

HA-RDet: Hybrid Anchor Rotation Detector for Oriented Object Detection

Phuc Nguyen

University of Information Technology, Vietnam

phucnda@gmail.com / 20520276@gm.uit.edu.vn

Abstract

Oriented object detection in aerial images poses a significant challenge due to their varying sizes and orientations. Current state-of-the-art detectors typically rely on either two-stage or one-stage approaches, often employing Anchor-based strategies, which can result in computationally expensive operations due to the redundant number of generated anchors during training. In contrast, Anchor-free mechanisms offer faster processing but suffer from a reduction in the number of training samples, potentially impacting detection accuracy. To address these limitations, we propose the Hybrid-Anchor Rotation Detector (HA-RDet), which combines the advantages of both anchor-based and anchor-free schemes for oriented object detection. By utilizing only one preset anchor for each location on the feature maps and refining these anchors with our Orientation-Aware Convolution technique, HA-RDet achieves competitive accuracies, including 75.41 mAP on DOTA-v1, 65.3 mAP on DIOR-R, and 90.2 mAP on HRSC2016, against current anchor-based state-of-the-art methods, while significantly reducing computational resources. Source code: <https://github.com/PhucNDA/HA-RDet>

1. Introduction

Object detection in aerial images (ODAI) stands as a pivotal challenge within the domains of computer vision and pattern recognition, boasting numerous practical applications [6, 16, 21–23, 32]. Furthermore, there exists a pressing need within the research community to explore advanced technologies capable of autonomously analyzing vast-scale remote sensing images. However, unlike scenes captured from horizontal perspectives, aerial images obtained from an overhead view often feature densely populated objects appearing in various scales, shapes, and arbitrary orientations. Such complexities pose formidable challenges for models, especially those relying on Horizontal Bounding Boxes (HBB) to accurately represent object instances in the image. To address this, Oriented Bounding Boxes (OBB) are utilized to offer a more precise depiction of object lo-

cations by incorporating additional directional information. Notably, several significant challenges emerge in the realm of detecting objects in aerial images, including:

- **Large aspect ratio:** Objects in aerial images often have irregular shapes, resulting in high aspect ratios, such as bridges, ships, harbors, etc.
- **Scale variations:** Varying ground sampling distance (GSD) of sensors leads to scale variation among objects captured by different sensors in the same scene [27].
- **Dense arrangement:** Aerial images usually contain multiple objects distributed densely, such as ships in a harbor or vehicles in a parking lot, which can pose a big challenge for object detection algorithms as they need to accurately identify and distinguish object instances.
- **Arbitrary orientations:** Objects in aerial images can appear in various orientations, requiring object detection models to have precise orientation estimation capabilities.

Existing oriented detectors [3, 5, 9, 29, 30] typically rely on proposal-driven frameworks utilizing a Region Proposal Network (RPN), which employs numerous anchors with varying scales and aspect ratios to generate Regions of Interest (RoIs). However, horizontal RoIs often cover multiple instances, causing ambiguity, while rotated RoIs improve recall but incur high computational costs. Hence, developing an efficient region proposal network capable of generating high-quality proposals is crucial to overcome accuracy and computational limitations in current state-of-the-art oriented detectors. We identify two key reasons for the lack of efficiency in region proposal-based oriented detectors: **(1)** Most oriented object detection methods adopt Anchor-based or Anchor-free training approaches, each with its own challenges and trade-offs in terms of performance and computational efficiency. Anchor-based methods achieve good performance but are sensitive to oriented anchor hyperparameters, while Anchor-free methods accelerate training and inference speed but suffer from decreased detection performance due to shortage number of preset training anchors. **(2)** The conventional convolution technique, designed with fixed receptive fields and aligned with axes, is ill-suited for objects with arbitrary orientations and varying scales in aerial images. Predefined anchor represen-

tations often fail to fully capture the shape and orientation of objects, leading to misalignment between fixed convolution features and object orientations, thus hindering effective object localization and classification in oriented object detection models.

This study presents a novel method, Hybrid-Anchor Rotation Detector (HA-RDet), for oriented object detection. HA-RDet combines the strengths of anchor-based and anchor-free paradigms through a hybrid anchor training scheme that employs a single anchor per location on the feature map. At the core of our design is the proposed Orientation-Aware Convolution (O-AwareConv), which serves as a bridge between one-stage and two-stage detection heads. O-AwareConv improves proposal quality by dynamically adapting to the object’s shape and orientation. Extensive experiments on standard oriented object detection benchmarks demonstrate the effectiveness of our approach. HA-RDet addresses key challenges such as insufficient positive training samples and over-encoding of geometric information by initially reducing anchor hyperparameters and subsequently refining proposals in later stages. Unlike prior methods, HA-RDet achieves a superior balance between accuracy and computational efficiency in both training and inference. It requires fewer resources while maintaining competitive performance, thereby accelerating the detection process. Overall, HA-RDet offers an elegant and efficient region proposal-based solution that bridges the gap between anchor-based and anchor-free approaches, providing a strong baseline for future research in oriented object detection.

Our contributions are summarized as follows:

- Presenting the Hybrid-Anchor Rotation Detector that innovatively combines anchor-based and anchor-free techniques for oriented object detection
- Introducing the novel O-AwareConv for enhanced orientation-sensitive feature extraction
- Extensive experiments are conducted on challenging oriented object detection datasets to demonstrate the effectiveness of our proposed method.

2. Related works

With the surge of machine learning, particularly deep learning, object detection has witnessed significant advancements, dividing into two main categories: two-stage detectors and one-stage detectors. Two-stage detectors begin by generating a sparse set of Regions of Interest (RoIs) in the first stage, followed by RoI-wise bounding box regression and object classification in the second stage. In contrast, one-stage detectors directly detect objects without the need for RoI generation. Furthermore, two fundamental training strategies, Anchor-free and Anchor-based methods, are crucial for object detection. While Anchor-free methods excel in rapid object detection, they suffer from a scarcity of

positive training samples, resulting in performance degradation. Conversely, Anchor-based methods generate numerous anchors, facilitating an abundant supply of positive training samples and achieving excellent performance, albeit at a higher computational cost. Our HA-RDet integrates both training strategies, utilizing only one preset anchor for each feature map’s location, thereby significantly reducing resource requirements while maintaining asymptotically state-of-the-art performance.

2.1. Anchor-Based Detectors

Most methods in object detection rely on predefined generated anchors, where multiple preset anchors with varying scales and ratios are assigned to each specific location on a feature map. For instance, ROI Transformer generates preset horizontal anchors and then applies RRoi Leaner and RRoi Wrapping for robust feature extraction. Oriented R-CNN proposed by Xie et al. [29] adopts an anchor-based approach with an Oriented Region Proposal Network, along with the Midpoint-offset representation. This representation encodes horizontal anchors into six distinguishable learnable parameters, ensuring a reasonable transformation from horizontal generated anchors to oriented proposals. Similarly, Faster R-CNN [23] introduces the Region Proposal Network with nine multi-scale heuristically tuned anchors for every location on a feature map. These densely generated anchors are assigned either positive or negative labels based on the Intersection over Union (IoU) with a ground-truth box. Additionally, SSD [17] enhances detector efficiency by directly predicting object location offsets and corresponding categories. While anchor-based methods are well-studied and yield remarkable performance by generating a large number of anchors with different ratios and scales for each location, they require significant computational resources for training and inference, making them impractical for oriented object detection in aerial images datasets.

2.2. Anchor-Free Detectors

The anchor-free mechanism differs from anchor-based strategies by utilizing only one key point for each location on a feature map, implying that each point on the feature map corresponds to a single anchor with a specific scale and ratio. S^2 A-Net [8] introduced the single-shot alignment network featuring the Feature Alignment Module and Oriented Detection Module, which ensure feature alignment using the Alignment Convolution specifically designed for oriented feature extraction. YOLO [22] divides the image into an $S \times S$ grid, with each grid cell responsible for determining the presence of an object’s center. CornerNet [12] predicts the final bounding box using a pair of key points (top-left and bottom-right) and undergoes a complex post-processing phase to combine key points of the

same object. FCOS [24] assigns all locations within an object’s ground truth as positive, incorporating a ”centerness” branch to suppress low-quality bounding boxes, thereby enhancing overall performance. Despite being faster, anchor-free methods may experience a performance decline compared to previous anchor-based methods.

3. Proposed Hybrid-Anchor Rotation Detector

Most oriented object detectors follow the anchor-based strategy for predicting the 5-D encoded oriented bounding boxes. They often generate multiple anchors of different ratios and scales to allocate many instances of various shapes and sizes *greedily*. This approach requires a large number of anchors, leading to high computational requirements. On the other hand, the anchor-free mechanism uses only a single anchor per location on a feature map, but it suffers from accuracy deterioration due to the lack of positive training samples. Many studies improve the performance of either method via post-refining, new assigner strategies, new anchor definitions, etc. Two strategies standstill and need to be fully exploited in oriented object detection.

In this section, we reason our proposals and explain how the strategy is applied and operated appropriately under different anchor encoder-decoder fashion.

3.1. Hybrid-Anchor RPN

Anchor encoder-decoder strategy. We begin by examining the vanilla baseline of HA-RDet, which builds on RetinaNet [16]—a detector originally developed for generic horizontal object detection—adapted here for oriented object detection by representing anchors in a 5D format (x, y, w, h, θ) . The Anchor-free head, guided by an Anchor-free assigner, is trained to generate oriented region proposals directly from feature maps produced by standard 2D convolution (2DConv). In the next stage, the Anchor-based head refines these anchors using a 5D encoder-decoder to regress offsets $(\delta x, \delta y, \delta w, \delta h, \delta \theta)$, resulting in the final oriented object proposals. However, this experimental setup proves ineffective for two main reasons, both of which compromise proposal quality and consequently degrade the performance of the anchor-based head: **(1) Limited positive training samples results in insufficient and unreliable proposals from the anchor-free generator;** **(2) Overly complex five-dimensional anchor encoding in the anchor-free head introduces noise and reduces proposal precision.**

Our proposed approach (depicted in Fig. 1) tackles these challenges by introducing the Hybrid-Anchor RPN, which simplifies anchor encoding from 5D (x, y, w, h, θ) to 4D (x, y, w, h) while ensuring consistency between the two heads. We achieve this by utilizing the Anchor-free assigner strategy on rectangularized oriented ground truths for horizontal anchors, significantly increasing the number of positive samples. This augmentation aids the model in learning

sufficient information for further refinement. Subsequently, the resulting horizontal anchors undergo processing in the Anchor-based stage, producing high-quality horizontal proposals for downstream tasks. This design guarantees an adequate number of positive samples for training and gradually enhances the quality of the final proposals. Detailed insights into the algorithmic implementation of the Hybrid-Anchor RPN are provided in Algorithm 1.

Anchor-Free head. We use one preset anchor with a single scale and ratio for each location on a feature map. In the Anchor-free head, the normal RetinaNet [16] architecture used as a backbone encodes and decodes the horizontal anchors (a_x, a_y, a_w, a_h) with the rectangularized ground truth target (t_x, t_y, t_w, t_h) as the following Eq. 1:

$$\begin{cases} \delta_x = (t_x - a_x)/a_w, & \delta_y = (t_y - a_y)/a_h \\ \delta_w = \log(t_w/a_w), & \delta_h = \log(t_h/a_h) \\ a'_x = \widehat{\delta}_x a_w + a_x, & a'_y = \widehat{\delta}_y a_h + a_y \\ a'_w = a'_w \exp(\widehat{\delta}_w), & a'_h = a'_h \exp(\widehat{\delta}_h) \end{cases} \quad (1)$$

where $\delta = (\delta_x, \delta_y, \delta_w, \delta_h)$ and $\widehat{\delta} = (\widehat{\delta}_x, \widehat{\delta}_y, \widehat{\delta}_w, \widehat{\delta}_h)$ are predicted encoded value and target encoded value, respectively. Having regressed $\widehat{\delta}$, we can easily decode the proposal (a'_x, a'_y, a'_w, a'_h) . With the classification branch omitted, the Anchor-free head aims to minimize the bounding box loss function L_{AF} :

$$L_{AF}(\delta, \widehat{\delta}) = w^{AF} \left[-\ln \left(\frac{\text{Intersection}(\delta, \widehat{\delta})}{\text{Union}(\delta, \widehat{\delta})} \right) \right] \quad (2)$$

where w^{AF} is the loss weight used to fine-tune the model. At the end of this stage, from a limited number of preset anchors, we obtain a preliminary set of horizontal candidates, prepared for further refinement.

Orientation-aware Convolution. To extract meaningful geometric cues from the set of predicted candidates, we propose orientation-aware convolution (O-AwareConv), which improves the alignment of global features and allows the extraction of orientation-sensitive features for horizontal anchor candidates. O-AwareConv adaptively incorporates the orientation of the ground truth object θ_{GT} during training. Specifically, it calculates a shape offset field for each horizontal anchor, in combination with an orientation offset derived from the rotation matrix $R(\theta_{GT})^T$, where T denotes the transposition of the matrix. These offsets guide the convolution to adjust its sampling locations, enabling it to align with the true orientation and geometry of objects. As illustrated in Fig. 2, this design transforms conventional convolutional features into orientation-aware representations, tailored to each anchor-ground-truth pair.

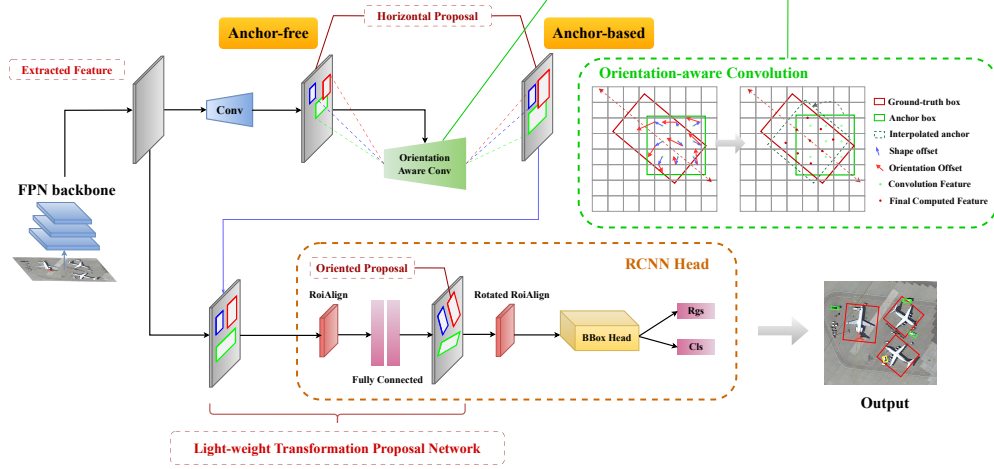


Figure 1. Illustration of the architecture of the proposed HA-RDet. At first, an aerial image is passed through the FPN backbone to extract deep features. Then, the extracted features are fed through the Hybrid-Anchored RPN to produce horizontal anchors via an Anchor-free assigner strategy on rectangularized oriented ground truths. The anchors are refined with the novel Orientation-aware Convolution and the later Anchor-based head to produce high-quality horizontal proposals. After that, the lightweight proposal transformation network learns to produce oriented proposals from predicted RoI Align processed horizontal proposals. Finally, these oriented proposals are refined through oriented bounding box heads for classification and bounding box regression.

Formally, for each spatial location p on a feature map X , O-AwareConv uses the anchor candidate $a^* = (x, y, w, h)$, stride S , kernel weights W , and kernel size k to calculate the position Ω of the final compute feature Y as defined in Eq. 3:

$$\begin{cases} Y(p) = \sum_{r \in \text{Kernels}; o \in \Omega} W(r) \cdot X(p + r + o) \\ \Omega = \frac{1}{S}((\mathbf{x}, \mathbf{y}) + \frac{1}{k}(\mathbf{w}, \mathbf{h}) \cdot r \cdot R(\theta_{GT})^T) - p - r \\ R(\alpha) = \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix} \end{cases} \quad (3)$$

In our approach, each anchor is assigned to at most one ground-truth object, while each ground truth can correspond to multiple anchors. To make full use of the training data, we retain all anchor candidates, including negative ones, in the previous anchor-free heads as a set of potential predefined candidates, anticipating that their quality will improve progressively during the refinement stage. During inference, when ground-truth annotations are not available, the orientation offset field is set to zero, ensuring deterministic feature offset computation.

Compared to standard 2D convolution, which assumes objects are axis-aligned, our method incorporates Deformable Convolution (DeformConv) [2] with augmented offsets to improve spatial adaptability. However, DeformConv can still sample from suboptimal locations, especially under weak supervision and in densely populated scenes.

To mitigate this, our O-AwareConv is used for feature extraction during training for each generated horizontal anchor, enabling robust orientation-aware alignment within the Hybrid-Anchored RPN. Unlike 2DConv and DeformConv, the offset field in O-AwareConv is explicitly derived from the object’s shape and orientation. As shown in Fig. 2, this mechanism allows for more accurate and geometry-aligned feature extraction. During inference, O-AwareConv effectively captures global features without relying on orientation supervision, thereby improving anchor prediction accuracy. At the end of this stage, we obtain a set of horizontal candidates from the Anchor-free head, along with their corresponding orientation-aware features Y extracted by O-AwareConv.

Anchor-Based head. The anchor-based head serves as a refinement module. It re-encodes the input according to Eq. 1, taking as input the anchors generated in the previous stage along with their features extracted by the orientation-aware convolution. Its goal is to produce refined, high-quality proposals. The Anchor-based bounding box loss function L_{AB} is derived from Eq. 2, albeit with a different weight w^{AB} . However, unlike Anchor-free, the classification branch with binary cross-entropy loss is used in this stage to estimate the objectness of the predicted region proposal. The overall loss of our Region Proposal Network is the sum of the Anchor-free loss L_{AF} and the Anchor-based loss L_{AB} , denoted as $L_{RPN} = L_{AF} + L_{AB}$. Subsequently, we obtain the final set of horizontal region proposals, which effectively capture orientation-sensitive objects and are ready to be transformed into oriented proposals.




Figure 2. The shape offsets are computed via anchor size; ground-truth orientation is adopted to perform orientation offsets calculation. Shape offsets and orientation offsets produce the final features used for training the Anchor-based refinement stage. The final computed features are orientation-awareness, surging HA-RDet performance significantly.

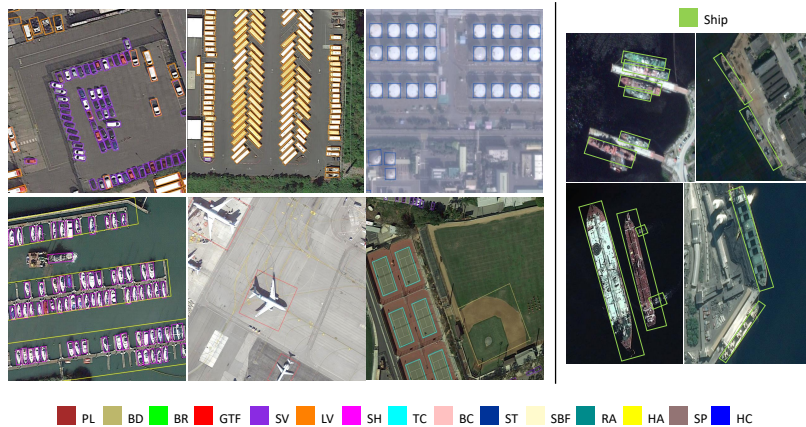


Figure 3. **Left:** Detection results on DOTA. **Right:** Detection results on HRSC2016

3.2. R-CNN oriented bounding box heads

Having regressed well-qualified horizontal proposals from the Hybrid-Anchor RPN, we implemented a lightweight multi-layer neural network to transform the horizontal to oriented region proposals. RoIAlign [10] with bilinear interpolation is adopted to our baseline to extract the region of interest features. We then use 5D offset value representation $(\delta x, \delta y, \delta w, \delta h, \delta \theta)$ to encode the oriented proposals from horizontal ones and perform simple regression using two fully-connected layers.

For the oriented bounding box head, we adopt the Rotated RoIAlign operation [29] to extract rotation-sensitivity features from each oriented proposal.

4. Experiments

4.1. Datasets

DOTA-v1.0 [28] is a large-scale dataset comprising high-quality aerial images for detecting objects, containing 2,806 high-defined images with resolutions 4000×4000 . It classifies annotated instances into 15 common categories: plane (PL), baseball diamond (BD), bridge (BR), ground track field (GTF), small vehicle (SV), large vehicle (LV), ship (SH), tennis court (TC), basketball court (BC), storage

tank (ST), soccer ball field (SBF), roundabout (RA), harbor (HA), swimming pool (SP) and helicopter (HC). During the training and testing process, the images are divided into cropped patches of size 1024×1024 pixels with a stride of 824 pixels, including random flipping is applied as a form of data augmentation. We submit our test set results to the evaluation website.

DIOR-R [1] is an extended version of the DIOR dataset [13] that shares the same 23,463 optimal remote sensing images but is labeled with oriented bounding boxes. The images are 800×800 pixels and vary in spatial resolution. Each object in the image is categorized into one of 20 classes. The training procedure for DIOR-R follows the same one implemented on the DOTA-v1.0 dataset.

HRSC2016 [18] dataset is a collection of 1061 high-resolution aerial images with oriented bounding box annotations for ship detection. The image sizes range from 300×300 to 1500×900 . This dataset encompasses over 25 categories of ships with a large diversity of shapes, scales, ratios, and orientations annotated. During the training process, the images are resized to 800×512 pixels while preserving their original aspect ratio. Additionally, random flipping is applied to augment the training data.

Algorithm 1: Hybrid-Anchor training strategy

Input:

- x : Multi-level features
- gt : Ground truth
- rgt : Rectangularized ground truth

Output:

- Region proposals

Anchor-free stage

```
1. A = anchor_generator(X)
2. target = assigner(A, rgt)['index']
3. X = 2Dconv(X)
4. bbox_pred = regression(X)
5. A = decode(A, bbox_pred)
6. loss_af = box_loss(A, rgt[target])
```

Anchor-based stage

```
7. target = assigner(A, rgt)['index']
8. offset = shape_offset(A)
9. or_offset = or_offset(A, gt[target])
10. offset = dot(offset, or_offset)
11. X = OAconv(X, offset)
12. bbox_pred = regression(X)
13. bbox_cls = classification(X)
14. A = decode(A, bbox_pred)
15. loss_rgs = box_loss(A, rgt[target])
16. loss_cls = objectness(bbox_cls, label[target])
17. loss_ab = loss_rgs + loss_cls
```

Return

```
19. loss_rpn = loss_af + loss_ab
20. proposals = refinement(A, bbox_cls)
```

APPENDIX

- $target$: Target ground-truth
 - $loss_{af}$: Anchor-free loss
 - $loss_{ab}$: Anchor-based loss
 - $decode(*)$: Bounding box decoder
 - $dot(*)$: Dot product mathematical operation
 - $assigner(*)$: Anchor assigner strategy
 - $OAconv(*)$: Orientation-aware Convolution
 - $shape_offset(*)$: Get Shape offset
 - $or_offset(*)$: Get Orientation offset
 - $refinement(*)$: Box refinement via NMS
-

4.2. Implementation details

We utilize two backbone architectures: ResNet50+FPN and ResNet101+FPN, as well as ResNext101-FPN-DCNv2, all pretrained on ImageNet [4]. Our optimization setup in-

cludes the SGD optimizer with an initial learning rate of 0.005 and a decay factor of 10 for each decay step. The weights w^{AF} and w^{AB} are both set to 7.0. We set the momentum to 0.9 and the weight decay to 0.0001. The anchor’s scale and ratio is set to 4.0 and 1.0, respectively. For the anchor assigner strategy, we employ the Ratio-based strategy (positive: 0.3; negative: 0.1) [26] for the Anchor-Free head and the Max-IoU strategy (positive: 0.7; negative: 0.3) for the Anchor-based head. Both assigning strategies are followed by random balanced sampling of 256 positive and 256 negative samples. We train HA-RDet for 12 epochs on the DOTA and DIOR-R datasets and 36 epochs on the HRSC2016 dataset. Training is conducted using a single A5000 GPU with a batch size of 2, while inference is performed on a single A5000 GPU.

4.3. Comparisons with State-of-the-Arts

The experimental results in Tab. 1 highlight the performance of HA-RDet on the DOTA-v1 dataset in comparison to prior detectors. While recent methods leveraging ViT-based backbones [15, 25, 33], end-to-end training paradigms [35] and different encoding-decoding [34] mechanisms have achieved superior performance—surpassing both traditional one-stage and two-stage detectors as well as our HA-RDet—this study focuses on fair comparisons under the same configuration (Oriented-RCNN, RoI Transformer, S^2ANet). **Specifically, we evaluate HA-RDet against methods using the same backbone architecture and data augmentation strategies.** Future work on these adjustments would be considered.

Under these settings, HA-RDet demonstrates prominent performance, achieving 75.41 mAP and 76.02 mAP with ResNet-50 and ResNet-101 backbones, respectively. The highest accuracy of 77.01 mAP is attained using ResNeXt101-DCNv2. HA-RDet surpasses most anchor-free methods and delivers results competitive with the well-established two-stage Oriented R-CNN. The incorporation of Orientation-Aware Convolution enables HA-RDet to better align features with object geometry, leading to significantly improved detection accuracy for irregularly shaped objects—such as bridges, storage tanks, and harbors—effectively addressing the challenges posed by high aspect ratios.

On the HRSC2016 dataset, as shown in Tab. 2, HA-RDet achieves competitive results of 90.20 mAP and 95.32 mAP on the VOC2007 and VOC2012 metrics, respectively. These results outperform other training methods that use multiple anchors and demonstrate competitiveness against the S^2A -Net, which follows the same training scheme of only using one anchor per location.

Tab. 3 presents the results obtained from the DIOR-R dataset, where our novel Hybrid-Anchor training schemes approach demonstrates superior performance. Using the

Model	Backbones	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	mAP
<i>One-stage</i>																	
RSDet [20]	ResNet50+FPN	89.3	82.7	47.7	63.9	66.8	62	67.3	90.8	85.3	82.4	62.3	62.4	65.7	68.6	64.6	70.8
SASM [11]	ResNet50+FPN	86.42	78.97	52.47	69.84	77.30	75.99	86.72	90.89	82.63	85.66	60.13	68.25	73.98	72.22	62.37	74.92
PSCD [34]	ResNet50+FPN	89.32	82.29	37.92	71.52	78.40	66.33	78.01	90.89	84.21	80.63	60.22	64.73	59.69	68.37	53.85	71.09
S^2 A-Net [8]	ResNet50+FPN	89.3	80.49	50.42	73.23	78.42	77.4	86.8	90.89	85.66	84.24	62.16	65.93	66.66	67.76	53.56	74.19
<i>Two-stage</i>																	
RoI Transformer [5]	ResNet101+FPN	88.64	78.52	43.44	75.92	68.81	73.68	83.59	90.74	77.27	81.46	58.39	53.54	62.83	58.93	47.67	69.56
ARC [19]	ARC-ResNet50+FPN	89.40	82.48	55.33	73.88	79.37	84.05	88.06	90.90	86.44	84.83	63.63	70.32	74.29	71.91	65.43	77.35
Oriented R-CNN [29]	ResNet50+FPN	89.35	81.41	52.6	75.02	79.03	82.41	87.82	90.9	86.4	85.3	63.36	65.7	68.28	70.48	57.23	75.69
ReDet [9]	ReResNet50+ReFPN	88.79	82.64	53.97	74.00	78.13	84.06	88.04	90.89	87.78	85.75	61.76	60.39	75.96	68.07	63.59	76.25
<i>Anchor-free</i>																	
Rotated RepPoints [32]	ResNet50+FPN	83.36	63.71	36.27	51.58	71.06	50.35	72.42	90.1	70.22	81.98	47.46	59.5	50.65	55.51	53.07	59.15
CFA [7]	ResNet101+FPN	89.26	81.72	51.81	67.17	79.99	78.25	84.46	90.77	83.4	85.54	54.86	67.75	73.04	70.24	64.96	75.05
AOPG [1]	ResNet50+FPN	89.27	83.49	52.50	69.97	73.51	82.31	87.95	90.89	87.64	84.71	60.01	66.12	74.19	68.30	57.80	75.24
<i>Ours</i>																	
HA-RDet	ResNet50+FPN	88.98	83.34	55.02	72.04	78.25	81.43	86.94	90.91	85.46	85.48	61.33	61.57	74.50	68.97	56.86	75.408
HA-RDet	ResNet101+FPN	89.12	84.95	54.81	69.33	77.67	82.91	87.54	90.87	86.06	86.03	64.19	61.90	74.25	70.99	65.43	76.02
HA-RDet	ResNeXt101_DCNv2+FPN	89.31	85.01	57.03	68.94	78.14	83.35	87.61	90.91	87.32	86.67	65.14	63.05	76.35	72.29	63.98	77.012

Table 1. Comparison of SOTA results on the oriented object detection DOTA 1.0 task. We report the results of methods using the same training configuration on DOTA 1.0

Model	R3Det [31]	AO2-DETR [3]	Gliding Vertex [30]	S^2 A-Net [8]	ReDet [9]	AOPG [1]	Oriented RepPoints [14]	Oriented R-CNN [29]	HA-RDet
#Anchors	21	-	20	1	20	1	9	20	1
mAP	87.14	88.12	88.20	90.14	90.2	90.34	90.38	90.4	90.20 / 95.32*

Table 2. Results on HRSC2016 with #Anchors represents the quantity of anchors generated per location on a feature map.

backbone ResNet50, our method achieves the highest accuracy at 65.3 mAP, surpassing the well-known two-stage Oriented R-CNN by 0.2 mAP, the one-stage S^2 A-Net by 3.6 mAP, and the Anchor-free Oriented Proposal Generator (AOPG) by 0.9 mAP.

We conduct further experiments to compare our proposed Hybrid-Anchor training scheme with the well-known anchor-free S^2 A-Net and anchor-based Oriented R-CNN methods in Tab. 5. We evaluate these methods beyond the mean Average Precision (mAP) measure, such as inference speed (FPS), computational resources (VRAM) for storing as well as processing the intermediate data during the training process, and the number of model parameters. The S^2 A-Net reached the detection result of 74.19 mAP with a fast inference speed of 15.5 FPS; it consumed 4.6 GB of VRAM due to its relatively low number of parameters. Meanwhile, Oriented R-CNN, with 20 anchors in the training stage, attained a slightly better accuracy than S^2 A-Net but with a trade-off regarding inference speed (13.5 FPS) and higher computational resource requirements (14.2 GB of VRAM). Our HA-RDet follows the Hybrid-Anchor training approach and generates only one anchor per location, similar to S^2 A-Net. It achieved better detection accuracy (75.41 mAP) while maintaining a stable inference speed of 14.9 FPS. Compared to Oriented R-CNN, our model has a higher parameter requirement for processing

but consumes approximately half the amount of VRAM (6.9 GB) due to the significant reduction in the number of generated anchors. Oriented R-CNN utilized dense generated anchors and a complex encoding strategy; therefore, its accuracy and inference speed is slightly higher than our model, but it comes at the cost of significantly higher computational resource requirements. Some visualization of HA-RDet are demonstrated in Fig. 3

4.4. Ablation studies

Effectiveness of O-AwareConv. As mentioned in section 3.1, we introduced an approach called O-AwareConv that enhances the awareness of ground-truth orientation. Tab. 6 compares O-AwareConv against alternative methods under the same configurations to demonstrate its effectiveness in oriented object detection. Our HA-RDet employs Standard Convolution and reaches only 74.1 mAP because the fixed receptive field fails to extract orientation-sensitive features. On the other hand, Deformable Convolution addresses feature misalignment, yet its performance degrades when detecting densely packed and arbitrarily oriented objects in aerial images. O-AwareConv significantly enhances the performance of our proposed detector, surpassing Deformable Convolution by 0.91 mAP and outperforming standard Convolution by 1.31 mAP.

Effectiveness of Hybrid-Anchor mechanism. To eval-

Model	R3Det [31]	CFA [7]	S^2 A-Net [8]	Rotated FCOS [24]	RoI Transformer [5]	AOPG [1]	Oriented R-CNN [29]	HA-RDet
mAP	57.2	60.8	61.7	62.4	63.9	64.4	65.1	65.3

Table 3. Experimental comparison of SOTA methods on DIOR-R, trained using **ResNet50+FPN**

Model (anchor scale)	Backbones	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	mAP
HA-RDet (4)	ResNet50 + FPN	88.98	83.34	55.02	72.04	78.25	81.43	86.94	90.91	85.46	85.48	61.33	61.57	74.50	68.97	56.86	75.408
HA-RDet (6)	ResNet50 + FPN	89.01	81.84	54.34	73.72	72.59	77.52	86.88	90.9	87.66	85.75	61.55	61.91	75.69	66.54	59.72	75.032
HA-RDet (8)	ResNet50 + FPN	88.97	80.79	53.18	70.64	70.96	76.49	78.98	90.9	85.57	84.58	59.29	60.83	75.91	66.49	53.28	73.122

Table 4. Models with anchor scales above 6 perform better on large objects, while smaller scales yield higher accuracy for small instances

	S^2 A-Net [8]	Oriented R-CNN [29]	HA-RDet (Ours)
#Anchors	1	20	1
mAP	74.19	75.69	75.41
FPS	15.5	13.5	14.9
VRAM (GB)	4.6	14.2	6.9
#params	38,599,976	41,139,754	55,715,449

Table 5. Evaluation results of HA-RDet with S^2 A-Net and Oriented R-CNN on the **ResNet50+FPN** on DOTA. Anchors are saved under CUDA tensors for fast computation

Convolution method	mAP
Standard Convolution	74.1
Deformable Convolution [2]	74.5 (\uparrow 0.4)
Orientation-aware Convolution (Ours)	75.41 (\uparrow 1.3)

Table 6. Comparison of Orientation-aware Convolution accuracy with other convolutions on the DOTA-v1 dataset.

uate the effectiveness of the Anchor-free head, Anchor-based head, and O-AwareConv in our RPN, we experiment with different HA-RDet settings where omitting certain branches, shown in Tab. 7. When we excluded the Anchor-free head and relied solely on the anchor-based mechanism, the accuracy slightly decreased to 73.153 mAP compared to using the entire component on DOTA-v1. The role of the Anchor-free head is to provide initial proposals not constrained by predefined anchor boxes. By removing this module, we argue that the detector loses its adaptability to various object sizes and aspect ratios, which can lead to a decrease in performance.

On the other hand, taking out the Anchor-based head indicates a terrible accuracy of 0.038 mAP. The Anchor-based head is responsible for refining the sparsely initial proposals and generating high-quality region proposals. Without this component, the detector lacks the refinement process, leading to imprecise and less reliable region proposals. Furthermore, without the Anchor-based head, the detector heavily relies on the initial proposals generated by the Anchor-free

Module	Our Hybrid-Anchor RPN		
Anchor-free head	✓	✓	-
Anchor-based head	✓	-	✓
O-AwareConv	✓	-	-
mAP	75.408	0.038	73.153

Table 7. Experimental evaluation of our Hybrid-Anchor RPN by taking out a couple of modules, using ResNet50+FPN on DOTA.

head, which may not be as effective in handling objects with different scales, aspect ratios, or complex backgrounds.

These configurations neglect the use of O-Aware Convolution, which bridges the Anchor-free and Anchor-based heads in our Hybrid-Anchor RPN. By disregarding one of these modules, the detector loses the refinement process, accurately predicting bounding box regression, and the ability to handle object shape and orientation information, resulting in a significant decline in performance.

Study on anchor scale. We use one preset anchor with one scale and ratio for each location on a feature map following the anchor-free mechanism Tab. 4. Then, we extend the training longer, and the diminishing of samples happens in models having anchor scale larger than 6 due to the overwhelming number of small-sized instances of the datasets. Finally, we decided to choose the anchor scale of 4 with ratio of 1.0 as the representative for our HA-RDet baseline given its stability, easy convergence, and suitability.

5. Discussion

We propose HA-RDet, a Hybrid-Anchor training scheme that integrates the novel O-AwareConv to tackle two critical challenges in oriented object detection: the scarcity of positive samples and the risk of over-encoding geometric information. HA-RDet bridges the gap between one-stage and two-stage detection paradigms, striking a balance between training/inference efficiency and detection accuracy. Our method is particularly well-suited for Earth Observation applications, where aerial and satellite imagery often contain densely packed, arbitrarily oriented objects.

References

- [1] Gong Cheng, Jiabao Wang, Ke Li, Xingxing Xie, Chunbo Lang, Yanqing Yao, and Junwei Han. Anchor-free oriented proposal generator for object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–11, 2022. 5, 7, 8
- [2] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017. 4, 8
- [3] Linhui Dai, Hong Liu, Hao Tang, Zhiwei Wu, and Pinhao Song. Ao2-detr: Arbitrary-oriented object detection transformer. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(5):2342–2356, 2023. 1, 7
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [5] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2849–2858, 2019. 1, 7, 8
- [6] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 1
- [7] Zonghao Guo, Chang Liu, Xiaosong Zhang, Jianbin Jiao, Xi-angyang Ji, and Qixiang Ye. Beyond bounding-box: Convex-hull feature adaptation for oriented and densely packed object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8792–8801, 2021. 7, 8
- [8] Jiaming Han, Jian Ding, Jie Li, and Gui-Song Xia. Align deep features for oriented object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–11, 2021. 2, 7, 8
- [9] Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotation-equivariant detector for aerial object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2786–2795, 2021. 1, 7
- [10] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 5
- [11] Liping Hou, Ke Lu, Jian Xue, and Yuqiu Li. Shape-adaptive selection and measurement for oriented object detection. In *Proceedings of the AAAI conference on artificial intelligence*, pages 923–932, 2022. 7
- [12] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018. 2
- [13] 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 journal of photogrammetry and remote sensing*, 159:296–307, 2020. 5
- [14] Wentong Li, Yijie Chen, Kaixuan Hu, and Jianke Zhu. Oriented reppoints for aerial object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1829–1838, 2022. 7
- [15] Yuxuan Li, Qibin Hou, Zhaohui Zheng, Ming-Ming Cheng, Jian Yang, and Xiang Li. Large selective kernel network for remote sensing object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16794–16805, 2023. 6
- [16] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 1, 3
- [17] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 2
- [18] Zikun Liu, Liu Yuan, Lubin Weng, and Yiping Yang. A high resolution optical satellite image dataset for ship recognition and some new baselines. In *International conference on pattern recognition applications and methods*, pages 324–331. SciTePress, 2017. 5
- [19] Yifan Pu, Yiru Wang, Zhuofan Xia, Yizeng Han, Yulin Wang, Weihao Gan, Zidong Wang, Shiji Song, and Gao Huang. Adaptive rotated convolution for rotated object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6589–6600, 2023. 7
- [20] Wen Qian, Xue Yang, Silong Peng, Xiujuan Zhang, and Junchi Yan. Rsdet++: Point-based modulated loss for more accurate rotated object detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 2022. 7
- [21] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017. 1
- [22] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 1, 2
- [24] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 3, 8
- [25] Di Wang, Qiming Zhang, Yufei Xu, Jing Zhang, Bo Du, Dacheng Tao, and Liangpei Zhang. Advancing plain vision transformer toward remote sensing foundation model. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–15, 2022. 6
- [26] Jiaqi Wang, Kai Chen, Shuo Yang, Chen Change Loy, and Dahua Lin. Region proposal by guided anchoring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2965–2974, 2019. 6
- [27] Zhenyu Wu and Haibin Yan. Adaptive dynamic networks for object detection in aerial images. *Pattern Recognition Letters*, 166:8–15, 2023. 1

- [28] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3974–3983, 2018. 5
- [29] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3520–3529, 2021. 1, 2, 5, 7, 8
- [30] 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 transactions on pattern analysis and machine intelligence*, 43(4):1452–1459, 2020. 1, 7
- [31] Xue Yang, Junchi Yan, Ziming Feng, and Tao He. R3det: Refined single-stage detector with feature refinement for rotating object. In *Proceedings of the AAAI conference on artificial intelligence*, pages 3163–3171, 2021. 7, 8
- [32] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9657–9666, 2019. 1, 7
- [33] Hongtian Yu, Yunjie Tian, Qixiang Ye, and Yunfan Liu. Spatial transform decoupling for oriented object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6782–6790, 2024. 6
- [34] Yi Yu and Feipeng Da. Phase-shifting coder: Predicting accurate orientation in oriented object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13354–13363, 2023. 6, 7
- [35] Cong Zhang, Jingran Su, Yakun Ju, Kin-Man Lam, and Qi Wang. Efficient inductive vision transformer for oriented object detection in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–20, 2023. 6