

# Normalization Matters in Weakly Supervised Object Localization

Jeesoo Kim<sup>1</sup> Junsuk Choe<sup>3\*</sup> Sangdoon Yun<sup>2</sup> Nojun Kwak<sup>1</sup>

<sup>1</sup> Department of Intelligence and Information, Seoul National University <sup>2</sup> NAVER AI Lab

<sup>3</sup> Department of Computer Science and Engineering, Sogang University

## Abstract

*Weakly-supervised object localization (WSOL) enables finding an object using a dataset without any localization information. By simply training a classification model using only image-level annotations, the feature map of the model can be utilized as a score map for localization. In spite of many WSOL methods proposing novel strategies, there has not been any de facto standard about how to normalize the class activation map (CAM). Consequently, many WSOL methods have failed to fully exploit their own capacity because of the misuse of a normalization method. In this paper, we review many existing normalization methods and point out that they should be used according to the property of the given dataset. Additionally, we propose a new normalization method which substantially enhances the performance of any CAM-based WSOL methods. Using the proposed normalization method, we provide a comprehensive evaluation over three datasets (CUB, ImageNet and Open-Images) on three different architectures and observe significant performance gains over the conventional min-max normalization method in all the evaluated cases (See Fig. 1).*

## 1. Introduction

Given nothing but the class information of an object, *weakly-supervised object localization (WSOL)* allows a convolutional neural network (CNN) to localize the object in a scene. Although many fully-supervised object detectors guarantee considerable performance in locating objects in an image, localization techniques in the absence of bounding box annotations are still in need.

WSOL using neural networks has been initially introduced by the class activation map (CAM) [30] approach. Training a convolutional neural network (CNN) model with a classification problem enables the model to generate an activation map from the last layer of it. After that, simply cutting out the activation map with a proper threshold en-

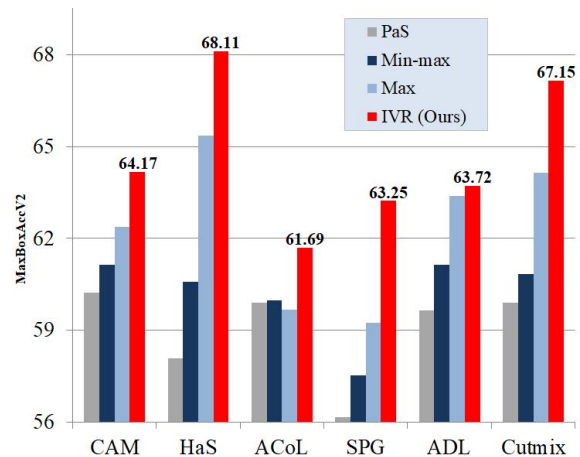


Figure 1: Comparison of several WSOL methods with different kinds of normalization methods for a class activation map. The accuracy has been evaluated under the evaluation metric suggested in [6] with CUB-200-2011 dataset. All scores in this figure are the average scores of ResNet50, VGG16, and InceptionV3. In all WSOL methods, the performance using our normalization method, IVR, is the best.

ables the localization of an object. In spite of the plausible axiom that features contributing better to a specific class are likely to represent the location of the object, the problem still exists that the discriminative parts of an object hoax the activation of the model to make an inaccurate localization of the target object. Many methods have been proposed to overcome this problem [21, 27, 28, 7, 25] and persuasive evidences of performance improvement have been demonstrated qualitatively and quantitatively.

However, prior works have been evaluated under different conditions and their hyperparameters have been chosen empirically. Usually, the feature from the last layer of the model is post-processed to be used as a class activation map. The normalization scheme used by every method differs each other and makes the comparison unfair. The work of [6] has proposed a new evaluation protocol and offered

\*Work done as a research scientist at NAVER AI Lab.

a thorough comparison of six previous WSOL methods. In their experiments, the min-max normalization has been applied to every method for the fair comparison and the best hyperparameters have been found using a random search [2]. According to it, all methods have shown almost no improvement compared to the original CAM [30].

In the work of [1], the authors suggest a thresholding strategy that excludes exceptionally high activation values in each image. In the sense that the valid value range in the activation map changes, this can be compared with other normalization methods. With this new method, some methods such as CAM [30] and HaS [21] have shown performance improvements in several datasets.

In this paper, we investigate the problem that can occur when using the min-max normalization. Even though the min-max normalization is the most popular scheme in recent WSOL works, we verify that the min-max scheme can deteriorate the performance of most WSOL methods. We revisit max normalization which has been prevalent for a long time and point out that it can resolve the problem above. Also, although other methods including *percentile as a standard for thresholding* (PaS) [1] have shown better results than the min-max normalization, they still suffer from problems which will be described in this paper. To resolve these problems, we propose a new normalization method *inferior value removal* (IVR). Through extensive experimental results, IVR has been shown to improve the localization performances of almost all WSOL methods. Enabling better exploitation of each WSOL method, a comprehensive re-evaluation and ordering of several WSOL methods have been conducted. The contribution of this paper is as follows:

- We provide a thorough investigation about problems of commonly used normalization methods in WSOL. The problem originating from using min-max normalization is explained qualitatively and quantitatively.
- We propose a new normalization method which can better exploit the performance of many WSOL methods. It can be used in any kind of WSOL methods which use a class activation map.
- A comprehensive evaluation with various kinds of normalization methods in three different datasets and three different architectures has been made. We provide a renewed benchmark of six WSOL methods.

## 2. Related works

Locating objects in images is one of the most important and frequently studied tasks in the field of computer vision. Depending on the type of supervision given to the model, semantic segmentation [31, 29, 5, 11, 8] and object detection [15, 13, 14, 12, 3, 26, 23] can be powerful localization methods.

There has been a constant demand on a localization method relying only on the image-level annotation. Class activation mapping (CAM) [30] is the first approach using the activated convolutional feature as a score map to locate an object in an image, and it has provided a great milestone for other following research. Hide-and-see (HaS) [21] makes a grid in an image and randomly erases multiple patches. The model struggles to make a correct decision with the corrupted image and this induce the feature of the model to be activated in the location of the target object. Adversarial complementary learning (ACoL) [27] uses two branches which adversarially get rid of the highlighted activated region from each other. This approach is different from HaS in that the feature is erased instead of the image itself. Self-produced guidance (SPG) [28] generates activation masks in each layer and uses them as a pseudo supervision for the preceding layer. Attention-based dropout layer (ADL) [7] adds attention modules to the model and adversarially drops highly activated regions. CutMix [25] has originally been designed to enhance the robustness of any CNN model. Patches from training images are cut and pasted to one another and this helps the model to capture less discriminative parts of the target object. Aside from CAM-based WSOL methods mentioned above, various gradient-based WSOL methods [18, 4, 19, 17] have been proposed as well. However, they are often heavy to be used practically and we do not consider these methods in this paper.

Meanwhile, a comprehensive and fair evaluation of CAM-based methods mentioned above has been made in [6]. More reasonable evaluation metrics have been proposed and all the considered WSOL methods have been re-evaluated with a thorough hyperparameter search and a unified min-max normalization method. As a result, it has been claimed that all WSOL methods after the emergence of CAM have turned out to be not significantly different in performance from CAM. To make a better use of the class activation map, a percentile based threshold modification method has been proposed [1]. This method assumes that large values must be treated as outliers and excluded from the class activation map. With the help of other techniques, several existing methods have succeeded to improve the performance. However, empirical choice of threshold violates the evaluation protocol in [6] that every WSOL method has its own optimal threshold. In this paper, we review the problems which most CAM-based WSOL methods experience and reset the benchmark for the future research.

## 3. Preliminary

WSOL is a task of finding an object belonging to a specific class in an image. While training a CNN with a classification problem, intermediate layers inside the model generate activation patterns to provide a correct output. Af-

ter a given image  $X \in \mathbb{R}^{W \times H \times 3}$  is processed by a CNN  $f$ , its generated feature map is usually average-pooled and handed to a classifier. A weight vector  $\mathbf{w}_i^c$  from the classifier is multiplied to the average-pooled feature and then the logit of the class  $c$  becomes available. On the other hand, by multiplying back this weight vector  $\mathbf{w}_i^c$  to the activated feature map  $f(X) \in \mathbb{R}^{W' \times H' \times K}$ , we can infer the location of an object belonging to class  $c$  within an image. After the multiplication, averaging along the channel dimension produces  $\mathbf{F} \in \mathbb{R}^{W' \times H'}$ , which we call as *class activation map* in this paper:

$$\mathbf{F}_c = \frac{1}{K} \sum_{i=1}^K \mathbf{w}_i^c \cdot f_i(X) \quad (1)$$

$$\mathbf{F}'_c = H(\mathbf{F}_c). \quad (2)$$

Here,  $K$  denote the number of channels. From now on, we omit the subscript  $c$  denoting the class index for brevity.

In all WSOL methods, values in  $\mathbf{F}$  typically lie in a certain range whose minimum value is a negative value in most cases and maximum value may be either a very small or large positive value. To apply a consistent level of threshold to all images,  $\mathbf{F}$  is normalized by a normalization function  $H$  into the range of  $[0, 1]$  producing a score map  $\mathbf{F}'$ . Not many of researchers have looked upon the importance of feature normalization. We discuss the effect of many different normalization functions  $H$  in this paper and suggest strategies for a better evaluation of WSOL.

## 4. Normalizing the class activation map

The context of images in a dataset or the portion of an object inside the image can induce a substantially different range of values in the score map  $\mathbf{F}'$ . Many normalization methods have been suggested but how the normalization can affect the performance has never been dealt with. Here in this section, we explain all the existing normalization methods up to our knowledge including the one we propose.

### 4.1. Min-max normalization

Min-max normalization is the most frequently used normalization method for WSOL which has been used in ACoL [27], SPG [28], ADL [7] and CutMix [25]. After the minimum value is subtracted from the class activation map, it is divided by its maximum value. This can be expressed as follows:

$$\mathbf{F}' = \frac{\mathbf{F} - \min(\mathbf{F})}{\max(\mathbf{F}) - \min(\mathbf{F})}. \quad (3)$$

By aligning the minimum value to zero and the maximum to one, all values in the score maps fall within the score range. This is described visually in (a) of Fig. 2. To

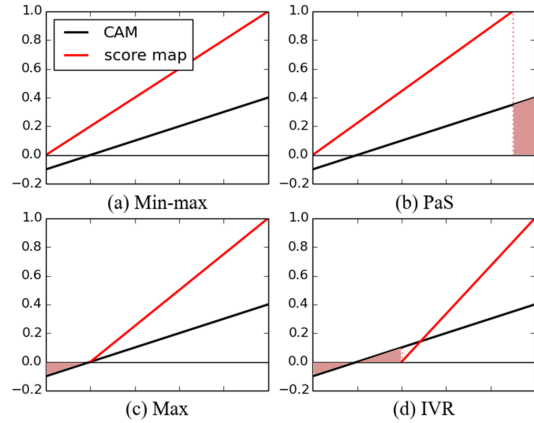


Figure 2: Illustration of how each normalization method aligns values in the class activation map. In this case, all methods share the same class activation map  $\mathbf{F}$  whose range is assumed to be  $[-0.1, 0.4]$ . Different methods map  $\mathbf{F}$  into different score maps  $\mathbf{F}'$  all ranging from 0 to 1. Consequently, we can assume that the region colored in red is excluded from the final score map.

provide a fair comparison under a consistent normalization method, the work of [6] suggests a well designed evaluation metric and compares every methods under min-max normalization. The conclusion of [6] is that the WSOL methods that have emerged after CAM actually do not show a significant performance enhancement. However, we observed that using min-max normalization experiences a problem that has not been considered in the work of [6]. Outlier values which are too large or small may distort other values during normalization. Within a dataset, class activation maps of some images (but not rare) often include exceptionally smaller minimum values and this often significantly drops the localization performance on those images. For most images which have an ordinary level of minimum value in the class activation map  $\mathbf{F}$ , this minimum value does not affect the overall score map  $\mathbf{F}'$ . However, when it comes to images that have an exceptionally smaller minimum value, it raises the activation level of the whole image and almost the entire image region is localized as the object in question (See Fig. 3, especially, CAM, HaS and CutMix).

### 4.2. Max normalization

Max normalization has been first used in CAM [30] and then in HAS [21]. Unlike min-max normalization, max normalization divides feature values by their maximum value and constrains the class activation map to be lower than one. This is described visually in (c) of Fig. 2. This can be ex-




Figure 3: Visualization of score maps and activated areas by the optimal thresholds in CAM, HaS, ADL, and CutMix on CUB dataset. Red, green and blue area are regions extracted by the optimal thresholds of IoU 0.3, 0.5 and 0.7.

pressed as follows:

$$\mathbf{F}' = \frac{\mathbf{F}}{\max(\mathbf{F})} \quad (4)$$

Negative values may exist after the normalization. Allowing values below zero, this method linearly transforms  $\mathbf{F}$  into  $\mathbf{F}'$ . As typically the optimal activating threshold  $\tau^*$  of a WSOL method is expected to be positive, max normalization can be considered simply as a method that ignores negative values. However, the value range of all images differ considerably and there is possibility that negative values may contribute to the quality of score maps.

### 4.3. Percentile as a Standard for Thresholding (PaS)

Percentile as a Standard for Thresholding (PaS) has been proposed in [1]. The authors say that exceptionally large values in the activation map make other relatively small values to be ignored. These large values must be excluded from the score map by using a percentile instead of their maximum value. This is described visually in (b) of Fig. 2. According to the code<sup>1</sup> available by the author, PaS uses min-max normalization except the maximum value is substituted with 90 percentile of all score map values. Therefore, it can be seen as a variety of the min-max normalization. This can

<sup>1</sup><https://github.com/won-bae/rethinkingCAM>

be expressed as follows:

$$\mathbf{F}' = \frac{\mathbf{F} - \min(\mathbf{F})}{\text{Pct}_p(\mathbf{F} - \min(\mathbf{F}))}. \quad (5)$$

$\text{Pct}_p(\cdot)$  denotes a function which finds  $p$  percentile from the given values. In the original paper, the optimal threshold  $\tau^*$  is fixed to a specific value for every dataset and  $p$  is set to be 90 heuristically. However, the evaluation metric suggested in the work of [6] validates all threshold values and picks the best threshold since WSOL methods are quite vulnerable to the un-optimized threshold. We re-evaluate PaS based on this metric and report the performance of this method. Also, we further apply PaS in all other WSOL methods and datasets mentioned in other papers [1, 6].

### 4.4. Inferior Value Removal (IVR)

In this paper, we propose a new normalization method, Inferior Value Removal (IVR), with a thorough investigation. The basis of IVR is that extremely small values in the class activation map unnecessarily raise the overall values in the score map and this disturbs the consistent thresholding among images. Unlike PaS, IVR is a normalization methods which restricts the minimum value of the activation map. Before dividing all values in  $F$  by their maximum value, a percentile value from its minimum value is subtracted. This is described visually in Fig. 2 (d). In other words, IVR uses

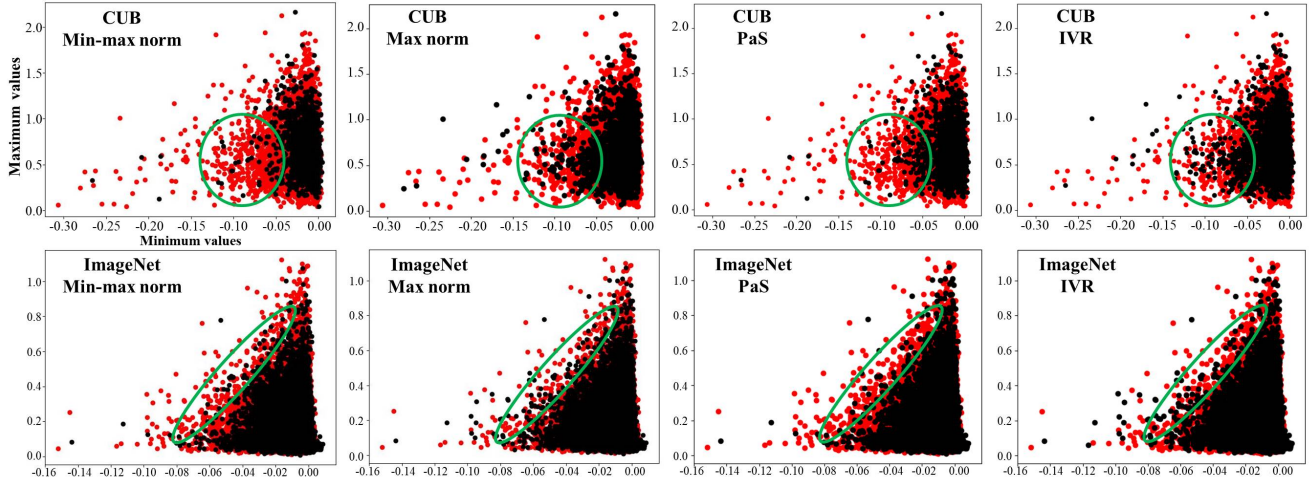


Figure 4: Distribution of minimum and maximum values from class activation map  $F$  of all test images. Each point represents a single image whose values from horizontal and vertical axes are minimum value and maximum value respectively. Black dots correspond to correctly localized images while red dots are not. Note that IVR shows overall higher density of positive samples (black dots) than other methods as highlighted with green circles.

max normalization except the minimum value is substituted with a certain percentile value of all activation map values. This can be expressed as follows:

$$\mathbf{F}' = \frac{\mathbf{F} - \text{Pct}_p(\mathbf{F})}{\max(\mathbf{F} - \text{Pct}_p(\mathbf{F}))}. \quad (6)$$

Values up to  $p$ -percentile are excluded from the activation map and the remaining values are re-arranged by their maximum value. The percentile parameter  $p$  is decided from the validation set and the performance tendency has been verified to be the same with the test set. By using IVR, we can use the original activation map which is relatively unaffected from the value shift by its minimum value while the range is calibrated according to the given values.

## 5. Experiments

### 5.1. Settings

**Dataset:** For the evaluation of WSOL, Caltech-UCSD Birds-200-2011 (CUB) [24], ImageNet [16] and OpenImages [10] are most frequently used. In the work of [6], both datasets are separated into three different divisions: `train-weaksup`, `train-fullsup` and `test`. We follow the details of all datasets described in [6]. In all three datasets, we select the best percentile for IVR from the `train-fullsup` set both in CUB and ImageNet.

**Evaluation metric:** In this paper, we use `MaxBoxAccV2` and `PxAP` as the evaluation metric as suggested in [6]. For CUB and ImageNet, `MaxBoxAccV2`( $\delta$ ) measures the performance in multiple intersection over union (IoU)

$\delta \in \{0.3, 0.5, 0.7\}$  to address the trade-off between precision and recall. `PxAP` measures the pixel-wise average precision when masks are available for evaluation. Unlike preceding metrics, `MaxBoxAccV2`( $\delta$ ) and `PxAP` choose the best operating threshold and report its score. In this regard, we can make sure that any performance improvement made in this paper is not dependent on the choice of the threshold. The evaluation has been conducted over VGG-GAP [20, 30], InceptionV3 [22], and ResNet50 [9].

### 5.2. Analysis on the class activation map

**Qualitative analysis:** Fig. 3 shows the class activation map and the region with values above the operating threshold  $\tau_{30}^*$ ,  $\tau_{50}^*$  and  $\tau_{70}^*$  (red, green and blue respectively). Note that  $\tau_{70}^*$  is usually smaller than  $\tau_{30}^*$  and  $\tau_{50}^*$ , capturing wider area in the images. In all methods, min-max normalization and PaS suffer from mishandled class activation maps. A region with exceptionally small values in the score map, which looks like a hole in black, raises the overall scores of all other pixels and makes the background more brighter than that of max normalization and IVR. Especially for PaS, values in the background is even more higher than in min-max normalization. Quite a lot of images possess sinkhole values in the class activation map and the optimal threshold becomes too low for them. Consequently, the bounding box simply holds the whole image leading to a performance collapse. Max normalization and Our IVR show localize the object tightly in all three IoU thresholds by leaving the background behind.

Even though the activated area in the class activation

Table 1: Evaluating WSOL using MaxBoxAccV2 with different normalization methods.

Method	Norm	ImageNet (MaxBoxAccV2)				CUB (MaxBoxAccV2)				OpenImages (PxAP)			
		VGG	Inception	ResNet	Mean	VGG	Inception	ResNet	Mean	VGG	Inception	ResNet	Mean
CAM	Minmax	60.02	63.40	63.65	62.36	63.71	56.68	62.98	61.13	58.30	63.23	58.49	60.01
	Max	60.85	64.77	64.88	63.50	64.54	59.10	63.47	62.37	59.65	64.97	59.49	<b>61.37</b>
	PaS	61.77	64.20	64.72	63.57	63.16	55.07	62.67	60.23	55.94	60.33	55.37	57.21
	IVR	61.47	65.49	65.57	<b>64.18</b>	65.27	60.76	66.83	<b>64.29</b>	59.25	63.66	58.97	60.62
HaS	Minmax	60.59	63.72	63.40	62.57	63.71	53.38	64.63	60.58	58.14	58.11	55.93	57.39
	Max	61.16	65.07	64.64	63.62	69.83	58.46	67.84	65.37	59.21	62.06	56.28	<b>59.18</b>
	PaS	62.05	64.53	64.57	63.72	61.16	51.27	61.85	58.09	55.97	57.79	52.85	55.54
	IVR	61.62	65.71	64.90	<b>64.08</b>	71.77	60.56	71.23	<b>67.85</b>	58.78	61.31	55.68	58.59
ACoL	Minmax	57.43	63.69	62.29	61.14	57.38	56.18	66.43	59.99	54.34	57.02	57.25	56.26
	Max	57.17	63.55	62.14	60.95	56.82	55.78	66.45	59.68	53.99	56.83	56.22	55.74
	PaS	58.10	63.75	62.70	61.51	57.29	56.02	66.38	59.90	51.29	52.71	52.43	52.14
	IVR	57.96	64.76	61.95	<b>61.55</b>	60.22	58.78	66.33	<b>61.78</b>	54.13	57.33	59.54	<b>57.00</b>
SPG	Minmax	59.92	63.27	63.27	62.15	56.28	55.91	60.37	57.52	58.31	62.31	56.71	59.11
	Max	60.44	64.63	64.04	63.04	60.21	56.66	60.85	59.24	59.20	64.50	57.42	<b>60.37</b>
	PaS	61.20	63.96	64.32	63.16	55.01	54.74	58.77	56.17	55.67	60.08	54.10	56.62
	IVR	60.86	65.49	64.59	<b>63.65</b>	60.22	58.41	66.56	<b>61.73</b>	58.79	64.08	56.73	59.87
ADL	Minmax	59.86	61.43	63.65	61.65	66.27	58.81	58.33	61.14	58.65	56.85	55.15	56.88
	Max	63.20	62.88	64.59	63.56	67.10	59.92	63.17	63.40	59.87	57.61	55.76	<b>57.75</b>
	PaS	60.86	62.30	64.49	62.55	64.97	57.32	56.67	59.65	55.99	54.42	53.06	54.49
	IVR	63.71	64.01	65.06	<b>64.26</b>	67.25	60.28	64.30	<b>63.94</b>	59.25	56.88	55.03	57.05
CutMix	Minmax	59.46	63.86	63.27	62.20	62.31	57.43	62.76	60.83	58.07	62.56	57.73	59.45
	Max	60.14	65.43	64.60	63.39	69.03	59.79	63.61	64.14	59.51	64.63	59.82	<b>61.32</b>
	PaS	61.07	64.40	64.42	63.30	60.94	55.90	62.58	59.90	55.72	60.12	55.47	57.10
	IVR	60.57	66.04	64.93	<b>63.84</b>	71.54	61.74	67.14	<b>66.80</b>	58.89	63.24	58.95	60.36

map  $\mathbf{F}$  is clearly visible in all normalization methods, we can understand how harmful the low values are in  $\mathbf{F}'$ . If not processed properly, the performance of any method will no longer be evaluated right no matter how well the feature is trained.

**Quantitative analysis:** Fig. 4 shows the distribution of minimum and maximum values within the class activation maps of all test images. Note that values expressed in Fig. 4 are those before normalization and the localization are done after each normalization method. HaS with VGG16 has been used for the comparison in CUB and ADL with VGG16 has been used for ImageNet. The minimum values are plotted on the x-axis while maximum values are plotted on the y-axis. In other words, a single point can express the values range in class activation map  $\mathbf{F}$  of an image. Applying the optimal threshold  $\tau^*$  of IoU 70, black and red dots represent positive and negative samples of localization. In both datasets, min-max normalization shows the most inefficient shape of distribution. Especially, the variation of maximum values does not seem to affect the localization performance in a great deal. Meanwhile, numerous number of samples with highly negative minimum values fail to be localized. This can be proven in that almost all samples with a minimum value under -0.07 fail in localization in CUB. In the graph of max normalization,

much more samples have successfully been localized. In this point, we can assume that highly negative values in the class activation map are unnecessary for localization. The area covered by PaS is a little bit wider than that of min-max normalization but still it does not completely restrain the influence of extremely small minimum values. As explained in the previous sections, IVR can be assumed as a variant of max normalization and the distribution resembles that of max normalization.

### 5.3. Evaluation of WSOL methods

All hyperparameters pioneered by [6] are available online<sup>2</sup>. Using these optimized values obtained from a hyperparameter search, we can avoid issues coming from selecting different hyperparameters. WSOL methods, architectures, and datasets are the same as in [6] as well.

Tab. 1 shows the performances of all concerned experiments. In ImageNet, min-max normalization shows the worst performance in all WSOL methods except ACoL. With min-max normalization, CAM still works better than all the other WSOL methods. When using max normalization, all methods except ACoL improve from min-max normalization slightly. In this case, ADL performs the best for ImageNet. A little bit higher overall performance than max

<sup>2</sup><https://github.com/clovaai/wsolevaluation>

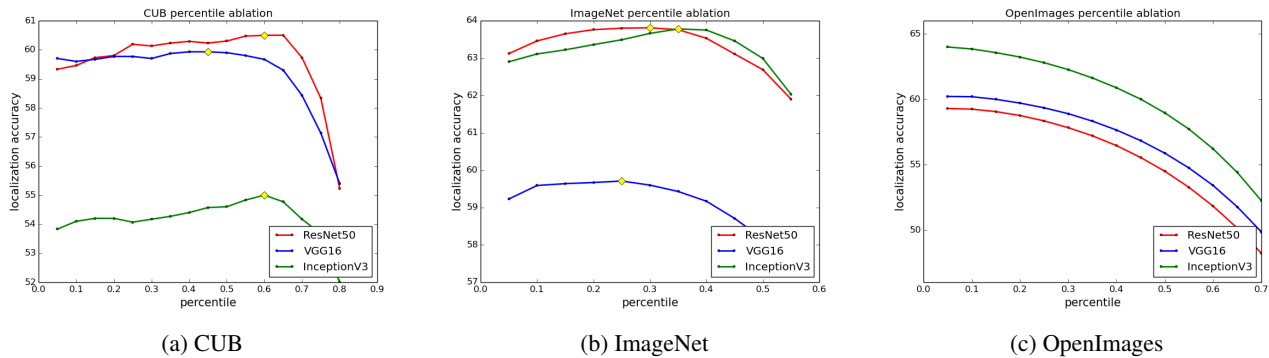


Figure 5: Localization accuracy measured with different percentile values when using IVR. The evaluation has been done only with CAM in the validation set of every datasets. To keep simplicity, the best percentile value in each architecture and dataset has been used in all other WSOL methods.

normalization can be seen in PaS except ADL. Note that ADL with PaS is still better than min-max normalization. IVR has enhanced all WSOL methods without any exception. In this case, ADL has shown the best performance when all architectures are averaged. When measured individually, CutMix with IVR has recorded 66.04% in Inception which is the best ImageNet performance in total.

In CUB, PaS shows lower performance than min-max normalization in all WSOL methods. Note that the reported scores in this paper are higher than those in [1]. For example, [1] reports 65.90 of IoU 50 localization accuracy in CAM with PaS. In our experiment, IoU 50 accuracy of CAM with PaS has recorded 71.61 since much more better hyperparameters from [6] are available. Meanwhile, IVR has shown a drastic improvement in all experiments. Especially, HaS and CutMix have exceeded 70% in VGG, which has never been reached by any other normalization methods. Unlike in ImageNet, the performance gap among WSOL methods becomes considerably noticeable.

In OpenImages, the performance of PaS rather decreases and even the performance of IVR is slightly better than min-max normalization. Max normalization shows an overall improvement and CAM is still the best method. IVR is the second best normalization method while PaS is the worst one. Even with  $P_{\times AP}$ , low values rather than high values affect every pixels in the class activation map leading to a performance degradation. Consequently, CAM still works better than or at least almost equal to the best method in both ImageNet and OpenImages. However, it has yielded its top position to HaS in CUB. The conclusion of [6] that many CAM-based WSOL methods have barely improved the original CAM is correct only in ImageNet and OpenImages.

We have also conducted experiments using PaS and IVR simultaneously. In all cases, the resulting score is much better than using PaS only but slightly worse than using IVR

alone. Therefore, we omit the result in the table to avoid redundancy.

#### 5.4. Percentile selection in IVR

In PaS, the choice of the maximum value percentile has been made empirically. To alleviate the ambiguity of choosing a percentile for IVR, we have used the validation set and chosen the best values. Fig. 5 shows how the percentile value affects the performance in CUB, ImageNet, and OpenImages.

In CUB, class activation maps extracted from images have shown a large variation in minimum values as described in Fig. 4. Consistently, removing much of the portion from an activation map has shown better performance. Using 45-th, 60-th, and 60-th percentile have peaked in VGG, ResNet, and Inception respectively. In all cases, these values are all positive and this implies that even positive values in the class activation map may not contribute to the localization performance. In ImageNet, 25-th, 30-th, and 35-th percentile work best in VGG, ResNet and Inception respectively. In other words, ImageNet uses a wider range in minimum value distribution than CUB. As in Fig. 4, this is because the variation of minimum values in ImageNet is relatively smaller than that in CUB. In OpenImages, using IVR has constantly degraded the performance from 0-th percentile in all architectures. Experiments in Tab. 1 uses 5-th percentile in all the cases.

**Evaluation with NWC [1]** In the previous sections, we have mainly discussed about the normalization of class activation maps and excluded the usage of Thresholded Average Pooling (TAP) and Negative Weight Clamping (NWC) proposed in [1] when using PaS. TAP alleviates the problem that the global averaging pooling layer does not reflect the difference in different channels. By applying a threshold during average pooling, the model can focus on important activations in each channel while training. As this is a tech-

Table 2: Localization accuracy when using Negative Weight Clamping additionally [1].

Method	Norm	Localization accuracy			
		IoU 30	IoU 50	IoU 70	Mean
HaS	Min-max	97.62	84.43	43.89	75.31
	Max	97.45	83.88	43.36	74.89
	PaS	94.74	75.47	39.75	69.99
	IVR	98.41	86.56	46.82	<b>77.26</b>
CutMix	Min-max	98.50	88.13	48.03	78.22
	Max	98.39	87.94	47.53	77.95
	PaS	97.39	82.91	42.87	74.39
	IVR	98.74	89.25	49.59	<b>79.19</b>

nique used in training time, we do not discuss about it in this paper. Meanwhile, NWC is a method presuming that negative weights from the classifier does not contribute to the localization. Negative weights are clamped to zero and low values in the class activation map may be excluded. The authors of [1] claim that PaS improves CAM-based WSOL methods by itself but verification of other methods with NWC is necessary. Tab. 2 shows the comparison of all normalization methods combined with NWC in CUB. Both in HaS and CutMix, PaS still shows the worst performance. Even considering the best localization accuracy 78.58% of HaS at IoU 50 reported in [1], scores of other normalization methods exceed at least 83%. IVR in CutMix has scored 79.19% in average and especially 89.25% at IoU 50. To the best of our knowledge, this is the best localization score reported in CUB.

## 6. Discussion

In experiments above, max normalization shows the best result in OpenImages. This can be attributed to the property of the dataset. Fig. 6 shows the minimum and maximum values of CAM with three different architectures' class activation maps in OpenImages. As in the figure, the variation of minimum values is much smaller than those of CUB and ImageNet. Numerically, the ratio of the standard deviation of all maximum values to the one of all minimum values in VGG has recorded approximately 11, 12.04, and 18.36 in CUB, ImageNet, and OpenImages respectively. In case of ResNet, this ratio has recorded up to 31.79 and the performance improvement of max normalization and IVR in ResNet is insignificant as in VGG. Also, Inception has recorded 8.19 and the performance improvement is better than other architectures. Depending on the dataset and architecture, we can see that this rate and the optimal percentile are inversely proportional. When the variation of maximum values is much more higher, using max normalization instead of any other normalization is likely to be a better choice. Therefore depending on the property of the dataset and architecture, we recommend to use an appropri-




Figure 6: Distribution of minimum and maximum values from class activation map  $F$  of all test images in OpenImages. Compared to CUB and ImageNet, the distribution of maximum values is more influential.

ate normalization method.

Also, all WSOL methods emerged after CAM have indeed improved CAM in CUB, contrary to [6], while they have hardly contributed in ImageNet and OpenImages as concluded in [6]. Because CUB is for a fine-grained classification problem, all the images share common features of a bird. Therefore, a classifier for CUB dataset has to discriminate fine details and the common features have little discriminative power. Recent WSOL methods have focused on compensating these relatively less discriminative features. According to our investigation, WSOL methods studied so far have successfully dealt with this problem. On the other hand, they have not focused on more general classification problems such as ImageNet and OpenImages and there remains much room for investigation for this problem.

## 7. Conclusion

For several years, an issue about whether many proposed weakly supervised object localization methods have actually derived improvement has been raised. Meanwhile, there has been few works dealing with the normalization effect of class activation maps. In this paper, we proposed a new and effective normalization method along with a solid evaluation with many other possible normalization methods. The new proposed normalization method achieves the new state-of-the-art performance in CUB-200-2011 and ImageNet dataset. Also, we point out that the normalization method should be selected according to the traits of the dataset.

For future works in WSOL, we suggest that even though many WSOL methods successfully improved the performance in a dataset like CUB, a new perspective which will also work in real world datasets such as ImageNet and OpenImages is still in need.

## Acknowledgements

This work was supported by NRF grant (2021R1A2C3 006659) and IITP grant (2021-0-00537), both funded by the Government of the Republic of Korea.



## References

- [1] Wonho Bae, Junhyug Noh, and Gunhee Kim. Rethinking class activation mapping for weakly supervised object localization. In *European Conference on Computer Vision*, pages 618–634. Springer, 2020.
- [2] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(2), 2012.
- [3] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.
- [4] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 839–847. IEEE, 2018.
- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [6] Junsuk Choe, Seong Joon Oh, Seungho Lee, Sanghyuk Chun, Zeynep Akata, and Hyunjung Shim. Evaluating weakly supervised object localization methods right. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3133–3142, 2020.
- [7] Junsuk Choe and Hyunjung Shim. Attention-based dropout layer for weakly supervised object localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2219–2228, 2019.
- [8] Jun Fu, Jing Liu, Yuhang Wang, Yong Li, Yongjun Bao, Jinhui Tang, and Hanqing Lu. Adaptive context network for scene parsing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6748–6757, 2019.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.
- [11] Di Lin, Dingguo Shen, Siting Shen, Yuanfeng Ji, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Zigzagnet: Fusing top-down and bottom-up context for object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7490–7499, 2019.
- [12] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [13] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [14] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [15] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015.
- [16] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [17] Wojciech Samek, Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, and Klaus-Robert Müller. Evaluating the visualization of what a deep neural network has learned. *IEEE transactions on neural networks and learning systems*, 28(11):2660–2673, 2016.
- [18] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [19] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [21] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *2017 IEEE international conference on computer vision (ICCV)*, pages 3544–3553. IEEE, 2017.
- [22] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [23] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10781–10790, 2020.
- [24] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010.
- [25] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019.
- [26] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and

- anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9759–9768, 2020.
- [27] Xiaolin Zhang, Yunchao Wei, Jiashi Feng, Yi Yang, and Thomas S Huang. Adversarial complementary learning for weakly supervised object localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1325–1334, 2018.
- [28] Xiaolin Zhang, Yunchao Wei, Guoliang Kang, Yi Yang, and Thomas Huang. Self-produced guidance for weakly-supervised object localization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 597–613, 2018.
- [29] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [30] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.
- [31] Yi Zhu, Karan Sapra, Fitsum A Reda, Kevin J Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8856–8865, 2019.