

Weakly Supervised Semantic Segmentation using Web-Crawled Videos

Seunghoon Hong[†] Donghun Yeo[†] Suha Kwak[‡] Honglak Lee[§] Bohyung Han[†]
[†]POSTECH [‡]DGIST [§]University of Michigan
Pohang, Korea Daegu, Korea Ann Arbor, MI, USA

{maga33, hanulbog, bhhan}@postech.ac.kr skwak@dgist.ac.kr honglak@umich.edu

Abstract

We propose a novel algorithm for weakly supervised semantic segmentation based on image-level class labels only. In weakly supervised setting, it is commonly observed that trained model overly focuses on discriminative parts rather than the entire object area. Our goal is to overcome this limitation with no additional human intervention by retrieving videos relevant to target class labels from web repository, and generating segmentation labels from the retrieved videos to simulate strong supervision for semantic segmentation. During this process, we take advantage of image classification with discriminative localization technique to reject false alarms in retrieved videos and identify relevant spatio-temporal volumes within retrieved videos. Although the entire procedure does not require any additional supervision, the segmentation annotations obtained from videos are sufficiently strong to learn a model for semantic segmentation. The proposed algorithm substantially outperforms existing methods based on the same level of supervision and is even as competitive as the approaches relying on extra annotations.

1. Introduction

Semantic segmentation has recently achieved prominent progress thanks to Deep Convolutional Neural Networks (DCNNs) [3, 21, 24, 32, 37, 41]. The success of DCNNs heavily depends on the availability of a large-scale training dataset, where annotations are given manually in general. In semantic segmentation, however, annotations are in the form of pixel-wise masks, and collecting such annotations for a large number of images demands tremendous effort and cost. Consequently, accurate and reliable segmentation annotations are available only for a small number of classes. Fully supervised DCNNs for semantic segmentation are thus limited to those classes and hard to be extended to many other classes appearing in real world images.

Weakly supervised approaches have been proposed to alleviate this issue by leveraging a vast amount of weakly annotated images. Among several types of weak supervision for semantic segmentation, image-level class label has been

widely used [17, 26, 28, 29, 30] as it is readily available from existing image databases [7, 10]. The most popular approach to generating pixel-wise labels from an image-level label is self-supervised learning based on the joint estimation of segmentation annotation and model parameters [6, 20, 29, 30]. However, since there is no way to measure the quality of estimated annotations, these approaches easily converge to suboptimal solutions. To remedy this limitation, other types of weak supervision have been employed in addition to image-level labels, *e.g.*, bounding box [6, 26], scribble [20], prior meta-information [28], and segmentation ground-truths of other classes [13]. However, they often require additional human intervention to obtain extra supervision [6, 13, 26] or employ domain-specific knowledge that may not be well-generalized to other classes [28].

The objective of this work is to overcome the inherent limitation in weakly supervised semantic segmentation without additional human supervision. Specifically, we propose to retrieve *videos* from the Web and use them as an additional source of training data, since temporal dynamics in video offers rich information to distinguish objects from background and estimate their shapes more accurately. More importantly, our video retrieval process is performed fully-automatically by using a set of class labels as search keywords and collecting videos from web repositories (*e.g.*, YouTube). The result of retrieval is a collection of weakly annotated videos as each video is given its query keyword as video-level class label. However, it is still not straightforward to learn semantic segmentation directly from weakly labeled videos due to ambiguous association between labels and frames. The association is temporally ambiguous since only a subset of frames in a video is relevant to its class label. Furthermore, although there are multiple regions exhibiting prominent motions, only a few among them might be relevant to the class label, which causes spatial ambiguity. These ambiguities are ubiquitous in videos crawled automatically with no human intervention.

The key idea of this paper is to utilize both weakly annotated images and videos to learn a single DCNN for semantic segmentation. Images are associated with clean class

labels given manually, thus they can be used to alleviate the ambiguities in web-crawled videos. Also, it is easier to estimate shape and extent of object in videos thanks to motion cues available exclusively in them. To exploit these complementary benefits of the two domains, we integrate techniques for discriminative object localization in images [42] and video segmentation [27] into a single framework based on DCNN, which generates reliable segmentation annotations from videos and learns semantic segmentation for image with the generated annotations.

The architecture of our DCNN is motivated by [13] and consists of two parts, each of which has its own role: an encoder for image classification and discriminative localization [42], and a decoder for image segmentation. The two parts of the network are trained separately with different data in our framework. The encoder is first learned from a set of weakly annotated images. It is in turn used to filter out irrelevant frames and identify discriminative regions in weakly annotated videos so that both temporal and spatial ambiguities of the videos are substantially reduced. By incorporating the identified discriminative regions together with color and motion cues, spatio-temporal segments of object candidates are obtained from the videos by a well-established graph-based optimization technique. The video segmentation results are then used as segmentation annotations to train the decoder of our network.

The contributions of this paper are three-fold as follows.

- We propose a weakly supervised semantic segmentation algorithm based on web-crawled videos. Our algorithm exploits videos to simulate strong supervision missing in weakly annotated images, and utilizes images to eliminate noises in video retrieval and segmentation processes.
- Our framework automatically collects video clips relevant to the target classes from web repositories so that it does not require human intervention to obtain extra supervision.
- We demonstrate the effectiveness of the proposed framework on the PASCAL VOC benchmark dataset, where it outperforms prior arts on weakly supervised semantic segmentation by a substantial margin.

The rest of the paper is organized as follows. We briefly review related work in Section 2 and describe the details of the proposed framework in Section 3. Section 4 introduces data collection process. Section 5 illustrates experimental results on benchmark datasets.

2. Related Work

Semantic segmentation has been rapidly improved in past few years, mainly due to emergence of powerful end-to-end learning framework based on DCNNs [3, 11, 21,

23, 24, 25, 41]. Built upon a fully-convolutional architecture [24], various approaches have been investigated to improve segmentation accuracy by integrating fully-connected CRF [3, 21, 23, 41], deep deconvolution network [25], multi-scale processing [3, 11], etc. However, training a model based on DCNN requires pixel-wise annotations, which involves expensive and time-consuming procedures to obtain. For this reasons, the task has been mainly investigated in small-scale datasets [10, 22].

Approaches based on weakly supervised learning have been proposed to reduce annotation efforts in fully-supervised methods [6, 13, 17, 26, 28, 29, 30]. Among many possible choices, image-level labels are of the form requiring the minimum annotation cost thus have been widely used [17, 28, 29, 30]. Unfortunately, their results are far behind the fully-supervised methods due to missing supervision on segmentation. This gap is reduced by exploiting additional annotations such as point supervision [2], scribble [20], bounding box [6, 26], masks from other class [13], but they lead to increased annotation cost that should be avoided in weakly supervised setting. Instead of collecting extra cues from human annotator, we propose to retrieve and exploit web-videos, which offers motion cue useful for segmentation without the need of any human intervention in collecting such data. The idea of employing videos for semantic segmentation is new and has not been investigated properly except [36]. Our work is differentiated from [36] by (i) exploiting complementary benefits in images and videos rather than directly learning from noisy videos, (ii) retrieving a large set of video clips from web repository rather than using a small number of manually collected videos. Our experimental results show that these differences lead to significant performance improvement.

Our work is closely related to webly-supervised learning [4, 5, 8, 18, 19, 31, 39], which aims to retrieve training examples from the resources on the Web. The idea has been investigated in various tasks, such as concept recognition [4, 5, 8, 39], object localization [5, 8, 19, 39], and fine-grained categorization [18]. The main challenge in this line of research is learning a model from noisy web data. Various approaches have been employed such as curriculum learning [4, 5], mining of visual relationship [8], semi-supervised learning with a small set of clean labels [39], etc. Our work addresses this issue using a model learned from another domain—we employ a model learned from a set of weakly annotated images to eliminate noises in web-crawled videos.

3. Our Framework

The overall pipeline of the proposed framework is described in Figure 1. We adopt a decoupled deep encoder-decoder architecture [13] as our model for semantic segmentation with a modification of its attention mechanism.

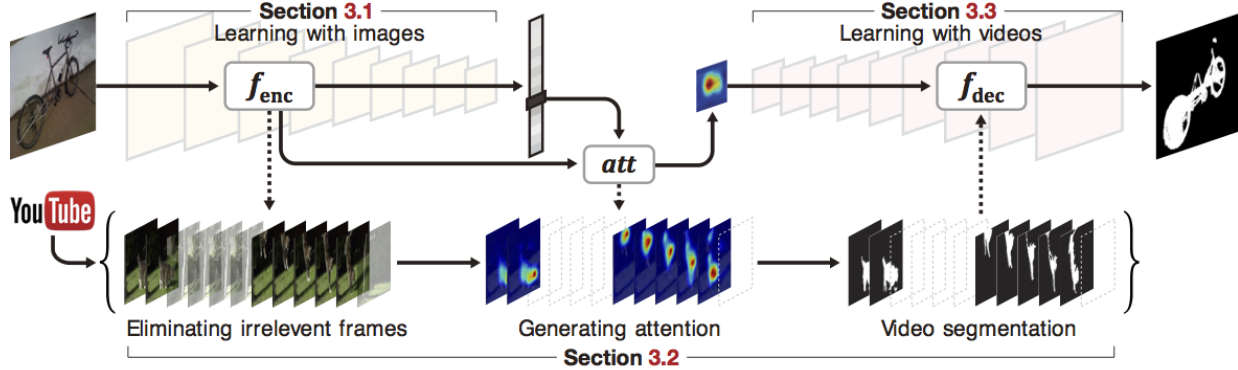


Figure 1. Overall framework of the proposed algorithm. Our algorithm first learns a model for classification and localization from a set of weakly annotated images (Section 3.1). The learned model is used to eliminate noisy frames and generate coarse localization maps in web-crawled videos, where the per-pixel segmentation masks are obtained by solving a graph-based optimization problem (Section 3.2). The obtained segmentations are served as annotations to train a decoder (Section 3.3). Semantic segmentation on still images is then performed by applying the entire network to images (Section 3.4).

In this architecture, the encoder f_{enc} generates class prediction and a coarse attention map that identifies discriminative image regions for each predicted class, and the decoder f_{dec} estimates a dense binary segmentation mask per class from the corresponding attention map. We train each component of the architecture using different sets of data through the procedure below:

- Given a set of weakly annotated images, we train the encoder under a classification objective (Section 3.1).
- We apply the encoder to videos crawled on the Web to filter out frames irrelevant to their class labels, and generate a coarse attention map of the target class per remaining frame. Spatio-temporal object segmentation is then conducted by solving an optimization problem incorporating the attention map with color and motion cues in each relevant interval of videos (Section 3.2).
- We train the decoder by leveraging the segmentation labels obtained in the previous stage as supervision (Section 3.3).
- Finally, semantic segmentation on still images is performed by applying the entire deep encoder-decoder network (Section 3.4).

We also introduce a fully automatic method to retrieve relevant videos from web repositories (Section 4). This method enables us to construct a large collection of videos efficiently and effectively, which was critical to improved segmentation performance. Following sections describe details of each step in our framework.

3.1. Learning to Attend from Images

Let \mathcal{I} be a dataset of weakly annotated images. An element of \mathcal{I} is denoted by $(\mathbf{x}, \mathbf{y}) \in \mathcal{I}$, where \mathbf{x} is an image and $\mathbf{y} \in \{0, 1\}^C$ is a label vector for C pre-defined classes.

We train the encoder f_{enc} to recognize visual concepts under a classification objective by

$$\min_{\theta_{enc}} \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{I}} e_c(\mathbf{y}, f_{enc}(\mathbf{x}; \theta_{enc})), \quad (1)$$

where θ_{enc} denotes parameters of f_{enc} , and e_c is a cross-entropy loss for classification. For f_{enc} , we employ the pre-trained VGG-16 network [34] except its fully-connected layers, and place a new convolutional layer after the last convolutional layer of VGG-16 for better adaptation to our task. On the top of them, two additional layers, global average pooling followed by a fully-connected layer, are added to produce predictions on the class labels. All newly added layers are randomly initialized.

Given the architecture and learned model parameters for f_{enc} , image regions relevant to each class are identified by Class Activation Mapping (CAM) [42]. Let $F(\mathbf{x}) \in \mathbb{R}^{(w \cdot h) \times d}$ be output of the last convolutional layer of f_{enc} given \mathbf{x} , and $W \in \mathbb{R}^{d \times C}$ the parameters for the fully-connected layer of f_{enc} , respectively, where w, h and d denote width, height and the number of channels of $F(\mathbf{x})$. Then for a class c , image regions relevant to the class are highlighted by CAM as follows:

$$\alpha^c = F(\mathbf{x}) \cdot W \cdot \mathbf{y}^c, \quad (2)$$

where \cdot is inner product and $\mathbf{y}^c \in \{0, 1\}^C$ means a one-hot encoded vector for class c . The output $\alpha^c \in \mathbb{R}^{w \cdot h}$ refers to an attention map for class c and highlights local image regions relevant to class c .

3.2. Generating Segmentation from Videos

Our next step is to generate object segmentation masks from a set of weakly annotated videos using the encoder trained in the previous section. Let \mathcal{V} be a set of weakly

annotated videos and $(\mathbf{V}, \mathbf{y}) \in \mathcal{V}$ an element in \mathcal{V} , where $\mathbf{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_T\}$ is a video composed of T frames and $\mathbf{y} \in \{0, 1\}^C$ is the label vector. As in the image case, each video is associated with a label vector \mathbf{y} , but in this case it is a one-hot encoded vector since a single keyword is used to retrieve each video.

Having collected from the Web, videos in \mathcal{V} typically contain many frames irrelevant to associated labels. Thus, segmenting objects directly from such videos may suffer from noises introduced by these frames. To address this issue, we measure class-relevance score of every frame \mathbf{v} in \mathcal{V} with the learned encoder by $\mathbf{y} \cdot f_{\text{enc}}(\mathbf{v}; \theta_{\text{enc}})$, and choose frames whose scores are larger than a threshold. If more than 5 consecutive frames are chosen, we consider them as a single relevant video. We construct a set of relevant videos $\hat{\mathcal{V}}$, and perform object segmentation only on videos in $\hat{\mathcal{V}}$.

The spatio-temporal segmentation of object is formulated by a graph-based optimization problem. Let s_i^t be the i -th superpixel of frame t . For each video $\mathbf{V} \in \hat{\mathcal{V}}$, we construct a spatio-temporal graph $G = (\mathcal{S}, \mathcal{E})$, where a node corresponds to a superpixel $s_i^t \in \mathcal{S}$, and the edges $\mathcal{E} = \{\mathcal{E}_s, \mathcal{E}_t\}$ connect spatially adjacent superpixels $(s_i^t, s_j^t) \in \mathcal{E}_s$ and temporally associated ones $(s_i^t, s_j^{t+1}) \in \mathcal{E}_t$.¹ Our goal is then reduced to estimating a binary label l_i^t for each superpixel s_i^t in the graph G , where $l_i^t = 1$ if s_i^t belongs to foreground (*i.e.*, object) and $l_i^t = 0$ otherwise. The label estimation problem is formulated by the following energy minimization:

$$\min_{\mathcal{L}} E(\mathcal{L}) = E_u(\mathcal{L}) + E_p(\mathcal{L}), \quad (3)$$

where E_u and E_p are unary and pairwise terms, respectively, and \mathcal{L} denotes labels of all superpixels in the video. Details of the two energy terms are described below.

Unary term. The unary term E_u is a linear combination of three components that take various aspects of foreground object into account, and is given by

$$\begin{aligned} E_u(\mathcal{L}) = & -\lambda_a \sum_{t,i} \log A_i^t(l_i^t) - \lambda_m \sum_{t,i} \log M_i^t(l_i^t) \\ & -\lambda_c \sum_{t,i} \log C_i^t(l_i^t), \end{aligned} \quad (4)$$

where A_i^t , C_i^t and M_i^t denote the three components based on attention, appearance, and motion of superpixel s_i^t , respectively. λ_a , λ_c , and λ_m are weight parameters to control relative importance of the three terms.

We use the class-specific attention map obtained by Eq. (2) to compute the attention-based term A_i^t . The attention map typically highlights discriminative parts of the object class, thus provides important evidences for video object segmentation. To be more robust against scale variation

¹We define a temporal edge between two superpixels from consecutive frames if they are connected by at least one optical flow [1].




Figure 2. Qualitative examples of attention map on video frame. (Top: video frame, Middle: attention with single scale input, Bottom: attention with multi-scale input.) Although the encoder is trained on images, its attention maps effectively identify discriminative object parts in videos. Also, multi-scale attention captures object parts and shapes better than its single scale counterpart.

of object, we compute multiple attention maps per frame by varying frame size. After resizing them to the original frame size, we merge the maps through max-pooling over scale to obtain a single attention map per frame. Figure 3.2 illustrates qualitative examples of such attention map. A_i^t is defined as attention over the superpixel s_i^t , and calculated by aggregating the max-pooled attention values within the superpixel.

Although the attention term described above provide strong evidences for object localization, it tends to favor local discriminative parts of object since the model is trained under the classification objective in Eq. (1). To better spread the localized attentions over the entire object area, we additionally take object appearance and motion into account. The appearance term C_i^t is implemented by a Gaussian Mixture Model (GMM). Specifically, we estimate two GMMs based on RGB values of superpixels in the video, one for foreground and another for background. During GMM estimation, we first categorize superpixels into foreground and background by thresholding their attention values, and construct GMMs from the superpixels with their attention values as sample weights. The motion term M_i^t returns higher value if the superpixel exhibiting more distinct motions is labeled as foreground. We utilize inside-outside map from [27], which identifies superpixels with distinct motion by estimating a closed curve following motion boundary.

Pairwise term. We employ the standard Potts model [27, 33] to impose both spatial and temporal smoothness on inferred labels by

$$\begin{aligned} E_p(\mathcal{L}) = & \sum_{(s_i^t, s_j^t) \in \mathcal{E}_s} [l_i^t \neq l_j^t] \phi_s(s_i^t, s_j^t) \phi_c(s_i^t, s_j^t) + \\ & \sum_{(s_i^t, s_j^{t+1}) \in \mathcal{E}_t} [l_i^t \neq l_j^{t+1}] \phi_t(s_i^t, s_j^{t+1}) \phi_c(s_i^t, s_j^{t+1}) \end{aligned} \quad (5)$$

where ϕ_s and ϕ_c denote similarity metrics based on spatial location and color, respectively, and ϕ_t is the percentage of pixels connected by optical flows between the two super-pixels.

Optimization. The Eq. (3) is optimized efficiently by the Graph-cut algorithm. The weight parameters are set to $\lambda_a = 2$, $\lambda_m = 1$, and $\lambda_c = 2$.

3.3. Learning to Segment from Videos

Given a set of generated segmentation annotations obtained in the previous section, we learn the decoder f_{dec} for segmentation by

$$\min_{\theta_{\text{dec}}} \sum_{\mathbf{v} \in \hat{\mathcal{V}}} \sum_{\mathbf{v} \in \mathcal{V}} e_s(\mathbf{z}_{\mathbf{v}}^c, f_{\text{dec}}(\alpha_{\mathbf{v}}^c; \theta_{\text{dec}})), \quad (6)$$

where θ_{dec} means parameters associated with the decoder, $\mathbf{z}_{\mathbf{v}}^c$ is a binary segmentation mask for class c of frame \mathbf{v} , and e_s is a cross-entropy loss between prediction and the generated segmentation annotation. Note that $\mathbf{z}_{\mathbf{v}}^c$ is computed from the segmentation labels \mathcal{L} estimated in the previous section.

We adopt the deconvolutional network [12, 13, 25] as our model for decoder f_{dec} , which is composed of multiple layers of deconvolution and unpooling. It takes the multi-scale attention map $\alpha_{\mathbf{v}}^c$ of frame \mathbf{v} as an input, and produces a binary segmentation mask of class c in the original resolution of the frame. Since our multi-scale attention $\alpha_{\mathbf{v}}^c$ already captures dense spatial configuration of object as illustrated in Figure 3.2, our decoder does not require the additional densified-attention mechanism introduced in [13]. Note that the decoder is shared by all classes as no class label is involved in Eq. (6).

The decoder architecture we adopt is well-suited to our problem for the following reasons. First, the use of attention as input makes the optimization in Eq. (6) robust against incomplete segmentation annotations. Because a video label identifies only one object class, segmentation annotations generated from the video ignore objects irrespective of the labeled class. The decoder will get confused during training if such ignored objects are considered as background since they may be labeled as non-background in other videos. By using the attention as input, the decoder does not care segmentation of such ignored objects and is thus trained more reliably. Second, our decoder learns class-agnostic segmentation prior as it is shared by multiple classes during training [12]. Since static objects (*e.g.*, chair, table) are not well-separated from background by motion, their segmentation annotations are sometimes not plausible for training. The segmentation prior learned from other classes is especially useful to improve the segmentation quality of such classes.

3.4. Semantic Segmentation on Images

Given encoder and decoder obtained by Eq. (1) and (6), semantic segmentation on still images is performed by the entire model. Specifically, given an input image \mathbf{x} , we first identify a set of class labels relevant to the image by thresholding the encoder output $f_{\text{enc}}(\mathbf{x}; \theta_{\text{enc}})$. Then for each identified label c , we compute attention map α^c by Eq. (2), and generate corresponding foreground probability map from the output of decoder $f_{\text{dec}}(\alpha^c; \theta_{\text{dec}})$. The final per-pixel label is then obtained by taking pixel-wise maximum of $f_{\text{dec}}(\alpha^c; \theta_{\text{dec}})$ for all identified classes.

4. Video Retrieval from Web Repository

This section describes details of the video collection procedure. Assume that we have a set of weakly annotated images \mathcal{I} , which is associated with predefined semantic classes. Then for each class, we collect videos from YouTube using the class label as a search keyword to construct a set of weakly annotated videos \mathcal{V} . However, videos retrieved from YouTube are quite noisy in general because videos are often lacking side-information (*e.g.* surrounding text) critical for text-based search, and class labels are usually too general to be used as search keywords (*e.g.* person). Although our algorithm is able to eliminate noisy frames and videos using the procedures described in Section 3.2, examining all videos requires tremendous processing time and disk space, which should be avoided to construct a large-scale video data.

We propose a simple, yet effective strategy that efficiently filters out noisy examples without looking at whole videos. To this end, we utilize thumbnails and key-frames, which are global and local summaries of a video, respectively. In this strategy, we first download thumbnails rather than entire videos of search results, and compute classification scores of the thumbnails using the encoder learned from \mathcal{I} . Since a video is likely to contain informative frames if its thumbnail is relevant to the associated label, we download the video if classification score of its thumbnail is above a predefined threshold. Then for each downloaded video, we extract key-frames² and compute their classification scores using the encoder to select only informative ones among them. Finally, we extract frames within two seconds around each of selected key-frames to construct a video for \mathcal{V} . Videos in \mathcal{V} may still contain irrelevant frames, which are handled by the procedure described in Section 3.2. We observe that videos collected by the above method are sufficiently clean and informative for learning.

²We utilize reference frames used to compress the video [38] as key-frames for computational efficiency. This enables selection and extraction of informative video intervals without decompressing a whole video.

5. Experiments

5.1. Implementation Details

Dataset. We employ the PASCAL VOC 2012 dataset [10] as the set of weakly annotated images \mathcal{I} , which contains 10,582 training images of 20 semantic categories. The video retrieval process described in Section 4 collects 4,606 videos and 960,517 frames for the raw video set \mathcal{V} when we limit the maximum number of videos to 300 and select up to 15 key-frames per video. The classification threshold for choosing relevant thumbnails and key-frames is set to 0.8, which favors precision more than recall.

Optimization. We implement the proposed algorithm based on Caffe [15] library. We use Adam optimization [16] to train our network with learning rate 0.001 and default hyper-parameter values proposed in [16]. The size of mini-batch is set to 14.

5.2. Results on Semantic Segmentation

This section presents semantic segmentation results on the PASCAL VOC 2012 benchmark [10]. We employ *comp6* evaluation protocol, and measure the performance based on mean Intersection Over Union (mIoU) between ground-truth and predicted segmentation.

5.2.1 Internal Analysis

We first compare variants of our framework to verify impact of each component in the framework. Table 1 summarizes results of the internal analysis.

Impact of Separate Training. We compare our approach with [36], which also employs weakly annotated videos, but unlike ours, learns a whole model directly from the videos. For fair comparison, we train our model using the same set of videos from the YouTube-object dataset [31], which is collected manually from YouTube for 10 PASCAL object classes. Under the identical condition, our method substantially outperforms [36] as shown in Table 1. This result empirically demonstrates that our separate training strategy successfully takes advantage of the complementary benefits of image and video domains, while [36] cannot.

Impact of Video Collection. Replacing a set of videos from [31] to the one collected from Section 4 improves the performance by 6% mIoU, although the videos are collected automatically with no human intervention. It shows that (i) our model learns better object shapes from a larger amount of data and (ii) our video collection strategy is effective in retrieving informative videos from noisy web repositories.

Impact of Domain Adaptation. Examples in \mathcal{I} and \mathcal{V} have different characteristics: (i) They have different biases and data distributions, and (ii) images in \mathcal{I} can be labeled by multiple classes while every video in \mathcal{V} is annotated by a single class (*i.e.*, search keyword). So we adapt our model

Table 1. Comparisons between variants of the proposed framework on the PASCAL VOC 2012 validation set. DA stands for domain adaptation on still images.

method	video set	DA	mIoU
MCNN [36]	[31]	Y	38.1
	[31]	N	49.2
Ours	YouTube	N	55.2
	YouTube	Y	58.1

trained on \mathcal{V} to the domain of \mathcal{I} . To this end, we apply the model to generate segmentation annotations of images in \mathcal{I} , and fine-tune the network using the generated annotations as strong supervision. By the domain adaptation, the model learns context among multiple classes (*e.g.* person rides bicycle) and different data distribution, which leads to the performance improvement by 3% mIoU.

5.2.2 Comparisons to Other Methods

The performance of our framework is quantitatively compared with prior arts on weakly supervised semantic segmentation in Table 2 and 3. We categorize approaches based on types of annotations used in training. *Ours* denote our methods described in 4th row of Table 1. Note that MCNN [36] utilizes manually collected videos [31] where associations between labels and videos are not as ambiguous as those in our case.

Our method substantially outperforms existing approaches based on image-level labels, improving the state-of-the-art result by more than 7% mIoU. Performance of our method is even as competitive as the approaches based on extra supervision, which rely on additional human intervention. Especially, our method outperforms some approaches based on relatively stronger supervision (*e.g.*, point supervision [2] and segmentation annotations of other classes [13]). These results show that segmentation annotations obtained from videos are sufficiently strong to simulate segmentation supervision missing in weakly annotated images. Note that our method requires the same degree of human supervision with image-level labels since video retrieval is conducted fully automatically in the proposed framework.

Figure 3 illustrates qualitative results. Compared to approaches based only on image labels, our method tends to produce more accurate predictions on object location and boundary.

5.3. Results on Video Segmentation

To evaluate the quality of video segmentation results obtained by the proposed framework, we compare our method with state-of-the-art video segmentation algorithms on the YouTube-object benchmark dataset [31]. We employed segmentation ground-truths from [14] for evaluation, which provides a binary segmentation masks at every 10 frames for selected video intervals. Following protocols in the pre-

Table 2. Evaluation results on the PASCAL VOC 2012 *validation* set.

Method	bkg	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbk	person	plant	sheep	sofa	train	tv	mean
Image labels:																						
EM-Adapt [26]	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 [28]	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 [17]	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
+Extra annotations:																						
Point supervision [2]	80.0	49.0	23.0	39.0	41.0	46.0	60.0	61.0	56.0	18.0	38.0	41.0	54.0	42.0	55.0	57.0	32.0	51.0	26.0	55.0	45.0	46.0
Bounding box [26]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	58.5
Bounding box [6]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	62.0
Scribble [20]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	63.1
Transfer learning [13]	85.3	68.5	26.4	69.8	36.7	49.1	68.4	55.8	77.3	6.2	75.2	14.3	69.8	71.5	61.1	31.9	25.5	74.6	33.8	49.6	43.7	52.1
+Videos (unannotated):																						
MCNN [36]	77.5	47.9	17.2	39.4	28.0	25.6	52.7	47.0	57.8	10.4	38.0	24.3	49.9	40.8	48.2	42.0	21.6	35.2	19.6	52.5	24.7	38.1
Ours	87.0	69.3	32.2	70.2	31.2	58.4	73.6	68.5	76.5	26.8	63.8	29.1	73.5	69.5	66.5	70.4	46.8	72.1	27.3	57.4	50.2	58.1

Table 3. Evaluation results on the PASCAL VOC 2012 *test* set.

Method	bkg	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbk	person	plant	sheep	sofa	train	tv	mean
Image labels:																						
EM-Adapt [26]	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 [28]	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 [17]	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
+Extra annotations:																						
Point supervision [2]	80.0	49.0	23.0	39.0	41.0	46.0	60.0	61.0	56.0	18.0	38.0	41.0	54.0	42.0	55.0	57.0	32.0	51.0	26.0	55.0	45.0	46.0
Bounding box [26]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	60.4
Bounding box [6]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	64.6
Transfer learning [13]	85.7	70.1	27.8	73.7	37.3	44.8	71.4	53.8	73.0	6.7	62.9	12.4	68.4	73.7	65.9	27.9	23.5	72.3	38.9	45.9	39.2	51.2
+Videos (unannotated):																						
MCNN [36]	78.9	48.1	17.9	37.9	25.4	27.5	53.4	48.8	58.3	9.9	43.2	26.6	54.9	49.0	51.1	42.5	22.9	39.3	24.2	50.2	25.9	39.8
Ours	87.2	63.9	32.8	72.4	26.7	64.0	72.1	70.5	77.8	23.9	63.6	32.1	77.2	75.3	76.2	71.5	45.0	68.8	35.5	46.2	49.3	58.7

Table 4. Evaluation results of video segmentation performance on the YouTube-object benchmark.

method	extra data	class avg.	video avg.
[35]	-	23.9	22.8
[27]	-	46.8	43.2
[40]	bounding box	54.1	52.6
[9]	bounding box	56.2	55.8
Ours	image label	58.6	57.1

vious work, we measure the performance based on mIoU over categories and videos.

The summary results are shown in Table 4. Our method substantially outperforms previous approaches based only on low-level cues such as motion and appearance, since the attention map we employ provides robust and semantically meaningful estimation of object location in video. Interestingly, our method outperforms approaches using object detector trained on bounding box annotations [9, 40] that require stronger supervision than image-level labels. This may be because attention map produced by our method provides more fine-grained localization of an object than coarse bounding box predicted by object detector.

Figure 4 illustrates qualitative results of the proposed approach. Our method generates accurate segmentation masks under various challenges in videos, such as occlusion, background clutter, objects of other classes, and so

on. More comprehensive qualitative results are available at our project webpage³.

6. Conclusion

We propose a novel framework for weakly supervised semantic segmentation based on image-level class labels only. The proposed framework retrieves relevant videos automatically from the Web, and generates fairly accurate object masks of the classes from the videos to simulate supervision for semantic segmentation. For reliable object segmentation in video, our framework first learns an encoder from weakly annotated images to predict attention map, and incorporates the attention with motion cues in videos to capture object shape and extent more accurately. The obtained masks are then served as segmentation annotations to learn a decoder for segmentation. Our method outperformed previous approaches based on the same level of supervision, and as competitive as the approaches relying on extra supervision.

Acknowledgments This work was partly supported by IITP grant (2014-0-00147 and 2016-0-00563), NRF grant (NRF-2011-0031648), DGIST Faculty Start-up Fund (2016080008), NSF CAREER IIS-1453651, ONR N00014-13-1-0762, and a Sloan Research Fellowship.

³http://cvlab.postech.ac.kr/research/weaksup_video/




Figure 3. Qualitative results on the PASCAL VOC 2012 validation images. SEC [17] is the state of the art among the approaches relying only on image-level class labels, and MCNN [36] exploits videos as an additional source of training data as ours does. Compared to these approaches, our method captures object boundary more accurately and covers larger object area.




Figure 4. Qualitative results of the proposed method on the YouTube-object dataset. Our method segments objects successfully in spite of challenges like occlusion (e.g., car, train), background clutter (e.g., bird, car), multiple instances (e.g., cow, dog), and irrelevant objects that cannot be distinguished from target object by motion (e.g. people riding horse and motorbike).

References

- [1] L. Bao, Q. Yang, and H. Jin. Fast edge-preserving patch-match for large displacement optical flow. *IEEE Transactions on Image Processing*, 23(12):4996–5006, 2014.
- [2] A. Bearman, O. Russakovsky, V. Ferrari, and L. Fei-Fei. What’s the Point: Semantic Segmentation with Point Supervision. In *ECCV*, 2016.
- [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 *ICLR*, 2015.
- [4] X. Chen and A. Gupta. Webly supervised learning of convolutional networks. In *ICCV*, 2015.
- [5] X. Chen, A. Shrivastava, and A. Gupta. Neil: Extracting visual knowledge from web data. In *ICCV*, 2013.
- [6] J. Dai, K. He, and J. Sun. BoxSup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *ICCV*, 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 *CVPR*, 2009.
- [8] S. Divvala, A. Farhadi, and C. Guestrin. Learning everything about anything: Webly-supervised visual concept learning. In *CVPR*, 2014.
- [9] B. Drayer and T. Brox. Object detection, tracking, and motion segmentation for object-level video segmentation. In *ECCV*, 2016.
- [10] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010.
- [11] G. Ghiasi and C. C. Fowlkes. Laplacian pyramid reconstruction and refinement for semantic segmentation. In *ECCV*, 2016.
- [12] S. Hong, H. Noh, and B. Han. Decoupled deep neural network for semi-supervised semantic segmentation. In *NIPS*, 2015.
- [13] S. Hong, J. Oh, H. Lee, and B. Han. Learning transferrable knowledge for semantic segmentation with deep convolutional neural network. In *CVPR*, 2016.
- [14] S. D. Jain and K. Grauman. Supervoxel-consistent foreground propagation in video. In *ECCV*, 2014.
- [15] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *MM*, pages 675–678. ACM, 2014.
- [16] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICRL*, 2015.
- [17] A. Kolesnikov and C. H. Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *ECCV*, 2016.
- [18] J. Krause, B. Sapp, A. Howard, H. Zhou, A. Toshev, T. Duerig, J. Philbin, and L. Fei-Fei. The unreasonable effectiveness of noisy data for fine-grained recognition. In *ECCV*, 2016.
- [19] K. Kumar Singh, F. Xiao, and Y. Jae Lee. Track and transfer: Watching videos to simulate strong human supervision for weakly-supervised object detection. In *CVPR*, 2016.
- [20] D. Lin, J. Dai, J. Jia, K. He, and J. Sun. Scribble-sup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016.
- [21] G. Lin, C. Shen, A. van den Hengel, and I. Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *CVPR*, 2016.
- [22] 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 *ECCV*. 2014.
- [23] Z. Liu, X. Li, P. Luo, C. C. Loy, and X. Tang. Semantic image segmentation via deep parsing network. In *ICCV*, 2015.
- [24] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [25] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *ICCV*, 2015.
- [26] G. Papandreou, L.-C. Chen, K. Murphy, and A. L. Yuille. Weakly-and semi-supervised learning of a DCNN for semantic image segmentation. In *ICCV*, 2015.
- [27] A. Papazoglou and V. Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, 2013.
- [28] D. Pathak, P. Krähenbühl, and T. Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *ICCV*, 2015.
- [29] D. Pathak, E. Shelhamer, J. Long, and T. Darrell. Fully convolutional multi-class multiple instance learning. *arXiv preprint arXiv:1412.7144*, 2014.
- [30] P. O. Pinheiro and R. Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, 2015.
- [31] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *CVPR*, 2012.
- [32] G.-J. Qi. Hierarchically gated deep networks for semantic segmentation. In *CVPR*, June 2016.
- [33] C. Rother, V. Kolmogorov, and A. Blake. ”grabcut”: Interactive foreground extraction using iterated graph cuts. In *SIGGRAPH*, 2004.
- [34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [35] K. Tang, R. Sukthankar, J. Yagnik, and L. Fei-Fei. Discriminative segment annotation in weakly labeled video. In *CVPR*, 2013.
- [36] P. Tokmakov, K. Alahari, and C. Schmid. Learning semantic segmentation with weakly-annotated videos. In *ECCV*, 2016.
- [37] R. Venulapalli, O. Tuzel, M.-Y. Liu, and R. Chellapa. Gaussian conditional random field network for semantic segmentation. In *CVPR*, June 2016.
- [38] T. Wiegand, G. J. Sullivan, G. Bjontegaard, and A. Luthra. Overview of the h.264/avc video coding standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(7):560–576, 2003.
- [39] T. Xiao, T. Xia, Y. Yang, C. Huang, and X. Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, 2015.
- [40] Y. Zhang, X. Chen, J. Li, C. Wang, and C. Xia. Semantic object segmentation via detection in weakly labeled video. In *CVPR*, 2015.
- [41] 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 *ICCV*, 2015.
- [42] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016.