

Robust Object Detection With Inaccurate Bounding Boxes

Chengxin Liu¹, Kewei Wang¹, Hao Lu¹, Zhiguo Cao^{1*}, and Ziming Zhang²

¹ School of Artificial Intelligence and Automation,
Huazhong University of Science and Technology, China

² Worcester Polytechnic Institute, USA
{cx.liu,zgcao}@hust.edu.cn

Abstract. Learning accurate object detectors often requires large-scale training data with precise object bounding boxes. However, labeling such data is expensive and time-consuming. As the crowd-sourcing labeling process and the ambiguities of the objects may raise noisy bounding box annotations, the object detectors will suffer from the degenerated training data. In this work, we aim to address the challenge of learning robust object detectors with inaccurate bounding boxes. Inspired by the fact that localization precision suffers significantly from inaccurate bounding boxes while classification accuracy is less affected, we propose leveraging classification as a guidance signal for refining localization results. Specifically, by treating an object as a bag of instances, we introduce an Object-Aware Multiple Instance Learning approach (OA-MIL), featured with object-aware instance selection and object-aware instance extension. The former aims to select accurate instances for training, instead of directly using inaccurate box annotations. The latter focuses on generating high-quality instances for selection. Extensive experiments on synthetic noisy datasets (*i.e.*, noisy PASCAL VOC and MS-COCO) and a real noisy wheat head dataset demonstrate the effectiveness of our OA-MIL. Code is available at <https://github.com/cxliu0/OA-MIL>.

Keywords: Object Detection, Inaccurate Bounding Boxes, Noisy Labels, Multiple Instance Learning

1 Introduction

Despite remarkable progress has been witnessed in the field of object detection in recent years, the success of modern object detectors largely relies on large-scale datasets like ImageNet [10] and MS-COCO [28]. However, acquiring precise annotations is no easy task in professional and natural contexts. In practical applications with professional backgrounds (*e.g.*, agricultural crop observation and medical image processing), domain knowledge is often required to annotate objects. This situation leads to a dilemma, *i.e.*, practitioners without computer

* Corresponding author.

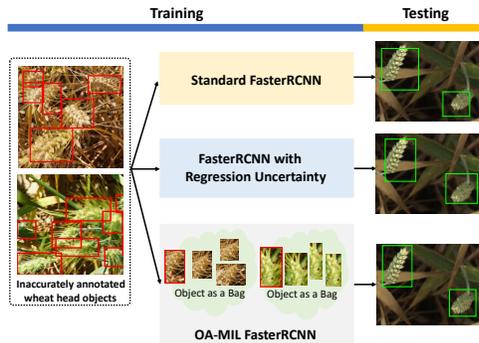


Fig. 1: Illustration of standard FasterRCNN [34], FasterRCNN with regression uncertainty [20], and our OA-MIL FasterRCNN on the ECCV wheat head detection challenge dataset. Given inaccurately annotated objects, we aim to learn a robust object detector by treating each object as a bag of instances. The inaccurate ground-truth boxes are in red and the predictions are in green.

vision background are not sure how to annotate high-quality boxes, while annotators without domain knowledge can also be difficult to annotate accurate object boxes. For example, recent wheat head detection challenge¹ that was hosted at the European Conference on Computer Vision (ECCV) workshop 2020 has shown that precise object bounding boxes are not easy to obtain, because in some domains the definition of the object is significantly different from generic objects in COCO, thus brings annotation ambiguities (Fig. 1). In these cases, there will be a demand calling for algorithms dealing with noisy bounding boxes. On the other hand, annotating a large amount of common objects in the natural context is expensive and time-consuming. To reduce the annotation cost [47], dataset producers may rely on social media platforms or crowd-sourcing platforms. Nevertheless, the above strategies would lead to low-quality annotations. Recent work [47] argues that the object detectors will suffer from the degenerated data. In addition, even large-scale datasets (*e.g.*, MS-COCO) are dedicated annotated, box ambiguities [20] still exist. Therefore, tackling noisy bounding boxes is a practical and meaningful task.

Recently, learning object detectors with noisy data have gained a surge of interest, several approaches [1,5,25,47] have made attempted to tackle noisy annotations. These approaches often assume that the noise occurs both on category labels and bounding box annotations, and devise a disentangled architecture to

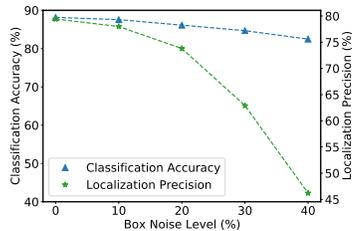


Fig. 2: Classification accuracy and localization precision of FasterRCNN on the simulated “noisy” PASCAL VOC 2007 dataset [14], where box annotations are randomly perturbed. With the box noise level increases, *i.e.*, ground-truth box becomes more and more inaccurate, localization precision drops significantly while the classification still maintains high accuracy.

¹ <https://www.kaggle.com/c/global-wheat-detection>

learn object detectors. Different from previous work, we focus on object detection with noisy bounding box annotations. The reasons are two-fold: i) due to the ambiguities of the objects [20] and the crowd-sourcing labeling process, box noise commonly exists in the real world; ii) object detection datasets [23] often involve object class verification, thus we consider noisy category labels are less severe than inaccurate bounding boxes.

Motivated by the observation that localization precision suffers significantly from inaccurate bounding boxes while classification accuracy is less affected (Fig. 2), we propose leveraging classification as a guidance signal for localization. Specifically, we present an Object-Aware Multiple Instance Learning approach by treating each object as a bag of instances, where the concept of the object bag is illustrated in Fig. 1. The idea is to select accurate instances from the object bags for training, instead of using inaccurate box annotations. Our approach is featured with object-aware instance selection and object-aware instance extension. The former is designed to select accurate instances and the latter to generate high-quality instances for selection. The optimization process involves jointly training the instance selector, the instance classifier, and the instance generator. To validate the effectiveness of our approach, we experiment on both synthetic noisy datasets (*i.e.*, noisy PASCAL VOC 2007 [14] and MS-COCO [28]) and real noisy wheat head dataset [8,9]. The main contributions are as follows:

- We contribute a novel view for learning object detectors with inaccurate bounding boxes by treating an object as a bag of instances;
- We present an Object-Aware Multiple Instance Learning approach, featured by object-aware instance selection and object-aware instance extension;
- OA-MIL exhibits generality on off-the-shelf object detectors and obtains promising results on the synthetic and the real noisy datasets.

2 Related Work

Learning With Noisy Labels. Training accurate DNNs under noisy labels has been an active research area. A major line of research focuses on the classification task, and develops various techniques to deal with noisy labels, such as sample selection [18,21], label correction [32,37], and robust loss functions [16,50]. Recently, much effort [1,5,25,33,47] has been devoted to the object detection task. Simon *et al.* [5] first investigate the impact of different types of label noise on object detection, and propose a per-object co-teaching method to alleviate the effect of noisy labels. On the other hand, Li *et al.* [25] propose a learning framework that alternately performs noise correction and model training to tackle noisy annotations, where the noisy annotations consist of noisy category labels and noisy bounding boxes. Xu *et al.* [47] further introduce a meta-learning based approach to tackle noisy labels by leveraging a few clean samples.

In contrast to previous works, we emphasize learning object detectors with inaccurate bounding boxes and contribute a novel Object-Aware MIL view to addressing this problem. In addition, we do not assume the accessibility to clean box annotations as previous work [47] does.

Weakly-Supervised Object Detection (WSOD). WSOD refers to learning object detectors with only image-level labels. The majority of previous works formulate WSOD as a multiple instance learning (MIL) problem [13], where each image is considered as a “bag” of instances (instances are tentative object proposals) with image-level label. Under this formulation, the learning process alternates between detector training and object location estimation. Since MIL leads to a non-convex optimization problem, solvers may get stuck in local optima. Accordingly, much effort [2,7,11,35,36,38] has been made to help the solution escape from local optima. Recently, deep MIL methods [3] emerge. However, the non-convexity problem remains. To address this problem, various techniques have been developed, including spatial regularization [3,12,44], context information [22,46], and optimization strategy [12,24,39,44,45]. For example, Zhang *et al.* [48] tackle noisy initialized object locations in WSOD and propose a self-directed localization network to identify noisy object instances.

In this work, we tackle learning object detectors with noisy box annotations, which is different from WSOD where only image-level labels are given. Despite we also formulate object detection as a MIL problem, we remark that our formulation is significantly different from WSOD in two aspects: i) we establish the concept of the bag on the object instead of on the image, which encodes object-level information; ii) we dynamically construct the object bag instead of using a fixed one as in WSOD, yielding a higher performance upper bound.

Semi-Supervised Object Detection (SSOD). SSOD aims to train object detectors with a large scale of image-level annotations and a few box-level annotations. Some prior works [40,41,42] address SSOD by knowledge transfer, where the information is transferred from source classes with bounding-box labels to target classes with only image-level labels. Other than the knowledge transfer paradigm, a recent work [15] adopts a training-mining framework and proposes a noise-tolerant ensemble RCNN to eliminate the harm of noisy labels.

However, previous SSOD methods generally assume the availability of clean bounding box annotations. In contrast, we only assume the accessibility to noisy bounding box annotations.

3 Object-Aware Multiple Instance Learning

In this work, we aim to learn a robust object detector with inaccurate bounding box annotations. Motivated by the observation that classification maintains high accuracy under noisy box annotations (Fig. 2), we suggest leveraging classification to guide localization. Intuitively, instead of using the inaccurate ground-truth boxes, we expect the classification branch to select more precise boxes for training. This idea derives the concept of object bag, where each object is formulated as a bag of instances for selection. Build upon the object bag, we present an Object-Aware Multiple Instance Learning approach that features object-aware instance selection and object-aware instance extension. In the following, we first introduce some preliminaries about MIL. Then, we present our Object-Aware

Multiple Instance Learning formulation. Finally, we show how to deploy our method on modern object detectors like FasterRCNN [34] and RetinaNet [27].

3.1 Preliminaries

Given image-level labels, an MIL method [7,45] in WSOD treats each image as a bag of instances, where instances are tentative object proposals. The learning process alternates between instance selection and instance classifier learning.

Formally, let $B_i \in \mathcal{B}$ denote the i^{th} bag (image) and \mathcal{B} denote the bag set (all training images). Each B_i is associated with a label $y_i \in \{1, -1\}$, where y_i indicates whether B_i contains positive instances. Also, let \mathbf{b}_i^j denote the j^{th} instance of bag B_i (*i.e.*, \mathbf{b}_i^j encodes the coordinates of an instance), where $j \in \{1, 2, \dots, N_i\}$ and N_i is the number of instances in B_i . With the definitions above, the instance selector $f(\mathbf{b}_i^j, \omega_f)$ with parameter ω_f is applied to a positive bag B_i to select the most positive instance $\mathbf{b}_i^{j^*}$, the index j^* is obtained by:

$$j^* = \arg \max_j f(\mathbf{b}_i^j, \omega_f), \quad (1)$$

where the instance selector $f(\mathbf{b}_i^j, \omega_f)$ takes an instance \mathbf{b}_i^j as input and outputs a confidence score that is in the range of $[-1, 1]$. Then, the selected instance $\mathbf{b}_i^{j^*}$ is used to train the instance classifier $g(\mathbf{b}_i^{j^*}, \omega_g)$ with parameter ω_g . The overall loss function is defined as:

$$L(\mathcal{B}, \omega_f, \omega_g) = \sum_i L_f(B_i, \omega_f) + L_g(B_i, \mathbf{b}_i^{j^*}, \omega_g), \quad (2)$$

where L_f and L_g are the loss of instance selector and instance classifier, respectively. Typically, L_f is defined as a standard hinge loss:

$$L_f(B_i, \omega_f) = \max(0, 1 - y_i \max_j f(\mathbf{b}_i^j, \omega_f)). \quad (3)$$

And the instance classifier loss L_g is defined as log-loss for classification.

3.2 Object-Aware MIL Formulation

Despite we formulate object detection as a MIL problem, we argue that the existing MIL paradigm in WSOD could not address the learning problem under noisy box annotations. First, since an image is defined as a bag in WSOD, the localization prior of objects is ignored. Second, the bags in WSOD are simply a collection of object proposals produced by off-the-shelf object proposal generators like selective search [43], which limits the detection performance.

Different from WSOD, in the context of our object bag, two challenges need to be solved: i) *how to select accurate instance in each object bag for training*; and ii) *how to generate high-quality instances for each object bag*.

To address the above challenges, we introduce an Object-Aware MIL formulation, which jointly optimizes the instance selector, the instance classifier, and the

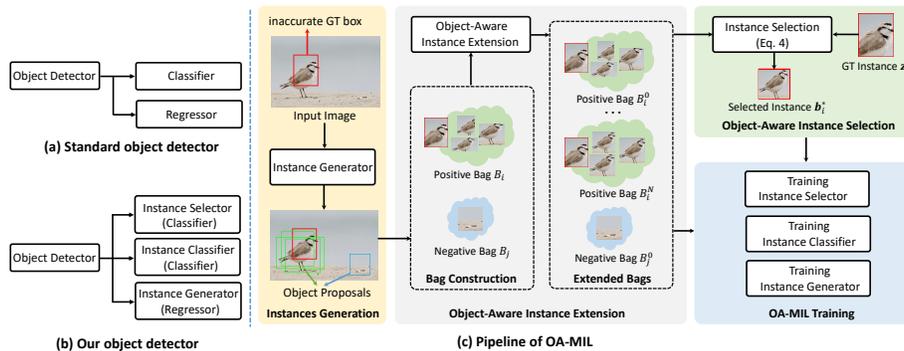


Fig. 3: **An overview of our OA-MIL formulation.** We augment the standard object detector (a) with an instance selector, forming our object detector (b). (c) illustrates the pipeline of OA-MIL. We first construct object bags based on the outputs of the instance generator (green and blue boxes), where the inaccurate ground-truth box (red box) is formulated as a positive bag B_i and the background box (blue box) is treated as a negative bag B_j . Then, object-aware instance extension is applied to obtain an extended bag set $\{B_i^0, B_i^1, \dots, B_i^N, B_j\}$. Based on this bag set and the noisy ground-truth instance \mathbf{z}_i , we adopt object-aware instance selection to select the best positive instance \mathbf{b}_i^* for training object detector (including instance selector, instance classifier, and instance generator).

instance generator. In the following, we first give the definition of object bag. Then, we introduce object-aware instance selection and object-aware instance extension, where the former is designed to select the accurate instance, while the latter aims to produce a set of high-quality instances for selection. Finally, we describe how to train the instance selector, the instance classifier, and the instance generator. Fig. 3 illustrates the pipeline of our OA-MIL formulation.

Bag Definition. We reuse the bag symbol in Sec. 3.1. But the definition of bag is different, *i.e.*, we treat each object as a bag. We denote $B_i \in \mathcal{B}$ as the i^{th} bag (object), and \mathcal{B} denotes the bag set (all objects in training images). A label $y_i \in \{1, -1\}$ is attached to each bag B_i . As illustrated in Fig. 3, we treat inaccurate ground-truth box as a positive bag B_i with $y_i = 1$, and the background box is treated as a negative bag B_j with $y_j = -1$. Suppose bag B_i contains N_i instances, we denote \mathbf{b}_i^j as the j^{th} instance of bag B_i , where $j \in \{1, \dots, N_i\}$. With object bag defined above, we naturally introduce Eq. (3) to train the instance selector.

Object-Aware Instance Selection. Since we treat each inaccurate ground-truth box as a bag of instances, the quality of the selected instance is essential for training an accurate object detector. Intuitively, we expect the selected instance covers the actual object as tight as possible. However, as the instance selector has poor discriminative ability in the early stage of training, the instance classifier and instance generator will inevitably suffer from the low-quality positive

instance. In some cases, poor instance initialization could render failure during training. As the inaccurate ground-truth box provides a strong prior of object localization, we jointly consider it and the selected instance to obtain a more suitable positive instance for training.

Specifically, we denote \mathbf{z}_i as the inaccurate ground-truth instance. We perform object-aware instance selection by merging \mathbf{z}_i and $\mathbf{b}_i^{j^*}$ as follows:

$$\mathbf{b}_i^* = \varphi(f(\mathbf{b}_i^{j^*}, \omega_f)) \cdot \mathbf{b}_i^{j^*} + (1 - \varphi(f(\mathbf{b}_i^{j^*}, \omega_f))) \cdot \mathbf{z}_i, \quad (4)$$

where $\mathbf{b}_i^{j^*}$ is the most positive instance selected by the instance selector, and φ is a mapping function, which adaptively assigns the coefficient for $\mathbf{b}_i^{j^*}$ and \mathbf{z}_i .

Recall that our goal is to select high-quality positive instances for training, thus we expect $\varphi(\cdot)$ to satisfy two conditions. First, higher weights should be assigned to $\mathbf{b}_i^{j^*}$ when $f(\mathbf{b}_i^{j^*}, \omega_f)$ has large value, because it indicates the confidence of the positive instance $\mathbf{b}_i^{j^*}$. Second, $\varphi(\cdot)$ should balance the weights of $\mathbf{b}_i^{j^*}$ and \mathbf{z}_i instead of relying on $\mathbf{b}_i^{j^*}$ when $f(\mathbf{b}_i^{j^*}, \omega_f)$ is close to 1. To satisfy the above conditions, we adopt a bounded exponential function as follows:

$$\varphi(x) = \min(x^\gamma, \theta), \quad (5)$$

where γ and θ are hyper-parameters, and $x \in [0, 1]$. A key property of Eq. (5) is that the ascent speed and the upper bound of φ are controllable.

Object-Aware Instance Extension. The quality of instances is another factor that affects the training process. In our formulation, the bags are dynamically constructed based on the outputs of the instance generator. Thus, the quality of bag instances can not always be guaranteed. Fortunately, the instances inside a positive bag are homogeneous, *i.e.*, instances are closely related to each other both on spatial location and class information. Therefore, it is possible to promote the quality of a positive bag by extending the positive instances.

We present two strategies for instance extension. The first strategy is to obtain new positive instances by recursively constructing positive bags. Specifically, we first obtain the initial object bag based on the noisy ground-truth boxes, then we use the most positive instance selected by Eq. (4) to construct a new positive bag. The process repeats until reaching the termination condition. This strategy is generic and applicable to any existing object detectors. The second strategy is to refine the positive instances in a multi-stage manner [4], which is suitable for object detectors that feature with a bounding box refinement module (*e.g.*, FasterRCNN [34]). The extended object bags are subsequently used to train the instance selector. Note that we do not extend negative bags.

Suppose we have conducted N times of instance extension, which produces a set of extended positive bags $\{B_i^0, B_i^1, \dots, B_i^N\}$, where B_i^0 denotes the initial object bag B_i (As shown in Fig. 3). We utilize the extended object bags to optimize the instance selector, the loss thus becomes:

$$L_f(\{B_i^k\}, \omega_f) = \sum_k L_f(B_i^k, \omega_f), \quad (6)$$

where $k \in \{0, 1, \dots, N\}$ only if B_i is a positive bag.

OA-MIL Training. Our OA-MIL involves jointly optimizing the instance selector, instance classifier, and the instance generator. The instance selector is trained using Eq. (6). As the instance classifier g (with parameter ω_g) is used to classify object, we adopt the binary-log-loss to train it:

$$L_g(B_i, \mathbf{b}_i^*, \omega_g) = - \sum_j \log(y_{i,j} \cdot (g(\mathbf{b}_i^j, \omega_g) - \frac{1}{2}) + \frac{1}{2}), \quad (7)$$

where $g(\mathbf{b}_i^j, \omega_g) \in (0, 1)$, which represents the probability of \mathbf{b}_i^j contains objects with positive class. $y_{i,j}$ is defined as:

$$y_{i,j} = \begin{cases} +1, & \text{if } y_i = 1 \text{ and } \text{IoU}(\mathbf{b}_i^j, \mathbf{b}_i^*) \geq 0.5 \\ -1, & \text{if } y_i = 1 \text{ and } \text{IoU}(\mathbf{b}_i^j, \mathbf{b}_i^*) < 0.5 \\ -1, & \text{if } y_i = -1 \end{cases}, \quad (8)$$

where IoU denotes the Intersection over Union between two instances. Note that the loss of each instance becomes $\log(g(\mathbf{b}_i^j, \omega_g))$ when $y_{i,j} = 1$, and $\log(1 - g(\mathbf{b}_i^j, \omega_g))$ otherwise.

One major difference between our MIL and WSOD MIL is that we jointly train a learnable instance generator, which is crucial for dealing with inaccurate bounding boxes. The loss function of the instance generator is as follows:

$$L_\phi(\mathcal{B}, \omega_\phi) = \sum_i \mathbb{1}(B_i) \cdot L_{reg}(B_i, \mathbf{b}_i^*, \omega_\phi), \quad (9)$$

where ω_ϕ is the parameters of instance generator, $\mathbb{1}(B_i)$ equals 1 if B_i is a positive bag, otherwise $\mathbb{1}(B_i)$ is 0 because a negative bag does not correspond to any actual objects, and L_{reg} is defined as:

$$L_{reg}(B_i, \mathbf{b}_i^*, \omega_\phi) = \sum_j \delta(y_{i,j}) \cdot \ell_{reg}(\mathbf{b}_i^j, \mathbf{b}_i^*), \quad (10)$$

where $\delta(\cdot)$ is the unit-impulse function ($\delta(y_{i,j})$ equals to 1 when $y_{i,j}$ is 1, otherwise 0), and ℓ_{reg} is a regression loss like ℓ_1 loss or smooth ℓ_1 loss [34].

To summarize, the overall loss function is formulated as:

$$L(\mathcal{B}, \omega_f, \omega_g, \omega_\phi) = \lambda \sum_i \sum_k L_f(B_i^k, \omega_f) + \sum_i L_g(B_i^0, \mathbf{b}_i^*, \omega_g) + \sum_i \mathbb{1}(B_i^0) \cdot L_{reg}(B_i^0, \mathbf{b}_i^*, \omega_\phi), \quad (11)$$

where λ is a balance parameter, \mathcal{B} is the extended object bags set, B_i^0 is the initial object bag, and \mathbf{b}_i^* is selected by Eq. (4).

3.3 Deployment to Off-the-Shelf Object Detectors

We remark that our formulation is general and is not limited to specific object detectors. To demonstrate the generality of our method, we apply our method on a two-stage detector—FasterRCNN [34] and a one-stage detector—RetinaNet [27]. Here we introduce the deployment procedure on FasterRCNN. It includes two

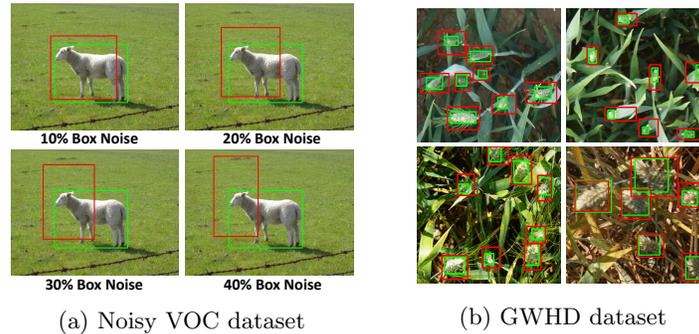


Fig. 4: **Examples of the inaccurate bounding boxes (red boxes) on VOC and GWHD dataset.** The clean ground-truth boxes are in green.

steps, the first step is to construct object bags and the second step is to apply OA-MIL on FasterRCNN. More details can be found in the supplementary.

Bag Construction. We construct object bags based on the outputs of the second stage of FasterRCNN. Specifically, we treat each inaccurately annotated object as a positive bag, where instances are positive anchors (object proposals) corresponding to a specific object.

OA-MIL Deployment. The instance selector and the instance classifier share the same classifier. The regressor in the second stage is treated as the instance generator. We perform object-aware instance extension by multi-stage refinement, which produces a set of extended object bags $\{B_i^0, B_i^1, \dots, B_i^N\}$. Then, object-aware instance selection is applied to select the best instance \mathbf{b}_i^* , where \mathbf{b}_i^* is used to train the instance generator (regressor), the instance selector (classifier), and the object detector (classifier). We follow Eq. (11) to train the second stage of FasterRCNN. Note that the training objective of RPN is the same as [34].

Implementation Details. We implement our method on FasterRCNN [34] with ResNet50-FPN [19,26] backbone. Following common practices [17], the model is trained with “1×” schedule. The hyper parameters are set as $\gamma = 7.5$, $\theta = 0.85$, $N = 4$, and λ is selected from $\{0.01, 0.1\}$ (depending on datasets and noise levels). Our implementation is based on MMDetection toolbox [6].

4 Results and Discussions

4.1 Datasets and Evaluation Metrics

Synthetic Noisy Dataset. Modern object detection datasets are delicately annotated and contain few inaccurate bounding boxes. Thus, we simulate noisy bounding boxes by perturbing the clean ones on two object detection datasets, including PASCAL VOC 2007 [14] and MS-COCO [28].

Box Noise Simulation. We simulate noisy bounding boxes by perturbing clean boxes. Specifically, let (cx, cy, w, h) denote the center x coordinate, center y

Table 1: Performance comparison on the PASCAL VOC 2007 test set. The evaluation metric is mAP@0.5 (%). The best performance is in **boldface**.

Model	Method	Box Noise Level			
		10%	20%	30%	40%
FasterRCNN	Noisy-FasterRCNN	76.3	71.2	60.1	42.5
	KL loss [20]	75.8	72.7	64.6	48.6
	Co-teaching [18]	75.4	70.6	60.9	43.7
	SD-LocNet [48]	75.7	71.5	60.8	43.9
	Ours	77.4	74.3	70.6	63.8
RetinaNet	Noisy-RetinaNet	71.5	67.5	57.9	45.0
	FreeAnchor [49]	73.0	67.5	56.2	41.6
	Ours	73.1	69.1	62.9	53.4
Clean Model	Clean-FasterRCNN	77.2	77.2	77.2	77.2
	Clean-RetinaNet	73.5	73.5	73.5	73.5

coordinate, width, and height of an object. We simulate an inaccurate bounding box by randomly shifting and scaling the box as follows:

$$\begin{cases} \hat{c}x = cx + \Delta_x \cdot w, & \hat{c}y = cy + \Delta_y \cdot h, \\ \hat{w} = (1 + \Delta_w) \cdot w, & \hat{h} = (1 + \Delta_h) \cdot h, \end{cases} \quad (12)$$

where Δ_x , Δ_y , Δ_w , and Δ_h follow the uniform distribution $U(-r, r)$, r is the box noise level. We simulate various box noise levels ranging from 10% to 40%. For example, when $r = 40\%$, Δ_x , Δ_y , Δ_w , and Δ_h are in the range of $(-0.4, 0.4)$. Note that Eq. (12) is performed on **every bounding box** in the training data. Fig. 4a shows examples of the synthetic inaccurate bounding boxes under different box noise level r 's on the VOC dataset, where r ranges from 10% to 40%.

Real Noisy Dataset. We also evaluate our approach on the Global Wheat Head Detection (GWHD) dataset [8,9]. This dataset includes 3.6K training images, 1.4K validation images, and 1.3K test images. It has two versions of training data, the first “noisy” challenge version¹ (inaccurate bounding box annotations) and the second “clean” version (calibrated clean annotations). Specifically, the “noisy” version contains around 20% noisy ground-truth boxes and the rest 80% boxes are the same as the “clean” version. We separately train on the “noisy” and “clean” training data to validate our approach. Fig. 4b shows some examples of the real inaccurate bounding boxes and the calibrated clean boxes.

Evaluation Metric. For VOC and COCO, we use mean average precision (mAP@.5) and mAP@[.5,.95] as evaluation metrics. Regarding GWHD dataset, we follow the GWHD Challenge 2021² to use Average Domain Accuracy (ADA) as the evaluation metric.

² <https://www.aicrowd.com/challenges/global-wheat-challenge-2021>

Table 2: Performance comparison on the MS-COCO dataset.

Method	20% Box Noise Level						40% Box Noise Level					
	AP	AP^{50}	AP^{75}	AP^S	AP^M	AP^L	AP	AP^{50}	AP^{75}	AP^S	AP^M	AP^L
FasterRCNN												
Clean-FasterRCNN	37.9	58.1	40.9	21.6	41.6	48.7	37.9	58.1	40.9	21.6	41.6	48.7
Noisy-FasterRCNN	30.4	54.3	31.4	17.4	33.9	38.7	10.3	28.9	3.3	5.7	11.8	15.1
KL loss [20]	31.0	54.3	32.4	18.0	34.9	39.5	12.1	36.7	3.7	6.2	13.0	17.4
Co-teaching [18]	30.5	54.9	30.5	17.3	34.0	39.1	11.5	31.4	4.2	6.4	13.1	16.4
SD-LocNet [48]	30.0	54.5	30.3	17.5	33.6	38.7	11.3	30.3	4.3	6.0	12.7	16.6
Ours	32.1	55.3	33.2	18.1	35.8	41.6	18.6	42.6	12.9	9.2	19.9	26.5
RetinaNet												
Clean-RetinaNet	36.7	56.1	39.0	21.6	40.4	47.4	36.7	56.1	39.0	21.6	40.4	47.4
Noisy-RetinaNet	30.0	53.1	30.8	17.9	33.7	38.2	13.3	33.6	5.7	8.4	15.9	18.0
FreeAnchor [49]	28.6	53.1	28.5	16.6	32.2	37.0	10.4	28.9	3.3	5.8	12.1	14.9
Ours	30.9	54.0	32.3	18.5	34.9	39.6	19.2	45.2	12.0	11.3	23.0	24.9

4.2 Comparison With State of the Art

We compare our method with several state-of-the-art approaches [18,20,48] on PASCAL VOC 2007 [14], MS-COCO [28], and GWHD [8,9] datasets. Note that, we denote Clean-FasterRCNN and Noisy-FasterRCNN as FasterRCNN models trained under clean and noisy training data with the default setting, respectively. Similarly, Clean-RetinaNet and Noisy-RetinaNet denote RetinaNet models trained under clean and noisy training data, respectively.

Results on the VOC 2007 dataset. Table 1 shows the comparison results on the VOC 2007 test set. For FasterRCNN, we observe that inaccurate bounding box annotations significantly deteriorate the detection performance of the vanilla model. On the contrary, our approach is more robust to noisy bounding boxes and outperforms other methods by a large margin under high box noise levels, *e.g.*, 30% and 40% box noise. In addition, Co-teaching and SD-LocNet only slightly improve the detection performance, which indicates that small-loss sample selection and sample weight assignment can not well tackle noisy box annotations. For RetinaNet, we compare our approach with the vanilla RetinaNet model and FreeAnchor [49]. As shown in Table 1, our approach still achieves consistent improvement over the vanilla model, which indicates that our approach is effective on both two-stage and one-stage detectors.

Results on the MS-COCO Dataset. The comparison results on the MS-COCO dataset are reported in Table 2. For FasterRCNN, our approach achieves considerable improvements over the vanilla model and performs favorably against state-of-the-art methods. For example, under 40% box noise, the vanilla model suffers from catastrophic performance drop, *e.g.*, AP^{50} drops from 58.1 to 28.9.

Table 3: Comparison results on the GWHD validation and test set, where models are trained on “noisy” and “clean” training data, respectively. The evaluation metric is ADA. FRCNN denotes FasterRCNN.

Model	Method	Trained on “Noisy” GWHD		Trained on “Clean” GWHD	
		Val ADA	Test ADA	Val ADA	Test ADA
FRCNN	Vanilla FRCNN	0.608	0.509	0.632	0.511
	KL Loss [20]	0.607	0.496	0.631	0.507
	Co-teaching [18]	0.624	0.491	0.631	0.504
	SD-LocNet [48]	0.621	0.498	0.626	0.512
	Ours	0.639	0.526	0.658	0.530
RetinaNet	Vanilla RetinaNet	0.607	0.494	0.622	0.503
	FreeAnchor [49]	0.619	0.517	0.635	0.525
	Ours	0.621	0.516	0.640	0.527

On the other hand, our approach significantly boosts the detection performance across all metrics, achieving 8.3%, 13.7%, and 9.6% improvements on AP , AP^{50} , and AP^{75} , respectively. Co-teaching and SD-LocNet, however, still can not well address inaccurate bounding box annotations, but KL Loss slightly improves the performance under 20% and 40% box noise. In addition, we observe that objects with different sizes suffer similarly under different noise levels. For RetinaNet, our approach also obtains consistent improvements. For example, our approach improves the performance of the vanilla RetinaNet by 5.9%, 11.6%, and 6.3% on AP , AP^{50} , and AP^{75} under 40% box noise, respectively.

Results on the GWHD Dataset. Here we report the results on both noisy and clean training data.

Results on the “Noisy” GWHD dataset. Table 3 shows the comparison results. Deploying our method on FasterRCNN boosts the Val ADA of the vanilla model from 0.608 to 0.639 and Test ADA from 0.509 to 0.526. Interestingly, our approach even performs better than the vanilla model that trained on clean training data (0.639 vs. 0.632 on Val ADA, 0.526 vs. 0.511 on Test ADA). In addition, Co-teaching and SD-LocNet improve the Val ADA but deteriorate the Test ADA, we infer that the large domain gap between validation and test data leads to the controversy. For RetinaNet, our approach obtains moderate improvements and performs favorably against FreeAnchor. Note that FreeAnchor performs well on “Noisy” GWHD because the clean ground-truth boxes dominate this dataset (around 80% boxes are clean), which is different from the synthetic VOC and COCO datasets where noisy ground-truth boxes are the majority.

Results on the “Clean” GWHD dataset. Table 3 shows that our approach can further improve the detection performance when trained on clean data. The reason may be that our OA-MIL exploits the information between object instances, thus strengthening the discriminative capability of the detection model. Specifically, we advance the performance of the vanilla FasterRCNN by 2.6% and 1.7%

Table 4: Ablation study on the VOC 2007 test set and COCO validation set. The evaluation metric is mAP@0.5 (%).

No.	Method	IS Loss	OA-IS	OA-IE	VOC 2007				COCO	
					10%	20%	30%	40%	20%	40%
B1	Vanilla FasterRCNN				76.3	71.2	60.1	42.5	54.3	28.9
B2	OA-MIL FasterRCNN	✓			77.1	73.3	66.9	56.0	54.6	32.6
B3		✓	✓		77.2	74.2	70.2	63.3	55.2	39.8
B4		✓	✓	✓	77.4	74.3	70.6	63.8	55.3	42.6

on Val ADA and Test ADA, respectively. Regarding RetinaNet, our approach outperforms the vanilla model by 1.8% on Val ADA and 2.4% on Test ADA, respectively. In addition, our approach can also cooperate with the runner-up solution [29,30] of the GWHD2021 challenge², which adopts the idea of dynamic network [31] to improve wheat head detection.

4.3 Ablation Study

Here we investigate the effectiveness of each component in our approach, including: (i) our object bag formulation, *i.e.*, training object detector with instance selection loss (IS Loss), where the loss is computed based on object bag; (ii) object-aware instance selection (OA-IS); (iii) object-aware instance extension (OA-IE). Table 4 shows the results. B1 is the performance of the vanilla FasterRCNN trained under different box noise levels. From B2 to B4, we gradually add IS Loss, OA-IS, and OA-IE into training. In addition, the analysis of parameter sensitivity (*e.g.*, γ and θ in Eq. (5)) can be found in the supplementary.

Effectiveness of Objec Bag Formulation. Interestingly, simply training under our object bag formulation significantly boosts the mAP performance of FasterRCNN on the VOC 2007 dataset across several box noise levels. For instance, our object bag formulation achieves 6.8% and 13.5% improvements under 30% and 40% box noise level, respectively. As for the COCO dataset, we still obtains moderate improvements. An intuitive explanation is that the instance selector is forced to select high-quality instances (*e.g.*, instance that covers the actual object more tightly) to minimize the loss function. As a consequence, the object detector benefits from the joint optimization process.

Effectiveness of OA-IS. Applying OA-IS further improves the detection performance on the VOC and COCO datasets, especially under high box noise levels. For example, under 40% box noise level, OA-IS boosts the performance from 56.0 to 63.3 on the VOC dataset and from 32.6 to 39.8 on the COCO dataset. To understand OA-IS more intuitively, we visualize the instances selected by OA-IS in Fig. 5. It is clear that the selected instances cover the objects more tightly than the noisy ground-truth boxes. Although the selected instances are not perfect, they provide more precise supervision signals for training the instance classifier and the instance generator.

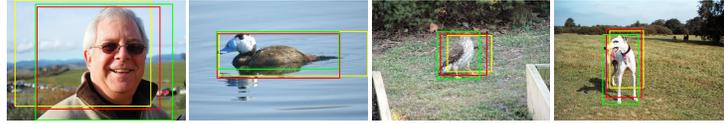


Fig. 5: **Examples of the selected instances (red boxes)**. Noisy ground-truth boxes are in yellow and the clean ground-truth boxes are in green.



Fig. 6: (a) Qualitative results of OA-MIL FasterRCNN (red boxes) and vanilla FasterRCNN (yellow boxes) on the COCO dataset. The ground-truth boxes are in green. (b) Failure cases, *e.g.*, missing detections on small/overlapped objects.

Effectiveness of OA-IE. OA-IE is designed to improve the quality of the bag instances. We observe that the impact of OA-IE is minor under low box noise levels. The reason is likely that the quality of bag instances is relatively high under low box noise situations. Nevertheless, OA-IE still brings improvement under high noise levels. For example, it improves the detection performance from 39.8 to 42.6 on the COCO dataset under 40% box noise.

Qualitative Results. Fig. 6a illustrates the qualitative results of the COCO dataset. The vanilla FasterRCNN tends to predict bounding boxes that cover object parts or include background areas. Instead, our method can predict more accurate bounding boxes. In addition, some failure cases are shown in Fig. 6b. Our approach may suffer from overlapped objects or small objects.

5 Conclusion

In this work, we tackle learning robust object detectors with inaccurate bounding boxes. By treating an object as a bag of instances, we present an Object-Aware Multiple Instance Learning method featured with object-aware instance selection and object-aware instance extension. Our approach is general and can easily cooperate with modern object detectors. Extensive experiments on the synthetic noisy datasets and real noisy GWHD dataset demonstrate that OA-MIL can obtain promising results with inaccurate bounding box annotations.

For future work, we plan to incorporate the attributes of the objects to address the limitation of OA-MIL.

Acknowledgement. This work was supported by the National Natural Science Foundation of China under Grant No. 61876211, No. U1913602, and No. 62106080.

References

1. Bernhard, M., Schubert, M.: Correcting imprecise object locations for training object detectors in remote sensing applications. *Remote Sensing* **13**(24) (2021)
2. Bilen, H., Pedersoli, M., Tuytelaars, T.: Weakly supervised object detection with convex clustering. In: *CVPR*. pp. 1081–1089 (2015)
3. Bilen, H., Vedaldi, A.: Weakly supervised deep detection networks. In: *CVPR*. pp. 2846–2854 (2016)
4. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. *CVPR* pp. 6154–6162 (2018)
5. Chadwick, S., Newman, P.: Training object detectors with noisy data. In: *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*. pp. 1319–1325 (2019)
6. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: MMDetection: Open mmlab detection toolbox and benchmark. *arXiv* (2019)
7. Cinbis, R.G., Verbeek, J., Schmid, C.: Multi-fold mil training for weakly supervised object localization. In: *CVPR*. pp. 2409–2416 (2014)
8. David, E., Madec, S., Sadeghi-Tehran, P., Aasen, H., Zheng, B., Liu, S., Kirchgessner, N., Ishikawa, G., Nagasawa, K., Badhon, M.A., Pozniak, C., de Solan, B., Hund, A., Chapman, S.C., Baret, F., Stavness, I., Guo, W.: Global wheat head detection (gwhd) dataset: A large and diverse dataset of high-resolution rgb-labelled images to develop and benchmark wheat head detection methods. *Plant Phenomics* **2020** (Aug 2020)
9. David, E., Serouart, M., Smith, D., Madec, S., Velumani, K., Liu, S., Wang, X., Pinto, F., Shafiee, S., Tahir, I.S.A., Tsujimoto, H., Nasuda, S., Zheng, B., Kirchgessner, N., Aasen, H., Hund, A., Sadhegi-Tehran, P., Nagasawa, K., Ishikawa, G., Dandrifosse, S., Carlier, A., Dumont, B., Mercatoris, B., Evers, B., Kuroki, K., Wang, H., Ishii, M., Badhon, M.A., Pozniak, C., LeBauer, D.S., Lillemo, M., Poland, J., Chapman, S., de Solan, B., Baret, F., Stavness, I., Guo, W.: Global wheat head detection 2021: An improved dataset for benchmarking wheat head detection methods. *Plant Phenomics* **2021** (Sep 2021)
10. Deng, J., Dong, W., Socher, R., Li, L., Kai Li, Li Fei-Fei: Imagenet: A large-scale hierarchical image database. In: *CVPR*. pp. 248–255 (2009)
11. Deselaers, T., Alexe, B., Ferrari, V.: Localizing objects while learning their appearance. In: *ECCV*. pp. 452–466 (2010)
12. Diba, A., Sharma, V., Pazandeh, A., Pirsiavash, H., Van Gool, L.: Weakly supervised cascaded convolutional networks. In: *CVPR*. pp. 5131–5139 (2017)
13. Dietterich, T., Lathrop, R., Lozano-Pérez, T.: Solving the multiple instance problem with axis-parallel rectangles. *Artificial Intelligence* **89**, 31–71 (03 1997)
14. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *IJCV* **88**(2), 303–338 (Jun 2010)
15. Gao, J., Wang, J., Dai, S., Li, L., Nevatia, R.: Note-rcnn: Noise tolerant ensemble rcnn for semi-supervised object detection. In: *CVPR*. pp. 9507–9516 (2019)
16. Ghosh, A., Kumar, H., Sastry, P.S.: Robust loss functions under label noise for deep neural networks. In: *AAAI*. p. 1919–1925 (2017)
17. Girshick, R., Radosavovic, I., Gkioxari, G., Dollár, P., He, K.: Detectron. <https://github.com/facebookresearch/detectron> (2018)
18. Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I.W., Sugiyama, M.: Co-teaching: Robust training of deep neural networks with extremely noisy labels. In: *NeurIPS*. p. 8536–8546 (2018)

19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
20. He, Y., Zhu, C., Wang, J., Savvides, M., Zhang, X.: Bounding box regression with uncertainty for accurate object detection. In: CVPR. pp. 2883–2892 (2019)
21. Jiang, L., Zhou, Z., Leung, T., Li, L., Fei-Fei, L.: Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In: ICML. pp. 2309–2318 (2018)
22. Kantorov, V., Oquab, M., Cho, M., Laptev, I.: Contextlocnet: Context-aware deep network models for weakly supervised localization. In: ECCV. pp. 350–365 (2016)
23. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., Duerig, T., Ferrari, V.: The open images dataset v4. IJCV **128**(7), 1956–1981 (Jul 2020)
24. Li, D., Huang, J., Li, Y., Wang, S., Yang, M.: Weakly supervised object localization with progressive domain adaptation. In: CVPR. pp. 3512–3520 (2016)
25. Li, J., Xiong, C., Socher, R., Hoi, S.C.H.: Towards noise-resistant object detection with noisy annotations. arXiv [abs/2003.01285](https://arxiv.org/abs/2003.01285) (2020)
26. Lin, T., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR. pp. 936–944 (2017)
27. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2999–3007 (2017)
28. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755 (2014)
29. Liu, C., Wang, K., Lu, H., Cao, Z.: Dynamic color transform for wheat head detection. In: ICCVW. pp. 1278–1283 (2021)
30. Liu, C., Wang, K., Lu, H., Cao, Z.: Dynamic color transform networks for wheat head detection. *Plant Phenomics* **2022** (Feb 2022)
31. Lu, H., Dai, Y., Shen, C., Xu, S.: Index networks. *IEEE TPAMI* **44**(1), 242–255 (2022)
32. Ma, X., Wang, Y., Houle, M.E., Zhou, S., Erfani, S., Xia, S., Wijewickrema, S., Bailey, J.: Dimensionality-driven learning with noisy labels. In: ICML. pp. 3355–3364 (2018)
33. Mao, J., Yu, Q., Aizawa, K.: Noisy localization annotation refinement for object detection. In: ICIP. pp. 2006–2010 (2020)
34. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE TPAMI* **39**(6), 1137–1149 (2017)
35. Siva, P., Tao Xiang: Weakly supervised object detector learning with model drift detection. In: ICCV. pp. 343–350 (2011)
36. Siva, P., Russell, C., Xiang, T.: In defence of negative mining for annotating weakly labelled data. In: ECCV. p. 594–608 (2012)
37. Song, H., Kim, M., Lee, J.G.: SELFIE: Refurbishing unclean samples for robust deep learning. In: ICML. pp. 5907–5915 (2019)
38. Song, H.O., Lee, Y.J., Jegelka, S., Darrell, T.: Weakly-supervised discovery of visual pattern configurations. In: NeurIPS. p. 1637–1645 (2014)
39. Tang, P., Wang, X., Bai, X., Liu, W.: Multiple instance detection network with online instance classifier refinement. In: CVPR. pp. 3059–3067 (2017)
40. Tang, Y., Wang, J., Gao, B., Dellandréa, E., Gaizauskas, R., Chen, L.: Large scale semi-supervised object detection using visual and semantic knowledge transfer. In: CVPR. pp. 2119–2128 (2016)

41. Tang, Y., Wang, J., Wang, X., Gao, B., Dellandréa, E., Gaizauskas, R., Chen, L.: Visual and semantic knowledge transfer for large scale semi-supervised object detection. *IEEE TPAMI* **40**(12), 3045–3058 (2018)
42. Uijlings, J.R.R., Popov, S., Ferrari, V.: Revisiting knowledge transfer for training object class detectors. In: *CVPR*. pp. 1101–1110 (2018)
43. Uijlings, J.R.R., van de Sande, K.E.A., Gevers, T., Smeulders, A.W.M.: Selective search for object recognition. *IJCV* **104**(2), 154–171 (Sep 2013)
44. Wan, F., Wei, P., Jiao, J., Han, Z., Ye, Q.: Min-entropy latent model for weakly supervised object detection. In: *CVPR*. pp. 1297–1306 (2018)
45. Wan, F., Liu, C., Ke, W., Ji, X., Jiao, J., Ye, Q.: C-MIL: continuation multiple instance learning for weakly supervised object detection. In: *CVPR*. pp. 2199–2208 (2019)
46. Wei, Y., Shen, Z., Cheng, B., Shi, H., Xiong, J., Feng, J., Huang, T.: Ts2c: Tight box mining with surrounding segmentation context for weakly supervised object detection. In: *ECCV*. pp. 454–470 (2018)
47. Xu, Y., Zhu, L., Yang, Y., Wu, F.: Training robust object detectors from noisy category labels and imprecise bounding boxes. *IEEE TIP* **30**, 5782–5792 (2021)
48. Zhang, X., Yang, Y., Feng, J.: Learning to localize objects with noisy labeled instances. In: *AAAI*. pp. 9219–9226 (2019)
49. Zhang, X., Wan, F., Liu, C., Ji, R., Ye, Q.: Freeanchor: Learning to match anchors for visual object detection. In: *NeurIPS*. pp. 147–155 (2019)
50. Zhang, Z., Sabuncu, M.R.: Generalized cross entropy loss for training deep neural networks with noisy labels. In: *NeurIPS*. p. 8792–8802 (2018)