

# Weakly-supervised Instance Segmentation via Class-agnostic Learning with Salient Images

Xinggong Wang<sup>†</sup> Jiapei Feng<sup>†</sup> Bin Hu<sup>†</sup> Qi Ding<sup>‡</sup> Longjin Ran<sup>‡</sup> Xiaoxin Chen<sup>‡</sup> Wenyu Liu<sup>†\*</sup>

<sup>†</sup>Huazhong University of Science and Technology <sup>‡</sup>VIVO Inc.

{xgwwang, fjpp, binhu\_19, liuwuy}@hust.edu.cn {qi.ding, lj.ran, xiaoxin.chen}@vivo.com

## Abstract

Humans have a strong class-agnostic object segmentation ability and can outline boundaries of unknown objects precisely, which motivates us to propose a box-supervised class-agnostic object segmentation (BoxCaseg) based solution for weakly-supervised instance segmentation. The BoxCaseg model is jointly trained using box-supervised images and salient images in a multi-task learning manner. The fine-annotated salient images provide class-agnostic and precise object localization guidance for box-supervised images. The object masks predicted by a pretrained BoxCaseg model are refined via a novel merged and dropped strategy as proxy ground truth to train a Mask R-CNN for weakly-supervised instance segmentation. Only using 7991 salient images, the weakly-supervised Mask R-CNN is on par with fully-supervised Mask R-CNN on PASCAL VOC and significantly outperforms previous state-of-the-art box-supervised instance segmentation methods on COCO. The source code, pretrained models and datasets are available at <https://github.com/hustvl/BoxCaseg>.

## 1. Introduction

Weakly-supervised instance segmentation (WSIS) is important in computer vision for at least two reasons: First, it reduces human annotation efforts so that it is useful in building label-efficient visual recognition systems. Second, humans can learn to segment objects without perfect/exact labels, and building a visual learning system similar to human perception is a long-term goal of computer vision; thus the research weakly-supervised instance segmentation is a way to approach human perception. In this paper, we investigate the problem of weakly-supervised instance segmentation using bounding box annotations rather than the precise mask annotations.

As reported in [3], labeling a mask takes about 79 seconds while labeling a bounding box takes only about 10 seconds on average. Thus, if a WSIS method can obtain

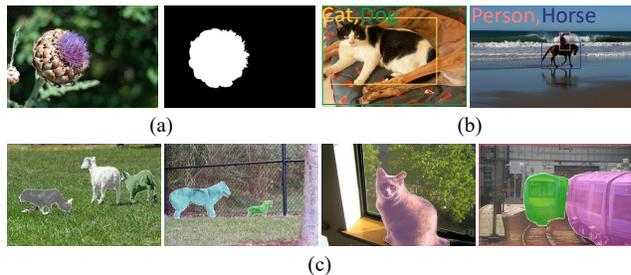


Figure 1. The top row shows a salient training image (a) and two box-supervised training images (b). Our model is jointly trained using both salient images and box-supervised images and obtains high quality instance segmentation results (c).

the segmentation precision on par with the state-of-the-art fully supervised instance segmentation methods (e.g., Mask R-CNN [13]), it is extremely useful and has great potential to be applied in real world applications. To solve this problem, previous methods have applied unsupervised image segmentation methods based on hand-crafted features [40, 38], pseudo labeling [7, 42, 2], multi-instance learning [16] and attention networks [23] in various WSIS frameworks.

We think the key problem in WSIS is box-supervised class-agnostic object segmentation (BoxCaseg), *i.e.*, given the object bounding box, we need to infer pixel-level object mask. Humans have strong ability on the task of BoxCaseg and can perceive object boundaries even the objects are from unknown semantic categories. To mimic this class-agnostic human learning process, we propose a joint training scheme, in which a salient object segmentation (also known as salient object detection) dataset that contains class-agnostic pixel-level object location information is applied to simulate humans' historical memory. The salient images are trained with box-annotated images (as shown in Fig. 1) in a single deep multi-task network and are mixed in the same mini-batches. The tasks include pixel-level labeling salient images and multi-instance classification [16] of both box-supervised images and salient images. During training, salient images can provide salient knowl-

\*Wenyu Liu is the corresponding author.

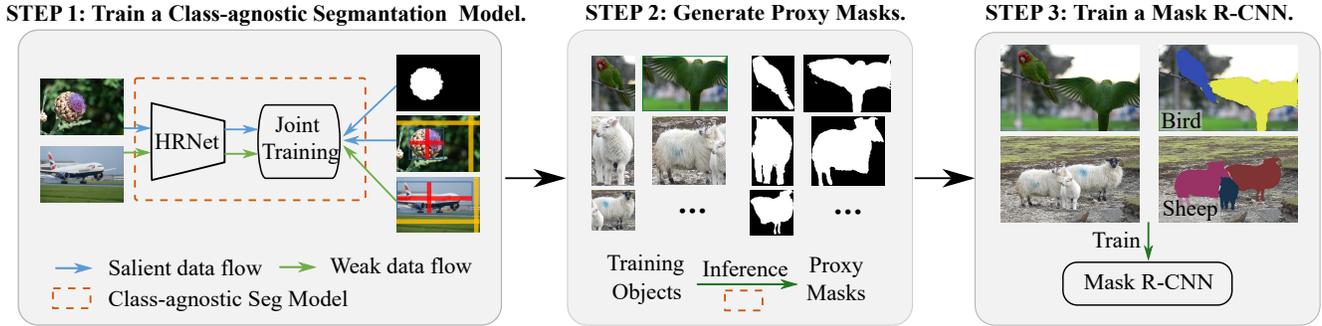


Figure 2. The pipeline of our approach. Our training process is divided into three steps. First, we obtain a class-agnostic segmentation model (BoxCaseg) through the joint training of box-supervised dataset and salient object segmentation dataset (Sec. 3.2). Then, we use the class-agnostic segmentation model to produce high quality predictions for training instances. Those segmentation predictions are merged to generate proxy instance masks for training images. Finally, we re-train a Mask R-CNN with proxy masks (Sec. 3.4).

edge to guide the segmentation of box-supervised images. Besides, inspired by Mask\* R-CNN [17], weight transfer is performed from the multi-instance learning task to the pixel-labeling task to enhance the joint learning.

The learned BoxCaseg model is applied to generate object masks for box-supervised training images. The masks are refined by our proposed merging and dropping strategies for training a Mask R-CNN model for WSIS. Experiments are carried out on the PASCAL VOC and COCO datasets. At last, our main contributions are highlighted as follows.

1. We propose a novel weakly-supervised instance segmentation method based on box-supervised class-agnostic object segmentation, in which class-agnostic precise salient object localization information is utilized as auxiliary memory to promote weakly-supervised instance segmentation during training.
2. We propose new mask merging and dropping strategies to obtain high-quality proxy masks for pseudo-training Mask R-CNN.
3. The class-agnostic segmentation model has great generalization ability. Only using 7991 salient images which are disjointed with PASCAL and COCO, for the first time, a box-supervised instance segmentation obtains similar performance with its fully-supervised counterpart. On COCO, our result is significantly better than previous state-of-the-art WSIS methods.

## 2. Related Work

Instance segmentation acquires fine-grained but expensive pixel-level annotations while many weakly supervised methods can use only coarser annotations, such as image tags [1, 6, 11, 25, 57, 58], points [24], scribble [30, 43], and bounding boxes [2, 16, 20, 44, 26].

Our setup considers using object bounding box as the weak annotation. To obtain the segmentation from the bounding box, there is a kind of methods [7, 36, 20] utilizing traditional image segmentation algorithms to segment

objects, such as CRF [22], GrabCut [40] and MCG [38]. However, these methods rely on low-level image features and cannot obtain stable WSIS performance. Dai *et al.* [7] first proposes to combine deep networks with hand-crafted object proposals for WSIS via iterative training and pseudo labeling. Then methods such as [42, 23] focus on the refinement of pseudo mask with the mean filling rates. Arun *et al.* [2] propose a conditional network and a prediction network which supervise each other. Unlike the above methods, BBTP [16] designs a multi-instance learning (MIL) formulation to train a weakly-supervised segmentation model. On the basis of BBTP, BoxInst [44] employs the projection loss to replace the MIL loss and improves the pairwise loss with the prior derived from proximal pixels' colors. Our weak segmentation head is also based on BBTP and does not rely on any traditional algorithms.

Our method is also based on the development of salient object detection (SOD, as known as salient object segmentation) [19, 49, 29, 4, 28, 15, 46], which aims at finding visually attractive objects in an image and segmenting them as a binary mask. It has achieved great success based on deep learning. Recent deep SOD methods are beneficial from prediction erasing via multi-level feature fusion [54], reverse attention [5], global and local context integration [33], cascaded partial decoder [47], edge-guided network [55], pooling-based global guidance [32] and nested U-Net (*i.e.*, U<sup>2</sup>Net) [39].

Salient object segmentation information had been applied as a cue for weakly-supervised instance/semantic segmentation, *e.g.*, [18, 10, 50, 56], in which [50, 56] are the most related works to our method. In [50] salient images are jointly trained with images with category labels for class-specific semantic segmentation and saliency maps are obtain via aggregating the class-specific prediction results. Conceptually different from [50], our method is for class-agnostic object segmentation and does not depend on the category intersection between salient images and box-supervised images. ShapeProp [56] performs pixel-wise



Figure 3. Two examples from box-supervised dataset (a) and salient object segmentation dataset (b). According to ground truth box (blue boxes in images), we can split each horizontal and vertical crossing lines of images into positive bags and negative bags. Red lines and yellow lines represent positive bags and negative bags respectively.

MIL and message passing within bounding boxes to obtain class-agnostic saliency information as prior for semi-supervised instance segmentation. In contrast, our motivation is to simulate human’s strong class-agnostic object segmentation ability and the domain gap between salient images and weakly-supervised images is much larger than that in ShapeProp. Besides, there are weakly-supervised learning methods for salient object segmentation, *e.g.*, [27, 51, 53]. In our work, the salient object information is utilized as an auxiliary memory for weakly-supervised segmenting generic images.

At last, class-agnostic object segmentation is becoming a hot research topic. Pinheiro *et al.* [37] first studies the problem of class-agnostic object candidate segmentation. Zhang *et al.* [52] perform class-agnostic segmentation via dense comparison for few-shot object segmentation. Sharma *et al.* [41] propose a novel class-agnostic segmentation loss function for better salient object detection and segmentation. Our work shows that class-agnostic learning is extremely helpful to WSIS.

### 3. Method

To describe the proposed method in a self-contained manner, we first introduce the MIL-based WSIS method [16], then formulate the proposed joint training scheme for BoxCaseg, and finally describe how to train a Proxy Mask R-CNN. The overall pipeline of our method is illustrated in Fig. 2.

#### 3.1. Weakly-supervised Instance Segmentation via Multi-instance Learning

In order to approach the instance segmentation task with bounding box annotations, Hsu *et al.* [16] propose a MIL-based method that utilizes the bounding box tightness prior (BBTP). The method replaces the original segmentation loss with a novel MIL loss in Mask R-CNN [13]. As shown in Fig. 3, the entire region of the object is in the bounding box and each row or each column in the bounding box must contain at least one pixel belonging to the object. BBTP regards these rows and columns as positive bags. Other rows

(or columns) outside the bounding box are treated as negative bags. A bag contains multiple pixels/features on the row (or column). As a result, either given the bounding box from box-supervised images or the box generated by salient object mask, we can split the image into positive bags and negative bags.

In Mask R-CNN, it is worth to note that there are multiple region proposals corresponding to one object. Thus the positive and negative bags generation process will be performed for each region proposal separately. Once the positive and negative bags are generated, its classification probabilities are generated by max-pooling pixel classification probabilities. The bag classification loss can be back-propagated to optimize pixel classifier to achieve pixel segmentation.

#### 3.2. Joint Training with Salient Images

The above MIL-based WSIS method can be end-to-end optimized under box supervision. However, it is hard to perceive precise object boundary information. We propose to utilize a relatively small extra salient object segmentation dataset as auxiliary knowledge. These data that contain salient knowledge are jointly trained with box-supervised images to provide boundary localization guidance during the whole learning process.

Our joint training network for box-supervised class-agnostic segmentation (BoxCaseg) is shown in Fig. 4. The BoxCaseg network takes both salient images and box-supervised images as input. HRNet [45] is applied for feature extraction. The salient image feature and box-supervised image feature are fed into three segmentation heads supervised by the pixel-labeling loss and the MIL loss. The details are described in the following sections.

##### 3.2.1 Salient Images

The salient images with mask annotations are derived from existing salient object segmentation/detection datasets, *e.g.*, the DUTS-TR dataset [46]. In each salient image, we make sure there is only one object. The salient objects are from a large range of categories and occupy a large region in their images. The salient images are used for providing generalizable fine-grained boundary localization information. To avoid over-fitting, we make sure the salient images used in our method have no overlap with the target WSIS datasets, *e.g.*, PASCAL and COCO.

##### 3.2.2 Data Augmentation

Different from BBTP which completes detection and segmentation simultaneously in Mask R-CNN, in which patch features are extracted using ROIAlign, we train a class-agnostic object segmentation model separately. Though ROIAlign is time-efficient to get patch features, it does not extract fine-grained features for precise object segmentation, as discussed in PointRend [21]. Therefore, instead of

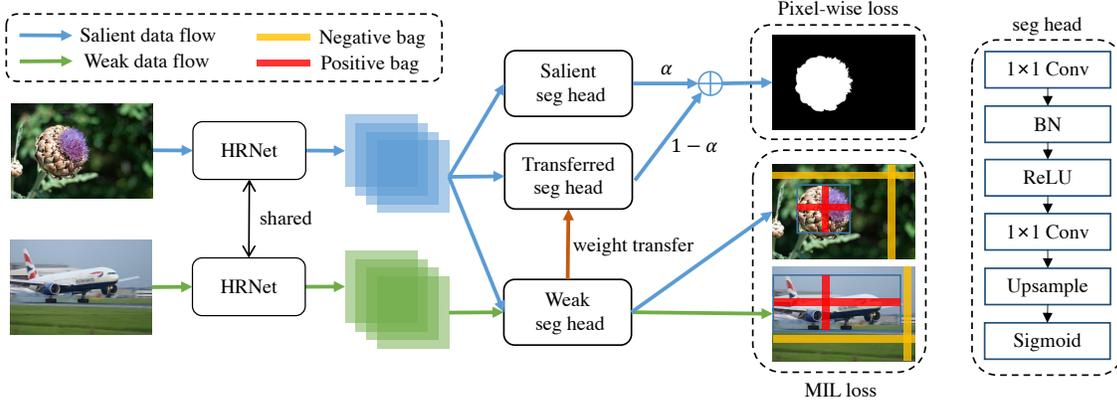


Figure 4. The joint training scheme for box-supervised class-agnostic object segmentation. A salient image and a box-supervised image within a mini-batch are fed into a backbone network (*i.e.*, HRNet [45]) to extract convolutional feature maps. The weak convolutional feature map is then fed into the weak segmentation head trained with the MIL loss described in Eq. (1). The salient convolutional feature map is simultaneously fed into the weak segmentation head, the transferred segmentation head, and the salient segmentation head. The last two heads are fused and supervised via a pixel-labeling loss (as described in Eq. (3)).

using ROIAlign, we crop patches according to object positions as the input of our BoxCaseg network. The patch cropping (*i.e.*, data augmentation) strategies are different for salient images and box-annotated images. For each salient image, it is firstly resized into  $320 \times 320$  and then randomly cropped into  $288 \times 288$  as the input. For box-annotated images, we find that the following more aggressive data augmentation strategy works better.

```

1 def box_augmentation(x, y, w, h):
2     x2 = x + random(-0.25, 0.25) * w
3     y2 = y + random(-0.25, 0.25) * h
4     w2 = random(0.5, 1.5) * w
5     h2 = random(0.5, 1.5) * h
6     return (x2, y2, w2, h2)

```

This data augmentation strategy shifts the position and shape of the bounding box and introduces background information. Once a new box is obtained, an image patch is cropped and resized into  $288 \times 288$  for joint training.

### 3.3. Multi-task Learning

As shown in Fig. 4, there are three tasks for training our BoxCaseg network, *i.e.*, MIL training for box-supervised images, MIL training for salient images, and pixel-labeling for salient images. There are three segmentation heads termed as salient segmentation head, transferred segmentation head, and weak segmentation head. The three heads have exactly the same network architecture as shown in the right side of Fig. 4.

#### 3.3.1 MIL loss for weak segmentation

The convolutional feature for the salient image is denoted as  $P_s$  and the convolutional feature for the weakly-supervised (*i.e.*, box-supervised) image is denoted as  $P_w$ .  $P_s$  and  $P_w$

are fed into the weak segmentation head to obtain one-channel probability maps  $S_s$  and  $S_w$  via a pixel-wise sigmoid function, respectively. Besides, the upsample operation makes sure that  $S_s$  and  $S_w$  have the same size as the input images. As described in Sec. 3.1, we can sample a set of positive bags  $\mathbf{B}^+$  and a set of negative bags  $\mathbf{B}^-$  from  $S_s$  and  $S_w$  respectively. The MIL loss is defined as follows.

$$\begin{aligned}
 L_{\text{mil}}(S, \mathbf{B}^+, \mathbf{B}^-) = & - \sum_{B \in \mathbf{B}^+} \log \left( \max_{p \in B} S(p) \right) \\
 & - \sum_{B \in \mathbf{B}^-} \log \left( 1 - \max_{p \in B} S(p) \right) \quad (1) \\
 & + 0.05 \times \varphi(S, \mathbf{B}^+, \mathbf{B}^-),
 \end{aligned}$$

where  $p$  is a spatial position in a bag  $B \in \{\mathbf{B}^+, \mathbf{B}^-\}$  and a score map in  $S \in \{S_s, S_w\}$ ,  $S(p)$  denotes the score at position  $p$ , and the last item  $\varphi(S, \mathbf{B}^+, \mathbf{B}^-)$  is a smooth term with a fixed weight of 0.05 defined as follows.

$$\varphi(S, \mathbf{B}^+, \mathbf{B}^-) = \sum_{p \in \{\mathbf{B}^+, \mathbf{B}^-\}} \sum_{p' \in \Omega(p)} \|S(p) - S(p')\|^2, \quad (2)$$

where  $\Omega(p)$  means eight-connected region of the pixel  $p$ . In Eq. (1), the unary term uses MIL for pixel classification and the pairwise term aims to keep local smoothness. Positive bags having at least one pixel of the object, we encourage the maximal score of any positive bag to be as close to 1 as possible, while suppressing maximal scores in negative bags.

#### 3.3.2 Pixel-wise loss for salient segmentation and transferred segmentation

The pixel-wise loss function is only applied to salient images. The precise annotations of salient images with the

pixel-wise loss are helpful for our BoxCaseg network to capture class-agnostic boundary information. As this loss cannot be applied to box-supervised images, we propose a transferred segmentation head to make the pixel-wise labeling parameters more suitable for box-supervised images. It is enabled by weight transfer, which is originally proposed in Mask\* RCNN [17]. As shown in Fig. 4, we transfer detached convolution parameters of the weak segmentation head to the transferred segmentation head via a weight transfer module. As for the weight transfer module, we simply use a two-layer multi-layer perceptron (MLP) with leaky ReLU as activation function. Note that the weight transfer is a single direction process, since it is detached from the weak segmentation head, *i.e.*, the gradients do not back-propagate to the weak segmentation head. During training, the parameters of the weight transfer MLP is optimized. Given a salient image, its ground-truth is a binary mask denoted as  $M$  and we denote its score maps output from the salient segmentation head and transferred segmentation head as  $S_a$  and  $S_t$  respectively. Then, the pixel-wise loss function is a standard binary cross-entropy loss defined as follows.

$$L_{\text{pix}}(M, S_a, S_t) = -M \log(\alpha S_a + (1 - \alpha) S_t) - (1 - M) \log(1 - (\alpha S_a + (1 - \alpha) S_t)), \quad (3)$$

where  $\alpha \in [0, 1]$  is a parameter to control how much transferred segmentation map used for calculating the pixel-wise loss.

The final loss function is summarized as follows.

$$L = L_{\text{mil}}(S, \mathbf{B}^+, \mathbf{B}^-) + \lambda L_{\text{pix}}(M, S_a, S_t), \quad (4)$$

where  $\lambda$  is equal to 0 if the input comes from box-supervised images, otherwise 1. Last but not least, during the training of the proposed network, there are both box-supervised images and salient images in a mini-batch. Their ratios will be detailed in the experiments. This mixed input strategy helps stabilize the training process and solve the data imbalance problem, *e.g.*, in the COCO dataset there are a huge number of box-supervised images while the number of salient images is limited.

### 3.4. Training a Proxy Mask R-CNN

To use jointly trained BoxCaseg network for instance segmentation, we need to obtain the positions of objects with a detector first in the testing phase. To simplify this process and avoid the time-consuming two-step testing procedure, we consider re-training a Mask R-CNN to perform box localization and mask prediction in one network. Our solution is straight-forward - we can use the BoxCaseg model to generate proxy masks on the training set. When generating proxy masks for training instances with the pre-trained BoxCaseg network, we just simply crop patches

from the image by ground truth boxes as inputs and only use the salient segmentation head and transferred segmentation head to predict object mask. In our method, the Mask R-CNN model is trained with the bounding box annotations and the proxy masks, thus called Proxy Mask R-CNN.

Though the jointly trained BoxCaseg model obtains high performance, there are inevitably segmentation errors. The main difficulty is caused by the fact that multiple objects are highly-overlapped in one bounding box, *e.g.*, motorbike with person and sofa with person in Fig. 5. Since these object pairs have large probabilities to appear together, weakly-supervised methods are hard to deal with them. To reduce the errors in the masks generated by BoxCaseg, we propose a *merge and drop*, *i.e.*, merge masks using the strategy of *smaller object better* and drop masks via the *proxy box agreement* rule, which are detailed as follows.

**Merging via smaller object better:** Our prior knowledge tells us that each pixel in an image can only belong to one object. When merging segmentation results, for those pixels belong to multiple predicted masks, we assign them to the smallest objects. The reason is that the smaller object is usually within the box of the large one.

**Dropping via proxy box agreement:** After the merging procedure, we further drop low quality masks. It is observed that the low quality masks have shapes significantly different from their ground truth (GT), however, we only have GT boxes but do not have GT masks. Thus, we find the bounding box of the proxy mask, called proxy bounding box, and calculate the IoU between the proxy bounding box and GT bounding box as the proxy box agreement. It is the lower bound of mask IoU between proxy mask and GT mask and can reflect the segmentation quality. Note that dropping the mask does not mean removing it from the annotation file. Instead, those deleted masks are neglected in the back propagation process of mask head training in Mask R-CNN.

## 4. Experiment

**Datasets** Extensive experiments are performed on the COCO 2017 dataset [31] and the augmented PASCAL VOC 2012 dataset [12]. And we adopt the DUTS-TR dataset [46] to provide salient images.

**PASCAL VOC:** It is a commonly-used evaluating dataset for WSIS. Following the common practice [20], augmented PASCAL VOC 2012 dataset is used, which is derived by combining the SBD dataset [12] and the PASCAL VOC 2012 dataset [9]. The final dataset contains 10,582 training images with 25,815 training instances and 1,449 validation images.

**COCO:** It is a large-scale instance segmentation dataset which contains 118,287 training images with 860,001 instances and 5,000 images for validation. It is also applied for evaluating WSIS in our experiments.

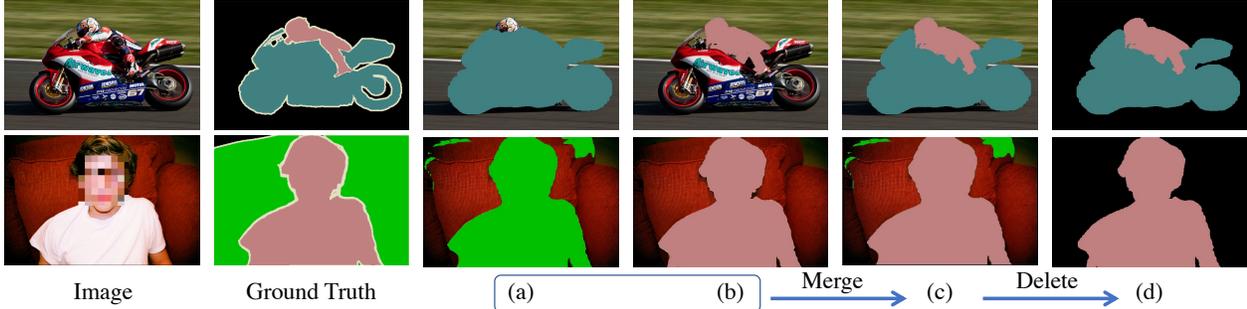


Figure 5. From the prediction results of BoxCaseg (a&b) to the proxy masks (d). Since there are two highly-overlapped objects in each image. BoxCaseg predicts two masks with errors. We first merge the two predicted masks using the *smaller object better* strategy and then drop bad masks (e.g., the sofa mask in bottom row) according to the *proxy box agreement* rule. In this case, the bounding box of the sofa mask in (c) has low agreement (i.e., low IoU) with its ground truth box, so the sofa mask is dropped.

**DUTS-Single:** The images in the DUTS-TR dataset [46] are collected from ImageNet DET [8] training/val/test sets and the SUN dataset [48]. DUTS-TR is widely-used for salient object segmentation/detection task. *We have carefully checked the image sources of DUTS-TR, they have no overlap with Pascal VOC and COCO dataset.* Usually, a bounding box contains only one object, so we filter out images which contain multiple objects to ensure that the foreground mask only belongs to a single object. The derived salient object segmentation dataset is named DUTS-Single, which contains 7991 images. The derived DUTS-Single dataset will be made publicly available.

**Evaluation metric** For evaluating the performance of class-agnostic object segmentation, we use  $\text{IoU}@k$ , which represents the percentage of predicted instance masks whose intersection over union (IoU) with its ground-truth mask is larger than  $k\%$ , where  $k \in \{50, 75\}$  in our experiments. Besides, we design a  $\text{mIoU}^*$  metric to evaluate our model’s class-related performance, which calculates IoUs of every instance prediction, per-class mean IoU and averages all the per-class mean IoUs.

After retraining the Proxy Mask R-CNN, we adopt the standard evaluation metric, the mean average precision (mAP). We follow the same evaluation scheme as [16] and report mAP with four IoU thresholds (including 0.25, 0.50, 0.70 and 0.75) for PASCAL VOC dataset. Besides, we also report the evaluation metric including AP, AP50, AP75 (averaged over IoU thresholds), and APS, APM, APL (AP at different scales) for COCO dataset.

**Implementation details** For joint training the BoxCaseg network, we apply the HRNetV2-W48 [45] as the backbone and its default hyper-parameters are applied, i.e., initial learning rate of  $4e-3$ , weight decay of  $1e-4$ , momentum of 0.9. The batch-size is set to be 16 per GPU. For generating proxy masks for training Mask R-CNN, each training object is cropped from its image according to its ground-truth bounding box and resized into  $320 \times 320$  as the input

of pre-trained BoxCaseg network. Proxy Mask R-CNN is trained using ResNet-101-FPN [14] as the backbone following the default settings in the maskrcnn benchmark [35]. Experiments are carried out on two Nvidia TITAN V GPUs (12GB RAM). Source code and pre-trained models will be made publicly available.

#### 4.1. Comparison with the State-of-the-art Methods

Tab. 1 presents the instance segmentation results on the PASCAL VOC 2012 validation set, including the fully-supervised Mask R-CNN baseline, weakly-supervised methods using image-level annotation and box-level annotation. In the metrics,  $AP_{50}$  is the most widely used one. Our  $AP_{50}$  is 67.6, which is on par with the fully-supervised Mask R-CNN. With the help of joint training with salient images, our method significantly outperforms other weakly-supervised instance segmentation methods. Comparing with the MIL-based BBTP method [16], joint training with salient images brings 8.7  $AP_{50}$  performance gain and 20.8  $AP_{75}$  performance gain. The results confirm that joint training with salient images is very helpful for precisely segmenting objects.

Tab. 2 presents the results on the COCO test-dev set. COCO is a large-scale instance segmentation dataset that contains 118,287 training images that contain about 860,000 training instances. Using only 7,991 (9.3%) salient objects for joint training, we improve the state-of-the-art baseline (BBTP) by 10.1  $AP$  and the results are getting close to the fully-supervised Mask R-CNN.

#### 4.2. Ablation Studies

To validate our proposed method and configurations’ usefulness, extensive ablation experiments are conducted on both the joint training BoxCaseg stage and the re-training stage.

##### 4.2.1 Less salient images for joint training

In order to investigate the influence of the number of salient images, we perform ablation experiments on the number of

Table 1. Instance segmentation results on the PASCAL VOC 2012 val set. In the supervision column, M denotes mask-level supervision, I denotes image-level supervision, B denotes box-level supervision and S denotes salient object supervision. † indicates the results are obtained by running their released codes using our settings.

Method	Sup.	$AP_{25}$	$AP_{50}$	$AP_{70}$	$AP_{75}$
ResNet-50 Backbone					
PRM [57]	I	44.3	26.8	-	9.0
IAM [58]	I	45.9	28.8	-	11.9
Label-PEnet [11]	I	49.1	30.2	-	12.9
CountSeg [6]	I	48.5	30.2	-	14.4
WISE [25]	I	49.2	41.7	-	23.7
IRN [1]	I	-	46.7	23.5	-
LIID [34]	I	-	48.4	-	24.9
BBTP† [16]	B+S	75.9	61.0	33.4	25.1
ShapeProp† [56]	B+S	75.1	61.6	37.8	30.3
BoxCaseg (ours)	B+S	<b>76.2</b>	<b>65.8</b>	<b>46.3</b>	<b>38.6</b>
ResNet-101 Backbone					
SDI [20]	B	-	46.4	-	18.5
AnnoCoIn [2]	B	73.1	57.7	33.5	31.2
BBAM [26]	B	76.8	63.7	39.5	31.8
BoxInst [44]	B	-	61.4	-	37.0
BBTP [16]	B	75.0	58.9	30.4	21.6
BBTP† [16]	B+S	77.0	61.4	34.4	26.6
ShapeProp† [56]	B+S	76.0	62.4	39.9	31.3
BoxCaseg (ours)	B+S	<b>77.7</b>	<b>67.6</b>	<b>49.4</b>	<b>42.4</b>
Mask R-CNN [13]	M	76.7	67.9	52.5	44.9

salient images in {100, 500, 1000, 2500, 4701, 7991/All} for joint training. The 100, 500, 1000, & 2500 salient images are randomly selected. *The 4701 salient images are manually picked to make sure that they do not contain any object belongs to any of the 20 PASCAL categories.*

As Tab. 3 shows, when there are too few salient images (100), our salient segmentation head and transferred segmentation head will perform awful and serve unstable performance, even worse than the performance of weak segmentation head. When the number of salient images comes to over 500, its training process becomes stable and the model will perform better with images number increase both in mIoU\* and AP evaluation metric. Training with the 4701 salient images, the results are very close to that uses all salient images, which confirms that BoxCaseg has a very strong category-wise transfer ability.

#### 4.2.2 The effectiveness of the three segmentation heads

Tab. 4 indicates the influence of three segmentation heads for the whole model respectively. First, only using a salient segmentation head, which is trained by salient images, will lead to an awful segmentation result on box-supervised images. This is due to the salient images are different from the box-supervised PASCAL images. So it is necessary to add

the weak segmentation head to help the backbone learn how to extract features of box-supervised images. Second, transferred segmentation head, which delivers box-supervised information from the weak segmentation head to pixel-level prediction, helps the salient segmentation head improve segmentation quality on box-supervised images. Finally, by using all three heads, we finally get a 71.8 mIoU\* instance segmentation result on Pascal VOC 2012 validation set. In fact, our BoxCaseg model has a strong salient object segmentation performance that is on-par with the SOTA U<sup>2</sup>Net[39].

#### 4.2.3 Sampling strategies for imbalanced training images

When there is an extreme unbalance between COCO (860,001 instances) and DUTS-Single (7991 instances), we discover that the model performs worse as the training progressing. Thus, we propose a simple yet efficient sample strategy that we use all salient images and randomly sample part of box-supervised images to form batches in an epoch. And all batches own a common fixed ratio of box-supervised images to salient images. Tab. 5 demonstrates our sample strategy is efficient, where  $r_{ws}$  is the ratio of box-supervised images to salient images in a batch and RS represents random sampling. Obviously, this strategy greatly solves the problem of data unbalance and makes training process more stable as well as enables model with better generalization ability. When  $r_{ws}$  is equal to 9 : 7, we can get 70.9 mIoU\* in COCO validation set.

#### 4.2.4 The weight of transferred segmentation head

We also conduct experiments on  $\alpha$ , which controls the weight of the transferred segmentation head when adding with the salient segmentation head for pixel-level prediction. As Tab. 6 shows, combining the prediction of transferred segmentation head and salient segmentation head, we have access to deriving better results than only using one of two heads by at least 1.3% improvement. What's more, it's relatively stable for segmentation result, when  $\alpha \in [0.3, 0.7]$  and proves our weight transfer method's usefulness. And we can derive the best result when  $\alpha$  is 0.7.

#### 4.2.5 The strategies for merging and deleting proxy masks

In Tab. 7, we firstly give a baseline of retraining results without any strategies. We can observe that the merge via smaller object better strategy improves the performance by 0.8  $AP_{75}$  and 1.1  $AP_{75}$  on the two datasets respectively. At the same time, the drop via proxy box agreement strategy obtains obvious improvement on the COCO dataset while little improvement on the PASCAL dataset. This is due to the fact that PASCAL is simpler than COCO dataset and bad cases are much fewer. In addition, we tried different

Table 2. Instance segmentation results on the COCO test-dev set.

Method	Supervision	Backbone	$AP$	$AP_{50}$	$AP_{75}$	$AP_S$	$AP_M$	$AP_L$
Mask R-CNN [13] (Oracle)	M	ResNet-101-FPN	35.7	58.0	37.8	15.5	38.1	52.4
BBTP [16]	B	ResNet-101-FPN	20.8	45.1	17.0	10.4	21.7	30.3
BBAM [26]	B	ResNet-101-FPN	25.7	50.0	23.3	-	-	-
BoxCaseg (ours)	B+S	ResNet-101-FPN	<b>30.9</b>	<b>54.3</b>	<b>30.8</b>	<b>12.1</b>	<b>32.8</b>	<b>46.3</b>

Table 3. Results on PASCAL val using different numbers of salient images for joint training.

Number	mIoU*	IoU@50	IoU@75	AP@50	AP@75
100	61.8	70.4	32.2	55.3	27.0
500	65.3	73.3	39.6	62.6	36.2
1000	66.4	77.4	47.3	65.1	37.2
2500	69.2	79.9	51.1	67.1	40.9
4701	70.1	80.1	50.2	66.7	38.9
All	71.8	83.8	56.8	67.6	42.4

Table 4. Results on PASCAL val when using different segmentation heads. Weak, Salient, and Trans denote the weak, salient, and transferred segmentation heads respectively. If only Trans or Salient is checked,  $\alpha$  in Eq. (3) is set to 0 or 1.

Weak	Salient	Trans.	mIoU*	IoU@50	IoU@75
✓			66.1	76.8	40.0
	✓		43.1	43.0	25.9
✓		✓	69.7	80.0	50.3
✓	✓		70.5	80.8	53.7
✓	✓	✓	<b>71.8</b>	<b>83.4</b>	<b>56.8</b>

Table 5. Results on COCO val using different image sampling methods. RS means randomly sampling from the union set of box-supervised images and salient images.  $r_{ws}$  is the ratio of box-supervised images to salient images in a batch.

$r_{ws}$	mIoU*	IoU@50	IoU@75
RS	58.1	62.2	27.9
8:8	70.7	84.5	48.0
9:7	<b>70.9</b>	<b>84.9</b>	<b>48.6</b>
10:6	70.0	82.7	45.4
11:5	70.7	83.6	47.8
12:4	70.4	82.7	47.2
13:3	68.1	77.6	44.1
14:2	69.2	81.1	46.4

Table 6. Results on PASCAL val set when the weight of salient segmentation head ( $\alpha$ ) varies.

$\alpha$	0.0	0.3	0.4	0.5	0.6	0.7	1.0
mIoU*	69.7	71.7	71.0	71.4	71.3	<b>71.8</b>	70.5
IoU@50	80.0	83.0	82.0	82.6	82.7	<b>83.4</b>	80.8
IoU@75	50.3	55.7	54.9	55.6	55.2	<b>56.8</b>	53.7

values for the dropping threshold to find the best performing threshold, shown in Tab. 8. The drop rate refers to the percentage of the ignored samples in the total samples. The

Table 7. Results on PASCAL val and COCO val when using the merging and dropping strategies or not.

Data	Merge	Drop	$AP$	$AP_{50}$	$AP_{75}$
PASCAL			39.6	67.5	41.3
PASCAL	✓		<b>40.2</b>	<b>67.6</b>	42.1
PASCAL	✓	✓	<b>40.2</b>	<b>67.6</b>	<b>42.4</b>
COCO			28.6	51.3	28.2
COCO	✓		29.5	51.9	29.3
COCO	✓	✓	<b>30.9</b>	<b>53.7</b>	<b>31.3</b>

results show that the drop strategy can pick out samples with better proxy masks and bring robust performance improvements. Using a very high IoU threshold of 0.95 obtains the best results among them, which reveals that have a few very high quality mask for training Mask R-CNN is better than having lots of medium quality masks.

Table 8. COCO val results and the ratios of dropped masks when the dropping threshold varies.

Drop thresh	0.0	0.75	0.85	0.90	0.95
Drop rate	0.00	0.27	0.49	0.64	0.80
COCO AP	29.5	30.2	30.4	30.7	<b>30.9</b>

## 5. Conclusion

We explore a new direction for WSIS, *i.e.*, box-supervised class-agnostic learning with salient images for object segmentation. The proposed BoxCaseg network uses the salient object information as an auxiliary memory during the whole training process. This design is to mimic the way humans perform class-agnostic object segmentation. Once the proxy masks are obtained via the pre-trained BoxCaseg model, we propose useful mask merging and dropping strategies for re-training Mask R-CNN. On PASCAL and COCO, state-of-the-art weakly-supervised instance segmentation results have been achieved. In the next step, we would like to develop new learning methods to push the performance of the BoxCaseg-based WSIS method to match the performance of fully-supervised Mask R-CNN on COCO using the 7991 salient images.

**Acknowledgement:** This work was in part supported by NSFC (No. 61876212 and No. 61733007) and the Zhejiang Laboratory under Grant 2019NB0AB02.

## References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2209–2218, 2019. 2, 7
- [2] Aditya Arun, CV Jawahar, and M Pawan Kumar. Weakly supervised instance segmentation by learning annotation consistent instances. *ECCV*, 2020. 1, 2, 7
- [3] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *European conference on computer vision*, pages 549–565. Springer, 2016. 1
- [4] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. Salient object detection: A benchmark. *IEEE transactions on image processing*, 24(12):5706–5722, 2015. 2
- [5] Shuhan Chen, Xiuli Tan, Ben Wang, and Xuelong Hu. Reverse attention for salient object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 234–250, 2018. 2
- [6] Hisham Cholakkal, Guolei Sun, Fahad Shahbaz Khan, and Ling Shao. Object counting and instance segmentation with image-level supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 12397–12405, 2019. 2, 7
- [7] Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1635–1643, 2015. 1, 2
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. IEEE, 2009. 6
- [9] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. 5
- [10] Ruochen Fan, Qibin Hou, Ming-Ming Cheng, Gang Yu, Ralph R. Martin, and Shi-Min Hu. Associating inter-image salient instances for weakly supervised semantic segmentation. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [11] Weifeng Ge, Sheng Guo, Weilin Huang, and Matthew R Scott. Label-penet: Sequential label propagation and enhancement networks for weakly supervised instance segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3345–3354, 2019. 2, 7
- [12] Bharath Hariharan, Pablo Arbeláez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *2011 International Conference on Computer Vision*, pages 991–998. IEEE, 2011. 5
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1, 3, 7, 8
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [15] Qibin Hou, Ming-Ming Cheng, Xiaowei Hu, Ali Borji, Zhuowen Tu, and Philip HS Torr. Deeply supervised salient object detection with short connections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3203–3212, 2017. 2
- [16] Cheng-Chun Hsu, Kuang-Jui Hsu, Chung-Chi Tsai, Yen-Yu Lin, and Yung-Yu Chuang. Weakly supervised instance segmentation using the bounding box tightness prior. In *Advances in Neural Information Processing Systems*, pages 6586–6597, 2019. 1, 2, 3, 6, 7, 8
- [17] Ronghang Hu, Piotr Dollár, Kaiming He, Trevor Darrell, and Ross Girshick. Learning to segment every thing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4233–4241, 2018. 2, 5
- [18] Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. Weakly-supervised semantic segmentation network with deep seeded region growing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7014–7023, 2018. 2
- [19] Huaizu Jiang, Jingdong Wang, Zejian Yuan, Yang Wu, Nanning Zheng, and Shipeng Li. Salient object detection: A discriminative regional feature integration approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2083–2090, 2013. 2
- [20] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017. 2, 5, 7
- [21] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9799–9808, 2020. 3
- [22] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in neural information processing systems*, pages 109–117, 2011. 2
- [23] Viveka Kulharia, Siddhartha Chandra, Amit Agrawal, Philip Torr, and Amrith Tyagi. Box2seg: Attention weighted loss and discriminative feature learning for weakly supervised segmentation. In *European Conference on Computer Vision*, pages 290–308. Springer, 2020. 1, 2
- [24] Issam H Laradji, Negar Rostamzadeh, Pedro O Pinheiro, David Vázquez, and Mark Schmidt. Instance segmentation with point supervision. *arXiv preprint arXiv:1906.06392*, 2019. 2
- [25] Issam H Laradji, David Vazquez, and Mark Schmidt. Where are the masks: Instance segmentation with image-level supervision. *arXiv preprint arXiv:1907.01430*, 2019. 2, 7
- [26] Jungbeom Lee, Jihun Yi, Chaehun Shin, and Sungroh Yoon. BBAM: Bounding box attribution map for weakly supervised semantic and instance segmentation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 7, 8

- [27] Guanbin Li, Yuan Xie, and Liang Lin. Weakly supervised salient object detection using image labels. *arXiv preprint arXiv:1803.06503*, 2018. 3
- [28] Guanbin Li and Yizhou Yu. Deep contrast learning for salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 478–487, 2016. 2
- [29] Xiaohui Li, Huchuan Lu, Lihe Zhang, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via dense and sparse reconstruction. In *Proceedings of the IEEE international conference on computer vision*, pages 2976–2983, 2013. 2
- [30] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3159–3167, 2016. 2
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 5
- [32] Jiang-Jiang Liu, Qibin Hou, Ming-Ming Cheng, Jiashi Feng, and Jianmin Jiang. A simple pooling-based design for real-time salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3917–3926, 2019. 2
- [33] Nian Liu, Junwei Han, and Ming-Hsuan Yang. Picanet: Learning pixel-wise contextual attention for saliency detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3089–3098, 2018. 2
- [34] Yun Liu, Yu-Huan Wu, Pei-Song Wen, Yu-Jun Shi, Yu Qiu, and Ming-Ming Cheng. Leveraging instance-, image- and dataset-level information for weakly supervised instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 7
- [35] Francisco Massa and Ross Girshick. maskrcnn-benchmark: Fast, modular reference implementation of Instance Segmentation and Object Detection algorithms in PyTorch. <https://github.com/facebookresearch/maskrcnn-benchmark>, 2018. Accessed: [2020/11/13]. 6
- [36] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015. 2
- [37] Pedro OO Pinheiro, Ronan Collobert, and Piotr Dollár. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*, pages 1990–1998, 2015. 3
- [38] Jordi Pont-Tuset, Pablo Arbelaez, Jonathan T Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping for image segmentation and object proposal generation. *IEEE transactions on pattern analysis and machine intelligence*, 39(1):128–140, 2016. 1, 2
- [39] Xuebin Qin, Zichen Zhang, Chenyang Huang, Masood Dehghan, Osmar R Zaiane, and Martin Jagersand. U2-net: Going deeper with nested u-structure for salient object detection. *Pattern Recognition*, 106:107404, 2020. 2, 7
- [40] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ” grabcut” interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics (TOG)*, 23(3):309–314, 2004. 1, 2
- [41] Angira Sharma, Naeemullah Khan, Ganesh Sundaramoorthi, and Philip Torr. Class-agnostic segmentation loss and its application to salient object detection and segmentation. *arXiv preprint arXiv:2010.14793*, 2020. 3
- [42] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3136–3145, 2019. 1, 2
- [43] Meng Tang, Abdelaziz Djelouah, Federico Perazzi, Yuri Boykov, and Christopher Schroers. Normalized cut loss for weakly-supervised cnn segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1818–1827, 2018. 2
- [44] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. BoxInst: High-performance instance segmentation with box annotations. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 7
- [45] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 3, 4, 6
- [46] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 136–145, 2017. 2, 3, 5, 6
- [47] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3907–3916, 2019. 2
- [48] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3485–3492, 2010. 6
- [49] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3166–3173, 2013. 2
- [50] Yu Zeng, Yunzhi Zhuge, Huchuan Lu, and Lihe Zhang. Joint learning of saliency detection and weakly supervised semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7223–7233, 2019. 2
- [51] Yu Zeng, Yunzhi Zhuge, Huchuan Lu, Lihe Zhang, Mingyang Qian, and Yizhou Yu. Multi-source weak supervision for saliency detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6074–6083, 2019. 3
- [52] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. Canet: Class-agnostic segmentation networks with it-

- erative refinement and attentive few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5217–5226, 2019. 3
- [53] Jing Zhang, Xin Yu, Aixuan Li, Peipei Song, Bowen Liu, and Yuchao Dai. Weakly-supervised salient object detection via scribble annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12546–12555, 2020. 3
- [54] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Xiang Ruan. Amulet: Aggregating multi-level convolutional features for salient object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 202–211, 2017. 2
- [55] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. Egnet: Edge guidance network for salient object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8779–8788, 2019. 2
- [56] Yanzhao Zhou, Xin Wang, Jianbin Jiao, Trevor Darrell, and Fisher Yu. Learning saliency propagation for semi-supervised instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10307–10316, 2020. 2, 7
- [57] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3791–3800, 2018. 2, 7
- [58] Yi Zhu, Yanzhao Zhou, Huijuan Xu, Qixiang Ye, David Doremann, and Jianbin Jiao. Learning instance activation maps for weakly supervised instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3116–3125, 2019. 2, 7