

Proposal Learning for Semi-Supervised Object Detection

Peng Tang[†] Chetan Ramaiah[†] Yan Wang[‡] Ran Xu[†] Caiming Xiong[†]

[†]Salesforce Research

[‡]Shanghai Key Laboratory of Multidimensional Information Processing, East China Normal University

{peng.tang, cramaiah, ran.xu, cxiong}@salesforce.com wyanny.9@gmail.com

Abstract

In this paper, we focus on semi-supervised object detection to boost performance of proposal-based object detectors (a.k.a. two-stage object detectors) by training on both labeled and unlabeled data. However, it is non-trivial to train object detectors on unlabeled data due to the unavailability of ground truth labels. To address this problem, we present a proposal learning approach to learn proposal features and predictions from both labeled and unlabeled data. The approach consists of a self-supervised proposal learning module and a consistency-based proposal learning module. In the self-supervised proposal learning module, we present a proposal location loss and a contrastive loss to learn context-aware and noise-robust proposal features respectively. In the consistency-based proposal learning module, we apply consistency losses to both bounding box classification and regression predictions of proposals to learn noise-robust proposal features and predictions. Our approach enjoys the following benefits: 1) encouraging more context information to be delivered in the proposals learning procedure; 2) noisy proposal features and enforcing consistency to allow noise-robust object detection; 3) building a general and high-performance semi-supervised object detection framework, which can be easily adapted to proposal-based object detectors with different backbone architectures. Experiments are conducted on the COCO dataset with all available labeled and unlabeled data. Results demonstrate that our approach consistently improves the performance of fully-supervised baselines. In particular, after combining with data distillation [39], our approach improves AP by about 2.0% and 0.9% on average compared to fully-supervised baselines and data distillation baselines respectively.

1. Introduction

With the giant success of Convolutional Neural Networks (CNNs) [25, 27], great leap forwards have been achieved in object detection [14, 15, 17, 29, 31, 40, 41].

However, training accurate object detectors relies on the availability of large scale labeled datasets [10, 30, 42, 44], which are very expensive and time-consuming to collect. In addition, training object detectors only on the labeled datasets may limit their detection performance. By contrast, considering that acquiring unlabeled data is much easier than collecting labeled data, it is important to explore approaches for the Semi-Supervised Object Detection (SSOD) problem, *i.e.*, training object detectors on both labeled and unlabeled data, to boost performance of current state-of-the-art object detectors.

In this paper, we focus on SSOD for proposal-based object detectors (a.k.a. two-stage object detectors) [14, 15, 41] due to their high performance. Proposal-based object detectors detect objects by 1) first generating region proposals that may contain objects and 2) then generating proposal features and predictions (*i.e.*, bounding box classification and regression predictions). Specially, we aim to improve the second stage by learning proposal features and predictions from both labeled and unlabeled data.

For labeled data, it is straightforward to use ground truth labels to get training supervisions. But for unlabeled data, due to the unavailability of ground truth labels, we cannot learn proposal features and predictions directly. To address this problem, apart from the standard fully-supervised learning for labeled data [41] shown in Fig. 1 (a), we present an approach named proposal learning, which consists of a self-supervised proposal learning module and a consistency-based proposal learning module, to learn proposal features and predictions from both labeled and unlabeled data, see Fig. 1.

Recently, self-supervised learning has shown its efficacy to learn features from unlabeled data by defining some pretext tasks [9, 16, 23, 53, 57]. Our self-supervised proposal learning module uses the same strategy of defining pretext tasks, inspired by the facts that context is important for object detection [2, 8, 20, 34] and object detectors should be noise-robust [32, 52]. More precisely, a proposal location loss and a contrastive loss are presented to learn context-aware and noise-robust proposal features re-

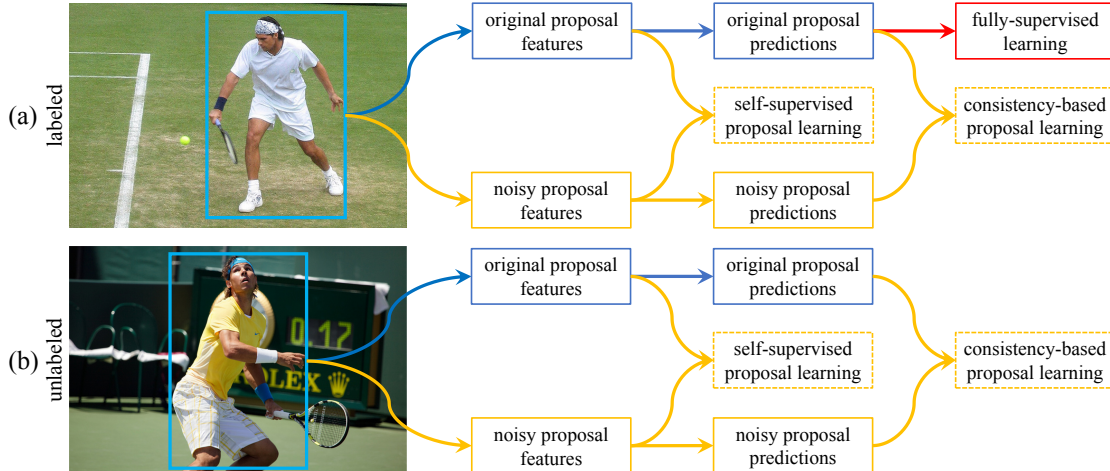


Figure 1. The design of our proposal learning approach. The proposed modules are highlighted in dashed boxes. Given an image, original/noisy proposal features and predictions are generated. (a) The standard fully-supervised learning is chosen for labeled data. (a) and (b) Our proposal learning approach learns proposal features and predictions from both labeled and unlabeled data.

spectively. Specifically, the proposal location loss uses proposal location prediction as a pretext task to supervise training, where a small neural network are attached after proposal features for proposal location prediction. This loss helps learn context-aware proposal features, because proposal location prediction requires proposal features understanding some global image information. At the same time, the contrastive loss learns noise-robust proposal features by a simple instance discrimination task [16, 53, 57], which ensures that noisy proposals features are closer to their original proposal features than to other proposal features. In particular, instead of adding noise to images to compute contrastive loss [16, 57], we add noise to proposal features, which shares convolutional feature computations for the entire image between noisy proposal feature computations and the original proposal feature computations for training efficiency [14].

To further train noise-robust object detectors, our consistency-based proposal learning module uses consistency losses to ensure that predictions from noisy proposal features and their original proposal features are consistent. More precisely, similar to consistency losses for semi-supervised image classification [33, 43, 54], a consistency loss for bounding box classification predictions enforces class predictions from noisy proposal features and their original proposal features to be consistent. In addition, a consistency loss for bounding box regression predictions enforces object location predictions from noisy proposal features and their original proposal features also to be consistent. With these two consistency losses, proposal features and predictions are robust to noise.

We apply our approach to Faster R-CNN [41] with feature pyramid networks [28] and RoIAlign [17], using different CNN backbones, where our proposal learning mod-

ules are applied to both labeled and unlabeled data, as shown in Fig. 1. We conduct elaborate experiments on the challenging COCO dataset [30] with all available labeled and unlabeled data, showing that our approach outperforms fully-supervised baselines consistently. In particular, when combining with data distillation [39], our approach obtains about 2.0% and 0.9% absolute AP improvements on average compared to fully-supervised baselines and data distillation based baselines respectively.

In summary, we list our main contributions as follows.

- We present a proposal learning approach to learn proposal features and predictions from both labeled and unlabeled data. The approach consists of 1) a self-supervised proposal learning module which learns context-aware and noise-robust proposal features by a proposal location loss and a contrastive loss respectively, and 2) a consistency-based proposal learning module which learns noise-robust proposal features and predictions by consistency losses for bounding box classification and regression predictions.
- On the COCO dataset, our approach surpasses various Faster R-CNN based fully-supervised baselines and data distillation [39] by about 2.0% and 0.9% respectively.

2. Related Work

Object detection is one of the most important tasks in computer vision and has received considerable attention in recent years [4, 14, 15, 17, 20, 21, 28, 29, 31, 40, 41] [45, 49, 52, 60]. One popular direction for recent object detection is proposal-based object detectors (a.k.a. two-stage object detectors) [14, 15, 17, 28, 41], which perform object

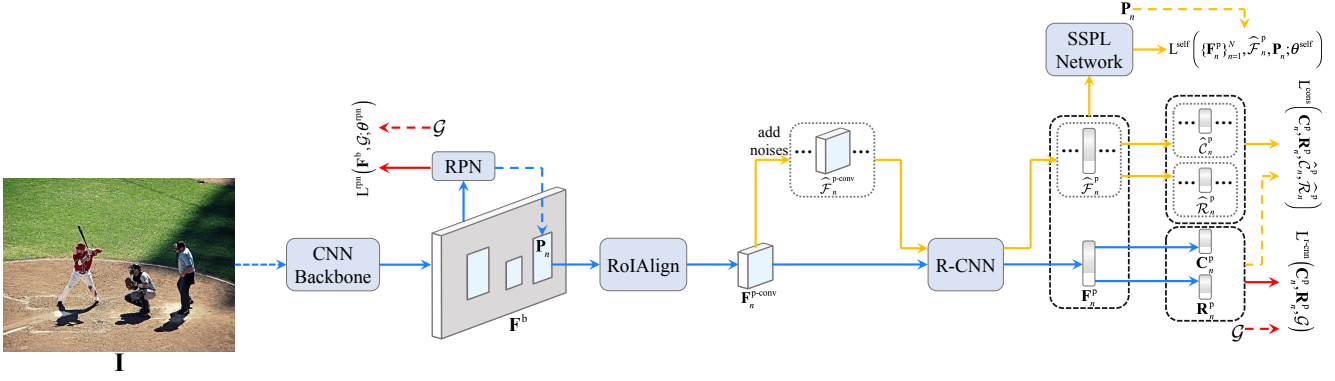


Figure 2. The overall framework of our proposal learning approach. All arrows have forward computations during training, only the solid ones have back-propagation computations, the red ones are only utilized for labeled data, and only the blue ones are utilized during inference. “RPN”: Region Proposal Network; “R-CNN”: Region-based CNN; “SSPL”: Self-Supervised Proposal Learning; “I”: input image; “ F^b ”: image convolutional feature maps; “ θ^{rpn} ”: parameters of the RPN; “ P_n ”: a proposal with its location; “ $F_n^{p\text{-conv}}$ ” and “ $\hat{F}_n^{p\text{-conv}}$ ”: the original and noisy convolutional feature maps of P_n ; “ F_n^p ” and “ \hat{F}_n^p ”: the original and noisy features of P_n ; “ C_n^p , R_n^p ” and “ \hat{C}_n^p , \hat{R}_n^p ”: the original and noisy predictions (bounding box classification and regression predictions) of P_n ; “ θ^{self} ”: parameters of the SSPL network; “ G ”: ground truth labels; “ $L^{rpn}(F^b, G; \theta^{rpn})$ ”: RPN loss; “ $L^{r\text{-cnn}}(C_n^p, R_n^p, G)$ ”: R-CNN loss; “ $L^{\text{self}}(\{F_n^p\}_{n=1}^N, \hat{F}_n^p, P_n; \theta^{\text{self}})$ ”: SSPL loss; “ $L^{\text{cons}}(C_n^p, R_n^p, \hat{C}_n^p, \hat{R}_n^p)$ ”: consistency-based proposal learning loss. See Section 3.2 for more details. Best viewed in color.

detection by first generating region proposals and then generating proposal features and predictions. Very promising results are obtained by these proposal-based approaches. In this work, we are also along the line of proposal-based object detectors. But unlike previous approaches training object detectors only on labeled data, we train object detectors on both labeled data and unlabeled data, and present a proposal learning approach to achieve our goal. In addition, Wang *et al.* [52] also add noise to proposal features to train noise-robust object detectors. They focus on generating hard noisy proposal features and present adversarial networks based approaches, and still train object detectors only on labeled data. Unlike their approach, we add noise to proposal features to learn noise-robust proposal features and predictions from both labeled and unlabeled data by our proposal learning approach.

Self-supervised learning learns features from unlabeled data by some defined pretext tasks [9, 13, 16, 24, 37, 57, 59]. For example, Doersch *et al.* [9] predicts the position of one patch relative to another patch in the same image. Gidaris *et al.* [13] randomly rotate images and predict the rotation of images. Some recent works [16, 53, 57] use an instance discrimination task to match features from noisy images with features from their original images. Please see the recent survey [23] for more self-supervised learning approaches. Our self-supervised proposal learning module applies self-supervised learning approaches to learn proposal features from both labeled and unlabeled data. Inspired by unsupervised feature learning task [9] and the instance discrimination task [16, 53, 57] of image level, we introduce self-supervised learning to SSOD by designing a proposal loca-

tion loss and contrastive loss on object proposals.

Semi-supervised learning trains models on both labeled and unlabeled data. There are multiple strategies for semi-supervised learning, such as self-training [56], co-training [3, 38, 61], label propagation [62], *etc.* Please see [5] for an extensive review. Recently, many works use consistency losses for semi-supervised image classification [1, 26, 33, 43, 47, 54], by enforcing class predictions from noisy inputs and their original inputs to be consistent, where noise is added to input images or intermediate features. Here we add noise to proposal features for efficiency and apply consistency losses to both class and object location predictions of proposals. Zhai *et al.* [58] also suggest to benefit semi-supervised image classification from self-supervised learning. Here we further apply self-supervised learning to SSOD by a self-supervised proposal learning module.

Semi-supervised object detection applies semi-supervised learning to object detection. There are some SSOD works with different settings [7, 11, 19, 35, 46, 50, 51]. For example, Cinbis *et al.* [7] train object detectors on data with either bounding box labels or image-level class labels. Hoffman *et al.* [19] and Tang *et al.* [46] train object detectors on data with bounding box labels for some classes and image-level class labels for other classes. Gao *et al.* [11] train object detectors on data with bounding box labels for some classes and either image-level labels or bounding box labels for other classes. Unlike their settings, in this work we explore the more general semi-supervised setting, *i.e.*, training object detectors on data which either have bounding box labels or are totally unlabeled, similar to the standard semi-supervised learning setting [5]. Jeong *et al.* [22]

also use consistency losses for SSOD by adding noise to images. Unlike their approach, we add noise to proposal features instead of images and present a self-supervised proposal learning module. In addition, all these works mainly conduct experiments on simulated labeled/unlabeled data by splitting a fully annotated dataset and thus cannot fully utilize the available labeled data [39]. Our work follows the setting in [39] which trains object detectors on both labeled and unlabeled data and uses all labeled COCO data during training. Unlike the data distillation approach presented in [39] which can be viewed as a self-training-based approach, our approach combines self-supervised learning and consistency losses for SSOD. In experiments, we will show that our approach and the data distillation approach are complementary to some extent.

3. Approach

In this section, we will first give the definition of our Semi-Supervised Object Detection (SSOD) problem (see Section 3.1), then describe our overall framework (see Section 3.2), and finally introduce our proposal learning approach, consisting of the self-supervised proposal learning module (see Section 3.3) and consistency-based proposal learning module (see Section 3.4). If not specified, the contents we described here are the training procedures since we aim to train object detectors under the SSOD setting.

3.1. Problem Definition

In SSOD, a set of labeled data $\mathcal{D}^l = \{(\mathbf{I}, \mathcal{G})\}$ and a set of unlabeled data $\mathcal{D}^u = \{\mathbf{I}\}$ are given, where \mathbf{I} and \mathcal{G} denote an image and ground truth labels respectively. In object detection, \mathcal{G} consists of a set of objects with locations and object classes. Our goal of SSOD is to train object detectors on both labeled data \mathcal{D}^l and unlabeled data \mathcal{D}^u .

3.2. The Overall Framework

The overall framework of our approach is shown in Fig. 2. As in standard proposal-based object detectors [28, 41], during the forward process, first, an input image \mathbf{I} is fed into a CNN backbone (e.g., ResNet-50 [18] with feature pyramid networks [28]) with parameters θ^b , which produces image convolutional feature maps $\mathbf{F}^b(\mathbf{I}; \theta^b)$. Then, a Region Proposal Network (RPN) with parameters θ^{rpn} takes $\mathbf{F}^b(\mathbf{I}; \theta^b)$ as inputs to generate region proposals $\mathcal{P}(\mathbf{F}^b(\mathbf{I}; \theta^b); \theta^{\text{rpn}})$. We use $\mathbf{F}^b, \mathcal{P}$ later for simplification, dropping the dependence on $\mathbf{I}, \theta^b, \mathbf{F}^b, \theta^{\text{rpn}}$. Next, an RoIAlign [17] layer takes each proposal $\mathbf{P}_n = (x_n, y_n, w_n, h_n) \in \mathcal{P}$ and \mathbf{F}^b as inputs to extract proposal convolutional feature maps $\mathbf{F}_n^{\text{p-conv}}$ (simplification of $\mathbf{F}^{\text{p-conv}}(\mathbf{F}^b, \mathbf{P}_n)$, dropping the dependence on $\mathbf{F}^b, \mathbf{P}_n$), where (x_n, y_n, w_n, h_n) denotes the location of the i^{th} proposal $\mathbf{P}_n, i \in \{1, 2, \dots, N\}$, and N is the number of proposals in \mathcal{P} . After that, $\mathbf{F}_n^{\text{p-conv}}$ is fed into a Region-based CNN

(R-CNN) to generate proposal features $\mathbf{F}^p(\mathbf{F}_n^{\text{p-conv}}; \theta^{\text{r-cnn}})$ and predictions (i.e., bounding box classification predictions $\mathbf{C}^p(\mathbf{F}^p(\mathbf{F}_n^{\text{p-conv}}; \theta^{\text{r-cnn}}); \theta^{\text{cls}})$ and bounding box regression predictions $\mathbf{R}^p(\mathbf{F}^p(\mathbf{F}_n^{\text{p-conv}}; \theta^{\text{r-cnn}}); \theta^{\text{reg}})$), where $\theta^{\text{r-cnn}}, \theta^{\text{cls}}$, and θ^{reg} denote parameters of the R-CNN to generate proposal features, bounding box classification predictions, and bounding box regression predictions, respectively. We use $\mathbf{F}_n^p, \mathbf{C}_n^p, \mathbf{R}_n^p$ later for simplification, dropping the dependence on $\mathbf{F}_n^{\text{p-conv}}, \theta^{\text{r-cnn}}, \mathbf{F}_n^p, \theta^{\text{cls}}, \theta^{\text{reg}}$.

For each labeled data $(\mathbf{I}, \mathcal{G}) \in \mathcal{D}^l$, it is straightforward to train object detectors according to the standard fully-supervised learning loss defined in Eq. (1), where the first and second terms denote the RPN loss and R-CNN loss respectively. This loss is optimized w.r.t. $\theta^b, \theta^{\text{rpn}}, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}$ to train object detectors during the back-propagation process. More details of the loss function can be found in [41].

$$\begin{aligned} L^{\text{sup}}(\mathbf{I}, \mathcal{G}; \theta^b, \theta^{\text{rpn}}, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}) \\ = L^{\text{rpn}}(\mathbf{I}, \mathcal{G}; \theta^b, \theta^{\text{rpn}}) + \frac{1}{N} \sum_n L^{\text{r-cnn}}(\mathbf{I}, \mathbf{P}_n, \mathcal{G}, \theta^b, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}) \\ = L^{\text{rpn}}(\mathbf{F}^b, \mathcal{G}; \theta^{\text{rpn}}) + \frac{1}{N} \sum_n L^{\text{r-cnn}}(\mathbf{C}_n^p, \mathbf{R}_n^p, \mathcal{G}). \end{aligned} \quad (1)$$

However, for unlabeled data $\mathbf{I} \in \mathcal{D}^u$, there is no available ground truth labels \mathcal{G} . Thus we cannot use Eq. (1) to train object detectors on \mathcal{D}^u . To train object detectors also on \mathcal{D}^u , we present a proposal learning approach, consisting of a self-supervised proposal learning module and a consistency-based proposal learning module, to also learn proposal features (i.e., \mathbf{F}_n^p) and predictions (i.e., $\mathbf{C}_n^p, \mathbf{R}_n^p$) from \mathcal{D}^u . It is possible to also benefit RPN from \mathcal{D}^u . We only focus on the R-CNN-related parts, because 1) the final object detection results are from R-CNN-related parts and thus improving the R-CNN-related parts will benefit object detectors directly; 2) gradients will also be back-propagated from the R-CNN-related parts to the CNN backbone to learn better \mathbf{F}^b , which could potentially improve RPN.

For the n^{th} proposal \mathbf{P}_n of image \mathbf{I} , during the forward process, we first generate a set of noisy proposal features $\hat{\mathcal{F}}_n^p = \{\hat{\mathbf{F}}_{nk}^p\}_{k=1}^K$ and predictions $\hat{\mathcal{C}}_n^p = \{\hat{\mathbf{C}}_{nk}^p\}_{k=1}^K, \hat{\mathcal{R}}_n^p = \{\hat{\mathbf{R}}_{nk}^p\}_{k=1}^K$, where K denotes the number of noisy proposal features for each \mathbf{P}_n . As we stated in Section 1, we add noise to proposal features to share convolutional feature computations for the CNN backbone (i.e., \mathbf{F}^b) between noisy proposal feature computations and the original proposal feature computations for efficiency. More specifically, we add random noise $\{\epsilon_{nk}\}_{k=1}^K$ to proposal convolutional feature maps $\mathbf{F}_n^{\text{p-conv}}$, which generates a set of noisy proposal convolutional feature maps $\hat{\mathcal{F}}_n^{\text{p-conv}} = \{\hat{\mathbf{F}}_n^{\text{p-conv}}(\mathbf{F}_n^{\text{p-conv}}, \epsilon_{nk})\}_{k=1}^K$, see Fig. 2. We use $\hat{\mathbf{F}}_{nk}^{\text{p-conv}}$ for simplification, dropping the dependence on $\mathbf{F}_n^{\text{p-conv}}, \epsilon_{nk}$. Similar to the procedure to generate $\mathbf{F}_n^p, \mathbf{C}_n^p, \mathbf{R}_n^p$, the noisy

proposal feature maps are fed into the R-CNN to generate the noisy proposal features $\hat{\mathcal{F}}_n^p = \{\hat{\mathbf{F}}_{nk}^p\}_{k=1}^K$ and predictions $\hat{\mathcal{C}}_n^p = \{\hat{\mathbf{C}}_{nk}^p\}_{k=1}^K$, $\hat{\mathcal{R}}_n^p = \{\hat{\mathbf{R}}_{nk}^p\}_{k=1}^K$ (we drop the dependence on $\hat{\mathbf{F}}_{nk}^{p-\text{conv}}$, $\theta^{\text{r-cnn}}$, $\hat{\mathbf{F}}_{nk}^p$, θ^{cls} , θ^{reg} for notation simplification).

For Self-Supervised Proposal Learning (SSPL), as shown in Fig. 2, during the forward process, we pass the original proposal features \mathbf{F}_n^p and the noisy proposal features $\hat{\mathcal{F}}_n^p$ through a small SSPL network with parameters θ^{self} . The outputs of the SSPL network and the proposal location $\mathbf{P}_n = (x_n, y_n, w_n, h_n)$ are taken as inputs to compute SSPL loss $L^{\text{self}}(\{\mathbf{F}_n^p\}_{n=1}^N, \hat{\mathcal{F}}_n^p, \mathbf{P}_n; \theta^{\text{self}})$ which is defined in later Eq. (6). Since this loss does not take any ground truth labels \mathcal{G} as inputs, by optimizing this loss w.r.t. $\mathbf{F}_n^p, \hat{\mathcal{F}}_n^p, \theta^{\text{self}}$, i.e., $\theta^b, \theta^{\text{r-cnn}}, \theta^{\text{self}}$, during the back-propagation process, we can learn proposal features also from unlabeled data. We will give more details about this module in Section 3.3.

For consistency-based proposal learning, as shown in Fig. 2, during the forward process the original proposal predictions $\mathbf{C}_n^p, \mathbf{R}_n^p$ and the noisy proposal predictions $\hat{\mathcal{C}}_n^p, \hat{\mathcal{R}}_n^p$ are taken as inputs to compute loss $L^{\text{cons}}(\mathbf{C}_n^p, \mathbf{R}_n^p, \hat{\mathcal{C}}_n^p, \hat{\mathcal{R}}_n^p)$ which is defined in later Eq. (9). Following [33, 54], this loss is optimized w.r.t. $\hat{\mathcal{C}}_n^p, \hat{\mathcal{R}}_n^p$ (not w.r.t. $\mathbf{C}_n^p, \mathbf{R}_n^p$), i.e., $\theta^b, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}$, during the back-propagation process. Computing this loss does not require any ground truth labels \mathcal{G} , and thus we can learn proposal features and predictions also from unlabeled data. We will give more details about this module in Section 3.4.

We apply the standard fully-supervised loss defined in Eq. (1) to labeled data \mathcal{D}^l , and the self-supervised proposal learning loss $L^{\text{self}}(\{\mathbf{F}_n^p\}_{n=1}^N, \hat{\mathcal{F}}_n^p, \mathbf{P}_n; \theta^{\text{self}})$ and the consistency-based proposal learning loss $L^{\text{cons}}(\mathbf{C}_n^p, \mathbf{R}_n^p, \hat{\mathcal{C}}_n^p, \hat{\mathcal{R}}_n^p)$ to unlabeled data \mathcal{D}^u . The object detectors are trained on $\mathcal{D}^l, \mathcal{D}^u$ by optimizing the loss Eq. (2) w.r.t. $\theta^b, \theta^{\text{rpn}}, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}, \theta^{\text{self}}$ during the back-propagation process.

$$\begin{aligned} & L(\mathbf{I}, \mathcal{G}; \theta^b, \theta^{\text{rpn}}, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}, \theta^{\text{self}}) \\ &= \frac{1}{|\mathcal{D}^l|} \sum_{(\mathbf{I}, \mathcal{G}) \in \mathcal{D}^l} L^{\text{sup}}(\mathbf{I}, \mathcal{G}; \theta^b, \theta^{\text{rpn}}, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}) \\ &+ \frac{1}{|\mathcal{D}^u|} \sum_{\mathbf{I} \in \mathcal{D}^u} \frac{1}{N} \sum_n L^{\text{self}}(\{\mathbf{F}_n^p\}_{n=1}^N, \hat{\mathcal{F}}_n^p, \mathbf{P}_n; \theta^{\text{self}}) \\ &+ \frac{1}{|\mathcal{D}^u|} \sum_{\mathbf{I} \in \mathcal{D}^u} \frac{1}{N} \sum_n L^{\text{cons}}(\mathbf{C}_n^p, \mathbf{R}_n^p, \hat{\mathcal{C}}_n^p, \hat{\mathcal{R}}_n^p). \end{aligned} \quad (2)$$

We can also apply the self-supervised and consistency-based proposal learning losses to both labeled and unlabeled data, following the semi-supervised learning works

[26, 33, 58]. Then the overall loss is written as Eq. (3).

$$\begin{aligned} & L(\mathbf{I}, \mathcal{G}; \theta^b, \theta^{\text{rpn}}, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}, \theta^{\text{self}}) \\ &= \frac{1}{|\mathcal{D}^l|} \sum_{(\mathbf{I}, \mathcal{G}) \in \mathcal{D}^l} L^{\text{sup}}(\mathbf{I}, \mathcal{G}; \theta^b, \theta^{\text{rpn}}, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}) \\ &+ \frac{1}{|\mathcal{D}^l| + |\mathcal{D}^u|} \sum_{\mathbf{I} \in \mathcal{D}^l, \mathcal{D}^u} \frac{1}{N} \sum_n L^{\text{self}}(\{\mathbf{F}_n^p\}_{n=1}^N, \hat{\mathcal{F}}_n^p, \mathbf{P}_n; \theta^{\text{self}}) \\ &+ \frac{1}{|\mathcal{D}^l| + |\mathcal{D}^u|} \sum_{\mathbf{I} \in \mathcal{D}^l, \mathcal{D}^u} \frac{1}{N} \sum_n L^{\text{cons}}(\mathbf{C}_n^p, \mathbf{R}_n^p, \hat{\mathcal{C}}_n^p, \hat{\mathcal{R}}_n^p). \end{aligned} \quad (3)$$

During inference, we simply keep the parts of the standard proposal-based object detectors, see the blue arrows in Fig. 2. Therefore, our approach does not bring any extra inference computations.

3.3. Self-Supervised Proposal Learning

Previous works have shown that object detectors can benefit from context [2, 8, 20, 34] and should be noise-robust [32, 52]. Our self-supervised proposal learning module uses a proposal location loss and a contrastive loss to learn context-aware and noise-robust proposal features respectively.

To compute proposal location loss, we use proposal location prediction as the pretext task, inspired by the approach in [9]. More specifically, we pass $\mathbf{F}_n^p, \hat{\mathcal{F}}_n^p$ through two fully-connected layers with parameters $\theta^{\text{self-loc}}$ and a sigmoid layer to compute location predictions $\mathbf{L}_n^p, \hat{\mathcal{L}}_n^p = \{\hat{\mathbf{L}}_{nk}^p\}_{k=1}^K$, where the numbers of the outputs of the two fully-connected layers are 1024 and 4 respectively. Here we drop the dependence on $\mathbf{F}_n^p, \hat{\mathcal{F}}_n^p, \theta^{\text{self-loc}}$ for notation simplification. Then we use ℓ_2 distance to compute proposal location loss, see Eq. (4), where $\tilde{\mathbf{P}}_n = (x_n/W, y_n/H, w_n/W, h_n/H)$ is a normalized version of \mathbf{P}_n , and W, H denote the width and height of image \mathbf{I} respectively.

$$\begin{aligned} & L^{\text{self-loc}}(\mathbf{F}_n^p, \hat{\mathcal{F}}_n^p, \mathbf{P}_n; \theta^{\text{self-loc}}) \\ &= L^{\text{self-loc}}(\mathbf{L}_n^p, \hat{\mathcal{L}}_n^p, \mathbf{P}_n) \\ &= \frac{1}{K+1} \left(\|\mathbf{L}_n^p - \tilde{\mathbf{P}}_n\|_2^2 + \sum_k \|\hat{\mathbf{L}}_{nk}^p - \tilde{\mathbf{P}}_n\|_2^2 \right). \end{aligned} \quad (4)$$

By optimizing this loss w.r.t. $\mathbf{F}_n^p, \hat{\mathcal{F}}_n^p, \theta^{\text{self-loc}}$, i.e., $\theta^b, \theta^{\text{r-cnn}}, \theta^{\text{self-loc}}$, we can learn context-aware proposal features because predicting proposal location in an image requires proposal features understanding some global information of the image. We do not use the relative patch location prediction task [9] directly, because images are large and there are always multiple objects in the same image for object detection, which makes relative patch location prediction hard to be solved.

To compute contrastive loss, we use instance discrimination as the pretext task, following [16, 53, 57]. More specifically, we first use a fully-connected layer with parameters $\theta^{\text{self-cont}}$ and an ℓ_2 normalization layer to project $\mathbf{F}_n^p, \hat{\mathbf{F}}_n^p$ to embedded proposal features $\mathbf{F}_{nk}^{\text{embed}}, \hat{\mathbf{F}}_n^{\text{embed}} = \{\hat{\mathbf{F}}_{nk}^{\text{embed}}\}_{k=1}^K$ (dropping the dependence on $\mathbf{F}_n^p, \hat{\mathbf{F}}_n^p, \theta^{\text{self-cont}}$), where the numbers of the outputs of the fully-connected layer is 128. Then the contrastive loss is written as Eq. (5), where τ is a temperature hyper-parameter.

$$\begin{aligned} & \mathcal{L}^{\text{self-cont}} \left(\{\mathbf{F}_n^p\}_{n=1}^N, \hat{\mathbf{F}}_n^p, \theta^{\text{self-cont}} \right) \\ &= \mathcal{L}^{\text{self-cont}} \left(\{\mathbf{F}_n^{\text{embed}}\}_{n=1}^N, \hat{\mathbf{F}}_n^{\text{embed}} \right) \\ &= -\frac{1}{K} \sum_k \log \frac{\exp((\hat{\mathbf{F}}_{nk}^{\text{embed}})^T \mathbf{F}_n^{\text{embed}} / \tau)}{\sum_{n'} \exp((\hat{\mathbf{F}}_{nk}^{\text{embed}})^T \mathbf{F}_{n'}^{\text{embed}} / \tau)}. \end{aligned} \quad (5)$$

By optimizing this loss w.r.t. $\mathbf{F}_n^p, \hat{\mathbf{F}}_n^p, \theta^{\text{self-cont}}$, *i.e.*, $\theta^b, \theta^{\text{r-cnn}}, \theta^{\text{self-cont}}$, noisy proposal features are enforced to be closer to their original proposal features than to other proposal features, which learns noise-robust proposal features and thus learns noise-robust object detectors.

By combining the proposal location loss in Eq. (4) and the contrastive loss in Eq. (5), the overall self-supervised proposal learning loss is written as Eq. (6), where $\lambda^{\text{self-loc}}, \lambda^{\text{self-cont}}$ are loss weights and $\theta^{\text{self}} = \{\theta^{\text{self-loc}}, \theta^{\text{self-cont}}\}$. This loss is optimized w.r.t. $\theta^b, \theta^{\text{r-cnn}}, \theta^{\text{self}}$ to learn proposal features.

$$\begin{aligned} & \mathcal{L}^{\text{self}} \left(\{\mathbf{F}_n^p\}_{n=1}^N, \hat{\mathbf{F}}_n^p, \theta^{\text{self}} \right) \\ &= \lambda^{\text{self-loc}} \mathcal{L}^{\text{self-loc}} \left(\mathbf{F}_n^p, \hat{\mathbf{F}}_n^p, \mathbf{P}_n; \theta^{\text{self-loc}} \right) \\ &+ \lambda^{\text{self-cont}} \mathcal{L}^{\text{self-cont}} \left(\{\mathbf{F}_n^p\}_{n=1}^N, \hat{\mathbf{F}}_n^p, \theta^{\text{self-cont}} \right). \end{aligned} \quad (6)$$

3.4. Consistency-Based Proposal Learning

To further train noise-robust object detectors, we apply consistency losses [33, 43, 54] to ensure consistency between noisy proposal predictions and their original proposal predictions. More precisely, we apply consistency losses to both bounding box classification and regression predictions.

For consistency loss for bounding box classification predictions $\mathbf{C}_n^p, \hat{\mathbf{C}}_n^p$, we use KL divergence as the loss to enforce class predictions from noisy proposals and their original proposals to be consistent, following [33, 54], see Eq. (7).

$$\mathcal{L}^{\text{cons-cls}} \left(\mathbf{C}_n^p, \hat{\mathbf{C}}_n^p \right) = \frac{1}{K} \sum_k \text{KL} \left(\mathbf{C}_n^p \parallel \hat{\mathbf{C}}_{nk}^p \right). \quad (7)$$

Unlike image classification containing only classification results, object detection also predicts object locations. To further ensure proposal prediction consistency, we compute consistency loss in Eq. (8) to enforce object location

predictions from noisy proposals and their original proposals to be consistent. Here we use the standard bounding box regression loss, *i.e.*, smoothed ℓ_1 loss [14]. We only selected the easiest noisy proposal feature to compute this loss for training stability.

$$\mathcal{L}^{\text{cons-reg}} \left(\mathbf{R}_n^p, \hat{\mathbf{R}}_n^p \right) = \min_k \left(\text{smooth}_{\ell_1} \left(\mathbf{R}_n^p - \hat{\mathbf{R}}_{nk}^p \right) \right). \quad (8)$$

By combining the consistency losses for bounding box classification predictions in Eq. (7) and for bounding box regression predictions in Eq. (8), the overall consistency-based proposal learning loss is written as Eq. (9), where $\lambda^{\text{cons-cls}}, \lambda^{\text{cons-reg}}$ are loss weights. Following [33, 54], this loss is optimized w.r.t. $\hat{\mathbf{C}}_n^p, \hat{\mathbf{R}}_n^p$ (not w.r.t. $\mathbf{C}_n^p, \mathbf{R}_n^p$), *i.e.*, $\theta^b, \theta^{\text{r-cnn}}, \theta^{\text{cls}}, \theta^{\text{reg}}$. Then we can learn more noisy-robust proposal features and predictions.

$$\begin{aligned} & \mathcal{L}^{\text{cons}} \left(\mathbf{C}_n^p, \mathbf{R}_n^p, \hat{\mathbf{C}}_n^p, \hat{\mathbf{R}}_n^p \right) \\ &= \lambda^{\text{cons-cls}} \mathcal{L}^{\text{cons-cls}} \left(\mathbf{C}_n^p, \hat{\mathbf{C}}_n^p \right) + \lambda^{\text{cons-reg}} \mathcal{L}^{\text{cons-reg}} \left(\mathbf{R}_n^p, \hat{\mathbf{R}}_n^p \right). \end{aligned} \quad (9)$$

4. Experiments

In this section, we will conduct thorough experiments to analyze our proposal learning approach and its components for semi-supervised object detection.

4.1. Experimental Setup

4.1.1 Dataset and evaluation metrics.

We evaluate our approach on the challenging COCO dataset [30]. The COCO dataset contains more than 200K images for 80 object classes. Unlike many semi-supervised object detection works conducting experiments on a simulated setting by splitting a fully annotated dataset into labeled and unlabeled subsets, we use all available labeled and unlabeled training data in COCO as our \mathcal{D}^l and \mathcal{D}^u respectively, following [39]. More precisely, we use the COCO train2017 set (118K images) as \mathcal{D}^l and the COCO unlabeled2017 set (123K images) as \mathcal{D}^u to train object detectors. In addition, we use the COCO val2017 set (5K images) for validation and ablation studies, and the COCO test-dev2017 set (20K images) for testing.

We use the standard COCO criterion as our evaluation metrics, including AP (averaged average precision over different IoU thresholds, the primary evaluation metric of COCO), AP₅₀ (average precision for IoU threshold 0.5), AP₇₅ (average precision for IoU threshold 0.75), AP_S (AP for small objects), AP_M (AP for medium objects), AP_L (AP for large objects).

Table 1. Experimental results of different components of our proposal learning approach on the COCO val2017 set. ResNet-50 is chosen as our CNN backbone here. No “✓”, “ $L_{\text{self-loc}}$ ”, “ $L_{\text{self-cont}}$ ”, “ $L_{\text{cons-cls}}$ ”, “ $L_{\text{cons-reg}}$ ”, “PLLD”, and “FSWA” denote the fully-supervised baseline, the proposal location loss in Eq. (4), the contrastive loss in Eq. (5), the consistency loss for bounding box classification predictions in Eq. (7), the consistency loss for bounding box regression predictions in Eq. (8), Proposal Learning for Labeled Data, and Fast Stochastic Weight Averaging, respectively.

$L_{\text{self-loc}}$	$L_{\text{self-cont}}$	$L_{\text{cons-cls}}$	$L_{\text{cons-reg}}$	PLLD	FSWA	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
						37.4	58.9	40.7	21.5	41.1	48.6
✓						37.6 \uparrow 0.2	58.9 \uparrow 0.0	40.7 \uparrow 0.0	21.4 \downarrow 0.1	41.1 \uparrow 0.0	49.2 \uparrow 0.6
	✓					37.6 \uparrow 0.2	59.0 \uparrow 0.1	41.2 \uparrow 0.5	21.2 \downarrow 0.3	41.0 \downarrow 0.1	49.0 \uparrow 0.4
✓	✓					37.7 \uparrow 0.3	59.2 \uparrow 0.3	40.6 \downarrow 0.1	21.8 \uparrow 0.3	41.3 \uparrow 0.2	49.1 \uparrow 0.5
		✓				37.8 \uparrow 0.4	59.2 \uparrow 0.3	41.0 \uparrow 0.3	21.6 \uparrow 0.1	41.2 \uparrow 0.1	50.1 \uparrow 1.5
			✓			37.6 \uparrow 0.2	59.0 \uparrow 0.1	40.8 \uparrow 0.1	21.0 \downarrow 0.5	41.3 \uparrow 0.2	49.2 \uparrow 0.6
		✓	✓			37.9 \uparrow 0.5	59.2 \uparrow 0.3	40.9 \uparrow 0.2	21.4 \downarrow 0.1	41.1 \uparrow 0.0	50.6 \uparrow 2.0
✓	✓	✓	✓			38.0 \uparrow 0.6	59.2 \uparrow 0.3	41.1 \uparrow 0.4	21.6 \uparrow 0.1	41.5 \uparrow 0.4	50.4 \uparrow 1.8
✓	✓	✓	✓	✓		38.1 \uparrow 0.7	59.3 \uparrow 0.4	41.2 \uparrow 0.5	21.7 \uparrow 0.2	41.2 \uparrow 0.1	50.7\uparrow2.1
					✓	37.5 \uparrow 0.1	59.0 \uparrow 0.1	40.7 \uparrow 0.0	22.2 \uparrow 0.7	41.1 \uparrow 0.0	48.6 \uparrow 0.0
✓	✓	✓	✓	✓	✓	38.4\uparrow1.0	59.7\uparrow0.8	41.7\uparrow1.0	22.6\uparrow1.1	41.8\uparrow0.7	50.6 \uparrow 2.0

4.1.2 Implementation details.

In our experiments, we choose Faster R-CNN [41] with feature pyramid networks [28] and RoIAlign [17] as our proposal-based object detectors, which are the foundation of many recent state-of-the-art object detectors. Different CNN backbones, including ResNet-50 [18], ResNet-101 [18], ResNeXt-101-32 \times 4d [55], and ResNeXt-101-32 \times 4d with Deformable ConvNets [63] (ResNeXt-101-32 \times 4d+DCN), are chosen.

We train object detectors on 8 NVIDIA Tesla V100 GPUs for 24 epochs, using stochastic gradient descent with momentum 0.9 and weight decay 0.0001. During each training mini-batch, we randomly sample one labeled image and one unlabeled image for each GPU, and thus the effective mini-batch size is 16. Learning rate is set to 0.01 and is divided by 10 at the 16th and 22nd epochs. We use linear learning rate warm up to increase learning rate from 0.01/3 to 0.01 linearly in the first 500 training iterations. In addition, object detectors are first trained only on labeled data using loss Eq. (1) for 6 epochs. We also use fast stochastic weight averaging for checkpoints from the last few epochs for higher performance, following [1].

Loss weights $\lambda^{\text{self-loc}}$, $\lambda^{\text{self-cont}}$, $\lambda^{\text{cons-cls}}$, $\lambda^{\text{cons-reg}}$ are set to 0.25, 1, 1, 0.5, respectively. We add two types of noise, DropBlock [12] with block size 2 and SpatialDropout [48] with dropout ratio 1/64, to proposal convolutional feature maps. Other types of noise are also possible and we find that these two types of noise work well. The number of noisy proposal features for each proposal is set to 4, *i.e.*, $K = 4$. The temperature hyper-parameter τ in Eq. (5) is set to 0.1, following [57]. Images are resized so that the shorter side is 800 pixels with/without random horizontal flipping for training/testing. Considering that most of the proposals mainly contain backgrounds, we only choose the

positive proposals for labeled data and the proposals having maximum object score larger than 0.5 for unlabeled data to compute proposal learning based losses, which ensures networks focusing more on objects than backgrounds.

Our experiments are implemented based on the PyTorch [36] deep learning framework and the MMDetection [6] toolbox.

4.2. Ablation Studies

In this part, we conduct elaborate experiments to analyze the influence of different components of our proposal learning approach, including different losses defined in Eq. (4), (5), (7), (8), applying proposal learning for labeled data, and Fast Stochastic Weight Averaging (FSWA). Without loss of generality, we only use ResNet-50 as our CNN backbone. For the first two subparts, we only apply proposal learning to unlabeled data, *i.e.*, using Eq. (2). For the first three subparts, we do not use FSWA. Results are shown in Table 1, where the first row reports results from fully-supervised baseline, *i.e.*, training object detectors only on labeled data using Eq. (1). For all experiments, we fix the initial random seed during training to ensure that the performance gains come from our approach instead of randomness.

Self-supervised proposal learning. We first discuss the influence of our self-supervised proposal learning module. As shown in Table 1, compared to the fully-supervised baseline, both the proposal location loss and the contrastive loss obtain higher AP (37.6% *vs.* 37.4% and 37.6% *vs.* 37.4% respectively). The combination of these two losses, which forms the whole self-supervised proposal learning module, improves AP from 37.4% to 37.7%, which confirms the effectiveness of the self-supervised proposal learning module. The AP gains come from that this module learns better proposal features from unlabeled data.

Consistency-based proposal learning. We then discuss the influence of our consistency-based proposal learning module. From Table 1, we observe that applying consistency losses to both bounding box classification and regression predictions obtains higher AP than the fully-supervised baseline (37.8% vs. 37.4% and 37.6% vs. 37.4% respectively). The combination of these two consistency losses, which forms the whole consistency-based proposal learning module, improves AP from 37.4% to 37.9%. The AP gains come from that this module learns better proposal features and predictions from unlabeled data. In addition, after combining the consistency-based proposal learning and self-supervised proposal learning modules, *i.e.*, using our whole proposal learning approach, the AP is improved to 38.0% further, which shows the complementarity between our two modules.

Proposal learning for labeled data. We also apply our proposal learning approach to labeled data, *i.e.*, using Eq. (3). From Table 1, we see that proposal learning for labeled data boosts AP from 38.0% to 38.1%. This is because our proposal learning benefits from more training data. The results also show that our proposal learning can improve fully-supervised object detectors potentially, but since we focus on semi-supervised object detection, we would like to explore this in the future.

Fast stochastic weight averaging. We finally apply FSWA to our approach and obtain 38.4% AP, as shown in Table 1. This result suggests that FSWA can also boost performance for semi-supervised object detection. To perform a fair comparison, we also apply FSWA to the fully-supervised baseline. FSWA only improves AP from 37.4% to 37.5%. The results demonstrate FSWA gives more performance gains for our approach compared to the fully-supervised baseline.

According to these results, in the rest of our paper, we use all components of our proposal learning on both labeled and unlabeled data (*i.e.*, using Eq. (3)), and apply FSWA to our approach.

4.3. Main Results

We report the result comparisons among fully-supervised baselines, Data Distillation (DD) [39], and our approach in Table 2 on the COCO *test-dev2017* set. As we can see, our approach obtains consistently better results compared to the fully-supervised baselines for different CNN backbones. In addition, both DD and our approach obtain higher APs than the fully-supervised baselines, which demonstrates that training object detectors on both labeled and unlabeled data outperforms training object detectors only on labeled data, confirming the potentials of semi-supervised object detection. Using our approach alone obtains comparable APs compared to DD.

In particular, we also evaluate the efficiency of our

Table 2. Experimental result comparisons among fully-supervised baselines (no “✓”), Data Distillation (DD) [39], and our approach (Ours) on the COCO *test-dev2017* set. Different CNN backbones are chosen. Results of DD are reproduced by ourselves and are comparable with or even better than the results reported in the original DD paper.

CNN backbone	DD	Ours	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
ResNet-50	✓		37.7	59.6	40.8	21.6	40.6	47.2
			38.5	60.4	41.7	22.5	41.9	47.4
		✓	38.6	60.2	41.9	21.9	41.4	48.9
		✓	39.6	61.5	43.2	22.9	42.9	49.8
ResNet-101	✓		39.6	61.2	43.2	22.2	43.0	50.3
			40.6	62.2	44.3	23.2	44.4	50.9
		✓	40.4	61.8	44.2	22.6	43.6	51.6
		✓	41.3	63.1	45.3	23.4	45.0	52.7
ResNeXt-101-32×4d	✓		40.7	62.3	44.3	23.2	43.9	51.6
			41.8	63.6	45.6	24.5	45.4	52.4
		✓	41.5	63.2	45.4	23.9	44.8	52.8
		✓	42.8	64.4	46.9	24.9	46.4	54.4
ResNeXt-101-32×4d+DCN	✓		44.1	66.0	48.2	25.7	47.3	56.3
			45.4	67.0	49.6	27.3	49.0	57.9
		✓	45.1	66.8	49.2	26.4	48.3	57.6
		✓	46.2	67.7	50.4	27.6	49.6	59.1

method by combining with DD. More specifically, we first train object detectors using our approach, then follow DD to label unlabeled data, and finally re-train object detectors using both fully-supervised loss and proposal learning losses. The combination of our approach and DD obtains 39.6% (ResNet-50), 41.3% (ResNet-101), 42.8% (ResNeXt-101-32×4d), and 46.2% (ResNeXt-101-32×4d+DCN) APs, which outperforms the fully-supervised baselines by about 2.0% on average and DD alone by 0.9% on average. The results demonstrate that our approach and DD are complementary to some extent.

5. Conclusion

In this paper, we focus on semi-supervised object detection for proposal-based object detectors (a.k.a. two-stage object detectors). To this end, we present a proposal learning approach, which consists of a self-supervised proposal learning module and a consistency-based proposal learning module, to learn proposal features and predictions from both labeled and unlabeled data. The self-supervised proposal learning module learns context-aware and noise-robust proposal features by a proposal location loss and a contrastive loss respectively. The consistency-based proposal learning module learns noise-robust proposal features and predictions by consistency losses for both bounding box classification and regression predictions. Experimental results show that our approach outperforms fully-supervised baselines consistently. It is also worth mentioning that we can further boost detection performance by combining our approach and data distillation. In the future, we will explore more self-supervised learning and semi-supervised learning ways for semi-supervised object detection, and explore how to apply our approach to semi-supervised instance segmentation.

References

- [1] Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. There are many consistent explanations of unlabeled data: Why you should average. *arXiv preprint arXiv:1806.05594*, 2018.
- [2] Sean Bell, C Lawrence Zitnick, Kavita Bala, and Ross Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2874–2883, 2016.
- [3] Avrim Blum and Tom Mitchell. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 92–100, 1998.
- [4] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018.
- [5] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. Semi-supervised learning. 2010.
- [6] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. MMDetection: Open MMLab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.
- [7] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Weakly supervised object localization with multi-fold multiple instance learning. *IEEE transactions on pattern analysis and machine intelligence*, 39(1):189–203, 2016.
- [8] Santosh K Divvala, Derek Hoiem, James H Hays, Alexei A Efros, and Martial Hebert. An empirical study of context in object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1271–1278, 2009.
- [9] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015.
- [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*, 111(1):98–136, 2015.
- [11] Jiyang Gao, Jiang Wang, Shengyang Dai, Li-Jia Li, and Ram Nevatia. NOTE-RCNN: Noise tolerant ensemble RCNN for semi-supervised object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9508–9517, 2019.
- [12] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Dropblock: A regularization method for convolutional networks. In *Advances in neural information processing systems*, pages 10727–10737, 2018.
- [13] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.
- [14] Ross Girshick. Fast R-CNN. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [15] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019.
- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Judy Hoffman, Sergio Guadarrama, Eric S Tzeng, Ronghang Hu, Jeff Donahue, Ross Girshick, Trevor Darrell, and Kate Saenko. LSDA: Large scale detection through adaptation. In *Advances in Neural Information Processing Systems*, pages 3536–3544, 2014.
- [20] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3588–3597, 2018.
- [21] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring R-CNN. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6409–6418, 2019.
- [22] Jisoo Jeong, Seungeui Lee, Jeessoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. In *Advances in neural information processing systems*, pages 10758–10767, 2019.
- [23] Longlong Jing and Yingli Tian. Self-supervised visual feature learning with deep neural networks: A survey. *arXiv preprint arXiv:1902.06162*, 2019.
- [24] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer. Revisiting self-supervised visual representation learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1920–1929, 2019.
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [26] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.
- [27] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [28] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [29] Tsung-Yi Lin, Priyanka Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE transactions on pattern analysis and machine intelligence*, 2018.

- [30] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European conference on computer vision*, pages 740–755, 2014.
- [31] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. SSD: Single shot multibox detector. In *European conference on computer vision*, pages 21–37, 2016.
- [32] Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S Ecker, Matthias Bethge, and Wieland Brendel. Benchmarking robustness in object detection: Autonomous driving when winter is coming. *arXiv preprint arXiv:1907.07484*, 2019.
- [33] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- [34] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 891–898, 2014.
- [35] Nhu-Van Nguyen, Christophe Rigaud, and Jean-Christophe Burie. Semi-supervised object detection with unlabeled data. In *VISIGRAPP (5: VISAPP)*, pages 289–296, 2019.
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8024–8035, 2019.
- [37] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [38] Siyuan Qiao, Wei Shen, Zhishuai Zhang, Bo Wang, and Alan Yuille. Deep co-training for semi-supervised image recognition. In *European conference on computer vision*, pages 135–152, 2018.
- [39] Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omniscient learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4119–4128, 2018.
- [40] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2017.
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [43] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in neural information processing systems*, pages 1163–1171, 2016.
- [44] 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 *Proceedings of the IEEE international conference on computer vision*, pages 8430–8439, 2019.
- [45] Peng Tang, Chunyu Wang, Xinggong Wang, Wenyu Liu, Wenjun Zeng, and Jingdong Wang. Object detection in videos by high quality object linking. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [46] Yuxing Tang, Josiah Wang, Boyang Gao, Emmanuel Delandrea, Robert Gaizauskas, and Liming Chen. Large scale semi-supervised object detection using visual and semantic knowledge transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2119–2128, 2016.
- [47] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017.
- [48] Jonathan Tompson, Ross Goroshin, Arjun Jain, Yann LeCun, and Christoph Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 648–656, 2015.
- [49] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggong Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *arXiv preprint arXiv:1908.07919*, 2019.
- [50] Keze Wang, Liang Lin, Xiaopeng Yan, Ziliang Chen, Dongyu Zhang, and Lei Zhang. Cost-effective object detection: Active sample mining with switchable selection criteria. *IEEE transactions on neural networks and learning systems*, 30(3):834–850, 2018.
- [51] Keze Wang, Xiaopeng Yan, Dongyu Zhang, Lei Zhang, and Liang Lin. Towards human-machine cooperation: Self-supervised sample mining for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1605–1613, 2018.
- [52] Xiaolong Wang, Abhinav Shrivastava, and Abhinav Gupta. A-Fast-RCNN: Hard positive generation via adversary for object detection. In *Proceedings of the IEEE Conference on computer vision and pattern recognition*, pages 2606–2615, 2017.
- [53] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018.

- [54] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation. *arXiv preprint arXiv:1904.12848*, 2019.
- [55] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [56] David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In *Proceedings of the 33rd annual meeting of the association for computational linguistics*, pages 189–196, 1995.
- [57] Mang Ye, Xu Zhang, Pong C Yuen, and Shih-Fu Chang. Unsupervised embedding learning via invariant and spreading instance feature. In *Proceedings of the IEEE Conference on computer vision and pattern recognition*, pages 6210–6219, 2019.
- [58] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S⁴L: Self-supervised semi-supervised learning. In *Proceedings of the IEEE international conference on computer vision*, pages 1476–1485, 2019.
- [59] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666, 2016.
- [60] Zhishuai Zhang, Siyuan Qiao, Cihang Xie, Wei Shen, Bo Wang, and Alan L Yuille. Single-shot object detection with enriched semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5813–5821, 2018.
- [61] Yuyin Zhou, Yan Wang, Peng Tang, Song Bai, Wei Shen, Elliot Fishman, and Alan Yuille. Semi-supervised 3D abdominal multi-organ segmentation via deep multi-planar co-training. In *IEEE winter conference on applications of computer vision*, pages 121–140, 2019.
- [62] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. 2002.
- [63] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable ConvNets v2: More deformable, better results. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9308–9316, 2019.