

Non-Salient Region Object Mining for Weakly Supervised Semantic Segmentation

Yazhou Yao^{1*}, Tao Chen^{1*}, Guo-Sen Xie^{2†}, Chuanyi Zhang¹,
Fumin Shen³, Qi Wu⁴, Zhenmin Tang¹, Jian Zhang⁵

¹Nanjing University of Science and Technology; ²Mohamed bin Zayed University of AI

³University of Electronic Science and Technology of China; ⁴University of Adelaide;

⁵University of Technology Sydney

Abstract

Semantic segmentation aims to classify every pixel of an input image. Considering the difficulty of acquiring dense labels, researchers have recently been resorting to weak labels to alleviate the annotation burden of segmentation. However, existing works mainly concentrate on expanding the seed of pseudo labels within the image’s salient region. In this work, we propose a non-salient region object mining approach for weakly supervised semantic segmentation. We introduce a graph-based global reasoning unit to strengthen the classification network’s ability to capture global relations among disjoint and distant regions. This helps the network activate the object features outside the salient area. To further mine the non-salient region objects, we propose to exert the segmentation network’s self-correction ability. Specifically, a potential object mining module is proposed to reduce the false-negative rate in pseudo labels. Moreover, we propose a non-salient region masking module for complex images to generate masked pseudo labels. Our non-salient region masking module helps further discover the objects in the non-salient region. Extensive experiments on the PASCAL VOC dataset demonstrate state-of-the-art results compared to current methods. The source codes are available at <https://github.com/NUST-Machine-Intelligence-Laboratory/nsrom>.

1. Introduction

Semantic segmentation is the task of classifying every pixel of an input image. It plays a vital role in many computer vision tasks, such as image editing and medical image analysis [32, 47]. Benefiting from the recent advances of deep learning, semantic segmentation has achieved remarkable progress. However, the training of deep convolu-




Figure 1. Comparison between the traditional methods and ours. (a) Input image. (b) The saliency map. (c) Results of the traditional methods. They mainly focus on expanding the seed of pseudo labels within the salient region of the image. They only obtain good segmentation results in conspicuous regions. (d) Our results. Our method also mines the object in the non-salient region and can get better results both in and outside the salient region. Best viewed in color.

tional neural networks (CNNs) usually requires large-scale datasets [56, 55, 59, 57, 58]. Moreover, obtaining precise pixel-wise annotations for semantic segmentation demands intensive labor efforts and is quite time-consuming. One promising approach to address the annotation problem for semantic segmentation is to learn from weak labels, such as image-level annotations [29, 51, 21, 8, 23, 2, 52, 24], bounding boxes [14, 26, 43], points [4], and scribbles [33, 46]. Among these weak supervisions, image-level labels are the easiest format to annotate and have been widely studied in various weakly supervised methods. However, semantic segmentation supervised with image-level labels remains a challenging task. Therefore, this paper follows the current trend and focuses on leveraging image-level labels to achieve weakly supervised semantic segmentation (WSSS).

To tackle the task of WSSS with only image-level labels, visualization-based approaches [63] have been widely

*Equal contribution.

†Corresponding author.

adopted to narrow the annotation gap between classification and segmentation [53, 54]. The typical methods train a classification network with image-level labels. Then they leverage class activation maps (CAMs) [63] to generate pseudo labels to train the segmentation network. However, these activation maps obtained from the classification network are sparse and incomplete. They can only locate the most discriminative part of objects. Many approaches have been proposed to enlarge the activated region to cover a large object area. For example, Jiang *et al.* [24] observed that the attention maps produced by the classification network focus on different object parts during training. Therefore, they proposed an online attention accumulation (OAA) strategy to combine the various activated regions. However, as shown in Fig. 1, the existing works mainly concentrate on enlarging the response maps for the salient region. Then they utilize the saliency maps to extract background. Few works focus on mining objects in the non-salient areas.

In this paper, we propose a non-salient region object mining method for weakly supervised semantic segmentation to make up for the shortcomings mentioned above. In contrast to the widely adopted center prior [5] for saliency detection, the non-salient region is usually scattered in corners or near the edge of the image. Such a characteristic of our protagonist requires the network to exploit the disjoint and distant surrounding information. While the traditional classification networks based on CNNs excel at modeling local relations, they are inefficient at capturing global relations between disjoint and distant regions. Therefore, we introduce a graph-based global reasoning unit [13] to strengthen the classification network’s capability in activating the object features outside the salient region.

On the other hand, though existing approaches can successfully enlarge activated regions for objects, they inevitably extend the object area to the background. These methods require the saliency maps to provide background clues. While the saliency maps can correct the pixel labels near conspicuous regions, they also remove the object labels outside the salient area. We notice that although the naive CAM, sparse and incomplete, does not have an accurate boundary, it can provide useful clues for the objects in the non-salient region. Therefore, we propose a potential object mining module to discover more objects that are outside the conspicuous region but activated in the naive CAM. Our potential object mining module aims to reduce the pseudo labels’ false-negative rate (in which case the object regions are falsely labeled as background). This improves the quality of pseudo labels and encourages the segmentation network to exert its self-correction ability. Such an ability of the network inspires us to further take advantage of the prediction of the segmentation network. Following [51], we divide the training images into simple and complex sets according to the number of categories in each image. The sim-

ple images with a single category of object(s) usually have a clean background. Their objects often exist in the conspicuous region and can be correctly segmented. In contrast, complex images (having two or more categories of objects) are more prone to having objects outside the salient area. Therefore, we propose a non-salient region masking module for complex images to generate masked pseudo labels. Our non-salient region masking module helps further discover objects in the non-salient region. Our contributions can be summarized as follows:

- For weakly supervised semantic segmentation, we leverage a global reasoning unit to capture global relations among disjoint and distant regions, helping the network activate object features outside salient areas.
- We propose a potential object mining module to discover more objects in the non-salient region, which improves the quality of pseudo labels by reducing the false-negative rate.
- We propose a non-salient region masking module with a dilation policy to generate masked pseudo labels, which leads to a more robust segmentation model to further discover objects outside the salient region.

2. Related Work

2.1. Semantic Segmentation

Semantic segmentation is an important computer vision task that assigns a semantic label to every pixel in an image. Since the adaption of the modern classification network into the fully convolutional network (FCN) [36, 45], deep learning has achieved great success in semantic segmentation [3, 10, 61, 35, 12, 11, 37]. To address the size issue caused by the down-sampling operation, early works [3] adopted an encoder-decoder architecture to recover the spatial resolution. Then dilated/atrous convolution [10] was proposed for the expansion of the receptive field without loss of resolution. Recently, the pyramid pooling module and context encoding [61] were introduced to capture the global semantic context of the scene. Auto-DeepLab [35] presented a network-level search space to allow efficient gradient-based architecture search for semantic segmentation.

2.2. Weakly Supervised Semantic Segmentation

Weakly supervised semantic segmentation attempts to learn a segmentation network with weaker annotation than pixel-wise labels. It aims to alleviate the annotation burden of segmentation tasks. Compared to bounding boxes [14, 26, 43], points [4], and scribbles [33, 46], image-level labels [29, 51, 21, 8, 23, 2, 52, 24] are the most widely used weak annotations due to their easy availability. They

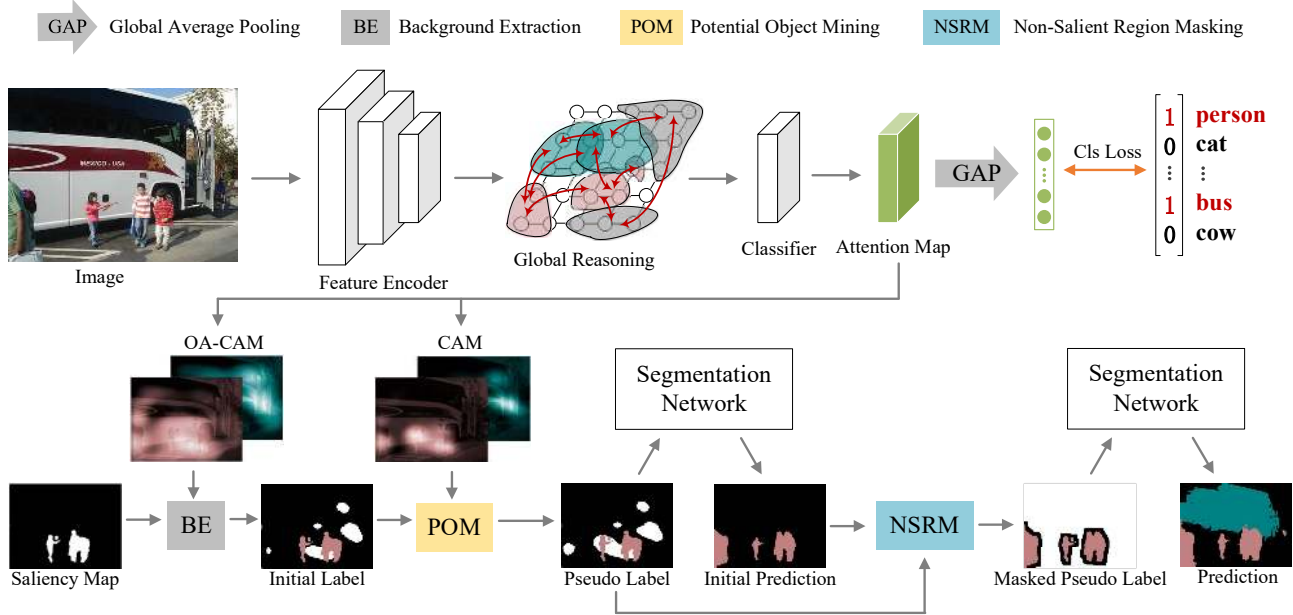


Figure 2. The architecture of our proposed approach. We train a classification network to generate the class activation maps (CAMs) and online accumulated class attention maps (OA-CAMs). A graph-based global reasoning unit is inserted into the classification network to help activate the objects outside the salient region. After obtaining the initial label with background extraction (BE), we utilize the potential object mining module (POM) to discover more objects missed in the initial label. We further leverage the non-salient region masking module (NSRM) to generate masked pseudo labels for the training of the segmentation network. Best viewed in color.

are already given in existing large-scale datasets (e.g. ImageNet [15]) or can be automatically generated through image retrieval techniques. Existing image-level label based approaches leverage the CAM to generate pixel-level seeds for training the segmentation model. Considering the initial seeds’ sparsity and incompleteness, researchers proposed many approaches to expand the seeds to integral object regions. For example, Kolesnikov and Lampert [29] introduced a new loss function for weakly supervised training based on three guiding principles: seed, expand, and constrain. Huang *et al.* [23] proposed to train a model starting from the discriminative regions and progressively increase the pixel-level supervision with the deep seeded region growing strategy. The work of RDC [52] leveraged the dilated convolution to enlarge the receptive fields of convolutional kernels. This helped transfer the object information to the non-discriminative region. AffinityNet [2] realized semantic propagation by a random walk with the semantic affinity between a pair of adjacent image coordinates. The recent work of SEAM [49] proposed a self-supervised equivariant attention mechanism to provide additional supervision for network learning. Apart from using intra-image information, Sun *et al.* [44] incorporated two neural co-attentions into the classifier to capture cross-image semantic relations for comprehensive object pattern mining. Zhang *et al.* [60] attributed the reason for the ambiguous boundaries of pseudo-masks to the confounding context.

They presented a causal inference framework to remove the confounding bias in image-level classification with an effective approximation for the backdoor adjustment.

3. The Proposed Approach

In this paper, we focus on the task of weakly supervised semantic segmentation with image-level labels. Our framework is illustrated in Fig. 2. Given a set of training images with image-level labels, we train a classification network. We leverage class activation maps to generate pseudo labels for learning a segmentation network. Unlike existing methods that mainly concentrate on refining pseudo labels in the salient area, we propose to discover more objects in the non-salient region for weakly supervised semantic segmentation. To achieve this, we insert a graph-based global reasoning unit into the classification network. This helps to activate the object features outside the salient region. We also adopt a potential object mining module (POM) and a non-salient region masking module (NSRM) to improve the quality of pseudo labels for non-salient region object mining.

3.1. CAM Generation

A classification network is first trained to generate class activation maps. As illustrated in Fig. 2, to strengthen the classification network’s ability to capture global relations among disjoint and distant regions, we introduce a graph-

based global reasoning unit [13] before the final classifier. The global reasoning module will help the network to activate the object parts outside the salient region. The features $X \in \mathbb{R}^{L \times K}$ generated by the encoder, with K being the feature dimension and $L = H \times W$ locations, is first projected from the coordinate space to a latent interaction space. The projection function $V = f(X) \in \mathbb{R}^{N \times K}$ is formulated as a linear combination:

$$v_i = b_i X = \sum_j b_{ij} x_j, \quad (1)$$

where $B = [b_1, \dots, b_N] \in \mathbb{R}^{N \times L}$ is the learnable projection weights, and N is the number of the features (nodes) in the interaction space.

Then a graph convolution [28] is applied to capture the relations between features in the new space:

$$Z = ((I - A_g) V) W_g. \quad (2)$$

A_g denotes the $N \times N$ node adjacency matrix learned by gradient descent during training. W_g denotes the state update function.

After obtaining the node-feature $Z \in \mathbb{R}^{N \times K}$, reverse projection $Y = g(Z) \in \mathbb{R}^{L \times K}$ is conducted to project the feature back to the original space:

$$y_i = d_i Z = \sum_j d_{ij} z_j, \quad (3)$$

where $D = [d_1, \dots, d_N] = B^T$.

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

$$L_{cls} = -\frac{1}{C} \sum_{c=1}^C y_c \log \sigma(p_c) + (1 - y_c) \log [1 - \sigma(p_c)]. \quad (4)$$

Here, p_c is the prediction of the network for the c -th class. $\sigma(\cdot)$ is the sigmoid function, and C is the total number of foreground classes. 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.

We obtain CAMs by selecting the class-specific feature maps generated by the final classifier. Following OAA [24], we also generate the online accumulated class attention maps (OA-CAMs) to have more entire regions and strengthen the lower attention values of target object regions with their integral attention model.

3.2. Potential Object Mining

After obtaining OA-CAMs, the work of OAA [24] uses them to extract object cues and saliency maps to extract background cues. The class label of each pixel is assigned by comparing the value of each OA-CAM. As shown in

Fig. 2, with the shape information provided by the saliency map, the initial label is derived with quite clear object boundaries after the background extraction (BE) process. However, the initial label misses many object parts outside the conspicuous area. Therefore, we propose to discover more objects in the non-salient region. Though the OA-CAM has a high recall of the object pixel, its precision is low. In contrast, the CAM, widely leveraged to generate initial seeds for proxy segmentation labels [29, 23], has low recall but high precision. Therefore, we propose a potential object mining (POM) module to discover the object region activated in the CAM. We mine the potential object with a class adaptive threshold T_c for class c that is present in the image:

$$T_c = \begin{cases} MED(v) & \text{if } \exists (i, j), \text{ s.t. } l_{ij} = c \\ TQ(v) & \text{otherwise} \end{cases}. \quad (5)$$

Here, v is the set of attention values of pixels in the CAM, whose locations p are selected as follows:

$$p = \begin{cases} \{(i, j) | l_{ij} = c\} & \text{if } \exists (i, j), \text{ s.t. } l_{ij} = c \\ \{(i, j) | a_{ij} > T_{bg}\} & \text{otherwise} \end{cases}, \quad (6)$$

where a_{ij} is the attention value in CAM at the position (i, j) . l_{ij} is the value in the initial label at the position (i, j) , which denotes the pseudo label of the pixel. As illustrated in Equation 5 and Equation 6, if the initial label contains class c , we select those pixels in its CAM and choose the median (MED) of their attention values as T_c . Otherwise, we select pixels in its CAM with an attention value greater than the background threshold T_{bg} and choose the top quartile (TQ) of their attention values as T_c .

We then adjust the initial label as follows:

$$l_{ij} = \begin{cases} 255 & \text{if } \exists c, (i, j), \text{ s.t. } l_{ij} = 0, a_{i,j}^c > T_c \\ l_{ij} & \text{otherwise} \end{cases}. \quad (7)$$

Here, a^c denotes the CAM for class c . As illustrated in Equation 7, the background pixels (labeled as 0) in the initial label with any CAM attention value greater than T are labeled as 255 and ignored for training. We do not label them for the corresponding potential class to avoid introducing wrong object labels. Such a strategy bypasses the necessity to locate the object boundary outside the salient region. We focus on reducing the false-negative rate of pseudo labels, which will help discard the gradients generated by the misleading information.

3.3. Non-Salient Region Masking

Our potential object mining strategy enriches pseudo labels with more ignored pixels. It allows the segmentation network to predict the correct labels for these potential object regions during training. The improved quality of




Figure 3. Our proposed non-salient region masking module.

pseudo labels can also encourage the segmentation network to fix the other incorrectly labeled regions. Therefore, we propose to further leverage the prediction of the segmentation model to generate pseudo labels of higher quality for retraining.

We notice that simple images with only one category of objects usually have a clean background. Objects in these images often exist in the salient region and can be correctly segmented by the segmentation network. However, complex images (with two or more categories of objects) are more prone to having objects outside the salient area. It remains challenging for the segmentation network to detect objects outside the salient region with pseudo labels only containing object labels in the salient area. Therefore, we propose a non-salient region masking (NSRM) module. It combines the object information in the segmentation network’s prediction and pseudo labels to generate masked labels for complex images.

Our proposed non-salient region masking module is illustrated in Fig. 3. Based on the assumption that object labels within the salient region are correct with high probability, we first expand the object region in the initial prediction with the guidance of our pseudo labels. Then we extract the object mask from the expanded prediction map. After that, we expand the object mask with a dilation operation. Finally, a masking operation is applied to the expanded prediction map to get the masked pseudo labels. Note that the dilation operation introduces a small portion of the background around the objects. It preserves the objects’ boundary information, which is of great importance for a successful segmentation network.

4. Experiments

4.1. Implementation Details

For the classification network, we adopt the VGG-16 model as our backbone, which is pre-trained on ImageNet [15]. Following [24], we add three convolutional layers on the top of the fully-convolutional backbone. A ReLU layer follows each convolutional layer for nonlinear transformation. A 1×1 convolutional layer of C channels is adopted

as the pixel-wise classifier to generate attention maps. The momentum and weight decay of the SGD [6] optimizer are 0.9 and 5×10^{-4} . The initial learning rate is set to 10^{-3} and is divided by 10 after every 5 epochs. Following the code of [24], we set the background threshold $T_{bg} = 0.3$ for fair comparison. We train the classification network for 14 epochs with batch size = 5.

For the segmentation network, following [7, 62, 18, 9], we adopt the DeepLab-v2 [10] framework. VGG-16 is pre-trained on ImageNet [15]. For ResNet-101 [20], we report results for models pre-trained on ImageNet [15] and MS-COCO [34], respectively. The momentum and weight decay of SGD optimizer are 0.9 and 10^{-4} . The initial learning rate is set to 2.5×10^{-4} and is decreased using polynomial decay with a power of 0.9. The segmentation network is trained for 10,000 iterations with batch size = 10.

4.2. Datasets and Evaluation Metrics

Following previous works, we evaluate our approach on the Pascal VOC 2012 dataset [16]. It contains 21 classes (20 object categories and the background) for semantic segmentation. There are 10,582 training images, which are expanded by [19], 1,449 validation images, and 1,456 test images. For all the experiments, we only adopt the image-level class labels for training. Standard mean intersection over union (mIoU) is taken as the evaluation metric for the semantic segmentation task.

4.3. Comparisons to the State-of-the-arts

Baselines. In this part, we compare our proposed method with the following state-of-the-arts approaches that leverage image-level labels for weakly supervised semantic segmentation: DCSM [41], SEC [29], AugFeed [39], STC [51], Roy *et al.* [40], Oh *et al.* [38], AE-PSL [50], WebS-i2 [25], Hong *et al.* [21], DCSP [8], TPL [27], GAIN [31], DSRG [23], MCOF [48], AffinityNet [2], RDC [52], SeeNet [22], OAA [24], ICD [17], BES [9], Fan *et al.* [18], Zhang *et al.* [62], MCIS [44], IRN [1], FickleNet[30], SSDD [42], SEAM [49], SCE [7], CONTA [60].

Experimental Results. We present our results for the backbone of VGG and ResNet in Table 1 and Table 2, respectively. As can be seen, our approach achieves better results than other state-of-the-art methods for both VGG and ResNet backbones. Specifically, for the VGG backbone, our segmentation results reach 65.5% and 65.3% on the validation and test set, respectively. For the ResNet backbone, we can get 68.3% on the validation set and 68.5% on the test set. Though the methods of STC [51], WebS-i2 [25] and Hong *et al.* [21] leverage additional training data, our method outperforms them on the validation set by 15.7%, 12.1% and 7.4%, respectively. Compared to DSRG [23] and CONTA [60], which also utilize the prediction of the segmentation network to get refined pseudo labels for train-

Methods	Publication	Val	Test
DCSM [41]	ECCV16	44.1	45.1
SEC [29]	ECCV16	50.7	51.7
AugFeed [39]	ECCV16	54.3	55.5
STC [51]	TPAMI17	49.8	51.2
Roy <i>et al.</i> [40]	CVPR17	52.8	53.7
Oh <i>et al.</i> [38]	CVPR17	55.7	56.7
AE-PSL [50]	CVPR17	55.0	55.7
WebS-i2 [25]	CVPR17	53.4	55.3
Hong <i>et al.</i> [21]	CVPR17	58.1	58.7
DCSP [8]	BMVC17	58.6	59.2
TPL [27]	ICCV17	53.1	53.8
GAIN [31]	CVPR18	55.3	56.8
DSRG [23]	CVPR18	59.0	60.4
MCOF [48]	CVPR18	56.2	57.6
AffinityNet [2]	CVPR18	58.4	60.5
RDC [52]	CVPR18	60.4	60.8
SeeNet [22]	NIPS18	63.1	62.8
OAA [24]	ICCV19	63.1	62.8
ICD [17]	CVPR20	64.0	63.9
BES [9]	ECCV20	60.1	61.1
Fan <i>et al.</i> [18]	ECCV20	64.6	64.2
Zhang <i>et al.</i> [62]	ECCV20	63.7	64.5
MCIS [44]	ECCV20	63.5	63.6
Ours	-	65.5	65.3

Table 1. Quantitative comparisons to previous state-of-the-art approaches with VGG backbone.

ing, our approach can improve their results by 6.9% and 2.2%, respectively. The work of ICD [17] uses the additional superpixel to help recover the object boundary information during training. Our approach can still outperform it by 1.5% for the VGG backbone and 0.5% for the ResNet backbone. Our results demonstrate the effectiveness of mining the objects in the non-salient region for the task of weakly supervised semantic segmentation. When training the ResNet based network with the COCO pre-trained weights, we can further reach 70.4% and 70.2% on the validation and test set, respectively.

4.4. Ablation Studies

Element-Wise Component Analysis. In this part, we demonstrate the contribution of each component proposed in our approach for weakly supervised semantic segmentation. The experimental results on the validation set of Pascal VOC are given in Table 3. We notice that, by leveraging the graph-based global reasoning unit (GR) to capture global relations among disjoint and distant regions, we can improve the segmentation result from 67.7% to 68.8%. By introducing our proposed potential object mining module (POM), we obtain another 0.2% performance gain. Note that if we directly retrain the segmentation network with its

Methods	Publication	Val	Test
DCSP [8]	BMVC17	60.8	61.9
DSRG [23]	CVPR18	61.4	63.2
MCOF [48]	CVPR18	60.3	61.2
AffinityNet [2]	CVPR18	61.7	63.7
SeeNet [22]	NIPS18	63.1	62.8
IRN [1]	CVPR19	63.5	64.8
FickleNet[30]	CVPR19	64.9	65.3
OAA [24]	ICCV19	65.2	66.4
SSDD [42]	ICCV19	64.9	65.5
SEAM [49]	CVPR20	64.5	65.7
SCE [7]	CVPR20	66.1	65.9
ICD [17]	CVPR20	67.8	68.0
Zhang <i>et al.</i> [62]	ECCV20	66.6	66.7
Fan <i>et al.</i> [18]	ECCV20	67.2	66.7
MCIS [44]	ECCV20	66.2	66.9
BES [9]	ECCV20	65.7	66.6
CONTA [60]	NIPS20	66.1	66.7
Ours	-	68.3	68.5
OAA* [24]	ICCV19	67.4	-
SEAM* [49]	CVPR20	63.2	-
Ours*	-	70.4	70.2

Table 2. Quantitative comparisons to previous state-of-the-art approaches with ResNet backbone. * denotes model is pre-trained on MS-COCO.

Methods	Val
baseline	67.7
+ GR	68.8
+ GR + retrain	68.7
+ GR + POM	69.0
+ GR + POM + retrain	69.7
+ GR + POM + NSRM	70.4

Table 3. Element-wise component analysis with ResNet backbone. GR: Global Reasoning, POM: Potential Object Mining, NSRM: Non-Salient Region Masking.

prediction, the performance drops from 68.8% to 68.7%. In contrast, with our potential object mining module, retraining the segmentation network can further improve the result to 69.7%. This highlights the importance of reducing the false-negative rate of pseudo labels. Our potential object mining module can directly improve the segmentation results and help to exert the self-correction ability of the segmentation network for a higher quality of pseudo labels. With our non-salient region masking module (NSRM), we further exploit the objects outside the conspicuous regions and improve the segmentation result to 70.4%.

Some qualitative segmentation examples on the PASCAL VOC 2012 validation set can be viewed in Fig. 4. As can be seen, with the graph-based global reasoning unit




Figure 4. Example results on PASCAL VOC 2012 validation set. For each (a) image, we show the (b) ground truth, the result of (c) baseline, (d) GR, (e) GR + POM, and (f) our full method. The first four columns of images contain objects of one category, and the other five columns of images contain objects of two or more categories. Best viewed in color.

Methods	Val
NSRM	70.4
NSRM all	68.8
NSRM - Object Expansion	70.2
NSRM - Masking	70.0
NSRM - dilation	68.5

Table 4. Ablation studies for NSRM with ResNet backbone.

(GR), the network can capture global relations and discover objects in disjoint and distant regions (*e.g.* the potted plant in the fifth column and person in the eighth column). As shown in the last column, our method with the POM module can further discover the car outside the salient region. Besides, our full method with the NSRM module exerts the segmentation network’s self-correction ability. It successfully predicts the bus and potted plants in the second and fourth columns, respectively. As shown in the fifth and eighth columns, our robust method further mines the objects outside the salient region.

We display the evolution of the labels we used for training the segmentation network in Fig. 5. As we can see, with our potential object mining module, we discover more object regions outside the salient area than the initial label. Our pseudo label can reduce the false gradients calculated for wrong annotations. By leveraging the initial

prediction, our non-salient region masking module generates high-quality masked labels, allowing the segmentation model to mine the objects in the non-salient region further.

Ablation Studies for NSRM. An in-depth study of our proposed NSRM module is presented in Table 4. As we can see, if we apply NSRM to all images without our simple and complex image division, the results drop from 70.4% to 68.8%. This highlights the importance of treating simple and complex images differently. When masking out the non-salient region of pseudo labels for complex images during training, we need to rely on the rich background information provided by simple images. We notice that removing the object expansion operation will cause a 0.2% performance drop. This shows that it is useful to utilize the pseudo labels to expand the object prediction within the salient region. Masking out the non-salient region of the pseudo label for training has a 0.4% performance gain. This shows that the masking operation can encourage the segmentation network to exert its self-correction ability. Note that if we do not conduct the dilation operation for the object mask, the performance directly drops to 68.5%. This highlights the importance of preserving the background area around the object. The background information, together with the object region, provides essential boundary knowledge for the network training.

Parameter Analysis. For the dilation operation in the NSRM module, we conduct experiments to study the effect




Figure 5. The evolution of labels for the PASCAL VOC 2012 training set. For each (a) image, we show the (b) ground truth, (c) saliency map, (d) initial label, (e) pseudo label, (f) initial prediction, (g) masked label, and (h) final prediction. Best viewed in color.




Figure 6. The parameter sensitivity of the dilation kernel size r in the non-salient region masking module. Results are reported on the validation set using the ResNet backbone.




Figure 7. The parameter sensitivity of the number of feature nodes N in the interaction space. Results are reported on the validation set using the ResNet backbone.

of the dilation kernel size r . As shown in Fig. 6, we vary the kernel size r over the range $\{2, 5, 10, 20, 30, 40, 50\}$. As we can see, we get better performance when the kernel size is between 5 and 30. A too large or small kernel size may not improve the results very much. We conjecture that a too large kernel size keeps too much background in the prediction, which hinders the object mining in the non-salient region. Meanwhile, a too small kernel size with little background blurs the boundary of objects, which impedes the training of the segmentation network. In our experiments, we empirically set $r = 30$.

For the graph-based global reasoning unit, we conduct

experiments to study the effect of the number of feature nodes N in the interaction space. As shown in Fig. 7, we vary N over the range $\{16, 32, 64, 96, 128\}$. We notice that a too large number of the nodes N may not improve the performance very much. In our experiments, we empirically set $N = 64$.

5. Conclusions

In this work, we proposed a non-salient region object mining approach for the task of weakly supervised semantic segmentation. Specifically, we introduced a graph-based global reasoning unit to help the classification network capture global relations among disjoint and distant regions. This can strengthen the network’s ability to activate the objects scattered in the corners or near the edge of the image. To further mine objects outside the non-salient region, we proposed to exert the segmentation network’s self-correction ability. A potential object mining module was proposed to reduce the false-negative rate in pseudo labels. Moreover, we proposed a non-salient region masking module for complex images to generate masked pseudo labels. Our non-salient region masking module helps further discover objects in the non-salient region. Extensive experiments on the PASCAL VOC 2012 dataset demonstrated the superiority of our proposed approach.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 61976116) and Fundamental Research Funds for the Central Universities (No. 30920021135).

References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *CVPR*, pages 2209–2218, 2019. 5, 6
- [2] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, pages 4981–4990, 2018. 1, 2, 3, 5, 6
- [3] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. In *IEEE TPAMI*, volume 39, pages 2481–2495, 2017. 2
- [4] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *ECCV*, pages 549–565, 2016. 1, 2
- [5] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. Salient object detection: A benchmark. In *IEEE TIP*, volume 24, pages 5706–5722, 2015. 2
- [6] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *COMPSTAT*, pages 177–186, 2010. 5
- [7] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation via sub-category exploration. In *CVPR*, pages 8991–9000, 2020. 5, 6
- [8] Arslan Chaudhry, Puneet Dokania, and Philip Torr. Discovering class-specific pixels for weakly-supervised semantic segmentation. In *BMVC*, 2017. 1, 2, 5, 6
- [9] Liyi Chen, Weiwei Wu, Chenchen Fu, Xiaojing Han, and Yun-Tao Zhang. Weakly supervised semantic segmentation with boundary exploration. In *ECCV*, 2020. 5, 6
- [10] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. In *IEEE TPAMI*, volume 40, pages 834–848, 2017. 2, 5
- [11] Tao Chen, Guo-Sen Xie, Yazhou Yao, Qiong Wang, Fumin Shen, Zhenmin Tang, and Jian Zhang. Semantically meaningful class prototype learning for one-shot image segmentation. In *IEEE TMM*, 2021. 2
- [12] Tao Chen, Jian Zhang, Guo-Sen Xie, Yazhou Yao, Xiaoshui Huang, and Zhenmin Tang. Classification constrained discriminator for domain adaptive semantic segmentation. In *ICME*, pages 1–6, 2020. 2
- [13] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Shuicheng Yan, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. In *CVPR*, pages 433–442, 2019. 2, 4
- [14] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *ICCV*, pages 1635–1643, 2015. 1, 2
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 3, 5
- [16] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. In *IJCV*, volume 88, pages 303–338, 2010. 5
- [17] Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, and Tieniu Tan. Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In *CVPR*, pages 4283–4292, 2020. 5, 6
- [18] Junsong Fan, Zhaoxiang Zhang, and Tieniu Tan. Employing multi-estimations for weakly-supervised semantic segmentation. In *ECCV*, 2020. 5, 6
- [19] Bharath Hariharan, Pablo Arbeláez, Lubomir D. Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, pages 991–998, 2011. 5
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [21] Seunghoon Hong, Donghun Yeo, Suha Kwak, Honglak Lee, and Bohyung Han. Weakly supervised semantic segmentation using web-crawled videos. In *CVPR*, pages 7322–7330, 2017. 1, 2, 5, 6
- [22] Qibin Hou, Peng-Tao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. In *NeurIPS*, pages 549–559, 2018. 5, 6
- [23] Zilong Huang, Xinggang Wang, Jiashi Wang, Wenyu Liu, and Jingdong Wang. Weakly-supervised semantic segmentation network with deep seeded region growing. In *CVPR*, pages 7014–7023, 2018. 1, 2, 3, 4, 5, 6
- [24] Peng-Tao Jiang, Qi-Bin Hou, Yuanpeng Cao, Ming-Ming Cheng, Yunchao Wei, and Hongkai Xiong. Integral object mining via online attention accumulation. In *ICCV*, pages 2070–2079, 2019. 1, 2, 4, 5, 6
- [25] Bin Jin, Maria Ortiz Segovia, and Sabine Süsstrunk. Webly supervised semantic segmentation. In *CVPR*, pages 3626–3635, 2017. 5, 6
- [26] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, pages 876–885, 2017. 1, 2
- [27] Dahun Kim, Donghyeon Cho, and Donggeun Yoo. Two-phase learning for weakly supervised object localization. In *ICCV*, pages 3534–3543, 2017. 5, 6
- [28] Thomas Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2016. 4
- [29] Alexander Kolesnikov and Christoph Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *ECCV*, pages 695–711, 2016. 1, 2, 3, 4, 5, 6
- [30] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *CVPR*, pages 5267–5276, 2019. 5, 6
- [31] Kunpeng Li, Ziyang Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. Tell me where to look: Guided attention inference network. In *CVPR*, pages 9215–9223, 2018. 5, 6
- [32] Xueyi Li, Tianfei Zhou, Jianwu Li, Yi Zhou, and Zhaoxiang Zhang. Group-wise semantic mining for weakly supervised semantic segmentation. 2021. 1
- [33] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, pages 3159–3167, 2016. 1, 2

- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Perona Perona, Deva Ramanan, Piotr Dollár, and Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014. 5
- [35] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan Yuille, and Li Fei-Fei. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In *CVPR*, pages 82–92, 2019. 2
- [36] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. 2
- [37] Haonan Luo, Guosheng Lin, Zichuan Liu, Fayao Liu, Zhenmin Tang, and Yazhou Yao. Segeqa: Video segmentation based visual attention for embodied question answering. In *ICCV*, pages 9667–9676, 2019. 2
- [38] Seong Joon Oh, Rodrigo Benenson, Anna Khoreva, Zeynep Akata, Mario Fritz, and Bernt Schiele. Exploiting saliency for object segmentation from image level labels. In *CVPR*, pages 5038–5047, 2017. 5, 6
- [39] Xiaojuan Qi, Zhengzhe Liu, Jianping Shi, Hengshuang Zhao, and Jiaya Jia. Augmented feedback in semantic segmentation under image level supervision. In *ECCV*, pages 90–105, 2016. 5, 6
- [40] Anirban Roy and Sinisa Todorovic. Combining bottom-up, top-down, and smoothness cues for weakly supervised image segmentation. In *CVPR*, pages 3529–3538, 2017. 5, 6
- [41] Wataru Shimoda and Keiji Yanai. Distinct class-specific saliency maps for weakly supervised semantic segmentation. In *ECCV*, pages 218–234, 2016. 5, 6
- [42] Wataru Shimoda and Keiji Yanai. Self-supervised difference detection for weakly-supervised semantic segmentation. In *ICCV*, pages 5208–5217, 2019. 5, 6
- [43] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In *CVPR*, pages 3136–3145, 2019. 1, 2
- [44] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *ECCV*, 2020. 3, 5, 6
- [45] Zeren Sun, Xian-Sheng Hua, Yazhou Yao, Xiu-Shen Wei, Guosheng Hu, and Jian Zhang. Crssc: salvage reusable samples from noisy data for robust learning. In *ACM MM*, pages 92–101, 2020. 2
- [46] Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *CVPR*, pages 7158–7166, 2017. 1, 2
- [47] Wenguan Wang, Tianfei Zhou, Fisher Yu, Jifeng Dai, Ender Konukoglu, and Luc Van Gool. Exploring cross-image pixel contrast for semantic segmentation. *arXiv preprint arXiv:2101.11939*, 2021. 1
- [48] Xiang Wang, Shaodi You, Xi Li, and Huimin Ma. Weakly-supervised semantic segmentation by iteratively mining common object features. In *CVPR*, pages 1354–1362, 2018. 5, 6
- [49] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*, pages 12275–12284, 2020. 3, 5, 6
- [50] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *CVPR*, pages 1568–1576, 2017. 5, 6
- [51] Yunchao Wei, Xiaodan Liang, Yunpeng Chen, Xiaohui Shen, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Stc: A simple to complex framework for weakly-supervised semantic segmentation. In *IEEE TPAMI*, volume 39, pages 2314–2320, 2016. 1, 2, 5, 6
- [52] Yunchao Wei, Huaxin Xiao, Humphrey Shi, Zequn Jie, Jiashi Feng, and Thomas Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *CVPR*, pages 7268–7277, 2018. 1, 2, 3, 5, 6
- [53] Yazhou Yao, Xiansheng Hua, Guanyu Gao, Zeren Sun, Zhibin Li, and Jian Zhang. Bridging the web data and fine-grained visual recognition via alleviating label noise and domain mismatch. In *ACM MM*, pages 1735–1744, 2020. 2
- [54] Yazhou Yao, Xian-sheng Hua, Fumin Shen, Jian Zhang, and Zhenmin Tang. A domain robust approach for image dataset construction. In *ACM MM*, pages 212–216, 2016. 2
- [55] Yazhou Yao, Fumin Shen, Guosen Xie, Li Liu, Fan Zhu, Jian Zhang, and Heng Tao Shen. Exploiting web images for multi-output classification: From category to subcategories. *IEEE TNNLS*, 31(7):2348–2360, 2020. 1
- [56] Yazhou Yao, Fumin Shen, Jian Zhang, Li Liu, Zhenmin Tang, and Ling Shao. Extracting multiple visual senses for web learning. *IEEE TMM*, 21(1):184–196, 2019. 1
- [57] Yazhou Yao, Fumin Shen, Jian Zhang, Li Liu, Zhenmin Tang, and Ling Shao. Extracting privileged information for enhancing classifier learning. *IEEE TIP*, 28(1):436–450, 2019. 1
- [58] Yazhou Yao, Jian Zhang, Fumin Shen, Xiansheng Hua, Jingsong Xu, and Zhenmin Tang. Exploiting web images for dataset construction: A domain robust approach. *IEEE TMM*, 19(8):1771–1784, 2017. 1
- [59] Yazhou Yao, Jian Zhang, Fumin Shen, Li Liu, Fan Zhu, Dongxiang Zhang, and Heng Tao Shen. Towards automatic construction of diverse, high-quality image datasets. *IEEE TKDE*, 32(6):1199–1211, 2020. 1
- [60] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xian-Sheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. In *NeurIPS*, volume 33, 2020. 3, 5, 6
- [61] Hang Zhang, Kristin Dana, Jianping Shi, Zhongyue Zhang, Xiaoang Wang, Amrith Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *CVPR*, pages 7151–7160, 2018. 2
- [62] Tianyi Zhang, Guosheng Lin, Weide Liu, Jianfei Cai, and Alex Kot. Splitting vs. merging: Mining object regions with discrepancy and intersection loss for weakly supervised semantic segmentation. In *ECCV*, 2020. 5, 6
- [63] Bolei Zhou, Aditya Khosla, Àgata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 1, 2