

Discriminative Region Suppression for Weakly-Supervised Semantic Segmentation

Beomyoung Kim, Sangeun Han, Junmo Kim

Korea Advanced Institute of Science and Technology (KAIST)
{qjadud1994, bichoomi, junmo.kim}@kaist.ac.kr

Abstract

Weakly-supervised semantic segmentation (WSSS) using image-level labels has recently attracted much attention for reducing annotation costs. Existing WSSS methods utilize localization maps from the classification network to generate pseudo segmentation labels. However, since localization maps obtained from the classifier focus only on sparse discriminative object regions, it is difficult to generate high-quality segmentation labels. To address this issue, we introduce discriminative region suppression (DRS) module that is a simple yet effective method to expand object activation regions. DRS suppresses the attention on discriminative regions and spreads it to adjacent non-discriminative regions, generating dense localization maps. DRS requires few or no additional parameters and can be plugged into any network. Furthermore, we introduce an additional learning strategy to give a self-enhancement of localization maps, named localization map refinement learning. Benefiting from this refinement learning, localization maps are refined and enhanced by recovering some missing parts or removing noise itself. Due to its simplicity and effectiveness, our approach achieves mIoU 71.4% on the PASCAL VOC 2012 segmentation benchmark using only image-level labels. Extensive experiments demonstrate the effectiveness of our approach.

Introduction

Recent developments in deep learning have achieved great success on semantic segmentation tasks with the help of deep convolutional neural networks (CNNs) and rich pixel-level annotations. However, collecting a large-scale pixel-level annotated dataset requires intensive human labor, which is both expensive and time-consuming. To end this limitation, weakly-supervised semantic segmentation (WSSS) using only image-level labels has recently attracted much attention.

One problem with using image-level annotations is that we have no information about the location of the target object; we only know whether the object is present in the image or not. This makes semantic segmentation learning challenging. To learn pixel-level semantic knowledge from image-level labels, it is common practice to use localization maps obtained from the classification network using class activa-

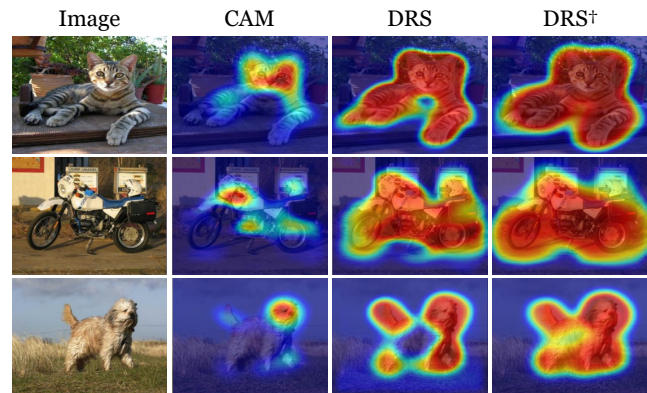


Figure 1: Visual comparisons of localization maps produced by original CAM, DRS, and DRS[†]. DRS[†] denotes that the refinement learning is applied.

tion maps (CAMs) (Zhou et al. 2016). Specifically, the discriminative region for each target class provided by CAMs is used as pixel-level supervision for segmentation network training. However, this discriminative region is usually very sparse and only covers a small part of the object, which is not enough for semantic segmentation learning as shown in the second column in Figure 1. Therefore, most studies in the weakly-supervised semantic segmentation field focus on expanding the object region to produce dense localization maps. One of the recent approaches is image-level and feature-level erasure of discriminative parts (Wei et al. 2017; Li et al. 2018; Hou et al. 2018). This approach strictly erases discriminative parts, letting the network focus on other non-discriminative parts. However, they not only tend to produce undesired true negative regions when most of the discriminative parts are erased but also require a lot of additional parameters for multiple classifiers or multiple branches.

In this paper, we propose discriminative region suppression (DRS) module, which is a simple and efficient yet effective and novel approach for generating dense localization maps. The goal of DRS is to suppress discriminative regions, not to erase them, so that attention spreads to adjacent non-discriminative regions; this mild approach helps the classifier effectively expand discriminative object regions. DRS module consists of three components as depicted in Figure

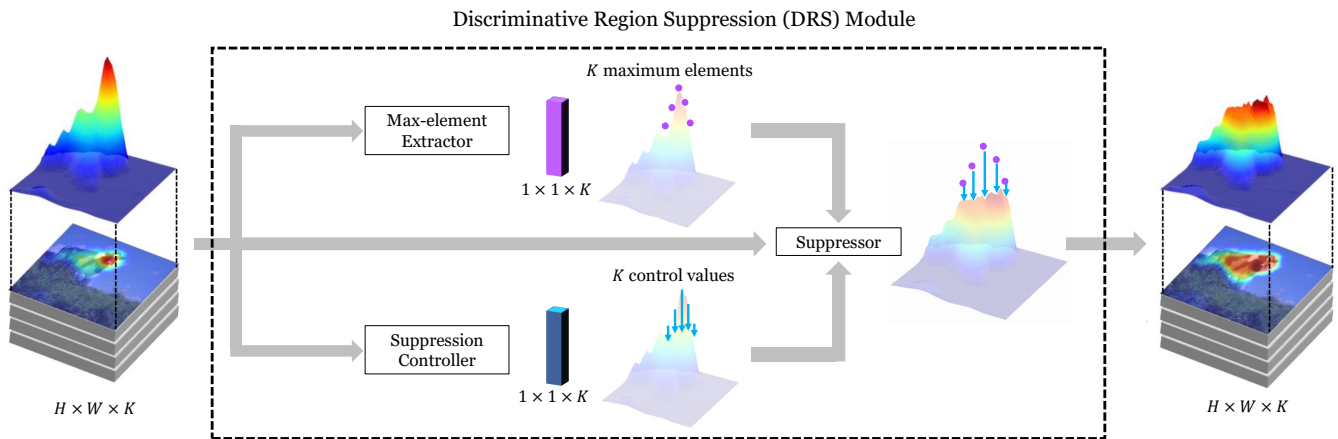


Figure 2: Diagram of our discriminative region suppression (DRS) module. DRS suppresses intermediate feature maps, diffusing the attention into adjacent non-discriminative parts. The max-element extractor extracts K maximum elements from intermediate feature maps. These K maximum elements are the maximum points of each discriminative region and are considered as starting points to be suppressed. For convenience, K maximum elements are illustrated in 5 purple points. The controller predicts control values, which determine how much to suppress feature maps from these K maximum elements. These K control values are illustrated in 5 blue arrows and the length of the arrow means how much suppress feature maps from the corresponding maximum element. Using these K maximum elements and K control values, the suppressor suppresses discriminative regions and spreads the attention into adjacent non-discriminative parts.

2: max-element extractor, suppression controller, and suppressor. These components work together to produce dense localization maps by reducing the attention gap between discriminative regions and adjacent non-discriminative regions. DRS not only effectively expands the object regions without generating much noise, but also can be plugged into any network with few or no additional parameters.

Although we can obtain dense segmentation labels from the classification network equipped with DRS, it does not recover missing parts or weak attention by itself because the objective of the classification network is classification, not localization. To address this issue, we introduce an additional training strategy, named localization map refinement learning, inspired by (Jiang et al. 2019) Localization map refinement learning induces self-enhancement of localization maps by recovering missing or weak attention region. In Figure 1, we compare some results of DRS and DRS \dagger , where DRS \dagger denotes that refinement learning is applied.

Following the convention, we generate pseudo segmentation labels from our dense localization maps and evaluate on weakly-supervised semantic segmentation task. On the PASCAL VOC 2012 segmentation benchmark, we achieve mIoU 71.4% on the testset using only image-level labels. In addition, extensive experiments demonstrate the effectiveness of our approach.

In summary, the contributions of our work are as follows:

- We introduce a simple, effective, and novel approach for weakly-supervised semantic segmentation named discriminative region suppression (DRS) module, which requires few or no additional parameters and can be easily plugged into any network.
- DRS effectively and efficiently suppresses discriminative

regions to generate dense localization maps, bridging the gap between discriminative regions and adjacent non-discriminative regions.

- For self-enhancement of localization maps, we introduce an additional training strategy, named localization map refinement learning.
- Extensive experiments and analyses demonstrate the effectiveness of our DRS module and we achieve competitive performance on Pascal VOC 2012 segmentation benchmark using only image-level labels.

Related Work

Most recent studies on semantic segmentation using image-level labels as weak supervision utilize CAMs (Zhou et al. 2016) to localize object regions and focus on expanding them to non-discriminative parts of the objects. To this end, AE-PSL (Wei et al. 2017), GAIN (Li et al. 2018), and SeeNet (Hou et al. 2018) propose erasing techniques to generate dense localization maps. However, these erasure-based approaches usually require multiple classifiers and complicated training procedures. Moreover, erasing most of the discriminative regions may introduce true negative regions and confuse the classifier.

To avoid the repetitive training procedures of AE-PSL (Wei et al. 2017), MDC (Wei et al. 2018) propose a multi-dilated convolution block in which the receptive fields of various sizes capture different patterns. As a more generalized approach, FickleNet (Lee et al. 2019) aggregate diverse localization maps produced by stochastic feature selection. Although they effectively expand the activated regions, some falsely labeled regions outside the object tend to be identified because the receptive fields of these methods

are not adaptive to object size. The recently proposed OAA (Jiang et al. 2019) accumulates attention maps at different training epochs and introduces integral attention learning to enhance attention maps. However, it may produce undesired attention regions due to training instability in the early stage.

Some other works (Ahn and Kwak 2018; Huang et al. 2018; Shimoda and Yanai 2019) adopt a region-growing technique to expand initial regions. More recently, RRM (Zhang et al. 2020) proposed a fully end-to-end network for joint training of classification and segmentation, and SGAN (Yao and Gong 2020) proposed a self-attention network guided by saliency priors that can produce dense and accurate localization maps from rich contextual information. BES (Chen et al. 2020) explores object boundaries to refine the semantic segmentation output. ICD (Fan et al. 2020) proposes an intra-class discriminator approach to separate foreground objects and the background within the same image-level class.

Method

The overview of our method is illustrated in Figure 3. We sequentially train three different networks for classification, refinement, and segmentation. After training the classification network with the discriminative region suppression (DRS) module, we produce dense localization maps. Using these dense localization maps as ground truth labels for refinement learning, we train the refinement network to produce refined localization maps. Then, pseudo segmentation labels are generated from the refined localization maps and used for training the semantic segmentation network. We measure the segmentation performance to evaluate the quality of our localization maps.

Observation

We produce localization maps from the class-specific feature maps of the last convolutional layer, which have been proven by (Zhang et al. 2018) to be mathematically equivalent to CAMs. We employ the VGG-16 (Simonyan and Zisserman 2014) as our classification network to produce localization maps. To be specific, we employ modified VGG-16, where all fully connected layers are removed. On top of it, three convolutional layers with 512 channels and kernel size 3, and a convolutional layer with C channels and kernel size 1 are added. Here C is the number of categories. This network produces output feature maps $F \in \mathbb{R}^{H_{out} \times W_{out} \times C}$ and classification score $P = \sigma(GAP(F))$ from input image. H_{out} and W_{out} are the height and width of output feature maps, respectively; $GAP(\cdot)$ is the global average pooling; and $\sigma(\cdot)$ is the sigmoid function. For each target category c , c -th localization map M^c is defined as the normalized c -th feature map F^c :

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

From the definition of M , we observe that discriminative object regions are identified with relatively high values on the feature maps F . Based on this observation, we regard the high-value areas on feature maps as discriminative regions.

Algorithm 1: Discriminative Region Suppression

Input: Intermediate feature maps $X \in \mathbb{R}^{H \times W \times K}$

Output: Suppressed feature maps : $\hat{X} \in \mathbb{R}^{H \times W \times K}$

$X_{max} \leftarrow extractor(X)$ // $X_{max} \in \mathbb{R}^{1 \times 1 \times K}$

$G \leftarrow controller(X)$ // $G \in [0, 1]^{1 \times 1 \times K}$

$\tau \leftarrow X_{max} \cdot G$ // upper bound, $\tau \in \mathbb{R}^{1 \times 1 \times K}$

$\tau \leftarrow expand\ to\ the\ same\ shape\ of\ X$

$\hat{X} \leftarrow min(X, \tau)$ // suppressor

Discriminative Region Suppression

To produce dense localization maps, we propose discriminative region suppression (DRS) module. The main problem of segmentation label generation using CAMs is that discriminative regions only appear partially and sparsely, as shown in the second column of Figure 1. To address this issue, DRS aims to spread the attention on discriminative regions to adjacent non-discriminative regions. Specifically, DRS suppresses the attention on discriminative regions, allowing the network to focus on non-discriminative regions.

Let $X \in \mathbb{R}^{H \times W \times K}$ be an intermediate feature map, where H , W , and K are the height, width, and the number of channels of X . DRS module consists of three parts: max-element extractor, suppression controller, and suppressor. The max-element extractor extracts K maximum elements from the intermediate feature map X using global max pooling. The output of the extractor is denoted as $X_{max} \in \mathbb{R}^{1 \times 1 \times K}$. Based on the observation, these K maximum elements are regarded as the criteria of discriminative regions and considered as starting points to be suppressed.

The suppression controller determines how much to suppress discriminative regions. In detail, it generates $G \in [0, 1]^{1 \times 1 \times K}$ and each k -th control value in G determines the amount of suppression in X with respect to the corresponding k -th maximum element.

Using the K maximum elements and K control values, the suppressor suppresses discriminative regions. Specifically, element-wise multiplication of X_{max} and G is regarded as the upper bound of X , denoted as $\tau = X_{max} \cdot G$, $\tau \in \mathbb{R}^{1 \times 1 \times K}$. The regions in X above this upper bound are regarded as discriminative regions to be suppressed. After the upper bound τ is expanded to the same shape of X , the element-wise minimum operation is applied on X and τ to suppress discriminative regions. For example, if the k -th control value is 0.7, X^k is suppressed until no element exceeds 70% of the k -th maximum value. In this way, the suppressor bridges the gap between discriminative regions and adjacent non-discriminative regions. The whole process of DRS is described in Algorithm 1 and illustrated in Figure 2.

For the suppression controller, there are two types of controller: learnable controller and non-learnable controller. If the suppression power is too strong, the discriminative feature extraction power is weakened. The learnable controller adaptively balances between discriminative feature extraction power and suppression power of the classification net-

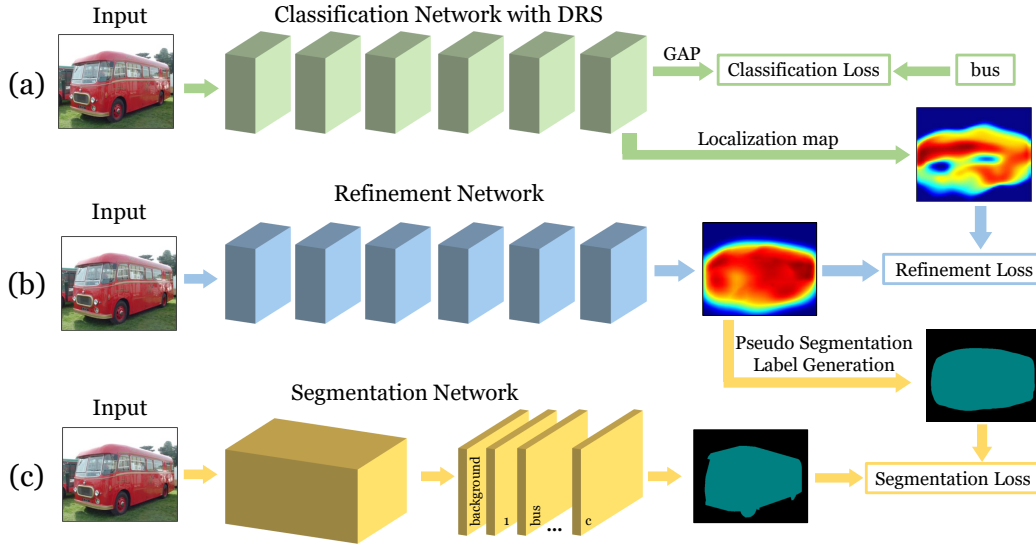


Figure 3: Overview of the proposed method. (a) Classification network with DRS for obtaining localization maps, (b) localization map refinement learning, and (c) weakly-supervised semantic segmentation. (a), (b), and (c) are executed sequentially, not simultaneously. Note that GAP in (a) means the global average pooling layer.

work. Formally, the output of the learnable controller is

$$G = \sigma(f(\text{GAP}(X); \theta)), \quad (2)$$

where f is a fully connected layer, θ is a learnable parameter of the controller, and $G \in [0, 1]^{1 \times 1 \times K}$. Since θ is trained with the classification objective, DRS with a learnable controller adaptively suppresses discriminative regions so as not to damage the discriminative feature extraction power much.

To produce even more dense localization maps at the expense of discriminative feature extraction power, we forcibly suppress discriminative regions; this is the goal of a non-learnable controller. For the non-learnable controller, each element of G is set to a constant value δ . We set the hyperparameter δ to a value between 0 and 1, and a lower δ means more intense suppression resulting in more dense localization maps. Compared to the learnable controller, the non-learnable controller does not require additional training parameters but requires a hyperparameter δ . In the experiment section, we analyze both learnable and non-learnable controller with quantitative and qualitative results.

Figure 3 (a) illustrates the process of obtaining dense localization maps from the classification network with DRS. As shown in the third column of Figure 1, DRS reduces the gap between the activation of discriminative regions and adjacent non-discriminative regions to obtain dense localization maps. Note that DRS can be plugged into any layer of a network.

Localization Map Refinement Learning

Although DRS helps produce dense localization maps, the DRS itself lacks the ability to recover missing parts of the target objects or enhance weak attention in adjacent non-discriminative regions because the goal of the classification network is essentially classification ability, not localization

map generation. Motivated by (Jiang et al. 2019), we introduce an additional learning strategy for localization map refinement to solve the above limitations. This learning strategy for self-enhancement of localization maps is called localization map refinement learning, denoted as DRS^\dagger . After training the classification network with DRS, we exploit the output localization maps $M \in [0, 1]^{H_{out} \times W_{out} \times C}$ as the ground truth localization maps for refinement learning.

The network for refinement learning, called refinement network, is based on the VGG-16; all fully-connected layers are removed and three convolutional layers with 512 channels and kernel size 3, and a convolutional layer with C channels and kernel size 1 are appended. The refinement network directly produces refined localization maps $N \in \mathbb{R}^{H_{out} \times W_{out} \times C}$, which have the same shape as M . We adopt the mean squared error (MSE) loss function as the refinement loss for the refinement network. Refinement learning is depicted in Figure 3 (b).

Benefiting from refinement learning, we can obtain more dense and high-quality localization maps through self-enhancement, as shown in Figure 1 (DRS^\dagger v.s. DRS)

Weakly-Supervised Semantic Segmentation

Using our dense localization maps obtained from the refinement network, we generate pseudo segmentation labels and use them as weak-supervision for the semantic segmentation network. We generate pseudo segmentation labels using object cues and background cues. We extract object cues from the localization maps by taking the pixels whose values are higher than α and extract background cues using salient object detection method (Liu et al. 2019), motivated by (Wei et al. 2017, 2018); the pixels with saliency values lower than β are taken as background. Those who belong to neither of the cues are ignored.

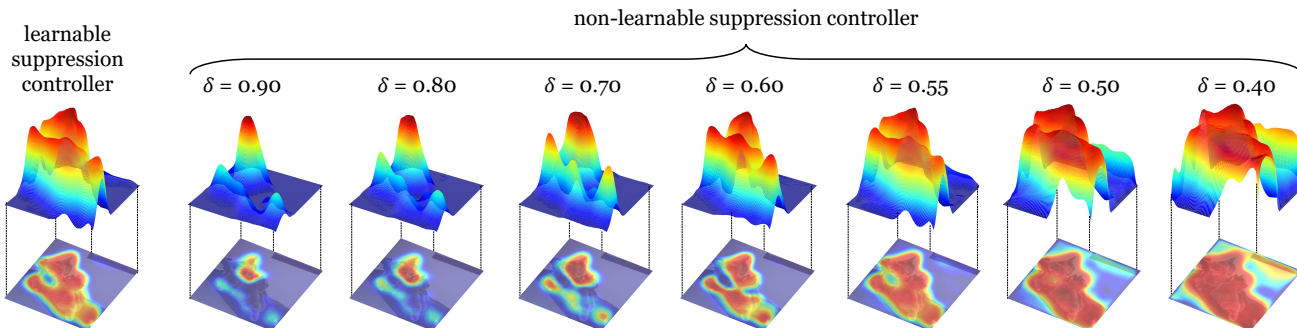


Figure 4: Visualization of localization maps of two types of controllers and the non-learnable controller with different δ .

Following the convention, we train the segmentation network such as (Chen et al. 2014, 2017) using the generated pseudo segmentation labels, as illustrated in Figure 3 (c). The segmentation performance is compared with other methods using the same segmentation network to evaluate the quality of pseudo segmentation labels.

Experiments

Dataset and Evaluation Metrics

We demonstrate the effectiveness of the proposed approach on the PASCAL VOC 2012 segmentation benchmark dataset (Everingham et al. 2014), which contains 20 object categories and one background category. Following the common practice in previous works, the training set is augmented to 10,582 images. We evaluate the performance of our model using the mean intersection-over-union (mIoU) metric and compare it with other state-of-the-art methods on the validation (1,449 images) and test set (1,456 images). For the test results, we submit the prediction outputs to the official PASCAL VOC evaluation server.

Implementation details

For the classification network, we adopt the modified VGG-16 with DRS plugged into every layer, as mentioned in the method section. Its parameters are initialized by the VGG-16 (Simonyan and Zisserman 2014) pre-trained on ImageNet (Deng et al. 2009) except for the additional convolutional layers. We train the classification network with binary cross-entropy loss using the SGD optimizer with a weight decay of $5e-4$ and a momentum of 0.9. The initial learning rate is set to $1e-3$ and is decreased by a factor of 10 at epoch 5 and 10. For data augmentation, we apply a random crop with 321×321 size, random horizontal flipping, and random color jittering. We use a batch size of 5 and train the classification network for 15 epochs.

We optimize the refinement network for the refinement learning with MSE loss using Adam (Kingma and Ba 2014) optimizer with a learning rate of $1e-4$. The batch size is 5, the total training epoch is 15, and the learning rate is dropped by a factor of 10 at epoch 5 and 10. We apply the data same augmentation strategy as in the classification network.

For the segmentation network, we experiment with three architectures: DeepLab-Large-FOV (Chen et al. 2014) with

suppression controller	mIoU
learnable	62.9%
non-learnable with $\delta=0.90$	51.9%
non-learnable with $\delta=0.80$	56.0%
non-learnable with $\delta=0.70$	58.7%
non-learnable with $\delta=0.60$	62.3%
non-learnable with $\delta=0.55$	62.8%
non-learnable with $\delta=0.50$	62.3%
non-learnable with $\delta=0.40$	59.6%

Table 1: Effect of the two types of controllers and the hyperparameter δ for the non-learnable controller.

VGG-16 and DeepLab-Large-FOV with ResNet-101 (He et al. 2016) backbones, and DeepLab-ASPP (Chen et al. 2017) with ResNet-101 backbone. When generating pseudo segmentation labels, we empirically choose $\alpha = 0.2$ for object cues and $\beta = 0.06$ for background cues. Our method is implemented on Pytorch (Paszke et al. 2017). We use the DeepLab-Large-FOV code¹ and DeepLab-ASPP code² implemented based on the Pytorch framework, following the same hyperparameter settings for training and the conditional random field (CRF) (Krähenbühl and Koltun 2011) as the original publications. All experiments are performed on NVIDIA TITAN XP.

Analysis

To analyze the effectiveness of the proposed method, we conduct several experiments. Following the convention of weakly-supervised semantic segmentation, we measure the mIoU score of the segmentation network outputs to evaluate the quality of our localization maps. For all experiments in this section, we adopt the DeepLab-Large-FOV with VGG-16 as the segmentation network and measure the mIoU score on the VOC 2012 validation set.

Suppression controller. In the method section, we introduced two types of suppression controller: learnable and non-learnable controller. We investigate both controllers and the effect of the parameter δ for the non-learnable controller using visualization and quantitative analysis. For this anal-

¹<https://github.com/wangleihitcs/DeepLab-V1-PyTorch>

²<https://github.com/kazuto1011/deeplab-pytorch>




Figure 5: Visualization of feature maps on each layer. Note that element-wise averaging and normalization are applied to feature maps of each layer for visualization.

ysis, we plugin the DRS to all layers of the classification network and skip the refinement learning procedure for precise effect analysis. In the case of a non-learnable controller, we set the same δ for all layers.

Firstly, we analyze the effect of δ for the non-learnable controller and compare the output localization maps of each δ in Figure 4. When δ is 0.90, the localization map is mostly activated in the head of the cat. Consequently, the gap between discriminative regions and adjacent non-discriminative regions is large, resulting in sparse localization maps. As the δ gets smaller, activation at the body of the cat becomes higher, and the activation gap between the head and the body of the cat is smaller. This indicates as the δ value decreases, the discriminative regions are further suppressed and the gap between discriminative regions and non-discriminative regions becomes smaller, resulting in dense localization maps. However, if the δ value is too low (*i.e.*, too much suppression), the gap between the background and the foreground becomes very small, resulting in a noisy localization map as shown in the rightmost result of Figure 4. Therefore, it is important to set an appropriate δ value for the non-learnable controller. The quantitative results in Table 1 support our arguments. The non-learnable controller with $\delta = 0.55$ achieves better performance than that of $\delta = 0.90$ (62.8% *v.s.* 51.9%), but in the case of over-suppression, *e.g.* the non-learnable controller with $\delta = 0.40$, the performance is rather worse than that of $\delta = 0.55$ (59.6% *v.s.* 62.8%). Through this experiment, we found that $\delta = 0.55$ yields the best mIoU performance.

In the case of a learnable controller, it suppresses without generating much noise, creating moderately dense localization maps as in the leftmost of Figure 4. Compared to the non-learnable controller with $\delta = 0.55$, the learnable controller produces similar mIoU performance (62.9% *v.s.* 62.8%) and localization maps. However, the classification accuracy of the learnable controller is much higher (72.6% *v.s.* 68.7%). From these results, we can notice that the learnable controller adaptively balances between the discriminative feature extraction power and the suppression power, whereas the non-learnable controller forcibly increases the suppression power at the expense of the feature extraction power. Note that the learnable controller is free from hyperparameter (*i.e.*, δ) tuning, but requires additional training parameters (from 21.8M to 24.4M training parameters).

layer1	layer2	layer3	layer4	layer5	layer6	mIoU
-	-	-	-	-	-	50.1%
-	-	-	-	-	✓	60.8%
-	-	-	-	✓	✓	62.2%
-	-	-	✓	✓	✓	62.8%
-	-	✓	✓	✓	✓	62.9%
-	✓	✓	✓	✓	✓	62.7%
✓	✓	✓	✓	✓	✓	62.9%
✓	✓	✓	✓	✓	-	58.2%
✓	✓	✓	✓	-	-	53.6%

Table 2: Effect of DRS in each layer. ✓ means DRS is applied.

suppression controller	without refine	with refine
learnable	62.9%	63.5%
non-learnable ($\delta=0.55$)	62.8%	63.6%

Table 3: Effect of localization map refinement learning.

Effect of DRS on each layer. To observe the effect of DRS on each layer, we employ two analytical methods: visualization and quantitative analysis. For this analysis, we use the DRS module with the learnable controller and skip refinement learning for precise effect analysis.

For visualization, we apply element-wise averaging and normalization from 0 to 1 on feature maps of every layer. Figure 5 shows the visualization results of the original CAM and our DRS-plugged classification network. In lower-level layers (*i.e.*, from layer1 to layer3), we notice that the effect of DRS is minor because a network mainly focuses on the local features (*e.g.*, edge) where the gap between the discriminative and adjacent non-discriminative regions tends to be extremely large. Meanwhile, in higher-level layers (*i.e.*, from layer4 to layer6), a network mostly focuses on the global features (*e.g.*, head of a bird) where the gap between the discriminative and adjacent non-discriminative regions is relatively small. In this case, the effect of DRS becomes significant because it suppresses the activation of discriminative regions and expands the attention to non-discriminative regions.

For quantitative analysis, we plug in and out DRS at each layer and evaluate the performance of each case. The results in Table 2 show that the more we plugin the DRS at higher-level layers, the higher performance (from 50.1% to 62.9%). On the other hand, applying DRS in lower-level layers has little effect (62.9% *v.s.* 62.7%). In addition, when we plug-out the DRS at higher-level layers, the performance significantly decreases (from 62.9% to 58.2% and 53.6%). From these results, we can conclude that DRS is more effective to produce dense localization maps when applied in higher-level layers.

Improvement through refinement learning. As mentioned in the method section, localization map refinement learning gives a self-enhancement effect to produce high-quality dense localization maps as in Figure 1. The improved mIoU performance is reported in Table 3 (+0.6% and +0.8%). In addition, Figure 6 shows some segmentation results on the PASCAL VOC 2012, where both DRS and DRS† show sat-

Method	\mathcal{S}	val	test
Segmentation Network : DeepLab-Large-FOV (VGG-16)			
AE-PSL (Wei et al. 2017)	✓	55.0%	55.7%
GAIN (Li et al. 2018)	✓	55.3%	56.8%
MCOF (Wang et al. 2018)	✓	56.2%	57.6%
AffinityNet (Ahn and Kwak 2018)	-	58.4%	60.5%
SeeNet (Hou et al. 2018)	✓	61.1%	60.7%
MDC (Wei et al. 2018)	✓	60.4%	60.8%
RRM (Zhang et al. 2020)	-	60.7%	61.0%
FickleNet (Lee et al. 2019)	✓	61.2%	61.8%
OAA (Jiang et al. 2019)	✓	63.1%	62.8%
ICD (Fan et al. 2020)	✓	64.0%	63.9%
BES (Chen et al. 2020)	-	60.1%	61.1%
Ours (learnable)	✓	63.5%	64.5%
Ours (non-learnable)	✓	63.6%	64.4%
Segmentation Network : DeepLab-Large-FOV (ResNet-101)			
MCOF (Wang et al. 2018)	✓	60.3%	61.2%
SeeNet (Hou et al. 2018)	✓	63.1%	62.8%
AffinityNet (Ahn and Kwak 2018)	-	61.7%	63.7%
FickleNet (Lee et al. 2019)	✓	64.9%	65.3%
RRM (Zhang et al. 2020)	-	66.3%	65.5%
OAA (Jiang et al. 2019)	✓	65.2%	66.4%
ICD (Fan et al. 2020)	✓	67.8%	68.0%
Ours (learnable)	✓	66.5%	67.5%
Ours (non-learnable)	✓	66.8%	67.4%
Segmentation Network : DeepLab-ASPP (ResNet-101)			
DSRG (Huang et al. 2018)	✓	61.4%	63.2%
BES (Chen et al. 2020)	-	65.7%	66.6%
SGAN (Yao and Gong 2020)	✓	67.1%	67.2%
Ours (learnable)	✓	70.4%	70.7%
Ours (non-learnable)	✓	71.2%	71.4%

Table 4: Comparison of state-of-the-art weakly-supervised semantic segmentation methods on the Pascal VOC 2012 dataset. \mathcal{S} means the saliency map is used for extra guidance.

isfactory results, but DRS \dagger leads to better segmentation results. Note that the learnable controller is used for Figure 1 and 6.

State-of-the-arts Comparison

We compare our approach (DRS) with other state-of-the-art weakly-supervised semantic segmentation methods that use only image-level labels as supervision. For comparison, we apply the DRS module to all layers of the classification network and perform refinement learning. We report the performances of both learnable controller and non-learnable controller with $\delta = 0.55$. Table 4 shows the mIoU performance comparison on the PASCAL VOC 2012 validation set and test set. We fairly compare the performance of each of the three architectures of the semantic segmentation network with other works using the same network. Note that \mathcal{S} in Table 4 indicates whether the saliency map is used as extra guidance.

As shown in Table 4, DRS outperforms erasing-based methods (*e.g.*, AE-PSL (Wei et al. 2017), GAIN (Li et al. 2018), SeeNet (Hou et al. 2018)), showing that suppression is more effective than erasing. Compared to the recent state-of-the-art methods, we achieve competitive performance despite our simplicity. In contrast to some works




Figure 6: Qualitative results on the PASCAL VOC 2012 validation set.

(*e.g.*, DSRG (Huang et al. 2018), FickleNet (Lee et al. 2019), AffinityNet (Ahn and Kwak 2018), BES (Chen et al. 2020)) where CRF in the training stage slows down the training process, our method does not apply CRF during learning, thereby achieving high performance with short training time. Although ICD (Fan et al. 2020) achieves higher mIoU scores using an intra-class discriminator approach for separating foreground and background within the same image-level class, it requires a careful training strategy for stable optimization. Unlike these methods, our approach enables fast and stable training procedure and is the simplest and the most effective way to achieve high segmentation performance.

The highlighted rows in Table 4 show that the learnable and non-learnable controllers are both effective, with only a marginal difference in performance. As mentioned in the analysis section, there is a trade-off between the two controller types, so we can choose depending on the situation.

Conclusion

In this paper, we propose a novel approach called DRS for enlarging the object regions highlighted by localization maps. DRS propagates the initial attention to non-discriminative regions, generating dense localization maps. The main advantage of our approach is that it is intuitive, efficient, and easily applicable to any classification network. Together with refinement learning, our proposed method successfully generates dense segmentation labels that cover the entire target objects. When applied to a weakly-supervised segmentation task, it achieves 71.4% mIoU on pascal VOC segmentation benchmark using only image-level labels as weak supervision.

References

- Ahn, J.; and Kwak, S. 2018. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4981–4990.
- Chen, L.; Wu, W.; Fu, C.; Han, X.; and Zhang, Y. 2020. Weakly Supervised Semantic Segmentation with Boundary Exploration. In *European Conference on Computer Vision*, 347–362. Springer.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2014. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* 40(4): 834–848.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Everingham, M.; Eslami, S. M. A.; Gool, L. V.; Williams, C. K. I.; Winn, J. M.; and Zisserman, A. 2014. The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision* 111: 98–136.
- Fan, J.; Zhang, Z.; Song, C.; and Tan, T. 2020. Learning Integral Objects With Intra-Class Discriminator for Weakly-Supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4283–4292.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hou, Q.; Jiang, P.; Wei, Y.; and Cheng, M.-M. 2018. Self-erasing network for integral object attention. In *Advances in Neural Information Processing Systems*, 549–559.
- Huang, Z.; Wang, X.; Wang, J.; Liu, W.; and Wang, J. 2018. Weakly-supervised semantic segmentation network with deep seeded region growing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7014–7023.
- Jiang, P.-T.; Hou, Q.; Cao, Y.; Cheng, M.-M.; Wei, Y.; and Xiong, H.-K. 2019. Integral Object Mining via Online Attention Accumulation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2070–2079.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krähenbühl, P.; and Koltun, V. 2011. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in neural information processing systems*, 109–117.
- Lee, J.; Kim, E.; Lee, S.; Lee, J.; and Yoon, S. 2019. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5267–5276.
- Li, K.; Wu, Z.; Peng, K.-C.; Ernst, J.; and Fu, Y. 2018. Tell me where to look: Guided attention inference network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9215–9223.
- Liu, J.-J.; Hou, Q.; Cheng, M.-M.; Feng, J.; and Jiang, J. 2019. A simple pooling-based design for real-time salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3917–3926.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch. In *NIPS-W*.
- Shimoda, W.; and Yanai, K. 2019. Self-supervised difference detection for weakly-supervised semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 5208–5217.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Wang, X.; You, S.; Li, X.; and Ma, H. 2018. Weakly-supervised semantic segmentation by iteratively mining common object features. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1354–1362.
- Wei, Y.; Feng, J.; Liang, X.; Cheng, M.-M.; Zhao, Y.; and Yan, S. 2017. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1568–1576.
- Wei, Y.; Xiao, H.; Shi, H.; Jie, Z.; Feng, J.; and Huang, T. S. 2018. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7268–7277.
- Yao, Q.; and Gong, X. 2020. Saliency Guided Self-Attention Network for Weakly and Semi-Supervised Semantic Segmentation. *IEEE Access* 8: 14413–14423.
- Zhang, B.; Xiao, J.; Wei, Y.; Sun, M.; and Huang, K. 2020. Reliability does matter: An end-to-end weakly supervised semantic segmentation approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12765–12772.
- Zhang, X.; Wei, Y.; Feng, J.; Yang, Y.; and Huang, T. S. 2018. Adversarial complementary learning for weakly supervised object localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1325–1334.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.