

# End-to-End Object Detection with Fully Convolutional Network

Jianfeng Wang<sup>1\*</sup> Lin Song<sup>2\*†</sup> Zeming Li<sup>1</sup> Hongbin Sun<sup>2</sup> Jian Sun<sup>1</sup> Nanning Zheng<sup>2</sup>  
<sup>1</sup>Megvii Technology <sup>2</sup>College of Artificial Intelligence, Xi'an Jiaotong University

wangjianfeng@megvii.com

stevengrove@stu.xjtu.edu.cn

{lizeming, sunjian}@megvii.com

{hsun, nnzheng}@mail.xjtu.edu.cn

## Abstract

Mainstream object detectors based on the fully convolutional network has achieved impressive performance. While most of them still need a hand-designed non-maximum suppression (NMS) post-processing, which impedes fully end-to-end training. In this paper, we give the analysis of discarding NMS, where the results reveal that a proper label assignment plays a crucial role. To this end, for fully convolutional detectors, we introduce a Prediction-aware One-To-One (POTO) label assignment for classification to enable end-to-end detection, which obtains comparable performance with NMS. Besides, a simple 3D Max Filtering (3DMF) is proposed to utilize the multi-scale features and improve the discriminability of convolutions in the local region. With these techniques, our end-to-end framework achieves competitive performance against many state-of-the-art detectors with NMS on COCO and CrowdHuman datasets. The code is available at <https://github.com/Megvii-BaseDetection/DeFCN>.

## 1. Introduction

Object detection is a fundamental topic in computer vision, which predicts a set of bounding boxes with pre-defined category labels for each image. Most of mainstream detectors [9, 22, 30, 54] utilize some hand-crafted designs such as anchor-based label assignment and non-maximum suppression (NMS). Recently, a quite number of methods [46, 58, 5] have been proposed to eliminate the pre-defined set of anchor boxes by using distance-aware and distribution-based label assignments. Although they achieve remarkable progress and superior performance, there is still a challenge of discarding the NMS post-processing, which hinders the fully end-to-end training.

To tackle this issue, Learnable NMS [13], Soft NMS [1] and other NMS variants [12, 24, 15], and CenterNet [5] are




Figure 1. As shown in the dashed box, most detectors based on the fully convolutional network adopt multiple predictions and NMS post-processing for each instance. With the proposed prediction-aware one-to-one label assignment and 3D Max Filtering, our end-to-end detector can directly perform a single prediction for each instance without post-processing.

proposed to improve the duplicate removal, but they still do not provide an effective end-to-end training strategy. Meanwhile, many approaches [43, 35, 26, 32, 36] based on recurrent neural networks have been introduced to predict the bounding box for each instance by using an autoregressive decoder. These approaches give naturally sequential modeling for the prediction of bounding boxes. But they are only evaluated on some small datasets without modern detectors, and the iterative manner makes the inference process inefficient.

Recently, DETR [3] introduces a bipartite matching based training strategy and transformers with the parallel decoder to enable end-to-end detection. It achieves competitive performance against many state-of-the-art detectors. However, DETR currently suffers from much longer training duration to coverage and relatively lower performance on the small objects. To this end, this paper explores a new perspective: *could a fully convolutional network achieve competitive end-to-end object detection?*

In this paper, we attempt to answer this question in two dimensions, *i.e.*, label assignment and network architecture. As shown in Fig. 1, most of fully convolutional detec-

\*Equal contribution.

†This work was done at Megvii Technology.

tors [22, 46, 50, 21] adopt a one-to-many label assignment rule, *i.e.*, assigning many predictions as foreground samples for one ground-truth instance. This rule provides adequate foreground samples to obtain a strong and robust feature representation. Nevertheless, the massive foreground samples lead to duplicate predicted boxes for a single instance, which prevents end-to-end detection. To demonstrate it, we first give an empirical comparison of different existing hand-designed label assignments. We find that the one-to-one label assignment plays a crucial role in eliminating the post-processing of duplicate removal. However, there is still a drawback in the hand-designed one-to-one assignment. The fixed assignment could cause ambiguity issues and reduce the discriminability of features, since the predefined regions of an instance may not be the best choice [17] for training. To solve this issue, we propose a prediction-aware one-to-one (POTO) label assignment, which dynamically assigns the foreground samples according to the quality of classification and regression simultaneously.

Furthermore, for the modern FPN based detector [46], the extensive experiment demonstrates that the duplicate bounding boxes majorly come from the nearby regions of the most confident prediction across adjacent scales. Therefore, we design a 3D Max Filtering (3DMF), which can be embedded into the FPN head as a differentiable module. This module could improve the discriminability of convolution in the local regions by using a simple 3D max filtering operator across adjacent scales. Besides, to provide adequate supervision for feature representation learning, we modify a one-to-many assignment as an auxiliary loss.

With the proposed techniques, our end-to-end detection framework achieves competitive performance against many state-of-the-art detectors. On COCO [23] dataset, our end-to-end detector based on FCOS framework [46] and ResNeXt-101 [49] backbone remarkably outperforms the baseline with NMS by **1.1%** mAP. Furthermore, our end-to-end detector is more robust and flexible for crowded detection. To demonstrate the superiority in the crowded scenes, we construct more experiments on CrowdHuman [37] dataset. Under the ResNet-50 backbone, our end-to-end detector achieves **3.0%** AP<sub>50</sub> and **6.0%** mMR absolute gains over FCOS baseline with NMS.

## 2. Related Work

### 2.1. Fully Convolutional Object Detector

Owing to the success of convolution networks [11, 40, 41, 39, 20, 51, 52], object detection has achieved tremendous progress during the last decade. Modern one-stage [22, 25, 31, 38, 29, 7] or two-stage detectors [33, 21, 2] heavily rely on the anchors or anchor-based proposals. In these detectors, the anchor boxes are made up of predefined sliding windows, which are assigned as foreground

or background samples with bounding box offsets. Due to the hand-designed and data-independent anchor boxes, the training targets of anchor-based detectors are typically sub-optimal and require careful tuning of hyper-parameters. Recently, FCOS [46] and CornerNet [18] give a different perspective for fully convolutional detectors by introducing an anchor-free framework. Nevertheless, these frameworks still need a hand-designed post-processing step for duplicate removal, *i.e.*, non-maximum suppression (NMS). Since NMS is a heuristic approach and adopts a constant threshold for all the instances, it needs carefully tuning and might not be robust, especially in crowded scenes. In contrast, based on the anchor-free framework, this paper proposes a prediction-aware one-to-one assignment rule for classification to discard the non-trainable NMS.

### 2.2. End-to-End Object Detection

To achieve end-to-end detection, many approaches are explored in the previous literature. Concretely, in the earlier researches, numerous detection frameworks based on recurrent neural networks [43, 35, 26, 32, 36] attempt to produce a set of bounding boxes directly. Albeit they allow end-to-end learning in principle, they are only demonstrated effectiveness on some small datasets and not against the modern baselines [46, 8]. Meanwhile, Learnable NMS [13] is proposed to learn duplicate removal by using a very deep and complex network, which achieves comparable performance against NMS. But it is constructed by discrete components and does not give an effective solution to realize end-to-end training. Recently, the relation network [14] and DETR [3] apply the attention mechanism to object detection, which models pairwise relations between different predictions. By using one-to-one assignment rules and direct set losses, they do not need any additional post-processing steps. Nevertheless, when performing massive predictions, these methods require highly expensive cost, making them not appropriate for the dense prediction frameworks. Due to the lack of image prior and multi-scale fusion mechanism, DETR also suffers from much longer training duration than mainstream detectors and lower performance on the small objects. Different from the approaches mentioned above, our method is the first to enable end-to-end object detection based on a fully convolutional network.

## 3. Methodology

### 3.1. Analysis on Label Assignment

To reveal the effect of label assignment on end-to-end object detection, we construct several ablation studies of conventional label assignments on COCO [23] dataset. As shown in Tab. 1, all the experiments are based on FCOS [46] framework, whose centerness branch is removed to achieve a head-to-head comparison. The results demonstrate the su-

Table 1. The comparison of different label assignment rules for end-to-end object detection on COCO *val* set.  $\Delta$  indicates the gap between with and without NMS. ‘Aux’ is the proposed auxiliary loss. All models are based on ResNet-50 backbone with 180k training iterations.

Assignment	Rule	Method	mAP			mAR		
			w/ NMS	w/o NMS	$\Delta$	w/ NMS	w/o NMS	$\Delta$
One-to-many	Hand-designed	FCOS [46] baseline *	40.5	12.1	-28.4	58.3	52.8	-5.5
One-to-one	Hand-designed	Anchor	37.2	35.8	-1.4	57.0	59.2	+2.2
		Center	37.2	33.6	-3.6	57.8	59.7	+1.9
	Prediction-aware	Foreground loss	38.3	37.1	-1.2	58.6	<b>61.4</b>	<b>+2.8</b>
		POTO	38.6	38.0	-0.6	57.9	60.5	+2.6
		POTO+3DMF	40.0	39.8	-0.2	58.8	60.9	+2.1
Mixture **	Prediction-aware	POTO+3DMF+Aux	<b>41.2</b>	<b>41.1</b>	<b>-0.1</b>	<b>58.9</b>	61.2	+2.3

\* We remove its centerness branch to achieve a head-to-head comparison.

\*\* We adopt a one-to-one assignment in POTO and a one-to-many assignment in the auxiliary loss, respectively.

priority of one-to-many assignment on feature representation and the potential of one-to-one assignment on discarding the NMS. The detailed analysis is elaborated in the following sections.

### 3.1.1 One-to-many Label Assignment

Since the NMS post-processing is widely adopted in dense prediction frameworks [21, 22, 58, 53, 46, 50, 57, 29, 7], one-to-many label assignment becomes a conventional way to assign training targets. The adequate foreground samples lead to a strong and robust feature representation. However, when discarding the NMS, due to the redundant foreground samples of one-to-many label assignment, the duplicate false-positive predictions could cause a dramatic drop in performance, *e.g.*, 28.4% mAP absolute drop on FCOS [22] baseline. In addition, the reported mAR in Tab. 1 indicates the recall rates for the predictions of the top 100 scores. Without NMS, the one-to-many assignment rule leads to numerous duplicate predictions with high scores, thus reducing the recall rate. Therefore, the detector is hard to achieve competitive end-to-end detection by relying only on the one-to-many assignment.

### 3.1.2 Hand-designed One-to-one Label Assignment

MultiBox [45] and YOLO [30] demonstrate the potential in applying the one-to-one label assignment to a dense prediction framework. In this paper, we evaluate two one-to-one label assignment rules to reveal the undergoing connection with discarding NMS. These rules are modified by two widely-used one-to-many label assignments: *Anchor* rule and *Center* rule. Concretely, *Anchor* rule is based on RetinaNet [22], each ground-truth instance is only assigned to the anchor with the maximum Intersection-over-Union (IoU). *Center* rule is based on FCOS [46], each ground-truth instance is only assigned to the pixel closest to the cen-

ter of the instance in the pre-defined feature layer. Besides, other anchors or pixels are set as background samples.

As shown in Tab. 1, compared with the one-to-many label assignment, the one-to-one label assignment allows the fully convolutional detectors without NMS to greatly reduce the gap between with and without NMS and achieve reasonable performance. For instance, the detector based on *Center* rule achieves 21.5% mAP absolute gains over the FCOS baseline. Besides, as it avoids the error suppression of the NMS in complex scenes, the recall rate is further increased. Nevertheless, there still exist two unresolved issues. First, when one-to-one label assignment is applied, the performance gap between detectors with and without NMS remains non-negligible. Second, due to the less supervision for each instance, the performance of the one-to-one label assignment is still inferior to the FCOS baseline.

## 3.2. Our Methods

In this paper, to enable competitive end-to-end object detection, we propose a mixture label assignment and a new 3D Max Filtering (3DMF). The mixture label assignment is made up of the proposed prediction-aware one-to-one (POTO) label assignment and a modified one-to-many label assignment (auxiliary loss). With these techniques, our end-to-end framework can discard the NMS post-processing and keep the strong feature representation.

### 3.2.1 Prediction-aware One-to-one Label Assignment

The hand-designed one-to-one label assignment follows a fixed rule. However, this rule may be sub-optimal for various instances in complex scenes, *e.g.*, *Center* rule for an eccentric object [17]. Thus if the assignment procedure is forced to assign the sub-optimal prediction as the unique foreground sample, the difficulty for the network to converge could be dramatically increased, leading to more false-positive predictions. To this end, we propose a new




Figure 2. The diagram of the head with 3D Max Filtering (3DMF) in a FPN stage. ‘POTO’ indicates the proposed Prediction-aware One-to-one Label Assignment rule to achieve end-to-end detection. ‘Conv +  $\sigma$ ’ denotes a convolution layer followed by a sigmoid function [10], which outputs coarsely classification scores. ‘Aux Loss’ is the proposed auxiliary loss to improve feature representation. The dotted lines are used to highlight the additional components in the training phase, which are abandoned in the inference phase.

rule named Prediction-aware One-To-One (POTO) label assignment by dynamically assigning samples according to the quality of predictions.

Let  $\Psi$  denotes the index set of all the predictions.  $G$  and  $N$  correspond to the number of ground-truth instances and predictions, respectively, where typically  $G \ll N$  in dense prediction detectors.  $\hat{\pi} \in \Pi_G^N$  indicates a  $G$ -permutation of  $N$  predictions. Our POTO aims to generate a suitable permutation  $\hat{\pi}$  of predictions as the foreground samples. The training loss is formulated as Eq. 1, which consists of the foreground loss  $\mathcal{L}_{fg}$  and the background loss  $\mathcal{L}_{bg}$ .

$$\mathcal{L} = \sum_i^G \mathcal{L}_{fg}(\hat{p}_{\hat{\pi}(i)}, \hat{b}_{\hat{\pi}(i)} | c_i, b_i) + \sum_{j \in \Psi \setminus \mathcal{R}(\hat{\pi})} \mathcal{L}_{bg}(\hat{p}_j), \quad (1)$$

where  $\mathcal{R}(\hat{\pi})$  denotes the corresponding index set of the assigned foreground samples. For the  $i$ -th ground-truth,  $c_i$  and  $b_i$  are its category label and bounding box coordinates, respectively. While for the  $\hat{\pi}(i)$ -th prediction,  $\hat{p}_{\hat{\pi}(i)}$  and  $\hat{b}_{\hat{\pi}(i)}$  correspond to its predicted classification scores and predicted box coordinates, respectively.

To achieve competitive end-to-end detection, we need to find a suitable label assignment  $\hat{\pi}$ . As shown in Eq. 2, previous works [6, 3] treat it as a bipartite matching problem by using foreground loss [22, 34] as the matching cost, which can be rapidly solved by the Hungarian algorithm [43].

$$\hat{\pi} = \arg \min_{\pi \in \Pi_G^N} \sum_i^G \mathcal{L}_{fg}(\hat{p}_{\pi(i)}, \hat{b}_{\pi(i)} | c_i, b_i). \quad (2)$$

However, foreground loss typically needs additional weights to alleviate optimization issues, *e.g.*, unbalanced training samples and joint training of multiple tasks. As

shown in Tab. 1, this property makes the training loss not the optimal choice for the matching cost. Therefore, as presented in Eq. 3 and Eq. 4, we propose a more clean and effective formulation (POTO) to find a better assignment.

$$\hat{\pi} = \arg \max_{\pi \in \Pi_G^N} \sum_i^G Q_{i, \pi(i)}, \quad (3)$$

$$\text{where } Q_{i, \pi(i)} = \underbrace{\mathbb{1}[\pi(i) \in \Omega_i]}_{\text{spatial prior}} \cdot \underbrace{\left(\hat{p}_{\pi(i)}(c_i)\right)^{1-\alpha}}_{\text{classification}} \cdot \underbrace{\left(\text{IoU}(b_i, \hat{b}_{\pi(i)})\right)^\alpha}_{\text{regression}}. \quad (4)$$

Here  $Q_{i, \pi(i)} \in [0, 1]$  represents the proposed matching quality of the  $i$ -th ground-truth with the  $\pi(i)$ -th prediction. It considers the spatial prior, the confidence of classification, and the quality of regression simultaneously.  $\Omega_i$  indicates the set of candidate predictions for  $i$ -th ground-truth, *i.e.*, spatial prior. The spatial prior is widely used in the training phase [21, 22, 58, 53, 46, 50]. For instance, the center sampling strategy is adopted in FCOS [46], which only considers the predictions in the central portion of the ground-truth instance as foreground samples. We also apply it in POTO to achieve higher performance, but it is not necessary for discarding NMS (more details refer to Sec. 4.2.2). To achieve balance, we define the quality by the weighted geometric mean of classification score  $\hat{p}_{\pi(i)}(c_i)$  and regression quality  $\text{IoU}(b_i, \hat{b}_{\pi(i)})$  in Eq. 4. The hyper-parameter  $\alpha \in [0, 1]$  adjusts the ratio between classification and regression, where  $\alpha = 0.8$  is adopted by default and more ablation studies are elaborated in Sec. 4.2.2. As shown in Tab. 1, POTO not only narrows the gap with NMS but also improves the performance.

Table 2. Comparison of different configurations for NMS post-processing on COCO *val* set. ‘Across scales’ indicates applying NMS to the multiple adjacent stages of the feature pyramid network. ‘Spatial range’ denotes the spatial range for duplicate removal in each scale.

Model	Across scales	Spatial range	mAP
FCOS [46]	✗	1 × 1	19.0
		3 × 3	37.4
		5 × 5	39.2
	✓	∞ × ∞	40.9

### 3.2.2 3D Max Filtering

In addition to the label assignment, we attempt to design an effective architecture to realize more competitive end-to-end detection. To this end, we first reveal the distribution of duplicate predictions. As shown in Tab. 2, for a modern FPN based detector [46], the performance has a noticeable degradation when applying the NMS to each scale separately. Moreover, we find that the duplicate predictions majorly come from the nearby spatial regions of the most confident prediction. Therefore, we propose a new module called 3D Max Filtering (3DMF) to suppress duplicate predictions.

Convolution is a linear operation with translational equivariance, which produces similar outputs for similar patterns at different positions. However, this property has a great obstacle to duplicate removal, since different predictions of the same instance typically have similar features [22] for the dense prediction detectors. Max filter is a rank-based non-linear filter [42], which could be used to compensate for the discriminant ability of convolutions in a local region. Besides, max filter has also been utilized in the key-point based detectors, *e.g.*, CenterNet [56] and CornerNet [18], as a new post-processing step to replace the non-maximum suppression. It demonstrates some potentials to perform duplicate removal, but the non-trainable manner hinders the effectiveness and end-to-end training. Meanwhile, the max filter only considers the single-scale feature, which is not appropriate for the widely-used FPN based detectors [22, 46, 50].

Therefore, we extend the max filter to a multi-scale version, called 3D Max Filtering, which transforms the features in each scale of FPN. The 3D Max Filtering is respectively adopted in each channel of a feature map.

$$\tilde{x}^s = \left\{ \tilde{x}^{s,k} := \text{Bilinear}_{x^s}(x^k) \mid \forall k \in \left[ s - \frac{\tau}{2}, s + \frac{\tau}{2} \right] \right\}. \quad (5)$$

Specifically, as shown in Eq. 5, given an input feature  $x^s$  in the scale  $s$  of FPN, we first adopt the bilinear operator [28] to interpolate the features from  $\tau$  adjacent scales as the same




Figure 3. The diagram of 3D Max Filtering. The detailed procedure of 3D max filtering is illustrated in the dashed box. ‘GN’ and ‘σ’ indicate the group normalization [47] and the sigmoid activation function, respectively.

size of input feature  $x^s$ .

$$y_i^s = \max_{k \in [s - \frac{\tau}{2}, s + \frac{\tau}{2}]} \max_{j \in \mathcal{N}_i^{\phi \times \phi}} \tilde{x}_j^{s,k}. \quad (6)$$

As shown in Eq. 6, for a spatial location  $i$  in scale  $s$ , the maximum value  $y_i^s$  is then obtained in a pre-defined 3D neighbour tube with  $\tau$  scales and  $\phi \times \phi$  spatial distance. This operation can be easily implemented by a highly efficient 3D max-pooling operator [27].

Furthermore, to embed the 3D Max Filtering into the existing frameworks and enable end-to-end training, we propose a new module, as shown in Fig. 3. This module leverages the max filtering to select the predictions with the highest activation value in a local region and could enhance the distinction with other predictions, which is further verified in Sec. 4.2.1. Owing to this property, as shown in Fig. 2, we adopt the 3DMF to refine the coarsely dense predictions and suppress the duplicate predictions. Besides, all the modules are constructed by simple differentiable operators and only have slightly computational overhead.

### 3.2.3 Auxiliary Loss

In addition, when using the NMS, as shown in Tab. 1, the performance of POTO and 3DMF is still inferior to the FCOS baseline. This phenomenon may be attributed to the fact that one-to-one label assignment provides less supervision, making the network difficult to learn the strong and robust feature representation [44]. It could further reduce the discrimination of classification, thus causing a decrease in performance. To this end, motivated by many previous works [44, 54, 55], we introduce an auxiliary loss based on




Figure 4. Visualization of the predicted classification scores from different approaches. The input image has three instances of different scales, *i.e.*, person, tie and pot. The heatmaps from left to right of each approach correspond to the score map in the FPN stage ‘P5’, ‘P6’ and ‘P7’, respectively. ‘Aux’ indicates the proposed auxiliary loss. Our POTO based detector significantly suppresses the duplicate predictions against the vanilla FCOS framework. The 3DMF enhances the distinctiveness of the local region across adjacent scales. Besides, the auxiliary loss can further improve the feature representation.

one-to-many label assignment to provide adequate supervision, which is illustrated in Fig. 2.

Similar to ATSS [50], our auxiliary loss adopts the focal loss [22] with a modified one-to-many label assignment. Specifically, the one-to-many label assignment first takes the top-9 predictions as candidates in each FPN stage, according to the proposed matching quality in Eq. 4. It then assigns the candidates as foreground samples whose matching qualities beyond a statistical threshold. The statistical threshold is calculated by the summation of the mean and the standard deviation of all the candidate matching qualities. In addition, different forms of one-to-many label assignment for the auxiliary loss are elaborately reported in the supplementary material.

## 4. Experiments

### 4.1. Implement Detail

As same as FCOS [46], our detector adopts a pair of 4-convolution heads for classification and regression, respectively. The output channel numbers of the first convolution and the second convolution in 3DMF are 256 and 1, respectively. All the backbones are pre-trained on the ImageNet dataset [4] with frozen batch normalizations [16]. In the training phase, input images are reshaped so that their shorter side is 800 pixels. All the training hyper-parameters are identical to the 2x schedule (180k iterations) in the DeTron2 [48] if not specifically mentioned.

## 4.2. Ablation Studies on COCO

### 4.2.1 Visualization

As shown in Fig. 4, we present the visualization of the classification scores from the FCOS baseline and our proposed framework. For a single instance, the FCOS baseline with one-to-many assignment rule outputs massive duplicate predictions, which are highly activated and have comparable activating scores with the most confident one. These duplicate predictions are evaluated as false-positive samples and greatly affect performance. In contrast, by using the proposed POTO rule, the scores of duplicate samples are significantly suppressed. This property is crucial for the detector to achieve direct bounding box prediction without NMS. Moreover, with the proposed 3DMF module, this property is further enhanced, especially in the nearby regions of the most confident prediction. Besides, since the 3DMF module introduces the multi-scale competitive mechanism, the detector can well perform unique predictions across different FPN stages, *e.g.*, an instance in the Fig. 4 has single highly activated scores in various stages.

### 4.2.2 Prediction-Aware One-to-One Label Assignment

**Spatial prior.** As shown in Tab. 3, for the spatial range of assignment, the center sampling strategy is relatively superior to the inside box and global strategies on the COCO dataset. It reflects that the prior knowledge of images is essential in the real world scenario.

**Classification vs. regression.** The hyper-parameter  $\alpha$ , as

Table 3. Results of POTO with different configurations of  $\alpha$  and spatial prior on COCO *val* set.  $\alpha = 0$  is equivalent to considering classification alone,  $\alpha = 1$  is equivalent to considering regression alone. ‘center sampling’ and ‘inside box’ both follow FCOS [46]. ‘/’ is used to distinguish between results without and with NMS.

$\alpha$	center sampling	inside box	global
0.0	33.5 / 33.6	24.1 / 24.2	1.9 / 2.1
0.2	33.7 / 33.9	28.8 / 28.8	19.4 / 19.5
0.4	35.0 / 35.2	32.7 / 32.8	28.3 / 28.4
0.6	36.6 / 36.9	35.3 / 35.5	34.7 / 34.9
0.8	<b>38.0 / 38.6</b>	37.4 / 37.9	37.3 / 37.9
1.0	11.8 / 29.7	4.5 / 13.0	non-convergence

Table 4. The effect of various quality functions on COCO *val* set. ‘/’ is used to distinguish between results without and with NMS. ‘Add’ and ‘Mul’ indicate two fusion functions.

Method	$\alpha$	mAP	AP <sub>50</sub>	AP <sub>75</sub>
Add	0.2	36.0 / 36.2	<b>55.7 / 57.0</b>	38.7 / 38.3
	0.5	37.3 / 37.8	54.9 / 57.4	40.5 / 40.4
	0.8	29.3 / 35.6	40.3 / 53.4	32.8 / 38.4
Mul	0.8	<b>38.0 / 38.6</b>	<b>55.2 / 57.6</b>	<b>41.4 / 41.3</b>

shown in Eq. 4, controls the ratio of the importance between classification and regression. As reported in Tab. 3, when  $\alpha = 1$ , the gap with NMS is not narrowed. It could be attributed to the misalignment between the best positions for classification and regression. When  $\alpha = 0$ , the assignment rule only relies on the predicted scores of classification. Under this condition, the gap with NMS is considerably eliminated, but the absolute performance is still unsatisfactory, which could be caused by overfitting the sub-optimal initialization. In contrast, with a proper fusion of classification and regression quality, the absolute performance is remarkably improved.

**Quality function.** We further explore the effect of different fusion methods on the quality function, *i.e.*, Eq. 4. As presented in Tab. 4, the method called ‘Add’ replaces the original quality function by  $(1 - \alpha) \cdot \hat{p}_{\pi(i)}(c_i) + \alpha \cdot \text{IoU}(b_i, \hat{b}_{\pi(i)})$ , which has a similar form to [19]. However, we find that the multiplication fusion, *i.e.*, ‘Mul’, is more suitable for the end-to-end detection, which achieves 0.7% mAP absolute gains over the ‘Add’ fusion method.

### 4.2.3 3D Max Filtering

**Components.** As shown in Tab. 5, without NMS post-processing, our end-to-end detector with POTO achieves 19.0% mAP absolute gains over the vanilla FCOS. By using the proposed 3DMF, the performance is further improved by 1.8% mAP, and the gap with NMS is narrowed to 0.2% mAP. As shown in Fig. 4, the result shows the crucial role of the multi-scale and local-range suppression for end-to-

Table 5. The effect of sub-modules in the proposed 3DMF module on COCO *val* set. ‘3DMF’ and ‘Aux Loss’ indicate using the 3D Max Filtering and the auxiliary loss, respectively. ‘/’ is used to distinguish between results without and with NMS.

Model	3DMF	Aux Loss	mAP
FCOS [46]	<b>×</b>	<b>×</b>	19.0 / 40.9
	<b>×</b>	<b>✓</b>	18.9 / <b>41.3</b>
	<b>✓</b> *	<b>×</b>	38.7 / 40.0
Ours	<b>×</b>	<b>×</b>	38.0 / 38.6
	<b>✓</b>	<b>×</b>	39.8 / 40.0
	<b>✓</b>	<b>✓</b>	<b>41.1 / 41.2</b>

\* We modify 3D Max Filtering as a post-processing.

Table 6. The effect of hyper-parameters in the proposed 3DMF module on COCO *val* set.  $\tau = 0$  is equivalent to applying 2D Max Filtering to transform features on a single scale. ‘/’ is used to distinguish between results without and with NMS.

	$\phi = 1$	$\phi = 3$	$\phi = 5$
$\tau = 0$	39.2 / 39.5	39.1 / 39.5	39.0 / 39.4
$\tau = 2$	39.0 / 39.3	<b>39.8 / 40.0</b>	39.3 / 39.5
$\tau = 4$	39.1 / 39.3	39.3 / 39.4	39.4 / 39.6

end object detection. The proposed auxiliary loss gives adequate supervision, making our detector obtain competitive performance against the FCOS with NMS.

**End-to-end.** To demonstrate the superiority of the end-to-end training manner, we replace the 2D Max Filtering of CenterNet [5] with the 3D Max Filtering as new post-processing for duplicate removal. This post-processing is further adopted to the FCOS detector. As shown in Tab. 5, the end-to-end manner achieves significant absolute gains by 1.1% mAP.

**Kernel size.** As shown in Tab. 6, we evaluate different settings of spatial range  $\phi$  and scale range  $\tau$  in the 3DMF. When  $\phi = 3$  and  $\tau = 2$ , our method obtains the highest performance on the COCO dataset. This phenomenon reflects the duplicate predictions majorly come from a local region across adjacent scales, which is similar to the observation in Sec. 3.2.2.

**Performance *w.r.t.* training duration.** As illustrated in Fig. 5(a), at the very beginning, the performance on COCO *val* set of our end-to-end detectors is inferior to the detectors with NMS. As the training progressed, the performance gap becomes smaller and smaller. After 180k training iterations, our method finally outperforms other detectors with NMS. This phenomenon also occurs on Crowd-Human *val* set, which is shown in Fig. 5(c). Moreover, due to the removal of hand-designed post-processing, Fig. 5(b) demonstrates the superiority of our method in the recall rate against the NMS based methods.




Figure 5. The comparison graphs of performance *w.r.t.* training duration. The value of the horizontal axis corresponds to the training iterations. All the models are based on the ResNet-50 backbone. The threshold of NMS is set to 0.6.

Table 7. The experiments of the proposed framework with larger backbone on COCO2017 *test-dev* set. The hyper-parameters of all the models follow the official settings.

Backbone	Model	Epochs	mAP
ResNet-101	RetinaNet [22]	36	41.0
	FCOS [46]	36	43.1
	DETR [3]	500	43.5
	Ours (w/o NMS)	36	<b>43.6</b>
ResNeXt-101+DCN	RetinaNet [22]	24	44.5
	FCOS [46]	24	46.5
	Ours (w/o NMS)	24	<b>47.6</b>

#### 4.2.4 Larger Backbone

To further demonstrate the robustness and effectiveness of our method, we provide experiments with larger backbones. The detailed results are reported in Tab. 7. Concretely, when using the ResNet-101 as the backbone, our method is slightly superior to FCOS by 0.5% mAP. But when introducing more stronger backbone, *i.e.*, ResNeXt-101 [49] with deformable convolutions [59], our end-to-end detector achieves 1.1% mAP absolute gains over the FCOS with NMS. It might be attributed to the flexibly spatial modeling of deformable convolutions. Moreover, the proposed 3DMF is efficient and easy to implement. As shown in Tab. 7, our 3DMF module only has a slightly computational overhead against the baseline detector with NMS.

#### 4.3. Evaluation on CrowdHuman

We evaluate our model on the CrowdHuman dataset [37], which is a large human detection dataset with various kinds of occlusions. Compared with the COCO dataset, CrowdHuman has more complex and crowded scenes, giving severe challenges to conventional duplicate removal. Our end-to-end detector is more robust and flexible in crowded scenes. As shown in Tab. 8 and Fig. 5, our method signif-

Table 8. The comparison of fully convolutional detectors on CrowdHuman *val* set. All models are based on the ResNet-50 backbone. ‘Aux’ indicates the auxiliary loss.

Method	Epochs	AP <sub>50</sub>	mMR	Recall
RetinaNet [22]	32	81.7	57.6	88.6
FCOS [46]	32	86.1	54.9	94.2
ATSS [50]	32	87.2	49.7	94.0
DETR [3]	300	72.8	80.1	82.7
Ground-truth (w/ NMS)	-	-	-	95.1
POTO	32	88.5	52.2	96.3
POTO+3DMF	32	88.8	51.0	<b>96.6</b>
POTO+3DMF+Aux	32	<b>89.1</b>	<b>48.9</b>	96.5

icantly outperforms several state-of-the-art detectors with NMS, *e.g.*, 3.0% mAP and 6.0% mMR absolute gains over the FCOS. Moreover, the recall rate of our method is even superior to the ground-truth boxes with NMS.

## 5. Conclusion

This paper has presented a prediction-aware one-to-one label assignment and a 3D Max Filtering to bridge the gap between fully convolutional network and end-to-end object detection. With the auxiliary loss, our end-to-end framework achieves superior performance against many state-of-the-art detectors with NMS on COCO and CrowdHuman datasets. Our method is also demonstrated great potential in complex and crowded scenes, which may benefit many other instance-level tasks.

## Acknowledgement

This research was supported by National Key R&D Program of China (No. 2017YFA0700800), National Natural Science Foundation of China (No. 61790563 and 61751401) and Beijing Academy of Artificial Intelligence (BAAI).

## References

- [1] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S Davis. Soft-nms—improving object detection with one line of code. In *IEEE International Conference on Computer Vision*, 2017. 1
- [2] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. *European Conference on Computer Vision*, 2020. 1, 2, 4, 8
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 6
- [5] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *IEEE International Conference on Computer Vision*, 2019. 1, 7
- [6] Dumitru Erhan, Christian Szegedy, Alexander Toshev, and Dragomir Anguelov. Scalable object detection using deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014. 4
- [7] Zheng Ge, Jianfeng Wang, Xin Huang, Songtao Liu, and Osamu Yoshie. Lla: Loss-aware label assignment for dense pedestrian detection. *arXiv preprint arXiv:2101.04307*, 2021. 2, 3
- [8] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V Le. Nas-fpn: Learning scalable feature pyramid architecture for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [9] Ross Girshick. Fast r-cnn. In *IEEE International Conference on Computer Vision*, 2015. 1
- [10] Jun Han and Claudio Moraga. The influence of the sigmoid function parameters on the speed of backpropagation learning. In *International Workshop on Artificial Neural Networks*, 1995. 4
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [12] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1
- [13] Jan Hosang, Rodrigo Benenson, and Bernt Schiele. Learning non-maximum suppression. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2
- [14] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [15] Xin Huang, Zheng Ge, Zequn Jie, and Osamu Yoshie. Nms by representative region: Towards crowded pedestrian detection by proposal pairing. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 6
- [17] Kang Kim and Hee Seok Lee. Probabilistic anchor assignment with iou prediction for object detection. *arXiv preprint arXiv:2007.08103*, 2020. 2, 3
- [18] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *European Conference on Computer Vision*, 2018. 2, 5
- [19] Hengduo Li, Zuxuan Wu, Chen Zhu, Caiming Xiong, Richard Socher, and Larry S Davis. Learning from noisy anchors for one-stage object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 7
- [20] Yanwei Li, Lin Song, Yukang Chen, Zeming Li, Xiangyu Zhang, Xingang Wang, and Jian Sun. Learning dynamic routing for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 3, 4
- [22] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision*, 2017. 1, 2, 3, 4, 5, 6, 8
- [23] 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*, 2014. 2
- [24] Songtao Liu, Di Huang, and Yunhong Wang. Adaptive nms: Refining pedestrian detection in a crowd. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1
- [25] 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*, 2016. 2
- [26] Eunbyung Park and Alexander C Berg. Learning to decompose for object detection and instance segmentation. *arXiv preprint arXiv:1511.06449*, 2015. 1, 2
- [27] 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*, 2019. 5
- [28] William H Press, Saul A Teukolsky, William T Vetterling, and Brian P Flannery. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007. 5
- [29] Han Qiu, Yuchen Ma, Zeming Li, Songtao Liu, and Jian Sun. Borderdet: Border feature for dense object detection. In *European Conference on Computer Vision*, 2020. 2, 3
- [30] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 1, 3
- [31] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2

- [32] Mengye Ren and Richard S Zemel. End-to-end instance segmentation with recurrent attention. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, 2015. 2
- [34] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 4
- [35] Bernardino Romera-Paredes and Philip Hilaire Sean Torr. Recurrent instance segmentation. In *European Conference on Computer Vision*, 2016. 1, 2
- [36] Amaia Salvador, Miriam Bellver, Victor Campos, Manel Baradad, Ferran Marques, Jordi Torres, and Xavier Giro-i Nieto. Recurrent neural networks for semantic instance segmentation. *arXiv preprint arXiv:1712.00617*, 2017. 1, 2
- [37] Shuai Shao, Zijian Zhao, Boxun Li, Tete Xiao, Gang Yu, Xiangyu Zhang, and Jian Sun. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*, 2018. 2, 8
- [38] Lin Song, Yanwei Li, Zhengkai Jiang, Zeming Li, Hongbin Sun, Jian Sun, and Nanning Zheng. Fine-grained dynamic head for object detection. *Advances in Neural Information Processing Systems*, 2020. 2
- [39] Lin Song, Yanwei Li, Zhengkai Jiang, Zeming Li, Xiangyu Zhang, Hongbin Sun, Jian Sun, and Nanning Zheng. Rethinking learnable tree filter for generic feature transform. *Advances in Neural Information Processing Systems*, 2020. 2
- [40] Lin Song, Yanwei Li, Zeming Li, Gang Yu, Hongbin Sun, Jian Sun, and Nanning Zheng. Learnable tree filter for structure-preserving feature transform. In *Advances in Neural Information Processing Systems*, 2019. 2
- [41] Lin Song, Shiwei Zhang, Gang Yu, and Hongbin Sun. Tacnet: Transition-aware context network for spatio-temporal action detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [42] Milan Sonka, Vaclav Hlavac, and Roger Boyle. *Image processing, analysis, and machine vision*. Cengage Learning, 2014. 5
- [43] Russell Stewart, Mykhaylo Andriluka, and Andrew Y Ng. End-to-end people detection in crowded scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 1, 2, 4
- [44] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 5
- [45] Christian Szegedy, Scott Reed, Dumitru Erhan, Dragomir Anguelov, and Sergey Ioffe. Scalable, high-quality object detection. *arXiv preprint arXiv:1412.1441*, 2014. 3
- [46] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *International Conference on Computer Vision*, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [47] Yuxin Wu and Kaiming He. Group normalization. In *European Conference on Computer Vision*, 2018. 5
- [48] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 6
- [49] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2, 8
- [50] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2, 3, 4, 5, 6, 8
- [51] Shiwei Zhang, Lin Song, Changxin Gao, and Nong Sang. Glnet: Global local network for weakly supervised action localization. *IEEE Transactions on Multimedia*, 2019. 2
- [52] Songyang Zhang, Shipeng Yan, and Xuming He. Latentgcn: Learning efficient non-local relations for visual recognition. In *International Conference on Machine Learning*, 2019. 2
- [53] Xiaosong Zhang, Fang Wan, Chang Liu, Rongrong Ji, and Qixiang Ye. Freeanchor: Learning to match anchors for visual object detection. In *Advances in Neural Information Processing Systems*, 2019. 3, 4
- [54] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 5
- [55] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems*, 2019. 5
- [56] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 5
- [57] Benjin Zhu, Jianfeng Wang, Zhengkai Jiang, Fuhang Zong, Songtao Liu, Zeming Li, and Jian Sun. Autoassign: Differentiable label assignment for dense object detection. *arXiv preprint arXiv:2007.03496*, 2020. 3
- [58] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1, 3, 4
- [59] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 8