

In-sample Contrastive Learning and Consistent Attention for Weakly Supervised Object Localization

Minsong Ki¹, Youngjung Uh^{2,3}, Wonyoung Lee³, and Hyeran Byun^{1,3}

¹ Department of Computer Science, Yonsei University, Seoul, Republic of Korea

² Department of Applied Information Engineering, Yonsei University, Seoul, Republic of Korea

³ Department of Artificial Intelligence, Yonsei University, Seoul, Republic of Korea
{kms2014,yj.uh,lwy8555,hrbyun}@yonsei.ac.kr

Abstract. Weakly supervised object localization (WSOL) aims to localize the target object using only the image-level supervision. Recent methods encourage the model to activate feature maps over the entire object by dropping the most discriminative parts. However, they are likely to induce excessive extension to the backgrounds which leads to over-estimated localization. In this paper, we consider the background as an important cue that guides the feature activation to cover the sophisticated object region and propose contrastive attention loss. The loss promotes similarity between foreground and its dropped version, and, dissimilarity between the dropped version and background. Furthermore, we propose foreground consistency loss that penalizes earlier layers producing noisy attention regarding the later layer as a reference to provide them with a sense of backgroundness. It guides the early layers to activate on objects rather than locally distinctive backgrounds so that their attentions to be similar to the later layer. For better optimizing the above losses, we use the non-local attention blocks to replace channel-pooled attention leading to enhanced attention maps considering the spatial similarity. Last but not least, we propose to drop background regions in addition to the most discriminative region. Our method achieves state-of-the-art performance on CUB-200-2011 and ImageNet benchmark datasets regarding **top-1 localization accuracy** and **MaxBoxAccV2**, and we provide detailed analysis on our individual components. The code will be publicly available online for reproducibility.

1 Introduction

Fully supervised approaches have demonstrated excellent performance by training convolution neural network (CNN) with human annotations, *e.g.*, bounding box for object localization, pixel-wise class labels for semantic segmentation [1–5]. However, they cost huge human labor to obtain accurate annotations. Therefore, weakly supervised approaches that use only image-level supervision have received significant attention over the various computer vision tasks [6–12].




Fig. 1. Comparison of methods for generating activation maps on the CUB [16] dataset. We display the final results obtained by ADL [9] (first row), SPG [8] (second row), and our method (last row). The red boxes are the ground-truth and the green boxes are the predicted ones. Activation maps are illustrated in heatmap color scale. ADL tries to activate more on less discriminative parts but ends in excessive extension to background. SPG tries to suppress background but still over-estimates the object regions. In contrast, our method covers the whole object delicately without extending to background.

Especially, weakly supervised object localization (WSOL) is a challenging task that pursues both classification and the localization of the target object where the training datasets provide only the class labels.

For example, Zhou et al. [6] generate class activation maps (CAM) using the classification model with a global average pooling (GAP). CAM highlights the class-specific discriminative regions in a given image [7–9, 13]. The crucial pitfall of the activation maps is that it focuses on discriminative parts (*e.g.*, the head of a bird) rather than including the full extent of the object. To mitigate this limitation, recent methods [9, 14, 15] propose to erase the most discriminative parts by thresholding to spread out the activations to less discriminative regions. However, they are likely to induce excessive extension to the backgrounds which over-estimates bounding boxes (Fig. 1).

In this paper, we propose four ingredients for more accurate attention over the entire object: contrastive attention loss, foreground consistency loss, non-local attention block and dropped foreground mask. The contrastive attention loss draws the foreground feature and its erased version close together, and pushes the erased foreground feature away from the background feature (Sec. 3.2). It helps the learned representation to reflect only the object region rather than the backgrounds which are usually helpful for classification but harmful to localization. The foreground consistency loss penalizes disagreement of attentions between layers to provide early layers with a sense of backgroundness (Sec. 3.3). While usual low-level features are activated on locally distinctive regions (*e.g.*,

edges) regardless of the presence of the objects, adding foreground consistency loss boosts the activations on the object regions while suppressing the activations on the background regions. Furthermore, we apply the non-local attention blocks to produce enhanced attention maps considering the similarity between locations in a feature map (Sec. 3.4). It allows boosting weights on the regions having similar features with the most discriminative parts to pursue correct activation. Last but not least, we propose a dropped foreground mask which drops the background region as well as the most discriminative region. It prohibits the model from excessively spreading attention to backgrounds.

Our method achieves state-of-the-art performance in terms of the conventional **top-1 localization accuracy** and the **MaxBoxAccV2** [17].

In summary, our main contributions are:

- We propose a contrastive attention loss that favors similarity between foreground feature and its dropped version and dissimilarity between the dropped foreground feature and background feature.
- We propose a foreground consistency loss that provides a sense of localization to earlier layers by guiding their features to be consistent with a high-level layer.
- We propose a dropped foreground mask which drops the background region and the most discriminative region.
- Our method achieves state-of-the-art performance on CUB-200-2011 and ImageNet benchmark datasets in terms of **top-1 localization accuracy** and **MaxBoxAccV2**.

2 Related Work

Weakly supervised object localization (WSOL). Given only the class labels with the images, most of the WSOL methods train a classifier and extract CAM [6]. CAM indicates the strength of activation in every location in the feature map to stimulate the corresponding class [7–9, 13]. Recent methods [6–10, 13] propose erasing the most discriminative region of the feature map to spread out the activations to the regions which are less discriminative but still in the object. Hide-and-Seek (HaS) [13] divides a training image into a grid of evenly-divided patches and selects a random patch to be hidden. Adversarial complementary learning (ACoL) [7] and attention-based dropout layer (ADL) [9] partially drop the most discriminative region by thresholding on the feature map. MEIL [18] runs two branches, one with erasing and one without erasing, and impose both branch with classification task. These approaches guide the models to discover previously neglected object regions. Our method steps further to consider background as a region to drop so that the model does not spread the activation excessively to the background.

Several methods have been proposed to suppress the background and localize the whole object. Zhang et al. [8] present self-produced guidance (SPG) that generates three pixel-wise masks (foreground, background, and undefined areas). Each mask is used as auxiliary supervision. However, it requires to find the

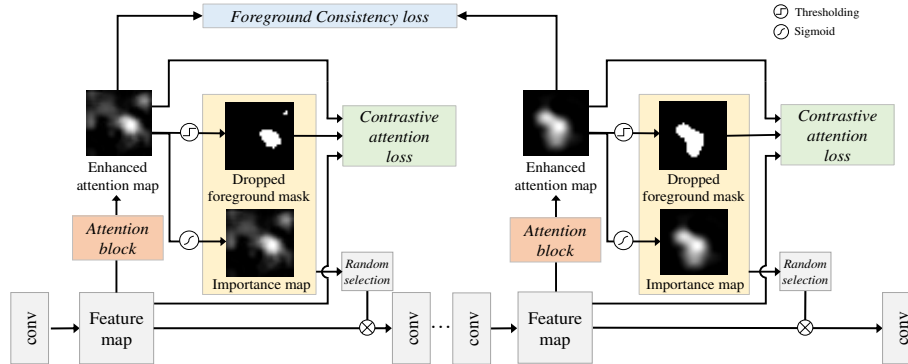


Fig. 2. Overview of the proposed method. The non-local attention block generates the enhanced attention map reflecting the similarity between locations. We create a dropped foreground mask and an importance map using thresholding and sigmoid activation, respectively. The selected map is multiplied with the input feature to feed the next layer. Foreground consistency loss encourages the consistency between the early and last layer. We calculate the contrastive attention loss at each convolution layer where our non-local attention block is inserted.

optimal six hyperparameters for producing the three masks. We also focus on the background but introduce simpler and more effective way.

Yang et al. [10] use a non-local block following every convolution-pooling block. While their non-local blocks are inserted within the main stem of the network, our non-local attention blocks are branch from the main stem and produce attention maps to be multiplied to the main convolutional features at chosen layers.

Contrastive visual representation learning. Contrastive learning [19] tries to distinguish similar and dissimilar pairs of samples by embedding the samples as feature representations. Recent self-supervised learning methods [20, 21] learn representations by maximizing agreement between differently augmented views of the same image. They also consider the different images to minimize the agreement for negative pairs.

Inspired by [20, 21], we define a contrastive prediction task for WSOL. Instead of building similar and dissimilar pairs of image samples, we regard the foreground region except for the most discriminative part (*i.e.*, the dropped foreground) as an anchor, and build the positive pair with the original foreground and the negative pair with the background. Our contrastive objective does not require a large batch size or large queue because it finds the pairs within an image. Separating the foreground representation and the background representation is suitable for WSOL task.

3 Proposed Method

This section describes elements of the proposed method and how we employ them on the networks.

3.1 Network overview

As shown in Fig. 2, we augment classification network with the non-local attention blocks (Sec. 3.4) and train it with the contrastive attention loss (Sec. 3.2) and the foreground consistency loss (Sec. 3.3). The non-local attention block receives a feature map \mathbf{F} and provides an enhanced attention map \mathbf{A} which becomes an importance map $\tilde{\mathbf{A}}$ through sigmoid activation and a dropped foreground mask \mathbf{M}_{dfg} by thresholding (Eq. 1). The dropped foreground mask or the importance map is randomly chosen based on a `drop_rate` and the chosen one is applied to the input feature by pixel-wise multiplication (element-wise multiplication with broadcasting over the channel dimension). The importance map is not to be dropped but to be applied to the feature map. The dropped foreground mask encourages activation of the input feature on less discriminative parts except background to maximize classification accuracy without losing localization accuracy, while the importance map rewards higher activation on the most discriminative part.

In the attention branch, the enhanced attention map and the dropped foreground mask from a non-local attention block are used to compute the contrastive attention loss. In addition, the enhanced attention maps from multiple non-local attention blocks are used to compute the foreground consistency loss.

The differences with ADL [9] in the forward process are that we use the dropped foreground mask instead of drop mask and the attention map is produced by our non-local attention block instead of vanilla convolutional feature. Fig. 3 illustrates the importance map, our dropped foreground mask, and the drop mask in [9]. Our dropped foreground mask \mathbf{M}_{dfg} is defined by:

$$\mathbf{M}_{\text{dfg}} = \mathbb{1}[\mathbf{A} < \theta_{\text{fg}}] \wedge \mathbb{1}[\mathbf{A} > \theta_{\text{bg}}], \quad (1)$$

where $\mathbb{1}$ denotes a matrix with the same shape with the input having ones according to the logical operation, \wedge denotes logical **and** operation, and θ 's are the pre-defined thresholds. Unlike the drop masks from ACoL [7] and ADL [9], our dropped foreground mask remedies excessive expansion of activation on the backgrounds by further erasing background regions in the mask.

The contrastive attention loss and the foreground consistency loss are computed wherever the attention maps are extracted.

3.2 Contrastive attention loss

Contrastive loss [20] is a function whose value is low when a query is similar to its equivalent instance and dissimilar to its different instances. Likewise, we design a contrastive attention loss whose value is low when a dropped foreground feature




Fig. 3. Examples of the importance map $\tilde{\mathbf{A}}$, the drop mask from [9], and our dropped foreground mask \mathbf{M}_{dfg} .

\mathbf{z}_{dfg} is similar to a foreground feature \mathbf{z}_{fg} and dissimilar to a background feature \mathbf{z}_{bg} (Fig. 4). The features \mathbf{z}_α are obtained by masked global average pooling of $\mathbf{F} \odot \mathbf{M}_\alpha$ where

$$\begin{aligned} \mathbf{M}_{\text{fg}} &= \mathbb{1}[\mathbf{A} > \theta_{\text{bg}}], \\ \mathbf{M}_{\text{bg}} &= \mathbb{1}[\mathbf{A} < \theta_{\text{bg}}], \end{aligned} \quad (2)$$

and the masked global average pooling is spatial average pooling of the pixels whose value on the mask is 1. Then, the contrastive attention loss is given by

$$\mathcal{L}_{\text{ca}} = \mathbb{E}_{\mathbf{x}} [|(d(\mathbf{z}_{\text{dfg}}, \mathbf{z}_{\text{fg}}) - d(\mathbf{z}_{\text{dfg}}, \mathbf{z}_{\text{bg}}) + m)|_+], \quad (3)$$

where $[\cdot]_+ = \max(\cdot, 0)$ and $d(\cdot, \cdot)$ denotes L_2 distance in auxiliary 128-dimensional embedding by 1×1 convolution. m denotes the margin.

Our contrastive attention loss guides the attention map to spread until it reaches boundary because including backgrounds in the attention map is penalized by the dissimilarity term. In addition, the similarity term favors homogenous features between the most discriminative part and less discriminative parts in the foreground region. Our contrastive attention loss does not require mining positive and negative samples as in triplet loss [22] nor managing large negative samples [20, 21]. Since we regard the masked features \mathbf{z}_{dfg} , \mathbf{z}_{fg} and \mathbf{z}_{bg} from one image as an anchor, a positive sample and a negative sample, respectively.

3.3 Foreground consistency loss

Attention maps roughly are the magnitude of activation on every location. Convolutions in early layers activate more on locally distinctive regions such as edges and corners [23], without inspecting the entire extent of objects due to their limited receptive field. To relieve this problem, we propose a foreground consistency loss that encourages attention maps from early layers to resemble later layers (Fig. 2).

Let A_i and A_j are the attention maps from early and later layers, respectively. Then we define the foreground consistency loss as:

$$\mathcal{L}_{fc} = \|A_i - A_j\|_2^2, \quad (4)$$

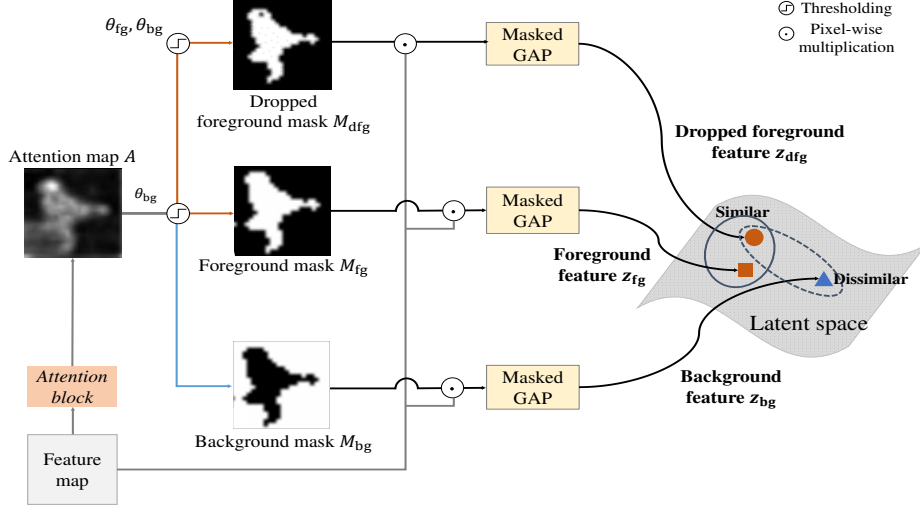


Fig. 4. The details of the contrastive attention loss, where A denotes an enhanced attention map from a non-local attention block. We generate three maps and features to compare the similarity in embedding space. \odot denotes pixel-wise multiplication. The contrastive attention loss is computed on an embedding space.

where $\|\cdot\|_2$ denotes L_2 norm of a matrix.

Gradients from the foreground consistency loss only run through the early layer to achieve the abovementioned goal. It reduces the noisy activations outside the object and boosts activations in the object.

3.4 Non-local attention block

In order to provide additional capacity for the network to produce a correct attention map, we employ non-local block [24] instead of average channel pooling of the convolutional features [9, 25]. Given a feature map, our non-local attention block embeds it into three different embeddings and outputs spatial summation of the third one weighted by similarity between the first two embeddings. Then the enhanced attention map is defined by its channel-pooled result.

Specifically, the block receives a feature map $x \in \mathbb{R}^{C \times H \times W}$ from a convolution layer. For simplicity, we omit the mini-batch dimension. We define $f(x), g(x) \in \mathbb{R}^{\tilde{C} \times H \times W}$, $z(x) \in \mathbb{R}^{C \times H \times W}$ that use 1×1 convolution layer for embedding. Then, $f(x), g(x)$ and $z(x)$ are reshaped to $f(x), g(x) \in \mathbb{R}^{\tilde{C} \times HW}$, $z(x) \in \mathbb{R}^{C \times HW}$, respectively.

The enhanced attention map A is given by:

$$A = \mathbb{E}_C [\text{Softmax}(f(x)^T g(x)) \odot z(x)], \quad (5)$$

where \mathbb{E}_C denotes average pooling over the channel dimension.

The non-local attention block produces the enhanced attention map regarding similarities between locations. It unleashes the receptive field of the layer and provides an additional clue for deciding where to attend. Our non-local attention block is different from [10] in that we organize it only when generating several enhanced attention maps. Yang et al. [10] apply the non-local module to all layers in the main branch with residual connection.

3.5 Training and Inference

We train the base network and non-local attention block with the full objective:

$$\mathcal{L}_{total} = \mathcal{L}_{cls} + \mathcal{L}_{ca} + \mathcal{L}_{fc} \quad (6)$$

We employ a GAP layer at the end of the network to produce softmax output \hat{y} and compute classification loss given the one-hot ground truth label y :

$$\mathcal{L}_{cls} = \text{CrossEntropy}(\hat{y}, y) \quad (7)$$

All network weights are updated wherever all losses send their gradients towards the input, while the foreground consistency loss does not convey gradients to its reference layer.

Our non-local attention block is applied only during training and deactivated in the testing phase. The input image goes through only the vanilla model to produce the class assignment. Then we follow [17] to extract the heatmap which leads to the bounding boxes by thresholding and its connected multi-contour.

4 Experiments

4.1 Experimental setup

Datasets. We evaluate the proposed method on two benchmark datasets: CUB-200-2011 [16] and ILSVRC [26] (ImageNet) for WSOL task, from which only the image-level labels are used in training. Many weak-supervision methods have used full supervision to some extent, directly or indirectly for hyperparameter tuning. Since the amount of full supervision used for hyperparameter tuning is not consistent, it has been ambiguous using the previous evaluation metric for a fair comparison. We follow the recent evaluation metric [17] which fixes the amount of full supervision only for hyperparameter search. Each dataset is divided into three subsets: **train-weaksup**, **train-fullsup** and **test**. The **train-weaksup** includes images only with the class labels for training. The **train-fullsup** contains images with full supervision, which has bounding boxes as well. It is left free for the users to use the **train-fullsup** for hyperparameter search. They collected five images per class (total 1,000 images) from Flickr for CUB experiments, and ten images per class (total 10,000 images) from ImageNetV2 [27] for ImageNet experiments, respectively. The **test** split for the

Table 1. Hyperparameters (`drop_rate`, γ_{fg} , γ_{bg}) for each backbone.

Backbone	<code>drop_rate</code>	γ_{fg}	γ_{bg}
VGG [29]	0.33	0.72	1.2
InceptionV3 [30]	0.69	0.86	1.2
ResNet50 [31]	0.85	0.95	1.2

final number is the same as the standard WSOL settings on CUB and ImageNet experiments [7–9, 17, 28]. In the CUB dataset, there are 5,994 images for training and 5,794 for testing from 200 bird species. ImageNet consists of 1.2M training images and 10K test images for 1,000 classes. All experimental analyses of the proposed method are conducted on the `test` split of the two abovementioned datasets.

Evaluation metrics. We use `top-1 classification accuracy` and `top-1 localization accuracy`, and `MaxBoxAccV2` [17].

`Top-1 classification accuracy` is the ratio of correct classification. The conventional `top-1 localization accuracy` measures ratio of the samples with the right class and the bounding box of IoU greater than 0.5.

`MaxBoxAcc` measures ratio of the samples with the correct box, while the correctness is defined by an IoU criterion δ at the optimal activation threshold. `MaxBoxAccV2` averages `MaxBoxAcc` at three IoU criterions $\delta \in \{30, 50, 70\}$ to address diverse demands for localization fineness. It is similar to the common GT-known metric but differs in that it evaluates on *three IoUs* by extracting the bounding box with *the optimal score map threshold*. We use % symbol as a *percent point* for mentioning differences on comparisons.

Implementation details. We build the proposed method upon three CNN backbones: VGG16 [29], InceptionV3 [30], and Resnet50 [31]. We need three hyperparameters: `drop_rate` for randomly choosing the importance map or the dropped foreground mask, θ_{fg} and θ_{bg} for thresholding. The threshold θ_{fg} is set to the maximum intensity of \mathbf{A} times pre-defined ratio γ_{fg} . The θ_{bg} is set to average intensity of \mathbf{A} times pre-defined ratio γ_{bg} . The specific values of the hyperparameters for each backbone are shown in Table 1.

The layers, from which the enhanced attention maps are extracted, are chosen to be the same with the baseline method [9]. We also calculate our contrastive attention loss and foreground consistency loss for all layers where attention maps are produced. We set the batch size to 32, weight decay to 0.0001, margin m to 1. The initial learning rate and the momentum of the SGD optimizer are set to 0.001 and 0.9, respectively. We start from loading weights from the model pre-trained on the ImageNet classification [26] and then fine-tuned the network. Our model is implemented using PyTorch and trained using two NVIDIA GeForce RTX 2080 Ti GPUs for approximately three hours. The input images are randomly cropped to 224×224 pixels after being resized to 256×256 pixels. During testing phase, we directly resize the input images to 224×224 .

4.2 Ablation study

We first show detailed experiments to validate effectiveness of each component. We fix the ResNet50 [31] as a backbone and add or remove each component. The experiments are performed on the CUB `test` split. The difference in performance in % represents percent points.

Ablation of the proposed losses. Table 2 shows that both the contrastive attention loss and the foreground consistency loss are the crucial element for the improved performance. Ours without the contrastive attention loss achieves 2% lower performance than the full setting. The loss has positive effect on all three IoU thresholds. The foreground consistency loss also plays an important role of improvement by 0.79%. It especially boosts the accuracy at IoU 0.7. We suggest that the loss helps precisely estimating the location of the object in the early layer by providing the hints from the later layer. Using the both losses leads to balanced improvements over all IoU thresholds. In addition, contrastive attention loss with the normalized temperature-scaled cross-entropy loss (*NT-Xent*) [20, 21] also shows improvements to some extent. Its result can be found in the supplementary material.

Effectiveness of the non-local attention block. If we use the vanilla attention map which is the channel-pooled result of the convolutional feature, the performance drops by 0.62% (the fourth row in Table 2). It shows that considering the relationship between pixels in feature map helps localizing where to attend.

Effectiveness of the dropped foreground mask. Here we validate the effectiveness of replacing the drop mask [9] with the dropped foreground mask \mathbf{M}_{dfg} . Without the replacement, the model achieves 1.12% lower performance than the full setting (the fifth row in Table 2). Also, only replacing the drop mask with the dropped foreground mask improves the performance of the baseline [9] by 0.87%. We suppose that the dropped foreground mask improves ours more than the baseline because the additional two losses and the non-local attention block

Table 2. The ablation study for each element of our method on Resnet50 [31] backbone in terms of `MaxBoxAccV2`. Contrastive: contrastive attention loss. \mathcal{L}_{fc} : foreground consistency loss. Non-local: non-local attention block. \mathbf{M}_{dfg} : dropped foreground mask. All elements contribute to the performance improvement.

Methods	MaxBoxAccV2@IoU (%)			
	0.3	0.5	0.7	Mean
Baseline [9]	91.82	64.78	18.43	58.34 (-4.86)
Ours w/o contrastive	94.79	70.84	17.98	61.20 (-2.00)
Ours w/o \mathcal{L}_{fc}	96.51	72.14	18.57	62.41 (-0.79)
Ours w/o non-local	95.80	71.62	20.32	62.58 (-0.62)
Ours w/o \mathbf{M}_{dfg}	96.42	72.21	17.62	62.08 (-1.12)
Ours (full)	96.18	72.79	20.64	63.20

Table 3. Performance comparison regarding at which layer to insert our attention block. The contrastive attention loss and the foreground consistency loss are in use for all cases. The `conv_5_3` layer is fixed as the reference layer and its performance is left empty because the foreground consistency loss requires at least two layers. We add the layers from later to earlier and report their performance in a cumulative setting.

Location	Top-1 classification	MaxBoxAccV2
<code>conv_5_3</code>	-	-
+ <code>pool_4</code>	74.23	64.67
+ <code>pool_3</code>	73.35	66.72
+ <code>pool_2</code>	66.26	63.25
+ <code>pool_1</code>	62.68	62.18

provide an extra guide for better importance map.

Location of our attention block. We investigate the influence of where to insert our non-local attention block on VGG16 [29], and report the results in Table 3. The `conv_5_3` layer is fixed as the reference layer for the foreground consistency and its preceding layers are added one by one cumulatively. The setting with top three layers, which is the same as the baseline [9], achieves the best performance in terms of `MaxBoxAccV2`. Adding the attention blocks on `pool_1` and `pool_2` layers decreases the performance. We suppose that the reason is their small receptive field which leads to noisy activations on extremely locally salient regions. Hence, we do not use the attention mechanism on the two earliest layers.

Table 4. `MaxBoxAccV2` [17] comparison with the state-of-the-art methods. The results for each backbone represent the average of the three IoU thresholds 0.3, 0.5, and 0.7. VGG: VGG16 [29]. Inc: InceptionV3 [30]. Res: ResNet50 [31]. The best and the second best entries in a column are marked in boldface and italic, respectively.

Methods	ImageNet				CUB-200-2011			
	VGG	Inc	Res	Mean	VGG	Inc	Res	Mean
CAM [6]	60.0	63.4	63.7	62.4	63.7	56.7	63.0	61.1
HaS [13]	60.6	63.7	63.4	<i>62.6</i>	63.7	53.4	64.7	60.6
ACoL [7]	57.4	63.7	62.3	61.2	57.4	56.2	66.5	60.0
SPG [8]	59.9	63.3	63.3	62.2	56.3	55.9	60.4	57.5
ADL [9]	59.8	61.4	63.7	61.7	66.3	58.8	58.4	<i>61.1</i>
CutMix [28]	59.4	63.9	63.3	62.2	62.3	57.5	62.8	60.8
Ours	61.3	62.8	65.1	63.1	66.7	60.3	63.2	63.4

Table 5. Detailed **MaxBoxAccV2** [17] comparison with the runner-up methods on each dataset. We compare ours and the second best methods on each dataset and each backbone in terms of **MaxBoxAccV2** including individual measures on the three IoU criterions. Mean indicates that the average value of the three IoU thresholds. VGG: VGG16 [29]. Inc: InceptionV3 [30]. Res: ResNet50 [31]. Bold texts denote the best performance in each column.

	Method Backbone		Top-1	MaxBoxAccV2@IoU (%)			Mean
				0.3	0.5	0.7	
CUB-200-2011	ADL	VGG	54.95	97.72	78.06	23.04	66.28
	Ours	VGG	73.35	96.20 (-1.5)	77.20 (-0.8)	26.75 (+3.7)	66.72 (+0.5)
	ADL	Inc	41.03	93.77	65.79	16.86	58.81
	Ours	Inc	64.01	95.89 (+2.1)	67.93 (+2.2)	17.20 (+0.4)	60.34 (+1.2)
	ADL	Res	66.60	91.82	64.76	18.43	58.34
	Ours	Res	80.35	96.18 (+4.3)	72.79 (+8.0)	20.64 (+2.2)	63.20 (+4.9)
ImageNet	HaS	VGG	68.26	80.72	62.15	38.89	60.59
	Ours	VGG	69.21	81.45 (+0.7)	63.20 (+1.1)	39.35 (+2.2)	61.33 (+0.8)
	HaS	Inc	69.07	83.95	66.27	40.94	63.72
	Ours	Inc	71.31	82.44 (-1.5)	65.21 (-1.0)	40.87 (-0.1)	62.84 (-0.9)
	HaS	Res	75.39	83.71	65.22	41.26	63.40
	Ours	Res	76.54	84.26 (+0.5)	67.62 (+2.4)	43.58 (+2.3)	65.15 (+1.7)

4.3 Comparison with state-of-the-art methods

We compare our method with the state-of-the-art WSOL methods in terms of the **MaxBoxAccV2** [17], **top-1 localization** and **top-1 classification accuracy**.

MaxBoxAccV2 [17]. Table 4 shows comparison of **MaxBoxAccV2** across all competitors on ImageNet and CUB. Our method outperforms all existing methods in terms of **MaxBoxAccV2** (Mean) and most of backbone choices. Table 5 shows detailed comparison with the runner-up methods of each dataset. Our method boosts performance especially when IoU criterions are 0.5 and 0.7 except when Inception network is the backbone. Our method exhibits the largest improvement when employed on ResNet backbone.

Top-1 localization accuracy. **Top-1 localization accuracy** on the ImageNet and CUB datasets is shown in Table 6. Our model outperforms the state-of-the-art methods on most settings. Note that we do not perform hyperparameter tuning using the **train-fullsup** split following the competitors for a fair comparison.

Top-1 classification accuracy. Table 7 compares our method with the state-of-the-art methods in terms of **top-1 classification accuracy**. While some other methods compromise classification accuracy for improving localization, our method achieves the best **MaxBoxAccV2** and localization accuracy without damaging the classification accuracy.

Table 6. Conventional Top-1 localization accuracy comparison with the state-of-the-art methods. The values are taken from their respective papers. Bold texts denote the best performance in each backbone network.

Methods	ImageNet			CUB-200-2011		
	VGG	Inc	Res	VGG	Inc	Res
CAM [6]	42.8	46.3	-	37.1	43.7	49.4
HaS [13]	-	-	-	-	-	-
ACoL [7]	45.8	-	-	45.9	-	-
SPG [8]	-	48.6	-	-	46.6	-
ADL [9]	44.9	48.7	-	52.4	53.0	-
CutMix [28]	43.5	-	47.3	-	52.5	54.8
MEIL [18]	46.8	49.5	-	57.5	-	-
Ours	47.2	49.3	48.4	57.5	56.1	56.1

Table 7. Top-1 classification performance of the state-of-the-art methods. Hyperparameters for each method are optimally selected for the localization performances on `train-fullsup` split. Bold texts denote the best performance. MEIL does not provide code for reproduction and its values are taken from the paper. Other values are reproduction from [17].

Methods	ImageNet				CUB-200-2011			
	VGG	Inc	Res	Mean	VGG	Inc	Res	Mean
CAM [6]	66.48	70.56	75.05	70.70	50.10	70.70	71.50	64.10
HaS [13]	68.26	69.07	75.39	70.91	75.90	64.50	69.70	70.10
ACoL [7]	64.55	71.81	73.09	69.82	71.80	71.50	71.10	71.40
SPG [8]	67.76	71.12	73.26	70.71	72.10	46.20	50.50	56.30
ADL [9]	67.58	61.17	71.99	66.91	55.00	41.00	66.60	54.20
CutMix [28]	66.36	69.16	75.71	70.41	48.40	71.00	73.00	64.10
MEIL [18]	70.27	73.31	-	-	74.77	74.55	-	-
Ours	69.21	76.54	71.31	72.35	73.40	64.00	80.40	72.60

4.4 Qualitative results

Fig. 1 compares activation maps and estimated bounding boxes from ADL [9], SPG [8] and ours. ADL excessively covers backgrounds because it simply encourages the model to use less discriminative parts, and SPG still over-estimates the bounding boxes although it tries to suppress background. In contrast, our method focuses on the entire object more accurately and estimates tighter bounding boxes. Fig. 5 illustrates more examples from our model. Our method not only spreads out of the most discriminative parts, but also restrains the activations in the object regions. Note that the water and mirrored image of the pelican does not earn large activation even though they are helpful cue for classification (the second row of the second column).

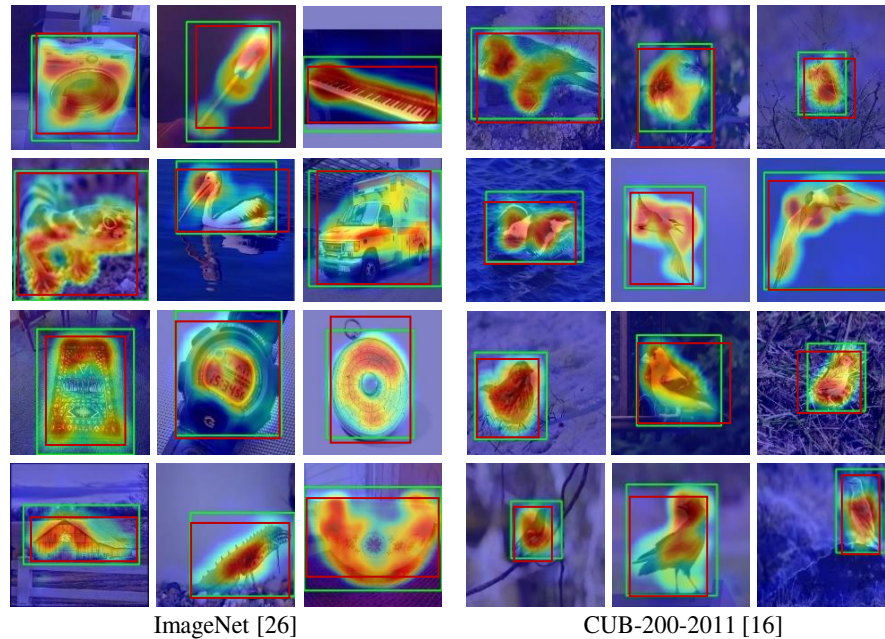


Fig. 5. Qualitative examples of activation map and localization produced by our model on the ImageNet and CUB `test` split. The red boxes are the ground-truth and the green boxes are the predicted ones. These maps output with colors ranging from red (higher importance) to blue (lower importance like a background).

5 Conclusion

In this paper, we consider the background as an important clue for localizing the entire object without excessive coverage and present two novel objective functions. The crucial weakness of the previous methods is that they focus on discriminative parts rather than localizing the whole object, or extend too much on the background. The proposed contrastive attention loss guides the model to spread the attention map within the objects. The foreground consistency loss decreases the activation to backgrounds in the early layers. The generated attention map not only better localizes the target object but also suppresses the background concurrently. In addition, our non-local attention block enhances the attention map with a larger capacity to better optimize the proposed losses. We achieve state-of-the-art performance on ImageNet and CUB-200-2011 datasets and provide detailed analysis on the effects of our individual components.

Acknowledgements. This work was supported by the National Research Foundation of Korea grant funded by Korean government (No. NRF-2019R1A2C2003760) and Artificial Intelligence Graduate School Program (YONSEI UNIVERSITY) under Grant 2020-0-01361. We thank Junsuk choe for his valuable discussion.

References

1. Chen, K., Li, J., Lin, W., See, J., Wang, J., Duan, L., Chen, Z., He, C., Zou, J.: Towards accurate one-stage object detection with ap-loss. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 5119–5127
2. Zhu, C., He, Y., Savvides, M.: Feature selective anchor-free module for single-shot object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 840–849
3. Chen, X., Girshick, R., He, K., Dollár, P.: Tensormask: A foundation for dense object segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 2061–2069
4. Robinson, A., Lawin, F.J., Danelljan, M., Khan, F.S., Felsberg, M.: Learning fast and robust target models for video object segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 7406–7415
5. Li, Z., Zhou, F., Yang, L.: Fast single shot instance segmentation. In: Asian Conference on Computer Vision, Springer (2018) 257–272
6. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2016) 2921–2929
7. Zhang, X., Wei, Y., Feng, J., Yang, Y., Huang, T.S.: Adversarial complementary learning for weakly supervised object localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 1325–1334
8. Zhang, X., Wei, Y., Kang, G., Yang, Y., Huang, T.: Self-produced guidance for weakly-supervised object localization. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 597–613
9. Choe, J., Shim, H.: Attention-based dropout layer for weakly supervised object localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 2219–2228
10. Yang, S., Kim, Y., Kim, Y., Kim, C.: Combinational class activation maps for weakly supervised object localization. In: The IEEE Winter Conference on Applications of Computer Vision. (2020) 2941–2949
11. Son, J., Kim, D., Lee, S., Kwak, S., Cho, M., Han, B.: Forget and diversify: Regularized refinement for weakly supervised object detection. In: Asian Conference on Computer Vision, Springer (2018) 632–648
12. Lee, P., Uh, Y., Byun, H.: Background suppression network for weakly-supervised temporal action localization. In: AAAI. (2020) 11320–11327
13. Kumar Singh, K., Jae Lee, Y.: Hide-and-peek: Forcing a network to be meticulous for weakly-supervised object and action localization. In: Proceedings of the IEEE International Conference on Computer Vision. (2017) 3524–3533
14. Wei, Y., Feng, J., Liang, X., Cheng, M.M., Zhao, Y., Yan, S.: 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. (2017) 1568–1576
15. Hou, Q., Jiang, P., Wei, Y., Cheng, M.M.: Self-erasing network for integral object attention. In: Advances in Neural Information Processing Systems. (2018) 549–559
16. Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., Perona, P.: Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology (2010)

17. Choe, J., Oh, S.J., Lee, S., Chun, S., Akata, Z., Shim, H.: Evaluating weakly supervised object localization methods right. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 3133–3142
18. Mai, J., Yang, M., Luo, W.: Erasing integrated learning: A simple yet effective approach for weakly supervised object localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 8766–8775
19. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). Volume 2., IEEE (2006) 1735–1742
20. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 9729–9738
21. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. arXiv preprint arXiv:2002.05709 (2020)
22. Chen, W., Chen, X., Zhang, J., Huang, K.: Beyond triplet loss: a deep quadruplet network for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2017) 403–412
23. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: European conference on computer vision, Springer (2014) 818–833
24. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2018) 7794–7803
25. Baek, K., Lee, M., Shim, H.: Psynet: Self-supervised approach to object localization using point symmetric transformation. In: AAAI. (2020) 10451–10459
26. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115** (2015) 211–252
27. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? arXiv preprint arXiv:1902.10811 (2019)
28. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE International Conference on Computer Vision. (2019) 6023–6032
29. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
30. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2016) 2818–2826
31. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2016) 770–778