This ICCV paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

DISCOBOX: Weakly Supervised Instance Segmentation and Semantic Correspondence from Box Supervision

Shiyi Lan^{1*} Zhiding Yu^{2†} Christopher Choy² Subhashree Radhakrishnan² Guilin Liu² Yuke Zhu^{2,3} Larry S. Davis¹ Anima Anandkumar^{2,4} ¹University of Maryland, College Park ²NVIDIA ³The University of Texas at Austin ⁴Caltech

Abstract

We introduce DiscoBox, a novel framework that jointly learns instance segmentation and semantic correspondence using bounding box supervision. Specifically, we propose a self-ensembling framework where instance segmentation and semantic correspondence are jointly guided by a structured teacher in addition to the bounding box supervision. The teacher is a structured energy model incorporating a pairwise potential and a cross-image potential to model the pairwise pixel relationships both within and across the boxes. Minimizing the teacher energy simultaneously yields refined object masks and dense correspondences between intra-class objects, which are taken as pseudo-labels to supervise the task network and provide positive/negative correspondence pairs for dense contrastive learning. We show a symbiotic relationship where the two tasks mutually benefit from each other. Our best model achieves 37.9% AP on COCO instance segmentation, surpassing prior weakly supervised methods and is competitive to supervised methods. We also obtain state of the art weakly supervised results on PASCAL VOC12 and PF-PASCAL with real-time inference.

1. Introduction

The ability to localize and recognize objects is at the core of human vision. This has motivated the vision community to study object detection [1] as a fundamental visual recognition task. Instance segmentation [2] is further introduced on top of detection to predict the foreground object masks, thus enabling localization with pixel-level accuracy. More recently, a growing number of works aim to lift the above tasks to the 3D space [3-7]. As a result, landmark [3,8] and (semantic) correspondence [9-30] have been widely studied to associate object parts across different views. These methods have become critical components in pose estimation [31-34] and reconstruction [35-38] because they help



Figure 1. Task overview. Given a pair of images, DISCOBOX simultaneously outputs detection, instance segmentation and semantic correspondence predictions. Best viewed in color.

to reduce uncertainties through additional constraints, such as determining camera poses and viewpoints [31, 36].

Among various correspondence tasks, semantic correspondence aims to establish the associations across different scenes and object instances, and is arguably the most challenging one due to large variations in appearance and pose. The literature of semantic correspondence and instance segmentation have largely remained decoupled. For instance, the main semantic correspondence benchmarks [19, 39–43] have been focusing on object-centric scenarios which de-emphasizes the role of object localization, while the latest instance segmentation methods do not make use of intra-class correspondences. However, these seemingly separate problems can benefit from each other because associating object parts requires understanding the object of interest a priori. Similarly, knowing the semantic parts of an object requires understanding the geometry of functional parts and can improve object localization [44, 45].

^{*}Work done during an internship at NVIDIA Research.

[†]Corresponding author: <zhidingy@nvidia.com>.

Even though the advantage of learning correspondences and instance segmentation jointly is clear, many state of the art methods do not make use of this approach due to the lack of large scale datasets with both masks and correspondences. To overcome this challenge, weakly supervised methods have been recently introduced to relax the need for costly supervision in both tasks [25–30, 46–49]. Our work is aligned with these efforts as we aim to address instance segmentation and semantic correspondence jointly with inexpensive bounding box supervision. This allows us to effectively push the boundaries with more data.

More importantly, box supervision presents a principled way to couple the above two tasks: First, instance segmentation greatly extends the capability of semantic correspondence to handle multi-object scenarios. This allows one to define a more generalized and challenging semantic correspondence task where the performance emphasizes both the quality of object-level correspondence and the accuracy of object localization. Second, multi-tasking provides the mutual constraints to overcome trivial solutions in box supervision. Indeed, our study shows a symbiotic relation where localization benefits correspondence via improved locality and representation, whereas correspondence in turn helps localization with additional cross-image information.

We propose DISCOBOX, a framework which instantiates the above targets as shown in Fig. 1. DISCOBOX leverages various levels of structured knowledge and self-supervision both within and across images to reduce the uncertainties.

Summary of contributions:

- Our work is the first to propose a unified framework for joint weakly supervised instance segmentation and semantic correspondence using bounding box supervision.
- We propose a novel self-ensembling framework where a teacher is designed to promote structured inductive bias and establish correspondences across objects. We show that the proposed framework allows us to jointly exploit both intra- and cross-image self-supervisions and leads to significantly improved task performance.
- We achieve state-of-the-art performance on weakly supervised instance segmentation. Our best model achieves 37.9% AP on COCO test-dev, surpassing competitive supervised methods such as YOLACT++ [50] (34.6% AP) and Mask R-CNN [51] (37.1% AP).
- We also achieve state-of-the-art performance on weakly supervised semantic correspondence, and are the first to propose a multi-object benchmark for this task.

Task network. Our task network contains an instance segmentation backbone with a multiple instance learning head. The module is supervised by bounding boxes which contain rich object information. Through multiple instance

learning, coarse object masks naturally emerge as network attention, and is taken by the teacher as initial predictions.

Teacher model. The teacher is defined by a Gibbs energy which comprises a unary potential, a pairwise potential and a cross-image potential. The unary potential takes the initial output from the student whereas the pairwise and cross-image potentials model the pairwise pixel relationships both within and across bounding boxes. Minimizing the teacher energy promotes contrast-sensitive smoothness while establishing dense correspondence across the objects. This allows one to consider cross-image self-supervision where correspondence provides positive and negative pairs for dense contrastive learning. We show that this in turn can improve the quality of instance segmentation.

Our promising results indicate the possibility to completely remove mask labels in future instance segmentation problems. We also envisage the wide benefit of DISCOBOX to many downstream applications, particularly 3D tasks.

2. Related Work

2.1. Object recognition and localization

Object detection. Object detection has been an active research area with rich literature. Training on large amounts of bounding box annotations with convolutional neural networks (CNNs) has become a standard paradigm [52]. Initial CNN based detectors tend to share a multi-stage design [52, 53] where the first stage gives redundant object proposals, followed by refinement by CNNs in the second stage. A recent trend of design aims to reduce the complexity by having one-stage architectures [54–57], and therefore achieves good trade-off between efficiency and performance. Our weakly supervised design allows DISCOBOX to be conveniently trained like any object detection algorithm on the increasingly large datasets [58–60], but output additional predictions beyond just bounding boxes.

Instance segmentation. Instance segmentation aims to produce more precise localization over detection by predicting the object segmentation masks. Bharath et al. [2] are the first to introduce an R-CNN-based framework with a precision-recall benchmark. Similar to R-CNN [52], their object proposal and mask generation [61] is not end-to-end learnable. Recent methods including Mask R-CNN [62-64] have largely followed this "detection-flavored" design and benchmarking, but introduce end-to-end learnable object proposal and mask prediction. Inspired by the one-stage detection, a number of one-stage instance segmentation methods have also been proposed [50,51,65-68]. These methods all require mask annotations during training, whereas DIS-COBOX only needs box labels. DISCOBOX is also agnostic to the choice of frameworks. In this work, we showcase DISCOBOX on both YOLACT++ [50] and SOLOv2 [68] by taking them as the base architectures for our method.

2.2. Weakly supervised segmentation

Weakly supervised semantic segmentation. A number of methods have been proposed to learn semantic segmentation with image-level class labels [69–73], points [74,75], scribbles [76–79] and bounding boxes [46,80–83]. Among them, box-supervised semantic segmentation is probably most related, and recent methods such as Box2Seg [83] have achieved impressive performance on Pascal VOC [58]. These methods often use MCG [61] and GrabCut [84] to obtain segmentation pseudo-labels for supervising subsequent tasks. However, they focus on semantic segmentation which does not distinguish different object instances.

Weakly supervised instance segmentation. Here, the term "weakly supervised" can either refer to the relaxed supervision on bounding box location [85, 86], or the absence of mask annotations [46-49]. The former can be viewed as an extension of weakly supervised object detection [87], whereas our work falls into the second category. Among the latter methods, Hsu et al. [47] leverages the fact that bounding boxes tightly enclose the objects, and proposes multiple instance learning framework based on this tightness prior. A pairwise loss is also imposed to maintain object integrity. However, their pairwise consistency is defined on all neighboring pixel pairs without distinguishing the pairwise pixel contrast. Arun et al. [48] proposes an annotation consistency framework which can handle weakly supervised instance segmentation with both image-level and bounding box labels. On COCO, the gap to supervised methods has remained large until recently BoxInst [49] reduced this gap significantly. DISCOBOX outperforms these methods while additionally targeting semantic correspondence.

2.3. Finding correspondence

Local features. Using local features to match the keypoints across different views has been widely used in 3D vision problems such as structure from motion. Over the past decade, the methods have evolved from hand-crafted ones [9–11] to learning-based ones using decision tree and deep neural networks [12–17] with extremely abundant literature. These methods primarily focus on multi-view association for the same object instance or scene, which differs from our task despite the underlying strong connection.

Semantic correspondence. Semantic correspondences has been a challenging problem. The problem probably dates back to SIFTFlow [18] which uses hand-crafted features to establish the correspondence. More recent methods have universally resorted to deep networks as powerful feature extractors [21–23]. The challenge of this task is further aggravated by the costly nature of correspondence annotation. Existing datasets [43, 88] are relatively smalling size, and only provide sparse correspondence ground truths since manually annotating dense ones is prohibitive. In light of this challenge, weakly supervised semantic correspondence correspondence correspondence correspondence is prohibitive.

respondence are proposed to learn correspondence without correspondence ground truths [25–30]. In addition, existing benchmarks and methods have predominantly focused on "object-centric" scenarios where each image is occupied by a major object. In this work, we further add challenge to the task by considering a more generalized multi-object scenario with object localization in the loop.

3. Method

We define the following notations for the variables in our problem, and use them throughout the rest of the paper. We denote the input image as I. Given any instance segmentation backbone, we assume that a set of box region proposals $R = \{r^n | n = 1, ..., N\}$ are generated. Each box proposal corresponds to an RoI feature map f^n of size $C \times H \times W$. Additionally, instance segmentation produces a set of object masks $M = \{m^n | n = 1, ..., N\}$, where each m^n is an $H \times W$ probability map associated with r^n . Fig. 2 illustrates an overview of the proposed framework.

3.1. Task network

DISCOBOX is agnostic to the task networks. We therefore base its design on YOLACT++ [50] and SOLOv2 [68], two recent one-stage instance segmentation frameworks.

YOLACT++. The architecture comprises the following components: 1) **Prediction head.** The framework adopts anchor-based one-stage detection where the prediction head outputs a set of box proposals containing the predicted coordinates and class probabilities. 2) **Mask head.** YOLACT++ proposes a PrototypeNet module to generate *D* latent segment proposals at image level, and use the prediction head to also predict a mask coefficient (a *D*-dim vector) for every box proposal. The mask activation of each proposal is therefore a result of the weighted combination of the segment proposals and the mask coefficient. 3) **Backbone.** Feature Pyramid Network (FPN) [89] is adopted as the backbone, where the pyramid features are scaled up by fusing the skip connected backbone features with higher resolutions.

SOLOv2. We also consider an alternative design based on SOLOv2 [68]. SOLOv2 is a state-of-the-art one-stage framework which directly predicts instance masks in a boxfree, grouping-free and fully convolutional manner. This is done by decoupling the object mask generation into a mask kernel prediction and mask feature learning, together with a parallelizable matrix non-maximum suppression algorithm. SOLOv2 also adopts FPN as the backbone, where the mask kernel is predicted at each pyramid level and a unified mask feature is obtained at 1/4 scale. A slight difference with our YOLACT++ based framework is that no box proposals are directly predicted. We therefore take the tightly enclosed box from each mask as our box proposal for cropping f^n .

Our framework is based on the original designs and implementations of YOLACT++ and SOLOv2, while Fig. 2



Figure 2. Overview of DISCOBOX. We design a self-ensembling framework where a structured teacher generates refined instance segmentation mask and establishes dense correspondence between intra-class box proposals to guide the task network. Best viewd in color.

gives an abstracted illustration. We kindly ask the readers to refer to [50, 68] for more details. We follow the same classification- and box-related training losses which we will jointly term \mathcal{L}_{det} in our paper. This involves $\mathcal{L}_{cls} + \mathcal{L}_{box}$ in YOLACT++, and \mathcal{L}_{cate} in SOLOv2. Since mask annotations are not available, we replace \mathcal{L}_{mask} with the following multiple instance learning (MIL) loss.

Multiple instance learning (MIL). MIL allows one to weakly supervise a task with inexact labels. We follow [47] which proposes an MIL framework leveraging the bounding box tightness prior. Given a box which tightly encloses an object, every row and column contains at least one foreground pixel and can be treated as positive bags. Negative bags can be similarly constructed if the rows and columns have zero overlap with the ground truth box. Denote b_i the set of mask probabilities of the pixel instances belonging to bag i of r^n , the MIL loss for YOLACT++ is defined as:

$$\mathcal{L}_{mil} = -\sum_{i} y_i \log(\max \boldsymbol{b}_i) + (1 - y_i) \log(1 - \max \boldsymbol{b}_i)$$

where $y_i = 1$ if bag *i* is positive, and $y_i = 0$ otherwise. For SOLOv2, \mathcal{L}_{mil} is similarly defined with Dice Loss [68].

3.2. Structured teacher.

The produced segmentation from MIL is still coarse in general. Our main idea is to consider self-ensembling [90] which imposes self-consistency between perturbed models as a self-supervision to improve the representation. Self-ensembling has been a key contributor to the recent success of semi-supervised learning [91, 92]. But unlike these methods which often use augmentation and random dropout to create a noisy student, our problem allows us to form a powerful perturbed teacher through modeling the structured relationships. Promoting contrast sensitive smoothness has

been an important structured inductive bias in segmentation [93]. Instead of achieving this in one shot through post processing [47], our key motivation is to guide the representation with structured inductive bias in a more gradual manner with a mean-field perturbed teacher.

We define a random field $X = \{X^n | n = 1, ..., N\}$ over a graph G = (V, E) where $x^n \in \{0, 1\}^{H \times W}$ is a labeling of X^n in the box proposal r^n . Each node $v_i, i \in r^n$ from box n is sparsely connected with its 8 immediate neighboring nodes $\{v_j | j \in N_p(i)\}$, and densely connected with all the nodes $\{v_k | k \in r^s, s \in N_c(n)\}$ from another intra-class box s. We then define the following Gibbs energy:

$$E(\boldsymbol{x}^n, \boldsymbol{T}_{ns}) = \tau_u(\boldsymbol{x}^n) + \tau_p(\boldsymbol{x}^n) + \sum_{s \in \boldsymbol{N}_c(n)} \tau_c(\boldsymbol{x}^n, \boldsymbol{T}_{ns})$$

where $\tau_u(\boldsymbol{x}^n) = \sum_i \psi(x_i^n)$ are the unary potentials taking the initial output \boldsymbol{m}^n from the instance segmentation head. $\tau_p(\boldsymbol{x}^n)$ are the pairwise potentials defined as:

$$\tau_p(\boldsymbol{x}^n) = \sum_{i \in \boldsymbol{r}^n, j \in \boldsymbol{N}_p(i)} w_1 \exp\Big(-\frac{|\boldsymbol{I}_i^n - \boldsymbol{I}_j^n|^2}{2\zeta^2}\Big) [x_i^n \neq x_j^n]$$

where I_i^n and I_j^n are the RGB colors of pixel *i* and *j* in box *n*, and $[x_i^n \neq x_j^n]$ is a label compatibility function given by the Potts model. Finally, $\tau_c(\boldsymbol{x}^n, \boldsymbol{x}^s, \boldsymbol{T}_{ns})$ are the crossimage potentials which simultaneously models the dense correspondence \boldsymbol{T}_{ns} and the cross-image pairwise labeling relationship. The term is defined with the following energy:

$$\tau_c(\boldsymbol{x}^n, \boldsymbol{T}_{ns}) = \\ - w_2 \sum_{i \in \boldsymbol{r}^n, k \in \boldsymbol{r}^s} \boldsymbol{T}_{ns}(i, k) (C_u(i, k) + C_g(i, k)) [x_i^n = x_k^s]$$

where T_{ns} is a soft assignment matrix of size $HW \times HW$ between box proposals r^n and r^s . In addition, $C_u(i,k)$ is a cost volume matrix which models the appearance similarity:

$$C_u(i,k) = \frac{\boldsymbol{f}_i^n \cdot \boldsymbol{f}_k^s}{|\boldsymbol{f}_i^n| \, |\boldsymbol{f}_k^s|}$$

where f_i^n and f_k^s represent the RoI features of pixel *i* in r^n and pixel *k* in r^s . And $C_g(i, k)$ is further defined as a pairwise smoothness regularization term aiming to impose geometric consistency:

$$C_g(i,k) = \sum_{j \in \mathbf{r}^n, l \in \mathbf{r}^s} \exp\left(-\frac{|\mathsf{off}_{i,k} - \mathsf{off}_{j,l}|^2}{2 * \gamma}\right) \mathbf{T}_{ns}(j,l)$$

where off_{*i*,*k*} represents the relative spatial offset between the pixel *i* in r^n and pixel *k* in r^s . The intuition is to smooth the pairwise offsets to avoid spurious correspondence.

3.3. Inference

We minimize the energy $E(x^n, T_{ns})$ with x^n and T_{ns} alternatively. Although the original $\tau_c(x^n, T_{ns})$ contains a different label compatibility function, its inference with x^n is exactly equivalent to the following energy:

$$\begin{aligned} \tau^*_c(\boldsymbol{x}^n,\boldsymbol{T}_{ns}) &= \\ \sum_{i,k} \boldsymbol{T}_{ns}(i,k) (C_u(i,k) + C_g(i,k)) [x^n_i \neq x^s_k] \end{aligned}$$

Thus, $E(x^n, T_{ns})$ can be minimized via standard meanfield. Kindly refer to Appendix A for more details.

When fixing x, we optimize T by solving an optimal transport problem [23] with the following energy:

$$\min_{\boldsymbol{T}_{ns}} \tau_{c}(\boldsymbol{x}^{n}, \boldsymbol{T}_{ns}) \\ \text{s.t.} \ \boldsymbol{T}_{ns} \boldsymbol{1}_{HW} = \mu_{n}, \ \boldsymbol{T}_{ns}^{\top} \boldsymbol{1}_{HW} = \mu$$

where μ_n , μ_s represent the pixel-wise importance in r^n and r^s , obtained by applying a step function to m^n and m^s .

One may use differentiable Hungarian (denoted \mathcal{H}) such as Sinkhorn's algorithm to solve an optimal transport problem. However, with the pairwise term, directly solving becomes very hard. We therefore approximate with iterated conditional mode where T_{ns} is optimized iteratively:

Initialize:
$$C_u \leftarrow \frac{\boldsymbol{f}^n \cdot \boldsymbol{f}^s}{|\boldsymbol{f}^n| |\boldsymbol{f}^s|}, C^0 \leftarrow C_u(i,k)$$

 $\textbf{Assign:} \quad \textbf{\textit{T}}_{ns}^t \leftarrow \mathcal{H}(C^t)$

 $\begin{array}{ll} \textbf{Update:} \quad C_g^t(i,k) \leftarrow \sum_{j,l} \exp(-\frac{|\mathsf{off}_{i,k} - \mathsf{off}_{j,l}|^2}{2*\gamma}) T_{ns}(j,l)^t \\ C^{t+1}(i,k) = C_u(i,k) + C_g^t(i,k) \end{array}$

An illustration of the above algorithm is shown in Fig. 3. Additional algorithm details can be found in Appendix B.



Figure 3. Illustration of the proposed iterated conditional mode algorithm. Given a pair of RoI features, we use them to initialize the cost volume C_u . Differentiable Hungarian is then iteratively applied to optimize T_{ns} and update C_g .

3.4. Learning

With the x and T inferred from teacher, we can define the following self-ensembling losses. We impose self-consistency between our task network and teacher:

$$\mathcal{L}_{con} = \frac{1}{|\boldsymbol{r}^n|} \sum_{i \in \boldsymbol{r}^n} \left[x_i^n \log(m_i^n) + (1 - x_i^n) \log(1 - m_i^n) \right]$$

For SOLOv2, \mathcal{L}_{con} is similarly defined with Dice Loss [68]. We also use the dense correspondence T_{ns} to obtain positive and negative pairs for dense contrastive learning:

$$\mathcal{L}_{nce} = \frac{1}{|\boldsymbol{r}^n|} \sum_{i \in \boldsymbol{r}^n} \log \frac{\exp(C_u(i, t_i)/\tau)}{\sum_{k \in \boldsymbol{r}^s} \exp(C_u(i, k)/\tau)}$$

where $t_i = \arg \max_k T_{ns}(i, k)$ and τ is the temperature. Our joint training loss therefore can be written as:

$$\mathcal{L} = \mathcal{L}_{det} + \alpha_{mil}\mathcal{L}_{mil} + \alpha_{con}\mathcal{L}_{con} + \alpha_{nce}\mathcal{L}_{nce}$$

3.5. Exponential moving averaged teacher

To strengthen the teacher with model-level augmentation, improved stability, and better consistency between iterations, we follow [90, 94] to obtain a mean teacher with exponential moving average (EMA). This is done by maintaining another network sharing the same architecture and update the parameters as follows:

$$\theta_t \leftarrow m\theta_t + (1-m)\theta_s$$

where θ_t , θ_s are the parameters of teacher network and task network. *m* is the momentum and is set to 0.999 following [94]. We do not train the teacher and only update it.

3.6. Object retrieval with memory bank

To conveniently obtain the object pairs for semantic correspondence, we construct a first-in-first-out (FIFO) queue for each category, where we push the RoI features f and masks m from each batch. This allows us to reuse the RoI features and masks and construct object pairs without much extra computation. During training, the model will retrieve similar intra-class objects from the object bank. After all losses are computed, we push all the objects except the object with area $< 32 \times 32$ into the object banks. Inter-class objects are stored in different queues. Only intra-class objects share the same object bank. In practice, we set the size of object bank of one category as 100.

4. Experiments

We conduct experiments on 4 datasets: PASCAL VOC 2012 (VOC12) [58], COCO [59], PF-PASCAL [88], PAS-CAL 3D+ [4]. We test instance segmentation on VOC12 and COCO, and semantic correspondence on the other two.

4.1. Datasets and metrics

COCO. COCO contains 80 semantic categories. We follow the standard partition which includes train2017 (118K images) and val2017 (5K images) for training and validation. We also report our results on the test-dev split. During training, we only use the box annotations.

VOC12. VOC12 consists of 20 categories with a training set of around 10,500 images and a validation set of around 5,000 images. Around 1,500 images of the validation set contain the instance segmentation annotations.

PF-PASCAL. The PF-PASCAL dataset contains a selected subset of object-centric images from PASCAL VOC. It contains around 1,300 image pairs with 700 pairs for the training set and 300 pairs for the validation set, and 300 image pairs for the test sets respectively. There is only one conspicuous object in the middle of the image. Each image pair contain two intra-class objects.

PASCAL 3D+. PASCAL 3D+ contains the annotations of object poses, landmarks and 3D CAD models in addition to bounding boxes, and consists of 12 rigid categories where each has 3,000 object instances on average. We evaluate multi-object correspondence on PASCAL 3D+ dataset. The bounding boxes, landmarks, and other 3D information make it an ideal dataset to evaluate multi-object semantic correspondence. We set the benchmark on the 12 rigid categories of PASCAL 3D+ and follow the VOC12 train/val split, with images only containing the 8 non-rigid classes removed from the validation set. For training, we still keep the full VOC12 training set and annotations (20 classes).

As PASCAL 3D+ does not provide image pairs, we need to generate image pairs and keypoint pairs on PASCAL for the correspondence evaluation. We enumerate all pairwise combinations of two images on the PASCAL 3D+ validation set. For any pairwise images, if both contain at least one intra-class object in common, we mark them as matched and keep this pair for evaluation. The second step is to generate the sparse correspondence ground truths on top of the matched image pairs using the provided keypoints. For any pairwise images, we find all combinations of intra-class object pairs and use the keypoint pairs between these object pairs as the correspondence ground-truth. Due to occlusion, some keypoints may be missing and are ignored in the evaluation. We also ignore any pairwise objects with the 3D orientation difference greater than 60 degrees, since this often results in very few valid keypoint pairs.

Multi-object correspondence metric. Similar to object detection, we introduce a precision-recall based metric with average precision (AP). We assume that there is a confidence associated with each predicted correspondence, and we define it as the multiplication of the pairwise box confidence in this work. This allows us to compute precision and recall by defining true positive (TP), false positive (FP) and false negative (FN). Since PASCAL 3D+ only provides sparse correspondence ground truths, the challenge here is to correctly ignore some of the correspondence prediction-s that are far away from any ground truth but are correct. To this end, we follow a keypoint transfer setting where we always define a *source* side s and a *target* side t for any pairwise objects. Given a ground truth (g_j^s, g_j^t) , a predicted correspondence (p_i^s, p_i^t) and a distance threshold α :

$$\begin{split} \mathrm{TP}_{i} &= \frac{\sum_{j} \mathbbm{1}[|\boldsymbol{p}_{i}^{s} - \boldsymbol{g}_{j}^{s}| \leq \alpha] \times \mathbbm{1}[|\boldsymbol{p}_{i}^{t} - \boldsymbol{g}_{j}^{t}| \leq \alpha]}{\sum_{j} \mathbbm{1}[|\boldsymbol{p}_{i}^{s} - \boldsymbol{g}_{j}^{s}| \leq \alpha] + \mathbbm{1}[\sum_{j} \mathbbm{1}[|\boldsymbol{p}_{i}^{s} - \boldsymbol{g}_{j}^{s}| \leq \alpha] = 0]} \\ \mathrm{FP}_{i} &= \frac{\sum_{j} \mathbbm{1}[|\boldsymbol{p}_{i}^{s} - \boldsymbol{g}_{j}^{s}| \leq \alpha] \times \mathbbm{1}[|\boldsymbol{p}_{i}^{t} - \boldsymbol{g}_{j}^{t}| > \alpha]}{\sum_{j} \mathbbm{1}[|\boldsymbol{p}_{i}^{s} - \boldsymbol{g}_{j}^{s}| \leq \alpha] + \mathbbm{1}[\sum_{j} \mathbbm{1}[|\boldsymbol{p}_{i}^{s} - \boldsymbol{g}_{j}^{s}| \leq \alpha] = 0]} \\ \mathrm{FN}_{i} &= \begin{cases} 1 & \text{if } \sum_{j} \mathbbm{1}[|\boldsymbol{p}_{i}^{s} - \boldsymbol{g}_{j}^{s}| \leq \alpha] = 0, \\ 0 & \text{otherwise} \end{cases} \end{split}$$

We term the average precision as AP@ α where α is a threshold relative to the box diagonal. We then define the final AP as: mean(AP@{0.75\%, 1\%, 1.5\%, 2\%, 3\%}).

4.2. Implementation details

Training. We use stochastic gradient descent (SGD) for network optimization. For loss weights, we set α_{mil} , α_{con} , α_{nce} as 10, 2, 0.1 on YOLACT++ and set α_{mil} , α_{con} , α_{nce} as 1, 1, 0.1 on SOLOv2. Kindly refer to Appendix C for additional implementation details.

4.3. Weakly supervised instance segmentation

Main results. We evaluate instance segmentation on COCO and VOC12, with the main results reported in Tab. 1 and 2, respectively. DISCOBOX outperforms BBTP [47] by 10.3% mAP on the COCO validation 2017 split with a smaller backbone (ResNet-50). DISCOBOX also outperforms BoxInst [49] which is the current state-of-the-art box-supervised method on both COCO and VOC12. Notably, BoxInst/ResNet-101-DCN also adopts BiFPN [95], an improved variant of FPN [89]. Fig. 4 and Appendix D additionally visualize the instance segmentation results.

Method	Backbone	Туре	Arch	Sup	FPS	AP	AP ₅₀	AP ₇₅	AP_S	AP_M	AP_L	Rat
MNC [62]	ResNet-101	Two-stage	MNC	Mask	<2.8	24.6	44.3	24.8	4.7	25.9	43.6	-
FCIS [63]	ResNet-101	Two-stage	FCIS	Mask	6.6	29.2	49.5	-	7.1	31.3	50.0	-
Mask R-CNN [64]	ResNet-101	Two-stage	Mask R-CNN	Mask	5	35.7	58.0	37.8	15.5	38.1	52.4	-
Mask R-CNN [64]	ResNeXt-101	Two-stage	Mask R-CNN	Mask	<5	37.1	60.0	39.4	16.9	38.9	53.5	-
PolarMask [51]	ResNet-101	One-stage	PolarMask	Mask	12.3	32.1	53.7	33.1	14.7	33.8	45.3	-
YOLACT-550++ [50]	ResNet-50-DCN	One-stage	YOLACT++	Mask	33.5	34.1	53.3	36.2	11.7	36.1	53.6	-
YOLACT-550++ [50]	ResNet-101-DCN	One-stage	YOLACT++	Mask	27.3	34.6	53.8	36.9	11.9	36.8	55.1	-
CondInst [67]	ResNet-101	One-stage	CondInst	Mask	10.3	39.1	60.9	42.0	21.5	41.7	50.9	-
SOLOv2 [68]	ResNet-101-DCN	One-stage	SOLOv2	Mask	10.3	41.7	63.2	45.1	18.0	45.0	61.6	-
SOLOv2 [68]	ResNeXt-101-DCN	One-stage	SOLOv2	Mask	7.4	42.4	64.0	45.8	19.1	46.3	63.4	-
BBTP† [47]	ResNet-101	Two-stgae	Mask R-CNN	Box	<5	21.1	45.5	17.2	11.2	22.0	29.8	-
DISCOBOX (Ours) [†]	ResNet-50	One-stage	SOLOv2	Box	18.5	31.4	52.6	32.2	11.5	33.8	50.1	-
BoxInst [49]	ResNet-101	One-stage	CondInst	Box	-	33.2	56.5	33.6	16.2	35.3	45.1	84.9
BoxInst [49]	ResNet-101-DCN	One-stage	CondInst	Box	-	35.0	59.3	35.6	17.1	37.2	48.9	-
DISCOBOX (Ours)	ResNet-50-DCN	One-stage	YOLACT++	Box	34.5	27.3	49.2	26.9	9.9	28.1	43.2	80.1
DISCOBOX (Ours)	ResNet-50	One-stage	SOLOv2	Box	18.5	32.0	53.6	32.6	11.7	33.7	48.4	82.5
DISCOBOX (Ours)	ResNet-101-DCN	One-stage	SOLOv2	Box	10.3	35.8	59.8	36.4	16.9	38.7	52.1	85.9
DISCOBOX (Ours)	ResNeXt-101-DCN	One-stage	SOLOv2	Box	7.4	37.9	61.4	40.0	18.0	41.1	53.9	89.4

Table 1. Main results on COCO. Results with "†" are on the val2017 split, and test-dev otherwise. DISCOBOX with SOLOv2/ResNet-50 outperforms BBTP [47] by 10.3%. Our best model achieves 37.9% AP on test-dev, surpassing supervised methods such as Mask R-CNN in absolute performance. The performance ratio between some box-supervised methods and their supervised counterparts are also reported.

Method	Backbone	Arch	AP ₂₅	AP ₅₀	AP ₇₀	AP ₇₅
SDI [46]	VGG-16	DeepLabv2	-	44.8	-	16.3
BBTP [47]	ResNet-101	Mask R-CNN	75.0	58.9	30.4	21.6
Arun et al. [48]	ResNet-101	Mask R-CNN	73.1	57.7	33.5	31.2
BoxInst [49]	ResNet-50	CondInst	-	59.1	-	34.2
BoxInst [49]	ResNet-101	CondInst	-	61.4	-	37.0
DISCOBOX	ResNet-50-DCN	YOLACT++	75.2	63.6	41.6	34.1
DISCOBOX	ResNet-50	SOLOv2	71.4	59.8	41.7	35.5
DISCOBOX	ResNet-101	SOLOv2	72.8	62.2	45.5	37.5

Table 2. Main results on the VOC12 validation set. DISCOBOX outperforms all previous methods with state-of-the-art results.

\mathcal{L}_{mil}	\mathcal{L}_{con}	\mathcal{L}_{nce}	EMA	Str	AP ₅₀ ^y	AP ₇₀ ^y	AP ₇₅	AP ^s ₅₀	AP ^s ₇₀	AP ^s ₇₅
\checkmark			\checkmark	\checkmark	43.3	18.3	17.0	42.1	18.0	17.3
\checkmark	\checkmark		\checkmark	\checkmark	62.0	40.1	33.5	58.1	40.9	34.9
\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	63.6	41.6	34.1	59.8	41.7	35.5
\checkmark	~	\checkmark	\checkmark		-	-	-	48.6	20.7	17.6
\checkmark	\checkmark	\checkmark		\checkmark	-	-	-	58.3	40.5	34.8

Table 3. Ablation study on VOC12, where 'y' denotes the results obtained by YOLOCT++/ResNet-50-DCN and 's' denotes the results obtained by SOLOv2/ResNet-50.

Analysis. We perform ablation study on VOC12 with $\mathcal{L}_{mil}, \mathcal{L}_{con}, \mathcal{L}_{nce}$, EMA, and structured teacher. The results in Tab. 3 show consistent improvements from \mathcal{L}_{con} and \mathcal{L}_{nce} , as well as the importance of structured teacher. We also conduct sensitivity analysis with the loss weights on both instance segmentation (VOC12)¹ and semantic correspondence (PASCAL 3D+, see Sec. 4.4). The results in Fig. 6 show that DISCOBOX is not sensitive to weight changes.

4.4. Weakly supervised semantic correspondence

PF-PASCAL (Object-Centric). We first evaluate DIS-COBOX on PF-PASCAL [88] using YOLACT++/ResNet-

Signal	Method	Backbone	PCK@0.05	PCK@0.1	PCK@0.15
none	PF _{HOG} [96]	VGG16	31.4	62.5	79.5
image	WeakAlign [25]	R101	49.0	74.8	84.0
image	RTNs [97]	R101	55.2	75.9	85.2
image	NC-Net [98]	R101	54.3	78.9	86.0
image	DCC-Net [99]	R101	55.6	82.3	90.5
image	DHPF [30]	R101	56.1	82.1	91.1
mask	SF-Net [28]	R101	53.6	81.9	90.6
box	DISCOBOX	R50-DCN	59.3	87.2	95.3

Table 4. Results on PF-PASCAL. DISCOBOX outperforms previous state-of-the-art methods on weakly supervised semantic correspondence without bells-and-whistles.

50-DCN, with the main results presented in Tab. 4. We do not directly train the DISCOBOX model on PF-PASCAL. Instead, we train it on the VOC12 training set, excluding those images that are present in the PF-PASCAL validation set. It is worth noting that many existing semantic correspondence methods can not be similarly trained on VOC12 without major changes, even though some of them do consider certain level of localization information such as attention. During inference, we use instance segmentation to obtain object masks, and use the structured teacher to produce dense pixel-wise correspondence by taking the masks as input. Our approach outperforms many competitive previous approaches on weakly supervised semantic correspondence with considerable margins. Such improvement can be attributed to three main factors: 1) The improved design of the structured teacher that renders good correspondence quality at the object level. 2) The box-supervised learning framework that makes it possible to scale up the training using more data and obtain improved correspondence representation. 3) The high quality object localization information as a result of the coupled learning framework that helps to guide the semantic correspondence learning.

¹Here, the AP follows COCO evaluation: mean(AP@{50,55,...,95}).



Figure 4. Visualization of instance segmentation on COCO (YOLACT++/ResNet-50-DCN).



Figure 5. Visualization of multi-object semantic correspondence on PASCAL 3D+ (YOLACT++/ResNet-50-DCN).



Figure 6. Sensitivity analysis with the loss weights on instance segmentation (VOC12) and semantic correspondence (PASCAL 3D+). Analysis conducted on YOLACT++/ResNet-50-DCN.

PASCAL 3D+ (Multi-Object). Finally, we benchmark DISCOBOX and several baselines on PASCAL 3D+. Tab. 5 lists the main results and Fig. 5 visualizes some predicted correspondence. The comparing methods in Tab. 5 are defined as follows: **Identity:** We align each pair of images only considering the positions of pixels. **SCOT:** A modified version of [23] by removing beam search and keeping their matching module on our RoI features. **DISCOBOX-:** Our model trained on VOC12 without dense NCE loss, but using teacher during inference for correspondence. **DISCOBOX:** Our full approach. We use YOLACT++/ResNet-50-DCN for all methods. Our method does not include beam search with the validation data and label [23], and is therefore purely box-supervised. The results show the effectiveness of our proposed teacher and dense contrastive learning.

Methods	AP	AP _{0.75}	AP_1	AP _{1.5}	AP_2	AP_3
Identity	26.6	10.5	16.3	26.2	34.2	46.0
SCOT [23]	29.3	13.2	19.8	29.8	36.0	47.3
DISCOBOX-	30.9	15.6	21.3	30.6	38.0	48.9
DISCOBOX	31.7	15.8	21.4	31.8	39.5	50.3

Table 5. Results of multi-object correspondence on PASCAL 3D+, where AP_{0.75}, AP₁, AP_{1.5}, AP₂, and AP₃ represent AP@ α with thresholds $\alpha \in \{0.75\%, 1\%, 1.5\%, 2\%, 3\%\}$ relative to the box diagonal. AP is defined as mean(AP_{0.75}, AP₁, AP_{1.5}, AP₂, AP₃).

5. Conclusions

We presented DISCOBOX, a novel framework able to jointly learn instance segmentation and semantic correspondence from box supervision. Our proposed self-ensembling framework with a structured teacher has led to significant improvement with state of the art performance in both tasks. We also proposed a novel benchmark for multi-object semantic correspondence together with a principled evaluation metric. With the ability to jointly produce high quality instance segmentation and semantic correspondence from box supervision, we envision that DISCOBOX can scale up and benefit many downstream 2D and 3D vision tasks.

Acknowledgement: We would like to sincerely thank Xinlong Wang, Zhi Tian, Shuaiyi Huang, Yashar Asgarieh, Jose M. Alvarez, De-An Huang, and other NVIDIA colleagues for the discussion and constructive suggestions.

References

- [1] Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR*, 2001. 1
- [2] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Simultaneous detection and segmentation. In *ECCV*, 2014. 1, 2
- [3] Mohsen Hejrati and Deva Ramanan. Analyzing 3d objects in cluttered images. In *NIPS*, 2012. 1
- [4] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond PASCAL: A benchmark for 3d object detection in the wild. In WACV, 2014. 1, 6
- [5] Yu Xiang, Wonhui Kim, Wei Chen, Jingwei Ji, Christopher Choy, Hao Su, Roozbeh Mottaghi, Leonidas Guibas, and Silvio Savarese. ObjectNet3D: A large scale database for 3d object recognition. In ECCV, 2016. 1
- [6] Abhijit Kundu, Yin Li, and James M Rehg. 3D-RCNN: Instance-level 3d object reconstruction via render-andcompare. In CVPR, 2018. 1
- [7] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh R-CNN. In *ICCV*, 2019.
- [8] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2D & 3D face alignment problem? (and a dataset of 230,000 3D facial landmarks). In *ICCV*, 2017. 1
- [9] David G Lowe. Distinctive image features from scaleinvariant keypoints. *IJCV*, 2004. 1, 3
- [10] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded up robust features. In ECCV, 2006. 1, 3
- [11] Engin Tola, Vincent Lepetit, and Pascal Fua. DAISY: An efficient dense descriptor applied to wide-baseline stereo. *IEEE Trans. PAMI*, 2009. 1, 3
- [12] Edward Rosten and Tom Drummond. Machine learning for high-speed corner detection. In ECCV, 2006. 1, 3
- [13] Yannick Verdie, Kwang Yi, Pascal Fua, and Vincent Lepetit. TILDE: A temporally invariant learned detector. In *CVPR*, 2015. 1, 3
- [14] Xufeng Han, Thomas Leung, Yangqing Jia, Rahul Sukthankar, and Alexander C Berg. MatchNet: Unifying feature and metric learning for patch-based matching. In *CVPR*, 2015. 1, 3
- [15] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. LIFT: Learned invariant feature transform. In *ECCV*, 2016. 1, 3
- [16] Yuki Ono, Eduard Trulls, Pascal Fua, and Kwang Moo Yi. LF-Net: Learning local features from images. 2018. 1, 3
- [17] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In CVPR, 2020. 1, 3
- [18] Ce Liu, Jenny Yuen, and Antonio Torralba. SIFT Flow: Dense correspondence across scenes and its applications. *IEEE Trans. PAMI*, 2010. 1, 3

- [19] Tinghui Zhou, Yong Jae Lee, Stella X Yu, and Alyosha A Efros. FlowWeb: Joint image set alignment by weaving consistent, pixel-wise correspondences. In CVPR, 2015. 1
- [20] Christopher B Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. In *NIPS*, 2016. 1
- [21] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *CVPR*, 2017. 1, 3
- [22] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Hyperpixel flow: Semantic correspondence with multilayer neural features. In *ICCV*, 2019. 1, 3
- [23] Yanbin Liu, Linchao Zhu, Makoto Yamada, and Yi Yang. Semantic correspondence as an optimal transport problem. In *CVPR*, 2020. 1, 3, 5, 8
- [24] Nilesh Kulkarni, Abhinav Gupta, and Shubham Tulsiani. Canonical surface mapping via geometric cycle consistency. In *ICCV*, 2019. 1
- [25] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Endto-end weakly-supervised semantic alignment. In CVPR, 2018. 1, 2, 3, 7
- [26] Yun-Chun Chen, Yen-Yu Lin, Ming-Hsuan Yang, and Jia-Bin Huang. Show, match and segment: Joint learning of semantic matching and object co-segmentation. *IEEE Trans. PAMI*, 2019. 1, 2, 3
- [27] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. SCOPS: Self-supervised co-part segmentation. In *CVPR*, 2019. 1, 2, 3
- [28] Junghyup Lee, Dohyung Kim, Jean Ponce, and Bumsub Ham. SFNet: Learning object-aware semantic correspondence. In *CVPR*, 2019. 1, 2, 3, 7
- [29] Sangryul Jeon, Dongbo Min, Seungryong Kim, Jihwan Choe, and Kwanghoon Sohn. Guided semantic flow. In ECCV, 2020. 1, 2, 3
- [30] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Learning to compose hypercolumns for visual correspondence. In ECCV, 2020. 1, 2, 3, 7
- [31] Shubham Tulsiani and Jitendra Malik. Viewpoints and keypoints. In CVPR, 2015. 1
- [32] Alvaro Collet, Manuel Martinez, and Siddhartha S Srinivasa. The MOPED framework: Object recognition and pose estimation for manipulation. *The international journal of robotics research*, 2011. 1
- [33] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-DoF object pose from semantic keypoints. In *ICRA*, 2017. 1
- [34] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *CoRL*, 2018. 1
- [35] Noah Snavely, Steven M Seitz, and Richard Szeliski. Photo Tourism: Exploring photo collections in 3D. In SIG-GRAPH. 2006. 1

- [36] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In ECCV, 2018. 1
- [37] David Novotny, Nikhila Ravi, Benjamin Graham, Natalia Neverova, and Andrea Vedaldi. C3DPO: Canonical 3d pose networks for non-rigid structure from motion. In *ICCV*, 2019. 1
- [38] Xueting Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Selfsupervised single-view 3D reconstruction via semantic consistency. 2020. 1
- [39] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 1
- [40] Tatsunori Taniai, Sudipta N Sinha, and Yoichi Sato. Joint recovery of dense correspondence and cosegmentation in two images. In CVPR, 2016. 1
- [41] David Novotny, Diane Larlus, and Andrea Vedaldi. I have seen enough: Transferring parts across categories. In *B-MVC*, 2016. 1
- [42] Bumsub Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow: Semantic correspondences from object proposals. *IEEE Trans. PAMI*, 2017. 1
- [43] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. SPair-71k: A large-scale benchmark for semantic correspondence. arXiv preprint arXiv:1908.10543, 2019. 1, 3
- [44] Xiangxin Zhu and Deva Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *CVPR*, 2012. 1
- [45] Lichao Huang, Yi Yang, Yafeng Deng, and Yinan Yu. DenseBox: Unifying landmark localization with end to end object detection. arXiv preprint arXiv:1509.04874, 2015. 1
- [46] Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *CVPR*, 2017. 2, 3, 7
- [47] 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 *NeurIPS*, 2019. 2, 3, 4, 6, 7
- [48] Aditya Arun, CV Jawahar, and M Pawan Kumar. Weakly supervised instance segmentation by learning annotation consistent instances. In ECCV, 2020. 2, 3, 7
- [49] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. Boxinst: High-performance instance segmentation with box annotations. *CVPR*, 2021. 2, 3, 6, 7
- [50] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT++: Better real-time instance segmentation. *IEEE Trans. PAMI*, 2020. 2, 3, 4, 7, 12
- [51] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polar-Mask: Single shot instance segmentation with polar representation. In *CVPR*, 2020. 2, 7

- [52] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In CVPR, 2014. 2
- [53] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 2
- [54] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 2
- [55] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. SSD: Single shot multibox detector. In *ECCV*, 2016. 2
- [56] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. arXiv preprint arXiv:1904.07850, 2019. 2
- [57] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, 2019. 2
- [58] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The PAS-CAL visual object classes (VOC) challenge. *IJCV*, 2010. 2, 3, 6
- [59] 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 ECCV, 2014. 2, 6
- [60] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *ICCV*, 2019. 2
- [61] Pablo Arbeláez, Jordi Pont-Tuset, Jonathan T Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping. In CVPR, 2014. 2, 3
- [62] Jifeng Dai, Kaiming He, and Jian Sun. Instance-aware semantic segmentation via multi-task network cascades. In *CVPR*, 2016. 2, 7
- [63] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, and Yichen Wei. Fully convolutional instance-aware semantic segmentation. In *CVPR*, 2017. 2, 7
- [64] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 2, 7
- [65] Min Bai and Raquel Urtasun. Deep watershed transform for instance segmentation. In *CVPR*, 2017. 2
- [66] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT: Real-time instance segmentation. In *IC-CV*, 2019. 2
- [67] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. arXiv preprint arXiv:2003.05664, 2020. 2, 7
- [68] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. *NeurIPS*, 2020. 2, 3, 4, 5, 7, 12, 13

- [69] Thibaut Durand, Taylor Mordan, Nicolas Thome, and Matthieu Cord. WILDCAT: Weakly supervised learning of deep convnets for image classification, pointwise localization and segmentation. In CVPR, 2017. 3
- [70] Bin Jin, Maria V Ortiz Segovia, and Sabine Susstrunk. Webly supervised semantic segmentation. In CVPR, 2017. 3
- [71] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In CVPR, 2018. 3
- [72] 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 ECCV, 2018. 3
- [73] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *ECCV*, 2020. 3
- [74] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. Whats the point: Semantic segmentation with point supervision. In ECCV, 2016. 3
- [75] Rui Qian, Yunchao Wei, Honghui Shi, Jiachen Li, Jiaying Liu, and Thomas Huang. Weakly supervised scene parsing with point-based distance metric learning. In *AAAI*, 2019.
 3
- [76] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. ScribbleSup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, 2016. 3
- [77] Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In CVPR, 2017. 3
- [78] Meng Tang, Abdelaziz Djelouah, Federico Perazzi, Yuri Boykov, and Christopher Schroers. Normalized cut loss for weakly-supervised cnn segmentation. In CVPR, 2018. 3
- [79] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised cnn segmentation. In *ECCV*, 2018. 3
- [80] Jifeng Dai, Kaiming He, and Jian Sun. BoxSup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *ICCV*, 2015. 3
- [81] 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 *ICCV*, 2015. 3
- [82] 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 *CVPR*, 2019. 3
- [83] Viveka Kulharia, Siddhartha Chandra, Amit Agrawal, Philip Torr, and Ambrish Tyagi. Box2Seg: Attention weighted loss and discriminative feature learning for weakly supervised segmentation. In ECCV, 2020. 3
- [84] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. "GrabCut": Interactive foreground extraction using iterated graph cuts. *SIGGRAPH*, 2004. 3

- [85] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In CVPR, 2018. 3
- [86] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *CVPR*, 2019. 3
- [87] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In CVPR, 2016. 3
- [88] Bumsub Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow. In CVPR, 2016. 3, 6, 7
- [89] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 3, 6
- [90] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, 2017. 4, 5
- [91] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semisupervised learning with consistency and confidence. In *NeurIPS*, 2020. 4
- [92] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *CVPR*, 2020. 4
- [93] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected CRFs with Gaussian edge potentials. In *NIPS*, 2011. 4
- [94] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In CVPR, 2020. 5
- [95] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficient-Det: Scalable and efficient object detection. In CVPR, 2020. 6
- [96] Kai Han, Rafael S Rezende, Bumsub Ham, Kwan-Yee K Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. SC-Net: Learning semantic correspondence. In *ICCV*, 2017.
 7
- [97] Seungryong Kim, Stephen Lin, Sang Ryul Jeon, Dongbo Min, and Kwanghoon Sohn. Recurrent transformer networks for semantic correspondence. *NeurIPS*, 2018. 7
- [98] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. arXiv preprint arXiv:1810.10510, 2018. 7
- [99] Shuaiyi Huang, Qiuyue Wang, Songyang Zhang, Shipeng Yan, and Xuming He. Dynamic context correspondence network for semantic alignment. In *ICCV*, 2019. 7
- [100] Philip A Knight. The Sinkhorn–Knopp algorithm: Convergence and applications. SIAM Journal on Matrix Analysis and Applications, 2008. 12