

BiSeg: Simultaneous Instance Segmentation and Semantic Segmentation with Fully Convolutional Networks

Viet-Quoc Pham
quocviet.pham@toshiba.co.jp

Satoshi Ito
satoshi13.ito@toshiba.co.jp

Tatsuo Kozakaya
tatsuo.kozakaya@toshiba.co.jp

Corporate Research and Development
Center, Toshiba Corporation
1, Komukai-Toshiba-cho, Saiwai-ku,
Kawasaki, 212-8582, Japan

Abstract

We present a simple and effective framework for simultaneous semantic segmentation and instance segmentation with Fully Convolutional Networks (FCNs). The method, called BiSeg, predicts instance segmentation as a posterior in Bayesian inference, where semantic segmentation is used as a prior. We extend the idea of position-sensitive score maps used in recent methods to a fusion of multiple score maps at different scales and partition modes, and adopt it as a robust likelihood for instance segmentation inference. As both Bayesian inference and map fusion are performed per pixel, BiSeg is a fully convolutional end-to-end solution that inherits all the advantages of FCNs. We demonstrate state-of-the-art instance segmentation accuracy on PASCAL VOC.

1 Introduction

Object detection and semantic segmentation are problems that have been intensively studied in the field of computer vision, and the accuracy of the proposed solutions has improved rapidly owing to the progress of deep learning [6, 10, 18]. In semantic segmentation, the task is to classify each pixel into a fixed set of categories without distinguishing among object instances. While on the other hand, object detection detects object instances at the bounding-box level. In this paper, we address the problem of instance segmentation, which aims to localize objects at the pixel level. Instance segmentation is challenging because it requires both high object detection accuracy and precise segmentation (see Fig. 1).

Early methods [4, 12, 16, 17] address instance segmentation using convolutional neural networks (CNNs), but require mask proposals from external modules [1, 3]. Different from these methods, DeepMask [2], SharpMask [28] and Instance FCN [8] learn to propose segment candidates.

DeepMask [2] is trained with two objectives: predicting a class-agnostic segmentation mask from a given image patch, and predicting how likely the patch is to contain an object. At test time, DeepMask is applied densely to an image and generates a set of object masks, and corresponding objectness scores. SharpMask [28] improves DeepMask by generating




Figure 1: Three important tasks in image recognition: object detection, semantic segmentation and instance segmentation. Instance segmentation requires both high object detection accuracy and precise segmentation.

higher quality masks using an additional top-down refinement step. Its idea is to first generate a coarse mask in a feedforward pass, then refine this mask in a top-down pass using features at successively lower layers.

Instance FCN is a fully convolutional approach for instance mask proposal generation [8]. It extends the translation invariant score maps in conventional FCNs to position-sensitive score maps, which are somewhat translation-variant. On top of these instance-sensitive score maps, an assembling module is able to output instance candidate at each position. The concept of position-sensitive score maps will be discussed in details in Section 2.

The common problem of the above three methods [8, 27, 28] is that they are only used for mask proposal generation, which is blind to semantic categories and requires a downstream network for category classification. In their implementation, the proposed segment candidates are classified by Fast R-CNN [13] in a different process. As these methods adopt a time-consuming image pyramid scanning to find instances at different scales, they are slow at inference time and less accurate.

In a different approach, Dai *et al.* [9] propose an end-to-end solution to address instance segmentation. Their multi-task network cascade (MNC) model consists of three stages: proposing box-level instances, regressing mask-level instances, and categorizing each instance. In the first stage, object instances are proposed in the form of bounding boxes, which are class-agnostic, and are predicted with an objectness score. The second stage takes the shared convolutional features and stage-1 boxes as input, and outputs a pixel-level segmentation mask for each box proposal. The third stage takes the predicted boxes and masks from the previous stages as input, and outputs category scores for each instance.

Recently, Li *et al.* [20] propose a fully convolutional end-to-end solution for instance segmentation, named FCIS, by extending the idea of position-sensitive score maps in [8]. The details will be stated in Section 2. The advantage of this method is that the underlying convolutional representation and the score maps are fully shared for the mask prediction and classification sub-tasks, via a joint formulation with no extra parameters. Furthermore, it operates on box proposals instead of sliding windows, enjoying the recent advances in object detection [29]. FCIS achieved state-of-the-art performance in both PASCAL VOC [17] and MS COCO [23]. It is also $6\times$ faster than MNC [9].

In this paper, we present a simple and effective framework for simultaneous semantic segmentation and instance segmentation with Fully Convolutional Networks (FCNs). The method, called BiSeg, predicts instance segmentation as a posterior in Bayesian inference, where semantic segmentation is used as a prior. We extend the idea of position-sensitive




Figure 2: Inside and outside position-sensitive score maps with 3×3 partitions. The inside/outside likelihood map for each ROI is produced by assembling its $k \times k$ partitions from the corresponding inside/outside score maps.

score maps used in [8, 20] to a fusion of multiple score maps at different scales and partition modes, and adopt it as a robust likelihood for instance segmentation inference.

There are advantages to combining semantic segmentation into the instance segmentation framework. First, semantic segmentation can be used to enhance instance segmentation owing to the strong correlation between the two tasks. In related work, Shrivastava and Gupta [61] add segmentation to Faster R-CNN [29] as a complementary task and use it to provide top-down information to guide region proposal generation and object detection. The intuition is that semantic segmentation captures contextual relationships between objects, and will essentially guide the region proposal module to focus attention in the right areas and learn detectors from them. The second advantage is that semantic segmentation can deal with stuff. Different from objects, stuff [25] (*e.g.*, sky, grass, water) is usually treated as the context in the image. Stuff mostly exhibits as colors or textures and has less well-defined shapes. It is thus inappropriate to use a single rectangular box or a single segment to represent stuff. Combination of object detection from instance segmentation and stuff detection from semantic segmentation is necessary for robust image recognition systems.

Finally, we verify the performance of our method on the PASCAL VOC dataset. BiSeg is a fully convolutional end-to-end solution that inherits all the advantages of FCNs. We achieve the best results, 67.3% and 54.4% which are nearly 2% higher than those of FCIS [20].

2 Preliminaries: FCIS

FCIS [20] is the first fully convolutional end-to-end solution for instance segmentation. It performs instance mask prediction and classification jointly by extending the ideas concerning position-sensitive score maps.

Inside/Outside Position-Sensitive Score Maps: In FCN [24], the method predicts $C + 1$ score maps representing the likelihood of each category at each pixel. Here, C is the number of object categories, 1 is for background. Instance FCN [8] uses k^2 position-sensitive score maps to encode the position information with respect to a relative spatial position from $k \times k$ partitions (*e.g.*, the top left part of an object). Based on the above ideas, FCIS [20] proposes a larger set of $2k^2 \times (C + 1)$ position-sensitive score maps to encode more detailed semantic information at each pixel. For each object category, each pixel has two scores in each partition, representing the likelihoods of being inside (or outside) the object boundary at a relative position (from $k \times k$ partitions). The score maps are illustrated in Fig. 2.

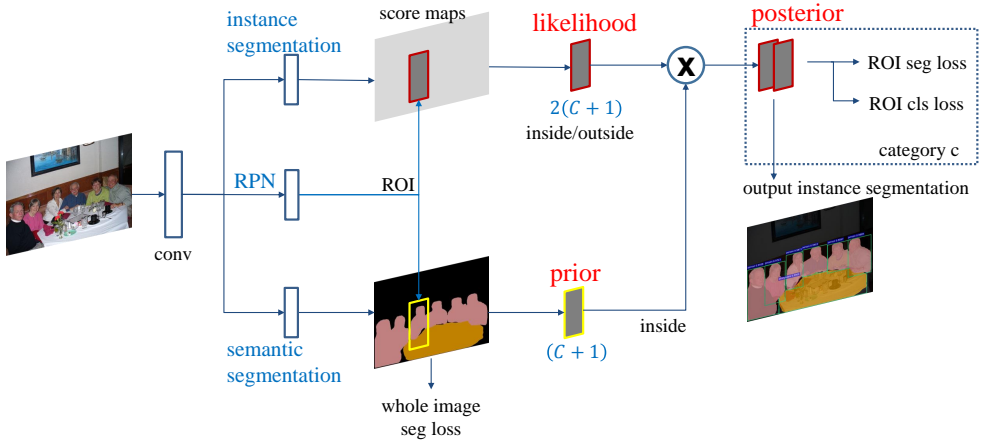


Figure 3: Network architecture of BiSeg. It takes an image of arbitrary size as input and outputs both semantic segmentation and instance segmentation results. RPN, semantic segmentation and instance segmentation sub-networks share the convolutional feature maps. The learnable weight layers are fully convolutional and computed on the whole image. For simplification, we do not show the bounding box regression sub-network here.

Joint Mask Prediction and Classification: For each region-of-interest (ROI), the method produces an inside/outside likelihood map by assembling its $k \times k$ partitions from the corresponding inside/outside score maps. On top of the two likelihood maps, the method applies a softmax operation to produce the foreground probability, and a max operation to produce the per-pixel likelihood of the object category, which is then used to infer the classification score by average pooling.

All the per-ROI components are parameter free. The score maps are computed by FCN, and does not involve any feature warping, resizing or fc layers.

3 Proposed Approach

In our BiSeg model, the network inputs an image of arbitrary size and outputs both semantic segmentation and instance segmentation results. The network has several sub-networks that share convolutional features: region proposal, bounding box regression, semantic segmentation and instance segmentation. The region proposal sub-network generates ROIs, following the work of Region Proposal Networks (RPN) [49]. The bounding box regression sub-network refines the initial input ROIs to produce more accurate detection results, as in [3]. The semantic segmentation sub-network infers the semantic segmentation probabilities. The instance segmentation sub-network estimates the likelihoods of instance segmentation. Finally, the semantic segmentation and instance segmentation sub-networks are combined to infer the instance segmentation probabilities. Our model is illustrated in Fig. 3.

Network Architecture: We use the ResNet model [18] as the convolutional backbone architecture for feature extraction. We add a 1×1 convolutional layer on top of the feature map to reduce the dimension from 2048 to 1024.

Following FCIS [2], we apply the “hole algorithm” [2] to reduce the feature stride of conv5 layers from 32 to 16. The RPN sub-network, which is fully convolutional, is added

on top of the conv4 layers, as in [9, 20].

The semantic segmentation sub-network is added on top of the conv5 feature maps. Our implementation is based on the FCN architecture in [24]. This sub-network produces $C + 1$ score maps S indicating the segmentation probabilities of $C + 1$ categories. To verify the effectiveness of Bayesian inference in our model, we try to keep the model of semantic segmentation as simple as possible. In this paper, we neither consider Conditional Random Field (CRF) [9, 21, 30, 54], higher order potential [9], nor domain transform [6]. The effect of higher semantic segmentation quality on the instance segmentation performance is an interesting topic that we intend to address in future work.

Fusion of Position-Sensitive Score Maps: In the instance segmentation sub-network, we aim to infer instance segmentation likelihoods L for each ROI. We extend the idea of inside/outside position-sensitive score maps in FCIS to a fusion of multiple score maps at different scales and partition modes. Although it was shown in [9] that 7×7 partitions gave the best overall accuracy, we found that other partition modes worked better in some cases. To reduce the dependence on the partition parameter, we make a fusion of multiple score maps with different partition parameters. Furthermore, to enlarge the variation of score maps, we also generate score maps from different scales, and employ the skip architecture in [24] to perform the map fusion.

In our implementation, we generate two sets of position-sensitive score maps from different feature maps. The first set with $2k_1^2 \times (C + 1)$ score maps is created from the conv5 layers by a 1×1 convolutional layer. The second set with $2k_2^2 \times (C + 1)$ score maps is created in the same way from conv3 layers, which is $2 \times$ larger than conv5 ones. In our experiments, we use $(k_1, k_2) = (7, 9)$ by default. Each ROI is projected into a $16 \times$ smaller region over the first set of score maps, and a $8 \times$ smaller region over the second set. Then, for each ROI, we apply the assembling operation on each set of score maps to produce two sets of $2(C + 1)$ ROI likelihood maps. In each set of likelihood maps, the first half consists of outside maps, and the second half consists of inside maps. Finally, we upsample the first set of ROI likelihood maps by $\times 2$, and sum the two sets to get the final likelihood maps. The process is illustrated in Fig. 4.

Bayesian Inference: In the next step, we aim to compute instance segmentation probabilities I for each ROI, which consists of $C + 1$ inside probability maps and $C + 1$ outside probability maps. We first crop the ROI rectangle from the $C + 1$ semantic segmentation score maps S to produce $C + 1$ ROI semantic segmentation probability maps. We then apply the element-wise product operation on these probability maps and the $C + 1$ ROI inside likelihood maps from L . The results are $C + 1$ ROI inside probability maps, as shown in Fig. 3. The intuition of the product operation here is that the ROI inside probability maps are predicted as posteriors in Bayesian inference, where the ROI semantic segmentation probability maps are used as priors, and the ROI inside likelihood maps are used as likelihoods. The ROI outside probability maps are equal to the ROI outside likelihood maps because there are no priors to predict outside probabilities.

To clarify our Bayesian inference formulation, first let $P(I_{ck}|X)$ be the posterior probability of instance segmentation for category c and ROI k , given the input image X . It can be expressed as:

$$P(I_{ck}|X) = \sum_{\text{category } j} P(I_{ck}, j|X). \quad (1)$$

As $P(I_{ck}, j) = 0$ for $j \neq c$, we can rewrite that

$$P(I_{ck}|X) = P(I_{ck}, c|X) = P(I_{ck}|c, X)P(c|X). \quad (2)$$




Figure 4: Fusion of multiple score maps at different scales and partition modes. The first set of score maps with $k_1 \times k_1$ partitions and the second set with $k_2 \times k_2$ partitions are generated from conv5 layers and conv3 layers, respectively. Assembled ROI likelihood maps from the first set are $\times 2$ upscaled and added to those from the second set.

Here we use the position-sensitive score map to approximate $P(I_{ck}|c, X)$, and the semantic segmentation score map to approximate the prior $P(c|X)$. As shown in our experiment results, the semantic segmentation priors significantly enhance the instance segmentation inference owing to the strong correlation between the semantic segmentation probabilities and the instance-aware inside likelihoods.

Following FCIS [20], on top of the inside and outside probability maps, we apply a softmax operation to produce the foreground probability, and a max operation to produce the per-pixel likelihood of the object category, which is then used to infer the classification score by average pooling.

Training: As in Fast R-CNN [13], an ROI is considered positive if it has Intersection over Union (IoU) with a ground-truth box of at least 0.5 and negative otherwise. We define a multi-task loss as follows:

$$L = L_{rpn} + L_{ss} + L_{cls} + L_{mask} + L_{bbox}. \quad (3)$$

The loss of the RPN sub-network L_{rpn} is defined as in Faster R-CNN [29]. The semantic segmentation loss L_{ss} is per-pixel multinomial cross-entropy loss, as in common FCNs [24]. The last three losses in Eq. (3) are defined on each ROI: the softmax classification loss L_{cls} over $C + 1$ categories, the binary cross-entropy loss L_{mask} over the foreground mask of the ground-truth category, and the bounding box regression loss L_{bbox} as in Fast R-CNN [13]. L_{mask} and L_{bbox} are defined only on positive ROIs.

We use the ImageNet pretrained model [8] to initialize the shared convolutional layers. We adopt image-centric training [13]: the shared convolutional layers and the semantic segmentation loss are computed on the entire image, while the ROIs are randomly sampled for computing the other losses. In our implementation, each mini-batch involves 1 image, and each image has 64 sampled ROIs. We use SGD optimization. The model is trained using a

method	mAP ^r @0.5(%)	mAP ^r @0.7(%)
SDS [16]	49.7	25.3
Hypercolumn [17]	60.0	40.4
CFM [9]	60.7	39.6
MPA [23]	61.8	-
MNC [9]	63.5	41.5
MNC, Instance FCN [8]	61.5	43.0
IIS [19]	63.6	43.3
CRF [9]	62.0	44.8
FCIS [20]	65.7	52.1
FCIS*	64.2	48.6
naive Multi-task	65.2	49.6
BiSeg (single PS score map)	66.4	50.5
BiSeg (fused PS score map)	67.3	54.4

Table 1: Comparisons of instance segmentation on the PASCAL VOC 2012 validation set.

learning rate of 0.001 for the first 20k, and 0.0001 for the last 10k iterations respectively. We train the model in one GPU, which holds 8 mini-batches.

Inference: At test time, we generate 300 ROIs from RPN. We run the bounding box regression branch on these proposals, followed by non-maximum suppression. The remaining ROIs are classified as the categories with highest classification scores. Following the mask voting scheme in [9], for each remaining ROI, we find its near-by instances which overlap with it by $\text{IoU} \geq 0.5$. Their foreground masks of the category are averaged pixel-by-pixel, weighted by their classification scores. The averaged masks are binarized to form the final output masks.

4 Experiments

Following previous work [9, 8, 9, 16, 17, 20], we evaluate our instance segmentation performance on the PASCAL VOC dataset [22]. We train the model on the VOC 2012 train set (5623 images), and perform evaluation on the VOC 2012 validation set (5732 images)¹, with the additional instance mask annotation from [15]. The evaluation metric we use is the mean Average Precision over regions, which is referred to as mAP^r [16]. We report mAP^r using IoU thresholds at 0.5 and 0.7. Different from the AP metric used in object detection, the IoU is computed over predicted and ground-truth regions instead of bounding-boxes.

Following FCIS [20], we use ResNet-101 model [18] pre-trained on ImageNet [11], and do not apply online hard example mining (OHEM) [5] on our model. We compare BiSeg with recent methods [9, 17, 8, 9, 16, 17, 19, 20, 23], and variants of our own method with different settings. The results are shown in Tab. 1.

FCIS*: This baseline is based on our own implementation of FCIS [20]². It is equivalent to our model that includes only the instance segmentation sub-network.

¹There is some confusion about the name of the evaluation dataset. For example, the SBD dataset used in [9] is equivalent to the dataset used in our experiments.

²The source code of FCIS has not been released at the time of writing this paper.

(k_1, k_2)	(7, 7)	(7, 9)	(7, 11)
mAP ^r @0.5(%)	66.9	67.3	66.8
mAP ^r @0.7(%)	54.1	54.4	54.4

Table 2: Comparisons of different combinations of partition modes (k_1, k_2) .

method	mean accuracy(%)	mean IU(%)
naive Multi-task	69.0	59.5
BiSeg	70.2	60.8

Table 3: Comparisons of semantic segmentation on the PASCAL VOC 2012 validation set.

naive Multi-task: To verify the importance of the combination of semantic segmentation and instance segmentation sub-networks, this baseline removes the combination module from the top of BiSeg. It can be considered a common multi-task learning method with two branches sharing the feature map.

BiSeg (single PS score map): This version of our model uses the traditional position-sensitive score maps [24] in the instance segmentation sub-network. The mAP^r scores are 66.4% and 50.5% at IoU thresholds of 0.5 and 0.7 respectively. They are 1% higher than **naive Multi-task**, and 2% than **FCIS***. Moreover, our AP^r at 0.5 achieves a 0.7% improvement compared with the previous state-of-the-art FCIS [24]. This verifies the effectiveness of enhancing instance segmentation by semantic segmentation.

BiSeg (fused PS score map): By using the fusion of position-sensitive score maps, we achieve the best results, 67.3% and 54.4%, which are nearly 2% higher than those of FCIS [24]. The AP^r at 0.7 achieves a 4% improvement compared with **BiSeg (single PS score map)**. This proves the importance of the proposed map fusion approach.

For a deeper understanding on the fusion of position-sensitive score maps, we evaluate BiSeg with different combinations of partition modes (k_1, k_2) . As stated in previous work [8], position-sensitive score maps with 7×7 partitions gave the best results. Therefore, we fix k_1 to 7, and test different values of $k_2 \geq 7$. As shown in Tab. 2, fusion of score maps with different partition modes $(k_1, k_2) = (7, 9)$ is more effective than score maps with the same partition modes $(k_1, k_2) = (7, 7)$. However, k_2 with a too large value ($k_2 = 11$) causes a decrease in mAP^r@0.5, due to the over-partition of the position-sensitive score maps.

We also evaluate the performance of semantic segmentation, as shown in Tab. 3. We report here the mean accuracy and the mean region intersection over union (mean IU). Our model outperforms **naive Multi-task**, which learns semantic segmentation and instance segmentation in separated branches. Once again, it verifies the importance of combining the two segmentation tasks in our model.

Finally, we show some of our results for both instance segmentation and semantic segmentation in Fig. 5. Our method produces high segmentation quality at borderlines without using CRF. Moreover, we confirm good instance segmentation results even for heavily occluded objects. Our approach can be applied to the problem of counting objects from crowded scenes [26].



Figure 5: Results of BiSeg. Our method produces high segmentation quality at borderlines without using CRF. For instance segmentation, we show all detected instances whose classification scores are larger than 0.5. We confirm good results of instance segmentation even for heavily occluded scenes (the last two rows).

5 Conclusion

We present a fully convolutional end-to-end solution for simultaneous semantic segmentation and instance segmentation. Our method predicts the instance segmentation mask as a posterior probability in Bayesian inference, where the semantic segmentation result is treated as prior, and fusion of multiple position-sensitive score maps at different scales and partition modes as likelihood. We achieve the best results, 67.3% and 54.4%, on the PASCAL VOC dataset, which are nearly 2% higher than those of FCIS. In future work, we intend to investigate other variants of context priors such as depth maps and instance-aware likelihoods such as human pose.

Acknowledgement

We would like to thank all lab members, in particular Susumu Kubota and Yuta Shirakawa for the helpful suggestions on making this paper better.

References

- [1] P. Arbelaez, J. Pont-Tuset, J. T. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. In *CVPR*, 2014.
- [2] A. Arnab and P.H.S. Torr. Pixelwise instance segmentation with a dynamically instantiated network. In *CVPR*, 2017. <https://arxiv.org/pdf/1704.02386.pdf>.
- [3] A. Arnab, S. Jayasumana, S. Zheng, and P.H.S. Torr. Higher order conditional random fields in deep neural networks. In *ECCV*, 2016.
- [4] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015.
- [5] L.-C. Chen, J. T. Barron, G. Papandreou, K. Murphy, and A. L. Yuille. Semantic image segmentation with task-specific edge detection using cnns and a discriminatively trained domain transform. In *CVPR*, 2016.
- [6] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv preprint*, 2016.
- [7] J. Dai, K. He, and J. Sun. Convolutional feature masking for joint object and stuff segmentation. In *CVPR*, 2015.
- [8] J. Dai, K. He, Y. Li, S. Ren, and J. Sun. Instance-sensitive fully convolutional networks. In *ECCV*, 2016.
- [9] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *CVPR*, 2016.
- [10] J. Dai, Y. Li, K. He, and J. Sun. R-fcn: Object detection via region-based fully convolutional networks. In *NIPS*, 2016.
- [11] 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.

- [12] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.
- [13] R. Girshick. Fast r-cnn. In *ICCV*, 2015.
- [14] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [15] B. Hariharan, P. Arbelaez, L. Bourdev, S. Maji, and J. Malik. Semantic contours from inverse detectors. In *ICCV*, 2011.
- [16] B. Hariharan, P. Arbelaez, R. Girshick, and J. Malik. Simultaneous detection and segmentation. In *ECCV*, 2014.
- [17] B. Hariharan, P. Arbelaez, R. Girshick, and J. Malik. Hypercolumns for object segmentation and fine-grained localization. In *CVPR*, 2015.
- [18] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [19] K. Li, B. Hariharan, and J. Malik. Iterative instance segmentation. In *CVPR*, 2016.
- [20] Y. Li, H. Qi, J. Dai, X. Ji, and Y. Wei. Fully convolutional instance-aware semantic segmentation. In *CVPR*, 2017. <https://arxiv.org/pdf/1611.07709v1.pdf>.
- [21] G. Lin, C. Shen, A. van den Hengel, and I. Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *CVPR*, 2016.
- [22] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollar, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [23] S. Liu, X. Qi, J. Shi, H. Zhang, and J. Jia. Multi-scale patch aggregation (mpa) for simultaneous detection and segmentation. In *CVPR*, 2016.
- [24] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [25] R. Mottaghi, X. Chen, X. Liu, N.-G. Cho, S.-W. Lee, S. Fidler, R. Urtasun, and A. Yuille. The role of context for object detection and semantic segmentation in the wild. In *CVPR*, 2014.
- [26] V. Q. Pham, T. Kozakaya, O. Yamaguchi, and R. Okada. Count forest: Co-voting uncertain number of targets using random forest for crowd density estimation. In *ICCV*, 2015.
- [27] P. O. Pinheiro, R. Collobert, and P. Dollar. Learning to segment object candidates. In *NIPS*, 2015.
- [28] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollar. Learning to refine object segments. In *ECCV*, 2016.
- [29] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [30] A. G. Schwing and R. Urtasun. Fully connected deep structured networks. *arXiv preprint*, 2015.
- [31] A. Shrivastava and A. Gupta. Contextual priming and feedback for faster r-cnn. In *ECCV*, 2016.
- [32] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016.

- [33] J. R. Uijlings, K. E. van de Sande, T. Gevers, and A. W. Smeulders. Selective search for object recognition. *IJCV*, 2013.
- [34] 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.