

C²AM: Contrastive learning of Class-agnostic Activation Map for Weakly Supervised Object Localization and Semantic Segmentation

Jinheng Xie, Jianfeng Xiang, Junliang Chen, Xianxu Hou, Xiaodong Zhao, Linlin Shen*

School of Computer Science & Software Engineering, Shenzhen University, China

Dept. of Computer Science, Wenzhou-Kean University, Wenzhou, China

Shenzhen Institute of Artificial Intelligence of Robotics of Society, Shenzhen, China

Guangdong Key Laboratory of Intelligent Information Processing, Shenzhen University, China

xiejinheng2020@email.szu.edu.cn, hxianxu@gmail.com, llshen@szu.edu.cn

Abstract

While class activation map (CAM) generated by image classification network has been widely used for weakly supervised object localization (WSOL) and semantic segmentation (WSSS), such classifiers usually focus on discriminative object regions. In this paper, we propose Contrastive learning for Class-agnostic Activation Map (C²AM) generation only using unlabeled image data, without the involvement of image-level supervision. The core idea comes from the observation that i) semantic information of foreground objects usually differs from their backgrounds; ii) foreground objects with similar appearance or background with similar color/texture have similar representations in the feature space. We form the positive and negative pairs based on the above relations and force the network to disentangle foreground and background with a class-agnostic activation map using a novel contrastive loss. As the network is guided to discriminate cross-image foreground-background, the class-agnostic activation maps learned by our approach generate more complete object regions. We successfully extracted from C²AM class-agnostic object bounding boxes for object localization and background cues to refine CAM generated by classification network for semantic segmentation. Extensive experiments on CUB-200-2011, ImageNet-1K, and PASCAL VOC2012 datasets show that both WSOL and WSSS can benefit from the proposed C²AM. Code will be available at <https://github.com/CVI-SZU/CCAM>.

1. Introduction

Massive image data and manual annotations are usually required to train deep neural networks for many vi-

*Corresponding Author

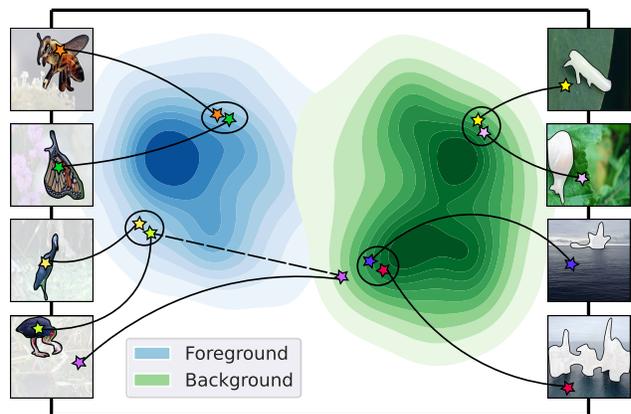


Figure 1. Feature manifold of foreground objects (blue) and backgrounds (green). As semantic information of foreground objects differs from that of backgrounds, the distribution of the representation of foreground objects (blue) is far away from backgrounds (green). Foreground objects with similar appearance or backgrounds with similar color/texture also have similar representations in the feature space. Based on these observations, positive and negative pairs can be formed for contrastive learning. t-SNE [34] is used to reduce the dimensionality of features.

sion tasks, e.g., object detection and semantic segmentation. However, it is time-consuming and labor-intensive to obtain bounding box or pixel-level annotations. In recent years, *weaker supervision*, e.g., image-level label, has been introduced in weakly supervised object localization (WSOL) and semantic segmentation (WSSS), which aims to achieve localization or segmentation with only image-level supervision, i.e., without bounding box or pixel-level annotations. Most previous WSOL and WSSS methods rely on class activation map (CAM) to estimate location of the target object. With image-level supervision, the classifier tries to find discriminative regions of target objects. Thus, though image-level labels enable CAM to indicate the correct location of target objects, they also limit the focus of CAM

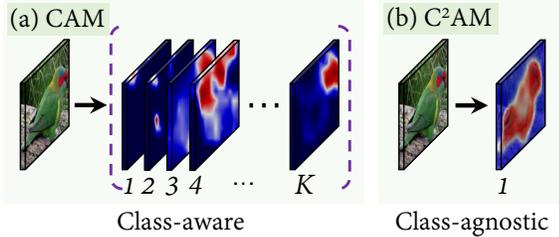


Figure 2. Difference between (a) the class activation map (CAM) and (b) class-agnostic activation map. CAM consists of K (number of class) activation maps and C^2AM only predicts one class-agnostic activation map for an image, which directly indicates the foreground and background regions. Best viewed in color.

on sparse and discriminative object regions. As a result, it could be very difficult to precisely estimate complete object regions by directly applying CAM to WSOL and WSSS.

Different methods have been introduced to mitigate the above problem of CAM. However, most of them are trained with image-level supervision, which would potentially affect the completeness of CAM. By contrast, this paper proposes to generate class-agnostic activation maps based on a novel *cross-image foreground-background contrastive learning*, without the requirement of image-level supervision. In contrast with K class activation maps (CAM) (including one target activation map), C^2AM only predicts one class-agnostic activation map to indicate the foreground and background regions in an image. The difference between CAM and C^2AM is illustrated in Figure 2. Compared with CAM, the class-agnostic activation maps learned from *cross-image foreground-background contrast*, without the involvement of any image-level supervision, produce more reliable foreground regions.

As shown in Figure 1, the semantic information of the foreground object differs from its background, which can be presented as the large distance between foreground and background representations in the feature space. This can be further extended to the cross-image cases where the semantic information of the foreground object from one image should also be largely different from that of background in other images. By contrast, the distance among foregrounds with similar appearance or backgrounds with similar color/texture shall be small. Based on the above observations, we propose the cross-image foreground-background contrastive loss to force the network to disentangle the foreground object and background in an image with a class-agnostic activation map. Specifically, as shown in Figure 3 and 4, the network first generates a class-agnostic activation map using an activation head, such that the image representation can be disentangled into the foreground and background representations, respectively. Then, the foreground and background representations form the negative pairs and the foreground-foreground or background-background representations form the positive pairs. As only

foregrounds with similar appearance or backgrounds with similar color/texture have similar representations in the feature space, pairs of foreground-foreground or background-background sharing less similar semantics might affect the learning of network. To mitigate this problem, we design a feature similarity based rank weighting to automatically reduce the influence of those dissimilar positive pairs. As the activation head is randomly initialized, the initial class-agnostic activation maps are random as well at beginning. When contrastive loss is applied to pull close and push apart the representations of positive and negative pairs, the class-agnostic activation map gradually separates the regions of foreground object and background in the image.

Extensive experiments conducted on the tasks of WSOL and WSSS show that the proposed C^2AM can replace or refine CAM for better performance. Specifically in WSOL, we follow [47] to divide WSOL into two tasks: class-agnostic object localization and object classification, and thus the class-agnostic activation maps can be used to extract class-agnostic object bounding boxes for localization. Background cues can also be extracted from the class-agnostic activation maps to refine the initial CAM such that false activation of background can be efficiently reduced to generate more reliable object regions. The refined CAM can thus greatly enhance the following segmentation performance in WSSS.

Collectively, the main contributions of this paper can be summarized as:

- We propose cross-image foreground-background contrastive learning to generate class-agnostic activation maps with unlabeled image data, which doesn't involve any image-level supervision and produces much more reliable foreground object regions than CAM.
- The class-agnostic activation maps can be used to extract class-agnostic bounding boxes for accurate object localization and serve as background cues to improve the quality of initial CAM.
- Extensive experiments show that both WSOL and WSSS tasks can benefit from our method. We provide an alternative solution to replace or refine CAM for improvements of weakly supervised learning. Besides, it might be further used in a lot of vision tasks to detect foreground regions without manual annotations.

2. Related Works

Class activation mapping (CAM) [50] produces a spatial class activation map that indicates the object regions responsible for the prediction of object category. However, the classifier tends to perform recognition based on the most discriminative regions of the object [12, 30], which affects the precision of CAM based localization or segmentation.

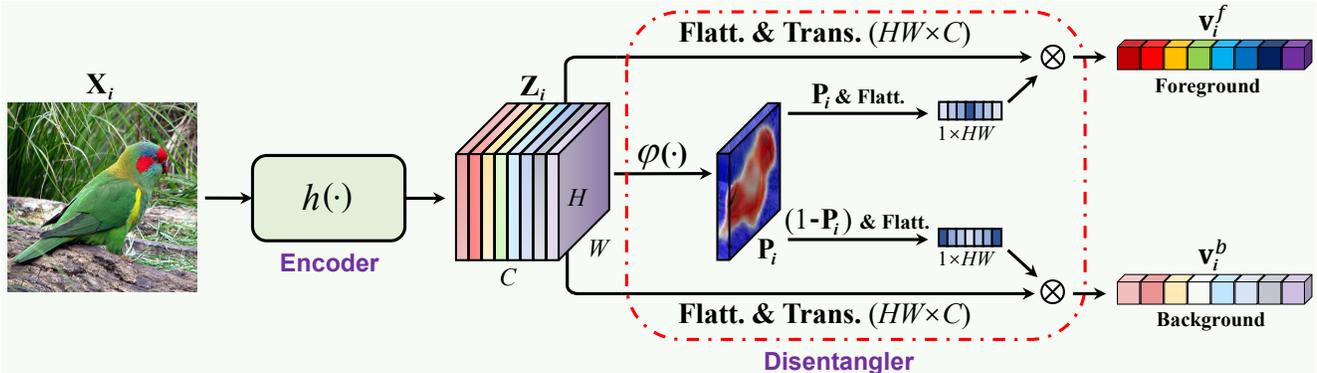


Figure 3. The overall network architecture of the proposed method. The encoder network $h(\cdot)$ maps image \mathbf{X}_i to the feature map \mathbf{Z}_i . In disentangler, the activation head $\varphi(\cdot)$ produces the class-agnostic activation map \mathbf{P}_i . Suppose \mathbf{P}_i activates the foreground regions and the background activation map can be derived as $(1 - \mathbf{P}_i)$. Based on the foreground and background activation maps, \mathbf{Z}_i can be disentangled into the foreground and background feature representations, i.e., \mathbf{v}_i^f and \mathbf{v}_i^b . In evaluation, only the trained $h(\cdot)$ and $\varphi(\cdot)$ are used for generating the class-agnostic activation map \mathbf{P}_i . **Flatt.**: matrix flattening; **Trans.**: matrix transpose; \otimes : matrix multiplication.

Weakly supervised object localization (WSOL). Most of existing WSOL works consider the image-level labels in the training dataset [28, 35]. To solve the above problem, various methods [6, 24, 31, 42–44, 48, 49] have been proposed. Baek *et al.* [2] introduce class-agnostic activation mapping, which can effectively generate heat map from a self-supervised model. Xie *et al.* [42] propose a low-level feature based two-stage learning framework to refine coarse activation maps. Differently, Zhang *et al.* [47] divide WSOL into two tasks: 1) class-agnostic object localization and 2) object classification. This allows a network to complete localization or classification separately. They use DDT [37] to extract class-agnostic object bounding boxes as pseudo labels to train a localization model and use state-of-the-art networks for classification.

Weakly supervised semantic segmentation (WSSS). Most WSSS works follow a three-stage learning process: initial CAM generation, pseudo masks generation, and segmentation model training. Hou *et al.* [13] propose two self-erasing strategies to focus attention only on the reliable regions and generate complete initial CAM. Chang *et al.* [3] propose to mine more object parts by investigating object sub-categories. Jungbeom *et al.* [19] propose an anti-adversarial manner to discover more regions of the target object in the activation map. Ahn and Kwak [1] propose pixel-level semantic affinity, which helps to generate fine-grained pseudo masks. Works [15, 20, 23, 32, 38] additionally use background cues from a fully supervised saliency detector to get better pseudo masks. In our work, the class-agnostic activation maps generated by C^2AM are used to extract background cues to refine the initial CAM, which increases the potential of initial CAM for the next stage.

Contrastive learning. The idea of contrastive learning is to pull close the samples from the positive pair and to push apart the samples from the negative pair [4, 7, 10, 11, 29]. Based on object class labels, the positive pair

is created with the samples from the same class, while the samples with different class labels form the negative pair [7, 10, 17, 29]. Unsupervised contrastive learning can be divided into two categories, i.e., instance-wise contrastive learning [4, 11, 39, 46] and clustering based contrastive learning [21, 41, 45]. In instance-wise contrastive learning, two samples augmented from the same instance form the positive pair, while samples augmented from different instances [4, 11] become the negative pair. For clustering based contrastive learning, the clustering algorithm is applied to generate the pseudo label for training samples, and then the supervised contrastive learning is applied [9].

3. Methodology

3.1. Architecture

The overall network architecture of C^2AM is shown in Figure 3. Given a batch of n images $\mathbf{X}_{1:n} = \{\mathbf{X}_i\}_{i=1}^n$, the encoder $h(\cdot)$ maps \mathbf{X} to high-level feature maps $\mathbf{Z}_{1:n} = \{\mathbf{Z}_i\}_{i=1}^n$, in which $\mathbf{Z}_i \in \mathbb{R}^{C \times H \times W}$. C and $H \times W$ denote the channel number and spatial dimension, respectively. The popular network, e.g., ResNet [12] or VGG [30], is used as the encoder $h(\cdot)$. Supervised or unsupervised pretraining, e.g., moco [11] and detco [40], on ImageNet-1K [28] can be adopted as initialization of $h(\cdot)$. Based on the extracted feature maps \mathbf{Z} , the disentangler employs an activation head $\varphi(\cdot)$ to produce the class-agnostic activation maps $\mathbf{P}_{1:n} = \{\mathbf{P}_i\}_{i=1}^n$, in which $\mathbf{P}_i \in \mathbb{R}^{1 \times H \times W}$. Specifically, $\varphi(\cdot)$ is a 3×3 convolution with a batch normalization layer. Suppose \mathbf{P}_i activates foreground regions, the background activation map of the i -th sample can be formulated as $(1 - \mathbf{P}_i)$. The foreground and background activation maps can finally disentangle the feature map $\mathbf{Z}_{1:n}$ into the foreground and background feature representations, i.e., $\mathbf{v}_{1:n}^f$ and $\mathbf{v}_{1:n}^b$, respectively. For the i -th sample, \mathbf{v}_i^f and

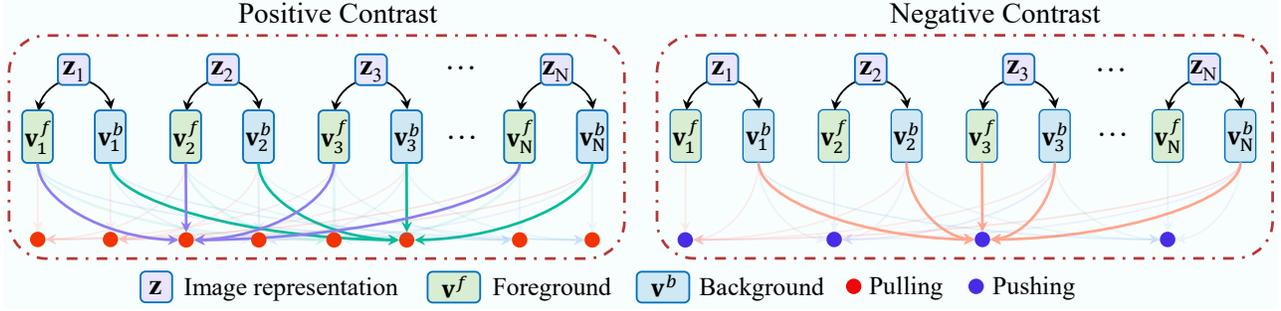


Figure 4. Illustration of cross-image foreground-background contrastive learning. Each image representation, i.e., \mathbf{Z}_i , is disentangled into the foreground and background representations, i.e., \mathbf{v}_i^f and \mathbf{v}_i^b . Two foreground or background representations are coupled into one positive pair, while a negative pair is formed with one foreground and one background representation. Contrastive learning is applied to pull close the representations from the positive pair and push apart the representations from the negative pair. Best viewed in color.

\mathbf{v}_i^b can be derived as:

$$\mathbf{v}_i^f = \mathbf{P}_i \otimes \mathbf{Z}_i^\top, \quad \mathbf{v}_i^b = (1 - \mathbf{P}_i) \otimes \mathbf{Z}_i^\top. \quad (1)$$

Here, \mathbf{P}_i and \mathbf{Z}_i are flattened, i.e., $\mathbf{P}_i \in \mathbb{R}^{1 \times HW}$ and $\mathbf{Z}_i \in \mathbb{R}^{C \times HW}$. $\mathbf{v}_i^f \in \mathbb{R}^{1 \times C}$ and $\mathbf{v}_i^b \in \mathbb{R}^{1 \times C}$. \otimes and \top indicate the matrix multiplication and transpose, respectively.

3.2. Foreground-background Contrast

Considering the fact that there is no label information in the training process, C²AM is proposed based on the idea of *cross-image foreground-background contrast*, which locates the foreground object regions by only utilizing the semantic information among foreground and background representations. As aforementioned, given an image, its foreground and background representations contain different semantic information and thus should have a large distance in the feature space. This observation is also true for cross-image cases. The distance between their foreground and background representations should also be large.

Based on these analyses, we propose to apply the contrastive learning to push apart foreground-background representations. Given n samples $\mathbf{X}_{1:n}$, the disentangler separates them into n foreground and n background representations, i.e., $\mathbf{v}_{1:n}^f$ and $\mathbf{v}_{1:n}^b$. The foreground-background representation pair, i.e., $(\mathbf{v}_i^f, \mathbf{v}_j^b)$, is treated as the *negative pair*. The negative contrastive loss is designed as:

$$\mathcal{L}_{NEG} = -\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \log(1 - s_{i,j}^{neg}), \quad (2)$$

$$s_{i,j}^{neg} = \text{sim}(\mathbf{v}_i^f, \mathbf{v}_j^b), \quad (3)$$

where $s_{i,j}^{neg}$ is the cosine similarity between \mathbf{v}_i^f and \mathbf{v}_j^b . \mathcal{L}_{NEG} considers the contrastive comparisons both within image ($i = j$) and cross-image ($i \neq j$).

3.3. Foreground-foreground and Background-background Contrast with Rank Weighting

The foreground and foreground or background and background representations from two different images form

the *positive pair*. However, only representations of foregrounds with similar appearance and background with similar color/texture have a smaller distance in the feature space and shall be pulled together. Positive pairs with large distances will affect the learning process, as there is fewer similar semantics in these two foreground objects or backgrounds. To address this issue, we design a feature similarity based rank weighting to automatically reduce the influence of those dissimilar positive pairs. We first calculate the cosine similarity between each candidate positive pair:

$$s_{i,j}^f = \text{sim}(\mathbf{v}_i^f, \mathbf{v}_j^f), \quad s_{i,j}^b = \text{sim}(\mathbf{v}_i^b, \mathbf{v}_j^b), \quad (4)$$

where $s_{i,j}^f$ and $s_{i,j}^b$ are the cosine similarity calculated from the foreground-foreground pair, i.e., $(\mathbf{v}_i^f, \mathbf{v}_j^f)$, and background-background pair, i.e., $(\mathbf{v}_i^b, \mathbf{v}_j^b)$, respectively. $s_{i,j}^f$ potentially indicates whether the foreground object from \mathbf{X}_i shares similar semantics with that from \mathbf{X}_j . Given the set of similarity from foreground-foreground representations $s^f = \{s_{1,2}^f, \dots, s_{i,j}^f, \dots\} (i \neq j)$ and the set of similarity from background-background representations $s^b = \{s_{1,2}^b, \dots, s_{i,j}^b, \dots\} (i \neq j)$, we then calculate a weight based on ranking for each positive pair as follow:

$$w_{i,j}^f = \exp(-\alpha \cdot \text{rank}(s_{i,j}^f)), \quad w_{i,j}^b = \exp(-\alpha \cdot \text{rank}(s_{i,j}^b)), \quad (5)$$

where α is a hyper-parameter controlling the smoothness of exponential function. $\text{rank}(s_{i,j}^f)$ and $\text{rank}(s_{i,j}^b)$ are the rank of $s_{i,j}^f$ and $s_{i,j}^b$ in the set of s^f and s^b , respectively. The rank weight $w_{i,j}$ ranges from 0 to 1. A large weight is assigned to positive pairs sharing similar semantics (e.g., similar appearance, color or texture) and a small one is assigned to the positive pairs with less similarity. It can reduce the influence of those dissimilar pairs to some extent for better contrastive learning. The final positive contrastive loss is formulated as:

$$\mathcal{L}_{POS} = \mathcal{L}_{POS}^f + \mathcal{L}_{POS}^b, \quad (6)$$

$$\mathcal{L}_{POS}^f = -\frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^n \mathbb{1}_{[i \neq j]} (w_{i,j}^f \cdot \log(s_{i,j}^f)), \quad (7)$$

$$\mathcal{L}_{POS}^b = -\frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j=1}^n \mathbb{1}_{[i \neq j]} (w_{i,j}^b \cdot \log(s_{i,j}^b)), \quad (8)$$

where $\mathbb{1}_{[i \neq j]} \in \{0, 1\}$ is an indicator function and it equals 1 if $i \neq j$. The overall contrastive loss \mathcal{L} is formulated as the summation of \mathcal{L}_{POS} and \mathcal{L}_{NEG} :

$$\mathcal{L} = \mathcal{L}_{POS} + \mathcal{L}_{NEG}. \quad (9)$$

When contrastive loss \mathcal{L} is applied to pull close and push apart the representations in positive and negative pairs, the class-agnostic activation map gradually separates the regions of foreground object and background in the image.

How to determine foreground regions. As the same contrastive loss is applied to foreground-foreground and background-background positive pairs, it can not guarantee that foreground or background regions are activated in \mathbf{P}_i . To solve this problem, we set a threshold to binarize the class-agnostic activation maps and detect the largest contour to determine the object regions.

3.4. Weakly Supervised Object Localization

We follow the route of PSOL [47] to adapt our C²AM to WSOL. Specifically, WSOL is divided into two tasks: class-agnostic object localization and object classification. PSOL uses DDT [37] to generate class-agnostic object bounding boxes on the training set. Given a group of images of the same category, DDT obtains those category-consistent regions to extract class-agnostic object bounding boxes. By contrast, the proposed C²AM directly learns class-agnostic activation maps from the whole dataset without any manual annotations. We set a threshold to binarize the class-agnostic activation maps and then extract class-agnostic object bounding boxes as pseudo labels (we present a fair comparison between DDT and C²AM in the next section). The localization model is trained with these pseudo labels for object bounding box predictions. The popular network, e.g., EfficientNet [33], is adopted for object classification.

3.5. Weakly Supervised Semantic Segmentation

We first use CAM-based methods to generate initial CAM for each image and then apply C²AM to refine it. Specifically, we use the background activation maps ($1 - \mathbf{P}$) as pseudo labels to further train a model to predict the background regions, i.e., background cues, in the image. As shown in Figure 5, we concatenate the predicted background cues with the initial CAM and perform the argmax process along the channel dimension to refine initial CAM. This helps to reduce the false activation of background and activate more foreground regions in initial CAM. More details of the pseudo training and refinement procedure are

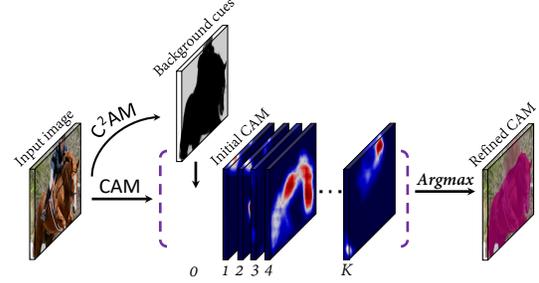


Figure 5. Refinement of initial CAM using background cues.

included in supplementary materials. We just use a simple way to demonstrate the effectiveness of our C²AM for refinement of initial CAM and leave more works in the future.

4. Experiments

4.1. Experimental Setup

Dataset. CUB-200-2011 [35] is a fine-grained classification dataset, which consists of 200 species of birds with 5,994 images for training and 5,794 images for test. ImageNet-1K [28] is a large scale visual recognition dataset with 1,000 classes, which contains 1,281,197 training images and 50,000 validation images. PASCAL VOC2012 [8] is a popular semantic segmentation dataset with 20 object categories. It consists of 1,464 images for training, 1,449 images for validation, and 1,456 images for test.

Evaluation Metrics. In WSOL, we follow [50] to use Top-1 localization accuracy (Top-1 Loc), Top-5 localization accuracy (Top-5 Loc), and GT-known localization accuracy (GT-known Loc) for evaluation. For GT-known Loc, a localization is correct when the predicted bounding box overlaps over 50% with one of the ground truth bounding boxes belonging to the same class. For Top-1 Loc, a prediction is correct when the Top-1 classification result and GT-known Loc are both correct. For Top-5 Loc, a prediction is correct when one of the Top-5 classification results and GT-known Loc are both correct. The MaxBoxAccV2 proposed in [5] is also adopted, which averages the performance across different bounding boxes ratios, i.e., 30%, 50%, and 70%. In WSSS, the Intersection over Union (IoU) and mean Intersection over Union (mIoU) is adopted as the evaluation metric.

The implementation details are provided in the supplementary materials.

5. Results and Analysis

We first present a visual comparison between CAM and the proposed C²AM to validate that more complete and correct foreground regions can be predicted in the class-agnostic activation maps generated by C²AM. Then, we apply C²AM to WSOL and show that CAM can be replaced by the class-agnostic activation maps for more accurate object localization. Besides, we also apply C²AM to

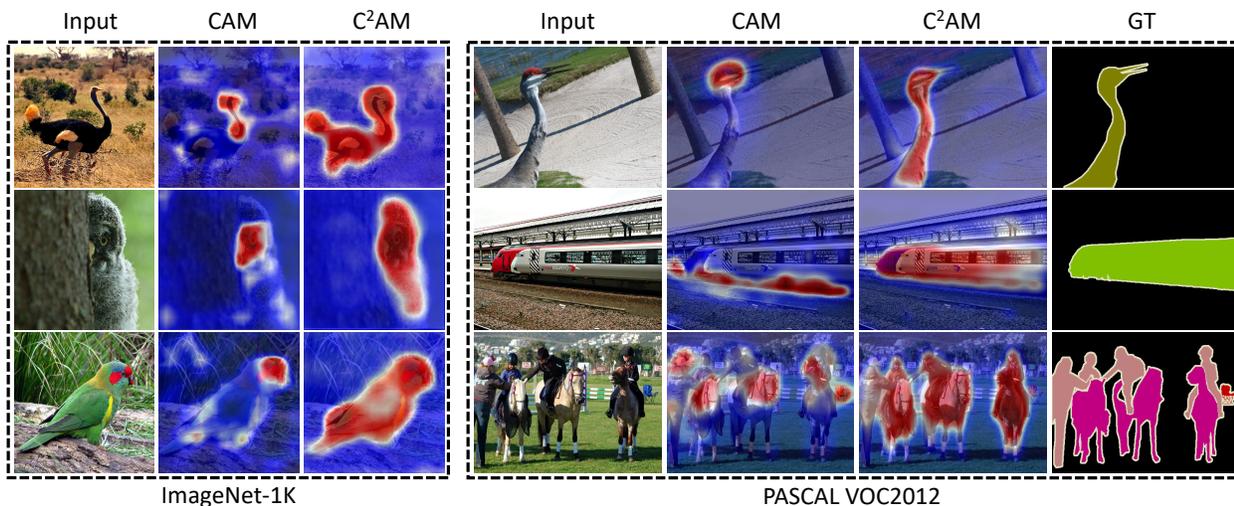


Figure 6. Visual comparison between CAM and the class-agnostic activation maps generated by C^2AM . Best viewed in color.

Table 1. Comparison of the performance between the proposed method and the state-of-the-art methods on CUB-200-2011 test set and ImageNet-1K validation set. Loc Back. denotes the localization backbone. Cls Back. denotes the backbone for classification. \dagger and $\dagger\dagger$ indicate the $h(\cdot)$ is initialized by the supervised and unsupervised (moco) pretraining, respectively. *: our re-implementation results.

Method	Loc Back.	Cls Back.	CUB-200-2011			ImageNet-1K		
			Top-1 Loc	Top-5 Loc	GT-known Loc	Top-1 Loc	Top-5 Loc	GT-known Loc
CAM _{CVPR'16} [50]		VGG-GAP	36.13	-	-	42.80	54.86	59.00
ADL _{CVPR'19} [6]		VGG-GAP	52.36	-	73.96	44.92	-	-
I ² C _{ECCV'20} [49]		InceptionV3	65.99	68.34	72.60	53.11	64.13	68.50
GC-Net _{ECCV'20} [24]		GoogLeNet	58.58	71.10	75.30	49.06	58.09	-
SPA _{CVPR'21} [26]		VGG16	60.27	72.5	77.29	49.56	61.32	65.05
FAM _{ICCV'21} [25]		VGG16	69.26	-	89.26	51.96	-	71.73
ORNet _{ICCV'21} [42]		VGG16	67.74	80.77	86.2	52.05	63.94	68.27
CAM _{CVPR'16} [50]		DenseNet161	29.81	39.85	-	39.61	50.40	52.54
PSOL _{CVPR'20} [47]	ResNet50	ResNet50	70.68	86.64	90.00	53.98	63.08	65.44
PSOL _{CVPR'20} [47]	DenseNet161	EfficientNet-B7	80.89*	89.97*	91.78*	58.00	65.02	66.28
$C^2AM(Ours)^\dagger$	DenseNet161	EfficientNet-B7	83.28	92.74	94.46	59.28	66.72	68.20
$C^2AM(Ours)^{\dagger\dagger}$	DenseNet161	EfficientNet-B7	81.76	91.11	92.88	59.56	67.05	68.53

WSSS and show that the class-agnostic activation maps can serve as background cues to reduce false-activation of background in CAM and segment more reliable object regions. Details are presented in the following sections.

Table 2. The GT-Known Loc of DDT and our C^2AM on the CUB-200-2011 and ImageNet-1K. Init. denotes the parameters initialization. sup-pre.: the parameters initialized from supervised pre-training. moco [11] and detco [40]: the unsupervised pretraining. *: the result is obtained with our re-implementation.

Method	Bac.	Init.	CUB-200-2011	ImageNet-1K
DDT	VGG16	sup-pre.	84.55	61.41
C^2AM	VGG16	sup-pre.	75.34	63.41
DDT	ResNet50	sup-pre.	72.39	59.92
C^2AM	ResNet50	sup-pre.	89.99	65.89
DDT*	ResNet50	moco	34.95	40.52
C^2AM	ResNet50	moco	89.90	66.51
DDT*	ResNet50	detco	35.86	41.23
C^2AM	ResNet50	detco	88.21	65.48

Table 3. Evaluation results in terms of MaxBoxAccV2.

Methods	ImageNet			CUB		
	VGG	Inception	ResNet50	VGG	Inception	ResNet50
Best WSOL	60.6	63.9	63.7	66.3	58.8	66.4
$C^2AM(Ours)$	66.3	65.8	66.8	81.4	82.4	83.8

5.1. Visual Comparison with CAM

Figure 6 presents the visual comparison between CAM and the proposed C^2AM on ImageNet-1K and PASCAL VOC2012 datasets. The second and fifth columns present the visualization of CAM. It can be seen that with image-level supervision, CAM can correctly indicate the location of the object but usually focus only on the most discriminative regions of the target object, e.g., the head of bird. Besides, CAM also activates class-related regions, e.g., the railroad, and ignores the complete body of the train. These issues degrade the performance of CAM in WSOL and WSSS. In contrast, the class-agnostic activation maps gen-

Table 4. Semantic segmentation performance (IoU(%)) comparison for 11 categories between the initial CAM and CAM refined by using the proposed C²AM. The initial CAM is generated using PSA, SC-CAM, SEAM, PuzzleCAM, and AdvCAM, respectively. †† indicates that unsupervised pretraining (moco) is adopted to initialize the backbone network $h(\cdot)$ of C²AM.

Method	bkg	areo	bike	bird	boat	bottle	bus	car	cat	chair	cow
PSA [1]	78.0	41.0	27.5	42.0	34.4	44.5	63.3	53.3	43.2	30.9	49.3
+ C ² AM††	88.0	75.4	45.1	79.5	52.6	64.8	78.0	73.5	81.4	27.0	78.2
SC-CAM [3]	78.6	42.1	29.2	44.5	37.0	56.1	69.9	58.2	59.8	27.8	52.5
+ C ² AM††	88.0	74.8	46.6	80.5	52.8	68.5	79.0	74.6	84.7	26.1	78.3
SEAM [36]	82.8	51.0	35.4	57.4	31.6	50.0	57.9	63.2	62.4	27.3	63.0
+ C ² AM††	87.1	71.9	44.2	77.9	45.6	61.9	75.5	70.4	82.5	22.5	78.0
PuzzleCAM [16]	78.5	43.2	32.3	36.7	23.1	51.7	67.6	61.4	76.7	16.7	60.4
+ C ² AM††	87.6	71.5	48.4	78.0	44.6	66.7	77.3	74.3	85.7	21.3	80.2
AdvCAM [19]	81.3	50.6	33.5	57.5	37.0	53.3	67.8	54.8	64.7	35.0	68.4
+ C ² AM††	87.5	72.9	46.4	78.9	50.7	60.5	77.8	71.2	84.5	26.8	79.5

erated by C²AM effectively mitigate the above issues. As shown in the left of Figure 6, C²AM can completely separate the foreground objects out of background regions, in which the whole body of bird is activated. Besides, as shown in the right of Figure 6, C²AM successfully discriminate the railroad from the train and thus the complete region of train can be activated.

5.2. Results of WSOL

Class-agnostic object bounding boxes. Table 2 compares the quality of class-agnostic bounding boxes generated by C²AM and DDT [37]. When supervised pretraining on ImageNet-1K are adopted as initialization of $h(\cdot)$, class-agnostic bounding boxes generated by C²AM have a higher GT-known Loc than DDT (except the backbone of VGG16). This indicates that our method has a good capability to generate higher quality class-agnostic bounding boxes than DDT. When considering unsupervised pretraining, e.g., moco [11] and detco [40], the GT-known Loc of DDT on two benchmarks has a large drop. By contrast, our C²AM maintains high performance, which is even better than that of DDT using supervised pretraining.

Results on CUB-200-2011. Table 1 compare the performance of C²AM and state-of-the-art methods on CUB-200-2011. Our method achieves Top-1 Loc of 83.28%, Top-5 Loc of 92.74%, and GT-known Loc of 94.46%. When compared with SPA and ORNet, our C²AM surpasses them by a large margin in terms of Top-1 Loc, Top-5 Loc, and GT-known Loc. Compared with PSOL, our C²AM with supervised pretraining surpasses it by 2.39%, 2.77%, and 2.68%, in terms of Top-1 Loc, Top-5 Loc, and GT-known Loc, respectively. When $h(\cdot)$ is initialized by moco, C²AM also achieves better performance than PSOL. We also provide the results of MaxBoxAccV2 metric in Table 3, where the result of Best WSOL is directly taken from [5]. One can observe that our C²AM achieves significantly better performance for different networks and datasets.

Results on ImageNet-1K. Table 1 also compares our method with other recently introduced weakly supervised

Table 5. Semantic segmentation performance (mIoU(%)) comparison among different methods. dCRF: dense CRF [18]. FSSD: background cues from fully supervised saliency detector [22]. ††: moco is adopted to initialize the backbone network $h(\cdot)$ of C²AM.

Method	CAMs	+ dCRF	+ FSSD	+ C ² AM††
PSA CVPR'18 [1]	48.0	-	62.0	65.5 (+17.5)
SC-CAM CVPR'20 [3]	50.9	55.3	62.1	66.0 (+15.1)
SEAM CVPR'20 [36]	55.4	56.8	61.5	63.9 (+8.5)
PuzzleCAM ICI'21 [16]	51.5	-	62.1	65.5 (+14.0)
AdvCAM CVPR'21 [19]	55.6	62.1	62.3	65.4 (+9.8)

object localization methods on ImageNet-1K validation set. Our method achieves Top-1 Loc of 59.56%, Top-5 Loc of 67.05%, and GT-known Loc of 68.53%, which achieves the state-of-the-art performance. When using the same localization (DenseNet161 [14]) and classification backbone (EfficientNet-B7 [33]), the performance of C²AM on the validation set significantly surpasses PSOL by 1.56%, 2.03%, and 2.25% in terms of Top-1 Loc, Top-5 Loc, and GT-known Loc, respectively.

5.3. Results of WSSS

Results on PASCAL VOC2012. As this section is designed to demonstrate the effectiveness of C²AM to refine initial CAM and the refined CAM can be further used in other stages of WSSS, the IoU and mIoU are reported on the training set of PASCAL VOC2012. Table 4 presents improvement of each category for existing CAM-based methods by using C²AM. It can be seen that the segmentation performances of many CAM-based methods for most categories, including PSA [1], SC-CAM [3], SEAM [36], PuzzleCAM [16], and AdvCAM [19], are largely improved by using C²AM. For example, there is a mean of 14.6%, 31.3%, 20.1% IoU improvements of these five methods for bike, bird, and cow, respectively.

To compare our C²AM with other CAM refinement approaches, Table 5 presents the segmentation performance of initial CAMs generated by literature works like PSA, SC-CAM, and PuzzleCAM etc., and the improvements after they are refined by dense CRF (+dCRF) [18], and background cues generated by a fully supervised saliency detector (+FSSD) [22], and our C²AM (+C²AM). The results of dense CRF are reported in their papers and the refinement process presented in Section 3.5 is used for both +FSSD and +C²AM. It can be clearly seen that the quality of initial CAM has been greatly improved by using our C²AM. In particular, the mIoU of initial CAM generated by PSA [1] and SC-CAM [3] have been improved by 17.5% and 15.1%, respectively. In particular, the increase of mIoU caused by background cues generated from C²AM is even higher than that generated by the fully supervised saliency detector.

We also provide a visualization of CAM refinement in Figure 7. It can be observed in the first column that, a lot of object regions, e.g., the body of the white cow and the feet

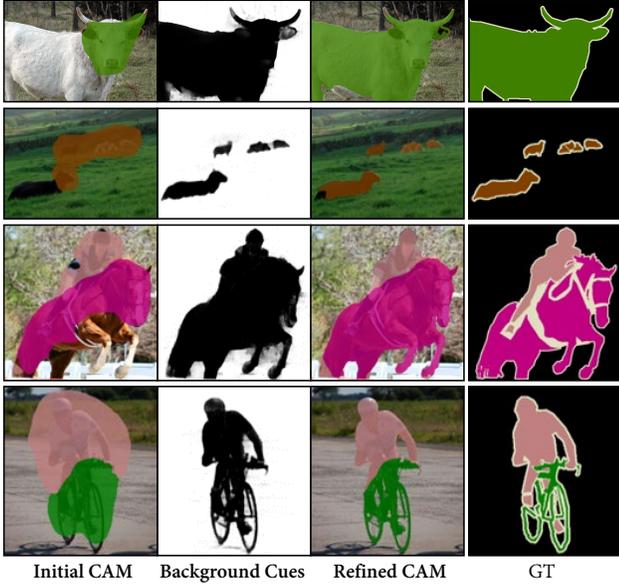


Figure 7. Illustration of CAM refinement using background cues extracted from C^2AM . First column: initial CAM. Second column: background cues extracted from C^2AM (255: background, 0: foreground). Third column: CAM refined using background cues. Last column: ground-truth masks. Best viewed in color.

Table 6. Sensitivity analysis of parameters initialization of 3×3 convolution in activation head. k_normal and $k_uniform$ are short for $kaiming_normal$ and $kaiming_uniform$ initialization in PyTorch [27]. Results are reported on ImageNet-1K validation set.

Metric	k_normal	$k_uniform$	normal	uniform
GT-known Loc	66.51	66.44	66.82	66.67

of the brown horse, are not activated in the initial CAM. Besides, backgrounds, e.g., the grass and ground, are usually falsely activated. To address these issues, we use the class-agnostic activation maps to extract background cues (as shown in the second column) to efficiently reduce the false activation of background and predict more reliable object regions. The results shown in the third column highly overlaps with the ground truth (in the last column).

5.4. Sensitivity Analysis

We conduct experiments on ImageNet-1K to examine the hyper-parameter sensitivity of C^2AM . The GT-known Loc is adopted as a performance evaluation.

Initialization of Activation Head. A 3×3 convolution with a batch normalization layer is adopted as the activation head $\varphi(\cdot)$. The backbone network $h(\cdot)$ can be initialized with supervised or unsupervised pretraining (e.g., moco, detco). How about the activation head? Does the initialization of activation head affect the performance of C^2AM ? We perform experiments with four different parameters initialization of activation head, i.e., $kaiming_normal$, $kaiming_uniform$, normal, and uniform initialization. The weights and bias in batch normalization layer are initial-

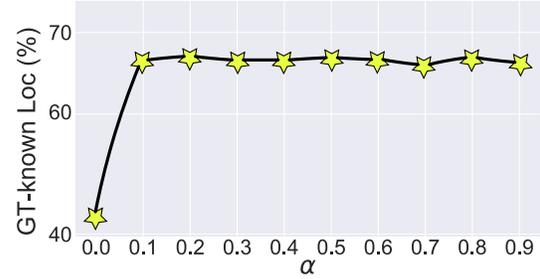


Figure 8. Sensitivity analysis of hyper-parameter α in Eq. 5.

ized with 1 and 0. Table 6 compares the performance of these four initialization choices. We can see that the performances are quite stable among these initialization methods, which shows that the activation head of C^2AM is not sensitive to the parameters initialization.

Hyper-parameter α . Figure 8 presents the GT-known Loc of C^2AM with different setting of α in Eq. 5. The results are reported on ImageNet-1K validation set. When α equals 0, every $w_{i,j}$ is equal to 1, and every positive pair including those foreground with different appearance and background with significant different color/texture are also pulled together. This violates the assumption that only representations with similar appearance, color or texture can be pulled close. Thus, the GT-known Loc is the lowest when α is set to 0. When α is larger than 0, different positive pairs are assigned with different weights based on the feature similarity. As seen, the Gt-Known Loc has been largely improved when α is larger than 0, demonstrating that the proposed rank weighting works well to automatically reduce the influence of those dissimilar positive pairs. When α varies from 0.1 to 0.9, the number of positive pairs involved for contrastive learning is gradually decreasing. The stable performances show that our approach is not sensitive to the number of positive pairs (i.e., the setting of α).

6. Conclusions and Discussions

We propose cross-image foreground-background contrast for class-agnostic activation maps generation using unlabeled image data. Class-agnostic activation map determines more reliable foreground regions, which can be used to replace or refine CAM for better WSOL and WSSS performance. Extensive experiments show that both WSOL and WSSS can benefit from our approach.

As the proposed C^2AM works well in the unlabeled training data, we believe that it can be further used to efficiently detect the foreground regions without manual annotations in a lot of vision tasks, e.g., saliency detection and skin lesion segmentation, etc.

Acknowledgments

The work was supported by the National Natural Science Foundation of China under grant 91959108.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, pages 4981–4990, 2018. 3, 7
- [2] Kyungjune Baek, Minhyun Lee, and Hyunjung Shim. Psynet: Self-supervised approach to object localization using point symmetric transformation. In *AAAI*, pages 10451–10459, 2020. 3
- [3] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation via sub-category exploration. In *CVPR*, pages 8988–8997, 2020. 3, 7
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020. 3
- [5] Junsuk Choe, Seong Joon Oh, Seungho Lee, Sanghyuk Chun, Zeynep Akata, and Hyunjung Shim. Evaluating weakly supervised object localization methods right. In *CVPR*, pages 3133–3142, 2020. 5, 7
- [6] Junsuk Choe and Hyunjung Shim. Attention-based dropout layer for weakly supervised object localization. In *CVPR*, pages 2219–2228, 2019. 3, 6
- [7] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, pages 539–546, 2005. 3
- [8] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 5
- [9] Yixiao Ge, Feng Zhu, Dapeng Chen, Rui Zhao, and Hongsheng Li. Self-paced contrastive learning with hybrid memory for domain adaptive object re-id. In *NeurIPS*, 2020. 3
- [10] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, pages 1735–1742, 2006. 3
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020. 3, 6, 7
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 2, 3
- [13] Qibin Hou, Peng-Tao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. In *NeurIPS*, pages 547–557, 2018. 3
- [14] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, pages 2261–2269, 2017. 7
- [15] Peng-Tao Jiang, Qibin Hou, Yang Cao, Ming-Ming Cheng, Yunchao Wei, and Hong-Kai Xiong. Integral object mining via online attention accumulation. In *ICCV*, pages 2070–2079, 2019. 3
- [16] Sanghyun Jo and In-Jae Yu. Puzzle-cam: Improved localization via matching partial and full features. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 639–643, 2021. 7
- [17] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020. 3
- [18] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NeurIPS*, pages 109–117, 2011. 7
- [19] Jungbeom Lee, Eunji Kim, and Sungroh Yoon. Anti-adversarially manipulated attributions for weakly and semi-supervised semantic segmentation. In *CVPR*, pages 4071–4080, 2021. 3, 7
- [20] 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 *CVPR*, pages 5495–5505, 2021. 3
- [21] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *ICLR*, 2021. 3
- [22] Jiangjiang Liu, Qibin Hou, Ming-Ming Cheng, Jiashi Feng, and Jianmin Jiang. A simple pooling-based design for real-time salient object detection. In *CVPR*, pages 3917–3926, 2019. 7
- [23] Yun Liu, Yu-Huan Wu, Pei-Song Wen, Yu-Jun Shi, Yu Qiu, and Ming-Ming Cheng. Leveraging instance-, image- and dataset-level information for weakly supervised instance segmentation. *TPAMI*, 2020. 3
- [24] Weizeng Lu, Xi Jia, Weicheng Xie, Linlin Shen, Yicong Zhou, and Jinming Duan. Geometry constrained weakly supervised object localization. In *ECCV*, pages 481–496, 2020. 3, 6
- [25] Meng Meng, Tianzhu Zhang, Qi Tian, Yongdong Zhang, and Feng Wu. Foreground activation maps for weakly supervised object localization. In *ICCV*, pages 3385–3395, 2021. 6
- [26] Xingjia Pan, Yingguo Gao, Zhiwen Lin, Fan Tang, Weiming Dong, Haolei Yuan, Feiyue Huang, and Changsheng Xu. Unveiling the potential of structure preserving for weakly supervised object localization. In *CVPR*, pages 11642–11651, 2021. 6
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, pages 8026–8037, 2019. 8
- [28] 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. *IJCV*, pages 211–252, 2015. 3, 5
- [29] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015. 3
- [30] Karen Simonyan and Andrew Zisserman. Very convolutional networks for large-scale image recognition. In *ICLR*, 2015. 2, 3

- [31] Krishna Kumar Singh and Yong Jae Lee. Hide-and-peek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *ICCV*, pages 3544–3553, 2017. [3](#)
- [32] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *ECCV*, pages 347–365, 2020. [3](#)
- [33] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *ICML*, pages 6105–6114, 2019. [5](#), [7](#)
- [34] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(11), 2008. [1](#)
- [35] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. [3](#), [5](#)
- [36] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *CVPR*, pages 12275–12284, 2020. [7](#)
- [37] Xiu-Shen Wei, Chen-Lin Zhang, Jianxin Wu, Chunhua Shen, and Zhi-Hua Zhou. Unsupervised object discovery and co-localization by deep descriptor transformation. *PR*, 88:113–126, 2019. [3](#), [5](#), [7](#)
- [38] 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 *CVPR*, pages 16765–16774, 2021. [3](#)
- [39] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, pages 3733–3742, 2018. [3](#)
- [40] Enze Xie, Jian Ding, Wenhai Wang, Xiaohang Zhan, Hang Xu, Peize Sun, Zhenguo Li, and Ping Luo. Detco: Unsupervised contrastive learning for object detection. In *ICCV*, pages 8392–8401, October 2021. [3](#), [6](#), [7](#)
- [41] Junyuan Xie, Ross B. Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *ICML*, pages 478–487, 2016. [3](#)
- [42] Jinheng Xie, Cheng Luo, Xiangping Zhu, Ziqi Jin, Weizeng Lu, and Linlin Shen. Online refinement of low-level feature based activation map for weakly supervised object localization. In *ICCV*, pages 132–141, 2021. [3](#), [6](#)
- [43] Haolan Xue, Chang Liu, Fang Wan, Jianbin Jiao, Xiangyang Ji, and Qixiang Ye. Danet: Divergent activation for weakly supervised object localization. In *ICCV*, pages 6589–6598, 2019. [3](#)
- [44] Haolan Xue, Chang Liu, Fang Wan, Jianbin Jiao, Xiangyang Ji, and Qixiang Ye. Danet: Divergent activation for weakly supervised object localization. In *ICCV*, pages 6589–6598, 2019. [3](#)
- [45] Jianwei Yang, Devi Parikh, and Dhruv Batra. Joint unsupervised learning of deep representations and image clusters. In *CVPR*, pages 5147–5156, 2016. [3](#)
- [46] Mang Ye, Xu Zhang, Pong C. Yuen, and Shih-Fu Chang. Supervised embedding learning via invariant and spreading instance feature. In *CVPR*, pages 6210–6219, 2019. [3](#)
- [47] Chen-Lin Zhang, Yun-Hao Cao, and Jianxin Wu. Rethinking the route towards weakly supervised object localization. In *CVPR*, pages 13457–13466, 2020. [2](#), [3](#), [5](#), [6](#)
- [48] Xiaolin Zhang, Yunchao Wei, Guoliang Kang, Yi Yang, and Thomas Huang. Self-produced guidance for weakly-supervised object localization. In *ECCV*, pages 597–613, 2018. [3](#)
- [49] Xiaolin Zhang, Yunchao Wei, and Yi Yang. Inter-image communication for weakly supervised localization. In *ECCV*, pages 271–287, 2020. [3](#), [6](#)
- [50] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. [2](#), [5](#), [6](#)