

TOOD: Task-aligned One-stage Object Detection

Chengjian Feng*
Intellifusion Inc.

feng.chengjian@intellif.com

Yujie Zhong*
Meituan Inc.

zhongyujie@meituan.com

Yu Gao
ByteDance Inc.

gaoyu.1520@bytedance.com

Matthew R. Scott
Malong LLC

msscott@malongtech.com

Weilin Huang[†]
Alibaba Group

weilin.hwl@alibaba-inc.com

Abstract

One-stage object detection is commonly implemented by optimizing two sub-tasks: object classification and localization, using heads with two parallel branches, which might lead to a certain level of spatial misalignment in predictions between the two tasks. In this work, we propose a Task-aligned One-stage Object Detection (TOOD) that explicitly aligns the two tasks in a learning-based manner. First, we design a novel Task-aligned Head (T-Head) which offers a better balance between learning task-interactive and task-specific features, as well as a greater flexibility to learn the alignment via a task-aligned predictor. Second, we propose Task Alignment Learning (TAL) to explicitly pull closer (or even unify) the optimal anchors for the two tasks during training via a designed sample assignment scheme and a task-aligned loss. Extensive experiments are conducted on MS-COCO, where TOOD achieves a **51.1 AP** at single-model single-scale testing. This surpasses the recent one-stage detectors by a large margin, such as ATSS [30] (47.7 AP), GFL [14] (48.2 AP), and PAA [9] (49.0 AP), with fewer parameters and FLOPs. Qualitative results also demonstrate the effectiveness of TOOD for better aligning the tasks of object classification and localization. Code is available at <https://github.com/fcjian/TOOD>.

1. Introduction

Object detection aims to localize and recognize objects of interest from natural images, and is a fundamental yet challenging task in computer vision. It is commonly formulated as a multi-task learning problem by jointly optimizing object classification and localization [4, 6, 7, 16, 22, 32]. The classification task is designed to learn discriminative features that focus on the key or salient part of an object,

*Equal contributions. [†] Corresponding author.

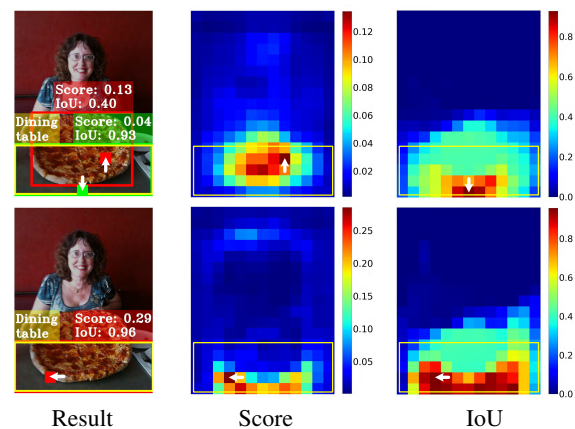


Figure 1. Illustration of detection results (‘Result’) and spatial distributions of classification scores (‘Score’) and localization scores (‘IoU’) predicted by ATSS [30] (top row) and the proposed TOOD (bottom row). Ground-truth is indicated by yellow boxes, and a white arrow means the main direction of the best anchor away from the center of an object. In the ‘Result’ column, a red/green patch is the location of the best anchor for classification/localization, while a red/green box means an object bounding box predicted from the anchor in the red/green patch (if they coincide, we only show the red patches and boxes).

while the localization task works on precisely locating the whole object with its boundaries. Due to the divergence of learning mechanisms for classification and localization, spatial distributions of the learned features by the two tasks can be different, causing a certain level of misalignment when predictions are made by using two separate branches.

Recent one-stage object detectors attempted to predict consistent outputs of the two separate tasks, by focusing on the center of an object [3, 10, 27, 30]. They assume that an anchor (*i.e.*, an anchor-point for an anchor-free detector, or an anchor-box for an anchor-based detector) at the center of the object is likely to give more accurate predictions for both classification and localization. For example, recent FCOS [27] and ATSS [30] both use a centerness branch

to enhance classification scores predicted from the anchors near the center of the object, and assign larger weights to the localization loss for the corresponding anchors. Besides, FoveaBox [10] regards the anchors inside a predefined central region of the object as positive samples. Such heuristic designs have achieved excellent results, but these methods might suffer from two limitations:

(1) Independence of classification and localization. Recent one-stage detectors perform object classification and localization independently by using two separate branches in parallel (*i.e.*, heads). Such a two-branch design might cause a lack of interaction between the two tasks, leading to an inconsistency in predictions when performing them. As shown in the ‘Result’ column in Figure 1, an ATSS detector recognizes an object of ‘Dining table’ (indicated by the anchor shown with a red patch), but localizes another object of ‘Pizza’ more accurately (red bounding box).

(2) Task-agnostic sample assignment. Most anchor-free detectors use a geometry-based assignment scheme to select anchor-points near the center of an object for both classification and localization [3, 10, 30], while anchor-based detectors often assign anchor-boxes by computing IoUs between the anchor boxes and ground truth [22, 23, 30]. However, the optimal anchors for classification and localization are often inconsistent, and may vary considerably depending on the shape and characteristics of the objects. The widely used sample assignment scheme is task agnostic, and thus may be difficult to make an accurate yet consistent prediction for the two tasks, as demonstrated in ‘Score’ and ‘IoU’ distributions of ATSS in Figure 1. The ‘Result’ column also illustrates that a spatial location of the best localization anchor (green patch) can be not at the center of the object, and it is not well aligned with the best classification anchor (red patch). As a result, a precise bounding box may be suppressed by the less accurate one during Non-Maximum Suppression (NMS).

To address such limitations, we propose a Task-aligned One-stage Object Detection (TOOD) that aims to align the two tasks more accurately by designing a new head structure with an alignment-oriented learning approach:

Task-aligned head. In contrast to the conventional head in one-stage object detection where classification and localization are implemented separately by using two branches in parallel, we design a Task-aligned head (T-head) to enhance an *interaction* between the two tasks. This allows the two tasks to work more collaboratively, which in turn aligns their predictions more accurately. T-head is conceptually simple: it computes task-interactive features, and makes predictions via a novel Task-Aligned Predictor (TAP). Then it aligns spatial distributions of the two predictions according to the learning signals provided by a task alignment learning, as described next.

Task alignment learning. To further overcome the misalignment problem, we propose a Task Alignment Learning (TAL) to *explicitly* pull closer the optimal anchors for the two tasks. It is performed by designing a sample assignment scheme and a task-aligned loss. The sample assignment collects training samples (*i.e.*, positives or negatives) by computing a degree of task-alignment at each anchor, whereas the task-aligned loss gradually unifies the best anchors for predicting both classification and localization during the training. Therefore, at inference, a bounding box with the highest classification score and jointly having the most precise localization can be preserved.

The proposed T-head and learning strategy can work collaboratively towards making predictions with high quality in both classification and localization. The main contributions of this work can be summarized as follows: (1) we design a new T-head to enhance the interaction between classification and localization while maintaining their characteristics, and further align the two tasks at the predictions; (2) we propose TAL to explicitly align the two tasks at the identified task-aligned anchors, as well as providing learning signals for the proposed predictor; (3) we conducted extensive experiments on MSCOCO [17], where our TOOD achieved a **51.1 AP**, surpassing recent one-stage detectors such as ATSS [30], GFL [14] and PAA [9], by a large margin. Qualitative results further validate the effectiveness of our task-alignment approaches.

2. Related Work

One-stage detectors. OverFeat [25] is one of the earliest CNN-based one-stage detectors. Afterward, YOLO [22] was developed to directly predict bounding boxes and classification scores, without an additional stage to generate region proposals. SSD [18] introduces anchors with multi-scale predictions from multi-layer convolutional features, and Focal loss [16] was proposed to address the problem of class imbalance for one-stage detectors like RetinaNet. Keypoint-based detection methods, such as [3, 11, 33], address the detection problem by identifying and grouping multiple key points of a bounding box. Recently, FCOS [27] and FoveaBox [10] were developed to locate objects of interest via anchor-points and point-to-boundary distances. Most mainstream one-stage detectors are composed of two FCN-based branches for classification and localization, which may lead to the misalignment between the two tasks. In this paper, we enhance the alignment between the two tasks via a new head structure and an alignment-oriented learning approach.

Training sample assignment. Most anchor-based detectors such as [22, 30], collect training samples by computing IoUs between proposals and ground truth, while an anchor-free detector regards the anchors inside the center region

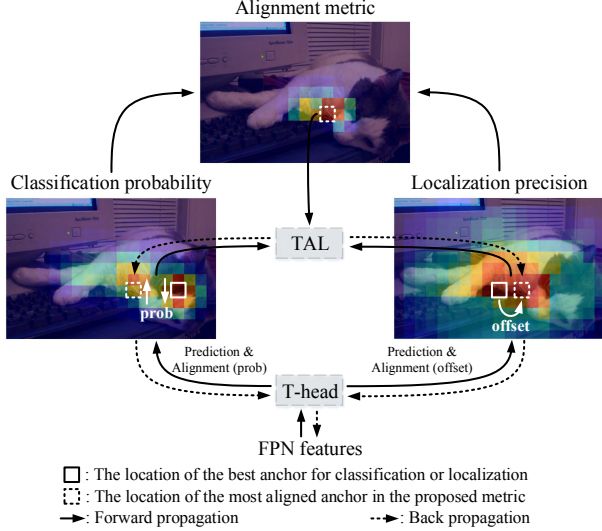


Figure 2. Overall learning mechanism of TOOD. First, predictions are made by T-head on the FPN features. Second, the predictions are used to compute a task alignment metric at each anchor point, based on which TAL produces learning signals for T-head. Lastly, T-head adjusts the distributions of classification and localization accordingly. Specifically, the most aligned anchor obtains a higher classification score through ‘prob’ (probability map), and acquires a more accurate bounding box prediction via a learned ‘offset’.

of an object as positive samples [3, 10, 27]. Recent studies attempted to train the detectors more effectively by collecting more informative training samples using output results. For example, FSAF [35] selects meaningful samples from feature pyramids based on the computed loss, and similarly, SAPD [34] provides a soft-selection version of FSAF by designing a meta-selection network. FreeAnchor [31] and MAL [8] identify the best anchor-box by computing the losses in an effort to improve the matching between anchors and objects. PAA [9] adaptively separates the anchors into positive and negative samples by fitting a probability distribution to the anchor scores. Different from the positive/negative sample assignment, PISA [1] re-weights the training samples according to the precision rank of the outputs. Noisy Anchor [12] assigns soft labels to the training samples, and re-weights the anchor-boxes using a cleanliness score to mitigate the noise incurred by binary labels. GFL [14] replaces the binary classification label with an IoU score to integrate the localization quality into classification. These excellent approaches inspired the current work to develop a new assignment mechanism from a task-alignment point of view.

3. Task-aligned One-stage Object Detection

Overview. Similar to recent one-stage detectors such as [14, 30], the proposed TOOD has an overall pipeline of ‘backbone-FPN-head’. Moreover, by considering effi-

ciency and simplicity, TOOD uses a single anchor per location (same as ATSS [30]), where the ‘anchor’ means an anchor point for an anchor-free detector, or an anchor box for an anchor-based detector. As discussed, existing one-stage detectors have limitations of task misalignment between classification and localization, due to the divergence of two tasks which are often implemented using two separate head branches. In this work, we propose to align the two tasks more explicitly using a designed Task-aligned head (T-head) with a new Task Alignment Learning (TAL). As illustrated in Figure 2, T-head and TAL can work collaboratively to improve the alignment of two tasks. Specifically, T-head first makes predictions for the classification and localization on the FPN features. Then TAL computes task alignment signals based on a new task alignment metric which measures the degree of alignment between the two predictions. Lastly, T-head automatically adjusts its classification probabilities and localization predictions using learning signals computed from TAL during back propagation.

3.1. Task-aligned Head

Our goal is to design an efficient head structure to improve the conventional design of the head in one-stage detectors (as shown in Figure 3(a)). In this work, we achieve this by considering two aspects: (1) increasing the interaction between the two tasks, and (2) enhancing the detector ability of learning the alignment. The proposed T-head is shown in Figure 3(b), where it has a simple feature extractor with two Task-Aligned Predictors (TAP).

To enhance the interaction between classification and localization, we use a feature extractor to learn a stack of *task-interactive* features from multiple convolutional layers, as shown by the blue part in Figure 3(b). This design not only facilitates the task interaction, but also provides multi-level features with multi-scale effective receptive fields for the two tasks. Formally, let $X^{fpn} \in \mathbb{R}^{H \times W \times C}$ denotes the FPN features, where H , W and C indicate height, width and the number of channels, respectively. The feature extractor uses N consecutive conv layers with activation functions to compute the *task-interactive* features:

$$X_k^{inter} = \begin{cases} \delta(\text{conv}_k(X^{fpn})), k = 1 \\ \delta(\text{conv}_k(X_{k-1}^{inter})), k > 1 \end{cases}, \forall k \in \{1, 2, \dots, N\}, \quad (1)$$

where conv_k and δ refer to the k -th conv layer and a *relu* function, respectively. Thus we extract rich multi-scale features from the FPN features using a single branch in the head. Then, the computed task-interactive features will be fed into two TAP for aligning classification and localization.

Task-aligned Predictor (TAP). We perform both object classification and localization on the computed task-interactive features, where the two tasks can well perceive the state of each other. However, due to the single branch

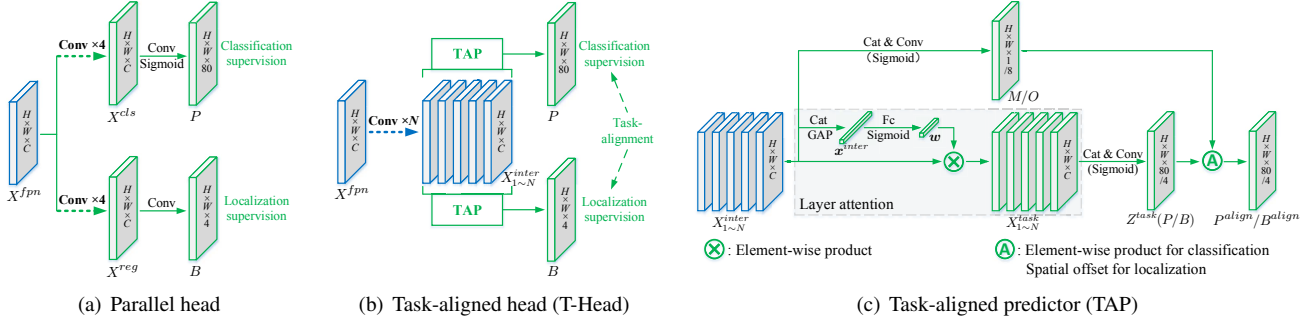


Figure 3. Comparison between the conventional parallel head and the proposed T-Head.

design, the task-interactive features inevitably introduce a certain level of feature conflicts between two different tasks, which have also been discussed in [26, 28]. Intuitively, the tasks of object classification and localization have different targets, and thus focus on different types of features (*e.g.*, different levels or receptive fields). Consequently, we propose a layer attention mechanism to encourage task decomposition by dynamically computing such *task-specific* features at the layer level. As shown in Figure 3(c), the task-specific features are computed separately for each task of classification or localization:

$$X_k^{task} = \mathbf{w}_k \cdot X_k^{inter}, \forall k \in \{1, 2, \dots, N\}, \quad (2)$$

where \mathbf{w}_k is the k -th element of the learned layer attention $\mathbf{w} \in \mathbb{R}^N$. \mathbf{w} is computed from the cross-layer task-interactive features, and is able to capture the dependencies between layers:

$$\mathbf{w} = \sigma(f_{c2}(\delta(f_{c1}(\mathbf{x}^{inter})))), \quad (3)$$

where f_{c1} and f_{c2} refer to two fully-connected layers. σ is a *sigmoid* function, and \mathbf{x}^{inter} is obtained by applying an average pooling to X^{inter} which are the concatenated features of X_k^{inter} . Finally, the results of classification or localization are predicted from each X^{task} :

$$Z^{task} = conv_2(\delta(conv_1(X^{task}))), \quad (4)$$

where X^{task} is the concatenated features of X_k^{task} , and $conv_1$ is a 1×1 conv layer for dimension reduction. Z^{task} is then converted into dense classification scores $P \in \mathbb{R}^{H \times W \times 80}$ using *sigmoid* function, or object bounding boxes $B \in \mathbb{R}^{H \times W \times 4}$ with *distance-to-bbox* conversion as applied in [27, 30].

Prediction alignment. At the prediction step, we further align the two tasks explicitly by adjusting the spatial distributions of the two predictions: P and B . Different from the previous works using a centerness branch [27] or an IoU branch [9] which can only adjust the classification prediction based on either classification features or localization

features, we align the two predictions by considering both tasks jointly using the computed task-interactive features. Notably, we perform the alignment method separately on the two tasks. As shown in Figure 3(c), we use a spatial probability map $M \in \mathbb{R}^{H \times W \times 1}$ to adjust the classification prediction:

$$P^{align} = \sqrt{P \times M}, \quad (5)$$

where M is computed from the interactive features, allowing it to learn a degree of consistency between the two tasks at each spatial location.

Meanwhile, to make an alignment on localization prediction, we further learn spatial offset maps $O \in \mathbb{R}^{H \times W \times 8}$ from the interactive features, which are used to adjust the predicted bounding box at each location. Specifically, the learned spatial offset enables the most aligned anchor point to identify the best boundary predictions around it:

$$B^{align}(i, j, c) = B(i + O(i, j, 2 \times c), j + O(i, j, 2 \times c + 1), c), \quad (6)$$

where an index (i, j, c) denotes the (i, j) -th spatial location at the c -th channel in a tensor. Eq.(6) is implemented by bilinear interpolation, and its computational overhead is negligible due to the very small dimension of B . Noteworthy, offsets are learned independently for each channel, which means each boundary of the object has its own learned offset. This allows for a more accurate prediction of the four boundaries because each of them can individually learn from the most precise anchor point near it. Therefore, our method not only aligns the two tasks, but also improves the localization accuracy by identifying a precise anchor point for each side.

The alignment maps M and O are learned automatically from the stack of interactive features:

$$M = \sigma(conv_2(\delta(conv_1(X^{inter})))) \quad (7)$$

$$O = conv_4(\delta(conv_3(X^{inter}))) \quad (8)$$

where $conv_1$ and $conv_3$ are two 1×1 conv layers for dimension reduction. The learning of M and O is performed by using the proposed Task Alignment Learning (TAL) which

will be described next. Notice that our T-head is an independent module and can work well without TAL. It can be readily applied to various one-stage object detectors in a plug-and-play manner to improve detection performance.

3.2. Task Alignment Learning

We further introduce a Task Alignment Learning (TAL) that further guides our T-head to make task-aligned predictions. TAL differs from previous methods [1, 8, 9, 12, 14, 31] in two aspects. First, it is designed from a task-alignment point of view. Second, it considers both anchor assignment and weighting simultaneously. It comprises a sample assignment strategy and new losses designed specifically for aligning the two tasks.

3.2.1 Task-aligned Sample Assignment

To cope with NMS, the anchor assignment for a training instance should satisfy the following rules: (1) a well-aligned anchor should be able to predict a high classification score with a precise localization jointly; (2) a misaligned anchor should have a low classification score and be suppressed subsequently. With the two objectives, we design a new anchor alignment metric to explicitly measure the degree of task-alignment at the anchor level. The alignment metric is integrated into the sample assignment and loss functions to dynamically refine the predictions at each anchor.

Anchor alignment metric. Considering that a classification score and an IoU between the predicted bounding box and the ground truth indicate the quality of the predictions by the two tasks, we measure the degree of task-alignment using a high-order combination of the classification score and the IoU. To be specific, we design the following metric to compute anchor-level alignment for each instance:

$$t = s^\alpha \times u^\beta, \quad (9)$$

where s and u denote a classification score and an IoU value, respectively. α and β are used to control the impact of the two tasks in the anchor alignment metric. Notably, t plays a critical role in the joint optimization of the two tasks towards the goal of task-alignment. It encourages the networks to dynamically focus on high-quality (*i.e.*, task-aligned) anchors from the perspective of joint optimization.

Training sample assignment. As discussed in [30, 31], training sample assignment is crucial to the training of object detectors. To improve the alignment of two tasks, we focus on the task-aligned anchors, and adopt a simple assignment rule to select the training samples: for each instance, we select m anchors having the largest t values as positive samples, while using the remaining anchors as negative ones. Again, the training is performed by computing new loss functions designed specifically for aligning the tasks of classification and localization.

3.2.2 Task-aligned Loss

Classification objective. To explicitly increase classification scores for the aligned anchors, and at the same time, reduce the scores of the misaligned ones (*i.e.*, having a small t), we use t to replace the binary label of a positive anchor during training. However, we found that the network cannot converge when the labels (*i.e.*, t) of the positive anchors become small with the increase of α and β . Therefore, we use a normalized t , namely \hat{t} , to replace the binary label of the positive anchor, where \hat{t} is normalized by the following two properties: (1) to ensure effective learning of hard instances (which usually have a small t for all corresponding positive anchors); (2) to preserve the rank between instances based on the precision of the predicted bounding boxes. Thus, we adopt a simple instance-level normalization to adjust the scale of \hat{t} : the maximum of \hat{t} is equal to the largest IoU value (u) within each instance. Then *Binary Cross Entropy* (BCE) computed on the positive anchors for the classification task can be rewritten as,

$$L_{cls.pos} = \sum_{i=1}^{N_{pos}} BCE(s_i, \hat{t}_i), \quad (10)$$

where i denotes the i -th anchor from the N_{pos} positive anchors corresponding to one instance. Following [16], we employ a focal loss for classification to mitigate the imbalance between the negative and positive samples during training. The focal loss computed on the positive anchors can be reformulated by Eq.(10), and the final loss function for the classification task is defined as follows:

$$L_{cls} = \sum_{i=1}^{N_{pos}} |\hat{t}_i - s_i|^\gamma BCE(s_i, \hat{t}_i) + \sum_{j=1}^{N_{neg}} s_j^\gamma BCE(s_j, 0), \quad (11)$$

where j denotes the j -th anchor from the N_{neg} negative anchors, and γ is the focusing parameter [16].

Localization objective. A bounding box predicted by a well-aligned anchor (*i.e.*, having a large t) usually has both a large classification score with a precise localization, and such a bounding box is more likely to be preserved during NMS. In addition, t can be applied for selecting high-quality bounding boxes by weighting the loss more carefully to improve the training. As discussed in [21], learning from high-quality bounding boxes is beneficial to the performance of a model, while the low-quality ones often have a negative impact on the training by producing a large amount of less informative and redundant signals to update the model. In our case, we apply the t value for measuring the quality of a bounding box. Thus, we improve the task alignment and regression precision by focusing on the well-aligned anchors (with a large t), while reducing the impact of the misaligned anchors (with a small t) during bounding

Method	Head	Head/full Params (M)	Head/full FLOPs (G)	AP	AP ₅₀	AP ₇₅
FoveaBox [10]	Parallel head	4.92/36.20	104.87/206.28	37.3	56.2	39.7
	T-head	4.82/36.10	100.79/202.20	38.0	56.8	40.5
FCOS w/ imprv [27]	Parallel head	4.92/32.02	104.91/200.50	38.6	57.2	41.7
	T-head	4.82/31.92	100.79/196.38	40.5	58.5	43.8
ATSS (anchor-based) [30]	Parallel head	4.92/32.07	104.87/205.21	39.3	57.5	42.8
	T-head	4.82/31.98	100.79/201.13	41.1	58.6	44.5
ATSS (anchor-free) [30]	Parallel head	4.92/32.07	104.87/205.21	39.2	57.4	42.2
	T-head	4.82/31.98	100.79/201.13	41.1	58.4	44.5

Table 1. Comparisons between different head structures in various detectors. FLOPs are measured on the input image size of 1280×800.

box regression. Similar to the classification objective, we re-weight the loss of bounding box regression computed for each anchor based on \hat{t} , and a $GIoU$ loss (L_{GIoU}) [24] can be reformulated as follows:

$$L_{reg} = \sum_{i=1}^{N_{pos}} \hat{t}_i L_{GIoU}(b_i, \bar{b}_i), \quad (12)$$

where b and \bar{b} denote the predicted bounding boxes and the corresponding ground-truth boxes. The total training loss for TAL is the sum of L_{cls} and L_{reg} .

4. Experiments and Results

Dataset and evaluation protocol. All experiments are implemented on the large-scale detection benchmark MSCOCO 2017 [17]. Following the standard practice [15, 16], we use the *trainval135k* set (115K images) for training and *minival* set (5K images) as validation for our ablation study. We report our main results on the *test-dev* set for comparison with the state-of-the-art detectors. The performance is measured by COCO Average Precision (AP) [17].

Implementation details. As with most one-stage detectors [10, 16, 27], we use the detection pipeline of ‘backbone-FPN-head’, with different backbones including ResNet-50, ResNet-101 and ResNeXt-101-64×4d pre-trained on ImageNet [2]. Similar to ATSS [30], TOOD tiles one anchor per location. Unless specified, we report experimental results of an anchor-free TOOD (an anchor-based TOOD can achieve a similar performance as shown in Table 3). The number of interactive layers N is set as 6 to make T-head have a similar number of parameters as the conventional parallel head, and the focusing parameter γ is set to 2 as used in [14, 16]. More implementation and training details are presented in Supplementary Material (SM).

4.1. Ablation Study

For an ablation study, we use the ResNet-50 backbone and train the model for 12 epochs unless specified. The performances are reported on COCO *minival* set.

Anchor assignment	Pos/neg	Weight	AP	AP ₅₀	AP ₇₅
IoU-based [16]	fixed	fixed	36.5	55.5	38.7
Center sampling [10]	fixed	fixed	37.3	56.2	39.3
Centerness [27]	fixed	fixed	37.4	56.1	40.3
ATSS [30]	fixed	fixed	39.2	57.4	42.2
PISA [1]	fixed	ada	37.3	56.5	40.3
NoisyAnchor [12]	fixed	ada	38.0	56.9	40.6
ATSS+QFL [14]	fixed	ada	39.9	58.5	43.0
FreeAnchor [31]	ada	fixed	39.1	58.2	42.1
MAL [8]	ada	fixed	39.2	58.0	42.3
PAA [9]*	ada	fixed	39.9	59.1	42.8
PAA+IoU pred. [9]*	ada	fixed	40.9	59.4	43.9
TAL	ada	ada	40.3	58.5	43.8
TAL*	ada	ada	40.9	59.3	44.3
TAL + TAP*	ada	ada	42.5	60.3	46.4

Table 2. Comparisons between different schemes of training sample assignments. ‘Pos/neg’: positive/negative anchor assignment. ‘Weight’: anchor weight assignment. ‘fixed’: fixed assignment. ‘ada’: adaptive assignment. Here TAP aligns the predictions based on both classification and localization features from the last head tower. * indicates the model is trained for 18 epochs.

On head structures. We compare our T-head with the conventional parallel head in Table 1. It can be adopted in various one-stage detectors in a plug-and-play manner, and consistently outperforms the conventional head by 0.7 to 1.9 AP, with fewer parameters and FLOPs. This validates the effectiveness of our design, and demonstrates that T-head can work more efficiently with higher performance, by introducing task interaction and prediction alignment.

On sample assignments. To demonstrate the effectiveness of TAL, we compare TAL with other learning methods using different sample assignment methods, as shown in Table 2. Training sample assignment can be divided into the fixed assignment and adaptive assignment according to whether it is a learning-based method. Different from the existing assignment methods, TAL adaptively assigns both positive and negative anchors, and at the same time, computes the weights of positive anchors more carefully, resulting in higher performance. To compare with PAA (+IoU pred.) which has an additional prediction structure, we inte-




Figure 4. Illustration of several detection results predicted from the best anchors for classification (in red) and localization (in green). The illustrated patches and bounding boxes correspond to that in Figure 1.

grate TAP into TAL, resulting in a higher AP of 42.5. More discussions on the differences between TAL and previous methods are presented in SM.

TOOD. We evaluate the performance of the complete TOOD (T-head + TAL). As shown in Table 3, the anchor-free TOOD and anchor-based TOOD can achieve similar performance, *i.e.*, 42.5 AP and 42.4 AP. Compared with ATSS, TOOD improves the performance of ~ 3.2 AP. To be more specific, the improvements on AP_{75} are significant, which yields ~ 3.8 points higher AP in TOOD. This validates that aligning the two tasks can improve the detection performance. Notably, TOOD brings a higher improvement (+3.3 AP) than the sum of the individual improvements by T-head + ATSS (+1.9 AP) and Parallel head + TAL (+1.1 AP), as shown in Table 6. It suggests that T-head and TAL can compensate strongly to each other.

On hyper-parameters. We first investigate the performance using different values of α and β for TAL, which control the influence of classification confidence and localization precision on anchor alignment metric via t . Through a coarse search shown in Table 4, we adopt $\alpha = 1$ and $\beta = 6$ for our TAL. We then conduct several experiments to study the robustness of the hyper-parameter m , which is used to select positive anchors. We use different values of m in [5, 9, 13, 17, 21], and achieve results in a range of 42.0 \sim 42.5 AP, which suggests the performance is insensitive to m . Thus, we adopt $m = 13$ in all our experiments.

4.2. Comparison with the State-of-the-Art

We compare our TOOD with other one-stage detectors on the COCO *test-dev* in Table 5. The models are trained with scale jitter (480-800) and for $2\times$ learning schedule (24

Type	Method	AP	AP ₅₀	AP ₇₅
Anchor-free	ATSS [30]	39.2	57.4	42.2
	TOOD	42.5	59.8	46.4
Anchor-based	ATSS [30]	39.3	57.5	42.8
	TOOD	42.4	59.8	46.1

Table 3. Performance of the complete TOOD (T-head + TAL).

α	β	AP	AP ₅₀	AP ₇₅
0.5	2	42.4	60.0	46.1
0.5	4	42.3	59.3	45.8
0.5	6	41.7	58.1	45.1
1.0	6	42.5	59.8	46.4
1.0	8	42.2	59.0	46.0
1.5	8	41.5	59.4	44.7

Table 4. Analysis of different hyper-parameters for t .

epochs) as the most current method [14]. For a fair comparison, we report results of single model and single testing scale. With ResNet-101 and ResNeXt-101-64 \times 4d, TOOD achieves 46.7 AP and 48.3 AP, outperforming the most current one-stage detectors such as ATSS [30] (by ~ 3 AP) and GFL [14] (by ~ 2 AP). Furthermore, with ResNet-101-DCN and ResNeXt-101-64 \times 4d-DCN, TOOD brings a larger improvement, comparing to other detectors. For example, it obtains an improvement of 2.8 AP (48.3 \rightarrow 51.1 AP) while ATSS has a 2.1 AP (45.6 \rightarrow 47.7 AP) improvement. This validates that TOOD can cooperate with Deformable Convolutional Networks (DCN) [36] more efficiently, by adaptively adjusting the spatial distribution of the learned features for task-alignment. Note that in TOOD, DCN is applied to the first two layers in the head tower. As shown in Table 5, TOOD achieves a new state-of-the-art result with 51.1 AP in one-stage object detection.

Method	Reference	Backbone	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
RetinaNet [16]	ICCV17	ResNet-101	39.1	59.1	42.3	21.9	42.7	50.2
FoveaBox [10]	-	ResNet-101	40.6	60.1	43.5	23.3	45.2	54.5
FCOS w/ imprv [27]	ICCV19	ResNet-101	43.0	61.7	46.3	26.0	46.8	55.0
Noisy Anchor [12]	CVPR20	ResNet-101	41.8	61.1	44.9	23.4	44.9	52.9
MAL [8]	CVPR20	ResNet-101	43.6	62.8	47.1	25.0	46.9	55.8
SAPD [34]	CVPR20	ResNet-101	43.5	63.6	46.5	24.9	46.8	54.6
ATSS [30]	CVPR20	ResNet-101	43.6	62.1	47.4	26.1	47.0	53.6
PAA [9]	ECCV20	ResNet-101	44.8	63.3	48.7	26.5	48.8	56.3
GFL [14]	NeurIPS20	ResNet-101	45.0	63.7	48.9	27.2	48.8	54.5
TOOD (ours)	-	ResNet-101	46.7	64.6	50.7	28.9	49.6	57.0
SAPD [34]	CVPR20	ResNeXt-101-64×4d	45.4	65.6	48.9	27.3	48.7	56.8
ATSS [30]	CVPR20	ResNeXt-101-64×4d	45.6	64.6	49.7	28.5	48.9	55.6
PAA [9]	ECCV20	ResNeXt-101-64×4d	46.6	65.6	50.8	28.8	50.4	57.9
GFL [14]	NeurIPS20	ResNeXt-101-32×4d	46.0	65.1	50.1	28.2	49.6	56.0
TOOD (ours)	-	ResNeXt-101-64×4d	48.3	66.5	52.4	30.7	51.3	58.6
SAPD [34]	CVPR20	ResNet-101-DCN	46.0	65.9	49.6	26.3	49.2	59.6
ATSS [30]	CVPR20	ResNet-101-DCN	46.3	64.7	50.4	27.7	49.8	58.4
PAA [9]	ECCV20	ResNet-101-DCN	47.4	65.7	51.6	27.9	51.3	60.6
GFL [14]	NeurIPS20	ResNet-101-DCN	47.3	66.3	51.4	28.0	51.1	59.2
TOOD (ours)	-	ResNet-101-DCN	49.6	67.4	54.1	30.5	52.7	62.4
SAPD [34]	CVPR20	ResNeXt-101-64×4d-DCN	47.4	67.4	51.1	28.1	50.3	61.5
ATSS [30]	CVPR20	ResNeXt-101-64×4d-DCN	47.7	66.5	51.9	29.7	50.8	59.4
PAA [9]	ECCV20	ResNeXt-101-64×4d-DCN	49.0	67.8	53.3	30.2	52.8	62.2
GFL [14]	NeurIPS20	ResNeXt-101-32×4d-DCN	48.2	67.4	52.6	29.2	51.7	60.2
GFLV2 [13] [†]	CVPR21	ResNeXt-101-32×4d-DCN	49.0	67.6	53.5	29.7	52.4	61.4
OTA [5] [†]	CVPR21	ResNeXt-101-64×4d-DCN	49.2	67.6	53.5	30.0	52.5	62.3
IQDet [19] [†]	CVPR21	ResNeXt-101-64×4d-DCN	49.0	67.5	53.1	30.0	52.3	62.0
VFNet [29] [†]	CVPR21	ResNeXt-101-64×4d-DCN	49.9	68.5	54.3	30.7	53.1	62.8
TOOD (ours)	-	ResNeXt-101-64×4d-DCN	51.1	69.4	55.5	31.9	54.1	63.7

Table 5. Detection results on the COCO *test-dev* set. [†] indicates the concurrent work.

Method	AP	PCC (top-50)	IoU (top-10)	#Correct boxes	#Redundant boxes	#Error boxes
Parallel head + ATSS [30]	39.2	0.408	0.637	30,261	25,428	92,677
T-head + ATSS [30]	41.1	0.440	0.644	30,601	21,838	79,189
Parallel head + TAL	40.3	0.415	0.643	30,506	15,927	72,320
T-head + TAL	42.5	0.452	0.661	30,734	15,242	69,013

Table 6. Analysis for task-alignment of TOOD with backbone ResNet-50.

4.3. Quantitative Analysis for Task-alignment

We quantitatively analyze the effect of the proposed methods on the alignment of two tasks. Without NMS, we calculate a Pearson Correlation Coefficient (PCC) between the rankings [20] of classification and localization by selecting top-50 confident predictions for each instance, and a mean IoU of the top-10 confident predictions, averaged over instances. As shown in Table 6, the mean PCC and IoU are improved by using T-head and TAL. Meanwhile, with NMS, the number of the correct boxes (IoU \geq 0.5) increases while those of the redundant (IoU \geq 0.5) and error boxes (0.1<IoU<0.5) decrease substantially when applying T-head and TAL. The statistics suggest that TOOD is more compatible with NMS, by preserving more correct boxes, and suppressing the redundant/error boxes significantly. At

last, detection performance is improved by 3.3 AP in total. Several detection examples are illustrated in Figure 4.

5. Conclusion

In this work, we illustrate the misalignment between classification and localization in the existing one-stage detectors, and propose TOOD to align the two tasks. In particular, we design a task-aligned head to enhance the interaction of two tasks, and then improve its ability of learning the alignment. Furthermore, a new task-aligned learning strategy is developed by introducing a sample assignment scheme and new loss functions, both of which are computed via an anchor alignment metric. With these improvements, TOOD achieved a 51.1 AP on MS-COCO, surpassing the state-of-the-art one-stage detectors by a large margin.

References

- [1] Yuhang Cao, Kai Chen, Chen Change Loy, and Dahua Lin. Prime sample attention in object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11583–11591, 2020.
- [2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [3] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6569–6578, 2019.
- [4] Chengjian Feng, Yujie Zhong, and Weilin Huang. Exploring classification equilibrium in long-tailed object detection. *arXiv preprint arXiv:2108.07507*, 2021.
- [5] Zheng Ge, Songtao Liu, Zeming Li, Osamu Yoshie, and Jian Sun. Ota: Optimal transport assignment for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 303–312, 2021.
- [6] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 580–587, 2014.
- [7] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yuning Jiang. Acquisition of localization confidence for accurate object detection. In *Proceedings of the European Conference on Computer Vision*, pages 784–799, 2018.
- [8] Wei Ke, Tianliang Zhang, Zeyi Huang, Qixiang Ye, Jianzhuang Liu, and Dong Huang. Multiple anchor learning for visual object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10206–10215, 2020.
- [9] Kang Kim and Hee Seok Lee. Probabilistic anchor assignment with iou prediction for object detection. In *Proceedings of the European Conference on Computer Vision*, 2020.
- [10] Tao Kong, Fuchun Sun, Huaping Liu, Yuning Jiang, Lei Li, and Jianbo Shi. Foveabox: Beyond anchor-based object detection. *IEEE Transactions on Image Processing*, 29:7389–7398, 2020.
- [11] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision*, pages 734–750, 2018.
- [12] Hengduo Li, Zuxuan Wu, Chen Zhu, Caiming Xiong, Richard Socher, and Larry S Davis. Learning from noisy anchors for one-stage object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10588–10597, 2020.
- [13] Xiang Li, Wenhai Wang, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11632–11641, 2021.
- [14] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. In *Advances in Neural Information Processing Systems*, 2020.
- [15] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017.
- [16] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2980–2988, 2017.
- [17] 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 *Proceedings of the European Conference on Computer Vision*, pages 740–755, 2014.
- [18] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Proceedings of the European Conference on Computer Vision*, pages 21–37, 2016.
- [19] Yuchen Ma, Songtao Liu, Zeming Li, and Jian Sun. Iqdet: Instance-wise quality distribution sampling for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1717–1725, 2021.
- [20] Kemal Oksuz, Baris Can Cam, Emre Akbas, and Sinan Kalkan. A ranking-based, balanced loss function unifying classification and localisation in object detection. In *Advances in Neural Information Processing Systems*, 2020.
- [21] Jiangmiao Pang, Kai Chen, Jianping Shi, Huajun Feng, Wanli Ouyang, and Dahua Lin. Libra r-cnn: Towards balanced learning for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 821–830, 2019.
- [22] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.
- [23] 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*, pages 91–99, 2015.
- [24] Hamid Rezaatoughi, 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 658–666, 2019.
- [25] Pierre Sermanet, David Eigen, Xiang Zhang, Michaël Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2013.
- [26] Guanglu Song, Yu Liu, and Xiaogang Wang. Revisiting the sibling head in object detector. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11563–11572, 2020.
- [27] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceed-*

- ings of the *IEEE International Conference on Computer Vision*, pages 9627–9636, 2019.
- [28] Yue Wu, Yinpeng Chen, Lu Yuan, Zicheng Liu, Lijuan Wang, Hongzhi Li, and Yun Fu. Double-head rcnn: Rethinking classification and localization for object detection. *arXiv preprint arXiv:1904.06493*, 2, 2019.
 - [29] Haoyang Zhang, Ying Wang, Feras Dayoub, and Niko Sunderhauf. Varifocalnet: An iou-aware dense object detector. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8514–8523, 2021.
 - [30] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9759–9768, 2020.
 - [31] 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*, pages 147–155, 2019.
 - [32] Yujie Zhong, Zelu Deng, Sheng Guo, Matthew R Scott, and Weilin Huang. Representation sharing for fast object detector search and beyond. In *Proceedings of the European Conference on Computer Vision*, pages 471–487, 2020.
 - [33] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.
 - [34] Chenchen Zhu, Fangyi Chen, Zhiqiang Shen, and Marios Savvides. Soft anchor-point object detection. In *Proceedings of the European Conference on Computer Vision*, 2020.
 - [35] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 840–849, 2019.
 - [36] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9308–9316, 2019.