Li, and X. Wang, "Activation modulation and recalibration scheme for weakly supervised semantic segmentation," in *Proc. AAAI*, 2022, vol. 36, no. 2, pp. 2117–2125.
- [87] L. Xu, W. Ouyang, M. Bennamoun, F. Boussaid, and D. Xu, "Multi-class token transformer for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 4310–4319.
- [88] L. Ru, Y. Zhan, B. Yu, and B. Du, "Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 16846–16855.
- [89] Q. Chen, L. Yang, J. Lai, and X. Xie, "Self-supervised image-specific prototype exploration for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 4288–4298.

- [90] S. Rossetti, D. Zappia, M. Sanzari, M. Schaerf, and F. Pirri, "Max pooling with vision transformers reconciles class and shape in weakly supervised semantic segmentation," in *Proc. ECCV*, 2022, pp. 446–463.
- [91] S.-H. Yoon, H. Kweon, J. Cho, S. Kim, and K.-J. Yoon, "Adversarial erasing framework via triplet with gated pyramid pooling layer for weakly supervised semantic segmentation," in *Proc. ECCV*, 2022, pp. 326–344.
- [92] L. Ru, H. Zheng, Y. Zhan, and B. Du, "Token contrast for weakly-supervised semantic segmentation," in *Proc. CVPR*, Jun. 2023, pp. 3093–3102.
- [93] S. Jo, I.-J. Yu, and K. Kim, "MARS: Model-agnostic biased object removal without additional supervision for weakly-supervised semantic segmentation," in *Proc. ICCV*, 2023, pp. 719–729.
- [94] F. Saleh et al., "Built-in foreground/background prior for weakly-supervised semantic segmentation," in *Proc. ECCV*, 2016, pp. 413–432.
- [95] Y. Wang, J. Zhang, M. Kan, S. Shan, and X. Chen, "Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 12275–12284.
- [96] Y. Li, Z. Kuang, L. Liu, Y. Chen, and W. Zhang, "Pseudo-mask matters in weakly-supervised semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 6964–6973.
- [97] H. Kweon, S.-H. Yoon, H. Kim, D. Park, and K.-J. Yoon, "Unlocking the potential of ordinary classifier: Class-specific adversarial erasing framework for weakly supervised semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 6994–7003.



**Tao Chen** (Member, IEEE) received the Ph.D. degree in pattern recognition and intelligent systems from the Nanjing University of Science and Technology, Nanjing, China, in 2022. He is currently a Postdoctoral Researcher with the School of Computer Science and Engineering, Nanjing University of Science and Technology. His research interests include weakly supervised semantic segmentation, domain adaptive semantic segmentation, and one-shot image segmentation.



**Yazhou Yao** (Member, IEEE) received the Ph.D. degree in computer science from the University of Technology Sydney, Australia, in 2018. From July 2018 to July 2019, he was a Research Scientist with the Inception Institute of Artificial Intelligence, Abu Dhabi, United Arab Emirates. He is currently a Professor with the School of Computer Science and Engineering, Nanjing University of Science and Technology. His research interests include multimedia processing and machine learning. He received the support of the China Scholarship Council for the Ph.D. degree.



Professor with the College of Instrumentation and Electrical Engineering, Jilin University.



ICIMCS in 2018. He serves as an Associate Editor for IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEM and *Information Sciences*.

**Xingguo Huang** received the bachelor's and master's degrees in geophysics from Jilin University, Changchun, China, in 2014 and 2017, respectively, and the Ph.D. degree in geophysics from the University of Bergen, Bergen, Norway, in 2020. He was a Visiting Ph.D. Student with the University of California at Santa Cruz, Santa Cruz, CA, USA, from July 2018 to January 2019 and from June to August 2019. He was a Research Fellow with the University of Bergen, Bergen, Norway, from May 2020 to September 2020. He is currently a

**Zechao Li** (Senior Member, IEEE) received the B.E. degree from the University of Science and Technology of China, China, in 2008, and the Ph.D. degree from the National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, in 2013. He is currently a Professor with the Nanjing University of Science and Technology. His research interests include big media analysis and computer vision. He was a recipient of the Best Paper Award from the ACM Multimedia Asia 2020 and the Best Student Paper Award from the



Mention in 2019, the SIGMM Rising Star in 2020, and the SIGIR Best Student Paper in 2021. He is the Area Chair of ACM MM 2018–2023. He is an Associate Editor of IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, IEEE TRANSACTIONS ON MULTIMEDIA, *ACM ToMM*, and *Information Sciences*.

**Liqiang Nie** (Senior Member, IEEE) received the B.E. degree from Xi'an Jiaotong University and the Ph.D. degree from the National University of Singapore (NUS). After the Ph.D. degree, he continued his research with NUS, as a Research Fellow for three years. He is currently a Professor with the Harbin Institute of Technology (Shenzhen). He has coauthored more than 200 articles and four books. His research interests include multimedia computing and information retrieval. He received many awards, such as ACM MM and SIGIR Best Paper Honorable



Conference on Multimedia (PCM) in 2011, the ACM International Conference on Internet Multimedia Computing and Service (ICIMCS) in 2011, and the ACM MM Asia 2020, the Best Paper Runner-Up Award from ACM MM 2015, and the Best Student Paper Awards from the International Conference on Multimedia Modeling (MMM) in 2016 and ICIMCS 2017. He served as an Associate Editor for IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, and IEEE TRANSACTIONS ON MULTIMEDIA.

**Jinhui Tang** (Senior Member, IEEE) received the B.E. and Ph.D. degrees from the University of Science and Technology of China, Hefei, China, in 2003 and 2008, respectively. He is currently a Professor with the Nanjing University of Science and Technology, Nanjing, China. He has authored more than 150 articles in top-tier journals and conferences. His research interests include multimedia analysis and computer vision. He was a recipient of the Best Paper Awards from the ACM International Conference on Multimedia (MM) in 2007, the Pacific-Rim