

# Inferring the Class Conditional Response Map for Weakly Supervised Semantic Segmentation

Weixuan Sun, Jing Zhang, Nick Barnes  
Australian National University

{weixuan.sun, jing.zhang, nick.barnes}@anu.edu.au

## Abstract

Image-level weakly supervised semantic segmentation (WSSS) relies on class activation maps (CAMs) for pseudo labels generation. As CAMs only highlight the most discriminative regions of objects, the generated pseudo labels are usually unsatisfactory to serve directly as supervision. To solve this, most existing approaches follow a multi-training pipeline to refine CAMs for better pseudo-labels, which includes: 1) re-training the classification model to generate CAMs; 2) post-processing CAMs to obtain pseudo labels; and 3) training a semantic segmentation model with the obtained pseudo labels. However, this multi-training pipeline requires complicated adjustment and additional time. To address this, we propose a class-conditional inference strategy and an activation aware mask refinement loss function to generate better pseudo labels without re-training the classifier. The class conditional inference-time approach is presented to separately and iteratively reveal the classification network’s hidden object activation to generate more complete response maps. Further, our activation aware mask refinement loss function introduces a novel way to exploit saliency maps during segmentation training and refine the foreground object masks without suppressing background objects. Our method achieves superior WSSS results without requiring re-training of the classifier. <https://github.com/weixuansun/InferCam>

## 1. Introduction

Recent work on 2D image semantic segmentation has achieved great progress via deep fully convolutional neural networks (FCN) [32]. The success of these models [51, 8, 7, 52, 9] comes from large training datasets with pixel-wise labels, which are laborious and expensive to obtain. To relieve the labeling burden, multiple types of weak labels have been studied, including image-level [23, 44, 2, 16], points [3], scribbles [40, 31, 39], and bounding boxes [10, 34, 25, 30, 36]. In this paper, we focus on weakly-supervised semantic segmentation with image-level labels, the lowest annotation-cost alternative.




Figure 1. Our object response maps compared with baseline CAMs. The baseline CAMs only highlight the most discriminative regions. Our proposed technique leads to response maps that integrally cover larger object regions without re-training.

The typical way to learn from image-level labels usually involves progressive steps: 1) an initial response map (the class activation map (CAM) [53]) is obtained to roughly locate the objects; 2) pseudo labels are generated based on the initial response map with post-processing techniques, e.g., denseCRF [26], random walk [33] or an additional network [1, 2, 41]; 3) a semantic segmentation network is trained with the pseudo labels as supervision. The quality of the initial response map plays an important role for image-level WSSS, as good response maps can fill the inherent gap between image-level labels and pixel-wise labels. However, as the CAM highlights discriminative regions of each category, the partial activation leads to unsatisfactory semantic segmentation. To refine the pseudo labels from CAMs, most recent state-of-the-art methods require additional training steps. *i.e.* re-train the classification model to encourage the CAMs to cover more object areas [4, 4, 44, 37, 50, 42], or train additional networks [2, 1, 41] to guide the CAMs to generate pseudo labels.

We observe that the partial activation of the CAM is because only the discriminative region is usually needed for effective object classification. So the network is prone to focus on the discriminative areas. However, we argue that this does not indicate that the classifier learns nothing about other less-discriminative patterns. We experimentally validate that, assuming there is sufficient data for each category, the trained classifier can generate activation on most areas

of objects but unevenly distributed. We describe this partial activation issue as an “unequal distribution of the activation”, and find that conventional inference fails to leverage the full power of the baseline classifier. We demonstrate that a basic classification network, pre-trained on the target dataset without modification is sufficient to generate uniform activation and cover most object areas with an effective inference strategy.

We propose an inference-time image augmentation method to reveal hidden activation of objects and generate better object response maps for WSSS. To prevent the risk of diverging from the well-trained classifier, we do not re-train the baseline classifier. Instead, our method adopts only a novel inference mechanism to deal with the unequal distribution issue of the CAM. Specifically, we first introduce a “split & unite” based image augmentation method to encourage the network to pay attention to different parts of objects and generate equal activation on each part. To further push the activation to other less discriminative areas, we present a “hide and re-inference” method, which iteratively mines activated regions of objects and aligns them to an equally-distributed response map. Finally, we integrate these two modules into a simple framework that can be used in the inference stage of existing classifiers. As shown in Fig. 1, our inference-time method generates more uniform object response across the entire object. We conduct extensive experiments on the PASCAL VOC 2012 dataset [12], and both qualitative and quantitative results demonstrate the effectiveness of our approach.

In addition, we explore a new method to leverage saliency maps in WSSS. We argue that, since salient object detection models are always trained by the class-agnostic objects with center bias, directly using saliency as background cues to generate pseudo labels can deteriorate the quality of the segmentation pseudo labels. To address this, we propose *activation aware mask refinement* to further refine the semantic segmentation, which uses saliency maps as a subsidiary supervision during semantic segmentation training along with pseudo labels. We can refine the foreground object boundaries and meanwhile inhibiting suppression on the activated objects in the background.

Our main contributions are summarized as follows: 1) We identify a core issue causing the unequal distribution of CAMs and explore a new option to refine initial response maps during inference. 2) We propose a *Class-conditional Inference-time Module* to obtain better object activation without any network modification or re-training. 3) We propose the *Activation Aware Mask Refinement Loss*, a new approach to incorporate saliency information in WSSS that can refine object boundaries, but also prevents suppression of background objects due to the saliency centre-bias. 4) Our inference-time method can also be treated as an add-on solution to the existing image-level based methods.

## 2. Related Work

**Weakly Supervised Semantic Segmentation:** A large number of WSSS methods have been proposed to achieve a trade-off between labeling efficiency and model accuracy, where the “weak” annotations can be image-level labels [23, 44, 2, 16, 34, 42, 5, 50, 6, 37, 48, 18, 47, 13, 28, 46, 45, 29], scribbles [40, 31, 39], points [3], or bounding boxes [10, 34, 25, 30, 36, 38]. We mainly focus on image-level label based weakly supervised models.

As the start point, the quality of the initial CAM is important for the semantic segmentation network. Two different methods have been widely studied to obtain a better initial response map, including network refinement based models and the data augmentation and erasing based techniques.

**Network Refinement:** [44] adopts dilated convolution with various dilation rates to enlarge the receptive field and transfer discriminative information to non-discriminative object regions. [5] performs clustering on image features to generate pseudo sub-category labels within parent classes, which were then used to train the classification model and generate CAMs that covered larger regions of the object. [50] introduces discrepancy loss and intersection loss to first mine regions of different patterns, and then merge the common regions of different response maps.

**Data Augmentation and Erasing:** These methods augment or erase input images to force the network to generate larger response map object coverage. [43] erases highly activated areas from the image and then retrains the classification network to discover new object regions for classification. [35] divides the image into a grid of patches, then randomly hides patches to force the network to focus on other relevant object parts. [21] leverages two self-erasing strategies to encourage the network to use reliable object and background cues to prohibit attention from spreading to unexpected background regions. [4, 18] utilize mixup data augmentation to calibrate the uncertainty in prediction, they randomly sample an image pair to mix them together and feed into the network, which forced the model to pay attention to other image regions.

In general, all above methods require “re-training” the classification model to obtain a refined initial response map. We introduce a new method for initial response map acquisition without re-training. One recent method that does not re-train is [15], which directly generates multiple CAMs for each image from the classifier using different input scales, backbones and post-processing. The segmentation model’s robustness is then leveraged to learn from the noisy CAMs, using a learnable per-pixel weighted sum of multiple CAMs. In this paper, we explore a new alternative option to refine the initial response maps during the CAM inference stage of a single network without re-training.

**Saliency Assisted WSSS:** Saliency maps are often adopted in WSSS [23, 21, 13, 14, 44, 24, 27, 37, 41, 43] to serve




Figure 2. Visualization of activated regions of the baseline CAM. Activation is unequally distributed in the object, where the highly activated region (CAM) is the discriminative region. However, we observe that the less discriminative regions are still activated covering most object areas (second and third column)

as background cue to generate pseudo labels. Recently, [45] proposes a pseudo label generation module, which uses saliency maps as background cues and chooses a predefined threshold to retain the activated objects in background. [29] directly utilizes saliency maps as supervision during classifier training to constrain object activation. [46] proposes potential object mining and Non-Salient Region Masking to explore objects outside salient regions. However, no existing WSSS methods directly use saliency maps during semantic segmentation training.

### 3. Our Approach

We first analyse the baseline CAM method, then propose the *class conditional response map* and a *class conditional inference-time approach* based on it. Finally, we introduce our *Activation Aware Mask Refinement Loss*, a novel approach to leverage pre-trained saliency maps in WSSS.

#### 3.1. Class Conditional Response Map

For a class  $c$ , a CAM [53] is a feature map indicating the discriminative region of the image that influence the classifier to make the decision that this object belong to class  $c$ . Given  $f_k(x, y)$ , the activation of unit  $k$  in the last convolutional layer at location  $(x, y)$ , and  $w_c$ , the weights from  $f_k(x, y)$  via global average pooling, the CAM  $M$  is a map of activation at locations  $(x, y)$  as described in [53]:

$$M_c(x, y) = \sum_k w_k^c f_k(x, y). \quad (1)$$

Consider a network  $h$  with parameters  $\theta$ . For input image  $I$ , suppose the output for class  $c$ ,  $h_\theta(c|I) > \tau$ , where  $\tau$  is the threshold probability. Recent high-performance networks (e.g., [22, 20]) trained on large datasets (e.g., ImageNet) yield highly accurate classification for well-represented classes. Hence, given positive classification for  $c$ , we can treat  $M_c$  as an approximation to a class conditional response map. Hence, we define a *class conditional*

*response map* for image  $I$  using network parameters  $\theta$  as:

$$R_c((x, y)|I; \theta) = M_c((x, y)|c, I; \theta) \quad (2) \\ \approx M_c((x, y)|I; \theta),$$

where  $M_c((x, y)|c, I; \theta)$  is the CAM for  $I$  in which class  $c$  appears, and  $M_c((x, y)|I; \theta)$  is the CAM for image  $I$ , given that class  $c$  appears with high probability in the image.

#### Over-complete Activation

Consider Fig. 2, column two shows the baseline CAM activation, where some discriminative areas are well-activated, but other visible object regions have weak activation, (e.g., the heads of the sheep versus the bodies). The third column shows all areas with activation greater than zero, which are large and generally include most object areas. Quantitatively, we obtain all activated masks for the PASCAL VOC 2012 training set and get a recall of 84% compared with the semantic segmentation ground truth, i.e. the majority of object regions are activated by the baseline classifier, not just discriminative regions. We can see that the baseline classifier learns most object features with sufficient training data. However, the response is over-complete and uneven, so extracting segmentation pseudo labels is difficult.

#### Uneven Distribution

For discriminative training, the loss is indifferent to extensive activation across an object, requiring only a sufficient global average value via pooling. Existing image-level weakly supervised methods observe that the CAM only shows high activation on an object’s most discriminative regions [43, 35, 44, 21, 4, 18, 5, 42, 50], but they disregard the less discriminative object regions where the activation is suppressed. By definition, the class conditional response map has a global average per-pixel response greater than a threshold. Then, we can infer that a unit  $f_k(x, y)$  at location  $(x, y)$  with high activation in  $R_c((x, y)|I; \theta)$  has an appearance pattern within its receptive field that is strongly associated with the presence of  $c$ .

#### Suppression of Broader Object Activation

We argue that for a particular input image, the presence of a highly discriminative region may suppress other less discriminative regions in  $R_c$ . Network processing is well-understood, but let’s consider network mechanisms that feed into  $R_c$ . The main mechanism for suppression in earlier layers in modern networks (e.g., ResNet) is via batch normalization and negative weights (ReLU output is non-negative). Each  $f_k(x, y)$  from Eq. 1 projects back into the penultimate layer by standard convolution:

$$f_k(x, y) = \sigma(\sum_{k' \in \mathcal{N}(x, y)} [w_{k'}^c f_{k'}(x, y)]), \quad (3)$$

where we use  $\sigma$  to represent a combined ReLU activation with Batch Normalization, and  $k'$  indexes over the convolution input units to  $f_k(x, y)$ . That is, the class conditional response at location  $(x, y)$  is a non-linear function of

weighted input units  $f_{k'}(x, y)$  from the prior layer. Hence, by a combination of negative  $w_{k'}$  and  $\sigma$ , a strong activation from some  $f_{k'}$  can lead to suppression of the corresponding  $R_c((x, y)|I'; \theta)$  in the class conditional map. Note that  $f_{k'}(x, y)$  can also be negative by the cascade of Eq. 3.

For networks with deep residual structures it would be difficult to trace back through the cascading and residual activation to find the exact pixels that have led to this suppression of CAM locations that may otherwise have positive features. Instead, we propose a class conditional inference-time approach to solve this issue. We aim to generate a more uniform distribution of activation across the visible object.

### 3.2. Class Conditional Response Map Inference

We assume an initially well-trained classification network  $h_\theta$ , which we do not re-train. Instead we propose *class conditional inference*, whereby we augment  $I$  at inference time to remove regions that may suppress the response of other object parts to mine the class-conditional object activation. Concretely, we perform this by computing the class conditional response map  $R_c(x, y|I'; \theta)$ , where  $I'$  is an augmented image. By removing pixels in  $I$  associated with high activation, we aim to explore other regions of being suppressed. A high response on  $R_c(x, y|I'; \theta)$  for some location where the activation was not high for  $R_c(x, y|I; \theta)$  is likely to indicate an appearance pattern that is strongly associated with the object, but was suppressed by the regions that visible in  $I$ , but not  $I'$ . In this section we describe our implementations of this by class conditional Split & Unite Augmentation, and Iterative Inference.

#### Split & Unite: Image Augmentation

To address the unequal distribution of object activation, we propose a class conditional split & unite inference strategy to investigate different parts of the object, and align their activation. A naive approach would be to randomly divide the image by a grid and perform inference on each grid cell. But it is likely that some cells would not contain the object, and have a higher chance of false positive responses on small regions.

Instead, we first apply conventional inference with the baseline classifier to obtain  $R_c((x, y)|I; \theta)$ . As we assume a well-trained model  $\theta$ , the initial high CAM activation will generally fall on the discriminative region of the target object. Then we calculate the centre of mass of the original CAM to obtain a center point about which we split the image into four patches, as shown in Fig. 3. We find the centre of mass generally falls on the class, in which case some part of the object appears in each of the four sub-images.

With the classifier fixed, we then compute  $R_c((x, y)|I_i; \theta)$  where  $i$  indexes each split. Although the unequal distribution of the object activation in the baseline CAM leads to only highly discriminative regions being highlighted, by separate inference, object discrim-

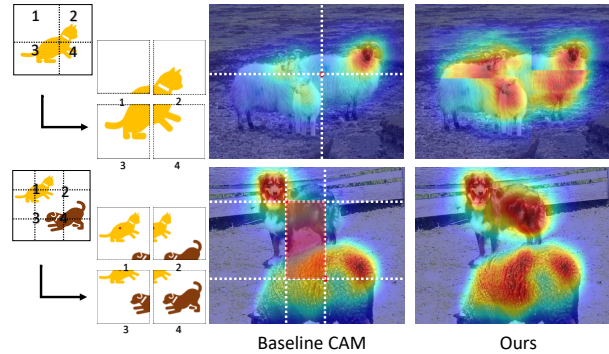


Figure 3. Our split & unite augmentation methods. During inference, we split the image into patches with each patch containing a part of object, and then we reprocess each patch to find its class conditional response map.

ination is computed individually in each patch with any suppressing elements in the other patches removed so we can locate more discriminative areas on different object parts. For example, see the first row of Fig. 3 showing the split for images with one object class. Without the highly activated sheep head in the patch #2, other parts of the sheep are highlighted in patch #1, #3 and #4. For images with two object classes, we refine our split strategy, as shown in second row of Fig. 3, we calculate centres of mass for both class’s CAM, then we use these two center points to obtain a rectangle inside the image (displayed as the red central area). We use the four corners of this rectangle as a split point to crop the original image into four patches. The central rectangle (red area) is retained in all four patches, as shown in the “Splits” column. Each split generally contains different parts of both object classes. Then we run inference on each patch in turn, and merge the four response maps, we take the max activation for the overlapping central rectangle area. For images with three or more classes, we split the images by the CAM mass center for each class separately.

Furthermore, our split & unite method could also be used as a common image augmentation method to assist re-training the classification model. Specifically, we feed each patch individually into the classification model to train the network with the same image-level ground truth as the original image. *i.e.* we use different object parts to train the classifier instead of the entire object, thus the classifier will naturally pay attention to more object parts during inference. We show further experimental results of in Sec. 4.3.

#### Refined Iterative Inference Module

In the Split & Unite module, the classifier still focuses on discriminative regions in each patch, so we extend our inference method by introducing the iterative erasing mechanism. Iterative training by erasing the high activation is an adopted technique in the WSSS. For example, [43] proposes an iterative training scheme, in each iteration the high acti-




Figure 4. We iteratively remove the highly activated regions from the image, and feed the new image back to the fixed classifier.

vation regions are erased then the image is fed back to train the classifier again. But it has the risk that all objects are erased while network is still updating with false positive, also the re-training time is greatly increased. In this section, we propose a refined inference-only iterative module to further improve the object activation maps. Our module is supported by our analysis in Sec. 3.1, which requires no re-training, avoid the risks of existing iterative erasing techniques and can achieve better results.

Refer to Fig. 4, we first feed the original image into the classifier and produce response map, revealing the highly discriminative areas. Then these highly discriminative areas are erased with mean colour value of the original image. The augmented image is then fed back to the classifier for next inference iteration. No training is performed as the pixels corresponding to the discriminative region are removed. With these object features absent, without suppression, weaker activation will be naturally driven to shift to high activation. We iterate this inference process and then add the newly generated activation map of every iteration together to obtain the final response map. As shown in Fig. 4, new object areas are activated progressively in each iteration without updating the network.

### Our Inference Module

Finally, refer to Fig.5, we integrate our split & unite and iterative inference together into a unified inference-only framework. We first split the image into 4 splits and feed them into the classifier in parallel to encourage more object activation on each split. Then we perform iterative inference on each of them respectively as shown in the green block. Finally, we combine each split together into our final object response map. The split & unite module and iterative inference module mutually benefit each other to balance activation across different object parts and densely cover larger object areas. Our class-conditional module can be seemed as an add-on module. It can be seamlessly utilized in the inference stage of any pre-trained classification networks. As the proposed inference module requires no re-training,

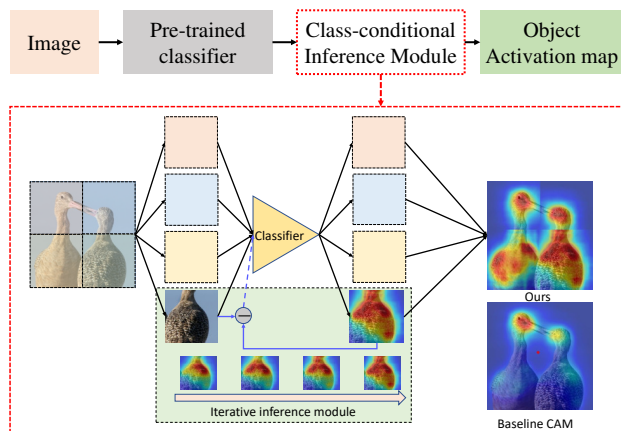


Figure 5. The framework of our proposed inference method. As shown, we split the image into 4 splits and do iterative inference (green block) on each split in parallel, then we combine the 4 splits together into our final object response map.

it saves all re-training time and is much easier to implement. In our experiments section, we compare with existing methods, and the results show that we achieve comparable performance with the state-of-the-art methods while most of them rely on re-training the classification model.

### 3.3. Activation Aware Mask Refinement Loss

Saliency information is used by many approaches in image-level WSSS to refine object boundaries and obtain a background mask [37, 24, 27, 16, 17, 21, 23, 44, 41]. Commonly, a pre-trained off-the-shelf salient object detection model is adopted to generate class-agnostic saliency masks on the segmentation dataset. Then the saliency masks are used during pseudo label generation, normally they are multiplied by a manually chosen parameter and concatenated onto the activation maps as background. High activation pixels are then used to obtain pixel-wise pseudo labels.

However, we observe that since saliency detection models are usually trained by class-agnostic objects with center bias, the saliency maps may falsely detect non-object salient areas in the foreground, and tend to ignore non-salient objects in the background (see Fig. 7). Thus, it may introduce errors into pseudo-labels and harm segmentation training. Some desired object regions are ignored and some non-object regions are falsely detected in the saliency maps. However, saliency maps are still required by WSSS to support finding accurate foreground object boundaries. To address this issue, we propose a new method to leverage saliency information in WSSS, including a new saliency loss. The use of CAMs to obtain pseudo-labels without saliency maps is unchanged. Then we use the pseudo labels together with the provided saliency maps to train our semantic segmentation model. As shown in Fig. 6, we keep




Figure 6. Top: the convention saliency guided WSSS training. Bottom: our activation-aware mask refinement training, where object activation and saliency information are visible to the network.

both labels visible to the network, so the saliency maps can refine foreground object boundaries, but we inhibit suppression of activated objects in the background. More formally, our segmentation loss is defined as:

$$L_{seg} = L_{seg} + \alpha(e^{\tau} - 1)L_{sal} \quad (4)$$

The first term  $L_{seg}$  is the cross entropy loss between segmentation prediction and our activation-generated pseudo labels. The second term  $L_{sal}$  denotes the binary cross entropy loss between the background channels of the predictions and the saliency maps. We incorporate a modulating factor  $\tau$  (called the *conflict temperature*) with tunable weight  $\alpha > 0$ . Intuitively, as two noisy signals (pseudo labels and saliency maps) in the supervision pair have conflict areas as shown in Fig. 7, they will compete with each other. We use  $\tau$  to control the competition between two supervision terms.  $\tau$  is defined as the mean intersection-over-union (MIoU) between our pseudo label background channel and the saliency maps. If  $\tau$  is low, it indicates the saliency map is not consistent with the pseudo background, and so the saliency maps will harm segmentation training, so the weight of the saliency loss  $L_{sal}$  is diminished. Contrarily, if MIoU is high, then the saliency loss is enhanced to refine the object boundaries. In summary, our activation-aware mask refinement loss proposes an adaptive mechanism to refine segmentation training. Saliency information is fully explored to refine the foreground object boundaries but not harm activated objects in the background.

## 4. Experiments

### 4.1. Implementation Details

In this section, we introduce implementation details of the proposed method and the following procedures to generate semantic segmentation results. We evaluate our approach on the PASCAL VOC 2012 dataset [12] with the background and 20 foreground object classes. The official dataset consists of 1446 training images, 1449 validation and 1456 test. We follow common practice, augmenting the



Figure 7. Conflicts between saliency maps and our activation-generated pseudo labels, where the ‘‘Pseudo label’’ is the activation-generated pseudo label. As shown in the first row, our activation correctly detects the TV monitor but the saliency map ignores it. In the second row, the saliency map falsely highlights the banner but ignores the bus.

Method	Training Set		Validation Set
	CAM	CAM + RW	CAM
Baseline [2]	48.0	58.1	46.8
Mixup-CAM [4]	50.1	61.9	x
SC-CAM [5]	50.9	63.4	49.6
Ours	52.2	64.2	50.7

Table 1. Performance comparison in mIoU(%) of the initial response maps on the PASCAL VOC training and validation set. Mixup-CAM [4] and SC-CAM [5] are two state-of-the-art methods, and x means the paper does not report this number.

Baseline	split & unite	Iterative Inference	mIoU
✓			48.0
✓	✓		49.9
✓		✓	50.2
✓	✓	✓	52.2

Table 2. The ablation study for each part of our method to validate effectiveness of the proposed strategies.

training set by adding images from the SBD dataset [19], to form a total of 10582 images.

#### Object Response Map Generation

In our pipeline, we use the weights of the baseline classification model provided by [2], pre-trained on ImageNet [11], and fine-tuned on PASCAL VOC 2012, without any re-training. Similar to [2] and others, our baseline classifier uses the ResNet-38 backbone with `output_stride = 8`, global average pooling, followed by a fully connected layer. In the iterative inference module, we empirically choose 0.7 as threshold for high activation, and remove high activation regions in each inference iteration. We stop iteration when the new activation is smaller than 1% of the image.

#### Semantic Segmentation Generation

After obtaining response maps using our method, following recent work [42, 5, 4, 37, 6], we adopt the AffinityNet random walk [2] to refine response maps into pixel-wise semantic segmentation pseudo labels. Also, we apply fully connected random fields [26] to refine the pseudo-label object boundaries. Finally, we use the generated pseudo labels and saliency maps as supervision to train the popular Deeplab semantic segmentation framework with ASPP [8] using the ResNet-101 backbone network. We use saliency




Figure 8. Sample results of initial response maps and ours. Baseline CAMs tend to only highlight discriminative areas. Our approach helps balance object activation across different object parts and densely cover larger object areas. For images with multiple classes, we can get better activation for all classes, and we combine them in the activation for convenience of reading.

Class number	1 class	2 classes	3+ classes	Total
Baseline	46.6	50.9	41.4	48.0
SC-CAM[5]	49.5	53.5	43.5	50.9
Ours	49.8	55.3	44.7	52.2

Table 3. Model performance with respect to number of categories.

maps from [24],  $\alpha$  in Eq. 4 is empirically set to 0.08.

## 4.2. Improvement on Initial Response Map

In Table 1, we show initial response map (CAM) performance, using best mean IoU, *i.e.* the best match between the response map and segmentation ground truth under all different background thresholds. We also report pseudo-label results after applying the random walk refinement (CAM + RW). Note that our results are obtained with the baseline classifier without any re-training, all improvements come from our proposed *class conditional inference*. As shown in Table 1, our initial response maps are significantly improved over baseline [2] on both training and validation sets. We also compare response maps generated by recent state-of-the-art methods [4, 5], and observe a clear margin.

Our better initial response maps lead to better performance of downstream tasks: generating pixel-wise pseudo-labels and final segmentation results. After refining the response map into pseudo labels that are used to train a semantic segmentation model, we also achieve significant improvement over the baseline and outperform competing methods, as shown in “CAM+RW” column in Table 1. This validates that, by direct inference without fine-tuning, our method can substantially improve object activation and generate better response maps than competing methods with re-trained classification models. In Fig. 8, we show qualitative examples compared with baseline CAMs, showing that ours can activate substantially more object parts and uniformly cover larger regions of the objects.

In Table 2, we show an ablation study, how each of our modules improves the initial response maps. The improvement mainly stems from more dense coverage of objects, we test each module independently. Our split & unite augmentation improves the baseline CAM by 1.9%. The iter-

	Training Set	Validation Set
Baseline	48.0	46.8
Ours (Direct inference)	52.2	50.7
Ours + split&unite re-training	52.7	51.8
Ours + SC-CAM[5]	53.3	51.5

Table 4. The effectiveness of our model as an “add-on” to re-trained classifier.

ative inference module has an improvement of 2.2% compared with baseline. It validates that both our modules provide manifest improvement over the baseline CAM. Integrating them together, we achieve a significant improvement over the baseline by 4.2% on our initial response map.

Finally, we report performance improvements for different numbers of classes appearing in an image in Table 3. As shown, we achieve consistent improvements over competing methods on images with all class numbers, demonstrating the effectiveness and generality of our method.

## 4.3. Ablation: Integrating with Re-training

Our method can be directly used as an add-on inference module with any existing re-trained classifier to obtain a better initial response map. First, as discussed in Sec. 3.2, our split & unite augmentation can be used in the training stage as a data augmentation method. We follow the methods described in Sec. 3.2 to augment the training set to feed different parts of the object into the classification model to train the network, so the network will update its weights to pay attention to more object parts. Refer to Table 4: fine-tuning the classifier with our split & unite augmentation further improves the initial response map performance.

In addition, we perform our inference-time data augmentation on the re-trained CAM of (SC-CAM [5]) to refine their produced CAM. SC-CAM [5] introduced a sub-category clustering method to force the network to activate in more categories, leading to enlarged activation regions. We then perform our inference-time approach on their re-trained classifier. As shown in Table 4, we (Ours + SC-CAM) improve their performance significantly by 2.4%. This validates that our inference method can be used as an add-on solution to integrate with existing re-trained classifiers to further refine the object activation maps. .




Figure 9. Qualitative results of our semantic segmentation results.

#### 4.4. Semantic Segmentation Performance

In Table 1, we show the refined pseudo labels generated by our initial response maps. We then use the pseudo labels and our *activation aware mask refinement loss* to train a segmentation network on PASCAL VOC dataset. As per common practice, we use the densely connected CRF [26] to refine the semantic segmentation predictions as post processing. We show final predictions in Fig. 9, clearly showing the effectiveness of our approach. In Table 5, we compare our method with recent work. We report performance on both the validation and test set of the PASCAL VOC 2012 dataset [12]. As shown, our method outperforms all others on the test set, and is comparable with state-of-the-art on the validation set. Note that most other methods require additional re-training steps to obtain better object activation, our method is easier and faster to implement. Further, we remove the saliency guidance in Fig. 6 from our framework, leading to the cheapest model (no re-training, no saliency), and achieve mIoU(%) on PASCAL VOC validation set and testing set as 66.2 and 66.3 respectively, which is comparable with re-training based model RRM [48] and slightly worse than AdvCAM [28].

#### 4.5. Discussion: Computation-time

We argue that current WSSS multi-training schemes are complicated and inelegant, so we propose to improve WSSS performance without further increasing computation time. First, our saliency loss is a subsidiary supervision calculating the binary cross entropy loss during segmentation training, the increased time complexity is negligible, and most SOTA approaches already incorporate saliency to generate pseudo labels. Second, although re-training time is saved, our *class conditional inference module* requires once-off additional inference-time computation to generate pseudo-labels, we give a quantitative analysis here. To obtain activation maps on the PASCAL VOC train set (1464 images) our method takes 40 minutes compared to 13 (factor of 3, e.g., an extra 3.5 hours for the 10582 image Augmented dataset). On the other hand, re-training the baseline classifier as performed by most others requires 6 hours on the same GPU settings (note that competing methods have ad-

Method	Re-training	Extra data	Val	Test
MDC[44] VGG16	✓	S	60.4	60.8
DSRG[23] ResNet101	✓	S	61.4	63.2
Affinity[2] ResNet38			61.7	63.7
SeeNet [21] ResNet101	✓	S	63.1	62.8
IRNet [1] ResNet50	✓		63.5	64.8
BDSSW [17] ResNet101		S	63.6	64.5
CIAN [16] ResNet101	✓	S	64.1	64.7
SEAM[42] ResNet38	✓		64.5	65.7
FickleNet[27] ResNet101	✓	S	64.9	65.3
OAA [24] ResNet101	✓	S	65.2	66.4
Mixup-CAM [4] ResNet101	✓		65.6	x
BES [6] ResNet101	✓		65.7	66.6
SC-CAM [5] ResNet101	✓		66.1	65.9
CONTA[49] ResNet101	✓		66.1	66.7
MCS [37] ResNet101	✓	S	66.2	66.9
RRM[48](two step) ResNet101	✓		66.3	66.5
EME[15] ResNet101		S	67.2	66.7
ICD[13] ResNet101	✓	S	67.8	68.0
AdvCAM [28] ResNet101	✓		68.1	68.0
NSRM[46] ResNet101	✓	S	68.3	68.5
EDAM[45] ResNet101	✓	S	70.9	70.6
EPS[29] ResNet101	✓	S	70.9	70.8
Ours ResNet101		S	70.8	71.8

Table 5. Semantic segmentation performance comparison on the PASCAL VOC 2012 val and test sets. We report the methods that re-train the classification model for better response maps. Also, the methods that utilize extra saliency masks to generate pseudo labels are marked with “S”.

ditional augmented images, and/or network enhancements meaning their re-training takes at least this long). Thus, our inference-time method still greatly saves overall time to obtain high-quality activation maps. As future work, we will refine our method to further reduce the inference time.

## 5. Conclusion

In this paper, we propose a novel inference method that helps generate better object response maps without re-training the baseline classifier, and a new method to utilize saliency information in WSSS. Specifically, we propose two inference-time modules to generate dense object response maps. Firstly, we develop an augmentation method and let the classifier inference on different image parts individually so as to shift the activation to more object areas. Secondly, we propose an iterative inference that encourages the classifier to progressively mine more object parts by hiding high activation areas during inference. Whereas most current state-of-the-art methods require multiple training steps, our method directly generates response maps using the baseline classifier. We show that our algorithm produces a better initial response map with less computation. In addition, our activation aware mask refinement loss provides a new way to incorporate saliency information in WSSS which further improves final semantic segmentation performance.

**Acknowledgments.** This research was in-part supported by the ANU-Optus Bushfire Research Center of Excellence.



## References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2209–2218, 2019.
- [2] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4981–4990, 2018.
- [3] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *Eur. Conf. Comput. Vis.*, pages 549–565, 2016.
- [4] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Mixupcam: Weakly-supervised semantic segmentation via uncertainty regularization. *arXiv preprint arXiv:2008.01201*, 2020.
- [5] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8991–9000, 2020.
- [6] Liyi Chen, Weiwei Wu, Chenchen Fu, Xiao Han, and Yuntao Zhang. Weakly supervised semantic segmentation with boundary exploration. In *Eur. Conf. Comput. Vis.*, pages 347–362. Springer, 2020.
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *ArXiv e-prints*, 2014.
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2017.
- [9] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *ArXiv e-prints*, 2017.
- [10] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Int. Conf. Comput. Vis.*, pages 1635–1643, 2015.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255, 2009.
- [12] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010.
- [13] Junsong Fan, Zhaoxiang Zhang, Chunfeng Song, and Tieniu Tan. Learning integral objects with intra-class discriminator for weakly-supervised semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4283–4292, 2020.
- [14] Junsong Fan, Zhaoxiang Zhang, and Tieniu Tan. Cian: Cross-image affinity net for weakly supervised semantic segmentation. *ArXiv e-prints*, 2018.
- [15] Junsong Fan, Zhaoxiang Zhang, and Tieniu Tan. Employing multi-estimations for weakly-supervised semantic segmentation. In *Eur. Conf. Comput. Vis.*, pages 332–348, 2020.
- [16] Junsong Fan, Zhaoxiang Zhang, Tieniu Tan, Chunfeng Song, and Jun Xiao. Cian: Cross-image affinity net for weakly supervised semantic segmentation. In *AAAI Conf. Art. Intell.*, 2020.
- [17] Ruochen Fan, Qibin Hou, Ming-Ming Cheng, Gang Yu, Ralph R Martin, and Shi-Min Hu. Associating inter-image salient instances for weakly supervised semantic segmentation. In *Eur. Conf. Comput. Vis.*, pages 367–383, 2018.
- [18] Hongyu Guo, Yongyi Mao, and Richong Zhang. Mixup as locally linear out-of-manifold regularization. In *AAAI Conf. Art. Intell.*, volume 33, pages 3714–3722, 2019.
- [19] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *Int. Conf. Comput. Vis.*, pages 991–998, 2011.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016.
- [21] Qibin Hou, PengTao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. In *Adv. Neural Inform. Process. Syst.*, pages 549–559, 2018.
- [22] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4700–4708, 2017.
- [23] Zilong Huang, Xinggong Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. Weakly-supervised semantic segmentation network with deep seeded region growing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7014–7023, 2018.
- [24] Peng-Tao Jiang, Qibin Hou, Yang Cao, Ming-Ming Cheng, Yunchao Wei, and Hong-Kai Xiong. Integral object mining via online attention accumulation. In *Int. Conf. Comput. Vis.*, pages 2070–2079, 2019.
- [25] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 876–885, 2017.
- [26] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Adv. Neural Inform. Process. Syst.*, pages 109–117, 2011.
- [27] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5267–5276, 2019.
- [28] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4071–4080, 2021.
- [29] Seungho Lee, Minhyun Lee, Jongwuk Lee, and Hyunjung Shim. Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In

- IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5495–5505, 2021.
- [30] Qizhu Li, Anurag Arnab, and Philip HS Torr. Weakly-and semi-supervised panoptic segmentation. In *Eur. Conf. Comput. Vis.*, pages 102–118, 2018.
- [31] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3159–3167, 2016.
- [32] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3431–3440, 2015.
- [33] László Lovász et al. Random walks on graphs: A survey. *Combinatorics, Paul erdos is eighty*, 2(1):1–46, 1993.
- [34] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Int. Conf. Comput. Vis.*, pages 1742–1750, 2015.
- [35] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *Int. Conf. Comput. Vis.*, pages 3544–3553. IEEE, 2017.
- [36] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3136–3145, 2019.
- [37] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. *arXiv preprint arXiv:2007.01947*, 2020.
- [38] Weixuan Sun, Jing Zhang, and Nick Barnes. 3d guided weakly supervised semantic segmentation. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [39] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised cnn segmentation. In *Eur. Conf. Comput. Vis.*, pages 507–522, 2018.
- [40] Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7158–7166, 2017.
- [41] Xiang Wang, Shaodi You, Xi Li, and Huimin Ma. Weakly-supervised semantic segmentation by iteratively mining common object features. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1354–1362, 2018.
- [42] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020.
- [43] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1568–1576, 2017.
- [44] Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7268–7277, 2018.
- [45] Tong Wu, Junshi Huang, Guangyu Gao, Xiaoming Wei, Xiaolin Wei, Xuan Luo, and Chi Harold Liu. Embedded discriminative attention mechanism for weakly supervised semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16765–16774, 2021.
- [46] Yazhou Yao, Tao Chen, Guo-Sen Xie, Chuanyi Zhang, Fumin Shen, Qi Wu, Zhenmin Tang, and Jian Zhang. Non-salient region object mining for weakly supervised semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2623–2632, 2021.
- [47] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Int. Conf. Comput. Vis.*, pages 6023–6032, 2019.
- [48] Bingfeng Zhang, Jimin Xiao, Yunchao Wei, Mingjie Sun, and Kaizhu Huang. Reliability does matter: An end-to-end weakly supervised semantic segmentation approach. In *AAAI Conf. Art. Intell.*, pages 12765–12772, 2020.
- [49] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. *arXiv preprint arXiv:2009.12547*, 2020.
- [50] 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 *Eur. Conf. Comput. Vis.*, 2020.
- [51] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2881–2890, 2017.
- [52] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Pscanet: Point-wise spatial attention network for scene parsing. In *Eur. Conf. Comput. Vis.*, pages 267–283, 2018.
- [53] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2921–2929, 2016.