

Scale-Aware Pixelwise Object Proposal Networks

Zequn Jie, Xiaodan Liang, Jiashi Feng, Wen Feng Lu, Eng Hock Francis Tay, and Shuicheng Yan

Abstract—Object proposal is essential for current state-of-the-art object detection pipelines. However, the existing proposal methods generally fail in producing results with satisfying localization accuracy. The case is even worse for small objects, which, however, are quite common in practice. In this paper, we propose a novel scale-aware pixelwise object proposal network (SPOP-net) to tackle the challenges. The SPOP-net can generate proposals with high recall rate and average best overlap, even for small objects. In particular, in order to improve the localization accuracy, a fully convolutional network is employed which predicts locations of object proposals for each pixel. The produced ensemble of pixelwise object proposals enhances the chance of hitting the object significantly without incurring heavy extra computational cost. To solve the challenge of localizing objects at small scale, two localization networks, which are specialized for localizing objects with different scales are introduced, following the divide-and-conquer philosophy. Location outputs of these two networks are then adaptively combined to generate the final proposals by a large-/small-size weighting network. Extensive evaluations on PASCAL VOC 2007 and COCO 2014 show the SPOP network is superior over the state-of-the-art models. The high-quality proposals from SPOP-net also significantly improve the mean average precision of object detection with Fast-Regions with CNN features framework. Finally, the SPOP-net (trained on PASCAL VOC) shows great generalization performance when testing it on ILSVRC 2013 validation set.

Index Terms—Object proposal, convolutional neural networks, deep learning.

I. INTRODUCTION

IN recent years, object proposal has become crucial for modern object detection methods as an important pre-processing step [1]–[3]. It aims to identify a small number (usually at the order of hundreds or thousands) of candidate regions that possibly contain class-agnostic objects of interest in an image. Compared with the exhaustive search scheme such as sliding windows [4], object proposal methods can significantly reduce the number of candidates to be examined and benefit object detection in following two aspects: they can reduce computation time and allow for applying more sophisticated classifiers.

Manuscript received January 15, 2016; revised June 11, 2016; accepted July 9, 2016. Date of publication July 19, 2016; date of current version August 5, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Chang-Su Kim.

Z. Jie, W. F. Lu, and E. H. F. Tay are with the Department of Mechanical Engineering, National University of Singapore, Singapore 119077 (e-mail: jiezequn@u.nus.edu; mpelwlf@nus.edu.sg; mpetayeh@nus.edu.sg).

X. Liang is with the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: xdliang328@gmail.com).

J. Feng is with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore 119077 (e-mail: jshfeng@gmail.com).

S. Yan is with the Artificial Intelligence Institute, Qihoo 360 Technology Company, Ltd., Beijing 100015, China (e-mail: eleyans@nus.edu.sg).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2016.2593342

Most of existing object proposal methods can be roughly divided into two categories: the classic low-level cues based ones and the modern convolutional neural network (CNN) based ones. The former category of methods mainly exploit low-level image features, including edge, gradient and saliency [5]–[10] to localize regions possibly containing objects. Typically they either follow a bottom-up paradigm *e.g.*, hierarchical image segmentation [8], [11] or examine densely distributed windows [5], [6]. However, it is difficult for them to balance well between localization quality and computation efficiency – they cannot provide object proposals of high quality without incurring expensive computational cost. On the other hand, CNN-based methods either directly predict the coordinates of all the objects in an image [12] or scan the image with a fully convolutional network (FCN) [13], [14] to find the regions of high objectness.¹ Although they can achieve high recall rate w.r.t. relatively loose overlap criteria, *e.g.* intersection over union (IoU) with a threshold value of 0.5, this type of methods usually fails to provide high recall rate under more strict criteria (*e.g.* IoU > 0.7), suggesting their poor localization quality.

Ideally, a generic object proposal generator should offer the following desired features: high recall rate on objects of various categories with only a few proposals, good localization quality for each specific object instance and high computation efficiency. In this work, we make an effort to develop the object proposal method toward these targets.

Our method is motivated by a statistical study on the scale of objects in a collection of natural images. As shown in Figure 2, we plot the distribution of objects with varying scales (measured by number of pixels) from the training and validation sets of the PASCAL VOC detection benchmark [15]. From the figure, one can observe that the objects of small scales (less than 2,000 pixels) actually dominate the distribution. Similar observations also hold in the ILSVRC 2013 and 2014 benchmark [16]. Unfortunately, most of existing methods perform poorly in localizing objects of such small sizes, in terms of the best overlap.² Based on these empirical observations, we argue that the quality of small objects localization is one main bottleneck for further improving the recall rate and average best overlap (ABO) for object proposal methods. Therefore, we focus on tackling such a challenging problem in this work.

In particular, we develop a novel CNN based object proposal method which contains a pixel-wise object proposal network, sharing the similar spirit with object segmentation networks [17]–[19]. Here the “pixel-wise” refers to: for *every*

¹“Objectness” measures membership to foreground objects vs. background

²Best overlap of a particular ground-truth object is defined as the maximal intersection over union (IoU) among all the given proposals w.r.t. this object. Throughout the paper, Average Best Overlap (ABO) is obtained by averaging the best overlap of all the ground-truth objects

pixel in an image, our proposed network model will predict a bounding box of the object containing this pixel. Such a pixel-level comprehensive object proposal strategy fully exploits the available annotations for object segmentation³ and substantially improves the quality of object proposals through enhancing the opportunities of accurately hitting the ground-truth object. As the receptive field of each pixel in CNN is a local region around the pixel, directly predicting the coordinates of the bounding box is challenging due to the various spatial displacements of objects. We thus propose to predict the *offset* of the bounding box w.r.t this pixel, for each pixel.

We then take a further step to focus on enhancing the localization precision for small-scale objects. We propose a *scale-aware* strategy for object proposal, which is inspired by the divide-and-conquer philosophy. Specifically, we train two independent networks, each of which predicts bounding box coordinates for objects at different scales (small or large). Then for each pixel, we will obtain two object proposals for choice. To adaptively fuse them, we introduce another object confidence network. The network consists of two branches – one for predicting objectness confidence and the other one for weighting the large-/small-size⁴ object localization networks. The objectness branch predicts the likelihood of each pixel coming from an object of interest, and the large-/small-size weighting branch trade-offs the contribution of the large-size and small-size networks to final prediction, by predicting the probability of the pixel belonging to an object of a large size. In the training phase, the size of an object can be easily inferred from its annotated segmentation mask, which is used for training the proposed network. For a new image without annotation, both the large-size and small-size object localization networks will predict the bounding box coordinates which are combined according to the weights from the confidence network. An overview of the proposed network model is presented in Figure 1.

Therefore, the scale-aware coordinates prediction can achieve outperforming localization quality for a wide range of object sizes as for various object sizes, the final result can always consider and fuse the bounding boxes predicted by two localization networks robustly based on a reliable large-/small-size weighting mechanism.

To further improve the performance of localizing small objects, we employ a multi-scale strategy for object proposal on a new image. This is inspired by the observation that by enlarging the challenging small object into a larger one, the coordinates prediction error of the small object will be scaled down, as in the case of zooming in on a small object to obtain a clearer view for humans or cameras. Finally, a superpixel based bounding box refinement operation is applied to fine tune the proposals.

In short, we make the following contributions to object proposal generation. Firstly, we introduce a segmentation-

³The segmentation annotations can be readily collected from many public benchmark datasets.

⁴Throughout the paper, we use “large-size network”/“small-size network” to refer to a localization network trained specifically for localizing objects of large/small sizes.

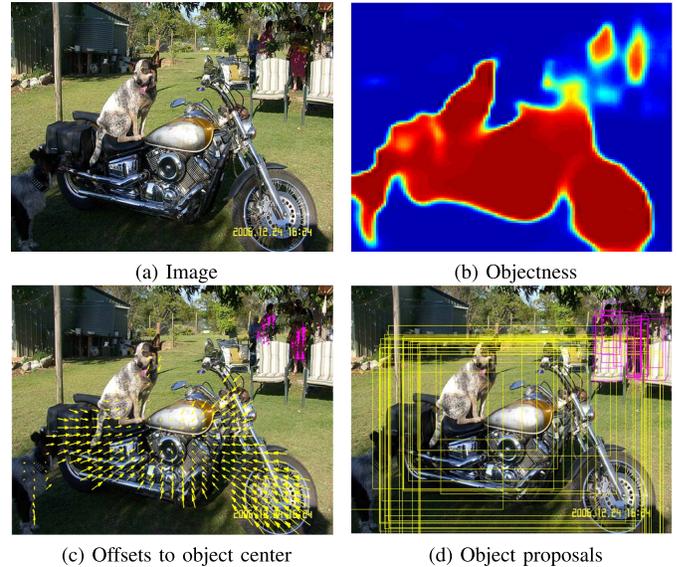


Fig. 1. Examples of predicted “objectness map” in (b), “offsets to object center” after weighted combination in (c) and “object proposals” in (d). “Offsets to object center” is indicated by the arrows pointing to $((x_{\min}^i + x_{\max}^i)/2 - x_i, (y_{\min}^i + y_{\max}^i)/2 - y_i)$ for each pixel i . Yellow and magenta colors in “offsets to object center” and “object proposals” indicate that the prediction is from a pixel with large-size confidence higher than 0.5 or less than 0.5. In the figure, only the predictions for the pixels with objectness higher than 0.5 are shown. (a) Image. (b) Objectness. (c) Offsets to object center. (d) Object proposals.

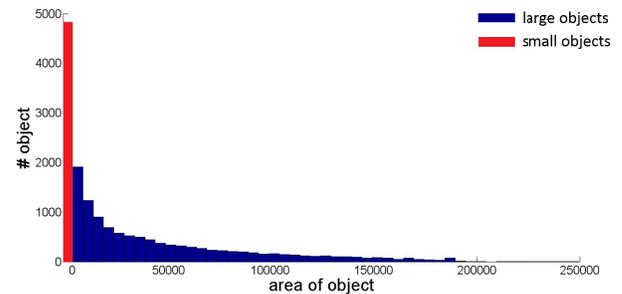


Fig. 2. Distribution of objects w.r.t. their areas (measured by number of contained pixels) on the PASCAL VOC 2012 benchmark. It can be seen that small objects occupy a large proportion of the collection.

like pixel-wise localization network to densely predict the object coordinates for each pixel. Secondly, we develop a scale-aware object localization strategy which combines the predictions from a large-size and a small-size network with a weighting mechanism to boost the coordinates prediction accuracy for a wide range of object sizes. Thirdly, we conduct extensive experiments on the PASCAL 2007, COCO 2014 and ILSVRC 2013 datasets. The results demonstrate that our proposed approach outperforms the state-of-the-art methods by a significant margin, verifying the superiority of the proposed scale-aware pixel-wise object proposal network.

The remainder of this paper is organized as follows. In Section II, we review the related works on object proposal generation. In Section III, we describe our scale-aware pixel-wise localization network. After showing the experimental results in Section IV, we draw the conclusion in Section V.

II. RELATED WORK

The existing object proposal generation methods can be classified into three types: *window scoring methods*, *segment grouping methods* and *CNN-based methods*.

Window scoring methods design different scoring strategies to predict the confidence of containing an object of interest for each candidate window. Generally, this type of methods first initializes a set of candidate window regions across scales and positions in an image, and then sorts them with a scoring model and selects the top ranked windows as proposals. Objectness [20] selects the initial proposals from the salient regions in an image and sorts them based on multiple low-level cues, such as color, edges, location size, etc. [10] proposed a cascade of SVMs trained on gradient features to estimate the objectness. BING [5] trains a simple linear SVM on image gradients and applies it in a sliding window scheme to find the highest scored windows as object proposals. Edge Boxes [6] is also performed in a sliding window manner, but relies on a carefully hand-designed scoring model which sums the edge strengths fully inside the window. Window scoring methods are usually computationally efficient as they do not output segmentation masks for the proposals. However, it seems difficult for them to achieve high recall rate under high overlap criteria (*e.g.* $\text{IoU} > 0.7$), which suggests the poor localization quality. This can probably be attributed to the discrete sampling of the sliding windows which are all in the pre-defined scales and positions.

Segment grouping methods are usually initialized with an oversegmentation to obtain superpixels for an image. Then different merging strategies are adopted to group the similar segments hierarchically to generate the object proposals of all scales. Generally, they follow a bottom-up scheme which relies on diverse low-level image cues including color, shape and texture. For example, Selective Search [8] iteratively merges the most similar segments to form proposals based on several low-level cues. Randomized Prim [9] learns a randomized merging strategy based on the superpixel connectivity graph. Multiscale Combinatorial Grouping (MCG) [11] utilizes multi-scale hierarchical segmentations based on the edge strength and the obtained proposals are then ranked using features including size, location, shape and contour. Geodesic object proposal [21] also depends on superpixels as initialization, and then computes a geodesic distance transform and selects certain level sets of the distance transform as object proposals. [22] proposes learning conditional random field (CRF) in multiscales to classify the superpixels into objects or background. Generally, compared with *window scoring methods*, *segment grouping methods* achieve more consistent and acceptable recall under both loose and strict overlap criteria, indicating a better localization ability. Nevertheless, these methods produce high quality proposals often by multiple segmentations in different scales and color spaces, thus are quite computationally expensive and time-consuming.

CNN-based methods follow the great success of Convolutional Neural Network in other vision tasks, [23]–[26], especially semantic segmentation [27]–[29]. They leverage the powerful discrimination ability of Convolutional Neural Network (CNN) to extract visual features as inputs of other

techniques to produce proposals or directly regress the coordinates of all the object bounding boxes in an image. Multi-Box [12] trains a network to directly predict a fixed number of proposals and their confidences in an image and ranks them with the obtained confidences. RPN [13] uses a Fully Convolutional Network (FCN) to densely generate the proposals in each local patch based on several pre-defined “anchors” in the patch. DeepProposal [30] hunts for the proposals in a sliding window manner by using the CNN features from the final to the beginning layers and training a cascade of linear classifiers to obtain the highest scored windows. Current CNN-based methods typically achieve high recall with only a small number (usually $< 1,000$) of proposals, under loose overlap criteria (*e.g.* $0.5 < \text{IoU} < 0.6$). But similar to window scoring methods, they can hardly achieve high recall rate under more strict overlap criteria (*e.g.* $\text{IoU} > 0.7$). To improve the object proposal localization quality, different from them, our approach predicts the object locations in a pixel-wise manner so that we have much more chances to localize each object with high precision. This also takes the full advantage of the publicly available segmentation masks annotations. This is similar to [31] which deals with object detection task in the object coordinates prediction part. In addition, our scale-aware prediction strategy provides adaptive accurate prediction for both large-size and small-size objects, which also distinguishes our method from others.

III. SCALE-AWARE PIXELWISE PROPOSAL NETWORK

The proposed Scale-aware Pixel-wise Object Proposal Network (SPOP-net) is based on a pixel-wise segmentation-like object coordinates prediction network, and includes a scale-aware localization mechanism for predicting the coordinates of objects of different sizes. In addition, a multi-scale prediction strategy is employed during testing to boost the small objects localization. Finally, a superpixel boundary based proposal refinement is introduced to further improve the proposal precision. We will elaborate all the components of SPOP-net in this section.

A. Pixelwise Localization Network

The proposed Scale-aware Pixel-wise Object Proposal Network (SPOP-net) takes an image of *any* size as input and predicts the location of the object w.r.t. each pixel in the image. More concretely, for each pixel, SPOP-net predicts the normalized coordinates of the bounding box of the object that contains the pixel. The predictions from the background pixels make no sense and will be ranked behind due to low objectness scores they obtain, thus making no difference to the recall performance of top-ranked proposals, which will be detailed later. In this subsection, we first explain the architecture of SPOP-net and then elaborate on how to train and apply the SPOP-net.

1) *Architecture*: Our SPOP-net is built upon a pre-trained DeepLab-LargeFOV segmentation network [17]. Its architecture is shown in Table I. The receptive field of our localization network in the last layer is 435×435 . This large receptive field enables SPOP-net to “see” a large region of the image in its last layer and predict the object bounding boxes effectively.

TABLE I
DETAILS OF DEEPLAB-LARGEFOV NETWORK ARCHITECTURE

layer	#channel	kernel size	stride	zero-padding	size	hole size	training map size	receptive field size	#weight
input image	3	-	-	-	-	-	513*513	435*435	-
conv1_1	64	3*3	1*1	1*1	-	-	513*513	433*433	1.75K
conv1_2	64	3*3	1*1	1*1	-	-	513*513	431*431	36K
pool1	64	3*3	2*2	1*1	-	-	257*257	215*215	-
conv2_1	128	3*3	1*1	1*1	-	-	257*257	213*213	72K
conv2_2	128	3*3	1*1	1*1	-	-	257*257	211*211	144K
pool2	128	3*3	2*2	1*1	-	-	129*129	105*105	-
conv3_1	256	3*3	1*1	1*1	-	-	129*129	103*103	288K
conv3_2	256	3*3	1*1	1*1	-	-	129*129	101*101	576K
conv3_3	256	3*3	1*1	1*1	-	-	129*129	99*99	576K
pool3	256	3*3	2*2	1*1	-	-	65*65	49*49	-
conv4_1	512	3*3	1*1	1*1	-	-	65*65	47*47	1.13M
conv4_2	512	3*3	1*1	1*1	-	-	65*65	45*45	2.25M
conv4_3	512	3*3	1*1	1*1	-	-	65*65	43*43	2.25M
pool4	512	3*3	1*1	1*1	-	-	65*65	41*41	-
conv5_1	512	3*3	1*1	2*2	2*2	2*2	65*65	37*37	2.25M
conv5_2	512	3*3	1*1	2*2	2*2	2*2	65*65	33*33	2.25M
conv5_3	512	3*3	1*1	2*2	2*2	2*2	65*65	29*29	2.25M
pool5	512	3*3	1*1	1*1	-	-	65*65	27*27	-
pool5a	512	3*3	1*1	1*1	-	-	65*65	25*25	-
fc6	1024	3*3	1*1	-	12*12	12*12	65*65	1*1	4.5M
fc7	1024	1*1	1*1	-	-	-	65*65	1*1	1M

2) *Training*: For each pixel, the pixel-wise localization network aims to predict the bounding box coordinates $\mathbf{t} = (x_{\min}/w, y_{\min}/h, x_{\max}/w, y_{\max}/h)$ of the object that contains this pixel. Here (x_{\min}, y_{\min}) and (x_{\max}, y_{\max}) denote the coordinates of the top-left and bottom-right corners of the object bounding box containing the pixel; h and w represent the height and the width of the image plane respectively. Therefore, for a single object, all the pixels inside this object are given the same ground-truth values $(x_{\min}/w, y_{\min}/h, x_{\max}/w, y_{\max}/h)$. We train the pixel-wise localization network to minimize the following localization error \mathcal{L} that is proportional to the Euclidean distance between the predicted coordinate vector \mathbf{t}_i and the ground-truth coordinate vector \mathbf{t}_i^* for all the foreground pixels. The loss function \mathcal{L} is defined as

$$\mathcal{L} = \sum_i p_i^* \|\mathbf{t}_i - \mathbf{t}_i^*\|^2, \quad (1)$$

where \mathbf{t}_i is the predicted 4-d object coordinate vector, and p_i^* is a binary variable indicating whether the pixel i is a foreground one: it takes 1 if the pixel i is from a foreground object and 0 otherwise. Such a filtered loss (through p_i^*) enables the localization network to concentrate on localizing foreground objects without being distracted by background pixels in the training phase. In the practical implementation, as the final layer has smaller size than the input image, we resize the ground-truth coordinate map to the same small size as the final layer.

However, due to the possible spatial displacement (e.g. two exactly the same objects could appear at different locations in an image), accurately predicting the absolute object bounding box coordinates is difficult. It is because these two objects have the same visual input for the model, but their locations the model needs to learn to predict are totally different. To solve this issue, for each pixel, we change its learning targets from the absolute object bounding box coordinates to the offsets from the pixel to the object bounding box. E.g. for object

bounding box coordinate x_{\min}/w , we change the target from x_{\min}/w to $(x_{\min} - x_{\text{self}})/w$, here x_{self} is the x coordinate of the pixel itself. Changing the coordinates to offsets can be conveniently achieved by element-wisely summing the output of the 2nd last layer and the spatial coordinate map (x or y values of all the pixels themselves). Then the absolute object bounding box coordinates can be used as learning targets for the final layer. In this way, applying the absolute coordinates learning targets to the final layer is equivalent to applying the following object coordinate offsets to the 2nd last layer.

$$\left(\frac{x_{\min} - x_{\text{self}}}{w}, \frac{y_{\min} - y_{\text{self}}}{h}, \frac{x_{\max} - x_{\text{self}}}{w}, \frac{y_{\max} - y_{\text{self}}}{h} \right)$$

Then we can directly obtain the absolute object proposal coordinates from the predictions of the final layer. After obtaining the output map from the final layer having a smaller size than the input image, all the subsequent procedures (e.g. refinement, ranking and NMS) are only based on the output map of smaller size. Because we just leverage pixel-level prediction of proposals for having higher chance to hit the ground-truth objects accurately instead of doing pixel-level classification as DeepLab. If resizing the smaller output map back into the original size, the subsequent refinement, ranking and NMS steps will bring much higher computation burden but not significant performance improvement.

B. Scale-Aware Localization

A fully trained pixel-wise localization network can predict the coordinates of object bounding boxes w.r.t. each pixel from an image. However, a single network model may not be able to well handle all the annotated objects that have quite diverse sizes and only offers inferior localization performance for objects of small sizes. To verify this point, we conduct the following preliminary experiments to evaluate the errors of bounding box prediction for large and small objects, using

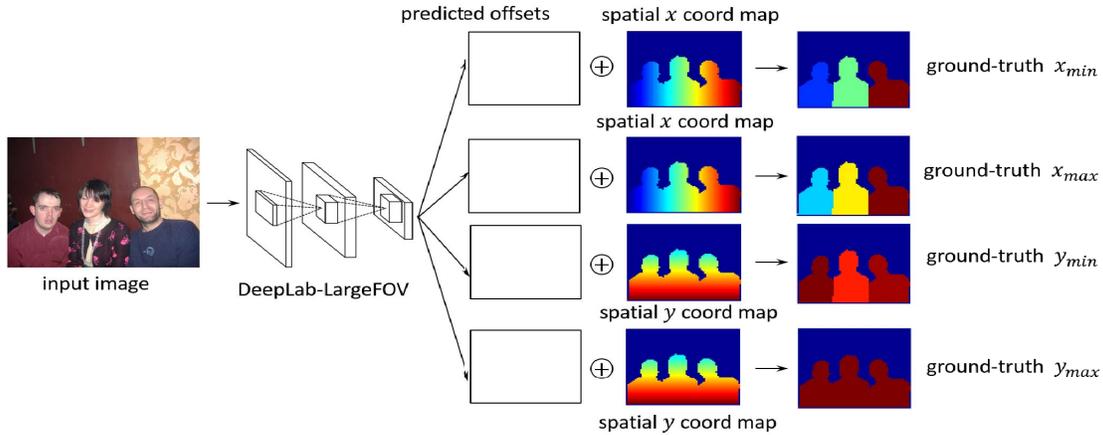


Fig. 3. An image passes through several layers to obtain four maps in the second last layer. In the second last layer, two maps are element-wise summed with spatial x coord map to produce the final prediction for the x_{min} and x_{max} of the corresponding objects for all the pixels, and the other two maps are element-wise summed with spatial y coord map to produce the final prediction for the y_{min} and y_{max} of the corresponding objects for all the pixels. In this way, the four maps in the second last layer in our fully trained network actually predict the four offsets between each pixel position and its corresponding object location, which makes it easier for the network to predict the object coordinates in the final layer. Different colors in the ground-truth maps and spatial coord maps represent different values. Note that we only show the foreground regions of spatial x and y maps for better view.

TABLE II

L_2 ERRORS OF NORMALIZED COORDINATES PREDICTION FOR BOTH LARGE ($\geq 2,000$ PIXELS) AND SMALL OBJECTS ($< 2,000$ PIXELS) IN VOC 2007 TESTING SET, BASED ON THE NETWORK TRAINED ON THE ANNOTATED OBJECTS OF ALL SIZES

	large objects	small objects
x_{min}^{err}/w	0.0090	0.0270
y_{min}^{err}/h	0.0064	0.0160
x_{max}^{err}/w	0.0080	0.0412
y_{max}^{err}/h	0.0088	0.0476

a single pixel-wise localization network trained on the annotated objects of all sizes. The evaluation results are shown in Table II.

From Table II, one can observe that the network trained on all the objects of different sizes produces an L_2 error for small objects that is about 3 to 6 times larger than the error for large objects. This demonstrates the poor localization ability of a single network model for small objects.

The difficulty of accurately localizing both large and small objects using a single network arguably lies in handling the highly diverse offsets of large and small objects. Apart from this, another difficulty comes from the extremely unbalanced training samples between the pixels from large and small objects. Such imbalance leads to the fact that training error of large objects dominates the training loss to minimize.

Also, we empirically verify the sample imbalance through statistics on the pixel-level distribution of the annotations in terms of the area of the object (see Figure 4) since our pixel-wise localization network is trained on pixel-level annotations.

To improve the localization accuracy for small objects, we propose a *scale-aware* localization strategy. Roughly, in the scale-aware strategy, two localization networks are trained – which share the same architecture – with two non-overlapped subsets of the objects. The large-size network is only trained on the pixels belonging to large objects and the small-size

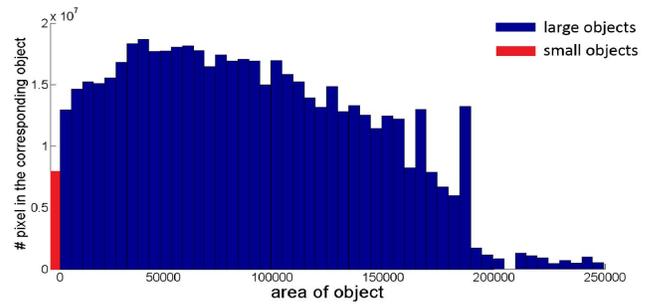


Fig. 4. The distribution of all the pixels w.r.t. the area of the object each pixel belongs to. It is shown that although the number of small objects is large according to Figure 2, the number of pixels belonging to small objects is still small, leading to the unbalanced pixel-level training samples.

network is only trained on the pixels belonging to small objects. The loss function to be optimized for the large-size and small-size network are shown in Eqn. (2) and Eqn. (3) below respectively:

$$L_l = \sum_i l_i^* \|\mathbf{t}_i - \mathbf{t}_i^*\|^2 \quad (2)$$

$$L_s = \sum_i s_i^* \|\mathbf{t}_i - \mathbf{t}_i^*\|^2 \quad (3)$$

where l_i^* and s_i^* are binary indicators showing whether the pixel i belongs to a large object or a small object. The effectiveness of training such scale-aware networks is validated by evaluating the L_2 errors of small objects location prediction with the small-size network. See Table III. During the testing phase, the two networks work simultaneously to output their own prediction for an image. Then, the predictions from two networks are combined with an adaptive weighting scheme.

The weight is output by a network trained for classifying large and small objects pixel-wisely and the weight is equal to the confidence of the pixel belonging to a large object obtained in the last layer of the network. Such a classification network is termed as “confidence network”.

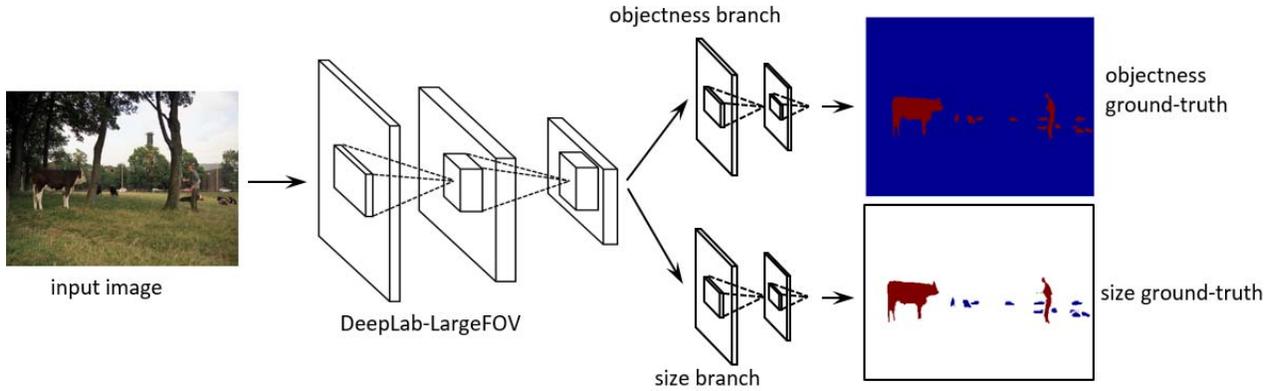


Fig. 5. Illustration of the “confidence network” which bifurcates into two branches to perform foreground/background classification and large/small object classification after all the layers of “DeepLab-LargeFOV” network. Both the sub-networks contain two convolution layers with 3×3 kernel size. The first layer outputs 1,024 feature maps while the second (also the last) layer produces two maps showing the final confidence of their own task. In the ground-truth map of the foreground/background classification branch, red pixels are in foreground objects while blue pixels are in background. In the ground-truth map of the large/small object classification branch, red pixels are in large objects, blue pixels are in small objects and white pixels are background pixels thus are not considered during training.

TABLE III

L_2 ERRORS OF NORMALIZED COORDINATES PREDICTION FOR SMALL OBJECTS ($< 2,000$ PIXELS) IN VOC 2007 TESTING SET, BASED ON THE NETWORK TRAINED ONLY ON SMALL OBJECTS

	small objects
x_{\min}^{err}/w	0.00058
y_{\min}^{err}/h	0.00040
x_{\max}^{err}/w	0.00068
y_{\max}^{err}/h	0.00086

The structure of the confidence network is illustrated in Figure 5. Apart from the large/small classification branch, the confidence network also outputs the objectness confidence in another branch aiming to classify all the pixels into two categories, *i.e.*, foreground pixels and background pixels.

In the confidence network, the two branches share the convolutional features in the lower layers. The last feature maps shared are then fed into the two branches. The intuition for dividing the confidence network into two branches at the higher layer is that for different tasks, the low-level features are usually common and can be shared [32], while the semantically high-level features extracted by the higher layers may be totally different for different tasks. For example, the foreground/background classification task prefers the common features that are insensitive to different sizes of objects, but the large/small classification task aims to extract the discriminative features between large and small objects. The large receptive field (*i.e.* 435×435) in the last layer of the “confidence network” provides a sufficient large view enabling the prediction of both foreground/background and large/small classifications.

The objective function to be optimized during training the confidence network is a multi-task cross-entropy loss:

$$L = \sum_i p_i^* \log(p_i) + (1 - p_i^*) \log(1 - p_i) + \sum_i p_i^* (z_i^* \log(z_i) + (1 - z_i^*) \log(1 - z_i)). \quad (4)$$

Here p_i^* and p_i are the ground-truth label of the foreground/background classification and the predicted confidence of being a foreground pixel for pixel i , respectively. z_i^* and z_i are the ground-truth label of the large/small object classification and the predicted confidence of being contained in a large object for pixel i , respectively. Note that the second term is only activated when p_i^* equals 1, indicating that the pixel belongs to a foreground object. After the large object confidence z_i for the pixel i is obtained, the final predicted coordinates of the object it belongs to are the weighted sum of the predictions by the large-size and small-size networks as follows.

$$\mathbf{t}_i = z_i \mathbf{t}_{l,i} + (1 - z_i) \mathbf{t}_{s,i} \quad (5)$$

where $\mathbf{t}_{l,i}$ and $\mathbf{t}_{s,i}$ are the predictions by the large-size and the small-size network respectively. Then we treat the predicted object coordinates by each pixel as an initial proposal to be passed to the later proposal refinement and non-maximum suppression (NMS) steps to obtain the final object proposals.

C. Multi-Scale Inference

To further enhance the accuracy of small objects localization, we propose to employ a multi-scale prediction strategy in the testing phase. The motivation is quite straightforward: by enlarging the challenging small object into a larger one, the coordinates prediction error of the small object will be scaled down, which is similar to zooming in on a small object to improve the localization accuracy. At the enlarged scale, all the proposals in the enlarged image will be mapped back to their corresponding positions at the original scale.

Therefore, given a testing image, in addition to its original scale, we resize it into a larger scale and run the prediction process as well. Specifically, both on the original scale and the enlarged scale, we simultaneously run the two localization networks (*i.e.* large-size and small-size) and the confidence network, and combine the both location predictions weighted by the large object confidence z_i of its own scale. As all the feed-forward computation of the networks is independent and

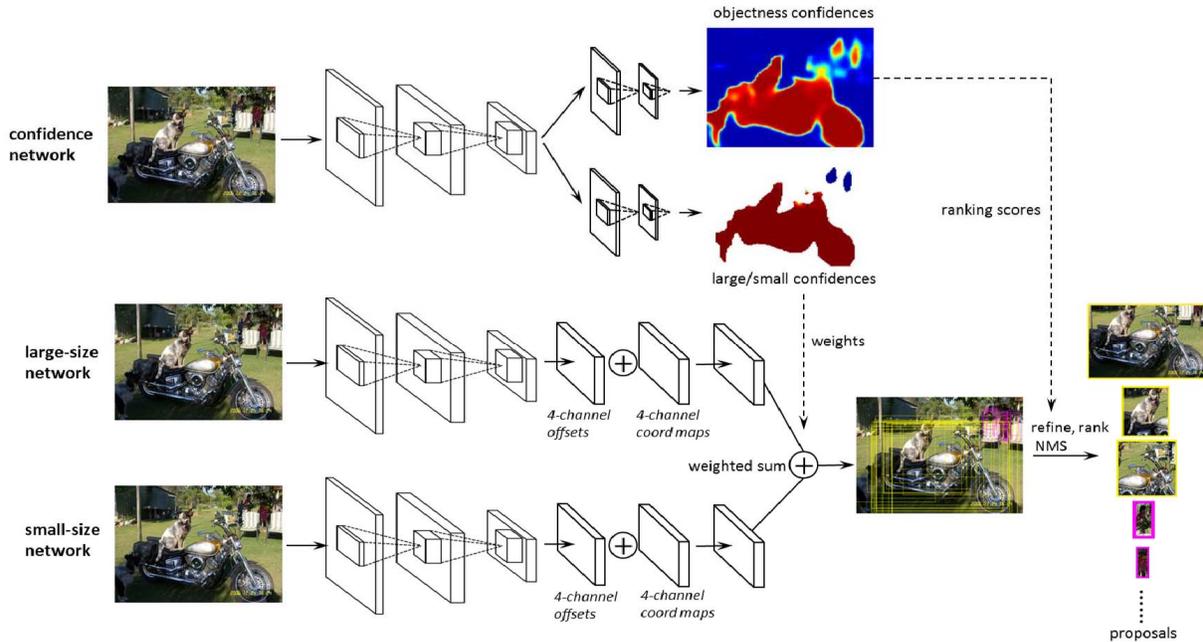


Fig. 6. Overview of our approach. An image passes the confidence network to obtain the pixel-wise objectness confidences and large/small size confidences (red color represents higher values, e.g., high objectness and high large size confidences). The image also passes two localization networks to obtain the predicted pixel-wise large and small object coordinates (x_{\min} , y_{\min} , x_{\max} , y_{\max}), respectively. The feed-forward computation of the three networks are independent and can be run in parallel. Then the final predicted object coordinates are the sum of the predictions by large-/small-size networks weighted by the large/small size confidences. Using the objectness confidences as ranking scores, the final proposals are produced after refinement, ranking and NMS. For multi-scale inference, all the above procedures are run for the enlarged input image as well. Then the proposals obtained by both the original and enlarged scales are mixed in the ranking and NMS.

can be performed in parallel, the computation time cost can remain relatively low.

D. Proposal Refinement

We then refine the two sets of proposals obtained in both original and enlarged scales. An inherent weakness for object localization by regressing the four coordinates with CNN is that the objectness and coordinates ground-truths only permit determining the most discriminative foreground windows. Therefore, even though the windows decided by the localization networks are likely to overlap with target objects, it cannot be ensured that they are able to delineate object boundaries well.

To take object boundaries into consideration, we utilize a superpixel boundary based window refinement method, similar to [33]. The main idea is to expand or shrink the proposals to align the four sides of the proposals with the boundaries of the superpixels better. The reason for using superpixels is that the boundaries of superpixels are informative indicators of object boundaries and superpixels can be generated efficiently with off-the-shelf algorithms (e.g. SLIC [34]). Specifically, for each proposal, we generate two versions of refined proposals, i.e. the minimum bounding rectangle of all the superpixels entirely inside this proposal and the minimum bounding rectangle of all the superpixels entirely inside this proposal or straddling this proposal (see Figure 7). As illustrated in Figure 7, expansion and shrinkage offer two possible ways of getting close to the ground-truth box for the proposals with different location biases to the ground-truth. Therefore, we pass all the two versions of refined proposals as well as the initial proposals to the later proposal ranking and NMS processing.

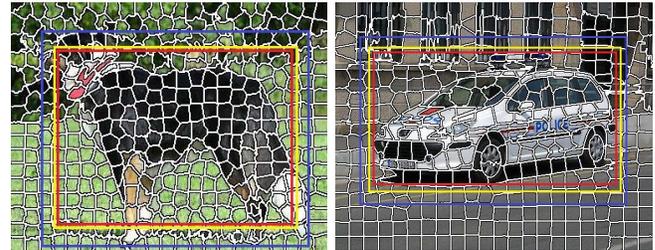


Fig. 7. Illustration of proposal refinement using superpixel boundary based expansion and shrinkage. Yellow boxes represent initial proposals; red boxes and blue boxes are the corresponding refined proposals after shrinkage and expansion respectively. In the left example, expansion finds a closer box to the ground-truth, but in the right example, shrinkage helps the proposal get close to the ground-truth.

In the stage of proposal ranking, we sort all the proposals (including the initial and the two refined ones in both original and enlarged scale) by their objectness confidence p_i . Recall p_i is the output from foreground/background classification branch of the confidence network. For each initial proposal, its two versions of refined proposals are assigned with the same objectness confidence p_i as itself. Finally, the standard non-maximum suppression (NMS) is employed to remove the highly overlapped redundant proposals.

IV. EXPERIMENTS AND DISCUSSION

A. Experimental Setups

The proposed Scale-aware Pixel-wise Object Proposal Network (SPOP-net) is trained on the SBD annotations [35] of PASCAL VOC 2012 trainval set, which provides 11,355 images with fine segmentation masks annotations. We manually label the objects containing more than 2,000 pixels as

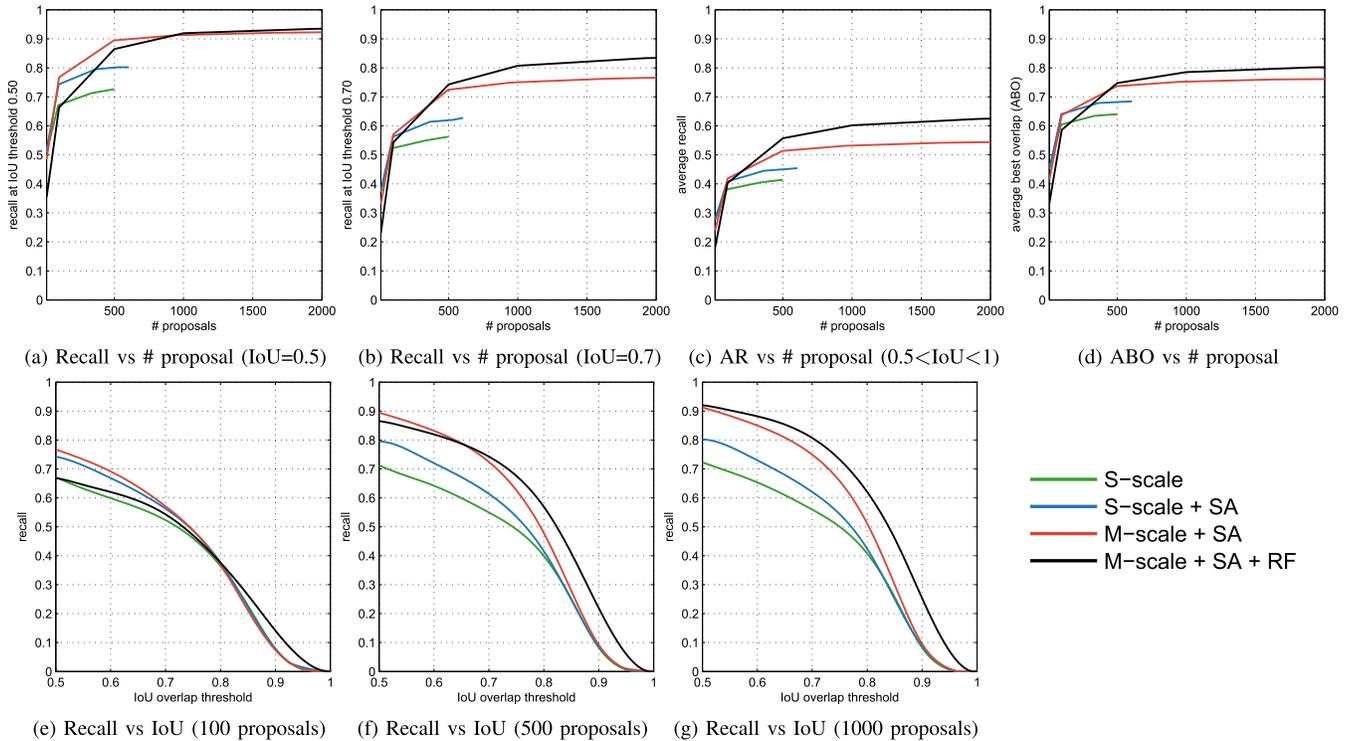


Fig. 8. Recall and average best overlap (ABO) comparison between different variants. S-scale, S-scale+SA, M-scale+SA, M-scale+SA+RF denote single-scale, single scale with scale-awareness, multi-scales with scale-awareness, multi-scales with scale-aware and refinement, respectively. “SA” and “RF” denote “scale-awareness” and “refinement”, respectively. (a) Recall vs # proposal (IoU=0.5). (b) Recall vs # proposal (IoU=0.7). (c) AR vs # proposal ($0.5 < \text{IoU} < 1$). (d) ABO vs # proposal. (e) Recall vs IoU (100 proposals). (f) Recall vs IoU (500 proposals). (g) Recall vs IoU (1000 proposals).

large objects and those containing less than 2,000 pixels as small ones. Considering the unbalanced pixel samples when training the large-/small-size weighting branch, for each large object, we randomly sample 100 pixels in it for training to balance the number of training pixels belonging to large and small objects. Both the “confidence network” and the two localization networks are trained using the published DeepLab code [17], which is based on the publicly available Deep Learning platform Caffe [36]. The weights in the newly added layers are all initialized with a zero-mean Gaussian distribution with the standard deviation 0.01 and the biases are initialized with 0. The initial learning rate is 0.001 for the pre-trained layers in the DeepLab-LargeFOV network and 0.01 for the newly-added layers. All of them are reduced by a scale of 10 after every 20 epochs. The mini-batch size is set as 8. We train the network for about 60 epochs. The overlap threshold for NMS in our experiments is set to 0.8 for a good trade-off between the recall at low IoU thresholds (*e.g.* 0.5) and high IoU thresholds (*e.g.* 0.8). The training images are all resized to 513*513. During testing, for original scale, all the images are directly fed into the networks without any scaling; for enlarged scale, all the images are enlarged by a factor of 2.

The proposed SPOP-net is then extensively evaluated on PASCAL VOC 2007 testing set which is the most widely used in comparison of object proposal algorithms. It contains 4,952 images with annotated objects (including “hard” objects) in bounding boxes. We are not able to evaluate on PASCAL VOC 2012 testing set because the ground-truths are not publicly released. Since the missed objects can never be recovered in the post-classification stage in a

proposal-based object detection pipeline, object recall rate is naturally regarded as the standard evaluation metric for object proposal algorithms. Also, we evaluate the localization quality measured by Average Best Overlap (ABO). In addition, the object detection performance using our proposals in Fast-Regions with CNN features (RCNN) [3] detection pipeline is evaluated to validate the effectiveness of our proposals in the object detection task. Finally, we conduct the generalization ability evaluation by testing the recall rate on ILSVRC 2013 validation set using our network which is trained on PASCAL VOC 2012.

B. Ablation Studies

We first study the effectiveness of the four components in our method: pixel-wise localization network (basic setting), scale-aware localization, multi-scale inference and proposal refinement. Several simplified variants of the SPOP-net are tested in terms of the object recall rate on PASCAL VOC 2007 testing set. Specifically, we use the prediction only at the original scale without scale-awareness and proposal refinement as our baseline, which is referred to as single scale. Without scale-awareness, only one localization network is trained on all of the foreground pixels including both large-size and small-size ones. Then, we accumulatively add scale-awareness, multi-scale inference, proposal refinement to the baseline to see the benefits of each component. Please note that multi-scale inference here indicates the prediction at two scales, namely the original image scale and the 2-time enlarged scale.

Figure 8 shows the recall and average best overlap (ABO) comparisons under different scenarios between the four

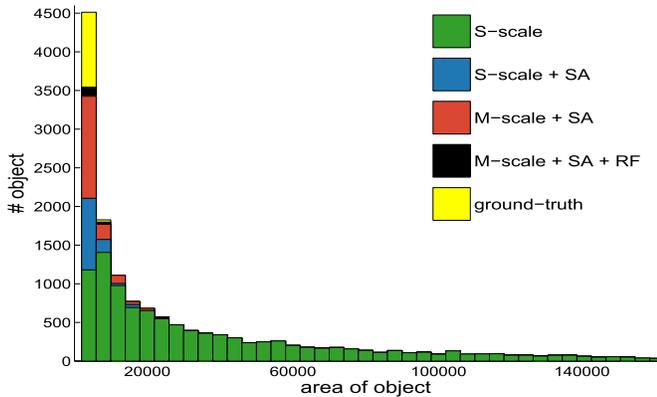


Fig. 9. Distribution of the detected objects w.r.t. the object areas (measured by number of contained pixels) on the PASCAL VOC 2007 testing set of the four variants of the SPOP-net. The IoU threshold is 0.5. 2,000 proposals are generated for each image.

variants, *i.e.* single scale, single scale with scale-awareness, multi-scales with scale-awareness, multi-scales with scale-awareness and refinement. The number of proposals of S-scale and S-scale+SA are around 500 due to that most proposals can be filtered after NMS as pixel-wise localization networks generate highly overlapped proposals (see Figure 14). From Figure 8(a), 8(b) and 8(c), 8(e), 8(f) and 8(g), we find that both scale-awareness and multi-scale inference improve the recall under both low IoU threshold (*e.g.* 0.5) and high IoU threshold (*e.g.* 0.7). As for proposal refinement, it is found that it harms the recall under low IoU thresholds (*e.g.* 0.5) when the number of proposals is less than 500. The reason probably lies in the large number of proposals after refinement, which is 3 times as big as that before refinement. Although this increases the opportunities of getting close to the ground-truths which can boost the recall for a large number of proposals, this also causes too many duplicate proposals to concentrate on a small area, which lowers down the recall under loose IoU criteria when only requiring a small number of proposals. For average best overlap, it shows a similar trend to the recall from Figure 8(d), suggesting the benefits of all three components in terms of localization quality.

We then study the contributions of all the components for different object areas. Figure 9 presents the distributions of the detected objects of both the four variants of SPOP-net and the ground-truths w.r.t. the object areas. It is found that the baseline variant, *i.e.* single scale without scale-awareness and refinement, can hit most of big objects but performs poor for small objects. Scale-aware weighted combination mechanism and multi-scale inference help improve the recall for small objects significantly, which shows the effectiveness of both the proposed scale-aware localization strategy and multi-scale inference.

To further verify the effectiveness of scale-awareness and multi-scale inference in small objects localization, we break up the SPOP-net into four building blocks, *i.e.* large-size network and small-size network in original scale, and large-size network and small-size network in enlarged scale, in order to investigate their respective contributions to the final localization. We evaluate the average best overlap (ABO) of the four building blocks for the ground-truth objects with

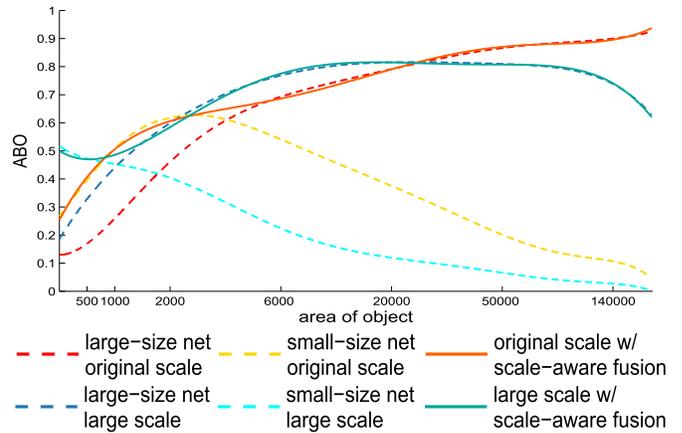


Fig. 10. Average best overlap (ABO) versus ground-truth object area for the four building blocks localization results: large-size network in original scale, small-size network in original scale, large-size network in enlarged scale and small-size network in enlarged scale. All the ABOs are computed given the top 1,000 proposals per image.

different areas. Figure 10 shows the ABO versus object area curves of the four building blocks. It can be seen that when the object becomes larger, the large-size network in original scale predicts more accurate localization results. The small-size network in original scale achieves the highest ABO when the object area is around 2,000, but it also performs poorly for those too small objects. Fortunately, the small-size network in enlarged scale covers this shortage, and gives the best performance for very small objects due to the enlarged view of small objects. As for the large-size network in enlarged scale, it performs the best for those middle-size objects containing 2,000 to 20,000 pixels, serving as the bridge between the large-size network in original scale and the small-size networks in both scales. The reason for the behavior of the large-size network in enlarged scale is probably that when the small objects are enlarged, they become “large objects” such that it becomes easier for the large-size network to predict, but original large objects become even larger which cannot be covered by the receptive field, making it difficult to precisely localize them. In both original scale and enlarged scale, the result after *scale-aware* fusion can achieve the maximal ABO among the two ABOs obtained by large-size and small-size networks, validating the effectiveness of the adaptive *scale-aware* fusion strategy.

By investigating the building blocks of the proposed SPOP-net, it is found that they can complement each other in localizing the objects with different areas and ensures the SPOP-net to perform well for a wide range of object sizes.

C. Comparisons on Object Recall

We compare our SPOP-net with the following state-of-the-art object proposal methods: BING [5], Edge Boxes [6], Geodesic Object Proposal [21], MCG [11], Objectness [20], Selective Search [8] and Region Proposal Network (RPN with VGG-16) [13]. We first evaluate object recall on PASCAL VOC 2007 testing set, which contains 4,952 images with about 15,000 annotated objects. Proposals of most state-of-the-art methods were provided by Hosang *et al.* [37] in a

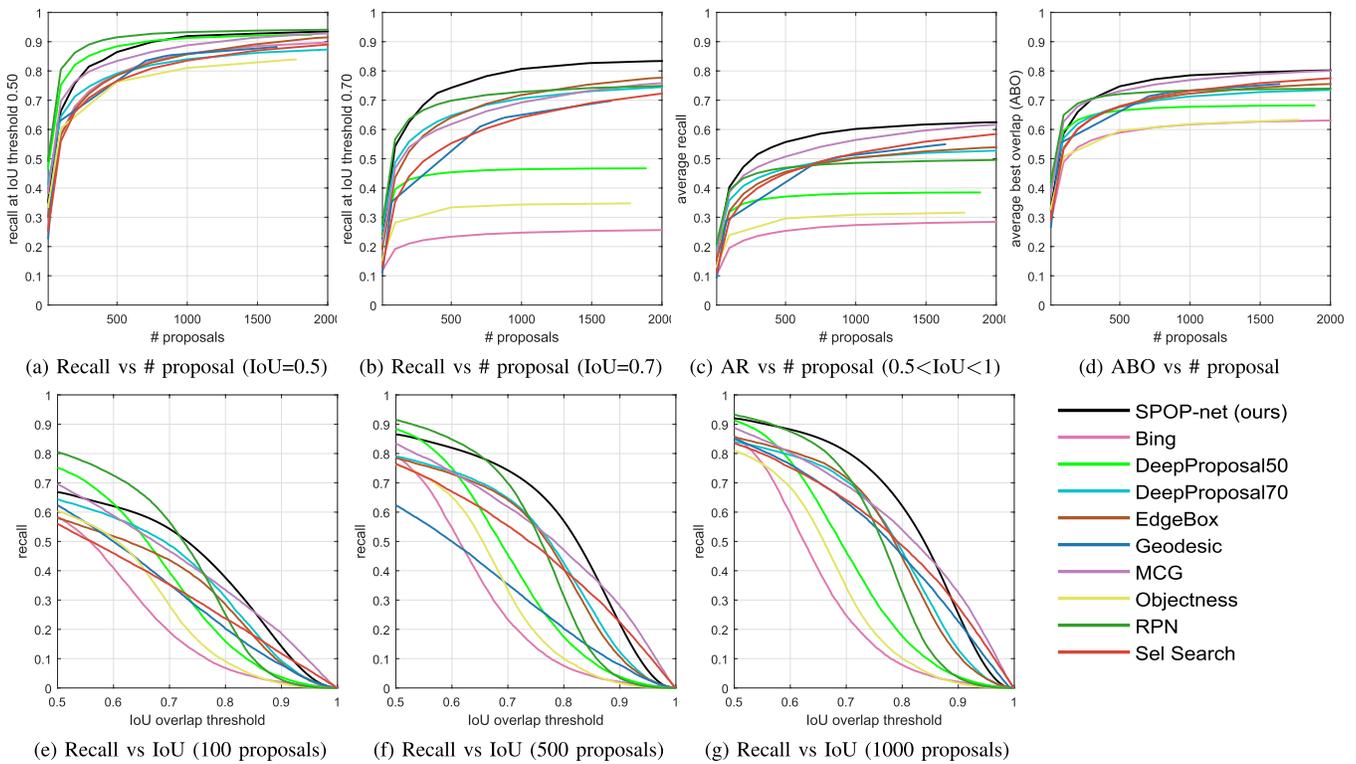


Fig. 11. Recall and average best overlap (ABO) comparison between our SPOP-net and other state-of-the-arts on PASCAL VOC 2007 testing set. (a) Recall vs # proposal (IoU=0.5). (b) Recall vs # proposal (IoU=0.7). (c) AR vs # proposal ($0.5 < \text{IoU} < 1$). (d) ABO vs # proposal. (e) Recall vs IoU (100 proposals). (f) Recall vs IoU (500 proposals). (g) Recall vs IoU (1000 proposals).

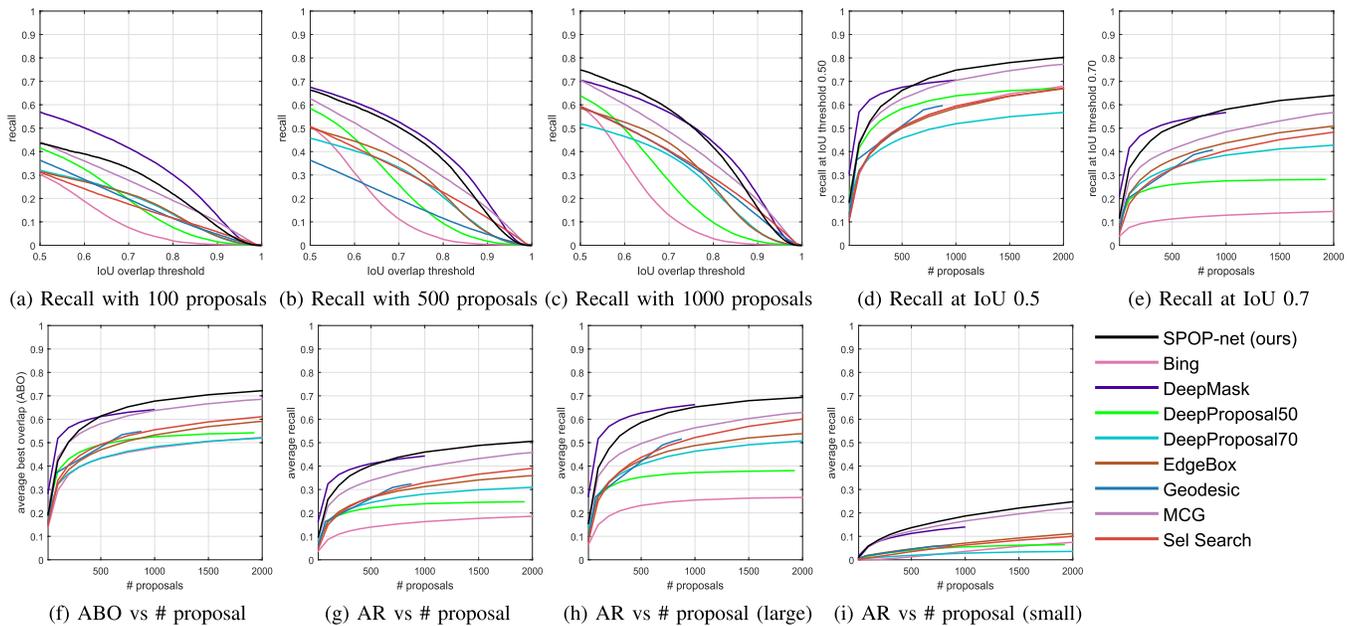


Fig. 12. Recall and average best overlap (ABO) comparison between our SPOP-net and other state-of-the-arts on MS COCO 2014 validation set. (a) Recall with 100 proposals. (b) Recall with 500 proposals. (c) Recall with 1000 proposals. (d) Recall at IoU 0.5. (e) Recall at IoU 0.7. (f) ABO vs # proposal. (g) AR vs # proposal. (h) AR vs # proposal (large). (i) AR vs # proposal (small).

standard format. As for DeepProposal approach, we directly downloaded the pre-computed proposals from the official website.⁵

Figure 11(a) and 11(b) show the recall when varying the number of proposals for different IoU thresholds. As can

be seen, under a loose 0.5 IoU threshold, RPN takes the lead all the time for both a small and a large number of proposals. DeepProposal 50 also performs well under low IoU thresholds (e.g. 0.5). Given a more strict IoU threshold 0.7, our SPOP-net almost keeps the best consistently. We also plot the average recall (AR) versus the number of proposals curves for all the methods in Figure 11(c). This is because

⁵<https://github.com/aghodrati/deepproposal>

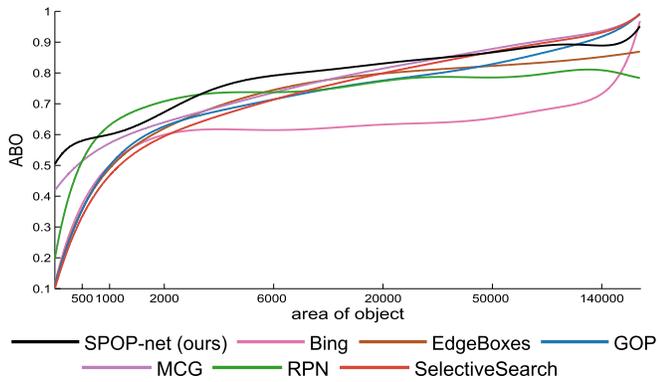


Fig. 13. Average best overlap (ABO) versus ground-truth object area for the SPOP-net and other state-of-the-art methods. All the ABO are computed given the top 1,000 proposals per image.

AR summarizes proposal performance across IoU thresholds and correlates well with object detection performance [37]. The proposed SPOP-net also takes the first place all the time regarding the number of proposals. Figure 11(d) shows the average best overlap (ABO) when changing the number of proposals. The proposed SPOP-net shows good localization quality, especially when the number of proposals is more than 500. Figure 11(e), 11(f) and 11(g) demonstrate the recall when the IoU threshold changes within the range [0.5, 1] for different numbers of proposals. It is found that RPN performs well with a small number of proposals when setting a low IoU threshold (< 0.7). When increasing the number of proposals from 100 to 1,000, our SPOP-net gradually shows its advantage. Especially for the top 1,000 proposals, the SPOP-net performs superiorly across a wide range of IoU thresholds from 0.5 to 0.85, which have the strongest correlation to object detection performance thus are typically desired in practice [37].

Figure 13 shows the average best overlap (ABO) of the proposed SPOP-net as well as several state-of-the-art methods for the ground-truth objects with different areas. For most object sizes, the SPOP-net shows outstanding performance. Especially for small objects whose area is less than about 1,000, the SPOP-net takes the first place, achieving an ABO higher than 0.5. RPN can achieve a good ABO around 0.7 for the objects whose areas are more than 2,000 pixels, but can hardly reach a higher ABO even if the object is large. This may explain why the recall of RPN is very high when setting a loose IoU threshold (e.g. 0.5) but decreases sharply with the increasing of IoU threshold when it exceeds 0.7. The classic low-level cues based methods (e.g. Selective Search, MCG, GOP) perform very well for large objects but have inferior performance for small ones compared with two CNN-based methods (i.e. SPOP-net, RPN).

For better understanding of the keys of enabling the SPOP-net to work well, we show the intermediate output maps of both the localization and confidence networks for visualization in Figure 14. For each image, we show its “objectness confidence map”, “offsets map” pointing to the object center, and its proposals. We argue that the first key is the reliable objectness prediction as the proposals predicted by the pixels obtaining low objectness confidence will be ranked behind.

Based on an accurate objectness confidence, for each ground-truth object, each pixel inside it predicts its own location of this object, as shown in the “offsets maps”, thus greatly increasing the chances of precise localization. Another advantage of this pixel-wise prediction is that most of predicted bounding box locations from the pixels within the same object are heavily overlapping, which can be easily filtered by NMS. Normally only a few proposals are remained after NMS, thus improving the recall of the top-ranked proposals. For small objects, to overcome the inherent weakness that less chances are available to propose the correct locations, a scale-aware prediction is adopted by relying on an accurate estimation of the object size (i.e. large or small) and combining the predictions of two networks.

The detailed running speed of the SPOP-net as well as other state-of-the-art methods is presented in Table IV. The detailed setting of parameters for each method is as follows. We choose the single color space (i.e. RGB) proposal computation for BING, and the “Fast” version for selective search. For the rest methods, we directly run their default codes. As can be seen, *window scoring methods* and *CNN-based methods* are much faster than *segment grouping methods*. Inference for an image of PASCAL VOC size (e.g. 300*500) takes 1.03s for our SPOP-net on a single TITAN X CPU. Specifically, testing one network of the original scale and the enlarged scale takes 0.11s and 0.23s on a single TITAN X GPU, respectively. However, as the computation within different CNNs of SPOP-net are independent of each other, training and testing SPOP-net can be accelerated by parallel computation over multiple GPUs. Although it is not one of the fastest object proposal methods (compared to BING, RPN and Edge Boxes), our approach is still competitive in speed among the proposal generators. We do, however, require use of the library Caffe [36] which is based on GPU computation for efficient inference like all CNN based methods. To further reduce the running time, some CNN speedup methods such as FFT, batch parallelization, or truncated SVD could be used in the future.

We also evaluate the proposed SPOP-net on MS COCO [38] 2014 validation set and the results are shown in Figure 12. The SPOP-net model here is trained on MS COCO training set which contains more than 80,000 pixel-level annotated images. To conduct fair comparisons with the state-of-the-art segmentation annotations based approach, i.e., DeepMask, we only evaluate on the first 5,000 images. Note that we directly used the public DeepProposal model to extract proposals on MS COCO images. It is observed that DeepMask performs well, especially for the cases with low IoU thresholds (e.g. 0.5) and a few proposals (e.g. 100 proposals). The performance of the proposed SPOP-net gradually increases and SPOP-net demonstrates its superiority as the number of proposals increases. Specifically, SPOP-net outperforms DeepMask in terms of recall at IoU 0.5 (Figure 12(d)), recall at IoU 0.7 (Figure 12(e)), ABO (Figure 12(f)) and average recall (Figure 12(g)) when the number of proposals is more than 500. Figure 12(h) and Figure 12(i) shows the average recall of all the methods on large and small objects, respectively. One can observe that SPOP-net performs best on

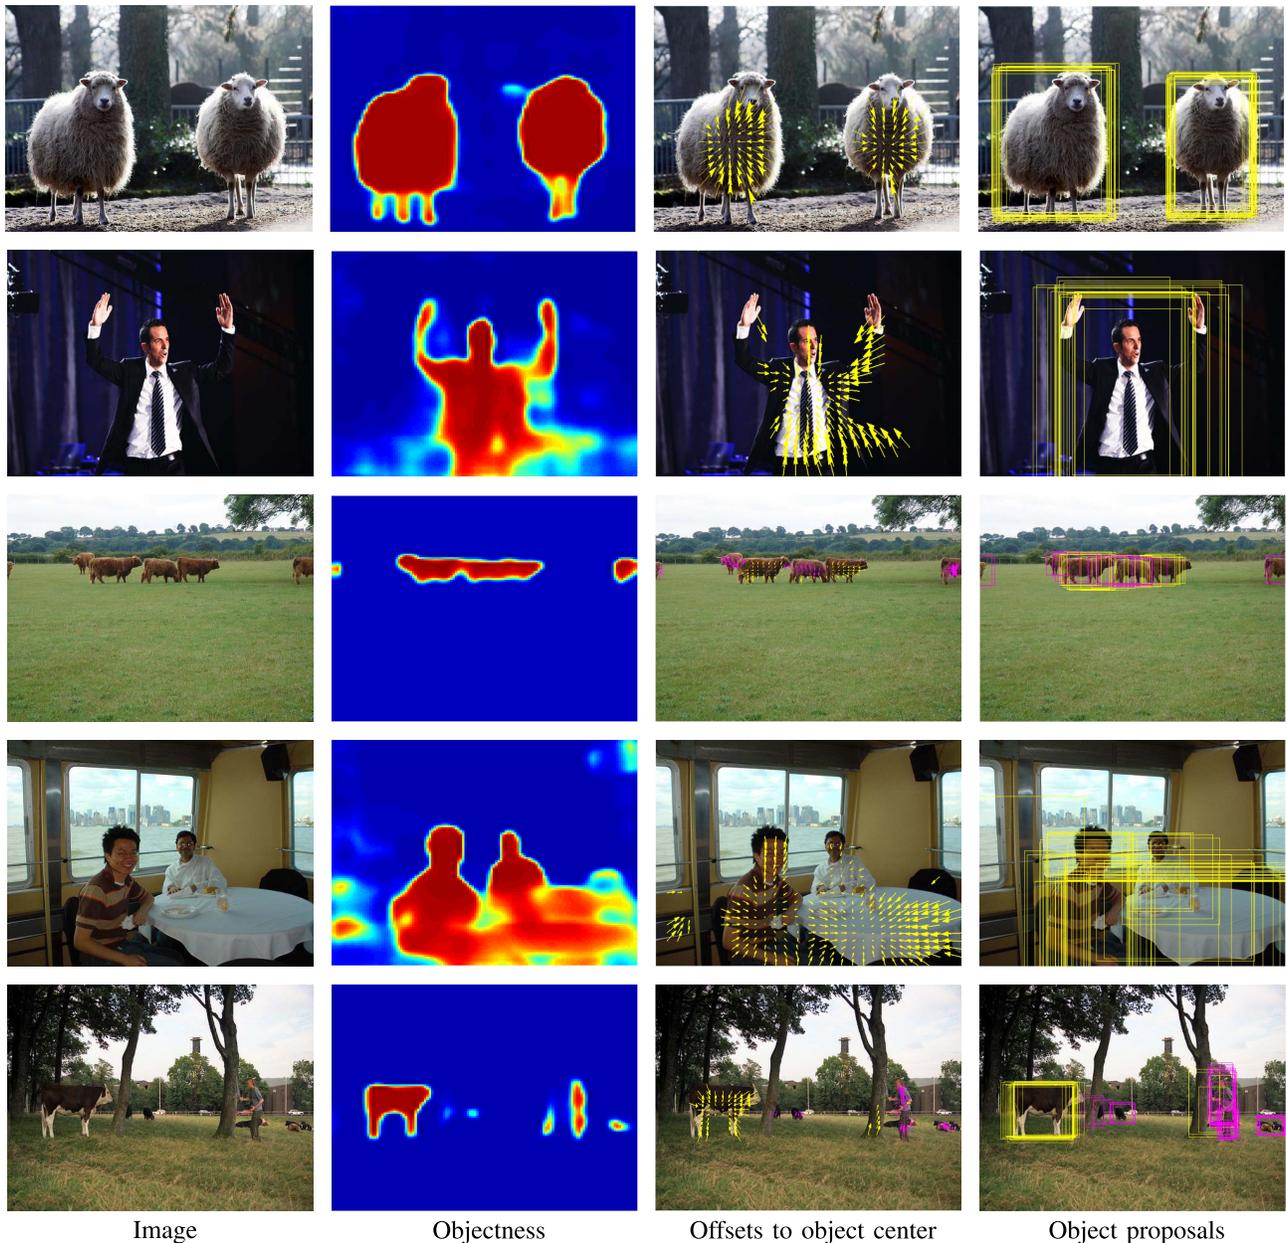


Fig. 14. Example results of predicted “objectness map” (second column), “offsets to object center” after weighted combination (third column) and “object proposals” (fourth column) for the input images (first column).

TABLE IV
TIME COST OF THE STATE-OF-THE-ARTS AND OUR METHOD

	Time cost per image
BING	0.01s
Edge Boxes	0.3s
Geodesic	1s
MCG	30s
Objectness	3s
Selective Search	10s
RPN	0.15s
SPOP-net (ours)	1.03s

detecting small objects in terms of AR, which clearly validates the superiority of SPOP-net in small objects localization.

D. Object Detection Performance

We conduct experiments analyzing object proposals for use with object detectors to evaluate the effects of proposals on

the detection quality. We utilize the standard Fast-RCNN [3] framework as the benchmark. We choose the publicly released VGG 16-layer [39] detector trained on VOC 2007 trainval set in all the evaluation experiments. The proposals of the proposed SPOP-net, Selective Search, Edge Boxes, MCG and RPN are evaluated. Please note that RPN itself integrates proposal generation and detection in a unified framework, called Faster-RCNN. To be fair, we do not adopt this unified detector for object detection with RPN proposals in our evaluation. This is because this unified detector has a weights sharing mechanism in 13 layers which are used for both proposal generation and object detection. These layers are trained on the class-specific annotations with object category information that is not employed in training other methods. For SPOP-net, Selective Search, Edge Boxes and MCG, we select the top 1,000 proposals to pass through the object detectors for

TABLE V

OBJECT DETECTION AVERAGE PRECISION FOR ALL THE 20 CATEGORIES AS WELL AS THE MEAN AVERAGE PRECISION (mAP) ON THE PASCAL VOC 2007 TESTING SET USING THE PUBLICLY AVAILABLE FAST-RCNN DETECTOR TRAINED ON VOC 2007 TRAINVAL SET

	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
Selective Search	76.1	77.3	65.3	53.9	37.8	76.9	78.2	80.9	40.6	74.0	67.2	79.4	82.4	74.9	66.1	33.3	66.0	67.3	73.3	67.3	66.9
Edge Boxes	62.8	77.3	66.2	53.6	42.9	80.6	77.7	81.5	41.4	73.5	65.3	78.1	79.5	76.2	67.8	36.7	64.5	62.4	70.3	67.9	66.3
MCG	69.3	72.3	62.4	54.4	39.2	77.8	70.1	80.4	40.1	67.2	68.7	77.3	75.0	68.8	60.7	34.1	59.5	64.7	70.6	68.2	64.0
RPN (1,000 props)	70.0	76.6	67.2	59.1	44.6	80.0	78.6	86.2	44.1	75.5	60.7	81.3	80.4	75.8	74.1	30.5	72.9	67.2	79.4	69.1	68.7
RPN (300 props)	71.8	77.4	68.0	58.9	46.3	81.8	79.0	86.6	45.6	79.4	60.2	81.7	81.1	75.9	74.5	31.7	73.6	67.2	79.5	70.6	69.5
SPOP-net (ours)	70.6	78.5	69.3	62.5	41.1	82.8	79.1	88.6	47.7	76.6	66.5	83.7	83.6	73.1	69.6	36.1	67.8	72.1	85.4	69.6	70.2

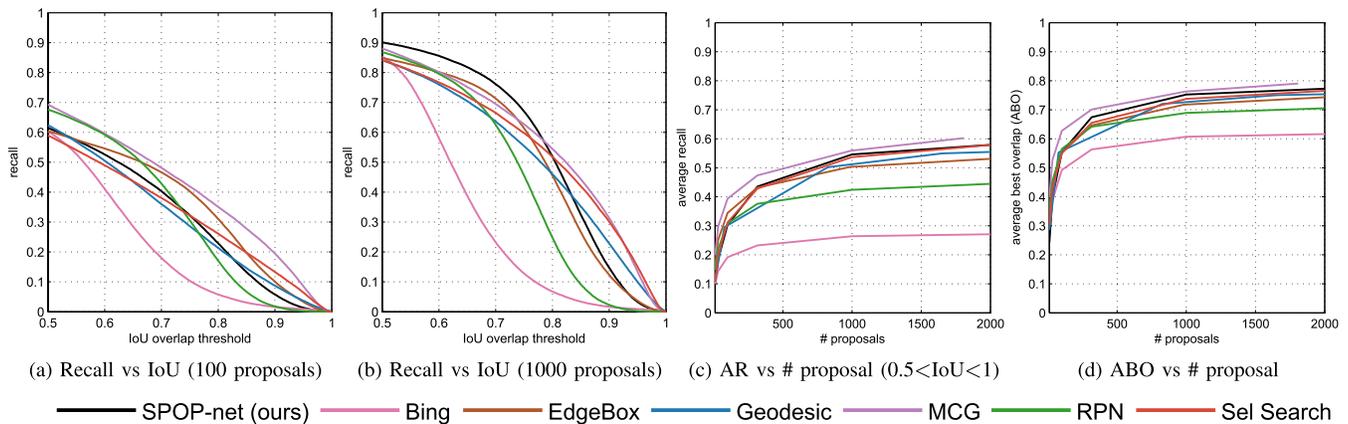


Fig. 15. Recall and average best overlap (ABO) comparison between our SPOP-net and other state-of-the-art methods on ILSVRC 2013 validation set. (a) Recall vs IoU (100 proposals). (b) Recall vs IoU (1000 proposals). (c) AR vs # proposal ($0.5 < \text{IoU} < 1$). (d) ABO vs # proposal.

post-classification. For the RPN method, considering that it only needs a small number of proposals to achieve high recall, and more proposals do not bring too many improvements to the recall but introduce more false positives, we conduct an extra setting which uses the top 300 proposals for detection, which is also claimed by [13].

The detection mean average precision (mAP) as well as the average precision of 20 categories is presented in Table V. It can be seen that the proposed SPOP-net wins on 11 categories among the 20 categories of PASCAL VOC 2007 and also achieves the best mAP 70.2%. Using 1,000 RPN proposals for detection, 68.7% mAP can be obtained. With only 300 proposals, RPN achieves a better mAP 69.5% than 1,000 proposals. This verifies the good performance of RPN when generating a small number of proposals.

E. Generalization to Unseen Categories

The high recall rate which our approach achieves on the PASCAL VOC 2007 testing set does not guarantee it to have learned the generic objectness notion or be able to predict the object proposals for the images containing novel objects in unseen categories. This is because it is possible that the model is highly tuned to the 20 categories of PASCAL VOC. To investigate whether it is capable of predicting the proposals

for the unseen categories beyond training, we evaluate our approach on the ImageNet ILSVRC 2013 validation set which contains more than 20,000 images with around 50,000 annotated objects in 200 categories.

From Figure 15, the overall trend of the SPOP-net remains consistent with that on the PASCAL VOC 2007. Specifically, with a small number of proposals (e.g. 100 proposals), the SPOP-net does not perform as well as MCG, RPN and Edge Boxes, but shows its superiority when the number of proposals reaches 1,000. See Figure 15(b). As for average recall (AR) and average best overlap (ABO), the SPOP-net is also one of the best methods across a broad range of proposal numbers. It is worth mentioning that RPN does not perform as well as on PASCAL VOC 2007. An obvious drop is seen under all the evaluation scenarios from Figure 15 compared to Figure 11. This may result from the category information employed when training the layers in the RPN network shared with class-specific detectors. Such class-awareness enables RPN to fit the 20 categories of PASCAL VOC 2007 better but affects its generalization ability to unseen categories.

Based on the high recall rate the SPOP-net remains when evaluated on ILSVRC 2013, no significant overfitting towards the PASCAL VOC categories is observed. In other words, the proposed approach has learned a generic

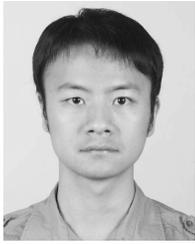
notion of objectness and can generalize well to the unseen categories.

V. SUMMARY AND CONCLUSIONS

In this paper, we developed an effective scale-aware pixel-wise localization network for object proposal generation. The network fully exploits the available pixel-wise segmentation annotations and predicts the proposals pixel-wisely. Each proposal combines two proposals predicted by two networks specialized for different sizes respectively. The combination follows a weighting mechanism utilizing the weighting confidence produced by a large-/small-size object classification model. This strategy is shown to enhance the accuracy of localization on small objects. Significant improvements over the state-of-the-art methods were achieved by the proposed SPOP-net on the PASCAL VOC 2007 testing set. The proposals of the SPOP-net used in Fast-RCNN detector also provide the highest mAP, benefiting from the high recall rate of the proposed model. In the future, we plan to extend our method to deal with both object proposal generation and bounding box regression step to achieve better localization performance.

REFERENCES

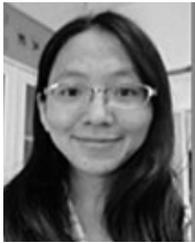
- [1] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 580–587.
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 346–361.
- [3] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1440–1448.
- [4] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 9, pp. 1627–1645, Sep. 2010.
- [5] M.-M. Cheng, Z. Zhang, W.-Y. Lin, and P. Torr, "BING: Binarized normed gradients for objectness estimation at 300fps," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 3286–3293.
- [6] C. L. Zitnick and P. Dollár, "Edge boxes: Locating object proposals from edges," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 391–405.
- [7] B. Alexe, T. Deselaers, and V. Ferrari, "What is an object?" in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2010, pp. 73–80.
- [8] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders, "Selective search for object recognition," *Int. J. Comput. Vis.*, vol. 104, no. 2, pp. 154–171, Apr. 2013.
- [9] S. Manen, M. Guillaumin, and L. Van Gool, "Prime object proposals with randomized Prim's algorithm," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2013, pp. 2536–2543.
- [10] Z. Zhang, J. Warrell, and P. H. S. Torr, "Proposal generation for object detection using cascaded ranking SVMs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2011, pp. 1497–1504.
- [11] P. Arbeláez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik, "Multi-scale combinatorial grouping," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 328–335.
- [12] D. Erhan, C. Szegedy, A. Toshev, and D. Anguelov, "Scalable object detection using deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 2155–2162.
- [13] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [14] Z. Jie, W. F. Lu, S. Sakhavi, Y. Wei, E. H. F. Tay, and S. Yan, "Object proposal generation with fully convolutional networks," *IEEE Trans. Circuits Syst. Video Technol.*, vol. PP, no. 99, to be published.
- [15] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL visual object classes challenge: A retrospective," *Int. J. Comput. Vis.*, vol. 111, no. 1, pp. 98–136, Jan. 2014.
- [16] O. Russakovsky *et al.*, "ImageNet large scale visual recognition challenge," *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, Dec. 2015.
- [17] 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 *Proc. Int. Conf. Learn. Represent.*, 2015. [Online]. Available: <https://arxiv.org/abs/1412.7062>
- [18] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 3431–3440.
- [19] X. Liang, Y. Wei, X. Shen, J. Yang, L. Lin, and S. Yan. (2015). "Proposal-free network for instance-level object segmentation." [Online]. Available: <http://arxiv.org/abs/1509.02636>
- [20] B. Alexe, T. Deselaers, and V. Ferrari, "Measuring the objectness of image windows," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2189–2202, Nov. 2012.
- [21] P. Krähenbühl and V. Koltun, "Geodesic object proposals," in *Computer Vision—ECCV*. Springer, 2014, pp. 725–739.
- [22] P. Krähenbühl and V. Koltun, "Learning to propose objects," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1574–1582.
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [24] Y. Wei *et al.*, "HCP: A flexible CNN framework for multi-label image classification," *IEEE Trans. Pattern Anal. Mach. Intell.*, to be published.
- [25] C. Szegedy *et al.*, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1–9.
- [26] X. Liang, S. Liu, Y. Wei, L. Liu, L. Lin, and S. Yan, "Towards computational baby learning: A weakly-supervised approach for object detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 999–1007.
- [27] Y. Wei *et al.* (2015). "STC: A simple to complex framework for weakly-supervised semantic segmentation," [Online]. Available: <http://arxiv.org/abs/1509.03150>
- [28] Y. Wei *et al.*, "Learning to segment with image-level annotations," *Pattern Recognit.*, to be published.
- [29] X. Liang *et al.*, "Reversible recursive instance-level object segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 633–641.
- [30] A. Ghodrati, A. Diba, M. Pedersoli, T. Tuytelaars, and L. Van Gool, "DeepProposal: Hunting objects by cascading deep convolutional layers," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 2578–2586.
- [31] L. Huang, Y. Yang, Y. Deng, and Y. Yu. (2015). "DenseBox: Unifying landmark localization with end to end object detection." [Online]. Available: <http://arxiv.org/abs/1509.04874>
- [32] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 818–833.
- [33] X. Chen, H. Ma, X. Wang, and Z. Zhao, "Improving object proposals with multi-thresholding straddling expansion," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 2587–2595.
- [34] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "SLIC superpixels compared to state-of-the-art superpixel methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 11, pp. 2274–2282, Nov. 2012.
- [35] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Malik, "Semantic contours from inverse detectors," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Nov. 2011, pp. 991–998.
- [36] Y. Jia *et al.*, "Caffe: Convolutional architecture for fast feature embedding," in *Proc. ACM Int. Conf. Multimedia*, 2014, pp. 675–678.
- [37] J. Hosang, R. Benenson, P. Dollár, and B. Schiele, "What makes for effective detection proposals?" *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 4, pp. 814–830, Apr. 2016.
- [38] T.-Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Computer Vision—ECCV*. Springer, 2014, pp. 740–755.
- [39] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Represent.*, 2015. [Online]. Available: <https://arxiv.org/abs/1409.1556>



Zequn Jie received the B.E. degree in mechanical engineering from the University of Science and Technology of China. He is currently pursuing the Ph.D. degree from the Vision and Machine Learning Group, directed by Prof. S. Yan and J. Feng of National University of Singapore. His current research interests mainly include object localization related topics in computer vision, such as object proposal, object detection.



Wen Feng Lu received the Ph.D. degree in mechanical engineering from the University of Minnesota, USA. He was a Faculty Member with the University of Missouri, USA, for ten years. He was also a Group Manager and a Senior Scientist with the Singapore Institute of Manufacturing Technology for six years. He is currently an Associate Professor with the Department of Mechanical Engineering, National University of Singapore (NUS). His research interests include IT in product design, sustainable design and manufacturing, 3-D printing, and intelligent manufacturing. He was a recipient of 1997 Ralph R. Teeter Educational Award from the Society of Automotive Engineers of USA and 2011 ASME Virtual Environments and Systems Technical Committee Best Paper Award.



Xiaodan Liang received the Ph.D. degree from Sun Yat-sen University, China. She is currently a Post-Doctoral Research Fellow with Carnegie Mellon University, under the supervision of Prof. E. Xing. Her research interests mainly include semantic segmentation, object/action recognition, and medical image analysis.



Eng Hock Francis Tay is currently an Associate Professor with the Department of Mechanical Engineering, Faculty of Engineering, National University of Singapore. He was also the Founding Director of the Microsystems Technology Initiative, and established the Microsystems Technology Specialization. He is currently the Deputy Director (Industry) for the Centre of Intelligent Products and Manufacturing Systems, where he takes charge of research projects involving industry and the Centre.



Jiashi Feng received the bachelor's degree in automation from the University of Science and Technology of China, Hefei, China, and the Ph.D. degree from the National University of Singapore (NUS), Singapore, in 2014. He was a Post-Doctoral Researcher with the University of California at Berkeley, Berkeley, CA, USA. He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, NUS. He is interested in computer vision and machine learning.



Shuicheng Yan is currently an Associate Professor with the National University of Singapore, and the Founding Lead of the Learning and Vision Research Group. He has authored/co-authored over 400 technical papers over a wide range of research topics, with Google Scholar citation over 12000 times. His research areas include machine learning, computer vision, and multimedia. He is ISI highly cited researcher 2014, and IAPR Fellow 2014. He received the best paper awards from ACM MM'13 (Best Paper and Best Student Paper), ACM MM'12 (Best Demo), PCM'11, ACM MM'10, ICME'10, and ICIMCS'09, the runner-up prize of ILSVRC'13, the winner prizes of the classification task in PASCAL VOC 2010-2012, the winner prize of the segmentation task in PASCAL VOC 2012, the honorable mention prize of the detection task in PASCAL VOC'10, 2010 TCSVT Best Associate Editor Award, 2010 Young Faculty Research Award, 2011 Singapore Young Scientist Award, and 2012 NUS Young Researcher Award. He has been serving as an Associate Editor of the IEEE TKDE, CVIU, and TCSVT.