

Weakly Supervised Object Detection With Segmentation Collaboration

Xiaoyan Li^{1,2} Meina Kan^{1,2} Shiguang Shan^{1,2,3} Xilin Chen^{1,2}

¹Key Laboratory of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS, Beijing 100190, China

²University of Chinese Academy of Sciences, Beijing 100049, China

³Peng Cheng Laboratory, Shenzhen, 518055, China

xiaoyan.li@vipl.ict.ac.cn {kanmeina, sgshan, xlchen}@ict.ac.cn

Abstract

Weakly supervised object detection aims at learning precise object detectors, given image category labels. In recent prevailing works, this problem is generally formulated as a multiple instance learning module guided by an image classification loss. The object bounding box is assumed to be the one contributing most to the classification among all proposals. However, the region contributing most is also likely to be a crucial part or the supporting context of an object. To obtain a more accurate detector, in this work we propose a novel end-to-end weakly supervised detection approach, where a newly introduced generative adversarial segmentation module interacts with the conventional detection module in a collaborative loop. The collaboration mechanism takes full advantages of the complementary interpretations of the weakly supervised localization task, namely detection and segmentation tasks, forming a more comprehensive solution. Consequently, our method obtains more precise object bounding boxes, rather than parts or irrelevant surroundings. Expectedly, the proposed method achieves an accuracy of 53.7% on the PASCAL VOC 2007 dataset, outperforming the state-of-the-arts and demonstrating its superiority for weakly supervised object detection.

1. Introduction

As the data-driven approaches prevail on object detection task in both academia and industry, the amount of data in an object detection benchmark is expected to be larger and larger. However, annotating object bounding boxes is both costly and time-consuming. In order to reduce the labeling workload, researchers hope to make object detectors work in a weakly-supervised fashion, *e.g.* learning a detector with only category labels rather than bounding boxes.

Recently, the most high-profile works on weakly supervised object detection all exploit the multiple instance learn-

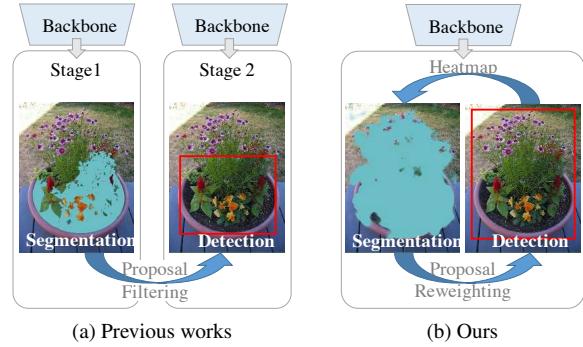


Figure 1: The schematic diagram of the previous works with segmentation utilization [7, 32] and our proposed collaboration approach. In [7, 32], a two-stage paradigm is used, in which proposals are first filtered and then detection is performed on these remaining boxes ([7] shares the backbone between two modules). In our approach, detection and segmentation modules instruct each other in a dynamic collaboration loop in the training process.

ing (MIL) paradigm [3, 5, 22, 21, 18, 14, 1, 23, 24, 2, 7]. Based on the assumption that the object bounding box should be the one contributing most to image classification among all proposals, the MIL based approaches work in an attention-like mechanism: automatically assign larger weights to the proposals consistent with the classification labels. Several promising works combining MIL with deep learning [2, 25, 32] have greatly pushed the boundaries of weakly supervised object detection. However, as noted in [25, 32], these methods are easy to over-fit on object parts, because the most discriminative classification evidence may derive from the entire object region, but may also from the crucial parts. The attention mechanism is effective in selecting the discriminative boxes, but does not guarantee the completeness of a detected object. For a more reasonable inference, a further elaborative mechanism is necessary.

Meanwhile, the completeness of a detected region is easier to ensure in weakly supervised segmentation. One common way to outline whole class-related segmentation re-

Task	Recall	Precision
Weakly supervised detection	62.9%	46.3%
Weakly supervised segmentation	69.7%	35.4%

Table 1: Pixel-wise recall and precision of detection and segmentation results on the VOC 2007 test set, following the same setting in Sec. 4.2. For a comparable pixel-level metric, the detection results are converted to the equivalent segmentation maps in a similar way described in Sec. 3.3.

gions is recurrently discovering and masking these regions in several forward passes [31]. These segmentation maps can potentially constrain the weakly supervised object detection, given that a proposal having low intersection over union (IoU) with the corresponding segmentation map is not likely to be an object bounding box. In [7, 32], weakly supervised segmentation maps are used to filter object proposals and reduce the difficulty of detection, as shown in Fig. 1a. However, these approaches adopt cascaded or independent models with relatively coarse segmentations to do “hard” delete on the proposals, inevitably resulting in a drop of the proposal recall. In a word, these methods underutilize the segmentation and limit the improvements.

The MIL based object detection approaches and semantic segmentation approaches focus on restraining different aspects of the weakly supervised localization and have opposite strengths and shortcomings. The MIL based object detection approaches are precise in distinguishing object-related regions and irrelevant surroundings, but incline to confuse entire objects with parts due to its excessive attention to the significant regions. Meanwhile, the weakly supervised segmentation is able to cover the entire instances, but tends to mix irreverent surroundings with real objects. This complementary property is verified in Table 1, that the segmentation can achieve a higher pixel-wise recall but lower precision, while the detection can achieve a higher pixel-level precision but lower recall. Rather than working independently, the two are naturally cooperative and can work together to overcome their intrinsic weaknesses.

In this work, we propose a segmentation-detection collaborative network (SDCN) for more precise object detection under weak supervision, as shown in Fig. 1b. In the proposed SDCN, the detection and segmentation branches work in a collaborative manner to boost each other. Specifically, the segmentation branch is designed as a generative adversarial localization structure to sketch the object region. The detection module is optimized in an MIL manner with the obtained segmentation map serving as spatial prior probabilities of the object proposals. Besides, the object detection branch also provides supervision back to the segmentation branch by a synthetic heatmap generated from all proposal boxes and their classification scores. Therefore, these two branches tightly interact with each other and form a dynamic cooperating loop. Overall, the entire network is

optimized under weak supervision of the classification loss in an end-to-end manner, which is superior to the cascaded or independent architectures in previous works [7, 32].

In summary, we make three contributions in this paper: 1) the segmentation-detection collaborative mechanism enforces deep cooperation between two complementary tasks and boosts valuable supervision to each other under the weakly supervised setting; 2) for the segmentation branch, the novel generative adversarial localization strategy enables our approach to produce more complete segmentation maps, which is crucial for improving both the segmentation and the detection branches; 3) as demonstrated in Section 4, we achieve the best performance on PASCAL VOC 2007 and 2012 datasets, surpassing the previous state-of-the-arts.

2. Related Works

Multiple Instance Learning (MIL). MIL [8] is a concept in machine learning, illustrating the essence of inexact supervision problem, where only coarse-grained labels are available [34]. Formally, given a training image \mathbf{I} , all instances in some form constitute a “bag”. *E.g.* object proposals (in detection task) or image pixels (in segmentation task) can be different forms of instances. If the image \mathbf{I} is labeled with class c , then the “bag” of \mathbf{I} is positive with regard to c , meaning that there is at least one positive instance of class c in this bag. If \mathbf{I} is not labeled with class c , the corresponding “bag” is negative to c and there is no instance of class c in this image. The MIL models aim at predicting the label of an input bag, and more importantly, finding positive instances in positive bags.

Weakly Supervised Object Detection. Recently, deep neural networks and MIL are incorporated and significantly improve the previous state-of-the-arts. Bilen *et al.* [2] proposed a Weakly Supervised Deep Detection Network (WSDDN) composing of two branches acting as a proposal selector and a proposal classifier, respectively. The idea, detecting objects by the attention-based selection, is proved to be so effective that most of the latter works follow it. *E.g.*, WSDDN is further improved by adding recursive refinement branches in [25]. [30, 28, 29] took advantage of the continuation optimization and progressively learned models from easy to difficult, which are very promising and effective. Besides these single-stage approaches, researchers have also considered the multiple-stage methods in which fully-supervised detectors are trained with the boxes detected by the single-stage methods as pseudo-labels. Zhang *et al.* [33] proposed a metric to estimate image difficulty with the proposal classification scores of WSDDN, and trained a Fast R-CNN with curriculum learning strategy. To speed up the weakly supervised object detectors, Shen *et al.* [19] used WSDDN as an instructor which guides a fast generator to produce similar detection results.

Weakly Supervised Object Segmentation. Another

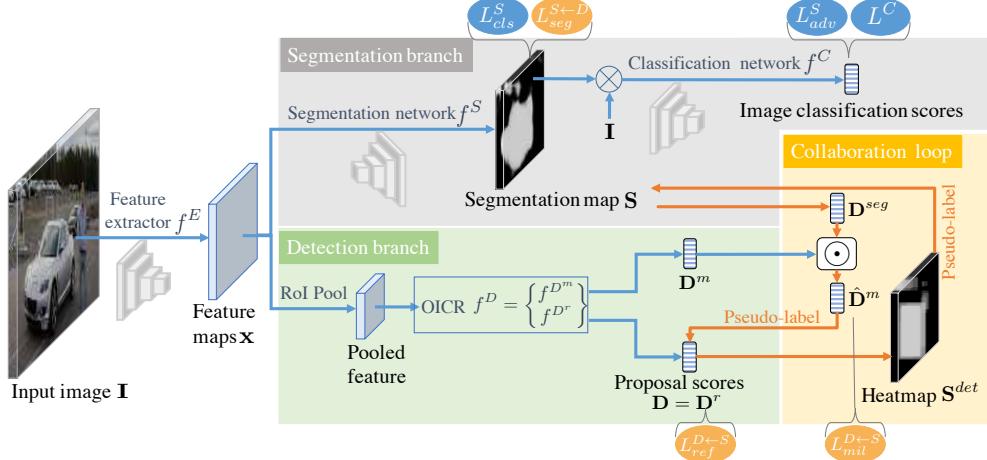


Figure 2: The overall architecture. The SDCN is composed of three modules: the feature extractor, the segmentation branch, and the detection branch. The segmentation branch is instructed by a classification network in a generative adversarial learning manner, while the detection branch employs a conventional weakly supervised detector OICR [25], guided by an MIL objective. These two branches further supervise each other in a collaboration loop. The solid ellipses denote the operations are denoted as blue arrows, while the collaboration loop is shown with orange ones.

route for localizing objects is semantic segmentation. To obtain weakly supervised segmentation map, in [17], Kolesnikov *et al.* took segmentation map as an output of the network and then aggregated it to a global classification prediction to learn with category labels. In [9], the aggregation function is improved to incorporate both negative and positive evidence, representing both the absence and presence of the target class. In [31], a recurrent adversarial erasing strategy is proposed to mask the response region of the previous forward passes and force to generate responses on other undetected parts during the current forward pass.

Utilization of Segmentation in Weakly Supervised Detection. Researchers have found that there are inherent relations between the weakly supervised segmentation and detection tasks. In [7], a segmentation branch generating coarse response maps is used to eliminate proposals unlikely to cover objects. In [32], the proposal filtering step is based on a new objectness rating TS2C defined with the weakly supervised segmentation map. Ge *et al.* [12] proposed a complex framework for both object segmentation and detection, where results from the weakly supervised segmentation models are used as both object proposal generator and filter for the latter detection models. These methods incorporate the segmentation to improve weakly supervised object detection, which are reasonable and promising given their superiorities over their baseline models. However, they ignore the mentioned complementarity of these tasks and only exploit one-way cooperation, as shown in Fig. 1a. The suboptimal manners in using the segmentation information limit the performance of their methods.

3. Method

The overall architecture of the proposed segmentation-detection collaborative network (SDCN) is shown in Fig. 2.

The network is mainly composed of three components: a backbone feature extractor f^E , a segmentation branch f^S , and a detection branch f^D . For an input image \mathbf{I} , its feature $\mathbf{x} = f^E(\mathbf{I})$ is extracted by the extractor f^E , and then feeds into f^S and f^D for segmentation and detection, respectively. The entire network is guided by the classification labels $\mathbf{y} = [y_1, y_2, \dots, y_N] \in \{0, 1\}^N$, (where N is the number of object classes), which is formatted as an adversarial classification loss and an MIL objective. Additional collaboration loss is designed for improving the accuracy of both branches in a manner of collaborative loop.

In 3.1, we first briefly introduce our detection branch, which follows the Online Instance Classifier Refinement (OICR) [25]. The proposed segmentation branch and collaboration mechanism are described in detail in 3.2 and 3.3.

3.1. Detection Branch

The detection branch f^D aims at detecting object instances in an input image, given only image category labels. The design of f^D follows the OICR [25], which works in a similar fashion to the Fast RCNN [13]. Specifically, f^D takes the feature \mathbf{x} from the backbone f^E and object proposals $\mathbf{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_B\}$ (where B is the number of proposals) from Selective Search [27] as input, and detects by classifying each proposal, formulated as below:

$$\mathbf{D} = f^D(\mathbf{x}, \mathbf{B}), \quad \mathbf{D} \in [0, 1]^{B \times (N+1)}, \quad (1)$$

where N denotes the number of classes with the $(N + 1)^{th}$ class as the background. Each element $D(i, j)$ indicates the probability of the i^{th} proposal \mathbf{b}_i belonging to the j^{th} class.

The detection branch f^D consists of two sub-modules, a multiple instance detection network (MIDN) f^{D^m} and an online instance classifier refinement module f^{D^r} . The

MIDN f^{D^m} serves as an instructor of the refinement module f^{D^r} , while f^{D^r} produces the final detection output.

The MIDN is the same as the mentioned WSDDN [2], which computes the probability of each proposal belonging to each class under the supervision of category label, with an MIL objective (in Eq. (1) of [25]) formulated as follows:

$$\mathbf{D}^m = f^{D^m}(\mathbf{x}, \mathbf{B}), \quad \mathbf{D}^m \in [0, 1]^{B \times N}, \quad (2)$$

$$L_{mil}^D = \sum_{j=1}^N L_{BCE} \left(\sum_{i=1}^B \mathbf{D}^m(i, j), \mathbf{y}(j) \right), \quad (3)$$

where $\sum_{i=1}^B \mathbf{D}^m(i, j)$ (denoted as ϕ_c in [25]) shows the probability of an input image belonging to the j^{th} category by summing up that of all proposals, and L_{BCE} denotes the standard multi-class binary cross entropy loss.

Then, the resulting probability \mathbf{D}^m from minimizing Eq. (3) is used to generate pseudo instance classification labels for the refinement module. This process is denoted as:

$$\mathbf{Y}^r = \kappa(\mathbf{D}^m), \quad \mathbf{Y}^r \in \{0, 1\}^{B \times (N+1)}. \quad (4)$$

Each binary element $\mathbf{Y}^r(i, j)$ indicates if the i^{th} proposal is labeled as the j^{th} class. κ denotes the conversion from the soft probability matrix \mathbf{D}^m to discrete instance labels \mathbf{Y}^r , where the top-scoring proposal and its highly overlapped ones are labeled as the image label and the rest are labeled as the background. Details are referred to Sec. 3.2 in [25].

The online instance classifier refinement module f^{D^r} performs detection proposal by proposal and further constrains the spatial consistency of the detection results with the generated labels \mathbf{Y}^r , which is formulated as below:

$$\mathbf{D}^r(i, :) = f^{D^r}(\mathbf{x}, \mathbf{b}_i), \quad \mathbf{D}^r \in [0, 1]^{B \times (N+1)}, \quad (5)$$

$$L_{ref}^D = \sum_{j=1}^{N+1} \sum_{i=1}^B L_{CE}(\mathbf{D}^r(i, j), \mathbf{Y}^r(i, j)), \quad (6)$$

where $\mathbf{D}^r(i, :) \in [0, 1]^{N+1}$ is a row of \mathbf{D}^r , indicating the classification scores for proposal \mathbf{b}_i . L_{CE} denotes the weighted cross entropy (CE) loss function in Eq. (4) of [25]. Here, L_{CE} is employed instead of L_{BCE} considering that each proposal has one and only one positive category label.

Eventually, the detection results are given by the refinement module, *i.e.* $\mathbf{D} = \mathbf{D}^r$, and the overall objective for the detection module is a combination of Eq. (3) and Eq. (6):

$$L^D = \lambda_{mil}^D L_{mil}^D + \lambda_{ref}^D L_{ref}^D, \quad (7)$$

where λ_{mil}^D and λ_{ref}^D are balancing factors for the loss.

After optimization according to Eq. (7), the refinement module f^{D^r} can do object detection independently by discarding the MIDN in testing.

3.2. Segmentation Branch

Generally, the MIL weakly supervised object detection module is subject to over-fitting on discriminative parts, since smaller regions with less variation are more likely

to have high consistency across the whole training set. To overcome this issue, the completeness of a detected object needs to be measured and adjusted, *e.g.* by comparing with a segmentation map. Therefore, a weakly supervised segmentation branch is proposed to cover the complete object regions with generative adversarial localization strategy.

In detail, the segmentation branch f^S takes the feature \mathbf{x} as input and predicts a segmentation map, as below,

$$\mathbf{S} = f^S(\mathbf{x}), \quad \mathbf{S} \in [0, 1]^{(N+1) \times h \times w}, \quad (8)$$

$$\mathbf{s}_k \triangleq \mathbf{S}(k, :, :), \quad k \in \{1, \dots, N+1\}, \quad \mathbf{s}_k \in [0, 1]^{h \times w} \quad (9)$$

where \mathbf{S} has $N+1$ channels. Each channel \mathbf{s}_k corresponds to a segmentation map for the k^{th} class with a size of $h \times w$.

To ensure that the segmentation map \mathbf{S} covers the complete object regions precisely, a novel generative adversarial localization strategy is designed as adversarial training between the segmentation predictor f^S and an independent image classifier f^C , severing as generator and discriminator respectively, as shown in Fig. 2. The training target of the generator f^S is to fool f^C into misclassifying by masking out the object regions, and the discriminator f^C aims to eliminate the effect of the erased regions and correctly predict the category labels. The f^S and f^C are optimized alternatively, given the other one fixed.

Here, we first introduce the optimization of the segmentation branch f^S , given the classifier f^C fixed. Overall, the objective of the segmentation branch f^S can be formulated as a sum of losses for each class,

$$L^S(\mathbf{S}) = L^S(\mathbf{s}_1) + L^S(\mathbf{s}_2) + \dots + L^S(\mathbf{s}_{N+1}). \quad (10)$$

$L^S(\mathbf{s}_k)$ is the loss for the i^{th} channel of the segmentation map, consisting of an adversarial loss L_{adv}^S and a classification loss L_{cls}^S , described in detail as following.

If the k^{th} class is a **positive foreground class**¹, the segmentation map \mathbf{s}_k should fully cover the region of the k^{th} class, but does not overlap with the regions of the other classes. In other words, for an accurate \mathbf{s}_k , only the object region masked out by \mathbf{s}_k should be classified as the k^{th} class, while its complementary region should not. Formally, this expectation can be satisfied by minimizing the function

$$\begin{aligned} L_{adv}^S(\mathbf{s}_k) = & L_{BCE}(f^C(\mathbf{I} * \mathbf{s}_k), \tilde{\mathbf{y}}) + \\ & L_{BCE}(f^C(\mathbf{I} * (1 - \mathbf{s}_k)), \hat{\mathbf{y}}), \end{aligned} \quad (11)$$

where $*$ denotes pixel-wise product. The first term represents that the object region covered by the generated segmentation map, *i.e.* $\mathbf{I} * \mathbf{s}_k$, should be recognized as the k^{th} class by the classifier f^C , but does not respond to any other classes with the label $\tilde{\mathbf{y}} \in \{0, 1\}^N$, where $\tilde{\mathbf{y}}(k) = 1$ and $\tilde{\mathbf{y}}(i \neq k) = 0$. The second term means that when the

¹A positive foreground class means that the foreground class presents in the current image, while a negative one means that it does not appear.

region related to the k^{th} class is masked out from the input, *i.e.* $\mathbf{I} * (1 - \mathbf{s}_k)$, the classifier f^C should not recognize the k^{th} class anymore without influence on the other classes, with the label $\hat{\mathbf{y}} \in \{0, 1\}^N$, where $\hat{\mathbf{y}}(k) = 0$ and $\hat{\mathbf{y}}(i \neq k) = \mathbf{y}(i \neq k)$. Here, we note that generally the mask can be applied to the image \mathbf{I} or the input of any layer of the classifier f^C , and since f^C is fixed, the loss function in Eq. (11) only penalizes the segmentation branch f^S .

If the k^{th} class is a **negative foreground class**, the \mathbf{s}_k should be all-zero, as no instance of this foreground class presents. This is restrained with a response constraint term. In this term, the top 20% response pixels of each map \mathbf{s}_k are pooled and averaged for a classification predication optimized with a binary cross entropy loss as below,

$$L_{cls}^S(\mathbf{s}_k) = L_{BCE}(\text{avgpool}_{20\%}\mathbf{s}_k, \mathbf{y}(k)). \quad (12)$$

If the k^{th} class is labeled as negative, $\text{avgpool}_{20\%}\mathbf{s}_k$ is enforced to be close to 0, *i.e.* all elements of the map \mathbf{s}_k should approximately be 0. However, the above loss is also applicable when the k^{th} class is positive, $\text{avgpool}_{20\%}\mathbf{s}_k$ should be close to 1, agreeing with the constraint in Eq. (11).

The **background** is taken as a special case. In Eq. (11), though the labels $\tilde{\mathbf{y}}$ and $\hat{\mathbf{y}}$ do not involve the background class, the background segmentation map \mathbf{s}_{N+1} is also applicable same as the other classes. When \mathbf{s}_{N+1} is multiplied as the first term in Eq. (11), the target label should be all-zero $\tilde{\mathbf{y}} = \mathbf{0}$; when $1 - \mathbf{s}_{N+1}$ is used as the mask in the second term of Eq. (11), the target label should be exactly the same as the original label $\hat{\mathbf{y}} = \mathbf{y}$. For Eq. (12), we assume that a background region always appears in any input image, *i.e.* $\mathbf{y}(N+1) = 1$ for all images.

Overall, the total loss of the segmentation branch in Eq. (10) can be summarized and rewritten as follows,

$$L^S = \lambda_{adv}^S \sum_{k \text{ if } \mathbf{y}(k)=1} L_{adv}^S(\mathbf{s}_k) + \lambda_{cls}^S \sum_{k=1}^{N+1} L_{cls}^S(\mathbf{s}_k), \quad (13)$$

where λ_{adv}^S and λ_{cls}^S denote balance weights.

After optimizing Eq. (13), following the adversarial manner, the segmentation branch f^S is fixed, and the classifier f^C is further optimized with the following objective,

$$L_{adv}^C(\mathbf{s}_k) = L_{BCE}(f^C(\mathbf{I} * (1 - \mathbf{s}_k)), \mathbf{y}), \quad (14)$$

$$L^C = L_{BCE}(f^C(\mathbf{I}), \mathbf{y}) + \sum_{k \text{ if } \mathbf{y}(k)=1} L_{adv}^C(\mathbf{s}_k). \quad (15)$$

The objective L^C consists of a classification loss and an adversarial loss L_{adv}^C . The target of the classifier f^C should always be \mathbf{y} , since it aims at digging out the remaining object regions, even if \mathbf{s}_k is masked out.

Our idea for designing the segmentation branch shares the same adversarial spirit with [31], but our design is more efficient compared with [31] that recurrently performs several forward passes for one segmentation map. Besides, we

do not have the trouble of deciding number recurrent steps as [31], which may vary with different objects.

3.3. Collaboration Mechanism

A dynamic collaboration loop is designed to complement both detection and segmentation for more accurate predictions, namely neither so large that cover the background nor so small that degenerate to object parts.

Segmentation Instructs Detection. As mentioned, the detection branch is easy to over-fit to discriminative parts, while the segmentation can cover the whole object region. So naturally, the segmentation map can be used to refine the detection results by making the proposal having a larger IoU with the corresponding segmentation map have a higher score. This is achieved by re-weighting the instance classification probability matrix \mathbf{D}^m in Eq. (2) in the detection branch by using a prior probability matrix \mathbf{D}^{seg} stemming from the segmentation map as follows,

$$\hat{\mathbf{D}}^m = \mathbf{D}^m \odot \mathbf{D}^{seg}, \quad (16)$$

where $\mathbf{D}^{seg}(i, k)$ denotes the overlap degree between the i^{th} object proposal and the connected regions from the k^{th} segmentation map. \mathbf{D}^{seg} is generated as below:

$$\mathbf{D}^{seg}(i, k) = \max_j \text{IoU}(\hat{\mathbf{s}}_{kj}, \mathbf{b}_i) + \tau_0. \quad (17)$$

Here, $\hat{\mathbf{s}}_{kj}$ denotes the j^{th} connected component under threshold T_c in the segmentation map \mathbf{s}_k , and $\text{IoU}(\hat{\mathbf{s}}_{kj}, \mathbf{b}_i)$ denotes the intersection over union between $\hat{\mathbf{s}}_{kj}$ and the object proposal \mathbf{b}_i . The constant τ_0 adds a fault-tolerance for the segmentation branch. Each column of \mathbf{D}^{seg} is normalized by its maximum value, to make it range within [0, 1].

With the re-weighting in Eq. (16), the object proposals only focusing on local parts are assigned with lower weights, while those proposals precisely covering the object stand out. The connected components are employed to alleviate the issue of multiple instance occurrences, which is a hard case for weakly supervised object detection. The recent TS2C [32] objectness rating designed for solving this issue is also tested in place of IoU with connected components, but no superiority shows in our case.

The re-weighted probability matrix $\hat{\mathbf{D}}^m$ replaces \mathbf{D}^m in Eq. (3) and further instructs the MIDN as in Eq. (18) and the refinement module as in Eq. (19):

$$L_{mil}^{D \leftarrow S} = \sum_j L_{BCE} \left(\sum_i \hat{\mathbf{D}}^m(i, j), \mathbf{y}(j) \right), \quad (18)$$

$$L_{ref}^{D \leftarrow S} = \sum_j \sum_i L_{CE} \left(\mathbf{D}^r(i, j), \hat{\mathbf{Y}}^r(i, j) \right), \quad (19)$$

where $\hat{\mathbf{Y}}^r$ denotes the pseudo labels deriving from $\hat{\mathbf{D}}^m$ as that in Eq. (4). Finally, the overall objective of the detection branch in Eq. (7) is reformulated as below,

$$L^{D \leftarrow S} = \lambda_{mil}^D L_{mil}^{D \leftarrow S} + \lambda_{ref}^D L_{ref}^{D \leftarrow S}. \quad (20)$$

Detection Instructs Segmentation. Though the detec-

tion boxes may not cover the whole object, they are effective for distinguishing an object from the background. To guide the segmentation branch, a detection heatmap $\mathbf{S}^{det} \in [0, 1]^{(N+1) \times h \times w}$ is generated, which can be seen as an analog of the segmentation map. Each channel $\mathbf{s}_k^{det} \triangleq \mathbf{S}^{det}(k, :, :)$ corresponds to a heatmap for the k^{th} class. Specifically, for the positive class k , each proposal box contributes its classification score to all pixels within this proposal and thus generates the \mathbf{s}_k^{det} by

$$\mathbf{s}_k^{det}(p, q) = \sum_{i \text{ if } (p, q) \in \mathbf{b}_i} \mathbf{D}(i, k), \quad (21)$$

while the other \mathbf{s}_k^{det} corresponding to negative classes are set to zero. Then, \mathbf{s}_k^{det} is normalized by its maximum response and the background heatmap \mathbf{s}_{N+1}^{det} can be simply calculated as the complementary set of the foreground, *i.e.*

$$\mathbf{s}_{N+1}^{det} = 1 - \max_{k \in \{1, \dots, N\}} \mathbf{s}_k^{det}. \quad (22)$$

To generate pseudo category label for each pixel, the soft segmentation map \mathbf{S}^{det} is first discretized by taking the arguments of the maxima at each pixel and then the top 10% pixels for each class are kept, while other ambiguous ones are ignored. The generated label is denoted by $\psi(\mathbf{S}^{det})$, and the instructive loss is formulated as below:

$$L_{seg}^{S \leftarrow D} = L_{CE}(\mathbf{S}, \psi(\mathbf{S}^{det})). \quad (23)$$

Therefore, the loss function of the whole segmentation branch in Eq. (13) is now updated to

$$L^{S \leftarrow D} = L^S + \lambda_{seg}^S L_{seg}^{S \leftarrow D}. \quad (24)$$

Overall Objective. With the updates in Eq. (20) and Eq. (24), the final objective for the entire network is

$$\operatorname{argmin}_{f^E, f^S, f^D} L = L^{S \leftarrow D} + L^{D \leftarrow S}. \quad (25)$$

Briefly, the above objective is optimized in an end-to-end manner. The image classifier f^C is optimized with the loss L^C alternatively, as most adversarial methods. The optimization can be easily conducted using gradient descent. For clarity, the training and the testing of our SDCN are summarized in Algorithm 1.

In the testing stage, as shown in Algorithm 1, only the feature extractor f^E and the refinement module f^{D^r} are needed, which make our method as efficient as [25].

4. Experiments

We evaluate the proposed segmentation-detection collaborative network (SDCN) for weakly supervised object detection to prove its advantages over the state-of-the-arts.

4.1. Experimental Setup

Datasets. The evaluation is conducted on two commonly used datasets for weakly supervised detection, including the PASCAL VOC 2007 [11] and 2012 [10]. The VOC 2007

Algorithm 1 Training and Testing SDCN

Input: training set with category labels $\mathbf{T}_1 = \{\mathbf{I}, \mathbf{y}\}$.

- 1: **procedure** TRAINING
- 2: forward SDCN $f^E(\mathbf{I}) \rightarrow \mathbf{x}$, $f^D(\mathbf{x}) \rightarrow \mathbf{D}$, $f^S(\mathbf{x}) \rightarrow \mathbf{S}$,
- 3: forward the classifier $f^C(\mathbf{s}_k * \mathbf{I})$ and $f^C((1 - \mathbf{s}_k) * \mathbf{I})$,
- 4: generate variables \mathbf{D}^{seg} and \mathbf{S}^{det} with \mathbf{S} and \mathbf{D} ,
- 5: compute $L^{D \leftarrow S}$ in Eq.(20) and $L^{S \leftarrow D}$ in Eq.(24),
- 6: backward the loss $L = L^{D \leftarrow S} + L^{S \leftarrow D}$ for SDCN,
- 7: compute and backward the loss L^C for f^C ,
- 8: continue until convergence.

Output: the optimized SDCN (f^E and f^D) for detection.

Input: test set $\mathbf{T}_2 = \{\mathbf{I}\}$.

- 1: **procedure** TESTING
- 2: forward SDCN $f^E(\mathbf{I}) \rightarrow \mathbf{x}$, $f^{D^r}(\mathbf{x}) \rightarrow \mathbf{D}$,
- 3: post-process for detected bounding boxes with \mathbf{D} .

Output: the detected object bounding boxes for \mathbf{T}_2 .

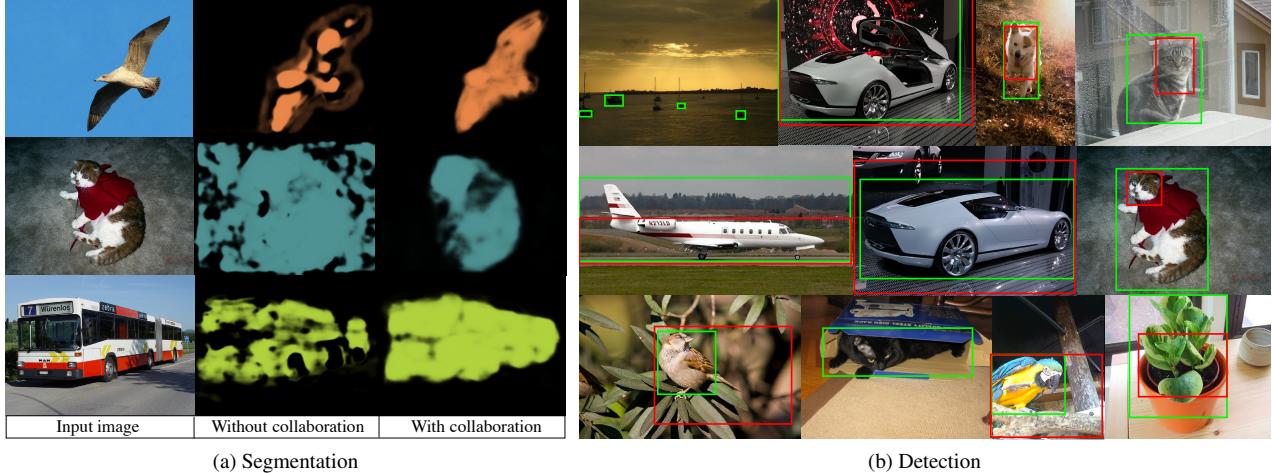
dataset includes 9,963 images with total 24,640 objects in 20 classes. It is divided into a *trainval* set with 5,011 images and a *test* set with 4,952 images. The more challenging VOC 2012 dataset consists of 11,540 images with 27,450 objects in *trainval* set and 10,991 images for test. In our experiments, the *trainval* split is used for training and the *test* set is for testing. The performance is reported in terms of two metrics: 1) correct localization (CorLoc) [6] on the *trainval* split and 2) average precision (AP) on the *test* set.

Implementation. For the backbone network f^E , we use the VGG-16 [20]. For f^D , the same architecture as that in OICR [32] is employed. For f^S , similar segmentation header to the CPN [4] is adopted. For the adversarial classifier f^C , ResNet-101 [15] is used and the segmentation masking operation is applied after the res4b22 layer.

We follow a three-step training strategy: 1) the classifier f^C is trained with a fixed learning rate 5×10^{-4} until its convergence; 2) the segmentation branch f^S and detection branch f^D are pre-trained without collaboration; 3) the entire architecture is trained following the end-to-end manner. The SDCN runs for 40k iterations with learning rate 10^{-3} , following 30k iterations with learning rate 10^{-4} . The same multi-scale training and testing strategies in OICR [25] are adopted. To achieve balanced impacts between detection and segmentation branches, the weights of the losses are simply set to make the gradients have similar scales, *i.e.* $\lambda_{adv}^S = 1$, $\lambda_{cls}^S = 0.1$, $\lambda_{seg}^S = 0.1$, $\lambda_{mil}^D = 1$ and $\lambda_{ref}^D = 1$, respectively. The constant τ_0 in Eq. (17) and the threshold T_c is empirically set to 0.5 and 0.1, respectively.

4.2. Ablation Studies

Our ablation study is conducted on VOC 2007 dataset. Five weakly supervised strategies are compared and the results are shown in Table 2. The baseline detection method without the segmentation branch is the same as



(a) Segmentation

(b) Detection

Figure 3: Visualization of the segmentation and the detection results without and with collaboration. In (a), the columns from left to right are the original images, the segmentation map obtained without and with the collaboration loop. In (b), the detection results of OICR[25] without consideration of collaboration, and the proposed method with collaboration loop are shown with red and green boxes, respectively. (Absence of boxes means no detected object given the detection threshold.)

the OICR[32]. Another naive consideration is directly including the detection and segmentation modules in a multi-task manner without any collaboration between them. The model where only segmentation branch instructs detection branch is also tested. Its mAP is the lowest, since the mean intersection over union (mIoU) between the segmentation results and the ground-truth drops from 37% to 25.1% without the guidance of detection branch, which proves that these two branches should not collaborate in one-way. The model can also be trained without the generative adversarial localization strategy, but its performance drops. Our full method achieves the highest mAP. It can be observed that the proposed method improves all baseline models by large margins, demonstrating the effectiveness and necessity of the generative adversarial localization strategy and the collaboration loop.

The segmentation masks and detection results without and with the collaboration are visualized in Fig. 3. As observed in Fig. 3a, with the instruction from the detection branch, the segmentation map becomes much more precise with fewer confusions between the background and the class-related region. Similarly, as shown in Fig. 3b, the baseline approach inclines to mix discriminative parts with target object bounding boxes, while with the guidance from segmentation the more complete objects are detected. The visualization clearly illustrates the benefits to each other.

4.3. Comparisons With State-of-the-Arts

All comparison methods are first evaluated on VOC 2007 as shown in Table 3 and Table 4 in terms of mAP and CorLoc. Among single-stage methods, our method outperforms others on the most categories, leading to a notable improvement on average. Especially, our method performs much better than the state-of-the-arts on “boat”, “cat”, “dog”, as

Det. branch	Seg. branch	Seg. → Det.	Det. → Seg.	Adv. Loc.	mAP
✓					41.2
✓	✓			✓	41.3
✓	✓	✓		✓	36.8
✓	✓	✓	✓		46.0
✓	✓	✓	✓	✓	50.2

Table 2: mAP (in %) of different weakly supervised strategies with the same backbone on the VOC 2007 dataset.

our approach leans to detect more complete objects, though in most cases instances of these categories can be identified by parts. Moreover, our method produces significant improvements compared with the OICR[25] with exactly the same architecture. The most competitive method [26] is designed for weakly supervised object proposal, which is not really competing but complementary to our method, and replacing the fixed object proposal in our method with [26] potentially improves the performance. Besides, the performance of our single-stage method is even comparable with the multiple-stage methods [25, 32, 33, 30], illustrating the effectiveness of the proposed dynamic collaboration loop.

Furthermore, all methods can be enhanced by training with multiple stages, as shown at the bottom of Table 3. Following [32], the top scoring detection bounding boxes from SDCN is used as the labels for training a Fast RCNN [13] with the backbone of VGG16, denoted as SDCN+FRCNN. By this simple multi-stage training strategy, the performance can be further boosted to 53.7%, which surpasses all the state-of-the-art multiple-stage methods, though [25, 26] use more complex ensemble models. It is noted that the approaches, *e.g.* HCP+DSD+OSSH3[16] and ZLDN-L[33], attempt to design more elaborate training mechanism by using self-paced or curriculum learning. We believe that the performance of our model SDCN+FRCNN can be further

Methods	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
Single-stage																					
WSDDN-VGG16 [2]	39.4	50.1	31.5	16.3	12.6	64.5	42.8	42.6	10.1	35.7	24.9	38.2	34.4	55.6	9.4	14.7	30.2	40.7	54.7	46.9	34.8
OICR-VGG 16 [25]	58.0	62.4	31.1	19.4	13.0	65.1	62.2	28.4	24.8	44.7	30.6	25.3	37.8	65.5	15.7	24.1	41.7	46.9	64.3	62.6	41.2
MELM-L+RL[30]	50.4	57.6	37.7	23.2	13.9	60.2	63.1	44.4	24.3	52.0	42.3	42.7	43.7	66.6	2.9	21.4	45.1	45.2	59.1	56.2	42.6
TS2C [32]	59.3	57.5	43.7	27.3	13.5	63.9	61.7	59.9	24.1	46.9	36.7	45.6	39.9	62.6	10.3	23.6	41.7	52.4	58.7	56.6	44.3
[26]	57.9	70.5	37.8	5.7	21.0	66.1	69.2	59.4	3.4	57.1	57.3	35.2	64.2	68.6	32.8	28.6	50.8	49.5	41.1	30.0	45.3
SDCN (ours)	59.4	71.5	38.9	32.2	21.5	67.7	64.5	68.9	20.4	49.2	47.6	60.9	55.9	67.4	31.2	22.9	45.0	53.2	60.9	64.4	50.2
Multiple-stage																					
WSDDN-Ens. [2]	46.4	58.3	35.5	25.9	14.0	66.7	53.0	39.2	8.9	41.8	26.6	38.6	44.7	59.0	10.8	17.3	40.7	49.6	56.9	50.8	39.3
HCP+DSD+OSSH3[16]	52.2	47.1	35.0	26.7	15.4	61.3	66.0	54.3	3.0	53.6	24.7	43.6	48.4	65.8	6.6	18.8	51.9	43.6	53.6	62.4	41.7
OICR-Ens.+FRCNN[25]	65.5	67.2	47.2	21.6	22.1	68.0	68.5	35.9	5.7	63.1	49.5	30.3	64.7	66.1	13.0	25.6	50.0	57.1	60.2	59.0	47.0
MELM-L2+ARL[30]	55.6	66.9	34.2	29.1	16.4	68.8	68.1	43.0	25.0	65.6	45.3	53.2	49.6	68.6	2.0	25.4	52.5	56.8	62.1	57.1	47.3
ZLDN-L[33]	55.4	68.5	50.1	16.8	20.8	62.7	66.8	56.5	2.1	57.8	47.5	40.1	69.7	68.2	21.6	27.2	53.4	56.1	52.5	58.2	47.6
TS2C+FRCNN [32]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	48.0
Ens.+FRCNN[26]	63.0	69.7	40.8	11.6	27.7	70.5	74.1	58.5	10.0	66.7	60.6	34.7	75.7	70.3	25.7	26.5	55.4	56.4	55.5	54.9	50.4
SDCN+FRCNN (ours)	59.8	75.1	43.3	31.7	22.8	69.1	71.0	72.9	21.0	61.1	53.9	73.1	54.1	68.3	37.6	20.1	48.2	62.3	67.2	61.1	53.7

Table 3: Average precision (in %) for our method and the state-of-the-arts on VOC 2007 *test* split.

Methods	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	CorLoc
Single-stage																					
WSDDN-VGG16 [2]	65.1	58.8	58.5	33.1	39.8	68.3	60.2	59.6	34.8	64.5	30.5	43.0	56.8	82.4	25.5	41.6	61.5	55.9	65.9	63.7	53.5
OICR-VGG16 [25]	81.7	80.4	48.7	49.5	32.8	81.7	85.4	40.1	40.6	79.5	35.7	33.7	60.5	88.8	21.8	57.9	76.3	59.9	75.3	81.4	60.6
TS2C [32] [26]	84.2	74.1	61.3	52.1	32.1	76.7	82.9	66.6	42.3	70.6	39.5	57.0	61.2	88.4	9.3	54.6	72.2	60.0	65.0	70.3	61.0
SDCN (ours)	85.0	83.9	58.9	59.6	43.1	79.7	85.2	77.9	31.3	78.1	50.6	75.6	76.2	88.4	49.7	56.4	73.2	62.6	77.2	79.9	68.6
Multiple-stage																					
HCP+DSD+OSSH3[16]	72.7	55.3	53.0	27.8	35.2	68.6	81.9	60.7	11.6	71.6	29.7	54.3	64.3	88.2	22.2	53.7	72.2	52.6	68.9	75.5	56.1
WSDDN-Ens. [2]	68.9	68.7	65.2	42.5	40.6	72.6	75.2	53.7	29.7	68.1	33.5	45.6	65.9	86.1	27.5	44.9	76.0	62.4	66.3	66.8	58.0
MELM-L2+ARL[30]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	61.4
ZLDNN-L[33]	74.0	77.8	65.2	37.0	46.7	75.8	83.7	58.8	17.5	73.1	49.0	51.3	76.7	87.4	30.6	47.8	75.0	62.5	64.8	68.8	61.2
OICR-Ens.+FRCNN[25]	85.8	82.7	62.8	45.2	43.5	84.8	87.0	46.8	15.7	82.2	51.0	45.6	83.7	91.2	22.2	59.7	75.3	65.1	76.8	78.1	64.3
Ens.+FRCNN[26]	83.8	82.7	60.7	35.1	53.8	82.7	88.6	67.4	22.0	86.3	68.8	50.9	90.8	93.6	44.0	61.2	82.5	65.9	71.1	76.7	68.4
SDCN+FRCNN (ours)	85.0	86.7	60.7	62.8	46.6	83.2	87.8	81.7	35.8	80.8	57.4	81.6	79.9	92.4	59.3	57.5	79.4	68.5	81.7	81.4	72.5

Table 4: CorLoc (in %) for our method and the state-of-the-arts on VOC 2007 *trainval* split.

Methods		mAP	CorLoc
Single-stage	OICR-VGG16 [25]	37.9	62.1
	TS2C [32]	40.0	64.4
	[26]	40.8	64.9
	SDCN (ours)	43.5	67.9
Multiple-stage	MELM-L2+ARL[30]	42.4	—
	OICR-Ens.+FRCNN [25]	42.5	65.6
	ZLDN-L[33]	42.9	61.5
	TS2C+FRCNN [32]	44.4	—
	Ens.+FRCNN[26]	45.7	69.3
	SDCN+FRCNN (ours)	46.7	69.5

Table 5: mAP and CorLoc (in %) for our method and the state-of-the-arts on VOC 2012 *trainval* split.

improved by adopting such algorithms. The comparison methods are further evaluated on the more challenging VOC 2012 dataset, as shown in Table 5. As expected, the proposed method achieves significant improvements with the same architecture as [25, 32], demonstrating its superiority.

Overall, our SDCN significantly improves the performance of weakly supervised object detection on average, benefitting from the deep collaboration of segmentation and detection. However, there are still several classes on which the performance is relatively lower as shown in Table 3, e.g. “chair” and “bottle”. The main reason is the large portion of occluded and overlapped samples for these classes, which leads to incomplete or connected responses on the segmen-

tation map and bad interaction with the detection branch, leaving room for further improvements.

Time Cost. Our training speed is roughly $2\times$ slower than that of the baseline OICR [25], but the testing time costs of our method and OICR are the same, since they share exactly the same architecture of the detection branch.

5. Conclusions and Future Work

In this paper, we present a novel segmentation-detection collaborative network (SDCN) for weakly supervised object detection. Different from the previous works, our method exploits a collaboration loop between segmentation task and detection task to combine the merits of both. Extensive experimental results safely reach the conclusion that our method successfully exceeds the previous state-of-the-arts, while it keeps efficiency in the inference stage. The design of SDCN may be more elaborate for densely overlapped or partially occluded objects, which is more challenging and left as future work.

Acknowledgement: This work is partially supported by the National Key R&D Program of China under contract No. 2017YFA0700800, and Natural Science Foundation of China under contract No. 61772496.

References

- [1] Hakan Bilen, Marco Pedersoli, and Tinne Tuytelaars. Weakly supervised object detection with posterior regularization. In *British Machine Vision Conference (BMVC)*, pages 1–12, 2014.
- [2] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2846–2854, 2016.
- [3] Matthew Blaschko, Andrea Vedaldi, and Andrew Zisserman. Simultaneous object detection and ranking with weak supervision. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 235–243, 2010.
- [4] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded Pyramid Network for Multi-Person Pose Estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7103–7112, 2018.
- [5] Thomas Deselaers, Bogdan Alexe, and Vittorio Ferrari. Localizing objects while learning their appearance. In *European Conference on Computer Vision (ECCV)*, pages 452–466, 2010.
- [6] Thomas Deselaers, Bogdan Alexe, and Vittorio Ferrari. Weakly supervised localization and learning with generic knowledge. *International Journal of Computer Vision (IJCV)*, 100(3):275–293, 2012.
- [7] Ali Diba, Vivek Sharma, Ali Mohammad Pazandeh, Hamed Pirsiavash, and Luc Van Gool. Weakly supervised cascaded convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 914–922, 2017.
- [8] Thomas G. Dietterich, Richard H. Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial Intelligence.*, 89(1-2):31–71, 1997.
- [9] Thibaut Durand, Taylor Mordan, Nicolas Thome, and Matthieu Cord. WILDCAT: Weakly Supervised Learning of Deep ConvNets for Image Classification, Pointwise Localization and Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 642–651, 2017.
- [10] 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 (IJCV)*, 111(1):98–136, 2015.
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision (IJCV)*, 88(2):303–338, 2010.
- [12] Weifeng Ge, Sibei Yang, and Yizhou Yu. Multi-evidence filtering and fusion for multi-label classification, object detection and semantic segmentation based on weakly supervised learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1277–1286, 2018.
- [13] Ross Girshick. Fast r-cnn. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1440–1448, 2015.
- [14] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Multi-fold mil training for weakly supervised object localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2409–2416, 2014.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [16] Zequn Jie, Yunchao Wei, Xiaojie Jin, Jiashi Feng, and Wei Liu. Deep self-taught learning for weakly supervised object localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1377–1385, 2017.
- [17] Alexander Kolesnikov and Christoph H. Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European Conference on Computer Vision (ECCV)*, pages 695–711, 2016.
- [18] Olga Russakovsky, Yuanqing Lin, Kai Yu, and Li Fei-Fei. Object-centric spatial pooling for image classification. In *European Conference on Computer Vision (ECCV)*, pages 1–15, 2012.
- [19] Yunhan Shen, Rongrong Ji, Shengchuan Zhang, Wangmeng Zuo, and Yan Wang. Generative adversarial learning towards fast weakly supervised detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5764–5773, 2018.
- [20] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [21] Parthipan Siva, Chris Russell, and Tao Xiang. In defence of negative mining for annotating weakly labelled data. In *European Conference on Computer Vision (ECCV)*, pages 594–608, 2012.
- [22] Parthipan Siva and Tao Xiang. Weakly supervised object detector learning with model drift detection. In *IEEE International Conference on Computer Vision (ICCV)*, pages 343–350, 2011.
- [23] Hyun Oh Song, Ross Girshick, Stefanie Jegelka, Julien Mairal, Zaid Harchaoui, and Trevor Darrell. On learning to localize objects with minimal supervision. In *International Conference on Machine Learning (ICML)*, pages 1611–1619, 2014.
- [24] Hyun Oh Song, Yong Jae Lee, Stefanie Jegelka, and Trevor Darrell. Weakly-supervised discovery of visual pattern configurations. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1637–1645, 2014.
- [25] Peng Tang, Xinggang Wang, Xiang Bai, and Wenyu Liu. Multiple instance detection network with online instance classifier refinement. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2843–2851, 2017.
- [26] Peng Tang, Xinggang Wang, Angtian Wang, Yongluan Yan, Wenyu Liu, Junzhou Huang, and Alan Yuille. Weakly supervised region proposal network and object detection. In *European Conference on Computer Vision (ECCV)*, pages 370–386, 2018.

- [27] Jasper R R Uijlings, Koen E A Van De Sande, Theo Gevers, and Arnold W M Smeulders. Selective search for object recognition. *International Journal of Computer Vision (IJCV)*, 104(2):154–171, 2013.
- [28] Fang Wan, Chang Liu, Wei Ke, Xiangyang Ji, Jianbin Jiao, and Qixiang Ye. C-MIL: Continuation multiple instance learning for weakly supervised object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2199–2208, 2019.
- [29] Fang Wan, Pengxu Wei, Zhenjun Han, Jianbin Jiao, and Qixiang Ye. Min-entropy latent model for weakly supervised object detection. *IEEE Transactions on Pattern Analysis & Machine Intelligence (TPAMI)*, 2019.
- [30] Fang Wan, Pengxu Wei, Jianbin Jiao, Zhenjun Han, and Qixiang Ye. Min-entropy latent model for weakly supervised object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1297–1306, 2018.
- [31] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1568–1576, 2017.
- [32] Yunchao Wei, Zhiqiang Shen, Bowen Cheng, Honghui Shi, Jinjun Xiong, Jiashi Feng, and Thomas Huang. TS²C: tight box mining with surrounding segmentation context for weakly supervised object detection. In *European Conference on Computer Vision (ECCV)*, pages 454–470, 2018.
- [33] Xiaopeng Zhang, Jiashi Feng, Hongkai Xiong, and Qi Tian. Zigzag learning for weakly supervised object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4262–4270, 2018.
- [34] Zhi-Hua Zhou. A brief introduction to weakly supervised learning. *National Science Review*, 5(1):44–53, 2018.