

# Measuring the Impact of Rotation Equivariance on Aerial Object Detection

Xiuyu Wu<sup>1</sup> Xinhao Wang<sup>1</sup> Xiubin Zhu<sup>1</sup> Lan Yang<sup>1</sup> Jiyuan Liu<sup>2</sup> Xingchen Hu<sup>2\*</sup>  
<sup>1</sup>Xidian University <sup>2</sup>National University of Defense Technology

{xiuyuwu, xhaowang}@stu.xidian.edu.cn, xbzhu@mail.xidian.edu.cn, yanglan@xidian.edu.cn,  
 liujiyuan13@nudt.edu.cn, xhu4@ualberta.ca

## Abstract

Due to the arbitrary orientation of objects in aerial images, rotation equivariance is a critical property for aerial object detectors. However, recent studies on rotation-equivariant aerial object detection remain scarce. Most detectors rely on data augmentation to enable models to learn approximately rotation-equivariant features. A few detectors have constructed rotation-equivariant networks, but due to the breaking of strict rotation equivariance by typical downsampling processes, these networks only achieve approximately rotation-equivariant backbones. Whether strict rotation equivariance is necessary for aerial image object detection remains an open question. In this paper, we implement a strictly rotation-equivariant backbone and neck network with a more advanced network structure and compare it with approximately rotation-equivariant networks to quantitatively measure the impact of rotation equivariance on the performance of aerial image detectors. Additionally, leveraging the inherently grouped nature of rotation-equivariant features, we propose a multi-branch head network that reduces the parameter count while improving detection accuracy. Based on the aforementioned improvements, this study proposes the **Multi-branch head rotation-equivariant single-stage Detector (MessDet)**, which achieves state-of-the-art performance on the challenging aerial image datasets *DOTA-v1.0*, *DOTA-v1.5* and *DIOR-R* with an exceptionally low parameter count.

## 1. Introduction

The primary difference between aerial object detection [3, 7, 65] and general object detection [18, 39, 66] lies in the bird’s-eye view perspective of aerial images, which causes objects to appear in arbitrary orientations. If horizontal bounding boxes (commonly used in general object detection) are applied to aerial images, they may encom-

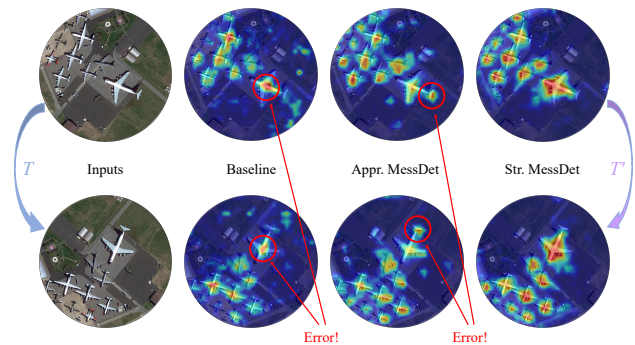


Figure 1. Visualization of features extracted by different methods. When the input is rotated, features extracted by conventional CNNs are easily affected, leading to degraded detection performance. In contrast, our MessDet framework includes both approximately rotation-equivariant (Appr. MessDet) and strictly rotation-equivariant (Str. MessDet) variants. While Appr. MessDet reduces rotation error during training, Str. MessDet eliminates such errors entirely, resulting in more robust detection performance.

pass excessive background and make it difficult to localize the target objects accurately. In recent years, several datasets [27, 33, 42, 50] have been introduced that utilize oriented bounding boxes (OBB) for annotating aerial images, which can more tightly enclose the target objects. However, bounding boxes now include angle information, increasing the data dimensionality from 4 to 5 when this approach is applied, which makes detection more challenging. When objects appear in arbitrary orientations, the detector should produce consistent outputs for classification tasks, regardless of orientation, requiring rotation invariance. In contrast, the detector needs to correctly adjust the output angles, necessitating rotation equivariance for bounding box regression. Most aerial object detectors implicitly learn rotation equivariance by increasing model parameters [12, 37, 44], enhancing data augmentation [5, 6, 14, 21, 51, 56, 57], proposing new bounding box representation methods [19, 24, 28, 45, 54, 55] or introduc-

\*Corresponding author.

Code page: <https://github.com/NuSance/MessDet>

ing customized loss functions [4, 58, 59, 61].

Cohen *et al.* [8] explicitly implement group-equivariant neural networks using group theory, which is further advanced in subsequent works such as [9, 46]. These developments lead to the creation of rotation-equivariant networks (RE-Nets) that are particularly well-suited for aerial object detection tasks. Some previous studies have applied RE-Nets to aerial object detection. ReDet [20] is the first to adopt a RE-Net as a backbone and neck. FRED [26] combines the anchor-free point-set representation from RepPoints [62] with RE-Nets to achieve a fully rotation-equivariant detector. Both ReDet and FRED achieved state-of-the-art performance at the time and have demonstrated the effectiveness of RE-Nets in aerial object detection.

However, most object detectors use typical downsampling as a trade-off between computational efficiency and accuracy, performing 2x downsampling with stride-2 convolutional kernels during forward propagation. Additionally, the object detection pipeline often resizes image dimensions to multiples of two to facilitate mapping predicted bounding box coordinates back to the original image. Mohamed *et al.* [36] has shown that applying stride-2 convolutional kernels to feature maps with even spatial dimensions breaks strict rotation equivariance. Previous research on RE-Nets in aerial object detection has not explored how strict rotation equivariance impacts model performance. For instance, ReDet only achieves approximate rotation equivariance. FRED ensures strict rotation equivariance by padding even-dimensional feature maps with zeros to make them odd-dimensional before downsampling. However, this approach may introduce misalignment issues.

This study proposes a new downsampling process that avoids single-sided padding. Instead, an additional tuning layer is introduced before the stride-2 downsampling layer to convert even-dimensional feature maps to odd dimensions to ensure strict rotation equivariance (see Fig. 1). Additionally, Edixhoven *et al.* [16] have shown that approximately rotation-equivariant models may relax strict equivariance constraints during training in image classification tasks. In certain scenarios, both strict and approximate rotation equivariance have their advantages and disadvantages. However, the impact of strict versus approximate rotation equivariance on model performance in aerial detection tasks remains an open question. This study will control the strictness of rotation equivariance by determining whether to include tuning layers in the network, thus comparing the effects of strict and approximate rotation equivariance on model performance. And we track the rotation equivariance error during training of the approximately rotation-equivariant model on the aerial object detection dataset to assess the importance of rotation equivariance.

Among the state-of-the-art CNN-based object detectors [48, 53] that have emerged in recent years, the channel

attention mechanism proposed by SENet [25] is a commonly used component. However, applying channel attention mechanisms in RE-Nets can break rotation equivariance. In this study, we implement a rotation-equivariant channel attention mechanism and develop a more advanced backbone and neck network based on RTMDet [35]. Due to the inherent grouping properties of rotation-equivariant features, the proposed model inputs feature maps from different orientation dimensions into separate branches of the head network, which are then combined for output. Thanks to the higher degree of weight sharing in RE-Nets and the multi-branch head network, the model proposed in this paper has an exceptionally low number of parameters. To our knowledge, it is currently the model with the fewest parameters while maintaining state-of-the-art performance. The multi-branch head rotation-equivariant single-stage object detector (MessDet) we propose achieves advanced performance on DOTA-v1.0 [15], DOTA-v1.5 [15] and DIOR-R [27].

In summary, the contributions of this study are as follows:

- A novel downsampling process is proposed and integrated with RE-Nets, enabling control over whether the network exhibits strict or approximate rotation equivariance.
- A rotation-equivariant channel attention mechanism is proposed to enable the implementation of rotation-equivariant backbone and neck networks with more advanced architectures. A multi-branch head network is introduced to make more efficient use of rotation-equivariant features.
- The variation of rotation equivariance error during training is quantitatively analyzed in the aerial detection task, highlighting the importance of rotation equivariance for aerial detection performance.

## 2. Related Work

This section briefly reviews the most pertinent studies and methodologies that have shaped the field with the aim of providing a foundation for the contributions of this work.

### 2.1. Aerial Object Detection

Due to the potential for objects in aerial images to appear at arbitrary angles, the detector's output must include not only the bounding box position, width, height, and class as in general object detection but also the angle information of the bounding box. This distinction allows research on aerial image object detectors to be roughly divided into the following categories. First, some studies focus on improving network structures or adding specialized modules to enable the model to better learn features related to orientation and enhance rotation equivariance. R<sup>3</sup>Det [57] and DRN [37] improve feature accuracy by adding feature re-

finement modules. SCRDet [56] and SCRDet++ [60] introduce instance-level denoising modules to reduce interference from cluttered background regions. RoI Transformer [14] transforms the feature regions corresponding to horizontal anchor boxes into those corresponding to oriented boxes. Oriented R-CNN [51] generates OBB with lower computational cost. S<sup>2</sup>ANet [21] introduces AlignConv, which shifts convolutional kernels, and constructs a feature alignment network using AlignConv. ReDet [20] is the first to introduce RE-Nets into aerial object detection tasks and improved RoI Align [22] by interpolating along the orientation dimension to align rotation-equivariant features, resulting in RiRoI Align. ARC [38] proposes an adaptive rotation convolutional module that allows filters to rotate for objects at different orientations. FRED [26] re-applies RE-Nets and combines them with point-set representation and deformable convolutions. All of the above studies can be viewed as methods to make networks more robust to rotational changes. In addition, LSKNet [29], PKINet [1], and Strip R-CNN [63] tackle challenges such as background noise and extreme aspect ratios in aerial images by employing large kernel and strip convolutions. This approach significantly reduces the parameter count while greatly enhancing detection performance.

Other studies use alternative representations for oriented boxes. CenterMap OBB [45] and Oriented RepPoints [28] use point-set representations for oriented boxes. Gliding Vertex [54] represents oriented boxes based on horizontal boxes by regressing four length ratios. CFA [19] uses convex hulls to represent oriented boxes. GWD [58] transforms the box representation into a two-dimensional Gaussian distribution, using Wasserstein Distance as the loss function. KLD [59] is similar to GWD but uses Kullback-Leibler Divergence as the loss function, addressing the issue of scale invariance that GWD lacks. These representations eliminate the need to predict extra rotation angles, preventing training instability caused by drastic regression loss changes when large objects with small angular shifts alter their aspect ratios. Additionally, Cheng *et al.* [4] add an extra regularization term in the loss function to enforce the model’s output features approximately equivariant to changes in image angles.

Similar to ReDet and FRED, this study aims to explicitly endow the model with rotation equivariance through RE-Nets, rather than enhancing the model’s robustness to angular variations by designing additional feature alignment methods or bounding box representation techniques.

## 2.2. Rotation-equivariant Networks

Cohen *et al.* [8] first introduce group-equivariant neural networks and have achieved equivariance for the discrete rotation group. Harmonic Networks [49] achieves 360-degree rotation equivariance by replacing conventional CNN fil-

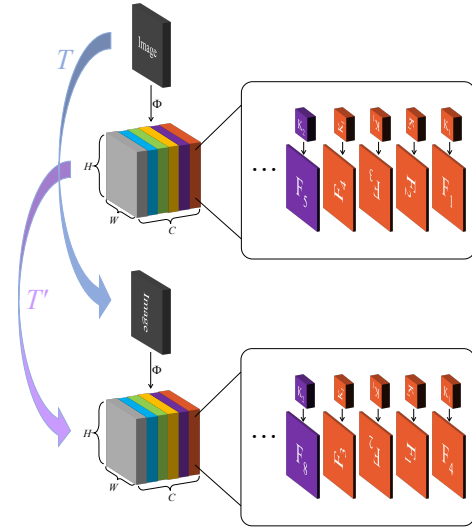


Figure 2. Rotation-equivariant feature. The upper half illustrates an input image processed by a RE-Net, which generates rotation-equivariant features. In the channel dimension, feature maps are shown in different colors, each representing a different kernel, such as  $K_1$  (Kernel 1). The RE-Net applies  $K_1$  at multiple orientations (four orientations are shown for illustration) and performs convolutions with the input feature maps to obtain different output feature maps, such as  $F_1$  and  $F_2$ . The lower half illustrates the resulting feature map when the input image is first rotated  $90^\circ$  clockwise before being processed by the RE-Net. The convolutional kernels and their orientations remain unchanged, but  $F_1$  shifts in the channel dimension and rotates  $90^\circ$  clockwise in the spatial dimension.

ters with circular harmonics. Subsequent works [9, 47] achieve better performance with convolutional kernels having fewer parameters. Spherical CNNs [10] accelerates the group-equivariant convolution process using generalized fast Fourier transform, while E2CNN [46] extends group-equivariant convolutional neural networks to the broader Euclidean group  $E(2)$  and experimentally demonstrate the performance advantages of  $E(2)$  convolutional kernels on multiple classification datasets. ReDet [20] is the first to apply RE-Nets to aerial object detection tasks, and FRED [26] implements a fully rotation-equivariant network, both achieving state-of-the-art performance at the time, highlighting the effectiveness of RE-Nets in aerial object detection.

## 3. Preliminaries

This section introduces concepts related to rotation equivariance in this study.

### 3.1. Rotation Equivariance

Let the signal be  $x$ , and after applying a certain transformation, it becomes  $Tx$ . If the output when  $Tx$  is input into the

network  $\Phi$  should be the same as the output when  $x$  is input directly into the network, after applying a certain transformation  $T'$ , that is,

$$\Phi(Tx) = T'\Phi(x), \quad (1)$$

the network  $\Phi$  is equivariant to the transformation  $T$  where  $T$  and  $T'$  do not need to be the same. In rotation-equivariant convolutional neural networks,  $T$  is defined as a transformation on the group  $G$ . If  $T' = 1$ , the network  $\Phi$  is said to be invariant to the transformation  $T$ , so invariance is a special case of equivariance. Ordinary convolutional neural networks exhibit equivariance to translation but not to rotation [8]. Rotation-equivariant convolutional neural networks (hereafter referred to as RE-Nets) achieve rotation equivariance by constructing cyclic groups, as shown in Fig. 2.

In Fig. 2, it can be observed that when the input image is rotated, the feature maps simultaneously shift within the cyclic group included in the channel dimension and undergo the corresponding rotation in the spatial dimension. So,  $\Phi(T_g x) = T'_g \Phi(x)$ , where  $g$  represents an element of the group  $G$ . When the input is a rotation-equivariant feature, the convolutional kernel's rotation involves both spatial rotations and translations within the cyclic group in the channel dimension.

### 3.2. Breaking Rotation Equivariance

In object detection tasks, it is common practice to repeatedly perform downsampling by a factor of 2, such as achieving  $32\times$  downsampling, to provide the network with a larger receptive field. Additionally, when mapping a certain region of the feature map back to the original image or converting coordinates on the feature map to original image coordinates, the image is often resized to a size that is a multiple of two to ensure feature alignment and precise boundary box localization. A typical downsampling method involves using a convolutional kernel with a stride of 2 on the feature map of even dimensions. However, Mohamed *et al.* [36] points out that doing this can result in different sampling points for the convolutional kernel before and after rotation, which leads to breaking the rotation equivariance. A specific example can be found in Sec. 7 of the supplementary material.

In classification tasks, Edixhoven *et al.* [16] experimentally demonstrated that for models that break rotation equivariance, the rotation equivariance error of shallow convolutional layers gradually decreases during the training process, resulting in approximate rotation equivariance. In some cases, the rotation equivariance error of deeper convolutional layers may even increase during training. Romero *et al.* [41] use  $2\times 2$  pooling for  $2\times$  downsampling to ensure rotation equivariance. However, this approach is not suitable for aerial object detection tasks, as  $2\times 2$  pooling

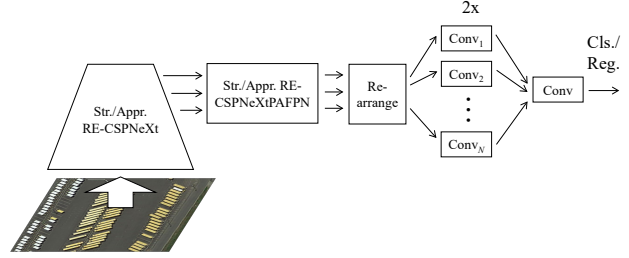


Figure 3. MessDet architecture. “Str.” and “Appr.” refer to strict rotation equivariance and approximate rotation equivariance, respectively. “ $N$ ” denotes the number of orientation dimensions.

causes significant loss of information for small objects. In detection tasks, FRED [26] uses padding to ensure that convolution with a stride of 2 always encounters odd-sized dimensions. However, single-side padding can also lead to misalignment of features, which may reduce the precision of the detector.

## 4. The Design of MessDet

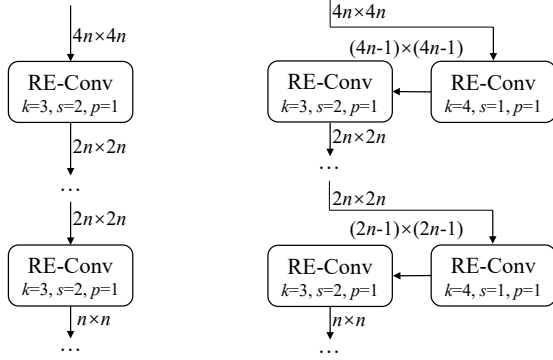
This study uses RTMDet [35] as the baseline and reimplements its backbone network CSPNeXt (an improved version of CSPNet [43]) and neck network CSPNeXtPAFPN (an improved version of PANet [32] and FPN [30]), incorporating the E2CNN [46]. Three improvements are made to handle rotation-equivariant features. These lead to the construction of a new aerial image object detector, MessDet. The network structure diagram is shown in Fig. 3. The following subsections will detail the design of MessDet.

### 4.1. Downsampling While Keeping Rotation Equivariance

If a RE-Net directly uses a convolutional layer with a stride of 2 for downsampling on feature maps with even spatial dimensions, it will lead to different sampling points. As a result, the rotation equivariance is broken and only achieves approximate rotation equivariance, as shown in Fig. 4a. This study proposes a novel downsampling process that maintains strict rotation equivariance, as shown in Fig. 4b. Assuming that the width and height of the input feature map are the same,  $S_{in} = H_{in} = W_{in}$ , the relationship between the spatial dimensions of the output feature map from the convolutional layer and the input feature map is given by:

$$S_{out} = \left\lfloor \frac{S_{in} + 2p - d \cdot (k - 1) - 1}{s} \right\rfloor + 1. \quad (2)$$

In the Formula 2,  $S_{out}$  and  $S_{in}$  represent the output size and input size, respectively,  $p$  is the padding size,  $d$  is the dilation (which is set to 1 in this study as dilation convolutions



(a) Common downsampling process. (b) Strictly rotation-equivariant downsampling process.

Figure 4. Comparison of different downsampling processes, where RE-Conv represents the rotation-equivariant convolutional module.

are not involved),  $k$  is the kernel size, and  $s$  is the stride of the convolutional kernel. To maintain strict rotation equivariance, feature maps with even spatial dimensions  $2n$  are first passed through a convolutional layer with  $k = 4$ ,  $p = 1$  and  $s = 1$ , referred to as the tuning layer. The output size of the tuning layer, calculated by Formula 2, is  $2n - 1$ . Then, it is passed through a downsampling convolutional layer with  $k = 3$ ,  $p = 1$ , and  $s = 2$ , and the output feature map size is given by:  $S_{out} = \lfloor ((2n - 1) - 1)/2 \rfloor + 1 = n$ . It can be seen that the inclusion of the tuning layer does not change the output size of the downsampling layer, but it ensures that the input feature map size for the 2x downsampling convolutional layer is odd, thereby guaranteeing strict rotation equivariance.

Users can easily alter whether the model exhibits strict rotation equivariance by adding a tuning layer before the two-times downsampling convolutional layer.

## 4.2. Rotation-Equivariant Channel Attention and Multi-branch Head

The channel attention mechanism introduced in SENet [25] is used in many state-of-the-art object detectors [35, 48, 53]. However, directly applying channel attention to rotation-equivariant features would break rotation equivariance. Therefore, this study proposes a rotation-equivariant channel attention (RE-CA) mechanism. Channel attention re-weights convolutional kernels, suppressing unimportant channels and enhancing important ones. The rotation-equivariant convolutional layer generates  $N$  convolutional kernels by rotating a single kernel in  $N$  orientations. Only  $C/N$  convolutional kernels are needed to generate a feature map with  $C$  channels. Thus, the RE-CA only needs to produce  $C/N$  weights, each of which is repeated  $N$  times, to weight the channels of the feature map. This ensures the

preservation of rotation equivariance.

For a rotation-equivariant feature  $\mathbf{X} \in \mathbb{R}^{C \times H \times W}$ , it can also be written as  $\mathbf{X} \in \mathbb{R}^{\frac{C}{N} \times N \times H \times W}$ . The process of RE-CA is as follows: the spatial information of the  $c$ -th channel of

$$\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_C]$$

is compressed into a scalar  $z_c$  using global average pooling,  $c \in [1, C]$ , resulting in the vector  $\mathbf{z} = [z_1, z_2, \dots, z_C]$ ,

$$z_c = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W \mathbf{x}_c(i, j). \quad (3)$$

Then,  $\mathbf{z}$  is passed through a fully connected layer  $\mathbf{W} \in \mathbb{R}^{\frac{C}{N} \times C}$  and an activation function  $\sigma$ ,

$$\mathbf{s} = \sigma(\mathbf{W} \cdot \mathbf{z}), \quad (4)$$

in the weight vector  $\mathbf{s} \in \mathbb{R}^{\frac{C}{N}}$ . Each element of the weight vector is repeated  $N$  times to obtain  $\mathbf{s}' \in \mathbb{R}^C$ , which is then multiplied with  $\mathbf{X}$  channel-wise,  $\tilde{\mathbf{x}}_c = \mathbf{s}' \cdot \mathbf{x}_c$ , to obtain the weighted rotation-equivariant feature  $\tilde{\mathbf{X}} = [\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \dots, \tilde{\mathbf{x}}_C]$ .

The rotation-equivariant features  $\tilde{\mathbf{X}} \in \mathbb{R}^{C \times H \times W}$  consist of feature maps generated by  $N$  convolutional kernels in different orientations, inherently exhibiting a grouping property. These features can be rearranged into  $N$  groups based on the orientation of the convolutional kernels that generate them,  $\mathbf{X} \in \mathbb{R}^{N \times \frac{C}{N} \times H \times W}$ , with each group having  $C/N$  channels. These groups are then fed into different branches of the head network, and the  $N$  feature maps are concatenated for the final output. Thanks to the multi-branch design, the parameter count of the head network is significantly reduced. Further details of rotation-equivariant channel attention and feature rearrangement can be found in Sec. 8 of the supplementary material.

## 5. Experiments

### 5.1. Datasets and Experiments Settings

We evaluate models on DOTA-v1.0 [15, 50], DOTA-v1.5 and DIOR-R [27]. DOTA-v1.0 contains 2,806 aerial images with OBB annotations. The resolutions of these images range from 800x800 to 4000x4000. It includes 188,282 instances across 15 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). The image data in DOTA-v1.5 is the same as that in DOTA-v1.0, but it includes instances smaller than 10 pixels and a new category, Container Crane (CC). DOTA-v1.5 contains 402,089 instances. DIOR-R [27] includes 23,463 aerial images and 192,518 annotations across 20 categories.

| Method                  | #P↓          | mAP↑         | PL           | BD           | BR           | GTF          | SV           | LV           | SH           | TC           | BC           | ST           | SBF          | RA           | HA           | SP           | HC           |
|-------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| <b>Two-stage</b>        |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |
| CenterMap[45]           | 41.1M        | 71.59        | 89.02        | 80.56        | 49.41        | 61.98        | 77.99        | 74.19        | 83.74        | 89.44        | 78.01        | 83.52        | 47.64        | 65.93        | 63.68        | 67.07        | 61.59        |
| SCRDet[56]              | 41.9M        | 72.61        | 89.98        | 80.65        | 52.09        | 68.36        | 68.36        | 60.32        | 72.41        | 90.85        | 87.94        | 86.86        | 65.02        | 66.68        | 66.25        | 68.24        | 65.21        |
| RoI Trans.[14]          | 55.1M        | 74.05        | 89.01        | 77.48        | 51.64        | 72.07        | 74.43        | 77.55        | 87.76        | 90.81        | 79.71        | 85.27        | 58.36        | 64.11        | 76.50        | 71.99        | 54.06        |
| FR-Est[17]              | -            | 74.20        | 89.63        | 81.17        | 50.44        | 70.19        | 73.52        | 77.98        | 86.44        | 90.82        | 84.13        | 83.56        | 60.64        | 66.59        | 70.59        | 66.72        | 60.55        |
| G. Vertex[54]           | 41.1M        | 75.02        | 89.64        | 85.00        | 52.26        | 77.34        | 73.01        | 73.14        | 86.82        | 90.74        | 79.02        | 86.81        | 59.55        | <b>70.91</b> | 72.94        | 70.86        | 57.32        |
| O-RCNN[51]              | 41.1M        | 75.87        | 89.46        | 82.12        | 54.78        | 70.86        | 78.93        | 83.00        | 88.20        | <b>90.90</b> | 87.50        | 84.68        | 63.97        | 67.69        | 74.94        | 68.84        | 52.28        |
| ReDet[20]               | 31.6M        | 76.25        | 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        |
| ARC[38]                 | 74.4M        | 77.35        | 89.40        | 82.48        | 55.33        | 73.88        | 79.37        | 84.05        | 88.06        | <b>90.90</b> | 86.44        | 84.83        | 63.63        | 70.32        | 74.29        | 71.91        | 65.43        |
| LSKNet[29]              | 31.0M        | 77.49        | 89.66        | <b>85.52</b> | <b>57.72</b> | 75.70        | 74.95        | 78.69        | 88.24        | 90.88        | 86.79        | 86.38        | 66.92        | 63.77        | 77.77        | 74.47        | 64.82        |
| PKINet[1]               | 30.8M        | 78.39        | <b>89.72</b> | 84.20        | 55.81        | <b>77.63</b> | 80.25        | 84.45        | 88.12        | 90.88        | 87.57        | 86.07        | 66.86        | 70.23        | 77.47        | 73.62        | 62.94        |
| <b>Single-stage</b>     |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |
| R <sup>3</sup> Det[57]  | 41.9M        | 69.70        | 89.00        | 75.60        | 46.64        | 67.09        | 76.18        | 73.40        | 79.02        | 90.88        | 78.62        | 84.88        | 59.00        | 61.16        | 63.65        | 62.39        | 37.94        |
| DRN[37]                 | -            | 70.70        | 88.91        | 80.22        | 43.52        | 63.35        | 73.48        | 70.69        | 84.94        | 90.14        | 83.85        | 84.11        | 50.12        | 58.41        | 67.62        | 68.60        | 52.50        |
| S <sup>2</sup> ANet[21] | 38.5M        | 74.12        | 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        |
| SASM[23]                | 36.6M        | 74.92        | 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        |
| FRED[26]                | -            | 75.56        | 89.37        | 82.12        | 50.84        | 73.89        | 77.58        | 77.38        | 87.51        | 90.82        | 86.30        | 84.25        | 62.54        | 65.10        | 72.65        | 69.55        | 63.41        |
| G-Rep[24]               | 36.6M        | 75.56        | 87.76        | 81.29        | 52.64        | 70.53        | 80.34        | 80.56        | 87.47        | 90.74        | 82.91        | 85.01        | 61.48        | 68.51        | 67.53        | 73.02        | 63.54        |
| O-RepPoints[28]         | 36.6M        | 75.97        | 87.02        | 83.17        | 54.13        | 71.16        | 80.18        | 78.40        | 87.28        | <b>90.90</b> | 85.97        | 86.25        | 59.90        | 70.49        | 73.53        | 72.27        | 58.97        |
| R3Det-GWD[58]           | 41.9M        | 76.34        | 88.82        | 82.94        | 55.63        | 72.75        | 78.52        | 83.10        | 87.46        | 90.21        | 86.36        | 85.44        | 64.70        | 61.41        | 73.46        | 76.94        | 57.38        |
| R3Det-KLD[59]           | 41.9M        | 77.36        | 88.90        | 84.17        | 55.80        | 69.35        | 78.72        | 84.08        | 87.00        | 89.75        | 84.32        | 85.73        | 64.74        | 61.80        | 76.62        | 78.49        | <b>70.89</b> |
| RTMDet[35]              | 52.3M        | 78.85        | 89.43        | 84.21        | 55.20        | 75.06        | 80.81        | 84.53        | 88.97        | <b>90.90</b> | 87.38        | 87.25        | 63.09        | 67.87        | 78.09        | 80.78        | 69.13        |
| Appr. MessDet           | <b>15.3M</b> | 78.45        | 89.28        | 84.69        | 56.20        | 66.18        | 81.37        | 85.33        | 88.86        | 90.78        | 88.28        | 87.62        | 64.47        | 65.59        | 78.18        | 82.10        | 67.82        |
| Str. MessDet            | 18.1M        | <b>79.12</b> | 88.08        | 85.15        | 55.81        | 72.09        | <b>81.54</b> | <b>85.85</b> | <b>88.98</b> | 90.84        | <b>88.36</b> | <b>88.70</b> | <b>68.78</b> | 66.85        | <b>78.84</b> | <b>82.66</b> | 64.26        |

Table 1. Comparison with state-of-the-art methods on the DOTA-v1.0 dataset [50] with single-scale training and testing. The MessDet backbones are pretrained on ImageNet-1K [13] for 300 epochs, similar to other methods [1, 29, 57].

| Method         | Strictly | #P↓   | mAP↑  |
|----------------|----------|-------|-------|
| MessDet        | ✓        | 19.0M | 78.51 |
| (Without head) | ✗        | 16.2M | 78.15 |
| RTMDet         | ✓        | 74.9M | 77.10 |
| (Without EMA)  | ✗        | 52.3M | 77.34 |

Table 2. Performance comparison of RTMDet and MessDet with two downsampling methods.

For a fair comparison, we follow the approach of major methods [21, 51], training models on both the training and validation sets and evaluating them on the test set. We adopt single-scale training and testing, using 1024×1024 patches for DOTA and 800×800 resized images for DIOR-R. Similar to mainstream methods [29, 57], we pretrain the strictly and approximately rotation-equivariant backbone networks used in the main results on ImageNet-1K [13] for 300 epochs, while the backbones in the ablation studies are pretrained for 100 epochs. Our model is implemented using the MMYOLO [11] and MMRotate [64] frameworks and trained for 36 epochs on DOTA-v1.0, DOTA-v1.5 and DIOR-R using the AdamW [34] optimizer. The experiments are conducted on 4 GPUs with a batch size of 8. Following the exploration of ReDet [20], this study set  $N$ , the number of orientation dimensions for the rotation-equivariant features in MessDet, to 8. More detailed experimental settings can be found in Sec. 9 of the supplementary

material.

## 5.2. Ablation Studies

This section presents ablation experiments to validate the effectiveness of the newly proposed downsampling process, rotation-equivariant channel attention (RE-CA) mechanism, and multi-branch head network. Unless otherwise stated, the reported results are based on the DOTA-v1.0 test set.

**Strictly Rotation-Equivariant Downsampling.** We have implemented rotation-equivariant backbone and neck networks with both approximately and strictly rotation-equivariant downsampling. To compare the impact of the new downsampling process on conventional convolutional networks, we re-implement RTMDet [35] with both approximately and strictly equivariant downsampling processes. For a fair comparison, MessDet uses the same head network as RTMDet, and RTMDet does not employ the Exponential Moving Average (EMA) strategy. Note that we only pretrain RTMDet’s two backbones for 100 epochs instead of the 600 epochs as in the original paper. The new downsampling process has little impact on conventional CNN but improves MessDet implemented with RE-Nets, as shown in Tab. 2.

**Rotation-Equivariant Channel Attention.** To demonstrate the effectiveness of the proposed RE-CA mechanism, we implement strictly and approximately rotation-equivariant backbones without the RE-CA, and the head

| Method              | #P↓          | mAP↑         | PL           | BD           | BR           | GTF          | SV           | LV           | SH           | TC           | BC           | ST           | SBF          | RA           | HA           | SP           | HC           | CC           |
|---------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| <b>Two-stage</b>    |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |
| FR-O[40]            | 41.1M        | 62.00        | 71.89        | 74.47        | 44.45        | 59.87        | 51.28        | 68.98        | 79.37        | 90.78        | 77.38        | 67.50        | 47.75        | 69.72        | 61.22        | 65.28        | 60.47        | 1.54         |
| Mask R-CNN[22]      | 44.4M        | 62.67        | 76.84        | 73.51        | 49.90        | 57.80        | 51.31        | 71.34        | 79.75        | 90.46        | 74.21        | 66.07        | 46.21        | 70.61        | 63.07        | 64.46        | 57.81        | 9.42         |
| HTC[2]              | 77.5M        | 63.40        | 77.80        | 73.67        | 51.40        | 63.99        | 51.54        | 73.31        | 80.31        | 90.48        | 75.12        | 67.34        | 48.51        | 70.63        | 64.84        | 64.48        | 55.87        | 5.15         |
| RoI Trans.[14]      | 55.1M        | 65.50        | 71.70        | 82.70        | 53.00        | 71.50        | 51.30        | 74.60        | 80.60        | 90.40        | 78.00        | 68.30        | 53.10        | 73.40        | 73.90        | 65.60        | 56.90        | 3.00         |
| ReDet[20]           | 31.6M        | 66.86        | 79.20        | 82.81        | 51.92        | 71.41        | 52.38        | 75.73        | 80.92        | 90.83        | 75.81        | 68.64        | 49.29        | 72.03        | 73.36        | 70.55        | 63.33        | 11.53        |
| LSKNet[29]          | 31.0M        | 70.26        | 72.05        | 84.94        | 55.41        | <b>74.93</b> | 52.42        | 77.45        | 81.17        | 90.85        | 79.44        | 69.00        | 62.10        | <b>73.72</b> | 77.49        | 75.29        | 55.81        | <b>42.19</b> |
| PKINet[1]           | 30.8M        | 71.47        | 80.31        | 85.00        | <b>55.61</b> | 74.38        | 52.41        | 76.85        | 88.38        | <b>90.87</b> | 79.04        | 68.78        | <b>67.47</b> | 72.45        | 76.24        | 74.53        | 64.07        | 37.13        |
| <b>Single-stage</b> |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |              |
| RetinaNet-O[31]     | 36.4M        | 59.16        | 71.43        | 77.64        | 42.12        | 64.65        | 44.53        | 56.79        | 73.31        | 90.84        | 76.02        | 59.96        | 46.95        | 69.24        | 59.65        | 64.52        | 48.06        | 0.83         |
| O-RepPoints[28]     | 36.6M        | 66.90        | 75.52        | 82.60        | 51.24        | 70.21        | 57.81        | 73.82        | 86.25        | 90.86        | 78.30        | 76.47        | 53.61        | 72.78        | 66.68        | 69.48        | 53.66        | 11.09        |
| FRED[26]            | -            | 68.30        | 79.60        | 81.44        | 52.60        | 72.57        | 58.07        | 74.82        | 86.12        | 90.81        | 82.13        | 74.84        | 53.37        | 72.93        | 69.51        | 69.91        | 54.82        | 19.27        |
| RTMDet[35]          | 52.3M        | 71.80        | 80.00        | 80.75        | 51.82        | 71.65        | 67.12        | 80.69        | 89.48        | 90.76        | 81.43        | 81.07        | 60.87        | 70.41        | 74.87        | 73.29        | <b>70.90</b> | 23.65        |
| Appr. MessDet       | <b>15.3M</b> | 72.38        | 80.61        | 84.99        | 54.56        | 65.36        | 73.60        | 82.69        | 89.73        | 90.83        | 82.46        | 82.03        | 59.35        | 72.96        | <b>77.66</b> | 75.14        | 65.09        | 21.12        |
| Str. MessDet        | 18.1M        | <b>73.14</b> | <b>80.83</b> | <b>86.17</b> | 55.58        | 63.28        | <b>74.24</b> | <b>82.73</b> | <b>89.79</b> | 90.85        | <b>83.47</b> | <b>82.45</b> | 63.97        | 71.65        | 77.44        | <b>76.06</b> | 66.17        | 25.54        |

Table 3. Comparison with state-of-the-art methods on the DOTA-v1.5 dataset [15] with single-scale training and testing. The MessDet backbones are pretrained on ImageNet-1K [13] for 300 epochs, similar to other methods [1, 20].

| Strictly | RE-CA | #P↓   | mAP↑         |
|----------|-------|-------|--------------|
| ✓        | ✓     | 19.0M | <b>78.51</b> |
|          | ✗     | 18.8M | 76.91        |
| ✗        | ✓     | 16.2M | <b>78.15</b> |
|          | ✗     | 16.0M | 77.47        |

Table 4. Performance comparison of MessDet with and without rotation-equivariant channel attention.

network without the multi-branch structure. The experimental results are shown in Tab. 4. The results indicate that the model with RE-CA performs significantly better than the model without it. Furthermore, due to the maintenance of rotation equivariance, the output dimension of the fully connected layer is reduced to  $1/N$  of the original. Hence, the parameter count required by the new channel attention mechanism is lower than that of conventional channel attention.

**Multi-branch Head.** We conduct a series of experiments to demonstrate that the multi-branch head network not only reduces the number of parameters but also improves model performance. To achieve the best possible performance, this study explores the impact of the number of convolutional modules in the multi-branch network on performance, using the Appr. MessDet as an example, as shown in Tab. 5. The number of convolutional modules includes the number of convolutional modules in each individual branch plus one convolutional module for the output aggregation. This shows that when the number of convolutional modules in the multi-branch head is set to 3, a balance between performance and parameter count is achieved.

### 5.3. Main Results

On DOTA-v1.0 [15], the performance comparison between MessDet and other state-of-the-art methods is re-

|                   | Num. | #P(head only)↓ | mAP↑         |
|-------------------|------|----------------|--------------|
| RTMDet Head       | 2    | 2.4M           | 78.15        |
|                   | 9    | 2.4M           | 78.19        |
|                   | 7    | 2.1M           | 77.97        |
| Multi-branch Head | 5    | 1.8M           | 77.62        |
|                   | 3    | 1.5M           | <b>78.45</b> |
|                   | 2    | 1.4M           | 77.46        |

Table 5. The impact of different head network designs on the performance of Appr. MessDet.

| Method                 | #P↓          | mAP↑         |
|------------------------|--------------|--------------|
| RetinaNet OBB[31]      | 36.4M        | 57.55        |
| Faster R-CNN OBB[40]   | 41.1M        | 59.54        |
| RoI Transformer[14]    | 55.1M        | 63.87        |
| LSKNet[29]             | 31.0M        | 65.90        |
| Oriented Reppoints[28] | 36.6M        | 66.71        |
| DCFL[52]               | -            | 66.80        |
| RTMDet[35]             | 52.3M        | 66.86        |
| PKINet[1]              | 30.8M        | 67.03        |
| Appr. MessDet          | <b>15.3M</b> | 67.42        |
| Str. MessDet           | 18.1M        | <b>68.19</b> |

Table 6. Comparison with state-of-the-art methods on the DIOR-R dataset [27]. The MessDet backbones are pretrained on ImageNet-1K [13] for 300 epochs, similar to other methods [1, 28, 52].

ported in Tab. 1. The approximately rotation-equivariant MessDet (Appr. MessDet) achieves 78.45% mAP, while the strictly rotation-equivariant MessDet (Str. MessDet) achieves 79.12% mAP. However, compared to other state-of-the-art models, MessDet has nearly half the parameter count of the lowest among them.

On DOTA-v1.5 [15], Tab. 3 presents the results, showing that Appr. MessDet attains 72.38% mAP, while Str. Mess-

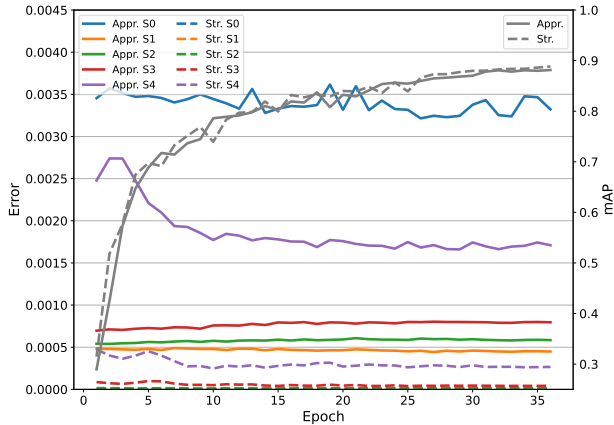


Figure 5. The variation of equivariance error across the five stages of the MessDet backbone network, where “Appr.” and “Str.” denote approximately and strictly rotation-equivariant MessDet respectively, and S0 denotes the stem, S1 denotes stage 1, and so on. mAP represents the model’s performance on the DOTA-v1.0 validation set.

Det achieves 73.14% mAP.

On DIOR-R [27], as detailed in Tab. 6, Appr. MessDet records an mAP of 67.03%, whereas Str. MessDet improves upon this with 68.19% mAP.

#### 5.4. Analysis

This section will analyze the behavior of rotation-equivariant networks in aerial object detection tasks, aiming to demonstrate the importance of rotation equivariance in such tasks.

**Equivariance Error Reduction in Detectors.** Edixhoven *et al.* [16] point out that in classification tasks, approximate rotation-equivariant models relax the rotation equivariance constraint, gradually increasing the rotation equivariance error during training to improve performance. Following this paper, we define the rotation equivariance error as follows:

$$\varepsilon = \frac{1}{ijk} \sqrt{\sum_{i,j,k} ([f(T(x))]_{ijk} - [T'(f(x))]_{ijk})^2}. \quad (5)$$

In Eq.(5),  $i$  and  $j$  are indices in the spatial dimensions, and  $k$  is the index in the channel dimension. To eliminate possible extraneous interference, we separately test the Appr. and Str. MessDet models without the multi-branch head network. The variation in rotation equivariance error across the five stages of the backbone network is shown in Fig. 5, based on the DOTA-v1.0 dataset.

As shown in Fig. 5, unlike in classification tasks, the rotation equivariance error in the approximately rotation-equivariant model decreases gradually during training on the DOTA-v1.0 aerial object detection dataset, proving that




Figure 6. The impact of various rotation angles on models performance.

aerial image object detection tasks indeed benefit more from rotation equivariance. Furthermore, the rotation equivariance error in the strictly RE-Net is naturally lower than in the approximately RE-Net, leading to improved detection performance.

**Robustness to Rotational Variations.** To demonstrate the robustness of RE-Nets to rotational variations, we test both Str. and Appr. MessDet, along with the baseline RTMDet [35], by applying rotations at various angles to the DOTA-v1.0 validation set. The effect of rotational changes on model accuracy is compared. Similarly, to eliminate any unnecessary interference, MessDet uses the same head network as RTMDet. The test results are shown in Fig. 6. As can be seen, even when RTMDet is rotated by multiples of 90° with no information loss, its performance decreases due to a lack of robustness to rotational variations, whereas MessDet exhibits greater robustness to such variations.

## 6. Conclusion

This paper proposes a novel downsampling process, implements a strictly rotation-equivariant network, and introduces a rotation-equivariant channel attention mechanism, enabling rotation-equivariant networks to adopt more advanced architectures. By combining a multi-branch head network, we present the MessDet. Experimental results show that we achieve performance comparable to state-of-the-art methods with an exceptionally low parameter count. Additionally, we quantitatively evaluate rotation equivariance error, demonstrating the necessity of rotation equivariance in aerial image object detection tasks.

**Acknowledgement.** This work was supported in part by the National Natural Science Foundation of China (Nos. 62302364, 62306324, 62376279, U24A20333), the Science and Technology Innovation Program of Hunan Province (No. 2024RC3128), the Research Foundation of National University of Defense Technology (No. ZK24-30), and the Key Laboratory of Target Cognition and Application Technology (No. 2023-CXPT-LC-005).

## References

- [1] Xinhao Cai, Qiuxia Lai, Yuwei Wang, Wenguan Wang, Zeren Sun, and Yazhou Yao. Poly kernel inception network for remote sensing detection. In *CVPR*, pages 27706–27716, 2024. 3, 6, 7
- [2] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *CVPR*, pages 4974–4983, 2019. 7
- [3] Gong Cheng and Junwei Han. A survey on object detection in optical remote sensing images. *ISPRS journal of photogrammetry and remote sensing*, 117:11–28, 2016. 1
- [4] Gong Cheng, Peicheng Zhou, and Junwei Han. Learning rotation-invariant convolutional neural networks for object detection in vhr optical remote sensing images. *IEEE TGRS*, 54(12):7405–7415, 2016. 2, 3
- [5] Gong Cheng, Jiabao Wang, Ke Li, Xingxing Xie, Chunbo Lang, Yanqing Yao, and Junwei Han. Anchor-free oriented proposal generator for object detection. *IEEE TGRS*, 60:1–11, 2022. 1
- [6] Gong Cheng, Yanqing Yao, Shengyang Li, Ke Li, Xingxing Xie, Jiabao Wang, Xiwen Yao, and Junwei Han. Dual-aligned oriented detector. *IEEE TGRS*, 60:1–11, 2022. 1
- [7] Gong Cheng, Xiang Yuan, Xiwen Yao, Kebin Yan, Qinghua Zeng, Xingxing Xie, and Junwei Han. Towards large-scale small object detection: Survey and benchmarks. *IEEE TPAMI*, 45(11):13467–13488, 2023. 1
- [8] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *ICML*, pages 2990–2999, 2016. 2, 3, 4
- [9] Taco Cohen and Max Welling. Steerable cnns. In *ICLR*, 2017. 2, 3
- [10] Taco Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. In *ICLR*, 2018. 3
- [11] MMYOLO Contributors. MMYOLO: OpenMMLab YOLO series toolbox and benchmark. <https://github.com/open-mmlab/mmyolo>, 2022. 6
- [12] Linhui Dai, Hong Liu, Hao Tang, Zhiwei Wu, and Pinhao Song. Ao2-detr: Arbitrary-oriented object detection transformer. *IEEE TCSVT*, 33(5):2342–2356, 2023. 1
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 6, 7
- [14] 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, 2019. 1, 3, 6, 7
- [15] Jian Ding, Nan Xue, Gui-Song Xia, Xiang Bai, Wen Yang, Michael Ying Yang, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Object detection in aerial images: A large-scale benchmark and challenges. *IEEE TPAMI*, 44(11):7778–7796, 2022. 2, 5, 7
- [16] Tom Edixhoven, Attila Lengyel, and Jan C. van Gemert. Using and abusing equivariance. In *ICCVW*, pages 119–128, 2023. 2, 4, 8
- [17] Kun Fu, Zhonghan Chang, Yue Zhang, and Xian Sun. Point-based estimator for arbitrary-oriented object detection in aerial images. *IEEE TGRS*, 59(5):4370–4387, 2021. 6
- [18] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, pages 580–587, 2014. 1
- [19] 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, 3
- [20] 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. 2, 3, 6, 7
- [21] Jiaming Han, Jian Ding, Jie Li, and Gui-Song Xia. Align deep features for oriented object detection. *IEEE TGRS*, 60:1–11, 2022. 1, 3, 6
- [22] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017. 3, 7
- [23] Liping Hou, Ke Lu, Jian Xue, and Yuqiu Li. Shape-adaptive selection and measurement for oriented object detection. In *AAAI*, pages 923–932, 2022. 6
- [24] Liping Hou, Ke Lu, Xue Yang, Yuqiu Li, and Jian Xue. Grep: Gaussian representation for arbitrary-oriented object detection. *Remote Sensing*, 15(3):757, 2023. 1, 6
- [25] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, pages 7132–7141, 2018. 2, 5
- [26] Chanho Lee, Jinsu Son, Hyounguk Shon, Yunho Jeon, and Junmo Kim. Fred: Towards a full rotation-equivariance in aerial image object detection. In *AAAI*, pages 2883–2891, 2024. 2, 3, 4, 6, 7
- [27] Ke Li, Gang Wan, Gong Cheng, Liqui 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. 1, 2, 5, 7, 8
- [28] Wentong Li, Yijie Chen, Kaixuan Hu, and Jianke Zhu. Oriented reppoints for aerial object detection. In *CVPR*, pages 1829–1838, 2022. 1, 3, 6, 7
- [29] Yuxuan Li, Qibin Hou, Zhaohui Zheng, Ming-Ming Cheng, Jian Yang, and Xiang Li. Large selective kernel network for remote sensing object detection. In *ICCV*, pages 16794–16805, 2023. 3, 6, 7
- [30] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 4
- [31] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 7
- [32] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *CVPR*, pages 8759–8768, 2018. 4
- [33] Zikun Liu, Liu Yuan, Lubin Weng, and Yiping Yang. A high resolution optical satellite image dataset for ship recognition and some new baselines. In *ICPRAM*, pages 324–331, 2017. 1
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6

- [35] Chengqi Lyu, Wenwei Zhang, Haiyan Huang, Yue Zhou, Yudong Wang, Yanyi Liu, Shilong Zhang, and Kai Chen. Rtmddet: An empirical study of designing real-time object detectors. *arXiv preprint arXiv:2212.07784*, 2022. 2, 4, 5, 6, 7, 8
- [36] Mirgahney Mohamed, Gabriele Cesa, Taco Cohen, and Max Welling. A data and compute efficient design for limited-resources deep learning. *arXiv preprint arXiv:2004.09691*, 2020. 2, 4
- [37] 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, 2020. 1, 2, 6
- [38] 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 *ICCV*, pages 6589–6600, 2023. 3, 6
- [39] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016. 1
- [40] 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, 2017. 7
- [41] David Romero, Erik Bekkers, Jakub Tomczak, and Mark Hoogendoorn. Attentive group equivariant convolutional networks. In *ICML*, pages 8188–8199, 2020. 4
- [42] Xian Sun, Peijin Wang, Zhiyuan Yan, Feng Xu, Ruiping Wang, Wenhui Diao, Jin Chen, Jihao Li, Yingchao Feng, Tao Xu, et al. Fair1m: A benchmark dataset for fine-grained object recognition in high-resolution remote sensing imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 184:116–130, 2022. 1
- [43] Chien-Yao Wang, Hong-Yuan Mark Liao, Yueh-Hua Wu, Ping-Yang Chen, Jun-Wei Hsieh, and I-Hau Yeh. Cspnet: A new backbone that can enhance learning capability of cnn. In *CVPRW*, pages 390–391, 2020. 4
- [44] 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 TGRS*, 61:1–15, 2023. 1
- [45] 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*, 59(5):4307–4323, 2021. 1, 3, 6
- [46] Maurice Weiler and Gabriele Cesa. General e(2)-equivariant steerable cnns. In *NeurIPS*, page 14334–14345, 2019. 2, 3, 4
- [47] Maurice Weiler, Fred A. Hamprecht, and Martin Storath. Learning steerable filters for rotation equivariant cnns. In *CVPR*, pages 849–858, 2018. 3
- [48] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *ECCV*, pages 3–19, 2018. 2, 5
- [49] Daniel E. Worrall, Stephan J. Garbin, Daniyar Turmukhambetov, and Gabriel J. Brostow. Harmonic networks: Deep translation and rotation equivariance. In *CVPR*, pages 5028–5037, 2017. 3
- [50] 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, 2018. 1, 5, 6
- [51] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *ICCV*, pages 3520–3529, 2021. 1, 3, 6
- [52] Chang Xu, Jian Ding, Jinwang Wang, Wen Yang, Huai Yu, Lei Yu, and Gui-Song Xia. Dynamic coarse-to-fine learning for oriented tiny object detection. In *CVPR*, pages 7318–7328, 2023. 7
- [53] Shangliang Xu, Xinxin Wang, Wenyu Lv, Qinyao Chang, Cheng Cui, Kaipeng Deng, Guanzhong Wang, Qingqing Dang, Shengyu Wei, Yuning Du, et al. Pp-yoloe: An evolved version of yolo. *arXiv preprint arXiv:2203.16250*, 2022. 2, 5
- [54] 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*, 43(4):1452–1459, 2021. 1, 3, 6
- [55] Xue Yang and Junchi Yan. Arbitrary-oriented object detection with circular smooth label. In *ECCV*, pages 677–694, 2020. 1
- [56] 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, 2019. 1, 3, 6
- [57] 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, 2021. 1, 2, 6
- [58] 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. 2, 3, 6
- [59] Xue Yang, Xiaojiang Yang, Jirui Yang, Qi Ming, Wentao Wang, Qi Tian, and Junchi Yan. Learning high-precision bounding box for rotated object detection via kullback-leibler divergence. In *NeurIPS*, pages 18381–18394, 2021. 2, 3, 6
- [60] Xue Yang, Junchi Yan, Wenlong Liao, Xiaokang Yang, Jin Tang, and Tao He. Scrdet++: Detecting small, cluttered and rotated objects via instance-level feature denoising and rotation loss smoothing. *IEEE TPAMI*, 45(2):2384–2399, 2023. 3
- [61] Xue Yang, Yue Zhou, Gefan Zhang, Jirui Yang, Wentao Wang, Junchi Yan, Xiaopeng Zhang, and Qi Tian. The kfiou loss for rotated object detection. In *ICLR*, 2023. 2
- [62] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *ICCV*, pages 9657–9666, 2019. 2
- [63] Xinbin Yuan, ZhaoHui Zheng, Yuxuan Li, Xialei Liu, Li Liu, Xiang Li, Qibin Hou, and Ming-Ming Cheng. Strip r-cnn: Large strip convolution for remote sensing object detection. *arXiv preprint arXiv:2501.03775*, 2025. 3
- [64] Yue Zhou, Xue Yang, Gefan Zhang, Jiabao Wang, Yanyi Liu, Liping Hou, Xue Jiang, Xingzhao Liu, Junchi Yan, Chengqi Lyu, et al. Mmrotate: A rotated object detection benchmark using pytorch. In *ACM MM*, page 7331–7334, 2022. 6

- [65] Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui-Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):8–36, 2017. [1](#)
- [66] Zhengxia Zou, Keyan Chen, Zhenwei Shi, Yuhong Guo, and Jieping Ye. Object detection in 20 years: A survey. *Proc. IEEE*, 111(3):257–276, 2023. [1](#)