Multiple Instance Curriculum Learning for Weakly Supervised Object Detection

Siyang Li¹

siyangli@usc.edu

Xiangxin Zhu²

xiangxin@google.com

Qin Huang¹

qinhuang@usc.edu

Hao Xu¹

hxu1@usc.edu

C.-C. Jay Kuo¹

cckuo@sipi.usc.edu

¹ Department of Electrical Engineering

University of Southern California

Los Angeles, CA, USA

² Google, Inc.

1600 Amphitheatre Pkwy

Mountain View, CA, USA

Abstract

When supervising an object detector with weakly labeled data, most existing approaches are prone to trapping in the discriminative object parts, e.g., finding the face of a cat instead of the full body, due to lacking the supervision on the extent of full objects. To address this challenge, we incorporate object segmentation into the detector training, which guides the model to correctly localize the full objects. We propose the multiple instance curriculum learning (MICL) method, which injects curriculum learning (CL) into the multiple instance learning (MIL) framework. The MICL method starts by automatically picking the easy training examples, where the extent of the segmentation masks agree with detection bounding boxes. The training set is gradually expanded to include harder examples to train strong detectors that handle complex images. The proposed MICL method with segmentation in the loop outperforms the state-of-the-art weakly supervised object detectors by a substantial margin on the PASCAL VOC datasets.

1 Introduction

Object detection is an important problem in computer vision. In recent years, a set of detectors based on convolutional neural networks (CNN) are proposed [15, 25, 30], which perform significantly better than traditional methods (e.g., [14]). Those detectors need to be supervised with fully labeled data, where both object category and location (bounding boxes) are provided. However, we argue that such data are expensive in terms of labeling efforts and thus are not scalable. With the increase of dataset size, it becomes extremely difficult to label the locations of all object instances.

In this work, we focus on object detection with weakly labeled data, where only image-level category labels are provided, while the object locations are unknown. This type of methods is attractive since image-level labels are usually much cheaper to obtain. For each

Figure 1: The training diagram of our MICL method. It iterates over re-localization and re-training. During re-localization, saliency maps are generated from the current detector. Segmentation seeds are obtained from the saliency maps, which later grow to segmentation masks. We use the segmentation masks to guide the detector to avoid being trapped in the discriminative parts. A curriculum is designed based on the segmentation masks and the current top scoring detections. With the curriculum, the multiple instance learning process can be organized in an easy-to-hard manner for the detector re-training.

image in a weakly labeled dataset, the image-level label tells both present and absent object categories. Thus, for each category, we have positive examples where at least one instance of that category is present as well as negative ones where no objects of that category exist.

Some previous methods [10, 11, 12] extract a set of object candidates via unsupervised object proposals (also called object candidates) [10, 13, 14, 15] and then identify the best proposals that lead to high image classification scores for existing categories. However, the best proposals for image classification do not necessarily cover the full objects. (For example, to classify an image as “cat”, seeing the face is already sufficient and even more robust than seeing the whole body, as the fluffy fur can be confused with other animals.) Specifically, the best proposals usually focus on discriminative object parts, which oftentimes do not overlap enough with the extent of full objects, and thus become false positives for detection.

To reduce the false positives due to trapping in the discriminative parts, we use segmentation masks to guide the weakly supervised detector in the typical relocalization-and-retraining loop [16, 17]. The segmentation process starts with a few seeds from the object saliency maps generated by the current detector. Then segmentation masks are obtained by expanding those seeds using the “Seed, Expand and Constrain (SEC)” method [18]. One may use all generated masks to directly supervise the detector. However, it will be misled by the hard and noisy examples where the segmentation network fails to produce reasonably good object masks. To overcome this challenge, we propose the multiple instance curriculum learning (MICL) approach, which combines the commonly used multiple instance learning (MIL) framework with the easy-to-hard learning paradigm - curriculum learning (CL) [9]. The work flow of the proposed MICL system is shown in Fig. 1. It learns from the “easy” examples only in the re-training step of MIL. The re-trained detector is later used to re-localize the segmentation seeds and object boxes. While this process iterates, the training set is grad-

ually expanded from easy to hard examples so the detector learns to handle more complex examples. We identify the easiness of an example by examining the consistency between the results from the detector and the segmenter, without additional supervision on “easiness” required by traditional CL methods. Once the proposed MICL process is finished, the detector is applied to test images directly.

The contributions of this work are summarized as following. First, we incorporate a semantic segmentation network to guide the detector to learn the extent of full objects and avoid being stuck at the discriminative object parts. Second, we propose an MICL method by combining MIL with CL so that the detector is not misled by hard and unreliable examples, and our CL process does not require any additional supervision on the “easiness” of the training examples. Third, we demonstrate the superior performance of the proposed MICL method as compared with the state-of-the-art weakly supervised detectors.

2 Related Work

Weakly supervised object detection. The weakly supervised object detection problem is oftentimes treated as a multiple instance learning (MIL) task [6, 8, 16, 32, 35, 36, 37, 38, 39]. Each image is considered as a bag of instances. An image is labeled as positive for one category if it contains at least one instance from that category, and an image is labeled as negative if no instance from that category appears in it. As a non-convex optimization problem, MIL is sensitive to model initialization and many prior arts focus on good initialization strategies [11, 38]. Recently, people start to exploit the powerful CNN to solve the weakly supervised object detection problem. Oquab *et al.* [27] convert the Alexnet [19] to a fully convolutional network (FCN) to obtain an object score map, but it gives only a rough location of objects. WSDNN [9] constructs a two-branch detection network with a classification branch and a localization branch in parallel. ContextLoc [17] is built on WSDNN and adds an additional context branch. DA [22] proposes a domain adaptation approach to identify proposals corresponding to objects and uses them as pseudo location labels for detector training. Size Estimate (SE) [32] trains a separate size estimator with additional object size annotation and the detector is trained by feeding images with a decreasing estimated size. Singh *et al.* [20] take the advantage of videos and transfer the object appearance to images for object localization. Most of the aforementioned approaches suffer from trapping in discriminative regions as no specific supervision is provided to learn the extent of full objects.

Weakly supervised image semantic segmentation. Similar with the location labels for detector training, the pixel-wise labels for semantic segmentation are also laborious to obtain. Consequently, many researchers focus on weakly supervised image semantic segmentation [8, 18, 24, 28], where only image-level labels or scribbles on objects are available for training. PointSup [8] trains a segmentation network with one point annotation per category (or object) only and uses objectness prior to improve the results. Similarly, ScribbleSup [24] trains the segmentation network with scribbles by alternating between the GraphCut [31] algorithm and the FCN training [26]. SEC [18] generates coarse location cues from image classification networks as “pseudo scribbles”, and constrains the training process with smoothness. In this work, SEC is used for segmentation network training with seeds.

Curriculum learning. The concept of curriculum learning is proposed by Bengio *et al.* [3], indicating that learning from easy to hard examples can be beneficial. It has been applied to various problems in computer vision [21, 29, 32, 40], where people present different definitions to “easy” examples. Some require human labelers to evaluate the difficulty level

of images [49] while others measure the easiness based on labeling time [40]. In our work, we determine the easiness by measuring the consistency between a trained detector and the segmentation masks, without requiring additional human labeling.

3 Proposed MICL Method

The proposed MICL approach starts with initializing a detector to find the most salient candidate (Sec. 3.1). Meanwhile, we obtain the saliency maps from a trained classifier/detector. The saliency maps are thresholded to obtain segmentation seeds. Then a segmentation network is trained to grow object masks from those seeds (Sec. 3.2). After that, we inject the curriculum learning (CL) paradigm into the commonly used re-localization/re-training multiple instance learning (MIL) framework [46] in the weakly supervised object detection problem, leading to the multiple instance curriculum learning (MICL) approach. With MICL we further train the detector to learn the extent of objects under the guidance of the segmentation network (Sec. 3.3).

3.1 Detector Initialization

To start the MIL process detailed in Sec. 3.3, we need an initial detector with certain localization capability. We achieve this by first training a whole image classifier and then tuning it into a detector to identify the most salient candidate. Similar with the methods described in [4, 12, 40], we first extract SS [42] object proposals. In order to find the most salient candidate (MSC) among the proposals, we then add a “saliency” branch composed of an FC layer in parallel with the classification branch, as shown in Fig. 2. This branch takes the region of

Figure 2: The detector with a saliency branch to find the most salient candidate (MSC) for image classification.

interest (RoI, denoted by r) features from the second last layer of the VGG16 network [33] as the input, and computes its saliency. Then the softmax operation is applied, so that the saliency scores of all RoIs for category c , denoted by $h(c; r)$, sum to one, i.e. $\sum_r h(c; r) = 1$. The saliency score is used to aggregate the RoI classification score $p(c; r)$ into image-level scores:

$$p(c) = \sum_r h(c; r)p(c; r) / \sum_r h(c; r) = \sum_r h(c; r)p(c; r). \quad (1)$$

Then, the whole network can be trained with image-level labels using the multi-label cross entropy loss. We rank the RoIs by the combined scores $h(c; r)p(c; r)$ and record the top scoring RoI for the curriculum learning process detailed in Sec. 3.3.

3.2 Segmentation-based Seed Growing (SSG)

In this module, we first obtain saliency maps from a classification (a single region of interest) or detection (multiple regions of interest) network and then train a segmentation network to expand object masks from the salient regions.

Saliency map for a single region. Several methods [64, 43, 44] are proposed to identify the saliency map from a trained image classification network automatically. The saliency map for category c , denoted by $A(x, y; c)$, describes the discriminative power of the cell, (x, y) , on the feature map $f(x, y, k)$ generated by the trained classification network. Here, we use CAM [44] to obtain the saliency maps for existing categories and Grad [64] for background as elaborated below.

Being applied to a GAP classification network¹, the CAM saliency map is defined as

$$A^{CAM}(x, y; c) = \sum_k f(x, y, k) w(k; c), \quad (2)$$

where (x, y) is the cell coordinates of the feature map from the last convolutional layer, $f(x, y, k)$ is the response at the k -th unit at this layer and $w(k; c)$ is the weight parameters of the fully connected (FC) layer corresponding to category c .

The Grad background saliency maps are defined as

$$A_{BG}^{Grad}(x, y) = 1 - \max_{c \in C^+} \frac{1}{Z(c)} \max_k \left| \frac{\partial p(c)}{\partial f(x, y, k)} \right|, \quad (3)$$

where $p(c)$ is the output of the classification network, indicating the probability that the input image belongs to category c , and $Z(c)$ is a normalization factor so that the maximum saliency for all existing categories, denoted by C^+ , is normalized to one. More specifically,

$$Z(c) = \max_{x, y, k} \left| \frac{\partial p(c)}{\partial f(x, y, k)} \right|. \quad (4)$$

Aggregated saliency map from multiple regions. The outputs of a detector are essentially a group of bounding boxes with classification scores. We propose a generalized saliency map mechanism that can be applied to detectors with the RoI pooling layer (e.g., Fast R-CNN [15]) as used in this work). As illustrated in Fig. 3(a), given a RoI r and its feature map, denoted by $f(x, y, k; r)$ (which has a fixed size because the classifier is composed of FC layers), obtained by the RoI pooling operator, one can compute the saliency map within r as

$$A(x, y; c, r) = \sum_k \frac{\partial p(c; r)}{\partial f(x, y, k; r)} \otimes f(x, y, k; r), \quad (5)$$

where $\otimes$ denotes element-wise multiplication and $p(c; r)$ is the classification score for r . To obtain the saliency maps for the entire image, the RoI saliency maps are aggregated via

$$A(x, y; c) = \sum_{r=1}^N \hat{A}(x, y; c, r) p(c; r), \quad (6)$$

where $\hat{A}(x, y; c, r)$ is obtained by resizing $A(x, y; c, r)$ to the RoI size using bilinear interpolation and then padding it to the image size, as shown in Fig. 3(b).

¹GAP refers to global average pooling. GAP networks refer to the architecture with only the top layer being fully connected (FC). Prior to the FC layer is a GAP layer which averages the 3-D feature map into a feature vector.

It is worthwhile to point out that Eq. (5) is reduced to the single region or the whole image saliency map for a classification network with no RoI inputs (i.e., the entire image as one RoI). When the image classification network is a GAP network, it can be further reduced to Eq. (2) as derived in the original CAM [14]. These relations are explained in the supplementary materials.

Figure 3: (a) Compute the saliency map in a RoI, and (b) aggregate RoI saliency maps to the image saliency maps.

Object mask growing from discriminative regions. Segmentation seeds are obtained by thresholding these saliency maps and the seeds from CAM and Grad are simply pooled together. We adopt the SEC method [15] to train the DeepLab [9] network to expand masks from those seeds. In the first round of training, seeds come from a classification network while in later rounds they are from the currently trained detector. The trained segmentation network is applied to all training images and a bounding box is drawn around the largest connected component in the mask. Thus one instance from each existing category is localized. The location information will guide the detector training process described in Sec. 3.3.

3.3 Multiple Instance Curriculum Learning

The commonly used MIL framework usually starts with an initial detector and then alternates between updating the boxes (re-localization) and updating the model (re-training). In re-localization, the current detector is applied to training images and the highest-scoring box is saved. In re-training, the detector is re-trained on the saved boxes.

To re-train the initialized detector from Sec. 3.1, one may use the highest-score boxes produced by itself, but it get stuck easily at the same box. Alternatively, one can use the bounding boxes from the SSG network, but those boxes may not be reliable due to inaccurate segmentation seeds. In other words, relying solely on the initial detector or the segmenter leads to sub-optimal results. To avoid misleading the detector by unreliable boxes, we guide the detector by organizing the MIL process on a curriculum that requires the detection boxes and segmentation masks to agree with each other. Details are elaborated below.

SSG-Guided Detector Training. The easy-to-hard learning principle proposed in [9] has been proved helpful in training a weakly supervised object detector [52, 41]. However, most previous methods require additional human supervision or additional object size information to determine the hardness of a training example, which is expensive to acquire. Instead of seeking additional supervision, we determine the hardness of an example by measuring the consistency between the outputs from the detector and the SSG network. The consistency is defined as the intersection over union (IoU) of the boxes from the SSG and the detector:

$$S(c; z) = IoU(B^{DET}(c; z), B^{SSG}(c; z)), \quad (7)$$

where z represents a positive example for category c ; $B^{DET}(c; z)$ and $B^{SSG}(c; z)$ stand for the predicted bounding boxes of the object by the detector and the SSG network, respectively. As shown in Fig. 1, an example is considered *easy* if

$$S(c; z) \geq T, \quad (8)$$

where T is a threshold to control the hardness of the selected examples. We argue the validity of this criterion from two perspectives. First, those examples are *easier* because the goal for the detector is to mimic the mask expansion ability of SSG, and one example would appear easier if the detector already produces something similar (i.e., the gap between achieved results and the learning target is small). Second, the object localization on those examples are confirmed by both the detector and SSG, meaning that those predicted locations are more reliable. In other words, if $B^{SSG}(c; z)$ significantly deviates from $B^{DET}(c; z)$, it tends to be unreliable. We verify the reliability of $B^{SSG}(c; z)$ on the selected examples in Sec. 4.3.

For an existing category c , one instance is localized by taking the average of $B^{DET}(c; z)$ and $B^{SSG}(c; z)$ on the selected examples. Those localized instances on easy examples are used for further detector training in a fully supervised manner. In this work, we use the popular Fast R-CNN detector [13] with Selective Search [42] as the object proposal generator.

Re-localization and re-training. The detector trained on easy examples lacks the ability to handle hard examples because it focuses on easy ones in the aforementioned training round. Thus, we gradually include more training examples by adopting the re-localization/re-training iterations in the MIL framework [16], illustrated in Fig. 1. In the re-localization step, the trained detector is applied to the whole training set and the highest scoring boxes for existing categories are recorded as the new $B^{DET}(c; z)$. Meanwhile, the outputs from the detector are used to re-localize segmentation seeds, based on which the SSG network is re-trained to generate new $B^{SSG}(c; z)$. By applying the same example selection criterion, another training subset, containing more examples because their results are more similar after learning from each other, is identified and the Fast R-CNN detector is re-trained. The MIL process alternates between re-training and re-localization until all training examples are included. After training is finished, the detector is applied directly to test images.

4 Experiments

4.1 Experiment Settings

Datasets. To evaluate the performance of the weakly supervised MICL detector, we conduct experiments on the PASCAL VOC 2007 and 2012 datasets (abbreviated as VOC07 and VOC12 below) [13], where 20 object categories are labeled. For the MICL detector training, we only use image-level labels, with no human labeled bounding boxes involved.

Evaluation metrics. Following the evaluation of fully supervised detectors, we use the average precision (AP) for each category on the test set as the performance metric. Besides, we choose another metric, the Correct Location (CorLoc) [16], to evaluate weakly supervised detectors, which is usually applied to training images. The CorLoc is the percentage of the true positives among the most confident predicted boxes for existing categories. A predicted box is a true positive if it overlaps sufficiently with one of the ground truth object boxes. The IoU threshold is set to 50% for both metrics.

Implementation Details. The backbone architecture for all modules is the VGG16 [53] network, which is pre-trained on the ImageNet dataset [9]. The whole system is implemented

Figure 4: Qualitative detection results, where the correctly detected ground truth objects are in green boxes and blue boxes represent correspondingly predicted locations. Objects that the model fails to detect are in yellow boxes and false positive detections are in red.

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	prsn	plant	sheep	sofa	train	tv	mAP
WSDNN-Ens[9]	46.4	58.3	35.5	25.9	14.0	66.7	53.0	39.2	8.9	41.8	26.6	38.6	44.7	59.0	10.8	17.3	40.7	49.6	56.9	50.8	39.3
WSDNN[9]	43.6	50.4	32.2	26.0	9.8	58.5	50.4	30.9	7.9	36.1	18.2	31.7	41.4	52.6	8.8	14.0	37.8	46.9	53.4	47.9	34.9
DA[22]	54.5	47.4	41.3	20.8	17.7	51.9	63.5	46.1	21.8	57.1	22.1	34.4	50.5	61.8	16.2	29.9	40.7	15.9	55.3	40.2	39.5
ConLoc[22]	57.1	52.0	31.5	7.6	11.5	55.0	53.1	34.1	1.7	33.1	49.2	42.0	47.3	56.6	15.3	12.8	24.8	48.9	44.4	47.8	36.3
Attn[22]	48.8	45.9	37.4	26.9	9.2	50.7	43.4	43.6	10.6	35.9	27.0	38.6	48.5	43.8	24.7	12.1	29.0	23.2	48.8	41.9	34.5
SE[22]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	37.2
MICL	61.2	51.9	47.1	13.5	10.1	52.1	56.9	71.0	7.6	36.4	49.7	64.5	63.0	57.8	27.9	16.6	30.4	53.8	41.1	40.3	42.6

Table 1: Comparison of mAP values on the VOC07 *test* set. Note that the gaps between the previous approaches and our method are particularly large on categories such as “cat”, “dog” and “horse”. The improvements are from the SSG network, which grows the bounding box to cover the full objects from the discriminative parts (i.e., faces).

with the TensorFlow[10] library. The saliency map thresholds in SSG are 0.2 for objects and 0.9 for background². The parameter T in Eq. 8 balances the number of selected examples and their easiness. Higher T means easier but fewer examples. It trades off between “overfitting to fewer clean examples” and “learning from large noisy data”. We find empirically that $T = 0.5$ gives good results. Other details are found in the supplementary materials.

4.2 Experimental Results

The average precision (AP) on the VOC07 *test* set is shown in Tab. 1. An mAP of 42.6% is achieved by our MICL method, with 3.1% higher than the second best method, DA [22]. Note that DA [22] needs to cache the RoI features from pretrained CNN models for MIL and demands highly on disk space. Our proposed method avoids feature caching and thus is more scalable. The third best method is from ensembles [9] (WSDNN-Ens in Tab. 1). If compared directly to the results without ensembles (WSDNN in Tab. 1), our method is 7.7% superior. Some visualized detection results are shown in Fig. 4.

Tab. 2 shows the CorLoc evaluation on the VOC07 *trainval* set. We achieve 2.5% higher in the CorLoc than WSDNN-Ens [9]. Also, as compared with that of DA [22], which ranks

²We reduce the background threshold to make background seeds take at least 10% of saliency map, as in [16].

the second in the mAP, our MICL method is superior by 8.1%. Note that in terms of both AP and CorLoc, our MICL detector performs much better on certain categories such as “cat”, “dog” and “horse”. Objects in these categories usually have very discriminative parts (i.e. faces). Improvements on those categories are from the SSG network, which grows the bounding boxes from the discriminative regions.

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	prsn	plant	sheep	sofa	train	tv	Avg.
WSDNN-Ens[1]	68.9	68.7	65.2	42.5	40.6	72.6	75.2	53.7	29.7	68.1	33.5	45.6	65.9	86.1	27.5	44.9	76.0	62.4	66.3	66.8	58.0
WSDNN[2]	65.1	63.4	59.7	45.9	38.5	69.4	77.0	50.7	30.1	68.8	34.0	37.3	61.0	82.9	25.1	42.9	79.2	59.4	68.2	64.1	56.1
DA[3]	78.2	67.1	61.8	38.1	36.1	61.8	78.8	55.2	28.5	68.8	18.5	49.2	64.1	73.5	21.4	47.4	64.6	22.3	60.9	52.3	52.4
ConLoc[4]	83.3	68.6	54.7	23.4	18.3	73.6	74.1	54.1	8.6	65.1	47.1	59.5	67.0	83.5	35.3	39.9	67.0	49.7	63.5	65.2	55.1
WSC[5]	83.9	72.8	64.5	44.1	40.1	65.7	82.5	58.9	33.7	72.5	25.6	53.7	67.4	77.4	26.8	49.1	68.1	27.9	64.5	55.7	56.7
MICL	85.3	58.4	68.5	30.4	20.9	67.2	77.1	84.6	24.3	69.5	51.5	80.0	79.8	85.3	46.2	44.9	52.1	64.6	61.3	60.0	60.5

Table 2: Comparison of the CorLoc on the VOC07 *trainval* set.

To our best knowledge, only [[1](#)] and [[2](#)] reported results on the VOC12 dataset. We follow the training/testing split in [[1](#)] and show the results in Table 3. We improve the mAP and the CorLoc by 3.6% and 8.4%, respectively. For performance benchmarking with [[2](#)], we estimate the mAP of our MICL method on the VOC12 *val* set by applying the detector trained on the VOC07 *trainval* set. The obtained mAP is 37.8%, which is 8% higher than [[2](#)]. We should point out that the VOC07 *trainval* set does not overlap with the VOC12 *val* set and its size is about the same as the VOC12 *train* set used in [[1](#)].

	Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	prsn	plant	sheep	sofa	train	tv	Avg.
AP	ConLoc[4]	64.0	54.9	36.4	8.1	12.6	53.1	40.5	28.4	6.6	35.3	34.4	49.1	42.6	62.4	19.8	15.2	27.0	33.1	33.0	50.0	35.3
	MICL	65.5	57.3	53.4	5.4	11.5	48.8	45.4	80.5	7.6	35.2	25.3	75.8	59.5	68.8	18.0	17.0	24.7	37.7	25.8	14.1	38.9
Cor-Loc	ConLoc[4]	78.3	70.8	52.5	34.7	36.6	80.0	58.7	38.6	27.7	71.2	32.3	48.7	76.2	77.4	16.0	48.4	69.9	47.5	66.9	62.9	54.8
	MICL	84.9	78.6	76.9	30.1	32.6	80.0	69.7	90.6	32.1	67.7	47.4	85.5	85.3	85.9	41.4	50.7	62.8	62.7	57.9	41.7	63.2

Table 3: Results on the VOC12 dataset, where the AP and the CorLoc are measured on the *test* and *trainval* set, respectively.

4.3 Analyses

On effectiveness of curriculum learning. To prove the effectiveness of curriculum learning, we set up three baseline Fast R-CNN detectors: (1) trained with pseudo location labels from the initial MSC detector, with no segmentation cue used, (2) trained with pseudo location labels from the SSG network, and (3) based on MIL yet without curriculum (where the subset of training examples is selected randomly in each round and the average of $B^{SSG}(c; z)$ and $B^{DET}(c; z)$ is used as pseudo location labels). The ablation study is conducted on the VOC07 dataset. The comparison of the MICL approach and the three baselines are illustrated in Tab. 4. Adding the segmentation cue into the detector training boosts the CorLoc by 10.9% and mAP by 4.4% from the initial MSC detector. In parallel, MIL without curriculum gives us even slightly larger improvement of 11.9% CorLoc and 6.3% mAP. By introducing curriculum learning (CL), we obtain a total improvement of 18.0% in the CorLoc and 10.6% in the mAP. We also compare the CorLoc change during training with and without CL in Fig. 5.

On easy example selection criterion. One reason that curriculum learning is effective is that the pseudo location labels on the selected subset are more reliable than the average, as shown in Tab. 5. This can be explained by the different preferences of the SSG network and the MSC detector when they localize objects: the SSG tends to group close instances from the same category due to the lack of instance information, resulting boxes larger than

	MSC	SSG	MIL	MICL
CorLoc	42.5	53.4	54.4	60.5
mAP	32.0	36.4	38.3	42.6

Table 4: Performance comparison of the three Fast R-CNN baselines and the one with the proposed multiple instance curriculum learning paradigm.

	Subset	All (SSG)	All (MSC)
CorLoc	72.9	53.2	42.4
mAP	38.0	36.3	32.0

Table 5: Comparison of the CorLoc on the selected training subset versus on the whole set and mAP on the *test* set achieved by the correspondingly trained detectors.

objects; the MSC detector, however, may focus on the most discriminative regions, usually smaller than the true objects. This intuition is confirmed by Fig. 6, where we analyze the localization errors from the SSG and the MSC. Among objects that are mis-localized, errors can be classified into three categories: too large, too small and others³. It is clear that the SSG tends to generate bounding boxes larger than the target while the MSC prefers smaller ones. Thus, if the results from the SSG and the MSC are consistent, the bounding box is neither too small nor too large, indicating reliable locations. We also train a Fast R-CNN detector on the easy subset only and compare with the detectors trained on the whole set with pseudo locations from SSG and MSC, respectively. As shown in Tab. 5, the one trained with easy examples achieves 1.7% higher in mAP than the second best, even though it sees fewer examples, which indicates the importance of the quality of pseudo location labels.

Figure 5: Comparison of the CorLoc performance of the Fast R-CNN detector with and without curriculum learning.

Figure 6: Percentages of three error types among the mis-localized objects from the SSG and the MSC.

5 Conclusion

In this work, we proposed an MICL method with a segmentation network injected to overcome the challenge that detectors often focus on the most discriminative regions when trained without manually labeled tight object bounding boxes. In the proposed MICL approach, where the segmentation-guided MIL is organized on a curriculum, the detector is trained to learn the extent of full objects from easy to hard examples and the easiness is determined automatically by measuring the consistency between the results from the current detector and the segmenter. The benefits from the segmentation network and the power of the easy-to-hard curriculum learning paradigm are demonstrated by extensive experimental results.

Acknowledgements

This research was partially supported by Ittiam Systems. We thank Yaguang Li, Junting Zhang and Heming Zhang for helpful discussions and proofreading.

³We provide details about how mis-localized objects are categorized in the supplementary materials.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [2] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European Conference on Computer Vision*, pages 549–565. Springer, 2016.
- [3] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM, 2009.
- [4] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [5] Hakan Bilen, Marco Pedersoli, and Tinne Tuytelaars. Weakly supervised object detection with posterior regularization. In *British Machine Vision Conference*, volume 3, 2014.
- [6] Hakan Bilen, Marco Pedersoli, and Tinne Tuytelaars. Weakly supervised object detection with convex clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1081–1089, 2015.
- [7] Joao Carreira and Cristian Sminchisescu. Constrained parametric min-cuts for automatic object segmentation. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 3241–3248. IEEE, 2010.
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li 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. IEEE, 2009.
- [10] Thomas Deselaers, Bogdan Alexe, and Vittorio Ferrari. Localizing objects while learning their appearance. In *European conference on computer vision*, pages 452–466. Springer, 2010.
- [11] Thomas Deselaers, Bogdan Alexe, and Vittorio Ferrari. Weakly supervised localization and learning with generic knowledge. *International journal of computer vision*, 100(3): 275–293, 2012.
- [12] Ali Diba, Vivek Sharma, Ali Pazandeh, Hamed Pirsiavash, and Luc Van Gool. Weakly supervised cascaded convolutional networks. *arXiv preprint arXiv:1611.08258*, 2016.
- [13] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The PASCAL visual object classes (VOC) challenge. *International Journal of Computer Vision (IJCV)*, 88(2):303–338, 2010.

- [14] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9):1627–1645, 2010.
- [15] Ross Girshick. Fast R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1440–1448, 2015.
- [16] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Multi-fold mil training for weakly supervised object localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2409–2416, 2014.
- [17] Vadim Kantorov, Maxime Oquab, Minsu Cho, and Ivan Laptev. Contextlocnet: Context-aware deep network models for weakly supervised localization. In *European Conference on Computer Vision*, pages 350–365. Springer, 2016.
- [18] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European Conference on Computer Vision*, pages 695–711. Springer, 2016.
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, 2012.
- [20] Krishna Kumar Singh, Fanyi Xiao, and Yong Jae Lee. Track and transfer: Watching videos to simulate strong human supervision for weakly-supervised object detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [21] Yong Jae Lee and Kristen Grauman. Learning the easy things first: Self-paced visual category discovery. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1721–1728. IEEE, 2011.
- [22] Dong Li, Jia-Bin Huang, Yali Li, Shengjin Wang, and Ming-Hsuan Yang. Weakly supervised object localization with progressive domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3512–3520, 2016.
- [23] Siyang Li, Heming Zhang, Junting Zhang, Yuzhuo Ren, and C.-C. Jay Kuo. Box refinement: Object proposal enhancement and pruning. In *Applications of Computer Vision (WACV), 2017 IEEE Winter Conference on*, pages 979–988. IEEE, 2017.
- [24] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3159–3167, 2016.
- [25] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European Conference on Computer Vision*, pages 21–37. Springer, 2016.
- [26] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.

- [27] Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 685–694, 2015.
- [28] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly- and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1742–1750, 2015.
- [29] Anastasia Pentina, Viktoriia Sharmanska, and Christoph H Lampert. Curriculum learning of multiple tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5492–5500, 2015.
- [30] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 91–99, 2015.
- [31] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM transactions on graphics (TOG)*, volume 23, pages 309–314. ACM, 2004.
- [32] Miaojing Shi and Vittorio Ferrari. Weakly supervised object localization using size estimates. In *European Conference on Computer Vision*, pages 105–121. Springer, 2016.
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [34] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [35] Parthipan Siva and Tao Xiang. Weakly supervised object detector learning with model drift detection. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 343–350. IEEE, 2011.
- [36] Parthipan Siva, Chris Russell, and Tao Xiang. In defence of negative mining for annotating weakly labelled data. In *European Conference on Computer Vision*, pages 594–608. Springer, 2012.
- [37] Parthipan Siva, Chris Russell, Tao Xiang, and Lourdes Agapito. Looking beyond the image: Unsupervised learning for object saliency and detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3238–3245, 2013.
- [38] Hyun Oh Song, Ross B Girshick, Stefanie Jegelka, Julien Mairal, Zaid Harchaoui, Trevor Darrell, et al. On learning to localize objects with minimal supervision. In *ICML*, pages 1611–1619, 2014.
- [39] Hyun Oh Song, Yong Jae Lee, Stefanie Jegelka, and Trevor Darrell. Weakly-supervised discovery of visual pattern configurations. In *Advances in Neural Information Processing Systems*, pages 1637–1645, 2014.

- [40] E Teh, Mrigank Rochan, and Yang Wang. Attention networks for weakly supervised object localization. *BMVC*, 2016.
- [41] Radu Tudor Ionescu, Bogdan Alexe, Marius Leordeanu, Marius Popescu, Dim P. Papadopoulos, and Vittorio Ferrari. How hard can it be? estimating the difficulty of visual search in an image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [42] Jasper RR Uijlings, Koen EA van de Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International Journal of Computer Vision (IJCV)*, 104(2):154–171, 2013.
- [43] Jianming Zhang, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-down neural attention by excitation backprop. In *European Conference on Computer Vision*, pages 543–559. Springer, 2016.
- [44] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2921–2929, 2016.
- [45] C Lawrence Zitnick and Piotr Dollár. Edge boxes: Locating object proposals from edges. In *European Conference on Computer Vision*, pages 391–405. Springer, 2014.