

Pseudo Mask Augmented Object Detection

Xiangyun Zhao*

Northwestern University

zhaoxiangyun915@gmail.com

Shuang Liang†

Tongji University

shuangliang@tongji.edu.cn

Yichen Wei

Microsoft Research

yichenw@microsoft.com

Abstract

In this work, we present a novel and effective framework to facilitate object detection with the instance-level segmentation information that is only supervised by bounding box annotation. Starting from the joint object detection and instance segmentation network, we propose to recursively estimate the pseudo ground-truth object masks from the instance-level object segmentation network training, and then enhance the detection network with top-down segmentation feedbacks. The pseudo ground truth mask and network parameters are optimized alternatively to mutually benefit each other. To obtain the promising pseudo masks in each iteration, we embed a graphical inference that incorporates the low-level image appearance consistency and the bounding box annotations to refine the segmentation masks predicted by the segmentation network. Our approach progressively improves the object detection performance by incorporating the detailed pixel-wise information learned from the weakly-supervised segmentation network. Extensive evaluation on the detection task in PASCAL VOC 2007 and 2012 [12] verifies that the proposed approach is effective.

1. Introduction

Recent years have seen significant progresses in object detection. Since the deep convolutional neural network has been firstly used in R-CNN [19], a lot of improvements have been made, and they improve the performance from many different aspects, *e.g.*, deeper networks and stronger features [44, 45, 24], better object proposals [38, 46], more discriminative and powerful features [29, 1], more accurate localization [36, 15], focusing on a set of hard examples [34, 42].

In this work, we investigate the object detection task from another important aspect, that is, how to exploit object segmentation to improve object detection. Although it has

been well recognized in the literature that the two tasks are closely related and detection could benefit from segmentation, most previous works, *e.g.*, [1, 8], share two common drawbacks.

First, they rely on accurate and pixel-wise ground truth segmentation masks for the segmentation problem. However, such mask annotation is very expensive to obtain. Instead, most large-scale object recognition datasets such as ImageNet [10] and PASCAL VOC [12] only provide bounding box level annotations. In addition, most of these methods only explore how to facilitate object detection with semantic image segmentation, which did not independently consider the characteristics of each instance. We argue that the instance-level segmentation task is more aligned with object detection by considering the object information from different granularity (pixel-level versus box-level). Recently, Mask-RCNN [23] unifies object detection and instance segmentation in a single network, and show that instance segmentation could help object detection. However, pixel-wise instance segmentation labeling is still required.

Second, most works have independent network structures for segmentation and detection tasks, *e.g.*, the state-of-the-art MNC [8] and Mask-RCNN [23]. Although the two tasks often share the same underlying convolutional features, the two networks do not directly interact with each other and the commonality between the two tasks may not be fully exploited. For the existing approaches, the benefits of jointing learning are mostly from the better learned deep feature representation as in a normal multi-task setting. It is seldom explored that how segmentation information can benefit detection directly and more closely in a deep learning framework.

In this work, we propose a novel approach that better addresses the above two issues, which augments the object detector with generated object masks from the bounding box annotation, named as Pseudo-mask Augmented Detection (PAD). It starts from a strong baseline network architecture that directly integrates the state-of-the-art Fast-RCNN [18] network for object detection and InstanceFCN [6] for object segmentation, in a normal multi-task setting.

Given the baseline network, we make two major con-

*This work is done when Xiangyun Zhao was an intern at Microsoft Research Asia.

†Corresponding author.

tributions. First, our PAD treats ground truth object segmentation masks as hidden variables as they are unknown, which are gradually refined by only using bounding box annotations as the supervision, called as *pseudo ground truth masks*. The pseudo masks of training images and the network parameters are optimized alternatively in an EM-like way. To make the alternative learning more effective, we propose two novel techniques. Between each iteration, the pseudo masks are progressively refined by embedding a graphical model, which improves the pixel-wise estimation with a graph-cut optimization with low-level appearance coherence and the ground truth bounding boxes as additional constraints. Beside the iteratively refined pseudo masks, we also incorporate a novel 1D box loss defined over the groundtruth box, as a supervision signal to help improve quality of pseudo masks learning, similar to LocNet [17].

Second, based on the commonality of segmentation and detection tasks, as well as the correlations of the network structures, we propose to connect the two networks such that the segmentation information provides a top-down feedback for detection network. In this way, the learning of detection network is improved as additional supervision signals are back propagated from the segmentation branch. The top-down segmentation feedback considers two contexts, on both the *the global* level and *instance* level. Their effectiveness on improving detection accuracy are both verified in experiments.

The proposed approach is validated using various state-of-the-art network architectures (VGG and ResNet) on several well-known object detection benchmarks (i.e., PASCAL VOC 2007 and 2012). The strong and state-of-the-art performance verifies its effectiveness.

2. Related Work

Joint Segmentation and Detection There exists quite a few works that integrate the object segmentation and detection tasks [14, 20, 11, 5, 1, 8, 23]. In spite of their various techniques, these methods have common limitations: 1) the pixel-level segmentation annotation is required, which is difficult to obtain, 2) the integration of segmentation and detection is usually loose due to the separately trained segmentation and detection network. Our work overcomes the two limitations in an integrated learning framework, where the top-down segmentation feedback is proposed to bridge the segmentation and detection network.

Using Graphical Models for Segmentation Graphical models are widely used for traditional image and object segmentation [3, 39, 32, 30, 41, 4, 27]. Compared to the feature representation learning by the CNNs, the graphical inferences possess the merits of effectively incorporating local and global image constraints (e.g. appearance consistencies, and structure priors) into a single optimization

framework. Recently, some recent works integrate graphical models (e.g., CRF/MRF) into the deep neural networks [47, 35] for a joint training.

In our approach, traditional graph cut based optimization [3] is embedded to refine the pseudo ground truth mask estimation during the iterative learning. It effectively refines the quality of pixel-wise pseudo masks to progressively improve the discriminative capability of detection and segmentation network.

Weakly Supervised Segmentation Due to the difficulty of obtaining large-scale pixel-wise segmentation ground truth, some works resort to weakly supervised learning of segmentation, such as using bounding box annotation [7, 26] or scribbles [33]. Such methods share some similarity with ours by using the iterative optimization to gradually refine the segmentation. They mainly focus on single image segmentation, while our approach jointly optimizes the detection and weakly-supervised object segmentation network. [7] does not adopt low-level color information to refine the segmentation. The most relevant work is [26]. It also iteratively refine the segmentation by graphical models (CRF). Different from it, our approach aims to improve object detection with weakly supervised segmentation.

3. Pseudo-mask Augmented Detection

We focus on facilitating the object detection with the instance-level object segmentation information, using only ground truth bounding box annotations. We denote the set of all ground truth boxes as $B^{gt} = \{B_o^{gt}\}$, where subscript o enumerates all objects in all training images. We use the former notation throughout the paper for its simplicity.

As motivated earlier, it is beneficial to estimate the per-pixel object segmentation as well. An auxiliary object segmentation task is added in a normal multi-task setting. That is, the two tasks share the same underlying convolutional feature maps. Since the ground truth binary object segmentation masks are unknown, we treat them as hidden variables, which are first initialized with B^{gt} , and then iteratively refined in our approach. We call them estimated object masks as *pseudo ground truth masks* from the bounding box annotation, denoted as $M^{pseudo} = \{M_o^{pseudo}\}$.

Let the network parameters be Θ , and the network output for object segmentation and detection be $M(\Theta)$ and $B(\Theta)$, respectively. The network parameters are learned to minimize the loss function

$$L_{seg}(M(\Theta)|M^{pseudo}, B^{gt}) + L_{det}(B(\Theta)|B^{gt}), \quad (1)$$

where the two loss terms are enforced on object segmentation and detection tasks, respectively. As defined the network optimization target, the performance of detection network heavily depends on the quality of estimated pseudo




Figure 1. An overview of our pseudo-mask augmented object detection, consisting of the network architecture and graph cut based pseudo mask refinement. For each image, the detection sub-network and instance-level object segmentation sub-network share convolutional layers (i.e. conv1-conv5 for VGG, conv1-conv4 for ResNet). For segmentation sub-network, position-sensitive score maps are generated by 1×1 convolutional layer, and it is then passed through position-sensitive pooling to obtain object masks. The predicted object masks and bounding box annotations are then combined to refine pseudo masks with a graph-cut refinement, which also provide the supervision signal for network training in next iteration. In each iteration, we explore the global feedback and instance top-down feedback from the instance segmentation sub-network to facilitate the object detection sub-network for better detection performance.

masks M^{pseudo} . That is, the poor estimation of M^{pseudo} leads to poor network learning of $M(\Theta)$, which in turn would cause negative chain effect on whole iterative framework for object detection. We propose an effective learning approach that progressively improves the quality M^{pseudo} from a coarse initialization using B^{gt} , as summarized in Algorithm 1. The detection network parameters Θ and pseudo masks M^{pseudo} are alternatively optimized following an EM-like way, with the other fixed in each iteration.

Note that Algorithm 1 only operates on training images. The learned network parameters Θ are applied on test images to generate detection and segmentation results.

The instance-level segmentation masks $M(\Theta)$ from the pixel-wise prediction of segmentation network are usually noisy and poor. This is partially because pseudo masks are not accurate enough, and the estimation is made in a pixel-wise manner, which does not consider the correlations between the pixels such as smoothness constraints used in most segmentation approaches. As shown in Algorithm 1, we thus propose two novel ingredients to achieve the effective iterative learning. First, in each object mask refinement step (Sec. 3.2), the pseudo ground truth mask for each object is improved using the traditional graphical inference. It is formulated as a global optimization problem that considers not only the current mask estimation from the network, but also the low level image appearance coherence and the ground truth bounding boxes, which is efficiently solved by graph cut [2].

Second, we notice that only using the pseudo mask

Algorithm 1 Iterative learning of network parameters Θ and pseudo ground truth masks M^{pseudo} .

- 1: **input:** ground truth bounding boxes B^{gt}
 - 2: initialize the pseudo masks M^{pseudo} from B^{gt} ;
 - 3: learn Θ_0 with loss in Eq. (1) ▷ Sec. 3.1
 - 4: **for** $t = 1$ **to** T **do**
 - 5: refine M^{pseudo} from $M(\Theta_{t-1})$ and B^{gt} ▷ Sec. 3.2
 - 6: learn Θ_t with loss in Eq. (1) ▷ Sec. 3.1
 - 7: **end for**
 - 8: **output:** final network parameters Θ_t
 - 9: **output:** pseudo ground truth masks M^{pseudo}
-

M^{pseudo} as 2D pixel-wise supervision signals may be not sufficient as the masks themselves are often noisy and not accurate enough. Thus, the 1D box loss (explained in Sec. 3.1) in Eq. (1) and (2) (Sec. 3.1) is incorporated to consider the additional constraints provided by the ground truth bounding box. The 1D loss complements the noisy 2D segmentation loss and performs better regularization on the segmentation network learning.

With the aforementioned two novel techniques, both pseudo masks M^{pseudo} and network parameters Θ are improved steadily, benefiting from each other. Based on the refined object masks, we add connections between the segmentation and detection sub-networks such that the segmentation features provide top-down feedback for the detection, leading to better results in the object detection (Sec. 3.1).




Figure 2. Example predicted pseudo masks by the iterative refinement and employing graph cut optimization.

3.1. Network Architecture and Training

As shown in Figure 1, following the common multi-task learning, we adopt two sub-networks for the object segmentation and detection tasks, which are built on the shared feature maps. We first extract the object proposals, or regions of interest (ROIs) from the Region Proposal Network (RPN) [38]. For simplicity, we do not use the complex training strategy in [38]. Instead, we pre-train the RPN and fix the ROIs throughout our experiments.

Object Segmentation with Pseudo Masks In general, this sub-network can adopt any instance-level object segmentation network such as DeepMask [37], MNC [8], etc. In this work, we adopt the similar technique as in InstanceFCN [6] and FCIS [31], which are state-of-the-art methods for instance segmentation. It applies 1×1 convolutional layer on the feature maps to produce k^2 position sensitive score maps. The k^2 (we use $k = 7$ in this work) maps encode the relative-position information between image pixels and ROIs (e.g., top-left, bottom-right). Given a ROI, its per-pixel score map is assembled by dividing the positive sensitive score maps into k^2 cells and copying each cell’s content from the corresponding k^2 maps. This step generates a fixed-size (28×28) per-ROI foreground probability map.

However, this strategy still has several limitations: first, it only runs on square sliding windows; second, it ignores object category information and is limited to object generate object segment proposals. We extend this approach in two aspects to seamlessly integrate it into the object detection network. First, we extend its k^2 score maps to $k^2 * C$ score maps, where C indicates the category number. In this way, the individual segmentation module for each category is optimized. Second, we employ generic object proposals [38] to replace the square sliding windows and the

position-sensitive ROI pooling layer in [9] on the proposals. During training, a sigmoid layer is applied on each category of the per-ROI score maps to generate the instance foreground probability maps.

Segmentation Loss Our approach estimate a corresponding pseudo mask for each object instance. For a ROI that has intersection-over-union (IOU) larger than 0.5 with a ground-truth object, we define a per-pixel sigmoid cross-entropy loss on the ROI’s foreground probability map, with respect to current pseudo mask of the object that is regarded as a hidden ground truth mask. We call this term *2D mask loss*.

Since the pseudo masks are quite noisy, it may damage our network if directly using it as the supervision. Since each ground truth bounding box tightly encloses the object, this implies that for a horizontal or vertical scan line in the box, at least one pixel on the line should be foreground. On the other hand, all pixels outside of the box should be background. Accordingly, we define a *1D box loss* term for each ROI. Specifically, the predicted foreground mask of each ROI is projected to two 1-d vectors along the horizontal and vertical directions, respectively, by applying a max operation on all values of a scan line. For each position on the 1-d vector, it is denoted as foreground when its corresponding line is inside the box. Otherwise, it is denoted as background. The 1D loss term summarizes the sigmoid cross-entropy loss on all positions of the 1-d vectors. Note that a similar 1D loss idea has also been utilized in LocNet [17] for bounding box regression. In this work, we use it for object segmentation.

In summary, by combining the 2D mask loss and 1D box loss, the segmentation loss in Eq. (1) is computed by

$$L_{\text{seg}} = L_{2D}(M(\Theta)|M^{\text{pseudo}}) + L_{1D}(M(\Theta)|B^{\text{gt}}). \quad (2)$$

Object Detection with Top-down Segmentation Feedback For the detection sub-network, we use the state-of-the-art Fast(er) R-CNN [18, 38]. It applies a ROI pooling layer for each ROI on the feature maps to obtain per-ROI feature maps, and then applies fully connected (FC) layers to output detection results.

In the common multi-task setting, the two sub-networks do not interact and only use the separate optimization target. However, the object segmentation and detection tasks are highly correlated to each other and the sub-networks also share similar structures, we connect the two so that the segmentation network provides top-down feedback information for detection network.

The feedback from segmentation consists of the *global feedback* and the *instance feedback* (the two red dotted arrows in Figure 1). In terms of the global level, the $k^2 * C$ position-sensitive score maps in the segmentation sub-network (before ROI pooling) encode the segmentation information on the whole image. They are of the same spatial dimension of the shared convolutional feature maps but different feature channels, in general. A 1×1 convolutional layer is applied on the score maps to change its channel number to match that of the shared feature maps. The two sets of maps are then element-wisely summed to produce the “rectified” global feature maps for the object detection sub-network.

In terms of instance level, the instance segmentation masks (i.e., per-ROI score maps) encode the specific pixel-wise characteristics of each object instances. As shown in Figure 1, after the ROI pooling step in both sub-networks, the per-ROI instance segmentation score maps from the segmentation branch are passed through a 1×1 convolutional layer and max pooling layer to obtain feature maps with the same dimension as the per-ROI feature maps of the detection branch. The score maps from two branches are then summed to produce the “rectified” instance feature maps.

Afterwards, several fully connected (FC) layers are used to generate the object classification scores and bounding box regression results, in the same way as Fast RCNN [18]. The detection loss in Eq. (1) includes the classification softmax loss and bounding box regression loss for all ROIs,

$$L_{det} = L_{cls}(B(\Theta)|B^{gt}) + L_{reg}(B(\Theta)|B^{gt}). \quad (3)$$

Training Given the estimated pseudo masks M_{pseudo} in each step, the instance-level segmentation network and object detection network are optimized by stochastic gradient descent, using image centric sampling [18]. In each mini-batch, two training images are randomly sampled. The loss gradients from Eq. (1) are back propagated to update all the network parameters jointly.

3.2. Pseudo Mask Refinement

Accurate pseudo mask is the key to bridge the object detection and instance-level segmentation networks. The estimated pseudo masks directly from the segmentation network are usually noisy and blurred. More importantly, as the pixels are considered individually by the convolutional network, the informative interactions between pixels are not fully exploited, such as the smoothness constraint used in traditional image and object segmentation.

In this work, we explore a graphical model to refine the pseudo mask estimation, which jointly incorporates the current mask probabilities from the instance-level segmentation network, the low level image appearance cues and the ground-truth bounding box information. The graphical model is defined on a graph constructed by the super-pixels generated by [13] for each object instance. For each graph, a vertex denotes a super-pixel while an edge is defined over neighboring super-pixels. Note that in this step the spatial range of the pseudo mask is enlarged by 20% from the ground truth bounding box, in order to include more boundary areas and thus improve segmentation quality.

Formally, for all super-pixels $\{x_i\}$ in the pseudo mask under consideration, we estimate their binary labels $\{y_i\}$, where $y_i = 1$ indicates foreground, and 0 for background. Similar to the traditional object segmentation approaches [3, 39], we define a global objective function in the form of

$$\sum_i U(y_i) + \sum_{i,j} V(y_i, y_j). \quad (4)$$

Unary Term. The unary term $U(y_i)$ measures the likelihood of the super pixel x_i being foreground. It considers both the foreground probabilities from the network and the ground truth bounding box b^{gt} , defined as

$$U(y_i) = \begin{cases} 0 & \text{if } y_i = 0 \text{ and } x_i \notin b^{gt} \\ \infty & \text{if } y_i = 1 \text{ and } x_i \notin b^{gt} \\ -\log(1 - \text{prob}_{fg}(x_i)) & \text{if } y_i = 0 \text{ and } x_i \in b^{gt} \\ -\log(\text{prob}_{fg}(x_i)) & \text{if } y_i = 1 \text{ and } x_i \in b^{gt}. \end{cases} \quad (5)$$

The first two cases ensure that x_i is background when it is outside the ground truth bounding box. The last two cases directly adopt the results from the current network estimation when x_i is inside, where $\log(\text{prob}_{fg}(x_i))$ is simply the summation of the pixel-wise log probability of all pixels in the super-pixel x_i . To obtain a pixel’s foreground probability, firstly we only consider the probability of the ground truth object category as the foreground probability. Secondly, the segmentation sub-network outputs all ROIs’ mask probability maps. To obtain the foreground probability on a single ground truth object, we find all ROIs with

method	w/ seg. loss	w/ mask refinement	w/ global feedback	w/ instance feedback	VGG-16	ResNet-50	ResNet-101
ION [1]					75.6		
CRAFT [46]					75.7		
HyperNet [29]					76.3		
RON [28]					75.4		
Faster-RCNN [38]					73.2	74.9	76.4
Ours (a)					74.5	78.1	79.4
Ours (b)	✓				74.9	78.2	79.6
Ours (c)	✓	✓			75.7	78.6	79.9
Ours (d)	✓	✓	✓		76.0	78.7	80.0
Ours (e)	✓	✓		✓	76.2	78.9	80.4
Ours (f) iter. 1	✓	✓	✓	✓	75.9	78.9	80.1
Ours (g) iter. 2	✓	✓	✓	✓	76.4	79.2	80.4
Ours (h) iter. 3	✓	✓	✓	✓	77.0	79.6	80.7

Table 1. Object detection results on VOC 2007 test training on the union set of VOC 2007, VOC 2012 train and validation dataset

IoU larger than 0.5 with the ground-truth object, and average their foreground probabilities together as the foreground probabilities for the object.

Pairwise Term. The pair-wise binary term $V(y_i, y_j)$ considers the local smoothness constraints, defined on all neighboring super pixels. It uses the low level image cues similarly as in [3, 39]. If the neighboring super-pixels are similar in appearance, the cost of assigning them different labels should be high. Otherwise, the cost is low.

We use both color and texture information to measure the similarity as in [33]. For a super-pixel x_i , its color histogram $h_c(x_i)$ is built on the RGB color space using 25 bins for each channel. The texture histogram $h_t(x_i)$ is built on the gradients at the horizontal and vertical orientations with 10 bins for each. The pair-wise binary term is defined as

$$V(y_i, y_j) = [y_i \neq y_j] \left\{ - \frac{\|h_c(x_i) - h_c(x_j)\|_2^2}{\delta_c^2} - \frac{\|h_t(x_i) - h_t(x_j)\|_2^2}{\delta_t^2} \right\}, \quad (6)$$

where $[\cdot]$ is 1 or 0 if the subsequent argument is true/false. δ_c and δ_t are set as 5 and 10, respectively.

The objective function in Eq. (4) is minimized by graph cut solver [2], for the pseudo mask of each object instance in each image. The resulting binary labels $\{y_i\}$ define the refined binary pseudo masks M_{pseudo} , which are then used as supervision signals to train the network in next iteration (Algorithm 1). Some exemplar predicted pseudo masks of object instances are illustrated in Figure 2.

4. Experiments

4.1. Implementation and Training Details

Our experiments are based on Caffe [25] and public Faster-RCNN code [38]. For simplicity, the region pro-

positional network (RPN) is trained once and the obtained object proposals are fixed. We evaluate the performance of the proposed PAD using three state-of-the-art network structures: VGG-16 [43], ResNet-50 [24] and ResNet-101 [24]. We use the publicly available pre-trained model on ILSVRC2012 [40] to initialize all network parameters.

Each mini-batch contains 2 randomly selected images, and we sample 64 region proposals per image leading to 128 ROIs for each network updating step. After training the baseline Faster R-CNN model [38] with OHEM [42] using the above settings, we actually obtain better accuracy than that reported in the original Faster R-CNN [38], as also revealed in Table 1.

We run SGD for 80k iterations with learning rate 0.001 and 40k iterations with learning rate 0.0001. The iteration number T in Algorithm 1 is set as 3 since no further performance increase is observed.

4.2. Ablation study on VOC 2007

Table 1 compares different strategies and variants in our proposed approach, as well as the results from representative state-of-the-art works as reference. Following the protocol in [18], all models are trained on the union set of VOC 2007 [12] *trainval* and VOC 2012 *trainval*, and are evaluated on VOC 2007 *test* set. We evaluate results using VGG-16 [44], ResNet-50 [24], and ResNet-101 models.

We start from the baseline where no pseudo mask is used (Table 1(a)). This is equivalent to our faster R-CNN implementation, which sets a strong and clean baseline. It achieves 74.5%, 78.1%, and 79.4% mAP scores by using VGG-16, ResNet-50, and ResNet-101 models, respectively. We then evaluate the naive pseudo mask baseline (Table 1(b)). This is equivalent to a simple multi-task baseline with coarse pseudo masks. It obtains slightly higher accuracies than (a). It indicates that multi-task learning is slightly helpful but limited, as pseudo mask quality is very

Method	Train	mAP
HyperNet [29]	07++12	71.4
RON [28]	07+12	73.0
CRAFT [46]	07++12	71.3
MR-CNN [16]	07++12	73.9
Faster-RCNN(VGG16) [38]	07++12	70.4
Faster-RCNN(ResNet100)	07++12	73.8
PAD (VGG-16)	07+12	74.4
PAD (ResNet-101)	07+12	79.5

Table 2. Detection results on VOC 2012 test. 07+12: 07 trainval + 12 trainval, 07++12: 07 trainvaltest + 12 trainval.

poor.

To evaluate iterative mask refinement, we report the results of our variant that does not use the mask feedback for object detection network, as Table 1(c). It differs from Table 1(b) in that the iterative mask refinement is used. After three iterations, the mAP scores are 75.7%, 78.6%, and 79.9%, respectively, which are 1.2%, 0.5%, and 0.5% higher than Table 1(a). It verifies that our approach is capable of generating more reasonable pseudo masks.

Furthermore, after performing the top-down segmentation feedback, the detection performance can be further improved by comparing the results of Table 1(h) and Table 1(c). As shown in Table 1(f), Table 1(g) and Table 1(h), the detection performance improve steadily over iterations. Our full model achieves mAP scores of 77.0%, 79.6%, and 80.7% using different networks, respectively, which are 2.5%, 1.5%, and 1.3% higher than Table 1(a). To disentangle the effectiveness of the global and instance feedbacks, we block one of them respectively in Table 1(d) and Table 1(e). We observe that both the feedbacks are effective for boosting the object detection accuracy, and combining them achieves the largest gain.

4.3. Detection Results on VOC 2012

The training data is the union of VOC 2007, VOC 2012 train and validation dataset, following [18]. As reported in Table 2, our approach obtains 74.4% and 79.5% with VGG-16 and ResNet-101, which are substantially better than currently leading methods.

4.4. Segmentation and Detection on VOC 2012 SDS

The Simultaneous Detection and Segmentation (SDS) task is widely used to evaluate instance-aware segmentation methods [22, 8]. Following the protocols in [22, 8], the model training and evaluation are performed on 5623 images from VOC 2012 *train*, and 5732 images from VOC 2012 *validation* sets, respectively. The ground-truth instance-level segmentation masks are provided by the addi-

method	train w/ gt mask?	mAP ^r (%)	mAP ^b (%)
Faster R-CNN		-	66.3
MNC [8]	✓	63.5	-
PAD w/ gt mask	✓	64.5	68.1
PAD w/ Grabcut mask		48.3	66.9
PAD w/o 1D box loss		49.1	66.9
PAD iter. 0		44.3	66.7
PAD iter. 1		52.1	67.0
PAD iter. 2		58.0	67.5
PAD		58.5	67.6

Table 3. Performance comparison on VOC 2012 SDS task.

tional annotations from [21]. The mask-level mAP^r scores and the box-level mAP^b scores are employed as the evaluation metrics for instance-level mask estimation and object detection performance measure, respectively.

First, we report the upper-bound result of training the PAD model using ground-truth object masks, i.e., PAD w/ gt mask in Table 3). It achieves 68.1% mAP^b and 64.5% mAP^r. We note that this “oracle” upper bound is strong. The segmentation accuracy is even higher than the state-of-the-art instance-aware segmentation method, MNC [8].

Second, we evaluate the superiority of alternative training for pseudo mask estimation and network. As shown in Table 3, PAD obtains mAP^b 67.6% and mAP^r 58.5%, which are slightly worse than the upper bound. The instance-level object segmentation results steadily improve with more iterative refinements (iter.1, iter.2). The iterative improvement can be observed in Figure 2. Finally, PAD w/ Grabcut mask shows the instance-level object segmentation results by using the pseudo masks obtained by Grabcut method [39], which is a traditional state-of-the-art object segmentation method. PAD w/o 1D box loss corresponds to results without using the 1D box loss as in Eq. 2. Their comparison with PAD in Table 3 demonstrate the effectiveness of iterative graph cut refinement and 1D segmentation loss, the key ingredients of PAD. Examples for object detection and segmentation are shown in Figure 3 and 4.

4.5. Complexity Analysis

On average, under ResNet-101, our Faster RCNN baseline requires 1.5 sec to process each image, during training. The PAD increases this to 1.9 sec. During testing, the baseline requires 0.42 sec and the PAD 0.49 sec. Note that the overall training time is non-trivially larger, due to the iterative training, but this has no impact in testing. The learned network is only slightly slower than the Faster RCNN. Regarding parameters (see Figure 1), the PAD has 3 additional 1×1 conv layers (1 in InstanceFCN, 1 for global feedback,




Figure 3. Example object detection results of our approach SDS validation set



Figure 4. Example object segmentation results of our approach on SDS validation set

1 for instance feedback). These add about 1M parameters. For reference, VGG has 134M and ResNet-101 has 42M. The increased capacity can be considered minor.

5. Conclusion

In this work, we present a novel Pseudo-mask Augmented object Detection (PAD) model to facilitate object detection with the instance-level segmentation information that are only supervised by bounding box annotation. Starting from the joint object detection and instance segmenta-

tion network, the proposed PAD recursively estimates the pseudo ground-truth object masks from the instance-level object segmentation network training, and then enhance the detection network with a top-down segmentation feedback.

Acknowledgements

Thanks Jifeng Dai the help and discussion for this work.

References

- [1] S. Bell, C. L. Zitnick, K. Bala, and R. Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *CVPR*, 2016.
- [2] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *TPAMI*, 2004.
- [3] Y. Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in nd images. In *ICCV*, 2001.
- [4] J. Carreira and C. 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.
- [5] X. Chen, A. Shrivastava, and A. Gupta. Enriching visual knowledge bases via object discovery and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2027–2034, 2014.
- [6] J. Dai, K. He, Y. Li, S. Ren, and J. Sun. Instance-sensitive fully convolutional networks. In *ECCV*, 2016.
- [7] J. Dai, K. He, and J. Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *ICCV*, 2015.
- [8] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *CVPR*, 2015.
- [9] J. Dai, Y. Li, K. He, and J. Sun. R-fcn: Object detection via region-based fully convolutional networks. In *NIPS*, 2016.
- [10] 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.
- [11] J. Dong, Q. Chen, S. Yan, and A. Yuille. Towards unified object detection and semantic segmentation. In *ECCV*. 2014.
- [12] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes (VOC) Challenge. *IJCV*, 2010.
- [13] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *IJCV*, 2004.
- [14] S. Fidler, R. Mottaghi, A. Yuille, and R. Urtasun. Bottom-up segmentation for top-down detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3294–3301, 2013.
- [15] S. Gidaris and N. Komodakis. Locnet: Improving localization accuracy for object detection. *arXiv preprint arXiv:1511.07763*, 2015.
- [16] S. Gidaris and N. Komodakis. Object detection via a multi-region & semantic segmentation-aware cnn model. *arXiv preprint arXiv:1505.01749*, 2015.
- [17] S. Gidaris and N. Komodakis. Locnet: Improving localization accuracy for object detection. In *CVPR*, 2016.
- [18] R. Girshick. Fast R-CNN. In *ICCV*, 2015.
- [19] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [20] R. Gokberk Cinbis, J. Verbeek, and C. Schmid. Segmentation driven object detection with fisher vectors. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2968–2975, 2013.
- [21] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik. Semantic contours from inverse detectors. In *ICCV*, 2011.
- [22] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Simultaneous detection and segmentation. In *ECCV*. 2014.
- [23] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2980–2988. IEEE, 2017.
- [24] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [25] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [26] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proc. CVPR*, 2017.
- [27] V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011.
- [28] T. Kong, F. Sun, A. Yao, H. Liu, M. Lu, and Y. Chen. Ron: Reverse connection with objectness prior networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, page 2, 2017.
- [29] T. Kong, A. Yao, Y. Chen, and F. Sun. Hypernet: Towards accurate region proposal generation and joint object detection. *arXiv preprint arXiv:1604.00600*, 2016.
- [30] J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the eighteenth international conference on machine learning, ICML*, volume 1, pages 282–289, 2001.
- [31] Y. Li, H. Qi, J. Dai, X. Ji, and Y. Wei. Fully convolutional instance-aware semantic segmentation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2359–2367, 2017.
- [32] Y. Li, J. Sun, C.-K. Tang, and H.-Y. Shum. Lazy snapping. In *ACM Transactions on Graphics (ToG)*, volume 23, pages 303–308. ACM, 2004.
- [33] D. Lin, J. Dai, J. Jia, K. He, and J. Sun. Scribble-sup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016.
- [34] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. *arXiv preprint arXiv:1708.02002*, 2017.
- [35] Z. Liu, X. Li, P. Luo, C.-C. Loy, and X. Tang. Semantic image segmentation via deep parsing network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1377–1385, 2015.
- [36] M. Najibi, M. Rastegari, and L. S. Davis. G-cnn: an iterative grid based object detector. *arXiv preprint arXiv:1512.07729*, 2015.
- [37] P. O. Pinheiro, R. Collobert, and P. Dollár. Learning to segment object candidates. In *NIPS*, 2015.
- [38] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [39] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics*, 2004.

- [40] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [41] J. Shotton, J. Winn, C. Rother, and A. Criminisi. Textonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation. In *European conference on computer vision*, pages 1–15. Springer, 2006.
- [42] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016.
- [43] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [44] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [45] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [46] B. Yang, J. Yan, Z. Lei, and S. Z. Li. Craft objects from images. *arXiv preprint arXiv:1604.03239*, 2016.
- [47] 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.