

# Oriented RepPoints for Aerial Object Detection

Wentong Li<sup>1</sup>, Yijie Chen<sup>1</sup>, Kaixuan Hu<sup>2</sup>, Jianke Zhu<sup>1,3\*</sup>

<sup>1</sup>Zhejiang University    <sup>2</sup>University of Electronic Science and Technology of China

<sup>3</sup>Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies

{liwentong, chen\_yj, jkzhu}@zju.edu.cn, 201922100522@std.uestc.edu.cn

## Abstract

In contrast to the generic object, aerial targets are often non-axis aligned with arbitrary orientations having the cluttered surroundings. Unlike the mainstreamed approaches regressing the bounding box orientations, this paper proposes an effective adaptive points learning approach to aerial object detection by taking advantage of the adaptive points representation, which is able to capture the geometric information of the arbitrary-oriented instances. To this end, three oriented conversion functions are presented to facilitate the classification and localization with accurate orientation. Moreover, we propose an effective quality assessment and sample assignment scheme for adaptive points learning toward choosing the representative oriented reppoints samples during training, which is able to capture the non-axis aligned features from adjacent objects or background noises. A spatial constraint is introduced to penalize the outlier points for robust adaptive learning. Experimental results on four challenging aerial datasets including DOTA, HRSC2016, UCAS-AOD and DIOR-R, demonstrate the efficacy of our proposed approach. The source code is available at: <https://github.com/LiWentong/OrientedRepPoints>.

## 1. Introduction

Being an important computer vision task [4, 34, 37], aerial object detection has recently attracted increasing attention, which plays the significant role in the remote image understanding. Different from the generic object detection, aerial target localization has its own difficulties, including non-axis aligned objects with arbitrary orientations [3, 9, 24] and densely packed distribution with complex context [7, 37, 45].

The mainstreamed approaches typically treat aerial object detection as a problem of rotated object localization [3, 8, 9, 24, 42, 45]. Among them, the direct angle-based orienta-

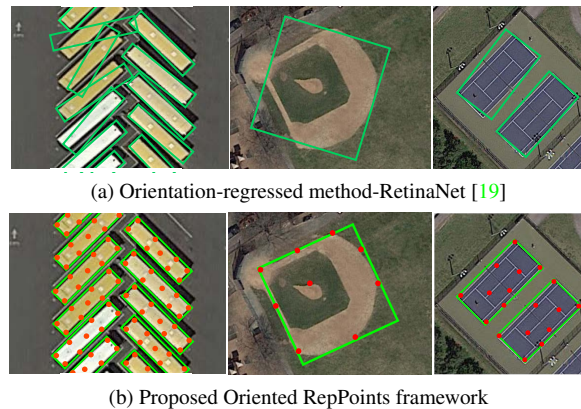


Figure 1. (a) denotes the common baseline-RetinaNet [19] of orientation regression-based methods, (b) is the baseline method of our Oriented RepPoints. Comparing to the direct orientation regression, our approach can estimate more accurate orientations by learning the adaptive points that are marked as red.

tion regression methods dominate this research area, which are derived from the general vanilla detectors [18, 19, 28, 53] with the extra orientation parameter. Although having achieved promising performance, the direct orientation prediction still has some issues including the discontinuity of loss and regression inconsistency [25, 40, 41, 43]. This is mainly due to the bounded periodic nature of the angular orientation and the orientation definition of the rotated bounding box. Despite of their attractive localization results, the orientation regression-based detectors may not accurately predict the orientations, as shown in Fig. 1-(a).

To effectively address the above issues, we revisit the representation for aerial objects in order to avoid the sensitive orientation estimation. As a fine-grained object representation, point set shows the great potential to capture the key semantic features in conventional generic detector like RepPoints [46]. However, its simple conversion functions only produce the upright-horizontal bounding boxes, which cannot precisely estimate the aerial objects' orientations. Moreover, RepPoints only regresses the key points

\*Corresponding author is Jianke Zhu.

according to the semantic features while ignoring to effectively measure the quality of learned points. This may lead to the inferior performance for the non-axis aligned objects with dense-distribution and complex scene in aerial images.

In this work, we proposed an oriented object detector for aerial images, named *Oriented RepPoints*, which introduces the adaptive points representation for diverse orientations, shapes and poses. In contrast to conventional orientation regression approach, our proposed method not only achieves the precise aerial detection with accurate orientation, but also captures the underlying geometric structure of the arbitrary-oriented aerial instances, as shown in Fig. 1. Specifically, the initial adaptive points are generated from the center point, which are further refined to adapt for the aerial object. To obtain the oriented bounding box, three oriented conversion functions are presented according to the layouts of the learned points. Moreover, an effective adaptive points assessment and assignment (APAA) scheme is proposed for point set learning, which measures the quality of oriented reppoints not only from the classification, localization but also from their orientation and point-wise feature correlation during training. Such scheme enables the detector to capture the non-axis aligned features from adjacent objects or background noises toward assigning the representative oriented reppoints samples. Furthermore, a spatial constraint is proposed to facilitate the vulnerable points to find their instance owner from the complex context in the aerial scene. Compared with the orientation regression-based methods, our framework obtains more precise detection performance with accurate orientation.

In summary, the main contributions of this paper are: (1) an effective aerial object detector named Oriented RepPoints, where the flexible adaptive points are introduced as the representation to achieve the oriented object detection; (2) a novel quality assessment and sample assignment scheme for adaptive points learning, which selects the points sample not only from the classification, localization but also from the orientation, point-wise feature correlation; (3) extensive experiments on four challenging datasets showing promising qualitative and quantitative results.

## 2. Related Work

Unlike most of generic object detectors with horizontal bounding boxes, the targets in aerial images are often arbitrary-oriented and dense-distributed. We discuss the related work in the following.

### 2.1. Oriented Object detection

The recent aerial object detection methods are mainly derived from the classical object detectors by introducing the orientation regression. SCRDet [45], CAD-Net [47], DRN [24], R3Det [42], ReDet [9] and Oriented R-CNN [35] achieve the promising performance by predicting

the rotation angles of bounding boxes. Gliding Vertex [36] and RSDet [25] improve the detection results through regressing quadrilateral. To address the boundary discontinuity in the angle-based orientation estimation, Yang *et al.* [41] transform the angular regression to angular classification [40]. Later, Yang *et al.* [43] convert the parameterization of the rotated bounding box into 2-D Gaussian distributions, which gains more robust results for the oriented object detection. These methods are devoted to improving orientation estimation using rotation angle representation. Alternatively, we introduce a more effective representation using adaptive points in this paper.

### 2.2. Non-axis Aligned Features Learning

Most of conventional object detection methods [26, 28, 30, 46, 48, 49, 53] focus on either upright or axis-aligned objects, which may have difficulties with the non-axis aligned targets densely distributed in the complex background. To address this issue, Ding *et al.* [3] adopt the spatial transformations on axis-align RoIs and learn the non-axis aligned representation under the supervision of oriented bounding box. SCRDet++ [44] enhances the non-axis aligned features and bring the higher object response to train the network. Han *et al.* [8] design a feature alignment module to alleviate the misalignment between axis-aligned convolutional features and arbitrary oriented objects. DRN [24] proposes the feature selection module to aggregate the non-axis aligned information from the different kernel sizes, shapes and orientations, and employs the dynamic filter generator for further regression. Guo *et al.* [7] employ the convex-hull representation to learn the irregular shapes and layouts, which intend to avoid the feature aliasing via learnable feature adaption. Our point set-based method aims to capture the key features for non-axis aligned aerial objects.

### 2.3. Samples Assignment for Object Detection

A large amount of detection methods adopt a simple way to set the IoU threshold for selecting the positive samples. However, such scheme cannot guarantee the overall quality of training samples due to the potential noise and hard cases [15, 22]. Some recent samples assignment methods in general object detection, such as the ATSS [50], FreeAnchor [51], PAA [13] and OTA [6], employ a learning-to-match optimization strategy [52] to choose the high-quality samples. In aerial scene, it is essential to select the high-quality samples for learning the oriented detector due to the diversity of the orientation and dense distribution. Ming *et al.* [23] introduce a matching degree measure to evaluate the spatial alignment based on the oriented anchors, which use the matching sensitive loss to enhance the correlation between classification and oriented localization. In this work, we suggest an effective quality assessment and sample assignment scheme to select the positive points samples.

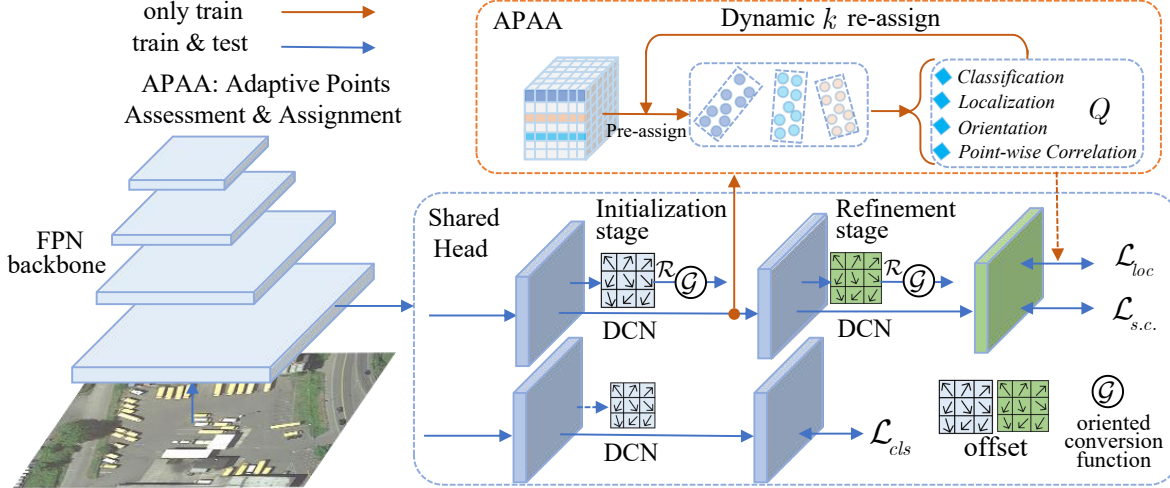


Figure 2. The framework of Oriented RepPoints. The proposed method is an anchor-free approach with the adaptive points as the representation, where a backbone with FPN network is employed for feature encoding. The structure of the shared head is same as RepPoints [46] for each scale of FPN, except of the proposed APAA and the oriented conversion function. Based on learning points from the initialization stage, the APAA scheme is performed only during training.

### 3. Oriented RepPoints

#### 3.1. Overview

Instead of directly regressing the orientations like the conventional methods [3, 8, 9, 45], we take advantage of the adaptive point set [46] as a fine-grained representation, which is able to capture the geometric structure of the aerial objects with sharp variety on orientation in the cluttered environments. To this end, we introduce the differentiable conversion functions, where the representative points are driven to adaptively move toward the appropriate positions over an oriented object. In order to effectively learn the high-quality adaptive points without direct points-to-points supervision, we suggest a quality measure scheme that selects the high-quality oriented reppoints at the training stage. To facilitate the robust adaptive point learning, the spatial constraints are employed to penalize the vulnerable outliers and find their instance owner from the complicated aerial context. Fig. 2 shows the overview of our proposed Oriented Reppoints approach.

#### 3.2. Adaptive Points Learning with Orientation

To facilitate the oriented detector with point set representation, the conversion function is introduced to transform the adaptive points into the oriented bounding box. Let  $\mathcal{G}$  denote the oriented conversion function as below:

$$OB = \mathcal{G}(\mathcal{R}) \quad (1)$$

where  $OB$  represents an oriented box converted from the learned point set  $\mathcal{R}$ . In this paper, we examine three oriented conversion functions:

- *MinAeraRect* intends to find the rotated rectangle with the minimum area from the learned point set over an oriented object.

- *NearestGTCorner* makes use of the ground-truth annotations. For each corner, we find the closest point from the learned point set as a predicted corner, where the selected corner points are used to build a quadrilateral as the oriented bounding box.

- *ConvexHull*. An oriented instance polygon can be defined as a convex hull of a set of points driven by the Jarvis March algorithm [7, 11], which is used by many contour-based methods.

Note that the *NearestGTCorner* and *ConvexHull* are differentiable functions while *MinAeraRect* is not. Thus, we employ *MinAeraRect* in the post-processing to get the standard rotated rectangle prediction, and the other two differentiable functions are used to optimize adaptive points learning during the training. Under the supervision of the oriented ground-truth annotations, the points move towards the semantic key and geometric features adaptively for each aerial object, which are driven by the classification and localization loss simultaneously.

The proposed framework consists of two stages. The initialization stage generates the adaptive point sets by refining from the object center point (feature map bins). The refinement stage further gains the accurate adjustment by minimizing the loss function as below:

$$\mathcal{L} = \mathcal{L}_{cls} + \lambda_1 \mathcal{L}_{s1} + \lambda_2 \mathcal{L}_{s2} \quad (2)$$

where  $\lambda_1$  and  $\lambda_2$  are balanced weighting.  $\mathcal{L}_{cls}$  denotes the

object classification loss:

$$\mathcal{L}_{cls} = \frac{1}{N_{cls}} \sum_i \mathcal{F}_{cls}(\mathcal{R}_i^{cls}(\theta), b_j^{cls}) \quad (3)$$

where  $\mathcal{R}_i^{cls}(\theta)$  represents the predicted class confidence based on the learned points, and  $b_j^{cls}$  is the assigned ground-truth class.  $\mathcal{F}_{cls}$  is the focal loss [19].  $N_{cls}$  denotes the total number of point sets.  $\mathcal{L}_{s1}$  and  $\mathcal{L}_{s2}$  represent the spatial localization loss at the initialization and refinement stage, respectively. For each stage,  $\mathcal{L}_s$  can be denoted as below:

$$\mathcal{L}_s = \mathcal{L}_{loc} + \mathcal{L}_{s.c.} \quad (4)$$

where  $\mathcal{L}_{loc}$  is localization loss based on converted oriented boxes, and  $\mathcal{L}_{s.c.}$ , denotes the spatial constraint loss.

Let  $N_{loc}$  denote the total number of positive point set samples.  $b_j^{loc}$  indicates the location of ground-truth box. Thus,  $\mathcal{L}_{loc}$  is defined as follows:

$$\mathcal{L}_{loc} = \frac{1}{N_{loc}} \sum_i [b_j^{cls} \geq 1] \mathcal{F}_{loc}(OB_i^{loc}(\theta), b_j^{loc}) \quad (5)$$

where  $\mathcal{F}_{loc}$  is the GIoU loss [29] for the oriented polygon.

Due to the diversity of different categories and the cluttered background in aerial images, a portion of learned points are susceptible to the background or adjacent objects with strong key features, which may move outside the ground-truth bounding box. To facilitate the vulnerable points to capture the geometric features on its instance owner, we introduce an effective spatial constraint to penalize the adaptive points outside the bounding box. Let  $\rho_{ij}$  denote the penalty function. The spatial loss  $\mathcal{L}_{s.c.}$  for each oriented object is defined as below:

$$\mathcal{L}_{s.c.} = \frac{1}{N_a} \frac{1}{N_o} \sum_{i=1} \sum_{j=1} \rho_{ij} \quad (6)$$

where  $N_a$  indicates the number of assigned positive point set samples for each object.  $N_o$  is the number of points outside the GT box in each point set.

Let  $\mathbf{p}_c$  denote the geometric center of the ground-truth bounding box. Given a sampled point  $\mathbf{p}_o$  outside the bounding box, the penalty term is defined as below:

$$\rho = \begin{cases} \|\mathbf{p}_o - \mathbf{p}_c\|, & \mathbf{p}_o \text{ is outside } GT \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

where  $GT$  denotes the ground-truth box.

### 3.3. Adaptive Points Assessment and Assignment

Due to the lack of direct supervision, learning high-quality points is essential to capture the geometric features adaptively for densely-packed and arbitrarily-oriented objects in aerial images. To this end, we propose an effective

assessment and assignment scheme to measure the quality of learned points, which towards assigning the representative samples of adaptive points as the positive samples at the training stage.

**Quality Measure of Adaptive Points.** Firstly, we define a quality measure  $Q$  to assess the learned adaptive points from four aspects, including classification and localization ability  $Q_{cls}$ ,  $Q_{loc}$ , orientation alignment  $Q_{ori}$ , and point-wise correlation  $Q_{poc}$  for each oriented point set. Thus,  $Q$  is derived as below:

$$Q = Q_{cls} + \mu_1 Q_{loc} + \mu_2 Q_{ori} + \mu_3 Q_{poc} \quad (8)$$

The classification ability  $Q_{cls}$  of a point set  $\mathcal{R}_i$  directly reflects its classification confidence  $\mathcal{R}_i^{cls}(\theta)$ , where the corresponding classification loss  $\mathcal{L}_{cls}$  measures the compatibility of the points feature with ground-truth class label  $b_j^{cls}$ . We define  $Q_{cls}$  as follows:

$$Q_{cls}(\mathcal{R}_i, b_j) = \mathcal{L}_{cls}(\mathcal{R}_i^{cls}(\theta), b_j^{cls}) \quad (9)$$

To evaluate the compatibility of the points position with the ground-truth  $b_j^{loc}$ , we employ the localization loss as the quality assessment measure, which is based on the IoU transformation. It indicates the spatial alignment when the center of a point set is near to the object's geometric center. Therefore,  $Q_{loc}$  is defined as below:

$$Q_{loc}(\mathcal{R}_i, b_j) = \mathcal{L}_{loc}(OB_i^{loc}(\theta), b_j^{loc}) \quad (10)$$

Since  $Q_{loc}$  can be regarded as a measure of spatial location distance, it is insensitive to the orientation variations, especially for the square-like objects in the aerial images. To account for the orientation alignment, we employ Chamfer distance [5] to assess the difference in orientation between the predicted point set and ground-truth box contour points. We firstly adopt the *MinAeaRect* conversion function to obtain four spatial corner points  $\{v_1, v_2, v_3, v_4\}$  from the learned point set. Then, an ordered point set  $\mathcal{R}^v$  (40 points by default) is sampled with the equal interval from two adjacent corner points. Similarly, the points  $\mathcal{R}^g$  are generated for the ground-truth corner points  $\{g_1, g_2, g_3, g_4\}$ . Therefore,  $Q_{ori}$  is defined as follows:

$$Q_{ori}(\mathcal{R}_i, b_j) = \mathcal{CD}(\mathcal{R}_i^v(\theta), \mathcal{R}_{b_j}^g) \quad (11)$$

where  $\mathcal{CD}$  denotes Chamfer distance between the above two group of sampling points:

$$\begin{aligned} \mathcal{CD}(\mathcal{R}^v, \mathcal{R}^g) &= \frac{1}{2n} \sum_{i=1}^n \min_j \|(x_i^v, y_j^v) - (x_i^g, y_j^g)\|_2 \\ &+ \frac{1}{2n} \sum_{j=1}^n \min_i \|(x_i^v, y_j^v) - (x_i^g, y_j^g)\|_2 \end{aligned} \quad (12)$$



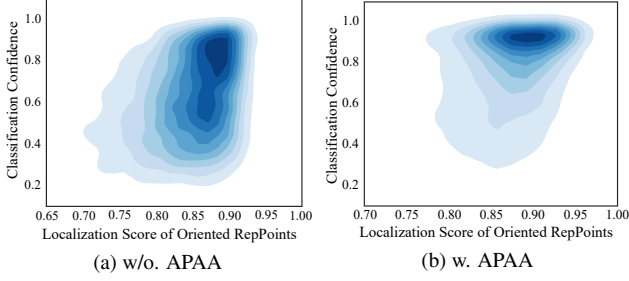


Figure 3. The correlation between the predicted classification confidence and localization score of the oriented reppoints with and without APAA scheme.

$(x_i^v, y_j^v) \in \mathcal{R}^v$  denotes the sampled points of predicted spatial corner points, and  $(x_i^g, y_j^g) \in \mathcal{R}^g$  denotes the sampled points generated from ground-truth corner points.

To measure the point-wise association upon a point set for an oriented object, we extract the point-wise features and employ the cosine similarity between the feature vectors as the correlation measure  $Q_{poc}$  for the learned adaptive points. Let  $e_{i,k}$  denote the  $k$ -th point-wise feature vector of  $i$ -th set of adaptive points.  $e_{i,k}^*$  and  $e_i^*$  represent the normalized embedding feature vector and their mean from the  $i$ -th point set:

$$e_{i,k}^* = \frac{e_{i,k}}{\|e_{i,k}\|_2} \quad (13)$$

$$e_i^* = \frac{1}{N_p} \sum_{k=1} e_{i,k}^* \quad (14)$$

where  $N_p$  denotes the num of points in a point set. The default setting is 9. Based on the above notations,  $Q_{poc}$  of  $i$ -th point sets can be formulated as the point-wise feature diversity as below:

$$\begin{aligned} Q_{poc} &= 1 - \frac{1}{N_p} \sum_k \cos \langle e_{i,k}^*, e_i^* \rangle \\ &= 1 - \frac{1}{N_p} \sum_k \frac{e_{i,k}^* \cdot e_i^*}{\|e_{i,k}^*\| \times \|e_i^*\|} \end{aligned} \quad (15)$$

**Dynamic  $k$  Label Assignment.** Based on the quality measure  $Q$ , we assign the oriented reppoints samples through an efficient and dynamic top  $k$  item selection scheme at different iteration. For each object, we sort all the point set samples from the initialization stage according to their quality scores. To retrieve high-quality adaptive point set samples, we set a sampling ratio  $\sigma$  to assign the top  $k$  samples at each iteration as the positive samples for training, which is calculated by:

$$k = \sigma * N_t \quad (16)$$

where  $N_t$  denotes the total number of point set samples at the initialization stage for each oriented object.

During the training, the points assigner [46] is used to get the sample assignment of center points at the initialization stage. At the refinement stage, the proposed adaptive points assessment and assignment (APAA) scheme is used to select the high-quality points samples according to the quality measure  $Q$ . Only the selected positive point sets are assigned with the ground-truth bounding box of target. As shown in Fig. 3, APAA scheme enables the detector to predict the high-quality oriented reppoints for improving both classification confidence and localization scores. It is worthy of mentioning that the presented scheme is only used for training, which does not incur the computational load at the inference stage.

## 4. Experiments

### 4.1. Evaluation Testbed

DOTA [34] is a large-scale dataset to evaluate the detection performance of oriented objects in aerial images, which contains 2806 images, 188,282 instances and 15 categories with a variety of orientations, scales, and shapes. The training set has 1411 images while the validation set contains 458 images. Testing set is made of 937 images. The image sizes range from  $800 \times 800$  to  $4000 \times 4000$ . In our experiments, both the training set and validation set are employed to train the proposed detector, and the testing set without annotations is used for evaluation. We crop the original images into the patches of  $1024 \times 1024$  with a stride of 824. At the training stage, we randomly resize and flip the images to avoid overfitting.

HRSC2016 [21] contains a large number of strip-like oriented objects with diverse appearances collected from several famous harbors for ship recognition. The entire dataset has 1061 images ranging from  $300 \times 300$  to  $1500 \times 900$ . For a fair comparison, the training set (436 images) and validation set (181 images) are employed for training, and the testing set (444 images) is used for evaluation.

UCAS-AOD [54] has 1510 images with 510 car images and 1000 airplane images. There are 14,596 instances in total. The entire dataset is randomly divided into 755 images for training, 302 images for validation and 453 images for testing with a ratio of 5:2:3. The size of all images is approximately  $1280 \times 659$ .

DIOR-R [2] gives the oriented bounding box annotations based on the DIOR dataset [16] for the oriented detection task. There are 23,463 images with the size of  $800 \times 800$  and 192,518 instances covering 20 object classes.

### 4.2. Implementation Details

We implement our proposed approach based on both ResNet-50 [10] and ResNet-101 backbone with FPN [18]. The FPN consists of  $P_3$  to  $P_7$  pyramid levels in our work. The stochastic gradient descent (SGD) optimizer is used in

training. The initial learning rate is set to 0.008 with the warming up for 500 iterations, and the learning rate is decreased by a factor of 0.1 at each decay step. The momentum is set to 0.9, and weight decay is  $10^{-4}$ . We train the models with 40 epochs, 40 epochs, 120 epochs and 120 epochs for DOTA, DIOR-R, HRSC2016 and UCAS-AOD, respectively. The scale jitter is employed during the training phase. The hyperparameters of focal loss are set to  $\alpha = 0.25$  and  $\gamma = 2.0$ . In Eq. (2), we set the balanced weights for each stage  $\lambda_1 = 0.3$  and  $\lambda_2 = 1.0$ , empirically. We set  $\mu_1 = 1.0$ ,  $\mu_2 = 0.3$  and  $\mu_3 = 0.1$  for quality evaluation  $Q$  in Eq. (8). A couple of experiments are performed to choose the appropriate values of the sampling ratio  $\sigma$  in Table 6.

We conduct the experiments on a server with 4 RTX 2080Ti GPUs using a total batch size of 8 (2 images per GPU) for training while a single RTX 2080Ti GPU is employed for inference.

### 4.3. Ablation Study

To examine the effectiveness of each component in our proposed framework, a series of ablation experiments are performed on DOTA dataset with ResNet-50-FPN.

**Evaluation on oriented conversion functions.** The conventional point set-based object detector RepPoints [46] obtains the upright bounding boxes by square conversion function like *min-max*, which cannot deal with the aerial objects having arbitrary orientations. To build a reasonable baseline, we compare the different conversion functions that map the adaptive points into the oriented box during the training and post-processing. Table 1 shows the experimental results. Based on the original RepPoints with *min-max* function both in training and post-processing, it is able to achieve 49.69% mAP. With the proposed oriented *MinAeraRect* function for post-processing to get the rotated rectangular boxes, it achieves the 53.21% mAP. With the differentiable oriented *NearestGTCorner* and *ConvexHull* functions, our Oriented RepPoints obtains the 66.97% mAP and 68.89% mAP, which shows that the oriented conversion functions are essential to aerial object detection.

Methods	training	post-processing	mAP
RepPoints [46]	<i>min-max</i>	<i>min-max</i>	49.69
	<i>min-max</i>	<i>MinAeraRect</i>	53.21
Oriented RepPoints	<i>NearestGTCorner</i>	<i>MinAeraRect</i>	66.97
	<i>ConvexHull</i>	<i>MinAeraRect</i>	<b>68.89</b>

Table 1. Comparison with different conversion functions.

**Comparison with angle-based detectors.** To examine the effectiveness of adaptive points representation, we compare our approach with the angle-based orientation regression on the anchor-based detector. As in S<sup>2</sup>A-Net [8], the angle-base detector presets one squared anchor for each

Methods	Backbone	mAP
angle-based detector	ResNet-50-FPN	67.50
Oriented RepPoints	ResNet-50-FPN	<b>68.89 (+1.39)</b>
angle-based detector	ResNet-101-FPN	68.73
Oriented RepPoints	ResNet-101-FPN	<b>70.19 (+1.46)</b>

Table 2. Comparisons between direct angle-based orientation regression and the Oriented RepPoints for oriented object detection.

Spatial Constraint	BD	BR	RA	HC	mAP
	76.85	41.72	67.09	41.55	68.89
✓	<b>79.99</b>	<b>45.33</b>	<b>71.39</b>	<b>51.87</b>	<b>70.11</b>
improvement	+3.14	+3.61	+4.30	<b>+10.02</b>	+1.22

Table 3. Performance evaluation on spatial constraints. BD, BR, RA and HC denote the categories of Baseball Diamond, Bridge, Roundabout and Helicopter, respectively.

feature map location at the initialization stage, where the predicted angle-based boxes are regarded as the refined anchors for the next stage to get the oriented bounding box. Table 2 shows the results of two detectors with different backbones. The Oriented RepPoints outperforms the angle-based orientation regression with +1.39% and +1.46% mAP improvement using ResNet-50-FPN and ResNet-101-FPN backbone, respectively.

**Evaluation on spatial constraint.** To investigate the effectiveness of spatial constraint, we compare it against the baseline method without using it. Table 3 shows the experimental results. It can be observed that our proposed spatial constraint is very effective, especially for the aerial objects with weak feature representation such as HC (Helicopter), and the objects similar to the background, e.g. BD (Baseball Diamond), BR (Bridge) and RA (Roundabout). This is because the spatial constraint enforces the adaptive points on their owner instance object.

**APAA scheme for adaptive points learning.** To study the proposed APAA scheme for adaptive points learning, we first report the performance of the quality measure term-by-term. Table 5 gives the results of different settings on quality assessment measure  $Q$ . The detection result is progressively improved, and the proposed approach achieves the best performance with 75.97% mAP and +5.86% gains using all four terms. It indicates that the quality assessment measure is effective to reflect the quality of adaptive points for aerial object detection. In APAA scheme, the number of assigned adaptive points samples is determined by the sampling ratio  $\sigma$ . As shown in Table 6, the model achieves the best performance when  $\sigma = 0.4$ . Additionally, we compare our APAA scheme with other sample assignment schemes for training the proposed detector, including Max-IoU [28], ATSS [50], PAA [13] and CFA [7]. As illustrated in Table 7,

Methods	Backbone	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	mAP
<i>Single-stage Methods</i>																	
RetinaNet-O [19]	R-50-FPN	88.67	77.62	41.81	58.17	74.58	71.64	79.11	90.29	82.18	74.32	54.75	60.60	62.57	69.67	60.64	68.43
DAL [23]	R-101-FPN	88.61	79.69	46.27	70.37	65.89	76.10	78.53	90.84	79.98	78.41	58.71	62.02	69.23	71.32	60.65	71.78
RSDet [25]	R-152-FPN	<b>90.10</b>	82.00	53.80	68.50	70.20	78.70	73.60	91.20	87.10	84.70	64.30	68.20	66.10	69.30	63.70	74.10
R <sup>3</sup> Det [42]	R-152-FPN	89.49	81.17	50.53	66.10	70.92	78.66	78.21	90.81	85.26	84.23	61.81	63.77	68.16	69.83	<b>67.17</b>	73.74
S <sup>2</sup> A-Net [8]	R-50-FPN	89.11	82.84	48.37	71.11	78.11	78.39	87.25	90.83	84.90	85.64	60.36	62.60	65.26	69.13	57.94	74.12
R <sup>3</sup> Det-DCL [40]	R-152-FPN	89.78	83.95	52.63	69.70	76.84	81.26	87.30	90.81	84.67	85.27	63.50	64.16	68.96	68.79	65.45	75.54
<i>Two-stage Methods</i>																	
Faster RCNN-O [28]	R-50-FPN	88.44	73.06	44.86	59.09	73.25	71.49	77.11	90.84	78.94	83.90	48.59	62.95	62.18	64.91	56.18	69.05
CAD-Net [47]	R-101-FPN	87.80	82.40	49.40	73.50	71.10	63.50	76.60	90.90	79.20	73.30	48.40	60.90	62.00	67.00	62.20	69.90
SCRDet [45]	R-101-FPN	89.98	80.65	52.09	68.36	68.36	60.32	72.41	90.85	<b>87.94</b>	<b>86.86</b>	<b>65.02</b>	66.68	66.25	68.24	65.21	72.61
FAOD [14]	R-101-FPN	<b>90.21</b>	79.58	45.49	76.41	73.18	68.27	79.56	90.83	83.40	84.68	53.40	65.42	74.17	69.69	64.86	73.28
RoI-Trans. [3]	R-101-FPN	88.65	82.60	52.53	70.87	77.93	76.67	86.87	90.71	83.83	82.51	53.95	67.61	74.67	68.75	61.03	74.61
Gliding Vertex [36]	R-101-FPN	89.64	<b>85.00</b>	52.26	<b>77.34</b>	73.01	73.14	86.82	90.74	79.02	<b>86.81</b>	59.55	<b>70.91</b>	72.94	70.86	57.32	75.02
MaskOBB [31]	R-50-FPN	89.61	<b>85.09</b>	51.85	72.90	75.28	73.23	85.57	90.37	82.08	85.05	55.73	68.39	71.61	69.87	<b>66.33</b>	74.86
CenterMap [32]	R-50-FPN	88.88	81.24	53.15	60.65	78.62	66.55	78.10	88.83	77.80	83.61	49.36	66.19	72.10	72.36	58.70	71.74
ReDet [9]	ReR-50-ReFPN [9]	88.79	82.64	53.97	74.00	78.13	<b>84.06</b>	<b>88.04</b>	90.89	<b>87.78</b>	85.75	61.76	60.39	<b>75.96</b>	68.07	63.59	76.25
Oriented R-CNN [35]	R-101-FPN	88.86	83.48	55.27	<b>76.92</b>	74.27	82.10	87.52	<b>90.90</b>	85.56	85.33	<b>65.51</b>	66.82	74.36	70.15	57.28	76.28
<i>Anchor-free Methods</i>																	
CenterNet-O [53]	DLA-34 [53]	81.00	64.00	22.60	56.60	38.60	64.00	64.90	90.80	78.00	72.50	44.00	41.10	55.50	55.00	57.40	59.10
PloU [11]	DLA-34	80.90	69.70	24.10	60.20	38.30	64.40	64.80	90.90	77.20	70.40	46.50	37.10	57.10	61.90	64.00	60.50
O <sup>2</sup> -DNNet [33]	H-104 [39]	89.31	82.14	47.33	61.21	71.32	74.03	78.62	90.76	82.23	81.36	60.93	60.17	58.21	66.98	61.03	71.04
DRN [24]	H-104	89.71	82.34	47.22	64.10	76.22	74.43	85.84	90.57	86.18	84.89	57.65	61.93	69.30	69.63	58.48	73.23
CFA [7]	R-101-FPN	89.26	81.72	51.81	67.17	79.99	78.25	84.46	90.77	83.40	85.54	54.86	67.75	73.04	70.24	64.96	75.05
Oriented RepPoints	R-50-FPN	87.02	83.17	54.13	71.16	<b>80.18</b>	78.40	<b>87.28</b>	<b>90.90</b>	85.97	86.25	59.90	<b>70.49</b>	73.53	72.27	58.97	75.97
Oriented RepPoints	R-101-FPN	89.53	84.07	<b>59.86</b>	71.76	79.95	80.03	87.33	90.84	87.54	85.23	59.15	66.37	75.23	<b>73.75</b>	57.23	<b>76.52</b>
Oriented RepPoints	Swin-T-FPN	89.11	82.32	<b>56.71</b>	74.95	<b>80.70</b>	<b>83.73</b>	<b>87.67</b>	90.81	87.11	85.85	63.60	68.60	<b>75.95</b>	<b>73.54</b>	63.76	<b>77.63</b>

Table 4. Comparison with state-of-the-art methods on DOTA dataset. All reported results are performed on the single-scale DOTA dataset. The results with red color denote the best results and with blue color present the second-best results in each column. ‘-O’ means the detection results with oriented bounding box (the same below).

Quality Measure $Q$ for Adaptive Points.				
$Q_{cls}$	✓	✓	✓	✓
$Q_{loc}$		✓	✓	✓
$Q_{ori}$			✓	✓
$Q_{poc}$				✓
mAP	70.11	72.34	74.46	<b>75.97</b>

Table 5. Performance evaluation on different settings of quality measure  $Q$  in APAA scheme.

$\sigma$	0.2	0.3	0.4	0.5
mAP	75.45	75.34	<b>75.97</b>	75.67

Table 6. Evaluation on various  $\sigma$  in the dynamic top  $k$  assignment of APAA scheme.

Methods	Max-IoU [28]	ATSS [50]	PAA [13]	CFA [7]	APAA
mAP	70.11	72.87	74.62	74.89	<b>75.97</b>

Table 7. Comparisons with different samples assignment methods on Oriented RepPoints detector.

our APAA scheme achieves the best performance without the complicated operations, which demonstrates that our proposed APAA is effective for adaptive points learning.

#### 4.4. Comparison with the State-of-the-art methods

**Results on DOTA.** We report full experimental results of single scale to make a fair comparison with the previous methods on DOTA dataset. With ResNet-50-FPN and ResNet-101-FPN as the backbones, our method obtains 75.97% and 76.52% mAP, respectively. It outperforms other methods with the corresponding backbones. Using the tiny version of Swin-Transformer [20] (Swin-T-FPN) as backbone with the random rotation and HSV transformation, we achieve the best performance with 77.63% mAP. Fig. 4 shows some visual results on DOTA test set.

**Results on HRSC2016.** To make a comprehensive comparison on HRSC2016, we report the results with both VOC2007 and VOC2012 metrics. Table 8 shows the experimental results. Our Oriented RepPoints achieves the best performance under VOC2012 metric and the second-best under VOC2007 metric with ResNet-50-FPN backbone.

**Results on UCAS-AOD.** UCAS-AOD datasets contains a large number of small objects with complex surrounding scenes. Table 9 shows the evaluation results with the recent methods on UCAS-AOD dataset. Our presented method achieves the best performance of 90.11% mAP.

**Results on DIOR-R.** DIOR-R datasets consists of 20 classes of aerial objects. Compared with the recent methods on this dataset, we achieve the best performance with 66.71% mAP and outperform other methods, as shown in Table 10.

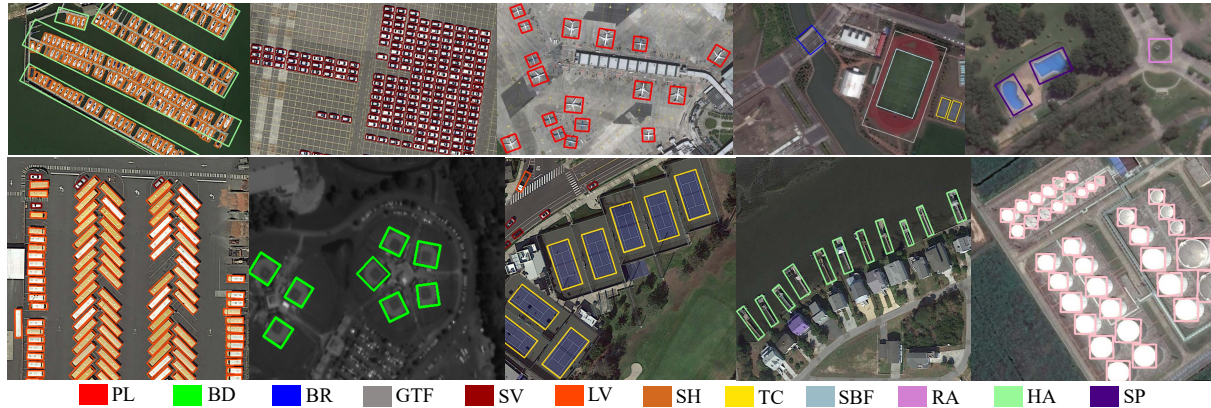


Figure 4. Example detection results of Oriented RepPoints on DOTA test set.

Methods	Backbone	mAP <sub>50</sub> (07)	mAP <sub>50</sub> (12)
R <sup>2</sup> CNN [12]	R-101-FPN	73.07	79.73
RRD [17]	VGG16	84.30	-
RoI-Trans. [3]	R-101-FPN	86.20	-
CenterNet-O [53]	DLA-34	87.89	-
Gliding Vertex [36]	R-101-FPN	88.20	-
DRN [24]	H-104	-	92.70
CenterMap-Net [32]	R-50-FPN	-	92.80
RetinaNet-O [19]	R-101-FPN	89.18	95.21
PIOU [1]	DLA-34	89.20	-
MFIAR-Net [38]	R-101-FPN	89.81	-
R <sup>3</sup> Det [42]	R-101-FPN	89.26	96.01
R <sup>3</sup> Det-DCL [40]	R-101-FPN	89.46	96.41
FPN-CSL [41]	R-101-FPN	89.62	96.10
DAL [23]	R-101-FPN	89.77	-
S <sup>2</sup> A-Net [8]	R-101-FPN	90.17	95.01
Oriented R-CNN [35]	R-50-FPN	<b>90.40</b>	96.50
Oriented RepPoints	R-50-FPN	90.38	<b>97.26</b>

Table 8. Results on HRSC2016 test set. mAP(07) and mAP(12) represent the results under VOC2007 and VOC2012 mAP metrics.

Methods	Car	Airplane	mAP
YOLOv3-O [27]	74.63	89.52	82.08
RetinaNet-O [19]	84.64	90.51	87.57
Faster R-CNN-O [34]	86.87	89.86	88.36
RoI Trans. [3]	87.99	89.90	88.95
DAL [23]	89.25	90.49	89.87
Oriented RepPoints	<b>89.51</b>	<b>90.70</b>	<b>90.11</b>

Table 9. Performance comparisons on UCAS-AOD dataset.

Methods	RetinaNet-O [19]	Faster RCNN-O [28]	Gliding Vertex [36]
mAP	57.55	59.54	60.06
Methods	RoI-Trans. [3]	AOPG [2]	Oriented RepPoints
mAP	63.87	64.41	<b>66.71</b>

Table 10. Detection accuracy on DIOR-R dataset. All experimental results are performed with ResNet-50-FPN backbone.

Methods	RetinaNet-O [19]	Faster RCNN-O [28]	RoI-Trans. [3]	S <sup>2</sup> A-Net [8]
mAOE	9.53	6.01	6.35	10.43
Methods	ReDet [9]	Oriented R-CNN [35]	Oriented RepPoints	
mAOE	6.35	7.53	<b>5.93</b>	

Table 11. Comparison on orientation errors with DOTA. All experiments are performed with *train* set for training, *val* set for testing.

#### 4.5. Evaluation on Orientation Accuracy

We further conduct experiment to evaluate the orientation accuracy of an oriented detector on DOTA dataset with ResNet-50-FPN backbone. We employ the mean Average Orientation Error (mAOE<sup>o</sup>) on all categories as the evaluation metric. As shown in Table 11, our proposed approach obtains the smallest orientation errors, which demonstrates that our point set-based approach is effective for precise oriented object detection, comparing to the conventional orientation regression-based methods.

## 5. Conclusion

This paper proposed an effective aerial object detector by taking advantage of the adaptive points as a fine-grained representation, which is able to capture the key geometric features for arbitrary-oriented, cluttered and non-axis aligned targets. To effectively learn the adaptive points, we introduced the quality assessment and sample assignment scheme to measure and select the high-quality points samples for training. Furthermore, a spatial constraint is used to penalize the points outside the oriented box for robust adaptive points learning. The extensive experiments have been performed on four testbeds, whose promising results demonstrate the efficacy of our proposed approach.

## Acknowledgments

This work is supported by National Natural Science Foundation of China under Grants (61831015).



## References

- [1] Zhiming Chen, Kean Chen, Weiyao Lin, John See, Hui Yu, Yan Ke, and Cong Yang. Piou loss: Towards accurate oriented object detection in complex environments. In *ECCV*, pages 195–211, 2020. 7, 8
- [2] Gong Cheng, Jiabao Wang, Ke Li, Xingxing Xie, Chunbo Lang, Yanqing Yao, and Junwei Han. Anchor-free oriented proposal generator for object detection. *arXiv preprint arXiv:2110.01931*, 2021. 5, 8
- [3] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *CVPR*, pages 2849–2858, Jun. 2019. 1, 2, 3, 7, 8
- [4] Jian Ding, Nan Xue, Gui-Song Xia, Xiang Bai, Wen Yang, Michael Ying Yang, Serge J. Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Object detection in aerial images: A large-scale benchmark and challenges. *IEEE TPAMI*, pages 1–1, 2021. 1
- [5] Haoqiang Fan, Hao Su, and Leonidas Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, pages 2463–2471, Jul. 2017. 4
- [6] Zheng Ge, Songtao Liu, Zeming Li, Osamu Yoshie, and Jian Sun. Ota: Optimal transport assignment for object detection. In *CVPR*, pages 303–312, 2021. 2
- [7] Zonghao Guo, Chang Liu, Xiaosong Zhang, Jianbin Jiao, Xiangyang Ji, and Qixiang Ye. Beyond bounding-box: Convex-hull feature adaptation for oriented and densely packed object detection. In *CVPR*, pages 8792–8801, 2021. 1, 2, 3, 6, 7
- [8] Jiaming Han, Jian Ding, Jie Li, and Gui-Song Xia. Align deep features for oriented object detection. *IEEE TGRS*, pages 1–11, Mar. 2021. 1, 2, 3, 6, 7, 8
- [9] Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotation-equivariant detector for aerial object detection. In *CVPR*, pages 2786–2795, 2021. 1, 2, 3, 7, 8
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [11] Ray A. Jarvis. On the identification of the convex hull of a finite set of points in the plane. *Information Processing Letters*, 2(1):18–21, 1973. 3
- [12] Yingying Jiang, Xiangyu Zhu, Xiaobing Wang, Shuli Yang, Wei Li, Hua Wang, Pei Fu, and Zhenbo Luo. R2cnn: Rotational region cnn for orientation robust scene text detection. *arXiv preprint arXiv: 1706.09579*, 2017. 8
- [13] Kang Kim and Hee Seok Lee. Probabilistic anchor assignment with iou prediction for object detection. In *ECCV*, pages 355–371, 2020. 2, 6, 7
- [14] Chengzheng Li, Chunyan Xu, Zhen Cui, Dan Wang, Tong Zhang, and Jian Yang. Feature-attentioned object detection in remote sensing imagery. In *ICIP*, pages 3886–3890, 2019. 7
- [15] Hengduo Li, Zuxuan Wu, Chen Zhu, Caiming Xiong, Richard Socher, and Larry S. Davis. Learning from noisy anchors for one-stage object detection. In *CVPR*, pages 10588–10597, 2020. 2
- [16] Ke Li, Gang Wan, Gong Cheng, Liqiu Meng, and Junwei Han. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS J. of Photogramm. Remote Sens.*, 159:296–307, 2020. 5
- [17] Minghui Liao, Zhen Zhu, Baoguang Shi, Gui song Xia, and Xiang Bai. Rotation-sensitive regression for oriented scene text detection. In *CVPR*, pages 5909–5918, Jun. 2018. 8
- [18] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 936–944, Jul. 2017. 1, 5
- [19] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. *IEEE TPAMI*, 42(2):318–327, Feb. 2020. 1, 4, 7, 8
- [20] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 7
- [21] Zikun Liu, Liu Yuan, Lubin Weng, and Yang Yiping. A high resolution optical satellite image dataset for ship recognition and some new baselines. pages 324–331, 2017. 5
- [22] Yuchen Ma, Songtao Liu, Zeming Li, and Jian Sun. Iqdet: Instance-wise quality distribution sampling for object detection. In *CVPR*, pages 1717–1725, 2021. 2
- [23] Qi Ming, Zhiqiang Zhou, Lingjuan Miao, Hongwei Zhang, and Linhao Li. Dynamic anchor learning for arbitrary-oriented object detection. In *AAAI*, pages 2355–2363, 2021. 2, 7, 8
- [24] Xingjia Pan, Yuqiang Ren, Kekai Sheng, Weiming Dong, Haolei Yuan, Xiaowei Guo, Chongyang Ma, and Changsheng Xu. Dynamic refinement network for oriented and densely packed object detection. In *CVPR*, pages 11207–11216, Jun. 2020. 1, 2, 7, 8
- [25] Wen Qian, Xue Yang, Silong Peng, Junchi Yan, and Yue Guo. Learning modulated loss for rotated object detection. In *AAAI*, pages 2458–2466, 2020. 1, 2, 7
- [26] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016. 2
- [27] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 8
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE TPAMI*, 39(6):1137–1149, Jun. 2017. 1, 2, 6, 7, 8
- [29] Hamid Rezaatofghi, 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 *CVPR*, pages 658–666, 2019. 4
- [30] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. 2
- [31] Jinwang Wang, Jian Ding, Haowen Guo, Wensheng Cheng, Ting Pan, and Wen Yang. Mask obb: A semantic attention-based mask oriented bounding box representation for multi-category object detection in aerial images. *Remote Sensing*, 11(24):2930, Dec. 2019. 7

- [32] Jinwang Wang, Wen Yang, Heng-Chao Li, Haijian Zhang, and Gui-Song Xia. Learning center probability map for detecting objects in aerial images. *IEEE TGRS*, pages 1–17, Jul. 2020. 7, 8
- [33] Haoran Wei, Yue Zhang, Zhonghan Chang, Hao Li, Hongqi Wang, and Xian Sun. Oriented objects as pairs of middle lines. *ISPRS J. of Photogramm. Remote Sens.*, 169:268–279, Nov. 2020. 7
- [34] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *CVPR*, pages 3974–3983, Jun. 2018. 1, 5, 8
- [35] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *ICCV*, pages 3520–3529, 2021. 2, 7, 8
- [36] Yongchao Xu, Mingtao Fu, Qimeng Wang, Yukang Wang, Kai Chen, Gui-Song Xia, and Xiang Bai. Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE TPAMI*, pages 1–1, Feb. 2020. 2, 7, 8
- [37] Fan Yang, Heng Fan, Peng Chu, Erik Blasch, and Haibin Ling. Clustered object detection in aerial images. In *ICCV*, pages 8311–8320, 2019. 1
- [38] Feng Yang, Wentong Li, Haiwei Hu, Wanyi Li, and Peng Wang. Multi-scale feature integrated attention-based rotation network for object detection in vhr aerial images. *Sensors*, 20(6):1686, 2020. 8
- [39] Jing Yang, Qingshan Liu, and Kaihua Zhang. Stacked hour-glass network for robust facial landmark localisation. In *CVPRW*, pages 2025–2033, 2017. 7
- [40] Xue Yang, Liping Hou, Yue Zhou, Wentao Wang, and Junchi Yan. Dense label encoding for boundary discontinuity free rotation detection. In *CVPR*, pages 15819–15829, 2021. 1, 2, 7, 8
- [41] Xue Yang and Junchi Yan. Arbitrary-oriented object detection with circular smooth label. In *ECCV*, pages 677–694, 2020. 1, 2, 8
- [42] Xue Yang, Junchi Yan, Ziming Feng, and Tao He. R3det: Refined single-stage detector with feature refinement for rotating object. In *AAAI*, pages 3163–3171, 2020. 1, 2, 7, 8
- [43] Xue Yang, Junchi Yan, Qi Ming, Wentao Wang, xiaopeng zhang, and Qi Tian. Rethinking rotated object detection with gaussian wasserstein distance loss. In *ICML*, pages 11830–11841, 2021. 1, 2
- [44] Xue Yang, Junchi Yan, Xiaokang Yang, Jin Tang, Wenlong Liao, and Tao He. Scrdet++: Detecting small, cluttered and rotated objects via instance-level feature denoising and rotation loss smoothing. *arXiv preprint arXiv:2004.13316*, 2020. 2
- [45] Xue Yang, Jirui Yang, Junchi Yan, Yue Zhang, Tengfei Zhang, Zhi Guo, Xian Sun, and Kun Fu. Scrdet: Towards more robust detection for small, cluttered and rotated objects. In *ICCV*, pages 8232–8241, Oct. 2019. 1, 2, 3, 7
- [46] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *ICCV*, pages 9657–9666, Oct. 2019. 1, 2, 3, 5, 6
- [47] Gongjie Zhang, Shijian Lu, and Wei Zhang. Cad-net: A context-aware detection network for objects in remote sensing imagery. *IEEE TGRS*, 57(12):10015–10024, Dec. 2019. 2, 7
- [48] Jialiang Zhang, Lixiang Lin, Jianke Zhu, Yang Li, Yun-chen Chen, Yao Hu, and Steven CH Hoi. Attribute-aware pedestrian detection in a crowd. *IEEE TMM*, 23:3085–3097, 2020. 2
- [49] Jialiang Zhang, Xiongwei Wu, Steven CH Hoi, and Jianke Zhu. Feature agglomeration networks for single stage face detection. *Neurocomputing*, 380:180–189, 2020. 2
- [50] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z. Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, pages 9759–9768, 2020. 2, 6, 7
- [51] Xiaosong Zhang, Fang Wan, Chang Liu, Rongrong Ji, and Qixiang Ye. Freeanchor: Learning to match anchors for visual object detection. In *NeurIPS*, volume 32, pages 147–155, 2019. 2
- [52] Xiaosong Zhang, Fang Wan, Chang Liu, Xiangyang Ji, and Qixiang Ye. Learning to match anchors for visual object detection. *IEEE TPAMI*, (1):1–1, 2021. 2
- [53] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 1, 2, 7, 8
- [54] Haigang Zhu, Xiaogang Chen, Weiqun Dai, Kun Fu, Qixiang Ye, and Jianbin Jiao. Orientation robust object detection in aerial images using deep convolutional neural network. In *ICIP*, pages 3735–3739, Sep. 2015. 5