

# IoU-aware Single-stage Object Detector for Accurate Localization

Shengkai Wu<sup>a</sup>, Xiaoping Li<sup>a,\*</sup>, Xinggang Wang<sup>b</sup>

<sup>a</sup>*State Key Laboratory of Digital Manufacturing Equipment and Technology, Huazhong University of Science and Technology, Wuhan, 430074, China.*

<sup>b</sup>*School of EIC, Huazhong University of Science and Technology, Wuhan, 430074, China.*

---

## Abstract

Due to the simpleness and high efficiency, single-stage object detectors have been widely applied in many computer vision applications. However, the low correlation between the classification score and localization accuracy of the predicted detections has severely hurt the localization accuracy of models. In this paper, IoU-aware single-stage object detector is proposed to solve this problem. Specifically, IoU-aware single-stage object detector predicts the IoU between the regressed box and the ground truth box. Then the classification score and predicted IoU are multiplied to compute the detection confidence, which is highly correlated with the localization accuracy. The detection confidence is then used as the input of NMS and COCO AP computation, which will substantially improve the localization accuracy of models. Sufficient experiments on COCO and PASCAL VOC dataset demonstrate the effectiveness of IoU-aware single-stage object detector on improving the localization accuracy. Without whistles and bells, the proposed method can substantially improve AP by 1.0% ~ 1.6% on COCO *test-dev* and 1.1% ~ 2.2% on PASCAL VOC2007 test compared with the baseline. The improvement for AP at higher IoU threshold(0.7 ~ 0.9) is 1.7% ~ 2.3% on COCO *test-dev* and 1.0% ~ 4.2% PASCAL VOC2007 test. The source code will be made publicly available.

*Keywords:* IoU prediction, IoU-aware detector, Accurate localization, Correlation between classification score and localization accuracy

---

\*Corresponding author

*Email addresses:* ShengkaiWu@hust.edu.cn (Shengkai Wu),  
lixiaoping@hust.edu.cn (Xiaoping Li), xgwang@hust.edu.cn (Xinggang Wang)

---

## 1. Introduction

As the development of deep convolutional neural networks, a large amount of object detection models have been proposed in recent years. Most of these models can be classified into single-stage object detectors [1, 2, 3, 4, 5, 6] and multi-stage object detectors [7, 8, 9, 10, 11, 12, 13]. For the multi-stage object detectors, multi-stage classification and localization are applied sequentially, which make these models more powerful on classification and localization tasks. Compared with single-stage object detectors, the multi-stage object detectors have achieved better average precision(AP), but the efficiency is hurt by the multi-stage classification and localization subnetworks. On the contrary, the single-stage detectors rely on a fully convolutional networks(FCN) for classification and localization, which is more simple and efficient. However, the AP of single-stage detectors generally lag behind that of the multi-stage detectors.

In this work, we aim to improve the AP of single-stage detectors while keeping their efficiency. We demonstrate that the low correlation between the classification score and localization accuracy of single-stage detectors have severely hurt the localization accuracy of the models. The low correlation is mostly caused by that the classification and localization subnetworks are trained with independent objective functions without knowing each other explicitly. After the models are converged, the classification subnetwork will predict classification score for each regressed anchor without knowing the localization accuracy, represented by IoU between the regressed anchor and the ground truth box. Thus, there will be many detections having the mismatch problem between the classification scores and their localization accuracy, such as detections with high classification scores but low IoU, detections with low classification scores but high IoU. These detections will hurt the localization accuracy of models in two ways during inference. Firstly, during standard non-maximum suppression(NMS), all the detections are ranked based on their classification scores and the detection with higher classification score will suppress other detections that have an overlap with it higher than a threshold. Consequently, the detections with low classification scores but high IoU will be suppressed by the detections with high classification scores but low IoU. Secondly, during computing average precision(AP), all the detections are also ranked based on their classification scores. To compute

the average precision, the precisions and recalls are computed based on these ranked detections and if the detections with high classification scores but low IoU rank before the detections with low classification scores but high IoU, the precision at high IoU threshold will be reduced, which results in lower AP at high IoU threshold. Both of these problems will hurt the localization accuracy of models.

To solve the above problem, we propose IoU-aware single-stage object detector based on RetinaNet [3]. A IoU prediction head parallel with the regression head is attached to the last layer of the regression branch to predict the IoU of each regressed anchor. During training, the IoU prediction head is trained jointly with the classification head and localization head. During inference, the detection confidence is computed by multiplying the classification score and predicted IoU for each detected box and then used to rank all the detections in the subsequent NMS and AP computation. Because the detection confidence is highly correlated with the localization accuracy, the problem mentioned above can be relieved and thus the localization accuracy of models can be substantially improved as the experiments show.

The rest of this paper is organized as follows. Section 2 introduces the related research work. Section 3 introduces the IoU-aware single-stage object detector in details. Section 4 presents extensive experiments on COCO and PASCAL VOC dataset to demonstrate the effectiveness of our method. Section 5 gives the conclusions.

## 2. Related Work

**Correlation between classification score and localization accuracy.** The low correlation between the classification score and localization accuracy hurts the models' localization accuracy severely and many methods have been proposed to solve this problem. Fitness NMS [14] improves DeNet [15] by dividing the localization accuracy into 5 levels and transforming the localization accuracy prediction task to the classification task. During inference, the fitness for each detected box is computed as the weighted sum of the predicted fitness probabilities and then multiplied with the classification score as the final detection score which is more correlated with the localization accuracy. Then this score is used as the input of NMS, denoted as Fitness NMS, to improve the localization accuracy of DeNet. IoU-Net [16] improves Faster R-CNN [7] by designing a IoU prediction head parallel with the R-CNN to predict the regressed IoU for each RoI. During inference, all

the detected boxes are ranked based on the predicted IoU and then IoU-guided NMS is applied to improve the localization accuracy. Similarly, MS R-CNN [17] improves Mask R-CNN [9] by attaching a MaskIoU head parallel with the Mask head to predict the IoU between the predicted mask and the corresponding ground truth mask. During inference, the predicted IoU is multiplied with the classification score as the final mask confidence and then used to rank the predicted mask when computing AP. All the above methods design additional subnetworks to predict the localization accuracy and are applied to multi-stage detectors. There also exists other research solving the problem by designing better loss functions without changing the models' architecture. PISA [18] assigns different weights to the positive examples in the classification loss based on their importance which is obtained by IoU Hierarchical Local Rank (IoU-HLR). In addition, the classification probabilities are used to reweight the contribution of each positive example to the regression loss, denoted as classification-aware regression loss. Both the improvement to the classification and regression loss can enhance the correlation between the classification score and localization accuracy. Similarly, IoU-balanced classification loss [19] uses the regressed IoU to reweight the classification loss for each positive example directly and aims to make the examples with higher IoU learn higher classification score, which thus enhances the correlation between classification score and localization accuracy. IoU-aware single-stage object detector aims to improve the single-stage detectors by designing a IoU prediction head to predict the IoU for each regressed anchor.

**Accurate object localization.** Accurate object localization is extremely challenging in complex scene such as COCO dataset and a large number of methods have been proposed to improve the localization accuracy of object detection models in recent years. Multi-region detector [20] finds that a single-stage regression is limited for accurate localization and thus a iterative bounding box regression procedure is proposed to refine the coordinates of detected boxes, followed by NMS and box voting. Cascade R-CNN [8] proposes a multi-stage object detection architecture which trains a sequence of R-CNN with increasing IoU thresholds. Thus the trained sequential R-CNN will be sequentially more powerful for accurate localization during inference. RefineDet [4] improves the localization accuracy of the single-stage detector by using two-step bounding regression. The anchor refinement module(ARM) firstly refines the human-designed anchors to improve the localization accuracy of human-designed anchors, then the object detection module(ODM)

uses these more accurate anchors for the second step bounding box regression, thus improving the localization accuracy of the final detections. Libra R-CNN [21] designs balanced L1 loss to promote the regression gradients from inliers (accurate samples) during training. Thus, the trained regression branch will be more powerful for accurate localization. Similarly, IoU-balanced localization loss [19] reweights the localization loss for each positive example based on their regressed IoU. This reweighting procedure can down-weight the gradients from outliers and up-weight the gradients from inliers, thus improving the localization accuracy of models. Differently, IoU-aware single-stage object detector improves the localization accuracy by suppressing the detections of low localization accuracy based on the computed detection confidence during NMS and AP computation.

**Anchor-free single-stage object detectors.** To overcome the drawbacks of anchor-based detectors, anchor-free single-stage object detectors have become more and more popular recently. FCOS [22] solves object detection in a per-pixel prediction fashion based on a fully convolutional neural networks. FCOS consists of three prediction head: classification head used for classification, regression head used for localization, centerness head used for predicting the centerness of each detected box. During inference, the predicted centerness of each detected box is multiplied with the corresponding classification score as the final score, which is used in the subsequent NMS and AP computation to suppress the poorly localized detections. PolarMask [23] modifies FCOS to realize the instance segmentation. Similarly, centerness head is also used to suppress the segmentations of low localization accuracy and improve the localization accuracy of models. IoU-aware single-stage object detector designs a IoU prediction head parallel with the regression head to predict the IoU of each detection and the predicted IoU can be used to suppress the poorly localized detections. Differently, IoU-aware single-stage object detector is a anchor-based detector and the IoU of each detected box is predicted.

### 3. Method

In this section, we will introduce the model architecture of IoU-aware single-stage object detector and different designing choices in details.

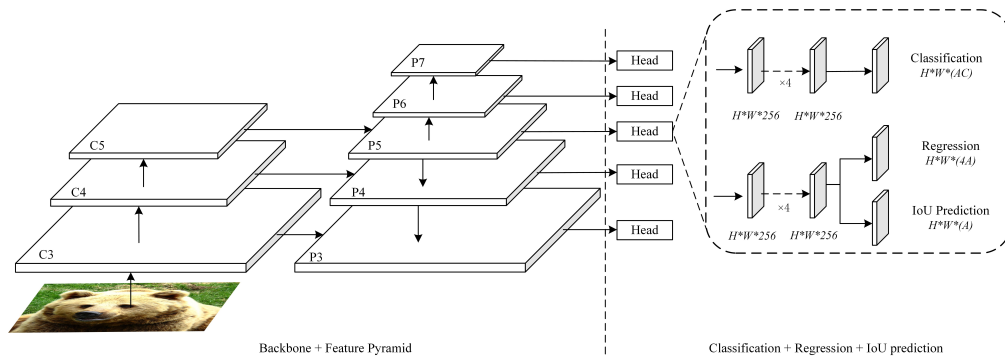


Figure 1: The model architecture of IoU-aware single-stage object detector. The same backbone and feature pyramid network(FPN) are adopted as RetinaNet. A IoU prediction head is designed parallel with the regression head at the last layer of regression branch to predict the IoU between the regressed anchors and ground truth boxes. The classification head, regression head and IoU prediction head all consist of only a single 3\*3 convolution layer.

### 3.1. IoU-aware single-stage object detector

**Model architecture.** IoU-aware single-stage object detector is mostly based on RetinaNet [3] and the same backbone and feature pyramid network(FPN) are adopted as Fig.1 shows. Different from the RetinaNet, we design a IoU prediction head parallel with the regression head in the last layer of regression branch to predict the IoU between each regressed anchor and the ground truth box while the classification branch is kept the same. To keep the model’s efficiency, the IoU prediction head consists of only a single 3\*3 convolution layer, followed by a sigmoid activation layer, ensuring the predicted IoU is in the range of [0, 1]. There are also many other choices about the design of IoU prediction head, such as designing an independent IoU prediction branch as the same as the classification branch and regression branch, but this kind of design will severely hurt the model’s efficiency. Our design brings little computation burden to the whole model and can still substantially improve the model’s AP.

**Training.** As the same to RetinaNet, focal loss is adopted for the classification loss and the smooth L1 loss is adopted for the regression loss as Equ.1,2 show. Because the predicted IoU is in the range of [0,1], binary cross-entropy loss is adopted for the IoU prediction loss as Equ.3 shows. During training, the IoU prediction head is trained jointly with the classification head and regression head. Other kinds of loss functions can also be considered, such

as L2 loss and L1 loss. These different loss functions will be compared in the following experiments.

$$L_{cls} = \frac{1}{N_{Pos}} \left( \sum_{i \in Pos} FL(p_i, \hat{p}_i) + \sum_{i \in Neg} FL(p_i, \hat{p}_i) \right) \quad (1)$$

$$L_{loc} = \frac{1}{N_{Pos}} \sum_{i \in Pos} \sum_{m \in \{cx, cy, w, h\}} \text{smooth}_{L1}(l_i^m - \hat{g}_i^m) \quad (2)$$

$$L_{IoU} = \frac{1}{N_{Pos}} \sum_{i \in Pos} CE(IoU_i, \hat{IoU}_i) \quad (3)$$

$$L_{total} = L_{cls} + L_{loc} + L_{IoU} \quad (4)$$

**Inference.** At inference, the classification score and the predicted IoU for each detected box is computed based on Equ.5 as the final detection confidence. The parameter  $\alpha$  is designed to control the contribution of the classification score and predicted IoU to the final detection confidence. This detection confidence can simultaneously represent the classification confidence and localization accuracy. Thus the detection confidence is more correlated with the localization accuracy compared with the classification score. Then the detection confidence is used to rank all the detections in the subsequent NMS and AP computation. The poorly localized detections will be suppressed in this procedure.

$$S_{det} = p_i^\alpha IoU_i^{(1-\alpha)} \quad (5)$$

## 4. Experiments

### 4.1. Experimental Settings

**Dataset and Evaluation Metrics.** Most of the experiments are evaluated on the challenging MS COCO [24] dataset. It consists of 118k images for training (*train-2017*), 5k images for validation (*val-2017*) and 20k images with no disclosed labels for test (*test-dev*). There exist totally over 500k annotated object instances from 80 categories in the dataset. To demonstrate the generalization ability of our method, we also conduct experiments on the PASCAL VOC [25] dataset in the ablation studies. VOC2007 consists of 5011 images for training (*VOC2007 trainval*) and 4952 for test (*VOC2007*

*test*). And VOC2012 consists of 17125 images for training (*VOC2012 trainval*) and 5138 for test (*VOC2012 test*). For all the experiments, the standard COCO-style Average Precision (AP) metrics are adopted which consists of AP (averaged AP at IoUs from 0.5 to 0.95 with an interval of 0.05), AP<sub>50</sub> (AP at IoU threshold 0.5), AP<sub>75</sub> (AP at IoU threshold 0.75), AP<sub>S</sub> (AP for objects of small scales), AP<sub>M</sub> (AP for objects of medium scales) and AP<sub>L</sub> (AP for objects of large scales).

**Implementation Details.** All the object detection models are implemented based on PyTorch and MMDetection [26]. As only 2 GPUs are available, linear scaling rule [27] is adopted to adjust the learning rate during training. For the main results, all the models are evaluated on COCO *test-dev*. The converged models provided by MMDetection are evaluated as the baselines. With the default setting in the MMDetection, IoU-aware single-stage object detectors are all trained for total 12 epochs with the image scale of [800, 1333]. Some papers report the main results obtained by training the models for total 1.5 longer time and with scale jitter. These tricks are not adopted in our experiments. In the ablation studies, IoU-aware single-stage object detector with ResNet50 as backbone is trained on COCO *train-2017* and evaluated on COCO *val-2017* using the image scale of [600, 1000]. For the experiments on PASCAL VOC, the models with different backbones are trained on the *VOC2007 trainval* and *VOC2012 trainval* and evaluated on *VOC2007 test* with the image scale of [600, 1000]. If not specified, all the other settings are kept the same as the default settings in the MMDetection.

#### 4.2. Main Results

In the main results shown by Table 1, the performance of IoU-aware single-stage object detectors with different backbones are compared with the state-of-the-art object detection models on the COCO *test-dev*. For fair comparison, the trained models provided by MMDetection [26] with different backbones are evaluated as the baselines. As Table 1 shows, IoU-aware RetinaNets with different backbones can substantially improve AP by 1.0% ~ 1.6% compared with the baselines. The performance for AP50 is increased or decreased marginally, but the performance at AP75 is largely improved by 1.7% ~ 2.3%, which demonstrates the effectiveness of IoU-aware RetinaNet on improving the models' localization accuracy. In addition, the performance of IoU-aware RetinaNets have surpassed the two-stage detector Faster R-CNN with the same backbone by 0.3% ~ 0.7% AP and the improvement mostly comes from the high localization accuracy of IoU-aware



Model	Backbone	Schedule	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
YOLOv2 [28]	DarkNet-19	-	21.6	44.0	19.2	5.0	22.4	35.5
YOLOv3 [29]	DarkNet-53	-	33.0	57.9	34.4	18.3	35.4	41.9
SSD300 [1]	VGG16	-	23.2	41.2	23.4	5.3	23.2	39.6
SSD512 [1]	VGG16	-	26.8	46.5	27.8	9.0	28.9	41.9
Faster R-CNN [7]	ResNet-101-FPN	-	36.2	59.1	39.0	18.2	39.0	48.2
Deformable R-FCN [30]	Inception-ResNet-v2	-	37.5	58.0	40.8	19.4	40.1	52.5
Mask R-CNN [9]	ResNet-101-FPN	-	38.2	60.3	41.7	20.1	41.1	50.2
Faster R-CNN*	ResNet-50-FPN	1x	36.2	58.5	38.9	21.0	38.9	45.3
Faster R-CNN*	ResNet-101-FPN	1x	38.8	60.9	42.1	22.6	42.4	48.5
Faster R-CNN*	ResNeXt-32x8d-101-FPN	1x	40.3	62.7	44.0	24.4	43.7	49.8
RetinaNet*	ResNet-50-FPN	1x	35.9	55.8	38.4	19.9	38.8	45.0
RetinaNet*	ResNet-101-FPN	1x	38.1	58.5	40.8	21.2	41.5	48.2
RetinaNet*	ResNeXt-32x8d-101-FPN	1x	39.0	59.7	41.9	22.3	42.5	48.9
IoU-aware RetinaNet	ResNet-50-FPN	1x	36.9	56.1	40.1	20.9	40.0	46.0
IoU-aware RetinaNet	ResNet-101-FPN	1x	39.2	58.2	42.9	22.1	42.7	50.0
IoU-aware RetinaNet	ResNeXt-32x8d-101-FPN	1x	40.6	60.1	44.2	23.4	43.9	51.8

Table 1: Comparison with the state-of-the-art methods on COCO *test-dev*. The symbol “\*” means the reimplementation results in MMDetection [26]. The training schedule is the same as Detectron [31]. “1x” means the model is trained for 12 epochs. Different from other research, the longer training schedule and scale jitters are not adopted in our experiments.

## RetinaNet.

IoU prediction loss	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
baseline	34.3	54.5	36.4	17.2	38.2	47.1
L2 loss	35.1	53.8	37.9	18.6	39.2	47.6
BCE loss	35.4	54.1	38.2	18.9	39.3	48.3

Table 2: The effectiveness of training IoU-aware RetinaNet-ResNet50 with different IoU prediction losses on COCO *val-2017*.

### 4.3. Ablation Studies

**IoU Prediction Loss.** Different IoU prediction losses are used to train IoU-aware RetinaNet. To investigate the effect of IoU prediction loss only, the detection confidence is computed by multiplying the classification score and predicted IoU directly without using the parameter  $\alpha$ . As shown in Table 2, training the model with binary cross-entropy loss can produce better performance than training the model with L2 loss. This may be caused by that the predicted IoU is more accurate when training the IoU prediction head with binary cross-entropy loss. Thus binary cross-entropy loss is adopted in all the subsequent experiments.

$\alpha$	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
baseline	34.3	54.5	36.4	17.2	38.2	47.1
none	35.4	54.1	38.2	18.9	39.3	48.3
1.0	34.5	54.1	36.7	17.4	38.4	46.8
0.9	34.8	54.3	37.1	17.8	38.6	47.0
0.8	35.0	54.4	37.4	18.3	38.8	47.6
0.7	35.2	54.4	37.7	18.5	39.0	47.9
0.6	35.4	54.3	38.0	18.8	39.2	48.2
0.5	35.5	54.1	38.2	18.9	39.4	48.3
0.4	35.5	53.7	38.4	18.9	39.5	48.4
0.3	35.4	53.0	38.5	18.8	39.5	48.4
0.2	34.8	51.3	38.4	18.5	39.1	48.1
0.1	32.5	46.4	36.2	16.7	37.0	45.9
0	0.4	0.5	0.4	0.2	0.7	0.8

Table 3: The effectiveness of computing the detection confidence without using the parameter  $\alpha$  and computing the detection confidence with varying the parameter  $\alpha$  on COCO *val-2017*.

**Detection Confidence Computation.** At inference, the detection confidence is computed based on Equ. 5 and the parameter  $\alpha$  is used to control the contribution of the classification score and predicted IoU to the final detection confidence. There are several observations from the experimental results in Table 3 and Table 4. Firstly, as Table 3 shows, when  $\alpha$  equals to 1.0, only the classification score is used as the detection confidence and the AP is improved by 0.2%. This demonstrates that multi-task training with IoU prediction loss is beneficial to the model’s performance. Secondly, when  $\alpha$  equals to 0.5 and 0.4, the best performance of AP 35.5% is obtained, which is 1.2% better than the baseline. The AP50 is marginally decreased by 0.4% ~ 0.8% while the AP70 and AP80 are improved by 2.0% ~ 2.7% as Table 4 shows, demonstrating the effectiveness of our method on improving the model’s localization accuracy. Thirdly, as the parameter  $\alpha$  is decreased to improve the contribution of the predicted IoU to the detection confidence, the AP50 decreases while the AP70 and AP80 increases as Table 4 shows. This demonstrates that the predicted IoU is very correlated with the localization accuracy and can bias the model to the detections with high localization accuracy. In addition, the detection confidence can also be computed by multiplying the classification score and predicted IoU directly without using the

$\alpha$	AP	AP <sub>50</sub>	AP <sub>60</sub>	AP <sub>70</sub>	AP <sub>80</sub>	AP <sub>90</sub>
baseline	34.3	54.5	49.2	41.3	29.8	10.9
none	35.4	54.1	50.0	43.3	31.9	12.2
1.0	34.5	54.1	49.5	42.0	30.1	11.0
0.9	34.8	54.3	49.7	42.4	30.4	11.2
0.8	35.0	54.4	49.9	42.6	30.7	11.4
0.7	35.2	54.4	50.1	43.0	31.2	11.7
0.6	35.4	54.3	50.2	43.2	31.5	12.0
0.5	35.5	54.1	50.0	43.3	31.9	12.2
0.4	35.5	53.7	49.8	43.3	32.5	12.5
0.3	35.4	53.0	49.3	43.2	32.8	12.8
0.2	34.8	51.3	48.0	42.5	32.9	13.1
0.1	32.5	46.4	43.7	39.4	31.7	13.3
0	0.4	0.5	0.4	0.4	0.4	0.3

Table 4: The impact of computing the detection confidence without using the parameter  $\alpha$  and computing the detection confidence with varying the parameter  $\alpha$  on AP at different IoU threshold on COCO *val-2017*.

parameter  $\alpha$ . As Table 3 shows, multiplying the classification score and predicted IoU directly without using the parameter  $\alpha$  can improve AP by 1.1%, which is slightly inferior than computing the detection confidence with the parameter  $\alpha$ . Thus, we choose to compute the detection confidence based on Equ. 5.

Model	Backbone	AP	AP <sub>50</sub>	AP <sub>60</sub>	AP <sub>70</sub>	AP <sub>80</sub>	AP <sub>90</sub>
RetinaNet	ResNet-50-FPN	51.4	78.8	74.3	63.6	44.9	15.2
RetinaNet	ResNet-101-FPN	55.1	81.1	77.2	67.5	50.4	20.1
IoU-aware RetinaNet	ResNet-50-FPN	53.6	79.0	75.1	66.1	48.7	19.4
IoU-aware RetinaNet	ResNet-101-FPN	56.2	80.5	76.8	68.5	52.4	22.8

Table 5: Experimental results on PASCAL VOC. All the models are trained on *VOC2007 trainval* and *VOC2012 trainval* and evaluated on *VOC2007 test* with the image scale of [600, 1000]. All the other settings are adopted the same as the default settings provided in the MMDetection.

**Ablation Studies on PASCAL VOC.** As Table 5 shows, IoU-aware RetinaNet can improve AP by 1.1%  $\sim$  2.2% compared with the baselines. In addition, the improvement for AP at higher IoU threshold(0.7,0.8,0.9) is 1.0%  $\sim$  4.2%, demonstrating that our method can substantially improve

the model’s localization accuracy. The observations in the experiments of PASCAL VOC dataset is the same as that in the experiments of COCO dataset, which demonstrates our method has generalization ability to other datasets and can be applied to different application scenes.

## 5. Conclusions

In this work, we demonstrate that the low correlation between the classification score and localization accuracy of the single-stage object detector can severely hurt the localization accuracy of models. Thus, IoU-aware single-stage object detector is designed by adding a IoU prediction head at the last layer of the regression branch to predict the IoU between each regressed anchor and the ground truth box. In this way, the model will be aware of the localization accuracy of each detection. At inference, the detection confidence is computed by multiplying the classification score and predicted IoU and then used to rank all the detections in the subsequent NMS and AP computation. Extensive experiments on MS COCO dataset and PASCAL VOC dataset have shown that IoU-aware single-stage object detector can substantially improve the model’s performance, especially the localization accuracy.

## References

- [1] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, A. C. Berg, Ssd: Single shot multibox detector, in: European conference on computer vision, Springer, 2016, pp. 21–37.
- [2] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: Unified, real-time object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 779–788.
- [3] T.-Y. Lin, P. Goyal, R. Girshick, K. He, P. Dollár, Focal loss for dense object detection, in: Proceedings of the IEEE international conference on computer vision, 2017, pp. 2980–2988.
- [4] S. Zhang, L. Wen, X. Bian, Z. Lei, S. Z. Li, Single-shot refinement neural network for object detection, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 4203–4212.

- [5] Z. Zhang, S. Qiao, C. Xie, W. Shen, B. Wang, A. L. Yuille, Single-shot object detection with enriched semantics, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 5813–5821.
- [6] B. Li, Y. Liu, X. Wang, Gradient harmonized single-stage detector, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, 2019, pp. 8577–8584.
- [7] S. Ren, K. He, R. Girshick, J. Sun, Faster r-cnn: Towards real-time object detection with region proposal networks, in: Advances in neural information processing systems, 2015, pp. 91–99.
- [8] Z. Cai, N. Vasconcelos, Cascade r-cnn: Delving into high quality object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 6154–6162.
- [9] K. He, G. Gkioxari, P. Dollár, R. Girshick, Mask r-cnn, in: Proceedings of the IEEE international conference on computer vision, 2017, pp. 2961–2969.
- [10] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, S. Belongie, Feature pyramid networks for object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 2117–2125.
- [11] J. Dai, Y. Li, K. He, J. Sun, R-fcn: Object detection via region-based fully convolutional networks, in: Advances in neural information processing systems, 2016, pp. 379–387.
- [12] R. Girshick, Fast r-cnn, in: Proceedings of the IEEE international conference on computer vision, 2015, pp. 1440–1448.
- [13] R. Girshick, J. Donahue, T. Darrell, J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2014, pp. 580–587.
- [14] L. Tychsen-Smith, L. Petersson, Improving object localization with fitness nms and bounded iou loss, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 6877–6885.

- [15] L. Tychsen-Smith, L. Petersson, Denet: Scalable real-time object detection with directed sparse sampling, in: Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 428–436.
- [16] B. Jiang, R. Luo, J. Mao, T. Xiao, Y. Jiang, Acquisition of localization confidence for accurate object detection, in: Proceedings of the European Conference on Computer Vision (ECCV), 2018, pp. 784–799.
- [17] Z. Huang, L. Huang, Y. Gong, C. Huang, X. Wang, Mask scoring r-cnn, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 6409–6418.
- [18] Y. Cao, K. Chen, C. C. Loy, D. Lin, Prime sample attention in object detection, arXiv preprint arXiv:1904.04821 (2019).
- [19] S. Wu, X. Li, IoU-balanced Loss Functions for Single-stage Object Detection, arXiv e-prints (2019) arXiv:1908.05641. arXiv:1908.05641.
- [20] S. Gidaris, N. Komodakis, Object detection via a multi-region and semantic segmentation-aware cnn model, in: Proceedings of the IEEE international conference on computer vision, 2015, pp. 1134–1142.
- [21] J. Pang, K. Chen, J. Shi, H. Feng, W. Ouyang, D. Lin, Libra r-cnn: Towards balanced learning for object detection, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 821–830.
- [22] Z. Tian, C. Shen, H. Chen, T. He, FCOS: Fully Convolutional One-Stage Object Detection, arXiv e-prints (2019) arXiv:1904.01355. arXiv:1904.01355.
- [23] E. Xie, P. Sun, X. Song, W. Wang, X. Liu, D. Liang, C. Shen, P. Luo, PolarMask: Single Shot Instance Segmentation with Polar Representation, arXiv e-prints (2019) arXiv:1909.13226. arXiv:1909.13226.
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft coco: Common objects in context, in: European conference on computer vision, Springer, 2014, pp. 740–755.

- [25] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, A. Zisserman, The pascal visual object classes (voc) challenge, *International journal of computer vision* 88 (2010) 303–338.
- [26] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, et al., Mmdetection: Open mmlab detection toolbox and benchmark, *arXiv preprint arXiv:1906.07155* (2019).
- [27] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, K. He, Accurate, large minibatch sgd: Training imagenet in 1 hour, *arXiv preprint arXiv:1706.02677* (2017).
- [28] J. Redmon, A. Farhadi, Yolo9000: better, faster, stronger, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7263–7271.
- [29] J. Redmon, A. Farhadi, Yolov3: An incremental improvement, *arXiv preprint arXiv:1804.02767* (2018).
- [30] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, Y. Wei, Deformable convolutional networks, in: *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 764–773.
- [31] R. Girshick, I. Radosavovic, G. Gkioxari, P. Dollár, K. He, Detectron, <https://github.com/facebookresearch/detectron>, 2018.