

RFLA: Gaussian Receptive Field based Label Assignment for Tiny Object Detection

Chang Xu¹, Jinwang Wang², Wen Yang^{1*}, Huai Yu¹, Lei Yu¹, and Gui-Song Xia³

¹ School of Electronic Information, Wuhan University

² Huawei Technologies Co., Ltd.

³ School of Computer Science, Wuhan University

{xuchangeis,yangwen,yuhuai,ly.wd,guisong.xia}@whu.edu.cn
wangjinwang3@huawei.com

Abstract. Detecting tiny objects is one of the main obstacles hindering the development of object detection. The performance of generic object detectors tends to drastically deteriorate on tiny object detection tasks. In this paper, we point out that either box prior in the anchor-based detector or point prior in the anchor-free detector is sub-optimal for tiny objects. Our key observation is that the current anchor-based or anchor-free label assignment paradigms will incur many outlier tiny-sized ground truth samples, leading to detectors imposing less focus on the tiny objects. To this end, we propose a Gaussian Receptive Field based Label Assignment (RFLA) strategy for tiny object detection. Specifically, RFLA first utilizes the prior information that the feature receptive field follows Gaussian distribution. Then, instead of assigning samples with IoU or center sampling strategy, a new Receptive Field Distance (RFD) is proposed to directly measure the similarity between the Gaussian receptive field and ground truth. Considering that the IoU-threshold based and center sampling strategy are skewed to large objects, we further design a Hierarchical Label Assignment (HLA) module based on RFD to achieve balanced learning for tiny objects. Extensive experiments on four datasets demonstrate the effectiveness of the proposed methods. Especially, our approach outperforms the state-of-the-art competitors with 4.0 AP points on the AI-TOD dataset. Codes are available at <https://github.com/Chasel-Tsui/mmdet-rfla>.

Keywords: tiny object detection · Gaussian receptive field · label assignment

1 Introduction

Tiny object, featured by its extremely limited amount of pixels (less than 16×16 pixels defined in AI-TOD [49]), is always *a hard nut to crack* in the computer vision community. Tiny Object Detection (TOD) is one of the most challenging

* Corresponding Author

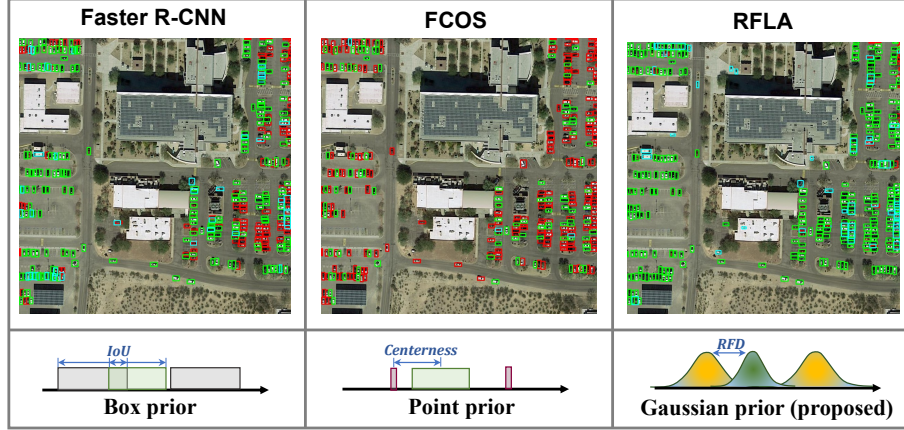


Fig. 1. Comparison between detection results of different label assignment schemes. The detection results are listed in the first row. The green, blue and red boxes denote true positive (TP), false positive (FP) and false negative (FN) predictions. The schematic diagram of different prior is in the second row, where the green region is gt , grey, red and yellow regions denote box, point and Gaussian prior respectively.

tasks, and generic object detectors usually fail to provide satisfactory results on TOD tasks [49,57], resulting from tiny object’s lack of discriminative features. Considering the particularity of tiny objects, several customized TOD benchmarks are proposed (*e.g.* AI-TOD [49], TinyPerson [57], and AI-TOD-v2 [51]), facilitating a series of downstream tasks including driving assistance, traffic management, and maritime rescue. Recently, TOD has gradually become a popular yet challenging direction independent of generic object detection [25,13].

Generic object detectors can be divided into two factions: the anchor-based and the anchor-free paradigms. For anchor-based detectors, prior boxes of discrete locations, scales, and aspect ratios are heuristically preset. Then, label assignment strategies (*e.g.* Max IoU Strategy [39], ATSS [58]) are constructed mainly based on IoU to find the appropriate matching relationship between anchors and ground truth (gt). Anchor-free detectors change the prior from boxes to points. Usually the point prior covered by gt is regarded as a positive sample (with the centerness in FCOS [46]), saving the effort of anchor box fine-tuning.

Despite the outstanding performance of the above two factions on generic object detection tasks, their performance on TOD tasks commonly suffers a drastic drop [49,57]. In this paper, we argue that the current prior box and point along with their corresponding measurement strategies are sub-optimal for tiny objects, which will further hinder the process of label assignment. Specifically, we take the individual prior box and point as instances and rethink them from

the perspective of distribution.

$$p(v|x, y) = \frac{\varepsilon(x - x_1)\varepsilon(x_2 - x)\varepsilon(y - y_1)\varepsilon(y_2 - y)}{(x_2 - x_1)(y_2 - y_1)} \quad (1)$$

where $p(v|x, y)$ is the probability density function of prior information, (x, y) is the location on the image, v is the weight of the corresponding location, $\varepsilon(\cdot)$ is a step function that equals to 1 when the input is larger than 0, otherwise equals to 0. $[(x_1, y_1), (x_2, y_2)]$ is the region of prior information, for anchor-based detectors $x_2 - x_1 = \text{width}$, $y_2 - y_1 = \text{height}$, while for anchor-free detectors $x_2 - x_1 = 1$, $y_2 - y_1 = 1$. The schematic diagram of different prior is shown in the second row of Fig. 1, existing prior information combining with its corresponding measurement strategy has the following problems for tiny objects.

First, the individual box prior and point prior both have a limited prior domain (where $p(v|x, y) > 0$), while existing label assignment metrics are highly dependent on the overlap of domain. In other words, when a particular gt has no overlap with a specific prior, their positional relationship cannot be solved by IoU or centerness. For tiny objects, it is often the case that the gt box has no overlap with almost all anchor boxes (*i.e.* $\text{IoU} = 0$) or does not contain any anchor points [52], leading to tiny objects' lack of positive samples [52]. To this end, heuristics are usually deployed to guarantee more positive samples for tiny objects [58,59]. However, the assigner often fails to compensate positive samples for tiny objects based on the zero-valued IoU or centerness. Therefore, the network will impose less attention on tiny object learning. Details about this point are analyzed in Sec. 4.5. Second, current prior region mainly follows a uniform distribution and treats each location inside the prior region equally ($v = \text{constant}$). However, prior information is essentially leveraged to assist the label assignment or feature point assignment process [58]. In this process, one implicit rule is assigning feature points with appropriate receptive field to gt [39,46]. As theoretically analyzed in previous work [29], when remapping the receptive field of feature point back onto the input image, the effective receptive field is actually Gaussian distributed. The gap between the uniformly distributed prior and the Gaussian distributed receptive field will lead to the mismatch between gt and the receptive field of the feature points assigned to it.

To mitigate the above problems, we introduce a novel prior based on Gaussian distribution and build a Gaussian Receptive Field based Label Assignment (RFLA) strategy that is more conducive to tiny objects. Specifically, we propose to directly measure the similarity between the Gaussian receptive field and gt region with a newly designed Receptive Field Distance (RFD). Leveraging the Gaussian receptive field as prior information can elegantly address the issues incurred by box and point prior. On the one hand, the Gaussian distribution is not step changed. The domain of each individual prior is the entire image, where the weight of each location gradually decays from the center to the periphery with a value higher than 0. It is thus feasible to model the positional relationship between any feature point and any gt on the whole image, making it possible to obtain balanced positive samples for different sized objects. On the other hand,

Gaussian prior can better fit the property of Gaussian effective receptive field, thereby alleviating receptive field mismatch problem, especially for tiny objects.

Moreover, since IoU and RFD are not in the same dimension, directly applying the new metric to the existing threshold-based label assignment structure is not rational. Instead, we introduce to rank the priority of each feature point *w.r.t.* their RFD scores, based on which we further design a Hierarchical Label Assigner (HLA) which progressively alleviates outlier *gt* samples and obtain sufficient training for tiny objects.

The contributions of this paper are summarized as follows:

- (1) We experimentally reveal that current anchor-based and anchor-free detectors exist scale-sample imbalance problem in tiny object label assignment.
- (2) To mitigate the above problem, we introduce a simple but effective Receptive Field-based Label Assignment (RFLA) strategy. The RFLA is easy to replace the standard box and point-based label assignment strategies in mainstream detectors, boosting their performance on TOD.
- (3) Extensive experiments on four datasets validate the performance superiority of our proposed method. The introduced method significantly outperforms the state-of-the-art competitors on the challenging AI-TOD dataset without additional costs in the inference stage.

2 Related Work

2.1 Object Detection

The mainstream object detection methods include anchor-based detectors and anchor-free detectors. Classic anchor-based detectors include Faster R-CNN [39], Cascade R-CNN [4], RetinaNet [24], YOLO series [37,38,3] *etc.* It is commonly believed that one fundamental defect of the anchor-based paradigm is its requirement of tuning *w.r.t.* the specific task [46]. Moreover, the IoU-based label assignment strategy [39] that is built upon anchor boxes also introduces additional hyper-parameters, showing a significant impact on the detection performance.

Anchor-free detectors get rid of the constraints of anchor-boxes, and seek to directly predict objects from center points like FCOS [46] and FoveaBox [19], or seek to predict objects from key-points such as CornerNet [20], Grid R-CNN [28] and RepPoints [55]. The recently published anchor-free detectors mainly follow the end-to-end paradigm, they merely preset a set of boxes without shape or location prior information, and then directly reason about the final predictions, such as DETR [5], Deformable DETR [64], and Sparse R-CNN [44]. Despite the success of the end-to-end paradigm on generic object detection tasks, their performance on TOD tasks require further investigation.

Unlike box and point prior-based detectors, we introduce another prior information based on the receptive field. Combining the Gaussian receptive field and its customized label assignment strategy can significantly alleviate the imbalance problem raised by existing prior and measurement for tiny objects.

2.2 Tiny Object Detection

Most of the existing tiny object detection methods can be roughly grouped into the following four classes: Data augmentation, Multi-scale learning, Customized training strategy for tiny objects, and Feature enhancement strategy.

Data augmentation. A simple yet effective way is to collect more tiny object data. Another way is to use simple data augmentations include rotating, image flipping, and up-sampling. Krisantal *et al.* [18] seek to enhance TOD performance by oversampling images that contain tiny objects and copy-pasting them.

Multi-scale learning. The multi-resolution image pyramid is a basic way of multi-scale learning. To reduce the computation cost, some works [27,23,61] propose to construct feature-level pyramid. After that, lots of methods attempt to further improve FPN, some of them are PANet [26], BiFPN [45], Recursive-FPN [35]. Besides, TridentNet [22] constructs multi-branch detection heads with different receptive fields to generate scale-specific feature maps. Multi-scale learning strategies commonly boost TOD performance with additional computation.

Customized training strategy for tiny objects. Object detectors usually cannot get satisfactory performance on tiny objects and large objects simultaneously. Inspired by this fact, SNIP [42] and SNIPER [43] are designed to selectively train objects within a certain scale range. In addition, Kim *et al.* [17] introduces a Scale-Aware Network (SAN) and maps the features of different spaces onto a scale-invariant subspace, making detectors more robust to scale variation.

Feature enhancement strategy. Some works propose to enhance the feature representation of small objects by super-solution or GAN. PGAN [21] makes the first attempt to apply GAN to small object detection. Moreover, Bai *et al.* [1] propose an MT-GAN which trains an image-level super-resolution model for enhancing the small RoI features. Feature-level super-resolution [32] is proposed to improve small object detection performance for proposal based detectors. Also, there are some other super-solution based methods including [8,2,36].

Most of the methods dedicated to TOD will bring about additional annotation or computation costs. In contrast, our proposed method attempts to push forward TOD from the perspective of label assignment, and our proposed strategy will not bring any additional cost in the inference stage.

2.3 Label Assignment in Object Detection

As revealed by ATSS [58], the essential difference between the anchor-free and anchor-based detector is the way of defining training samples. The selection of positive and negative (*pos/neg*) training samples will notably affect the detector’s performance. Recently, many works have been proposed for better label assignment in generic object detection tasks. FreeAnchor [60] decides positive anchors based on a detection-customized likelihood. PAA [16] proposes to use GMM to model the distribution of anchors and divide *pos/neg* samples based on the center of GMM. OTA [14] models the label assignment process as an optimal transport problem and seeks to solve the optimal assignment strategy. ATSS [58] adaptively adjusting the *pos/neg* samples *w.r.t.* their statistics characteristics.

AutoAssign [62] and IQDet [30] reweight and sample high-quality regions based on the predicted IoU and confidence.

Unlike the above-mentioned general object detection strategies, this paper focuses on the design of prior information and its corresponding label assignment strategy for TOD.

3 Method

3.1 Receptive Field Modelling

One basic principle that mainstream object detectors obey is *dividing and conquering*, namely detecting objects of different scales on the different layers of FPN [7,23]. Specifically, anchor-based detectors tile prior boxes of different scales on different layers of FPN to assist label assignment, and objects of different scales are thus detected on the different layers of FPN. For anchor-free detectors, they group objects in different scale ranges (*e.g.* $[0, 64]$ for P_3) onto different levels of FPN for detection. Despite label assignment strategy varies, one common ground of anchor-based and anchor-free detectors is to assign feature points of an appropriate receptive field to objects of different scales [39,46]. Thus, the receptive field can directly serve as a founded and convincing prior for label assignment without the designing of heuristic anchor box preset or scale grouping.

In this paper, we propose to directly measure the matching degree between the Effective Receptive Field (ERF) and the *gt* region for label assignment, getting rid of the box or point prior that deteriorates TOD. Previous work has pointed out that the ERF can be theoretically derived as Gaussian distribution [29]. In this work, we follow this paradigm and seek to model the ERF of each feature point as Gaussian distribution, and we first derive the Theoretical Receptive Field (TRF) of the n -th layer on a standard convolution neural network [15] by the following formula as tr_n :

$$tr_n = tr_{n-1} + (k_n - 1) \prod_{i=1}^{n-1} s_i \quad (2)$$

where tr_n denotes the TRF of each point on the n -th convolution layer, k_n and s_n denotes the kernel size and stride of the convolution operation on the n -th layer.

As studied in [29], the ERF and TRF have the same center points but the ERF of each feature point only occupies part of the full TRF. Therefore, we use the location of each feature point (x_n, y_n) as the mean vector of a standard 2-D Gaussian distribution. As it is hard to get the precise ERF, we approximate the ERF radius er_n with half the radius of TRF. The square of er_n serves as the co-variance of 2-D Gaussian distribution for a standard square-like convolution kernel. To sum up, we model the range of ERF into a 2-D Gaussian distribution $N_e(\boldsymbol{\mu}_e, \boldsymbol{\Sigma}_e)$ with

$$\boldsymbol{\mu}_e = \begin{bmatrix} x_n \\ y_n \end{bmatrix}, \boldsymbol{\Sigma}_e = \begin{bmatrix} er_n^2 & 0 \\ 0 & er_n^2 \end{bmatrix}. \quad (3)$$

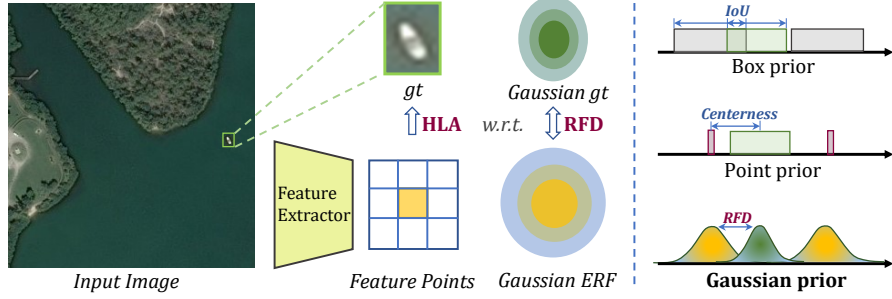


Fig. 2. The process of RFLA. In RFLA, we directly calculate the RFD between Gaussian ERF and gt , then assign labels with HLA.

3.2 Receptive Field Distance

Obtaining the Gaussian ERF, the following key step is to measure the matching degree between the ERF of feature points and a certain gt . As discussed in the introduction, the step-changed uniform distribution is not conducive to tiny objects, it is also necessary to model gt into another distribution.

Observing that the main body of object is aggregated in the center of bounding box [50,48], we also model the gt box (x_g, y_g, w_g, h_g) into a standard 2-D Gaussian distribution $N_g(\mu_g, \Sigma_g)$, where the center point of each annotated box serves as the mean vector of Gaussian and the square of half side length serves as the co-variance matrix, namely,

$$\mu_g = \begin{bmatrix} x_g \\ y_g \end{bmatrix}, \Sigma_g = \begin{bmatrix} \frac{w_g^2}{4} & 0 \\ 0 & \frac{h_g^2}{4} \end{bmatrix}. \quad (4)$$

In this paper, we investigate three types of classic distances between Gaussian distributions as Receptive Field Distance Candidates (RFDC). These distance measurement include Wasserstein distance [34,53], K-L divergence [11,54] and J-S divergence [12]. The J-S divergence between Gaussian distributions has no closed-form solution [12,31], enormous computation will be introduced when approximating its solution, thus, the J-S divergence is not used. Herein, we will first analyze their closed form solutions in our task, then discuss their pros and cons for the TOD task.

Wasserstein distance. The Wasserstein Distance comes from Optimal Transport theory [34]. Given the Gaussian ERF $n_e = N_e(\mu_e, \Sigma_e)$ and Gaussian gt $n_g = N_g(\mu_g, \Sigma_g)$, the 2^{nd} Wasserstein distance can be simplified as Eq. 5 [53].

$$W_2^2(n_e, n_g) = \left\| \left([x_n, y_n, er_n, er_n]^T, \left[x_g, y_g, \frac{w_g}{2}, \frac{h_g}{2} \right]^T \right) \right\|_2^2. \quad (5)$$

The main advantage of Wasserstein Distance is that it can measure two non-overlapping distributions [34]. It is always the case that the gt box has no overlap with most prior box and points, and the assigner fails to rank the priority of these

candidates to a certain gt . Thus, it is easy to say the characteristic of Wasserstein distance is conducive to TOD, which can consistently reflect the matching degree between all feature points and a certain gt box, making the assigner feasible to compensate more positive samples for tiny objects according to a rational priority. However, the Wasserstein distance is not scale invariant and might be sub-optimal when the dataset contains objects of large-scale variance [54].

Kullback-Leibler divergence. Kullback-Leibler Divergence (KLD) is a classic statistical distance which measures how one probability distribution is different from another. KLD between two Gaussian distributions also has a closed form solution, the KLD between ERF n_e and gt region n_g is as follows:

$$D_{KL}(n_e||n_g) = \frac{1}{2}(\text{tr}(\Sigma_g^{-1}\Sigma_e) + (\mu_g - \mu_e)^\top \Sigma_g^{-1}(\mu_g - \mu_e) + \ln \frac{|\Sigma_g|}{|\Sigma_e|} - 2), \quad (6)$$

Eq. 6 can be further simplified as:

$$D_{KL}(n_e||n_g) = \frac{er_n^2}{8w_g^2} + \frac{er_n^2}{8h_g^2} + \frac{2(x_n - x_g)^2}{w_g^2} + \frac{2(y_n - y_g)^2}{h_g^2} + \ln \frac{2w_g}{er_n} + \ln \frac{2h_g}{er_n} - 1. \quad (7)$$

As demonstrated by the work [54], KLD has the property of scale invariance between two 2-D Gaussian distributions, and the scale invariance is crucial for the detection [56]. While the main disadvantage of KLD is that it cannot consistently reflect the distance between two distributions when their overlap is negligible. Hence, the KLD between ERF and gt is selected as another RFDC in this paper.

In summary, we investigate three classic ways of probability distribution measurement, while Wasserstein distance and KLD are selected as RFDC. Then, we apply a non-linear transformation into RFDC and get the RFD with a normalized value range between (0, 1) as follows:

$$\text{RFD} = \frac{1}{1 + \text{RFDC}} \quad (8)$$

3.3 Hierarchical Label Assignment

Some anchor-based detectors set a threshold based on IoU to decide *pos/neg* samples [39,24,4], while anchor-free detectors mainly divide *pos/neg* samples by the spatial location between point prior and gt region. Since tiny objects are usually unwelcome in both threshold-based and gt region-based strategies, we propose to hierarchically assign labels to tiny objects by score ranking.

To guarantee that the positional relationship between any feature point and any gt can be solved, the proposed Hierarchical Label Assignment (HLA) strategy is built on the proposed RFD. Before assigning, an RFD score matrix between feature points and gt is computed based on the above method. In the first stage, we rank each feature point to its RFD score with a certain gt . Then, positive labels are assigned to feature points with top k RFD scores with a certain gt . Finally, we get the assigning result r_1 and the corresponding mask m of features that have been assigned, where m is binary-valued (0/1). In the second stage, to improve the overall recall and alleviate outliers, we slight decay the effective

radius er_n by multiplying a stage factor β , then repeat the above ranking strategy and supplement one positive sample to each gt , getting the assigning result r_2 . We obtain the final assigning result r by the following rule:

$$r = r_1m + r_2(1 - m), \quad (9)$$

where the mask operation m is taken to avoid introducing too many low-quality samples for those gt which have already been assigned with sufficient samples. Not that the occluded sample will be assigned to the smaller gt . Combining the RFD with the HLA strategy, we can get the complete Receptive Field based Label Assignment (RFLA) strategy for TOD.

3.4 Application to Detectors

The proposed RFLA strategy can easily be applied to anchor-based and anchor-free frameworks. Without losing generality, we take the classic Faster R-CNN [39] and FCOS [46] as examples. Concretely, for Faster R-CNN, RFLA can be used to replace the standard anchor tiling and MaxIoU anchor assigning process. For FCOS, we remove the constraint of limiting feature points inside gt box because the tiny box only covers an extremely limited region which commonly holds much fewer feature points than the large object. Then, it is easy to replace the point based assigning with RFLA for balanced learning. Note that we modify the centerness [46] loss into the following formula to avoid gradient explosion:

$$\text{centerness}^* = \sqrt{\frac{\varepsilon[\min(l^*, r^*)] \min(l^*, r^*) + c}{\max(l^*, r^*)} \times \frac{\varepsilon[\min(t^*, b^*)] \min(t^*, b^*) + c}{\max(t^*, b^*)}}, \quad (10)$$

where l^*, t^*, r^*, b^* are regression targets defined in FCOS, $\varepsilon(\cdot)$ is a step function same as that in Eq. 1, c is a factor set to 0.01 to avoid gradient vanishing problem when the center point of regression target is outside the gt box. In the following part, extensive experiments will show RFLA’s outstanding robustness to TOD.

4 Experiment

4.1 Dataset

Experiments are conducted on four datasets. The main experiments are performed on the challenging AI-TOD [49] dataset, which has the smallest average absolute object size of 12.8 pixels and contains 28,036 images. Furthermore, we test the proposed method on TinyPerson [57], VisDrone2019 [10] and DOTA-v2.0 [9]. Note that the selected datasets all contain a great amount of tiny objects (smaller than 16×16 pixels).

4.2 Experiment settings

All the experiments are conducted on a computer with 1 NVIDIA RTX 3090 GPU, and the model training is based on PyTorch [33], the core codes are built

Table 1. Comparison between different ways of measuring receptive-field distance. Note that in this experiment, all distances are built on HLA.

Distance	AP	AP _{0.5}	AP _{vt}	AP _t
GIoU	17.9	45.1	5.5	16.7
WD	21.1	52.2	6.6	21.5
KLD	21.1	51.6	9.5	21.2

Table 2. Influence of different designs. Note that RFD denotes only using the first stage of HLA, HLA means using all stages of HLA.

RFD	HLA	AP	AP _{0.5}	AP _{vt}	AP _t
		11.1	26.3	0.0	7.2
✓		20.7	50.6	7.6	20.5
✓	✓	21.1	51.6	9.5	21.2

Table 3. Influence of the stage factor β in the Hierarchical Label Assigner (HLA). The KLD is used as RFD.

β	AP	AP _{0.5}	AP _{vt}	AP _t
0.95	21.1	51.1	8.4	21.1
0.9	21.1	51.6	9.5	21.2
0.85	20.8	51.4	7.4	21.0
0.8	19.7	49.0	5.6	19.2

Table 4. Gaussian anchor and receptive anchor. GA means Gaussian Anchor and RA means Receptive Anchor.

Method	AP	AP _{0.5}	AP _{vt}	AP _t
GA	19.6	49.2	8.2	19.7
RA	18.9	47.5	6.1	19.1
baseline	11.1	26.3	0.0	7.2

upon MMDetection [6]. The ImageNet [41] pre-trained model is used as the backbone. All models are trained using the Stochastic Gradient Descent (SGD) optimizer for 12 epochs with 0.9 momenta, 0.0001 weight decay, and 2 batch size. The initial learning rate is set to 0.005 and decays at the 8th and 11th epochs. Besides, the number of RPN proposals is set to 3000. In the inference stage, the confidence score is set to 0.05 to filter out background bounding boxes, and the NMS IoU threshold is set to 0.5 with top 3000 bounding boxes. All the other parameters are set the same as default in MMDetection. The evaluation metric follows AI-TOD benchmark [49] except for experiments on TinyPerson. The above parameters are used in all experiments unless specified otherwise.

4.3 Ablation study

Effectiveness of different RFD. In this part, we respectively apply Wasserstein distance (WD) and Kullback-Leibler divergence (KLD) to measure the distance between Gaussian ERF and *gt* region as discussed in Sec. 3.2. We also test the performance of GIoU [40] by setting the prior as ERF sized box. Note that all experiments are conducted based on Faster R-CNN w/ HLA since RFD and HLA are interdependent. As shown in Tab. 1, it can be seen that GIoU is inferior to RFD since it fails to distinguish the locations of mutually inclusive boxes, while the performance of WD and KLD are comparable. The KLD surpasses WD in AP_{vt}, while slightly lower than WD under AP_t metric. As mentioned in Sec. 3.2, the KLD is scale-invariant, thus it is better for *very tiny* objects. Note that in the following experiments, we use KLD as the default RFD.

Effectiveness of individual component. The core designs in this paper are interdependent, whilst they can be separated into two parts: the Hierarchical Label Assignment (HLA) strategy and the Receptive Field Distance (RFD) built

Table 5. Main results on AI-TOD. Note that models are trained on the **trainval set** and validated on the **test set**. Note that FCOS* means using P2-P6 of FPN.

Method	Backbone	AP	AP _{0.5}	AP _{0.75}	AP _{vt}	AP _t	AP _s	AP _m
TridentNet [22]	ResNet-50	7.5	20.9	3.6	1.0	5.8	12.6	14.0
Faster R-CNN [39]	ResNet-50	11.1	26.3	7.6	0.0	7.2	23.3	33.6
Cascade RPN [47]	ResNet-50	13.3	33.5	7.8	3.9	12.9	18.1	26.3
Cascade R-CNN [4]	ResNet-50	13.8	30.8	10.5	0.0	10.6	25.5	26.6
DetectoRS [35]	ResNet-50	14.8	32.8	11.4	0.0	10.8	28.3	28.0
DotD [52]	ResNet-50	16.1	39.2	10.6	8.3	17.6	18.1	22.1
DetectoRS w/ NWD [48]	ResNet-50	20.8	49.3	14.3	6.4	19.7	29.6	38.3
SSD-512 [27]	ResNet-50	7.0	21.7	2.8	1.0	4.7	11.5	13.5
RetinaNet [24]	ResNet-50	8.7	22.3	4.8	2.4	8.9	12.2	16.0
PAA [16]	ResNet-50	10.0	26.5	6.7	3.5	10.5	13.1	22.1
ATSS [58]	ResNet-50	12.8	30.6	8.5	1.9	11.6	19.5	29.2
RepPoints [55]	ResNet-50	9.2	23.6	5.3	2.5	9.2	12.9	14.4
OTA [14]	ResNet-50	10.4	24.3	7.2	2.5	11.9	15.7	20.9
AutoAssign [62]	ResNet-50	12.2	32.0	6.8	3.4	13.7	16.0	19.1
FCOS [46]	ResNet-50	12.6	30.4	8.1	2.3	12.2	17.2	25.0
M-CenterNet [49]	DLA-34	14.5	40.7	6.4	6.1	15.0	19.4	20.4
FCOS*	ResNet-50	15.4	36.3	10.9	6.0	17.6	18.5	20.7
RetinaNet w/ RFLA	ResNet-50	9.1	23.1	5.2	4.1	10.5	10.5	12.3
AutoAssign w/ RFLA	ResNet-50	14.2	37.8	6.9	6.4	14.9	17.4	21.8
FCOS* w/ RFLA	ResNet-50	16.3	39.1	11.3	7.3	18.5	19.8	21.8
Faster R-CNN w/ RFLA	ResNet-50	21.1	51.6	13.1	9.5	21.2	26.1	31.5
Cascade R-CNN w/ RFLA	ResNet-50	22.1	51.6	15.6	8.2	22.0	27.3	35.2
DetectoRS w/ RFLA	ResNet-50	24.8	55.2	18.5	9.3	24.8	30.3	38.2

upon the HLA. Note that the validation of RFD requires using the first stage of HLA, we do not assign labels based on the threshold of RFD since the original threshold in the baseline detector is designed for IoU, which is not in the same dimension as RFD. We progressively apply RFD and HLA into the Faster R-CNN. Results are listed in Tab. 2, AP improves progressively, the individual effectiveness is thus verified. When switching the IoU-based assignment strategy to the RFD-based one, a notable improvement of 9.6 AP points is obtained. This can be explained that the limited domain of box prior leads to the remarkably low IoU between anchors and *gt*, many *gt* fail to match with any anchor. With Gaussian prior and RFD, the assigner is capable of measuring the priority (RFD score) of all feature points to a particular *gt*. Thus, even though the *gt* has no overlap with any box prior, some positive samples can be compensated for the *gt* with a rational receptive field, leading to the sufficient training of tiny objects.

Performance of different decay factor β . As in Sec. 3.3, in the HLA, we design a stage factor β to the ERF for mitigating the outlier effect. In Tab. 3, we keep all other parameters fixed and experimentally show that 0.9 is the best choice. Setting β to a lower value will introduce too many low-quality samples.

Performance of different k . In the HLA, the hyper-parameter k is designed to adjust the number of positive samples assigned to each instance. Herein, we keep all other parameters fixed and set k from 1 to 4. Their performance is 20.7, 21.1,

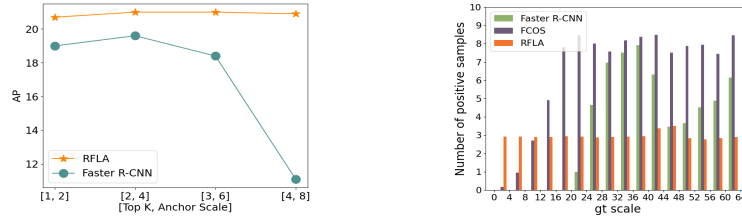


Fig. 3. Comparison between top k in Fig. 4. Scale-sample imbalance problems of different detectors. Base anchor scale tuning in Faster R-CNN baseline. is set to 8 for Faster R-CNN.

21.1, and 20.9 AP, respectively. When setting k to 2 or 3, the best performance can be attained. Thus 3 is recommended as the default setting. Moreover, the AP only waves by a small margin under the tested k . Moreover, we compare the AP under different k with the result of anchor size tuning, as shown in Fig. 3. It is easy to find that the performance of box prior based detector is quite sensitive to the box size on TOD tasks, while in our design, the performance is quite robust to the chosen of k , which consistently keeps a high level over box prior.

Gaussian anchor and receptive anchor. We directly model the anchors into Gaussian distributions, calculate the RFD score between gt , and then assign labels with HLA. Results are shown in Tab. 4. The results show a great advantage of Gaussian prior and its combination with HLA. The Gaussian prior has a broader domain, making sample compensation possible. In addition, we change the anchor scale to the ERF scale, then assign labels with the MaxIoU strategy. The improvement over baseline further indicates the sensitivity of the box prior to detection performance for TOD. It also reveals that the current anchor will potentially introduce the receptive field mismatch problem for tiny objects.

4.4 Main result

We compare our method with other state-of-the-art detectors on AI-TOD benchmark [49]. As shown in Tab 5, DetectoRS w/ RFLA achieves 24.8 AP, which has 4.0 AP above the state-of-the-art competitors. Notably, the improvement of RFLA with multi-stage anchor-based detectors is particularly significant. We think it mainly owes to the multi-stage detectors' mechanism of *looking and thinking twice*. In the first stage, the combination of the proposed RFLA with RPN can improve the recall of tiny objects to a great extent. In the second stage, proposals are refined for precise location and classification. Besides, improvements can also be expected on one-stage anchor-based or anchor-free detectors, and the improvement in AP_{vt} is more obvious, 1.7 points for RetinaNet and 1.3 points for FCOS*. The gap between one-stage and multi-stage detectors is common for TOD [57,63,10]. It mainly results from the lack of multi-stage regression, which is crucial for TOD.

Table 6. Results on TinyPerson.

Method	AP_{50}^{tiny}	AP_{50}^{tiny1}	AP_{50}^{tiny2}	AP_{50}^{tiny3}	AP_{50}^{small}	AP_{25}^{tiny}	AP_{75}^{tiny}
FCOS	23.4	9.8	22.7	34.7	39.2	43.8	1.7
Faster R-CNN	48.7	32.3	54.5	58.8	64.7	68.9	6.0
FCOS w/ RFLA	$26.5^{+3.1}$	$10.0^{+0.2}$	$24.4^{+1.7}$	$40.6^{+5.9}$	$50.5^{+11.3}$	$50.0^{+6.2}$	$2.9^{+1.2}$
Faster R-CNN w/ RFLA	$50.1^{+1.4}$	$32.8^{+0.5}$	$55.6^{+1.1}$	$60.6^{+1.8}$	$65.3^{+0.6}$	$69.9^{+1.0}$	$5.9^{-0.1}$

Table 7. Results on VisDrone2019. The **Table 8.** Results on DOTA-v2.0. The **train**, **val** sets are used for training and validation. FR, DR denote Faster R-CNN, validation. FR, DR denote Faster R-CNN, DetectoRS, * means *with RFLA*. DetectoRS, * means *with RFLA*.

Method	AP	$AP_{0.5}$	AP_{vt}	AP_t	Method	AP	$AP_{0.5}$	AP_{vt}	AP_t
FCOS	14.1	25.5	0.1	2.1	FCOS	31.8	55.4	0.3	4.0
FR	22.3	38.0	0.1	6.2	FR	35.6	59.5	0.0	7.1
DR	25.7	41.7	0.5	7.6	DR	40.8	62.6	0.0	7.0
FCOS*	$15.1^{+1.0}$	$27.3^{+1.8}$	$0.4^{+0.3}$	$3.8^{+1.7}$	FCOS*	$32.1^{+0.3}$	$55.6^{+0.2}$	$0.7^{+0.4}$	$6.8^{+2.8}$
FR*	$23.4^{+1.1}$	$41.4^{+3.4}$	$4.8^{+4.7}$	$11.7^{+5.5}$	FR*	$36.3^{+0.7}$	$61.5^{+2.0}$	$1.9^{+1.9}$	$11.7^{+4.6}$
DR*	$27.4^{+1.7}$	$45.3^{+3.6}$	$4.5^{+4.0}$	$12.9^{+5.3}$	DR*	$41.3^{+0.5}$	$64.2^{+1.6}$	$2.1^{+2.1}$	$10.8^{+3.8}$

4.5 Analysis

We conduct a group of analysis experiments to delve into different prior designs and assigners for tiny objects. In the first step, we respectively deploy the way of prior tiling in Faster R-CNN [39], FCOS [46] and RFLA. In the second step, we randomly generate different gt in different locations of the image and simulate the process of label assignment for statistics. Concretely, the gt scales are randomly picked from 0 to 64. After that, we divide the scale range into 16 intervals, as shown in Fig. 4, and calculate the average number of positive samples assigned to each gt in different scale ranges. Observations in Fig. 4 indicate severe scale-sample imbalance problems for existing detectors. For anchor-based detectors, objects in the tiny scale and the interval between box scales become outliers. Anchor-free detectors somewhat alleviate this problem. However, tiny objects are still outliers since tiny object covers an extremely limited region. The number of prior points inside gt is much smaller than that of large objects. The scale-sample imbalance problem will mislead the network towards unbalanced optimization, where less focus is imposed on outlier samples. In contrast, the number of positive samples assigned to gt in different scale ranges is greatly reconciled with RFLA, achieving a balanced optimization for tiny objects.

4.6 Experiment on more datasets

We conduct experiments on another TOD dataset TinyPerson [57]. The dataset setting and evaluation all follow TinyPerson benchmark [57]. The results are in Tab. 6, 3.1 and 1.4 AP_{50}^{tiny} improvement can be obtained when applying RFLA

into FCOS and Faster R-CNN. We also tested the RFLA on AI-TOD-v2 [51], whose performance is listed in GitHub. We further verify the effectiveness of RFLA on datasets which simultaneously hold a large scale variance and contain many tiny objects (*i.e.* VisDrone2019, DOTA-v2.0). The results are in Tab. 7 and Tab. 8. The improvement of AP_{vt} and AP_t is quite obvious for both datasets. The consistent improvement on various datasets indicates RFLA’s generality. Finally, visualization results on the AI-TOD dataset are shown in Fig. 5. When applying RFLA into Faster R-CNN, FN predictions can be greatly eliminated.

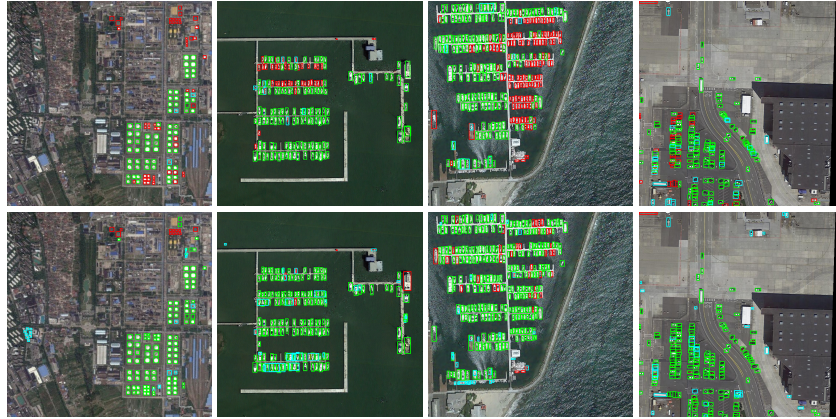


Fig. 5. Visualization results on AI-TOD. The first row is the result of Faster R-CNN and the second row is the result of Faster R-CNN w/ RFLA.

5 Conclusion

In this paper, we point out that box and point prior do not work well for TOD, leading to scale-sample imbalance problems when assigning labels. To this end, we introduce a new Gaussian receptive field prior. Then, we further design a new Receptive Field Distance (RFD), which measures the similarity between ERF and gt to conquer the shortages of IoU and centerness on TOD. The RFD works with the HLA strategy, obtaining balanced learning for tiny objects. Experiments on four datasets show the superiority and robustness of the RFLA.

Acknowledgement

This work was partly supported by the Fundamental Research Funds for the Central Universities under Grant 2042022kf1010, and the National Natural Science Foundation of China under Grant 61771351 and 61871297. The numerical calculations were conducted on the supercomputing system in the Supercomputing Center, Wuhan University.

References

1. Bai, Y., Zhang, Y., Ding, M., Ghanem, B.: Sod-mtgan: Small object detection via multi-task generative adversarial network. In: European Conference on Computer Vision. pp. 206–221. Springer (2018)
2. Bashir, S.M.A., Wang, Y.: Small object detection in remote sensing images with residual feature aggregation-based super-resolution and object detector network. Remote Sensing **13**(9), 1854 (2021)
3. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4: Optimal speed and accuracy of object detection. CoRR **abs/arXiv:2004.10934** (2020), <https://arxiv.org/abs/2004.10934>
4. Cai, Z., Vas., N.: Cascade r-cnn: Delving into high quality object detection. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 6154–6162 (2018)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020)
6. Chen, K., Wang, J., Pang, J., *et al.*: MMDetection: Open mmlab detection toolbox and benchmark. CoRR **abs/arXiv:1906.07155** (2019), <https://arxiv.org/abs/1906.07155>
7. Chen, Q., Wang, Y., Yang, T., Zhang, X., Cheng, J., Sun, J.: You only look one-level feature. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 13039–13048 (2021)
8. Courtrai, L., Pham, M.T., Lefèvre, S.: Small object detection in remote sensing images based on super-resolution with auxiliary generative adversarial networks. Remote Sensing **12**(19), 3152 (2020)
9. Ding, J., Xue, N., Xia, G.S., Bai, X., Yang, W., Yang, M.Y., Belongie, S., Luo, J., Datcu, M., Pelillo, M., *et al.*: Object detection in aerial images: A large-scale benchmark and challenges. IEEE Transactions on Pattern Analysis and Machine Intelligence p. in press (2021)
10. Du, D., Zhu, P., Wen, L., *et al.*: Visdrone-det2019: The vision meets drone object detection in image challenge results. In: IEEE International Conference on Computer Vision Workshops. pp. 213–226 (2019)
11. Duchi, J.: Derivations for linear algebra and optimization. Berkeley, California **3**(1), 2325–5870 (2007)
12. Endres, D.M., Schindelin, J.E.: A new metric for probability distributions. IEEE Transactions on Information Theory **49**(7), 1858–1860 (2003)
13. Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. International Journal of Computer Vision **111**(1), 98–136 (2015)
14. Ge, Z., Liu, S., Li, Z., Yoshie, O., Sun, J.: Ota: Optimal transport assignment for object detection. IEEE Conference on Computer Vision and Pattern Recognition (2021)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
16. Kim, K., Lee, H.S.: Probabilistic anchor assignment with iou prediction for object detection. In: European Conference on Computer Vision. pp. 355–371. Springer (2020)

17. Kim, Y., Kang, B.N., Kim, D.: San: Learning relationship between convolutional features for multi-scale object detection. In: European Conference on Computer Vision. pp. 316–331. Springer (September 2018)
18. Kisantal, M., Wojna, Z., Murawski, J., Naruniec, J., Cho, K.: Augmentation for small object detection. arXiv preprint arXiv:1902.07296 (2019)
19. Kong, T., Sun, F., Liu, H., Jiang, Y., Li, L., Shi, J.: Foveabox: Beyond anchor-based object detection. *IEEE Transactions on Image Processing* **29**, 7389–7398 (2020)
20. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: European Conference on Computer Vision. pp. 734–750. Springer (2018)
21. Li, J., Liang, X., Wei, Y., Xu, T., Feng, J., Yan, S.: Perceptual generative adversarial networks for small object detection. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 1222–1230 (2017)
22. Li, Y., Chen, Y., Wang, N., Zhang, Z.: Scale-aware trident networks for object detection. In: IEEE International Conference on Computer Vision. pp. 6054–6063 (2019)
23. Lin, T.Y., Dollar, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 2117–2125 (2017)
24. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. In: IEEE International Conference on Computer Vision. pp. 2980–2988 (2017)
25. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European Conference on Computer Vision. pp. 740–755. Springer (2014)
26. Liu, S., Qi, L., Qin, H., Shi, J., Jia, J.: Path aggregation network for instance segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 8759–8768 (2018)
27. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: Single shot multibox detector. In: European Conference on Computer Vision. pp. 21–37. Springer (2016)
28. Lu, X., Li, B., Yue, Y., Li, Q., Yan, J.: Grid r-cnn. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 7363–7372 (2019)
29. Luo, W., Li, Y., Urtasun, R., Zemel, R.: Understanding the effective receptive field in deep convolutional neural networks. *Advances in Neural Information Processing Systems* **29** (2016)
30. Ma, Y., Liu, S., Li, Z., Sun, J.: Iqdet: Instance-wise quality distribution sampling for object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1717–1725 (2021)
31. Nielsen, F.: On the jensen–shannon symmetrization of distances relying on abstract means. *Entropy* **21**(5) (2019). <https://doi.org/10.3390/e21050485>
32. Noh, J., Bae, W., Lee, W., Seo, J., Kim, G.: Better to follow, follow to be better: Towards precise supervision of feature super-resolution for small object detection. In: IEEE International Conference on Computer Vision. pp. 9725–9734 (2019)
33. Paszke, A., Gross, S., Massa, F., Lerer, A., et al.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems. pp. 8024–8035 (2019)
34. Peyré, G., Cuturi, M., et al.: Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning* **11**(5-6), 355–607 (2019)

35. Qiao, S., Chen, L.C., Yuille, A.: Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In: IEEE Conference on Computer Vision and Pattern Recognition (2021)
36. Rabbi, J., Ray, N., Schubert, M., Chowdhury, S., Chao, D.: Small-object detection in remote sensing images with end-to-end edge-enhanced gan and object detector network. *Remote Sensing* **12**(9), 1432 (2020)
37. Redmon, J., Farhadi, A.: YOLO9000: Better, faster, stronger. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 7263–7271 (2017)
38. Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018)
39. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems*. pp. 91–99 (2015)
40. Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 658–666 (2019)
41. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015)
42. Singh, B., Davis, L.S.: An analysis of scale invariance in object detection snip. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 3578–3587 (2018)
43. Singh, B., Najibi, M., Davis, L.S.: Sniper: Efficient multi-scale training. In: *Advances in Neural Information Processing Systems*. pp. 9310–9320 (2018)
44. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., Luo, P.: Sparse r-cnn: End-to-end object detection with learnable proposals. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 14454–14463 (2021)
45. Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 10781–10790 (2020)
46. Tian, Z., Shen, C., Chen, H., He, T.: FCOS: Fully convolutional one-stage object detection. In: IEEE International Conference on Computer Vision. pp. 9627–9636 (2019)
47. Vu, T., Jang, H., Pham, T.X., Yoo, C.: Cascade rpn: Delving into high-quality region proposal network with adaptive convolution **32**, 1432–1442 (2019)
48. Wang, J., Xu, C., Yang, W., Yu, L.: A normalized gaussian wasserstein distance for tiny object detection. *arXiv preprint arXiv:2110.13389* (2021)
49. Wang, J., Yang, W., Guo, H., Zhang, R., Xia, G.S.: Tiny object detection in aerial images. In: International Conference on Pattern Recognition. pp. 3791–3798 (2021)
50. Wang, J., Yang, W., Li, H.c., Zhang, H., Xia, G.S.: Learning center probability map for detecting objects in aerial images. *IEEE Transactions on Geoscience and Remote Sensing* **59**(5), 4307–4323 (2021)
51. Xu, C., Wang, J., Yang, W., Yu, H., Yu, L., Xia, G.S.: Detecting tiny objects in aerial images: A normalized wasserstein distance and a new benchmark. In: *ISPRS Journal of Photogrammetry and Remote Sensing*. vol. 190, pp. 79–93 (2022)
52. Xu, C., Wang, J., Yang, W., Yu, L.: Dot distance for tiny object detection in aerial images. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 1192–1201 (2021)

53. Yang, X., Yan, J., Ming, Q., Wang, W., Zhang, X., Tian, Q.: Rethinking rotated object detection with gaussian wasserstein distance loss. In: International Conference on Machine Learning. vol. 139, pp. 11830–11841 (2021)
54. Yang, X., Yang, X., Yang, J., Ming, Q., Wang, W., Tian, Q., Yan, J.: Learning high-precision bounding box for rotated object detection via kullback-leibler divergence. *Advances in Neural Information Processing Systems* **34** (2021)
55. Yang, Z., Liu, S., Hu, H., Wang, L., Lin, S.: Reppoints: Point set representation for object detection. In: IEEE International Conference on Computer Vision. pp. 9657–9666 (2019)
56. Yu, J., Jiang, Y., Wang, Z., Cao, Z., Huang, T.: Unitbox: An advanced object detection network. pp. 516–520 (2016)
57. Yu, X., Gong, Y., Jiang, N., Ye, Q., Han, Z.: Scale match for tiny person detection. In: IEEE Workshops on Applications of Computer Vision. pp. 1257–1265 (2020)
58. Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 9759–9768 (2020)
59. Zhang, S., Zhu, X., Lei, Z., Shi, H., Wang, X., Li, S.Z.: S3fd: Single shot scale-invariant face detector. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 192–201 (2017)
60. Zhang, X., Wan, F., Liu, C., Ji, X., Ye, Q.: Learning to match anchors for visual object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
61. Zhao, Q., Sheng, T., Wang, Y., Tang, Z., Chen, Y., Cai, L., Ling, H.: M2det: A single-shot object detector based on multi-level feature pyramid network. In: AAAI Conference on Artificial Intelligence. pp. 9259–9266 (2019)
62. Zhu, B., Wang, J., Jiang, Z., Zong, F., Liu, S., Li, Z., Sun, J.: Autoassign: Differentiable label assignment for dense object detection. *arXiv preprint arXiv:2007.03496* (2020)
63. Zhu, P., Wen, L., Du, D., Bian, X., Ling, H., Hu, Q., Nie, Q., Cheng, H., Liu, C., Liu, X., et al.: Visdrone-det2018: The vision meets drone object detection in image challenge results. In: European Conference on Computer Vision Workshops. pp. 437–468. Springer (2018)
64. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: International Conference on Learning Representations (2021)