

Learning Pixel-level Semantic Affinity with Image-level Supervision for Weakly Supervised Semantic Segmentation

Jiwoon Ahn
DGIST, Korea
jyun@dgist.ac.kr

Suha Kwak
POSTECH, Korea
suha.kwak@postech.ac.kr

Abstract

The deficiency of segmentation labels is one of the main obstacles to semantic segmentation in the wild. To alleviate this issue, we present a novel framework that generates segmentation labels of images given their image-level class labels. In this weakly supervised setting, trained models have been known to segment local discriminative parts rather than the entire object area. Our solution is to propagate such local responses to nearby areas which belong to the same semantic entity. To this end, we propose a Deep Neural Network (DNN) called AffinityNet that predicts semantic affinity between a pair of adjacent image coordinates. The semantic propagation is then realized by random walk with the affinities predicted by AffinityNet. More importantly, the supervision employed to train AffinityNet is given by the initial discriminative part segmentation, which is incomplete as a segmentation annotation but sufficient for learning semantic affinities within small image areas. Thus the entire framework relies only on image-level class labels and does not require any extra data or annotations. On the PASCAL VOC 2012 dataset, a DNN learned with segmentation labels generated by our method outperforms previous models trained with the same level of supervision, and is even as competitive as those relying on stronger supervision.

1. Introduction

Recent development of Deep Neural Networks (DNNs) has driven the remarkable improvements in semantic segmentation [2, 3, 4, 19, 22, 25, 32, 39]. Despite the great success of DNNs, however, we still have a far way to go in achieving semantic segmentation in an uncontrolled and realistic environment. One of the main obstacles is *lack of training data*. Due to the prohibitively expensive annotation cost of pixel-level segmentation labels, existing datasets often suffer from lack of annotated examples and class diversity. This makes the conventional approaches limited to a small range of object categories predefined in the datasets.

Weakly supervised approaches have been studied to resolve the above issue and allow semantic segmentation models more scalable. Their common motivation is to utilize annotations like bounding boxes [6, 12, 28] and scribbles [18, 36] that are weaker than pixel-level labels but readily available in a large amount of visual data or easily obtainable thanks to their low annotation costs. Among various types of weak annotations for semantic segmentation, image-level class labels have been widely used [11, 14, 17, 26, 29, 30, 37] since they are already given in existing large-scale image datasets (e.g., ImageNet [7]) or automatically annotated for image retrieval results by search keywords. However, learning semantic segmentation with the image-level label supervision is a significantly ill-posed problem since such supervision indicates only the existence of a certain object class, and does not inform object location and shape that are essential for learning segmentation.

Approaches in this line of research have incorporated additional evidences to simulate the location and shape information absent in the supervision. A popular choice for the localization cue is the Class Activation Map (CAM) [40], which highlights local discriminative parts of target object by investigating the contribution of hidden units to the output of a classification DNN. The discriminative areas highlighted by CAMs are in turn used as seeds that will be propagated to cover the entire object area. To recover the object area accurately from the seeds, previous approaches have utilized image segmentation [17, 30], motions in video [35], or both [11], all of which are useful to estimate object shape. For the same purpose, a class-agnostic salient region is estimated and incorporated with the seeds in [26]. However, they demand extra data (i.e., videos) [11, 35], additional supervision (i.e., object bounding box) [26], or off-the-shelf techniques (i.e., image segmentation) that cannot take advantage of representation learning in DNNs [11, 17, 30].

In this paper, we present a simple yet effective approach to compensate the missing information for object shape with no external data or additional supervision. The key component of our framework is AffinityNet, which is a DNN that takes an image as input and predicts semantic

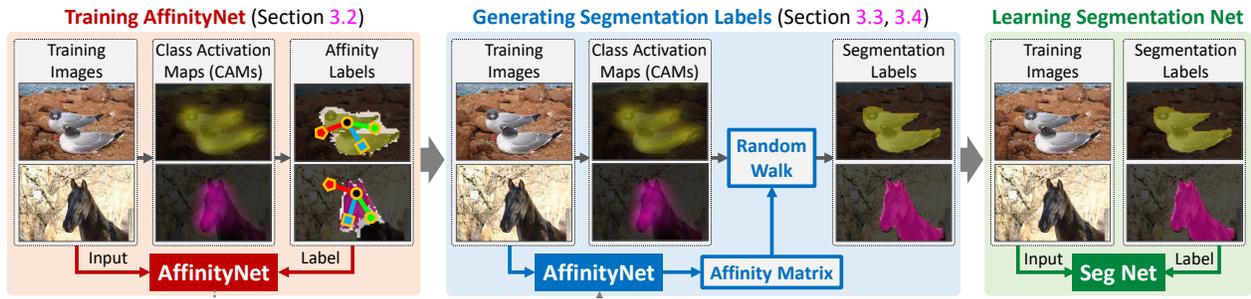


Figure 1. Illustration of our approach. Salient areas for object classes and background are first localized in training images by CAMs [40] (Section 3.1). From the salient regions, we sample pairs of adjacent coordinates and assign binary labels to them according to their class consistency. The labeled pairs are then used to train AffinityNet (Section 3.2). The trained AffinityNet in turn predicts semantic affinities within local image areas, which are incorporated with random walk to revise the CAMs (Section 3.3) and generate their segmentation labels (Section 3.4). Finally, the generated annotations are employed as supervision to train a semantic segmentation model.

affinities of pairs of adjacent image coordinates. Given an image and its CAMs, we first build a neighborhood graph where each pixel is connected to its neighbors within a certain radius, and estimate semantic affinities of pairs connected in the graph through AffinityNet. Sparse activations in CAMs are then diffused by random walk [23] on the graph, for each class: The affinities on edges in the graph encourage random walk to propagate the activations to nearby and semantically identical areas, and penalize propagation to areas of the other classes. This semantic diffusion revises CAMs significantly so that fine object shapes are recovered. We apply this process to training images for synthesizing their segmentation labels by taking the class label associated to the maximum activation of the revised CAMs at each pixel. The generated segmentation labels are used to train a segmentation model for testing.

The remaining issue is how to learn AffinityNet without extra data or additional supervision. To this end, the initial CAMs of training images are utilized as sources of supervision. Because CAMs often miss some object parts and exhibit false alarms, they are incomplete as a supervision for learning semantic segmentation whose goal is to predict the entire object masks accurately. However, we found that they are often locally correct and provide evidences to identify semantic affinities within a small image area, which are the objective of AffinityNet. To generate reliable labels of the local semantic affinities, we disregard areas with relatively low activation scores on the CAMs so that only confident object and background areas remain. A training example is then obtained by sampling a pair of adjacent image coordinates on the confident areas, and its binary label is 1 if its coordinates belong to the same class and 0 otherwise.

The overall pipeline of the proposed approach is illustrated in Figure 1. First, CAMs of training images are computed and utilized to generate semantic affinity labels, which are used as supervision to train AffinityNet. We then apply the trained AffinityNet to each training image to compute the semantic affinity matrix of its neighborhood graph,

which is employed in random walk to revise its CAMs and obtain synthesized segmentation labels. Finally, the generated segmentation labels are used to train a semantic segmentation DNN, which is the only network that will be used at test time. Our contribution is three-fold:

- We propose a novel DNN named AffinityNet that predicts high-level semantic affinities in a pixel-level, but is trained with image-level class labels only.
- Unlike most previous weakly supervised methods, our approach does not rely heavily on off-the-shelf techniques, and takes advantage of representation learning through end-to-end training of AffinityNet.
- On the PASCAL VOC 2012 [8], ours achieves state-of-the-art performance among models trained under the same level of supervision, and is competitive with those relying on stronger supervision or external data. Surprisingly, it even outperforms FCN [22], the well-known fully supervised model in the early days.

The rest of this paper is organized as follows. Section 2 reviews previous approaches closely related to ours, and Section 3 describes each step of our framework in details. Then we empirically evaluate the proposed framework on the public benchmark in Section 5, and conclude in Section 6 with brief remarks.

2. Related Work

Various types of weak supervision: Weakly supervised approaches have been extensively studied for semantic segmentation to address the data deficiency problem. Successful examples of weak supervision for semantic segmentation include bounding box [6, 12, 28], scribble [18, 36], point [1], and so on. However, these types of weak supervision still require a certain amount of human intervention during annotation procedure, so it is costly to annotate these weak labels for a large number of visual data.

Image-level labels as weak supervision: Image-level class labels have been widely used as weak supervision for semantic segmentation since they demand minimum or no human intervention to be annotated. Early approaches have tried to train a segmentation model directly from image-level labels [28, 29], but their performance is not satisfactory since the labels are too coarse to teach segmentation. To address this issue, some of previous arts incorporate segmentation seeds given by discriminative localization techniques [27, 40] with additional evidences like superpixels [11, 17, 30], segmentation proposals [30], and motions in video [11, 35], which are useful to estimate object shapes and obtained by off-the-shelf unsupervised techniques.

Our framework based on AffinityNet has a clear advantage over the above approaches. AffinityNet learns from data how to propagate local activations to entire object area, while the previous methods cannot take such an advantage. Like ours, a few methods improve segmentation quality without off-the-shelf preprocessing. Wei *et al.* [37] propose to progressively expand segmentation results by searching for new and complementary object regions sequentially. On the other hand, Kolesnikov and Lampert [14] learn a segmentation model to approximate the output of dense Conditional Random Field (dCRF) [15] applied to the segmentation seeds given by CAMs.

Learning pixel-level affinities: Our work is also closely related to the approaches that learn to predict affinity matrices in pixel-level [2, 5, 36]. Specifically, a pixel-centric affinity matrix of an image is estimated by a DNN trained with segmentation labels in [2, 5]. Bertasius *et al.* [2] incorporate the affinity matrix with random walk, whose role is to refine the output of a segmentation model like dCRF. Cheng *et al.* [5] design a deconvolution network, in which unpooling layers leverage the affinity matrix to recover sharp boundaries during upsampling. Both of the above methods aim to refine outputs of fully supervised segmentation models in a pixel-level. On the contrary, our goal is to recover object shape from coarse and noisy responses of object parts with a high-level semantic affinity matrix, and AffinityNet has a totally different architecture accordingly. Vernaza and Chandraker [36] employ scribbles as weak supervision, and propose to learn a segmentation network and the random walk affinity matrix simultaneously so that the output of the network and that of random walk propagation of the scribbles become identical. Our approach is different from this work in the following three aspects. First, our framework is trained with image-level labels, which are significantly weaker than the scribbles employed in [36]. Second, in our approach random walk can jump to any other positions within a certain radius, but in [36] it is allowed to move only to four nearest neighbors. Third, AffinityNet learns pairwise semantic affinity explicitly, but the model in [36] learns it implicitly.

Learning with synthetic labels: We adopt the disjoint pipeline that first generates synthetic labels and train a segmentation model with the labels in a fully supervised manner. Such a pipeline has been studied for object detection [34] as well as semantic segmentation [6, 11, 12, 17, 26, 35] in weakly supervised settings. A unique feature of our approach is the AffinityNet, an end-to-end trainable DNN improving the quality of synthetic labels significantly, when compared to previous methods that adopt existing optimization techniques (*e.g.*, GraphCut, GrabCut, and dCRF) and/or off-the-shelf preprocessing steps aforementioned.

3. Our Framework

Our approach to weakly supervised semantic segmentation is roughly divided into two parts: (1) Synthesizing pixel-level segmentation labels of training images given their image-level class labels, and (2) Learning a DNN for semantic segmentation with the generated segmentation labels. The entire framework is based on three DNNs: A network computing CAMs, AffinityNet, and a segmentation model. The first two are used to generate segmentation labels of training images, and the last one is the DNN that performs actual semantic segmentation and is trained with the synthesized segmentation annotations. The remainder of this section describes characteristics of the three networks and training schemes for them in details.

3.1. Computing CAMs

CAMs play an important role in our framework. As in many other weakly supervised approaches, they are considered as segmentation seeds, which typically highlight local salient parts of object and later propagate to cover the entire object area. Furthermore, in our framework they are employed as sources of supervision for training AffinityNet.

We follow the approach of [40] to compute CAMs of training images. The architecture is a typical classification network with global average pooling (GAP) followed by a fully connected layer, and is trained by a classification criteria with image-level labels. Given the trained network, the CAM of a groundtruth class c , which is denoted by M_c , is computed by

$$M_c(x, y) = \mathbf{w}_c^\top f^{\text{cam}}(x, y), \quad (1)$$

where \mathbf{w}_c is the classification weights associated to the class c and $f^{\text{cam}}(x, y)$ indicates the feature vector located at (x, y) on the feature map before the GAP. M_c is further normalized so that the maximum activation equals 1: $M_c(x, y) \rightarrow M_c(x, y) / \max_{x, y} M_c(x, y)$. For any class c' irrelevant to the groundtruths, we disregard $M_{c'}$ by making its activation scores zero. We also estimate a background activation map, which is given by

$$M_{\text{bg}}(x, y) = \left\{ 1 - \max_{c \in C} M_c(x, y) \right\}^\alpha \quad (2)$$

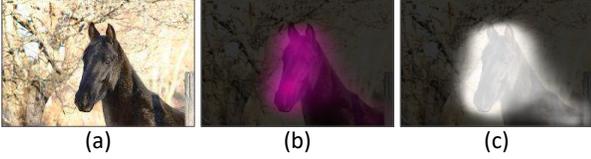


Figure 2. Visualization of CAMs obtained by our approach. (a) Input image. (b) CAMs of object classes: Brighter means more confident object region. (c) CAMs of background: Darker means more confident background region.

where C is the set of object classes and $\alpha \geq 1$ denotes a hyper-parameter that adjusts the background confidence scores. Qualitative examples of CAMs obtained by our approach are visualized in Figure 2.

3.2. Learning AffinityNet

AffinityNet aims to predict class-agnostic semantic affinity between a pair of adjacent coordinates on a training image. The predicted affinities are used in random walk as transition probabilities so that random walk propagates activation scores of CAMs to nearby areas of the same semantic entity, which improves the quality of CAMs significantly.

For computational efficiency, AffinityNet is designed to predict a convolutional feature map f^{aff} where the semantic affinity between a pair of feature vectors is defined in terms of their L_1 distance. Specifically, the semantic affinity between features i and j is denoted by W_{ij} and defined as

$$W_{ij} = \exp \left\{ - \left\| f^{\text{aff}}(x_i, y_i) - f^{\text{aff}}(x_j, y_j) \right\|_1 \right\}, \quad (3)$$

where (x_i, y_i) indicates the coordinate of the i^{th} feature on the feature map f^{aff} . In this way, a large number of semantic affinities in a given image can be computed efficiently by a single forward pass of the network. Figure 3 illustrates the AffinityNet architecture and the way it computes f^{aff} . Training this architecture requires semantic affinity labels for pairs of feature map coordinates, *i.e.*, labels for W_{ij} in Eq. (3). However, such labels are not directly available in our setting where only image-level labels are given. In the remaining part of this section, we present how to generate the affinity labels and train AffinityNet with them.

3.2.1 Generating Semantic Affinity Labels

To train AffinityNet with image-level labels, we exploit CAMs of training images as incomplete sources of supervision. Although CAMs are often inaccurate as shown in Figure 2, we found that by carefully manipulating them, reliable supervision for semantic affinity can be obtained.

Our basic idea is to identify confident areas of objects and background from the CAMs, and sample training examples only from these areas. By doing so, the semantic equivalence between a pair of sampled coordinates can be

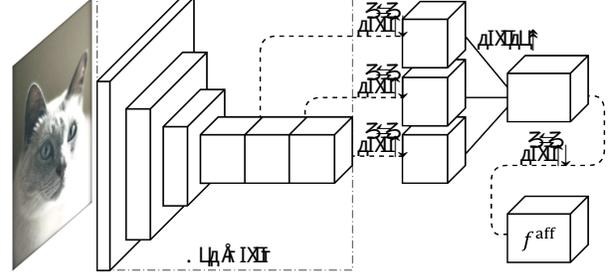


Figure 3. Overall architecture of AffinityNet. The output feature map f^{aff} is obtained by aggregating feature maps from multiple levels of a backbone network so that f^{aff} can take semantic information at various field-of-views. Specifically, we first apply 1×1 convolutions to the multi-level feature maps for dimensionality reduction, concatenate the results as a single feature map, and employ one more 1×1 convolution for adaptation to the target task. More details of the architecture is described in Section 4.

determined reliably. To estimate confident areas of objects, we first amplify M_{bg} by decreasing α in Eq. (2) so that background scores dominate insignificant activation scores of objects in the CAMs. After applying dCRF to the CAMs for their refinement, we identify confident areas for each object class by collecting coordinates whose scores for the target class are greater than those of any other classes including the amplified background. Also, in the opposite setting (*i.e.*, increasing α to weaken M_{bg}), confident background areas can be identified in the same manner. Remaining areas in the image are then considered as *neutral*. A result of this procedure is visualized in Figure 4(a).

Now a binary affinity label can be assigned to every pair of coordinates according to their class labels determined by the confident areas. For two coordinates (x_i, y_i) and (x_j, y_j) that are not *neutral*, their affinity label W_{ij}^* is 1 if their classes are the same, and 0 otherwise. Also, if at least one of the coordinates is *neutral*, we simply ignore the pair during training. This scheme, which is illustrated in Figure 4(b), enables us to collect a fairly large number of pairwise affinity labels which are also reliable enough.

3.2.2 AffinityNet Training

AffinityNet is trained by approximating the binary affinity labels W_{ij}^* with the predicted semantic affinities W_{ij} of Eq. (3) in a gradient descent fashion. Especially, affinities of only sufficiently adjacent coordinates are considered during training due to the following two reasons. First, it is difficult to predict semantic affinity between two coordinates too far from each other due to the lack of context. Second, by addressing pairs of adjacent coordinates only, we can reduce computational cost significantly. Thus the set of coordinate pairs used in training, denoted by \mathcal{P} , is given by

$$\mathcal{P} = \{(i, j) \mid d((x_i, y_i), (x_j, y_j)) < \gamma, \forall i \neq j\} \quad (4)$$

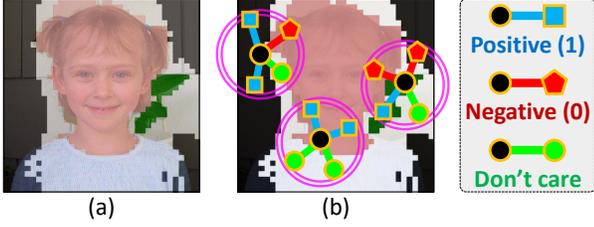


Figure 4. Conceptual illustration of generating semantic affinity labels. (a) Confident areas of object classes and background: peach for *person*, green for *plant*, and black for *background*. The *neutral* area is color-coded in white. (b) Coordinate pairs sampled within a small radius for training AffinityNet. Each pair is assigned label 1 if its two coordinates come from the same class, and label 0 otherwise. When at least one of the two coordinates belongs to the *neutral* area, the pair is ignored during training.

where $d(\cdot, \cdot)$ is the Euclidean distance and γ is a search radius that limits the distance between a selected pair.

However, learning AffinityNet directly from \mathcal{P} is not desirable due to the class imbalance issue. We observed that in \mathcal{P} the class distribution is significantly biased to positive ones since negative pairs are sampled only around object boundaries. Also in the subset of positive pairs, the number of background pairs is notably larger than that of object pairs as background is larger than object areas in many photos. To address this issue, we divide \mathcal{P} into three subsets, and aggregate losses obtained from individual subsets. Specifically, we first divide \mathcal{P} into two subsets of positive and negative pairs:

$$\mathcal{P}^+ = \{(i, j) \mid (i, j) \in \mathcal{P}, W_{ij}^* = 1\}, \quad (5)$$

$$\mathcal{P}^- = \{(i, j) \mid (i, j) \in \mathcal{P}, W_{ij}^* = 0\}, \quad (6)$$

and further break \mathcal{P}^+ into $\mathcal{P}_{\text{fg}}^+$ and $\mathcal{P}_{\text{bg}}^+$ for objects and background, respectively. Then the cross-entropy loss is computed per subset as follows:

$$\mathcal{L}_{\text{fg}}^+ = -\frac{1}{|\mathcal{P}_{\text{fg}}^+|} \sum_{(i,j) \in \mathcal{P}_{\text{fg}}^+} \log W_{ij}, \quad (7)$$

$$\mathcal{L}_{\text{bg}}^+ = -\frac{1}{|\mathcal{P}_{\text{bg}}^+|} \sum_{(i,j) \in \mathcal{P}_{\text{bg}}^+} \log W_{ij}, \quad (8)$$

$$\mathcal{L}^- = -\frac{1}{|\mathcal{P}^-|} \sum_{(i,j) \in \mathcal{P}^-} \log(1 - W_{ij}). \quad (9)$$

Finally, the loss for training the AffinityNet is defined as

$$\mathcal{L} = \mathcal{L}_{\text{fg}}^+ + \mathcal{L}_{\text{bg}}^+ + 2\mathcal{L}^-. \quad (10)$$

Note that the loss in Eq. (10) is class-agnostic. Thus the trained AffinityNet decides the class consistency between two adjacent coordinates while not aware of their classes explicitly. This class-agnostic scheme allows AffinityNet

to learn a more general representation that can be shared among multiple object classes and background, and enlarges the set of training samples per class significantly.

3.3. Revising CAMs Using AffinityNet

The trained AffinityNet is used to revise CAMs of training images. Local semantic affinities predicted by AffinityNet are converted to a transition probability matrix, which enables random walk to be aware of semantic boundaries in image, and encourages it to diffuse activation scores within those boundaries. We empirically found that random walk with the semantic transition matrix significantly improves the quality of CAMs and allows us to generate accurate segmentation labels consequently.

For an input image, AffinityNet generates a convolutional feature map, and semantic affinities between features on the map are computed according to Eq. (3). Note that, as in training of AffinityNet, the affinities are computed between features within local circles of radius γ . The computed affinities form an affinity matrix W , whose diagonal elements are 1. The transition probability matrix T of random walk is derived from the affinity matrix as follows:

$$T = D^{-1}W^{\circ\beta}, \quad \text{where } D_{ii} = \sum_j W_{ij}^{\circ\beta}. \quad (11)$$

In the above equation, the hyper-parameter β has a value greater than 1 so that $W^{\circ\beta}$, the Hadamard power of the original affinity matrix, ignores immaterial affinities in W . Thus using $W^{\circ\beta}$ instead of W makes our random walk propagation more conservative. The diagonal matrix D is computed for row-wise normalization of $W^{\circ\beta}$.

Through random walk with T , a single operation of the semantic propagation is implemented by multiplying T to the CAMs. We perform this propagation iteratively until the predefined number of iterations is reached. Then M_c^* , the revised CAM of class c is given by

$$\text{vec}(M_c^*) = T^t \cdot \text{vec}(M_c) \quad \forall c \in C \cup \{\text{bg}\}, \quad (12)$$

where $\text{vec}(\cdot)$ means vectorization of a matrix, and t is the number of iterations. Note that the value of t is set to a power of 2 so that Eq. (12) performs matrix multiplication only $\log_2 t + 1$ times.

3.4. Learning a Semantic Segmentation Network

The revised CAMs of training images are then used to generate segmentation labels of the images. Since CAMs are smaller than their input image in size, we upsample them to the resolution of the image by bilinear interpolation, and refine them with dCRF. A segmentation label of a training image is then obtained simply by selecting the class label associated with the maximum activation score at every pixel in the revised and upsampled CAMs. Note that

the *background* class can be also selected since we compute CAMs for background as well as object classes.

The segmentation labels obtained by the above procedure are used as supervision to train a segmentation network. Any fully supervised semantic segmentation model can be employed in our approach as we provide segmentation labels of training images.

4. Network Architectures

In this section, we present details of DNN architectures adopted in our framework. Note that our approach can be implemented with any existing DNNs serving the same purposes, although we carefully design the following models to enhance segmentation performance.

4.1. Backbone Network

The three DNNs in our framework are all built upon the same backbone network. The backbone is a modified version of Model A1 [38], which is also known as ResNet38 and has 38 convolution layers with wide channels. To obtain the backbone network, the ultimate GAP and fully connected layer of the original model are first removed. Then the convolution layers of the last three levels¹ are replaced by atrous convolutions with a common input stride of 1, and their dilation rates are adjusted so that the backbone network will return a feature map of stride 8. The atrous convolution has been known to enhance segmentation quality by enlarging receptive field without sacrificing feature map resolution [4]. We empirically observed that it works also in our weakly supervised models, CAM and AffinityNet, as it enables the models to recover fine shapes of objects.

4.2. Details of DNNs in Our Framework

Network computing CAMs: We obtain this model by adding the following three layers on the top of the backbone network in the order: a 3×3 convolution layer with 512 channels for a better adaptation to the target task, a global average pooling layer for feature map aggregation, and a fully connected layer for classification.

AffinityNet: This network is designed to aggregate multi-level feature maps of the backbone network in order to leverage semantic information acquired at various field-of-views when computing affinities. For the purpose, the feature maps output from the last three levels of the backbone network are selected. Before aggregation, their channel dimensionalities are reduced to 128, 256, and 512, for the first, second, and third feature maps, respectively, by individual 1×1 convolution layers. Then the feature maps are concatenated to be a single feature map with 896 channels. We finally add one more 1×1 convolution layer with 896 channels on the top for adaptation.

¹A level is a group of residual units that share the same output stride.

Kwak <i>et al.</i> [17]		Ours		
CAM	SPN	CAM	CAM+RW	CAM+RW+dCRF
30.5	43.8	48.0	58.1	59.7

Table 1. Accuracy of synthesized segmentation labels in mIoU, evaluated on the PASCAL VOC 2012 *training* set. SPN: Superpixel Pooling Net of [17], RW: random walk with AffinityNet.

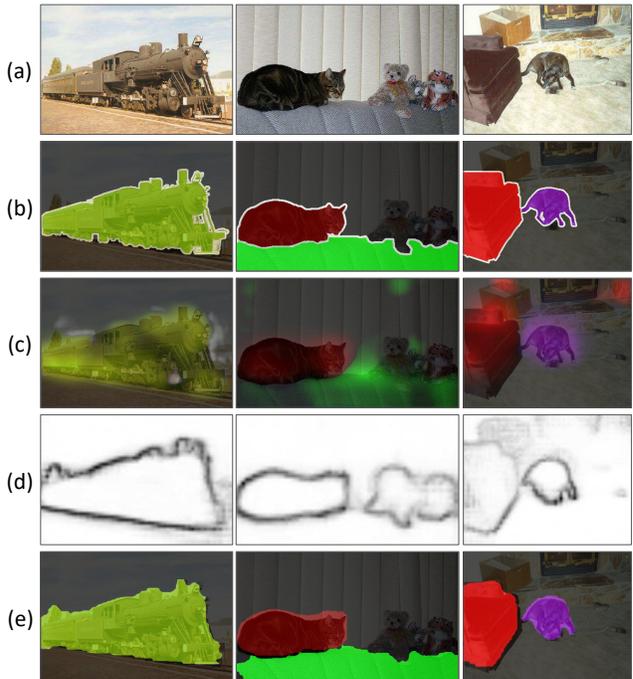


Figure 5. Qualitative examples of synthesized segmentation labels of training images in the PASCAL VOC 2012 benchmark. (a) Input images. (b) Groundtruth segmentation labels. (c) CAMs of object classes. (d) Visualization of the predicted semantic affinities. (e) Synthesized segmentation annotations.

Segmentation model: We strictly follow [38] to build our segmentation network. Specifically, we put two more atrous convolution layers on the top of the backbone. They have the same dilation rate of 12, while the numbers of channels are 512 for the first one and 21 for the second. The resulting network is called “Ours-ResNet38” in the next section.

5. Experiments

This section demonstrates the effectiveness of our approach with comparisons to current state of the art in weakly supervised semantic segmentation on the PASCAL VOC 2012 segmentation benchmark [8]. For a performance metric, we adopt Intersection-over-Union (IoU) between groundtruth and predicted segmentation.

5.1. Implementation Details

Dataset: All DNNs in our framework are trained and evaluated on the PASCAL VOC 2012 segmentation benchmark,

Method	bkg	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbk	person	plant	sheep	sofa	train	tv	mean
EM-Adapt [28]	67.2	29.2	17.6	28.6	22.2	29.6	47.0	44.0	44.2	14.6	35.1	24.9	41.0	34.8	41.6	32.1	24.8	37.4	24.0	38.1	31.6	33.8
CCNN [29]	68.5	25.5	18.0	25.4	20.2	36.3	46.8	47.1	48.0	15.8	37.9	21.0	44.5	34.5	46.2	40.7	30.4	36.3	22.2	38.8	36.9	35.3
MIL+seg [30]	79.6	50.2	21.6	40.9	34.9	40.5	45.9	51.5	60.6	12.6	51.2	11.6	56.8	52.9	44.8	42.7	31.2	55.4	21.5	38.8	36.9	42.0
SEC [14]	82.4	62.9	26.4	61.6	27.6	38.1	66.6	62.7	75.2	22.1	53.5	28.3	65.8	57.8	62.3	52.5	32.5	62.6	32.1	45.4	45.3	50.7
AdvErasing [37]	83.4	71.1	30.5	72.9	41.6	55.9	63.1	60.2	74.0	18.0	66.5	32.4	71.7	56.3	64.8	52.4	37.4	69.1	31.4	58.9	43.9	55.0
Ours-DeepLab	87.2	57.4	25.6	69.8	45.7	53.3	76.6	70.4	74.1	28.3	63.2	44.8	75.6	66.1	65.1	71.1	40.5	66.7	37.2	58.4	49.1	58.4
Ours-ResNet38	88.2	68.2	30.6	81.1	49.6	61.0	77.8	66.1	75.1	29.0	66.0	40.2	80.4	62.0	70.4	73.7	42.5	70.7	42.6	68.1	51.6	61.7

Table 2. Performance on the PASCAL VOC 2012 *val* set, compared to weakly supervised approaches based only on image-level labels.

Method	bkg	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbk	person	plant	sheep	sofa	train	tv	mean
EM-Adapt [28]	76.3	37.1	21.9	41.6	26.1	38.5	50.8	44.9	48.9	16.7	40.8	29.4	47.1	45.8	54.8	28.2	30.0	44.0	29.2	34.3	46.0	39.6
CCNN [29]	70.1	24.2	19.9	26.3	18.6	38.1	51.7	42.9	48.2	15.6	37.2	18.3	43.0	38.2	52.2	40.0	33.8	36.0	21.6	33.4	38.3	35.6
MIL+seg [30]	78.7	48.0	21.2	31.1	28.4	35.1	51.4	55.5	52.8	7.8	56.2	19.9	53.8	50.3	40.0	38.6	27.8	51.8	24.7	33.3	46.3	40.6
SEC [14]	83.5	56.4	28.5	64.1	23.6	46.5	70.6	58.5	71.3	23.2	54.0	28.0	68.1	62.1	70.0	55.0	38.4	58.0	39.9	38.4	48.3	51.7
AdvErasing [37]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	55.7
Ours-DeepLab	88.0	61.1	29.2	73.0	40.5	54.1	75.2	70.4	75.1	27.8	62.5	51.4	78.4	68.3	76.2	71.8	40.7	74.9	49.2	55.0	48.3	60.5
Ours-ResNet38	89.1	70.6	31.6	77.2	42.2	68.9	79.1	66.5	74.9	29.6	68.7	56.1	82.1	64.8	78.6	73.5	50.8	70.7	47.7	63.9	51.1	63.7

Table 3. Performance on the PASCAL VOC 2012 *test* set, compared to weakly supervised approaches based only on image-level labels.

Method	Sup.	Extra Data	<i>val</i>	<i>test</i>
TransferNet [10]	\mathcal{I}	MS-COCO [20]	52.1	51.2
Saliency [26]	\mathcal{I}	MSRA [21], BSDS [24]	55.7	56.7
MCNN [35]	\mathcal{I}	YouTube-Object [31]	38.1	39.8
CrawlSeg [11]	\mathcal{I}	YouTube Videos	58.1	58.7
What’sPoint [1]	\mathcal{P}	-	46.0	43.6
RAWK [36]	\mathcal{S}	-	61.4	-
ScribbleSup [18]	\mathcal{S}	-	63.1	-
WSSL [28]	\mathcal{B}	-	60.6	62.2
BoxSup [6]	\mathcal{B}	-	62.0	64.6
SDI [12]	\mathcal{B}	BSDS [24]	65.7	67.5
FCN [22]	\mathcal{F}	-	-	62.2
DeepLab [3]	\mathcal{F}	-	67.6	70.3
ResNet38 [38]	\mathcal{F}	-	80.8	82.5
Ours-DeepLab	\mathcal{I}	-	58.4	60.5
Ours-ResNet38	\mathcal{I}	-	61.7	63.7

Table 4. Performance on the PASCAL VOC 2012 *val* and *test* sets. The supervision types (Sup.) indicate: \mathcal{P} –point, \mathcal{S} –scribble, \mathcal{B} –bounding box, \mathcal{I} –image-level label, and \mathcal{F} –segmentation label.

for a fair comparison to previous approaches. Following the common practice, we enlarge the set of training images by adopting segmentation annotations presented in [9]. Consequently 10,582 images in total are used as training examples and 1,449 images are kept for validation.

Network parameter optimization: The backbone network of our DNNs is pretrained on the ImageNet [7]. The entire network parameters are then finetuned on the PASCAL VOC 2012 by Adam [13]. The following data augmentation techniques are commonly used when training all the three DNNs: horizontal flip, random cropping, and color jittering [16]. Also, for the networks except AffinityNet, we randomly scale input images during training, which is useful to impose scale invariance on the networks.

Parameter setting: α in Eq. (2) is 16 by default, and changed to 4 and 24 to amplify and weaken background activations, respectively. We set γ in Eq. (4) to 5 and β in Eq. (11) to 8. Also, t in Eq. (12) is fixed by 256. For dCRF, we used the default parameters given in the original code.

5.2. Analysis of Synthesized Segmentation Labels

The performance of our label synthesis method is measured in mIoU between groundtruth and generated segmentation labels as summarized in Table 1. For ablation study, our method is divided into three parts: CAM, RW (random walk with AffinityNet), and dCRF. To demonstrate the advantage of the proposed method, we also report the score of Superpixel Pooling Net (SPN) [17] that incorporates CAM with superpixels as additional clues to generate segmentation labels with image-level label supervision. As shown in Table 1, even our CAM outperforms SPN in terms of the quality of generated segmentation labels without using off-the-shelf techniques like superpixel. We believe this is because of the various data augmentation techniques and the more powerful backbone network with atrous convolution layers. Moreover, through random walk with the learned semantic affinity, the quality of segmentation annotations is improved remarkably, which demonstrates the effectiveness of AffinityNet. Finally, dCRF further improves the label quality slightly, and we employ this last version as the supervision for learning the segmentation network.

Examples of the synthesized segmentation labels are shown in Figure 5, where one can see that random walk with AffinityNet handles false positives and missing areas in CAMs effectively. To illustrate the role of AffinityNet in this process, we also visualized predicted semantic affinities of the images by detecting edges on the feature map f^{aff} , and observed that AffinityNet has capability to detect semantic boundaries although it is trained with image-level labels. As such boundaries penalize random walk propagation between semantically different objects, the synthesized segmentation labels can recover accurate shapes of objects.

5.3. Comparisons to Previous Work

We first quantitatively compare our approach with previous methods based only on image-level class labels. The

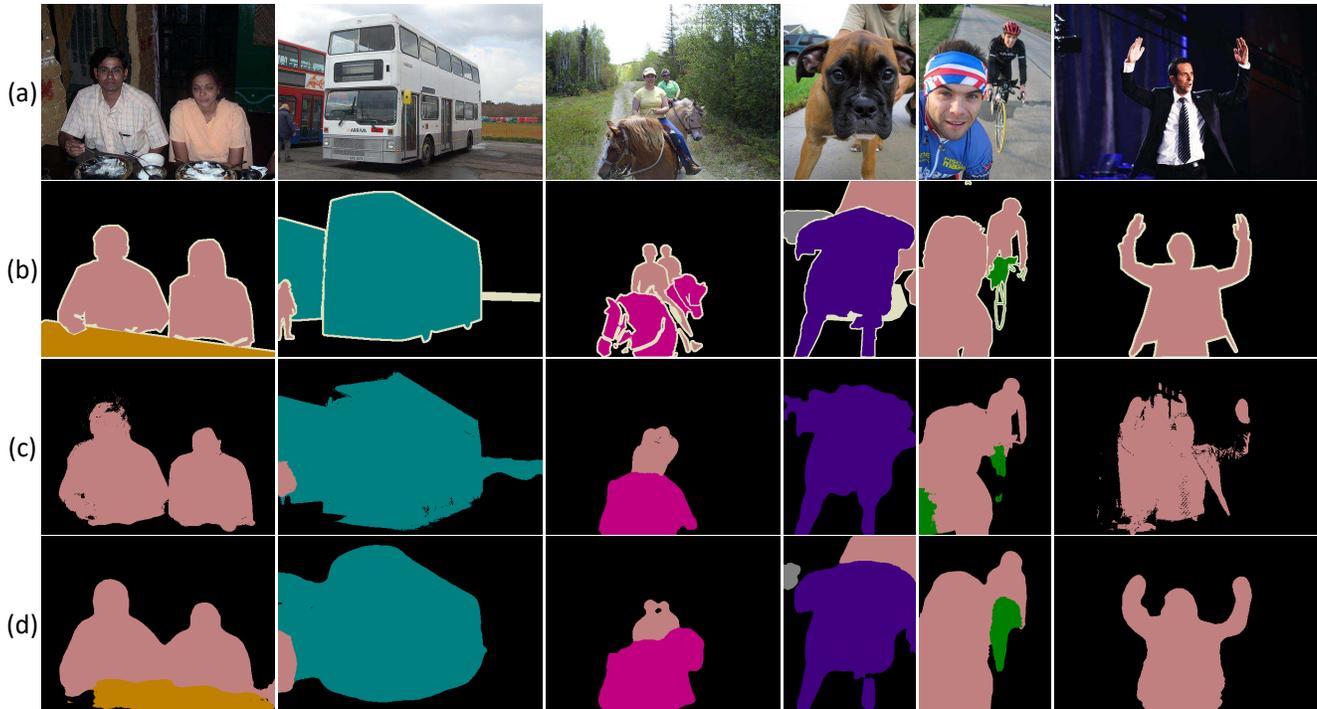


Figure 6. Qualitative results on the PASCAL VOC 2012 *val* set. (a) Input images. (b) Groundtruth segmentation. (c) Results obtained by CrawlSeg [11]. (d) Results of Ours-ResNet38. Compared to CrawlSeg, which is the current state-of-the-art model based on image-level label supervision, our method better captures larger object areas and less prone to miss objects. The object boundaries of our results are smoother than those of CrawlSeg as we do not apply dCRF to the final results. More results can be found in the supplementary material.

results on the PASCAL VOC 2012 are summarized in Table 2 and 3. Note that we also evaluate DeepLab [4] trained with our synthetic labels, called “Ours-DeepLab”, for fair comparisons to other models whose backbones are VGG16 [33]. Both of our models outperform the current state of the art [37] by large margins in terms of mean accuracy on both *val* and *test* sets of the benchmark, while Ours-ResNet38 is slightly better than Ours-DeepLab thanks to the more powerful representation of ResNet38. Our models are also compared to the approaches based on extra training data or stronger supervision in Table 4. They substantially outperform the approaches based on the same level of supervision but with extra data and annotations like segmentation labels in MS-COCO [20], class-agnostic bounding boxes in MSRA Saliency [21], and YouTube videos [31]. They are also competitive with previous arts relying on stronger supervision like scribble and bounding box. Surprisingly, Ours-ResNet38 outperforms even FCN [22], the well-known early work on fully supervised semantic segmentation. These results show that segmentation labels generated by our method are sufficiently strong, and can substitute for extra data or stronger supervision. We finally compare our models with their fully supervised versions, DeepLab [4] and ResNet38 [38], which are the upper

bounds we can achieve. Specifically, Ours-DeepLab recovers 86% of its bound, and Ours-ResNet38 achieves 77%.

Figure 6 presents qualitative results of Ours-ResNet38 and compares them to those of CrawlSeg [11], which is the current state of the art using image-level supervision. Our method relying only on image-level label supervision tends to produce more accurate results even though CrawlSeg exploits extra video data to synthesize segmentation labels.

6. Conclusion

To alleviate the lack of annotated data issue in semantic segmentation, we have proposed a novel framework based on AffinityNet to generate accurate segmentation labels of training images given their image-level class labels only. The effectiveness of our approach has been demonstrated on the PASCAL VOC 2012 benchmark, where DNNs trained with the labels generated by our method substantially outperform the previous state of the art relying on the same level of supervision, and is competitive with those demanding stronger supervision or extra data.

Acknowledgement: This work was supported by Kakao, Korea Creative Content Agency (KOCCA), and Ministry of Culture, Sports, and Tourism (MCST) of Korea.

References

- [1] A. Bearman, O. Russakovsky, V. Ferrari, and L. Fei-Fei. What's the Point: Semantic Segmentation with Point Supervision. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 549–565, 2016.
- [2] G. Bertasius, L. Torresani, S. X. Yu, and J. Shi. Convolutional random walk networks for semantic image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [3] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected CRFs. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [4] L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, PP(99):1–1, 2017.
- [5] Y. Cheng, R. Cai, Z. Li, X. Zhao, and K. Huang. Locality-sensitive deconvolution networks with gated fusion for rgb-d indoor semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [6] J. Dai, K. He, and J. Sun. BoxSup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1635–1643, 2015.
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: a large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [8] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision (IJCV)*, 88(2):303–338, 2010.
- [9] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik. Semantic contours from inverse detectors. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 991–998, 2011.
- [10] S. Hong, J. Oh, B. Han, and H. Lee. Learning transferrable knowledge for semantic segmentation with deep convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3204 – 3212, 2016.
- [11] S. Hong, D. Yeo, S. Kwak, H. Lee, and B. Han. Weakly supervised semantic segmentation using web-crawled videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7322–7330, 2017.
- [12] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 876–885, 2017.
- [13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [14] A. Kolesnikov and C. H. Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 695–711, 2016.
- [15] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Proceedings of the Neural Information Processing Systems (NIPS)*, pages 109–117. 2011.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Proceedings of the Neural Information Processing Systems (NIPS)*, 2012.
- [17] S. Kwak, S. Hong, and B. Han. Weakly supervised semantic segmentation using superpixel pooling network. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 4111–4117, 2017.
- [18] D. Lin, J. Dai, J. Jia, K. He, and J. Sun. Scribble-sup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3159–3167, 2016.
- [19] G. Lin, C. Shen, A. van den Hengel, and I. Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3194 – 3203, 2016.
- [20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 740–755, 2014.
- [21] T. Liu, J. Sun, N. N. Zheng, X. Tang, and H. Y. Shum. Learning to detect a salient object. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, June 2007.
- [22] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431 – 3440, 2015.
- [23] L. Lovasz. Random walks on graphs: A survey, 1993.
- [24] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.
- [25] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1520 – 1528, 2015.
- [26] S. J. Oh, R. Benenson, A. Khoreva, Z. Akata, M. Fritz, and B. Schiele. Exploiting saliency for object segmentation from image level labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4410–4419, 2017.
- [27] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [28] G. Papandreou, L.-C. Chen, K. Murphy, and A. L. Yuille. Weakly-and semi-supervised learning of a DCNN for seman-

- tic image segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1742 – 1750, 2015.
- [29] D. Pathak, P. Krähenbühl, and T. Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1742 – 1750, 2015.
- [30] P. O. Pinheiro and R. Collobert. From image-level to pixel-level labeling with convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1713 – 1721, 2015.
- [31] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3282 – 3289, 2012.
- [32] G.-J. Qi. Hierarchically gated deep networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.
- [34] P. Tang, X. Wang, X. Bai, and W. Liu. Multiple instance detection network with online instance classifier refinement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3059–3067, July 2017.
- [35] P. Tokmakov, K. Alahari, and C. Schmid. Weakly-supervised semantic segmentation using motion cues. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 388–404, 2016.
- [36] P. Vernaza and M. Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [37] Y. Wei, J. Feng, X. Liang, M.-M. Cheng, Y. Zhao, and S. Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [38] Z. Wu, C. Shen, and A. van den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *arXiv preprint arXiv:1611.10080*, 2016.
- [39] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1529 – 1537, 2015.
- [40] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2921 – 2929, 2016.