# Projecting Points to Axes: Oriented Object Detection via Point-Axis Representation

Zeyang Zhao<sup>1</sup>, Qilong Xue<sup>2</sup>, Yuhang He<sup>1</sup>, Yifan Bai<sup>2</sup>, Xing Wei<sup>2</sup>( $\boxtimes$ ), and Yihong Gong<sup>1</sup>

<sup>1</sup> Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University <sup>2</sup> School of Software Engineering, Xi'an Jiaotong University ⊠ Corresponding Author {zeyang, xql1119, yfbai}@stu.xjtu.edu.cn {heyuhang, weixing, ygong}@mail.xjtu.edu.cn https://PointAxis.github.io/



Fig. 1: Robust oriented object detection results with point-axis representation.

Abstract. This paper introduces the point-axis representation for oriented object detection, as depicted in aerial images in Figure 1, emphasizing its flexibility and geometrically intuitive nature with two key components: points and axes. 1) Points delineate the spatial extent and contours of objects, providing detailed shape descriptions. 2) Axes define the primary directionalities of objects, providing essential orientation cues crucial for precise detection. The point-axis representation decouples location and rotation, addressing the loss discontinuity issues commonly encountered in traditional bounding box-based approaches. For effective optimization without introducing additional annotations, we propose the max-projection loss to supervise point set learning and the cross-axis loss for robust axis representation learning. Further, leveraging this representation, we present the Oriented DETR model, seamlessly integrating the DETR framework for precise point-axis prediction and end-to-end detection. Experimental results demonstrate significant performance improvements in oriented object detection tasks.

**Keywords:** Oriented Object Detection, Aerial Object Detection, Point-Axis Representation, Detection Transformer



**Fig. 2:** Mainstream oriented object representations. (a) is the *rotated bounding box* representation, (b) is the *point set* representation, and (c) is the *point-axis* representation we propose.

## 1 Introduction

Oriented object detection [1-5] is an important task in computer vision and has raised remarkable research interests in the community. Compared to vanilla object detection [6-11] which only outputs the horizontal bounding box aligned with the image's axis, oriented object detection requires a more flexible and accurate approach to handle non-axis aligned objects.

Current methods for oriented object detection predominantly rely on rotated bounding boxes to represent objects, as illustrated in Figure 2 (a). Despite variations in their definitions, these methods all encounter a common challenge known as loss discontinuity, stemming from abrupt changes in rotation degree or heightwidth definitions. Angle-based methods [12–15] define a rotated bounding box as  $(x, y, w, h, \theta)$ , where x and y denote the center point location, w and h represent the width (short edge) and height (long edge), and  $\theta$  denotes the orientation angle between the long edge of the box and the horizontal axis. Although these methods are effective in most cases, they suffer from a critical issue: when the lengths of the long and short edges are similar, the angle  $\theta$  can switch between  $\theta$ and  $\theta \pm 90^{\circ}$ , leading to discontinuities in the loss function. This switching behavior makes it difficult for the model to learn stable and consistent representations, thereby affecting detection performance.

Quadrilateral methods [16–18], on the other hand, determine a rotated bounding box using an outer horizontal box and offsets from its four corners, denoted  $(x, y, w, h, l_1, l_2, l_3, l_4)$ . While these methods provide more flexibility in representing objects' shapes and orientations, they also encounter ambiguous definitions. Specifically, when an object approaches a horizontal orientation, it encounters ambiguous definitions that the offset could change abruptly, as is the case shown in the bottom of Figure 2 (a). Other variant definitions, such as those based on box boundary vectors [19], middle lines [20], or Gaussian distributions [21,22], also face various representation issues, for example, the square problem [23,24].

Recent research has drawn attention to object representations based on point set [25, 26] as a promising approach to address these challenges. In these approaches, each object is represented as a collection of points. These methods then determine the object's bounding boxes by computing the minimum area rectangle enclosing these points. While these point set-based approaches excel at capturing the detailed locations of targets, their effectiveness heavily relies on the quality and quantity of the predicted points. Moreover, they often disregard the primary directionalities of objects, which can hinder their ability to accurately detect targets with complex shapes. In such cases, the computed rectangles may fail to adequately enclose the targets, as illustrated in Figure 2 (b).

To overcome existing limitations, we present an innovative point-axis representation that offers a divide-and-conquer approach for describing oriented objects, as illustrated in Figure 2 (c). Unlike methods that rely on direct orientation regression, our technique leverages a more flexible and geometrically intuitive framework centered on point sets and axes.

This representation emphasizes two pivotal aspects.

- Points for Shape Descriptions: The use of points to delineate the spatial extent and contours of objects provides a detailed shape representation. This aspect is crucial, especially when dealing with irregularly shaped objects that traditional bounding box-based methods might struggle to accurately describe.
- Axes for Orientation Cues: Beyond shape information, knowing the primary directionalities of objects is essential for precise detection. Axes provide these orientation cues, which are particularly important when keypoints can appear in near-circular shapes.

Our point-axis representation effectively disentangles location and rotation, facilitating a more flexible and continuous modeling approach for both aspects. Additionally, the axis representation is axis-order invariant, implying that it neither prioritizes one axis over another nor defines an object's long side. This attribute addresses the boundary discontinuities encountered in rotated bounding box-based methods, bolstering our method's resilience when handling objects with near-square or near-circular shapes. Consequently, this leads to more robust optimization and consistent prediction. For effective optimization, we devise two specialized loss functions.

- Max-Projection Loss: This loss function is introduced to supervise point set learning and facilitate object convergence without explicit keypoint annotations. It encourages the model to learn the optimal set of points that best represent the shape and contours of the object. By focusing on the projection of points onto the object's boundary, this loss function helps the model converge toward a more accurate representation of the object's spatial extent.
- Cross-Axis Loss: Designed for robust axis representation learning, crossaxis loss transforms axis representation into label encoding. This is achieved

4 Z. Zhao et al.

by discretizing angles into bins and applying smoothing to enhance robustness. The resulting label encoding yields a cross-axis shape with four peaks, representing the primary directionalities of the object. We utilize crossentropy loss to supervise this learning process.

Building on this flexible representation, we propose the Oriented-DETR model. This model extends the end-to-end detection Transformer framework with point-axis prediction. Specifically, our model introduces conditioned point queries and a points detection decoder to predict points. It leverages group self-attention to facilitate information interaction between point queries, enabling the capture of relationships between points and iterative refinement through multiple layers. Experimental results demonstrate effectiveness in main datasets, showing significant performance improvements in oriented object detection tasks.

## 2 Related Work

## 2.1 Oriented Object Detection

Oriented object detection faces the challenge of accurately representing objects with arbitrary orientations, whose axes often do not align with the image axes. While numerous studies [12, 27–30] have emphasized feature extraction techniques for such objects, the representation of oriented objects remains a fundamental issue. One common approach is the use of rotated bounding boxes. This representation has led to numerous research efforts aimed at enhancing its robustness. Some methods [13, 15, 21, 22, 31] utilize  $(x, y, w, h, \theta)$  to define each rotated bounding boxes. They focus on addressing the boundary issues caused by the periodicity of angles, such as CSL [13], which resolves the angular periodicity problem by transforming angle regression into a classification problem. However, these methods overlook the fundamental flaws in definition boundaries, which also lead to discontinuous loss. Some methods use different descriptions. Gliding Vertex [17] and RSDet [16] propose to describe objects using the outer horizontal bounding box along with the offsets of the four vertices. However, They also encounter the issue of vertices regression order. Oriented Reppoints [26] and OS-KDet [32] introduce point representation for oriented objects. Although avoiding the issue of abrupt definition changes, these point-based methods struggle to accurately describe both the position and orientation of shape-irregular objects. STD [33] employs a strategy of decoupling spatial transforms sequentially while CrackDet [34] partitioning the 180-degree regression task into four sub-tasks. Compared to these methods, our Point-Axis representation introduces, for the first time a decoupled representation of position and orientation for oriented object detection.

### 2.2 Detection Transformers

The Transformer framework [35] has emerged as a promising approach for various computer vision tasks [36–40], especially for object detection [41–45]. Compared



**Fig. 3:** The overall framework of the point-axis representation. (a) provides a visual depiction of the representation method. (b) outlines the overall process of loss constraints. (c) introduces the key loss, max projection loss, proposed by us for supervising the point set learning.

to traditional CNN-based detectors [6,8,9,46,47], DETR offers a simpler design, end-to-end optimization, and improved accuracy and scalability [48–52]. Several studies [53–55] have attempted to adapt the DETR framework for oriented object detection. Typically, these methods model object queries as rotated bounding boxes and refine them using a transformer decoder. However, they neglect the fact that the iterative updating of queries by the previous DETR decoder [43,48] is based on the premise that horizontal boxes are aligned with the image axes. Furthermore, they suffer from the previously discussed representation issues of oriented objects. In this study, we seamlessly integrate the point-axis representation with the DETR framework to resolve the aforementioned issues.

# **3** Oriented DETR with Point-Axis Representation

In this section, we introduce the point-axis representation and elaborate on its implementation, along with the corresponding loss constraints. Using this representation as a foundation, we present the Oriented DETR model and detail its architectural framework, emphasizing the design of the point-axis prediction.

#### 3.1 Point-Axis Representation

We propose the point-axis representation, as illustrated in Figure 3 (a). In this representation, each oriented object *i* is defined by a combination of a set of points  $\mathcal{P}_i$  and an axis representation  $\mathcal{A}_i$ . The set of points, denoted as  $\mathcal{P}_i = \{p_i^j\}_{j=1,2,...,K}$ , comprises *K* points, with the *K*-th point designated as the object's center. For the axis representation, we discretize directions into bins and apply Gaussian smoothing to generate a four-peak label encoding that represents the primary directionalities of the object. Notably, our representation exhibits the following advantages:

- 6 Z. Zhao et al.
- The decoupling of orientation from the bounding box prevents abrupt changes at boundaries that occur due to the coupled definition in methods based on rotated bounding boxes. This separation ensures a smoother transition between orientations, enhancing the stability and reliability of the representation for oriented objects.
- Employing circular labels ensures the consistency of angles at both the beginning and end of the defined range, thereby avoiding abrupt losses caused by angle periodicity. For instance, labels for 0° and 360° are identical, minimizing loss between 0° and 359°. This approach maintains continuity and smoothness in representing angles, enhancing the robustness and effectiveness of the orientation encoding process.

The point-axis representation provides a detailed depiction of oriented objects by incorporating positional information via point sets and directional details through axes. In the following sections, we delve into the loss constraints associated with this representation, elucidating how they contribute to the overall effectiveness of our approach.

**Point-Axis Loss Functions.** Toward effective model optimization, we propose two specialized loss functions as depicted in Figure 3 (b). For each object *i*, we initially transform its ground truth bounding box, defined by four coordinates  $\{(x_i^j, y_i^j)\}_{j=1,2,3,4}$ , into a central point  $C_i = \{(x_i^c, y_i^c)\}$ , accompanied by radial vectors extending from the central point to the edges of the bounding boxes  $\mathcal{V}_i = \{v_i^j\}_{j=1,2,3,4}$ . From these vectors, we derive  $\mathcal{A}_i$ . Subsequently, our model generates a set of points  $\hat{\mathcal{P}}_i$  and predicted axis encoding  $\hat{\mathcal{A}}_i$ . Finally, we enforce constraints on these predictions using max-projection loss and cross-axis loss.

The overall Point-Axis loss function is defined as:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^{N} (\lambda_1 \mathcal{L}_{proj}(\hat{\mathcal{P}}_i, \mathcal{V}_i, C_i) + \lambda_2 \mathcal{L}_{ca}(\hat{\mathcal{A}}_i, \mathcal{A}_i)),$$
(1)

where N is the total number of instances,  $\lambda_1$  and  $\lambda_2$  are balancing coefficients,  $\mathcal{L}_{proj}$  denotes our proposed max-projection loss, and  $\mathcal{L}_{ca}$  represents the cross-axis loss. We will now detail each loss function.

**Max-Projection Loss.** The max-projection loss is concise and easy to implement, with the computation process illustrated in Figure 3 (c). For the *i*-th object with its predicted point set  $\hat{\mathcal{P}}_i$ , we first convert it into a vector representation relative to  $C_i$ :  $\hat{\mathcal{V}}_i = {\hat{v}_i^j}_{j=1,2,...,K}$ . We then project each vector onto the elements of  $\mathcal{V}_i$ , selecting the element with the maximum projection value for optimization. The optimization objective is intuitively depicted in Figure 3 (c) and mathematically expressed as:

minimize 
$$\sum_{j=1}^{4} \left| \max_{m=1,2,\dots,K-1} \left( \frac{(\hat{v}_{i}^{m} - v_{i}^{j}) \cdot v_{i}^{j}}{\|v_{i}^{j}\|} \right) \right| + \|\hat{v}_{i}^{K}\|.$$
 (2)



Fig. 4: The architecture of Oriented DETR.

It is worth noting that the max-projection loss does not impose auxiliary constraints on non-projected local maxima in the set, which reduces the ambiguity in optimization direction and enhances the flexibility of point set description.

**Cross-Axis Loss.** The cross-axis loss is implemented using the cross-entropy loss, the mathematical expression of the optimization objective is as follows:

minimize 
$$\frac{1}{N_{bins}} \sum_{j=1}^{N_{bins}} \left[ \mathcal{A}_i^j \log \hat{\mathcal{A}}_i^j + (1 - \mathcal{A}_i^j) \log(1 - \hat{\mathcal{A}}_i^j) \right].$$
(3)

where  $N_{bins}$  is the number of bins for discretizing directions, with a default setting of 360. During inference, we take the argmax of these bins to obtain the principal direction, and then expand to the other three directions at intervals of 90 degrees.

#### 3.2 Architecture

The model architecture of Oriented DETR is depicted in Figure 4. Following the typical DETR framework, given an input image  $x_{img} \in \mathbb{R}^{3 \times H \times W}$ , it initially extracts features using a backbone and enriches them with global context through a transformer encoder. Subsequently, based on the scores corresponding to the features, it selects the top N as object queries  $Q_o \in \mathbb{R}^{N \times dim}$ , where N is the predefined maximum instance numbers. Our key design involves initializing point queries using an object-to-point query conversion function and enabling interaction among point queries using a points detection decoder. Finally, a prediction head is employed to generate the output.

**Object-to-Point Query Conversion.** For the *i*-th object query  $Q_o^i$  with its reference point  $(x_{ref}^i, y_{ref}^i)$ , we convert it into K point queries  $\{Q_p^{(i,j)}\}_{j=1,2,...,K}$  that are conditioned on  $Q_o^i$ . Among these,  $Q_p^{(i,K)}$  is defined as the center point

8 Z. Zhao et al.

query corresponding to the definition of  $\mathcal{P}_i$ . Following the common practice in DETR frameworks, we decompose the query into positional embeddings and content embeddings. For the center point query  $Q_p^{(i,K)}$ , we predict the offsets  $(\Delta x^i, \Delta y^i)$  of the center point relative to the reference point  $(x_{ref}^i, y_{ref}^i)$  using a multi-layer perceptron (MLP) layer:

$$\{\Delta x^i, \Delta y^i\} = MLP(Q_o^i). \tag{4}$$

The positional embedding  $P_i^K$  for the center point query is then generated by concatenating the positional encodings of the updated center point coordinates:

$$P_i^K = Concat(PE(x_{ref}^i + \Delta x^i), PE(y_{ref}^i + \Delta y^i)), \tag{5}$$

where PE is a positional encoding function.

For the other point queries, we generate their positional embeddings based on the polar coordinates. Initially, we define K-1 equidistant angles  $\{\theta_j\}_{j=1,2,...,K-1}$  and establish a polar coordinate system with  $(x_{ref}^i, y_{ref}^i)$  as the origin. Subsequently, we utilize an MLP layer to predict the polar coordinates of the object boundary in these directions and generate  $P_i^j$  as follows:

$$\{\Delta r_j^i\}_{j=1,2,\dots,K-1} = MLP(Q_o^i),$$
  

$$P_i^j = Concat(PE(x_{ref}^i + \Delta r_j^i \cdot \cos \theta_j), PE(y_{ref}^i + \Delta r_j^i \cdot \sin \theta_j)).$$
(6)

Finally, we use  $Q_o^i$  as the content embedding for all point queries to construct the initial point queries conditioned on *i*-th instance.

**Points Detection Decoder.** In the DETR framework, the decoder plays a pivotal role in facilitating information exchange between different components. It typically consists of self-attention mechanisms to enable interaction among queries and cross-attention mechanisms to promote interaction between queries and image features. Oriented DETR extends this concept by introducing a points detection decoder specifically designed to handle point queries. The point detection decoder in Oriented DETR comprises two key modules: the point-to-point attention module and the object-to-object attention module, as illustrated in Figure 4 (b).

In the point-to-point attention module, we divide point queries into N groups according to their corresponding instances, and each group contains K point queries. Then we apply a parameter-shared self-attention layer within each group to capture the relationships between points. In object-to-object attention, we extract the center point query from each instance to form a new group  $G_c = \{Q_p^{(i,K)}\}_{i=1,2,...,N}$  and apply self-attention within this new group to capture relationships between objects. Afterward, all queries are passed through a deformable cross-attention module and a FFN module for further processing.

**Prediction Head.** In the prediction head, for each instance *i*, each point query  $Q_i^j$  is mapped to a 2D point coordinate  $\hat{p}_i^j$ , while its class  $\hat{c}_i$  and axis  $\hat{\mathcal{A}}_i$  are mapped from all its conditioned point queries  $\{Q_p^{(i,j)}\}_{j=1,2,...,K}$ .

## 4 Experiment

#### 4.1 Evaluation Datasets

**DOTA** [1] is a large oriented object detection dataset comprising 2,806 images and 188,282 instances distributed across 15 categories, with a substantial variation in orientations, shapes, and scales. The image sizes range from  $800 \times 800$  to  $4000 \times 4000$ . The split ratio for training, validation, and testing sets is 3:1:2. According to the official toolkit, DOTA has two evaluation protocols: 1) single-scale training and testing 2) multi-scale training and testing. For single-scale training and testing, the training and testing images are cropped into  $1024 \times 1024$  patches with a stride of 824. For multi-scale training and testing, the raw pictures are first resized at three scales (0.5, 1.0, and 1.5) and then cropped into  $1024 \times 1024$ patches with a stride of 524.

**DIOR-R** [3] is constructed based on DIOR [56], with the addition of annotations for rotated bounding boxes. It comprises 23,463 images with a size of  $800 \times 800$  pixels, totaling 190,288 oriented instances across 20 categories. The dataset is divided into training, validation, and testing sets with a ratio of 1:1:2.

**HRSC2016** [57] is a renowned dataset for ship detection, comprising 1,061 images collected from ports with sizes ranging from  $300 \times 300$  to  $1500 \times 900$ . The training, validation, and testing sets consist of 436, 181 and 444 images, respectively.

In addition to the above datasets for oriented object detection, we also extend our model to handle general object detection tasks. To this end, we evaluate our approach on the **COCO2017** [58] dataset, which is a widely used benchmark for general object detection. It allows us to demonstrate the versatility and generalizability of our proposed approach.

#### 4.2 Implementation Details

When comparing with other methods on DOTA, DIOR-R, and HSRC2016, we use training and validation set for training and testing set for evaluation. The extension experiments on COCO are conducted by training on the training set and evaluating on the validation set. For DOTA, DIOR, and COCO, we train the model for 36 epochs, with a learning rate reduction at epochs 27 and 33. For HSRC2016, we train the model for 50 epochs, with learning rate reductions at epochs 33 and 45. We use the AdamW optimizer with an initial learning rate of 0.0001 and weight decay of 0.0001. Training is performed on 4 RTX 4090 GPUs with a total batch size of 8. Inference is performed on a single RTX 4090 GPU. The complete hyperparameters can be found in the supplementary file.

## 4.3 Ablation Study and Analyses

The key components of Oriented DETR are the point-axis representation and the points detection decoder. In our ablation study, we analyze these crucial components with ResNet-50 as backbone, training on the training set and evaluate on the evaluating on the validation set. Default settings are marked in gray.



Fig. 6: Visualization of results under different numbers of points.

Axis Representation Learning. To understand the behavior of axis predictions, we visualize their distributions for various objects in the DOTA dataset. As shown in Figure 1, for most objects, the axis predictions accurately capture the directional information. However, we also identify some anomalous distributions. Upon closer inspection, these anomalies often correspond to objects that lack an overall clear directional definition, such as certain swimming pools, as illustrated in Figure 5. It's worth noting that even in the DOTA dataset, annotations for such objects exhibit ambiguity. Nevertheless, our model is capable of learning the directional information for these objects and providing predictions that cover all possible directions through distributions. This analysis highlights the robustness of our approach in handling objects with varying degrees of directional clarity and demonstrates its ability to provide meaningful predictions even in the presence of annotation ambiguities. Additional visualizations of predicted axis distributions can be found in the supplementary file.

Number of Points. We investigate the influence of K (the number of points representing each object) on model performance. Table 1 and Figure 6 respectively present quantitative metrics and visual results. Statistically, increasing the number of points from 5 to 13 yielded a modest improvement of 0.49% for mAP<sub>50</sub> and 0.50% for mAP<sub>75</sub>. Visually, augmenting the number of points facilitates more precise localization. Nonetheless, even with only four boundary descriptor points at K = 5, predictions can still be accurately translated into rotated bounding boxes, predominantly owing to the auxiliary axis predictions.

**Point Loss Constraints.** In our method, the max-projection loss only imposes constraints on the maximum projection value in each direction. We seek to explore whether incorporating additional constraints could enhance supervision. Initially, we experiment by calculating the loss for all points located outside the ground-truth box and incorporating this into the penalty term. Furthermore, we

#### Point-Axis Representation 1

 Table 2: Performance comparison

with different point constraints.

**Table 1:** Performance comparisonwith different numbers of points.

Values of K	Б	0	12	16	10	Loss	Max-Projection	Max-Projection variants				
values of A	5	9	13	10	15	1055	india i rojection	w.penalty	$_{top-2}$	top-3		
mAP <sub>50</sub>	74.86	74.93	75.35	75.53	75.10	mAP <sub>50</sub>	75.35	75.20	74.77	73.20		
mAP <sub>75</sub>	49.64	49.98	50.14	50.03	49.60	mAP <sub>75</sub>	50.14	50.02	49.36	47.88		

Table 3: Ablation study on the design of the Points Detection Decoder.

	Baseline		Decoder components					
Point Queries?		<ul> <li>✓</li> </ul>	$\checkmark$	$\checkmark$				
Group Self-Attention?			$\checkmark$	$\checkmark$				
Decouple Cross-Attention?				$\checkmark$				
mAP <sub>50</sub>	72.80	70.98	74.21	$75.35\;(\uparrow\;2.55)$				
$mAP_{75}$	45.25	44.06	48.30	50.14 († 4.89)				

also attempt to constrain the top-k values of the projection, rather than solely focusing on the maximum value. However, the experimental result in Table 2 indicates that introducing additional penalties and increasing loss constraints don't result in improved accuracy. Specifically, when considering the top-k variant, we observe a notable decrease in accuracy when k = 2. We hypothesize that imposing excessive constraints may diminish the flexibility of the point set representation, ultimately leading to ambiguous optimization directions.

Points Detection Decoder. We conduct ablation experiments on different components of the points detection decoder, with the results shown in Table 3. Starting from a two-stage deformable DETR [42] (object queries in the decoder are generated from encoder output features) and retaining the object queries design, we utilize each object query to predict a set of points and an orientation axis, which serves as our baseline. This yields a  $mAP_{50}$  of 72.80% and a mAP<sub>75</sub> of **45.25**%. Subsequently, we model queries as dynamic points and use a vanilla self-attention layer to handle interactions between point queries, resulting in decreases of 1.82% for mAP<sub>50</sub> and 1.19% for mAP<sub>75</sub>. We attribute this to the ambiguity in the interaction between point queries of different objects. Therefore, we adopt group self-attention, initially grouping point queries by instances and then performing self-attention operations within each group, which significantly improves detection accuracy. Further, we decouple the deformable cross-attention modules for center point queries and boundary point queries to facilitate specialized feature learning, further improving its performance. Ultimately, the design of the points detection decoder leads to increases in accuracy of 2.55% for mAP<sub>50</sub> and 4.89% for mAP<sub>75</sub>.

#### 4.4 Comparison with State-of-the-Art Methods

**Results on DOTA.** We firstly conduct a fair comparison with state-of-the-art methods using the single-scale training and testing protocol, as commonly em-

11



Fig. 7: Visualization of results from Oriented DETR in DOTA.

ployed in most papers. The results are presented in Table 4. Using ResNet50 [59] and Swin-T [60] as backbones, Oriented DETR achieved mAP<sub>50</sub> of **79.1**% and **79.8**% respectively, outperforming all previous methods. Using the same backbone, Oriented DETR outperforms the point-based representation method Oriented Reppoints by **2.2**%. Compared to end-to-end DETR-based methods EMO2-DETR, ARS-DETR, and AO2-DETR, it achieves performance gains of **7.5**%, **4.3**%, and **1.4**%, respectively. Table 5 shows a systematic comparison (multi-scale training and testing) with others. Using Swin-L as the backbone, we achieve an mAP<sub>50</sub> of **82.26**%, surpassing the previous leading method, LSKNet, by **0.41**%, and outperforming the best DETR-based method, AO2-DETR, by **3.04**%. The visualization results are shown in Figure 7.

**Results on DIOR.** Table 6 shows our results on the DIOR dataset which has a richer variety of targets in terms of quantity and categories. using Swin-T as the backbone, Oriented DETR achieves an mAP<sub>50</sub> of **74.26**%, surpassing the previous best method by **3.21**%. Using ResNet50 as the backbone, Oriented DETR achieves an mAP<sub>50</sub> of **66.80**%, surpassing the best method with the same backbone by **2.89**%.

Table 4: Comparison with state-of-the-art methods on DOTA under single scale train-
ing and testing protocol. Best in <b>bold</b> .

Methods	Backbone	$_{\rm PL}$	$^{\mathrm{BD}}$	$\mathbf{BR}$	$\operatorname{GTF}$	SV	LV	$\operatorname{SH}$	$\mathrm{TC}$	$_{\rm BC}$	ST	$\operatorname{SBF}$	$\mathbf{R}\mathbf{A}$	$_{\rm HA}$	$^{\rm SP}$	HC	$\mathrm{mAP}_{50}$
one-stage methods																	
R <sup>3</sup> Det [63]	R-152	89.5	81.2	50.5	66.1	70.9	78.7	78.2	90.8	85.3	84.2	61.8	63.8	68.2	69.8	67.2	73.7
SASM [64]	R-50	86.4	79.0	52.5	69.8	77.3	76.0	86.7	90.9	82.6	85.7	60.1	68.3	74.0	72.2	62.4	74.9
GGHL [65]	Darknet53	89.7	85.6	44.5	77.5	76.7	80.5	86.2	90.8	88.2	86.3	67.1	69.4	73.4	68.5	70.1	77.0
Orient-Rep [26]	Swin-T	89.1	82.3	56.7	75.0	80.7	83.7	87.7	90.8	87.1	85.9	63.6	68.6	75.9	73.5	63.8	77.6
DHRec [66]	R-101	88.8	82.1	55.7	72.0	77.8	79.1	88.0	90.9	84.0	85.8	58.0	66.6	74.5	70.6	57.3	75.4
DCFL [67]	ReR101	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	75.4
PSC [15]	Darknet53	89.9	86.0	54.9	62.0	81.9	85.5	88.4	90.7	86.9	88.8	63.9	69.2	76.8	82.8	63.2	78.1
	two-stage methods																
SCRDet [68]	R-101	90.0	80.6	52.1	68.4	68.4	60.3	72.4	90.9	87.9	86.9	65.0	66.7	66.3	68.2	65.2	72.6
RoI-Trans [12]	R-101	88.6	82.6	52.5	70.9	77.9	76.7	86.9	90.7	83.8	82.5	53.9	67.6	74.7	68.8	61.0	74.6
G. Vertex [17]	R-101	89.6	85.0	52.3	77.3	73.0	73.1	86.8	90.7	79.0	86.8	59.6	70.9	72.9	70.9	57.3	75.0
CSL [13]	R-152	90.3	85.5	54.6	75.3	70.4	73.5	77.6	90.8	86.2	86.7	69.6	68.0	73.8	71.1	68.9	76.2
O-RCNN [27]	R-101	88.9	83.5	55.3	76.9	74.3	82.1	87.5	90.9	85.6	85.3	65.5	66.8	74.4	70.2	57.3	76.3
ReDet [28]	ReR50	88.8	82.6	54.0	74.0	78.1	84.1	88.0	90.9	87.8	85.8	61.8	60.4	76.0	68.1	63.6	76.3
OSKDet [32]	R-101	90.1	87.1	54.2	75.6	72.6	76.9	87.6	90.8	79.1	86.9	59.9	71.3	75.2	71.7	66.7	76.4
RVSA [69]	ViTAE-B	89.4	84.3	59.4	73.2	80.0	85.4	88.1	90.9	88.5	86.5	58.9	72.2	77.3	79.6	71.2	79.0
ARC [29]	ARC-R101	89.4	83.6	57.5	75.9	78.8	83.6	88.1	90.9	85.9	85.4	64.0	68.7	75.6	72.0	65.7	77.7
COBB [70]	ResNet50	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	76.6
PKINet [71]	PKINet-S	89.7	84.2	55.8	77.6	80.3	84.5	88.1	90.9	87.6	86.1	66.9	70.2	77.5	73.6	62.9	78.4
			DI	$\mathbf{TR}$	based	end-	to-ei	nd m	etho	ls							
AO2-DETR [54]	R-50	89.3	85.0	56.7	74.9	78.9	82.7	87.4	90.5	84.7	85.4	62.0	70.0	74.7	72.4	71.6	77.7
EMO2-DETR [53]	Swin-T	89.0	79.6	48.7	60.2	77.3	76.4	84.5	90.8	84.8	85.7	48.9	67.6	66.3	71.5	53.5	72.3
ARS-DETR [55]	Swin-T	87.7	76.5	50.6	69.9	79.8	83.9	87.9	90.3	86.2	85.1	54.6	67.0	75.6	73.7	63.4	75.5
Our Oriented-DETR	R-50	89.2	86.4	57.7	75.3	81.1	84.7	89.1	90.9	86.1	87.0	59.5	70.3	79.3	81.5	68.8	79.1
Our Oriented-DETR	Swin-T	89.4	85.1	57.8	75.0	81.2	86.1	89.1	90.9	88.7	87.0	62.9	69.1	80.7	82.8	71.0	79.8

Table 5: System comparison with state-of-the-art methods on DOTA.

Methods	EMO2-DETR [53]	SASM [64]	AO2-DETR [54]	ReDet [28]	R <sup>3</sup> Det-KLD [22]	AOPG [72]	O-RCNN [27]
mAP <sub>50</sub>	78.46	79.17	79.22	80.10	80.63	80.66	80.87
Methods	KFIoU [14]	RVSA [69]	RTMDet-L [73]	ARC [29]	LSKNet-S [30]	Ours(Swin-T)	Ours(Swin-L)
mAP <sub>50</sub>	80.93	81.24	81.33	81.77	81.85	81.78	82.26

**Results on HRSC2016.** On the HRSC2016 dataset, Oriented DETR achieves competitive results, as shown in Table 7. Oriented DETR achieves the highest accuracy of **98.02**% under the PASCAL VOC 2007 [61] metrics and achieves the second-best accuracy of **90.56**% under the PASCAL VOC 2012 [62] metrics, with a slight disadvantage compared to the previous best method (**90.60**%). On the HRSC2016 dataset, objects tend to have large aspect ratios and simple shapes, making the boundary issues almost non-existent for methods based on rotated bounding boxes of long-edge definition. Despite not having a significant advantage in accuracy, we have demonstrated the generalization capability of Oriented DETR across different scenarios.

#### 4.5 Extension to General Object Detection Task

We extend Oriented DETR to general object detection which only needs to output horizontal detection boxes and use COCO for evaluation. Concretely, we consider the axes of objects to be fixed in the horizontal and vertical directions, and remove the prediction and supervision for axes. Table 8 lists the experimental results, showing that Oriented DETR still maintains a precision advantage over 14 Z. Zhao et al.

**Table 6:** Performance comparisons onthe DIOR-R dataset.

**Table 7:** Performance comparisons on theHSRC2016 dataset.

Methods	Backbone	${}_{\mathrm{mAP}_{50}}$
Gliding Vertex [17]	R50	60.06
TIOE-Det [74]	R50	61.98
ROI-Trans [12]	R50	63.87
AOPG [72]	R50	64.41
O-RCNN [27]	R50	64.30
ARS-DETR [55]	R50	66.12
AO2-DETR [54]	R50	66.41
GGHL [65]	R50	66.48
Oriented RepPoints [26]	R50	66.71
DCFL [67]	R50	66.80
DCFL [67]	ReR101	71.03
RVSA [69]	ViTAE-B	71.05
Our Oriented-DETR	R50	69.69
Our Oriented-DETR	Swin-T	74.26

Methods	Backbone	mAP(07)	mAP(12)
ROI-Trans [12]	R101	86.20	-
AO2-DETR [54]	R50	88.12	97.47
R <sup>3</sup> Det [63]	R101	89.26	96.01
GWD [75]	R101	89.85	97.37
OSKDet [32]	R101	89.98	-
PSC [15]	DarkNet53	90.06	-
AOPG [72]	R101	90.34	96.22
Oriented Reppoints [26]	R50	90.38	97.26
ARC [29]	ARC-R50	90.41	-
ReDet [28]	ReR50	90.46	97.63
Oriented-RCNN [27]	R101	90.50	97.60
RTMDet [73]	RTM-L	90.60	97.10
Our Oriented-DETR Our Oriented-DETR	R50 Swin-T	90.52 90.56	97.73 98.02

Table 8: Results on COCO val.

Methods	Backbone	Epoches	AP	$\mathrm{AP}_{50}$	$\mathrm{AP}_{75}$	$\mathrm{AP}_{\mathrm{S}}$	$\mathrm{AP}_{\mathrm{M}}$	$\mathrm{AP}_{\mathrm{L}}$			
baseline methods											
DETR [41]	R-50	500	43.3	63.1	45.9	22.5	47.3	61.1			
Deformable-DETR 42	R-50	50	43.8	62.6	47.7	26.4	47.1	58.0			
$Deformable-DETR^{\dagger}$ [42]	R-50	50	46.2	65.2	50.0	28.8	49.2	61.7			
DAB-DETR [43]	R-50	50	45.7	66.2	49.0	26.1	49.4	63.1			
DAB-Deformable-DETR <sup>†‡</sup> [43]	R-50	50	49.7	68.2	54.3	32.0	52.9	65.3			
$\mathbf{Our} \ \mathbf{Oriented} \textbf{-} \mathbf{DETR}$	R-50	36	50.5	68.1	55.1	33.2	53.5	65.0			
methods based on training strategies											
DN-Deformable-DETR [76]	R-50	50	48.6	67.4	52.7	31.0	52.0	63.7			
$\mathcal{H}$ -Deformable-DETR [49]	R-50	36	50.0	-	-	32.9	52.7	65.3			
DINO-4scale [48]	R-50	36	50.9	69.0	55.3	34.6	54.1	64.6			
Group-DINO-4scale [77]	R-50	36	51.3	-	-	34.7	54.5	65.3			

<sup>†</sup>: two-stage.

<sup>‡</sup>: better result from detrex [78].

other DETR baseline methods. Future work could explore combining Oriented DETR with more advanced training strategies, such as those proposed in [48, 49, 76, 77], to further enhance its performance and versatility.

# 5 Conclusions

This paper proposes the point-axis representation, which provides a stable representation for aerial-oriented object detection by decoupling the location and orientation information, accompanied by the point-axis loss functions to supervise both the point set and the axis without introducing additional annotations. Leveraging this representation, we present the Oriented DETR model which integrates the DETR framework for precise point-axis prediction and end-to-end detection, showing significant performance improvements in aerial-oriented object detection tasks. In the future, we hope this representation could be extended to detecting arbitrary polygonal bounding boxes. The code will be made publicly available to facilitate easy reproduction.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grant No. U21B2048 and No. 62302382, Shenzhen Key Technical Projects under Grant CJGJZD2022051714160501, the Fundamental Research Funds for the Central Universities No. xxj032023020, and sponsored by the CAAI-MindSpore Open Fund, developed on OpenI Community.

## References

- 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, 2018.
- Jian Ding, Nan Xue, Gui-Song Xia, Xiang Bai, Wen Yang, Michael Ying Yang, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, et al. Object detection in aerial images: A large-scale benchmark and challenges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7778–7796, 2021.
- 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.
- 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, 2022.
- Shengxi Gui, Shuang Song, Rongjun Qin, and Yang Tang. Remote sensing object detection in the deep learning era—a review. *Remote Sensing*, 16(2):327, 2024.
- 6. Ross Girshick. Fast r-cnn. In ICCV, 2015.
- Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In CVPR, 2016.
- Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional onestage object detection. In *ICCV*, 2019.
- Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *CVPR*, 2023.
- Zhengxia Zou, Keyan Chen, Zhenwei Shi, Yuhong Guo, and Jieping Ye. Object detection in 20 years: A survey. *Proceedings of the IEEE*, 2023.
- 11. Xiang Song, Yuhang He, Songlin Dong, and Yihong Gong. Non-exemplar domain incremental object detection via learning domain bias. In AAAI, 2024.
- Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In CVPR, 2019.
- 13. Xue Yang and Junchi Yan. Arbitrary-oriented object detection with circular smooth label. In *ECCV*, 2020.
- Xue Yang, Yue Zhou, Gefan Zhang, Jitui Yang, Wentao Wang, Junchi Yan, Xiaopeng Zhang, and Qi Tian. The kfiou loss for rotated object detection. In *ICLR*, 2023.
- 15. Yi Yu and Feipeng Da. Phase-shifting coder: Predicting accurate orientation in oriented object detection. In *CVPR*, 2023.
- 16. Wen Qian, Xue Yang, Silong Peng, Junchi Yan, and Yue Guo. Learning modulated loss for rotated object detection. In AAAI, 2021.

- 16 Z. Zhao et al.
- Yongchao Xu, Mingtao Fu, Qimeng Wang, Yukang Wang, Kai Chen, Gui-Song Xia, and Xiang Bai. Gliding vertex on the horizontal bounding box for multioriented object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(4):1452–1459, 2020.
- Qi Ming, Lingjuan Miao, Zhiqiang Zhou, Xue Yang, and Yunpeng Dong. Optimization for arbitrary-oriented object detection via representation invariance loss. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2021.
- Jingru Yi, Pengxiang Wu, Bo Liu, Qiaoying Huang, Hui Qu, and Dimitris Metaxas. Oriented object detection in aerial images with box boundary-aware vectors. In WACV, 2021.
- Haoran Wei, Yue Zhang, Zhonghan Chang, Hao Li, Hongqi Wang, and Xian Sun. Oriented objects as pairs of middle lines. *ISPRS Journal of Photogrammetry and Remote Sensing*, 169:268–279, 2020.
- Xue Yang, Junchi Yan, Qi Ming, Wentao Wang, Xiaopeng Zhang, and Qi Tian. Rethinking rotated object detection with gaussian wasserstein distance loss. In *ICML*, 2021.
- 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. *NeurIPS*, 2021.
- Yi Yu and Feipeng Da. On boundary discontinuity in angle regression based arbitrary oriented object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- Zikai Xiao, Guoye Yang, Xue Yang, Taijiang Mu, Junchi Yan, and Shimin Hu. Theoretically achieving continuous representation of oriented bounding boxes. In *CVPR*, 2024.
- 25. Yihong Chen, Zheng Zhang, Yue Cao, Liwei Wang, Stephen Lin, and Han Hu. Reppoints v2: Verification meets regression for object detection. *NeurIPS*, 2020.
- Wentong Li, Yijie Chen, Kaixuan Hu, and Jianke Zhu. Oriented reppoints for aerial object detection. In CVPR, 2022.
- 27. Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *ICCV*, 2021.
- 28. Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotationequivariant detector for aerial object detection. In *CVPR*, 2021.
- 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*, 2023.
- 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*, 2023.
- 31. Xue Yang, Liping Hou, Yue Zhou, Wentao Wang, and Junchi Yan. Dense label encoding for boundary discontinuity free rotation detection. In *CVPR*, 2021.
- 32. Dongchen Lu, Dongmei Li, Yali Li, and Shengjin Wang. Oskdet: Orientationsensitive keypoint localization for rotated object detection. In *CVPR*, 2022.
- 33. Hongtian Yu, Yunjie Tian, Qixiang Ye, and Yunfan Liu. Spatial transform decoupling for oriented object detection. In AAAI, 2024.
- Zhuangzhuang Chen, Jin Zhang, Zhuonan Lai, Guanming Zhu, Zun Liu, Jie Chen, and Jianqiang Li. The devil is in the crack orientation: A new perspective for crack detection. In *ICCV*, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017.

- 36. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020.
- Xing Wei, Yuanrui Kang, Jihao Yang, Yunfeng Qiu, Dahu Shi, Wenming Tan, and Yihong Gong. Scene-adaptive attention network for crowd counting. arXiv preprint arXiv:2112.15509, 2021.
- Dahu Shi, Xing Wei, Liangqi Li, Ye Ren, and Wenming Tan. End-to-end multiperson pose estimation with transformers. In CVPR, 2022.
- Xing Wei, Yifan Bai, Yongchao Zheng, Dahu Shi, and Yihong Gong. Autoregressive visual tracking. In CVPR, 2023.
- 40. Yifan Bai, Zeyang Zhao, Yihong Gong, and Xing Wei. Artrackv2: Prompting autoregressive tracker where to look and how to describe. In *CVPR*, 2024.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In ECCV, 2020.
- Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2021.
- 43. Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. In *ICLR*, 2022.
- 44. Gongjie Zhang, Zhipeng Luo, Yingchen Yu, Kaiwen Cui, and Shijian Lu. Accelerating detr convergence via semantic-aligned matching. In *CVPR*, 2022.
- Gongjie Zhang, Zhipeng Luo, Zichen Tian, Jingyi Zhang, Xiaoqin Zhang, and Shijian Lu. Towards efficient use of multi-scale features in transformer-based object detectors. In CVPR, 2023.
- Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In CVPR, 2018.
- 47. Xiyang Dai, Yinpeng Chen, Bin Xiao, Dongdong Chen, Mengchen Liu, Lu Yuan, and Lei Zhang. Dynamic head: Unifying object detection heads with attentions. In CVPR, 2021.
- Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. In *ICLR*, 2023.
- Zhuofan Zong, Guanglu Song, and Yu Liu. Detrs with collaborative hybrid assignments training. In *ICCV*, 2023.
- 50. Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv, 2023.
- Feng Li, Hao Zhang, Huaizhe Xu, Shilong Liu, Lei Zhang, Lionel M Ni, and Heung-Yeung Shum. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *CVPR*, 2023.
- Feng Li, Qing Jiang, Hao Zhang, Tianhe Ren, Shilong Liu, Xueyan Zou, Huaizhe Xu, Hongyang Li, Chunyuan Li, Jianwei Yang, et al. Visual in-context prompting. arXiv, 2023.
- 53. Zibo Hu, Kun Gao, Xiaodian Zhang, Junwei Wang, Hong Wang, Zhijia Yang, Chenrui Li, and Wei Li. Emo2-detr: Efficient-matching oriented object detection with transformers. *IEEE Transactions on Geoscience and Remote Sensing*, 2023.

- 18 Z. Zhao et al.
- 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, 2022.
- 55. Ying Zeng, Yushi Chen, Xue Yang, Qingyun Li, and Junchi Yan. Ars-detr: Aspect ratio-sensitive detection transformer for aerial oriented object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–15, 2024.
- 56. 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.
- 57. 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*, volume 2, pages 324– 331. SciTePress, 2017.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016.
- 60. Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021.
- Mark Everingham and John Winn. The pascal visual object classes challenge 2007 (voc2007) development kit. Int. J. Comput. Vis, 88(2):303–338, 2010.
- Mark Everingham and John Winn. The pascal visual object classes challenge 2012 (voc2012) development kit. Pattern Anal. Stat. Model. Comput. Learn., Tech. Rep, 2007(1-45):5, 2012.
- 63. Xue Yang, Junchi Yan, Ziming Feng, and Tao He. R3det: Refined single-stage detector with feature refinement for rotating object. In AAAI, 2021.
- 64. Liping Hou, Ke Lu, Jian Xue, and Yuqiu Li. Shape-adaptive selection and measurement for oriented object detection. In AAAI, 2022.
- 65. Zhanchao Huang, Wei Li, Xiang-Gen Xia, and Ran Tao. A general gaussian heatmap label assignment for arbitrary-oriented object detection. *IEEE Transactions on Image Processing*, 31:1895–1910, 2022.
- 66. Guangtao Nie and Hua Huang. Multi-oriented object detection in aerial images with double horizontal rectangles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4932–4944, 2023.
- 67. 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, 2023.
- 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*, 2019.
- 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.
- Zikai Xiao, Guoye Yang, Xue Yang, Taijiang Mu, Junchi Yan, and Shimin Hu. Theoretically achieving continuous representation of oriented bounding boxes. In *CVPR*, 2024.
- Xinhao Cai, Qiuxia Lai, Yuwei Wang, Wenguan Wang, Zeren Sun, and Yazhou Yao. Poly kernel inception network for remote sensing detection. In CVPR, 2024.

- 72. 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.
- Chengqi Lyu, Wenwei Zhang, Haian Huang, Yue Zhou, Yudong Wang, Yanyi Liu, Shilong Zhang, and Kai Chen. Rtmdet: An empirical study of designing real-time object detectors. arXiv, 2022.
- Qi Ming, Lingjuan Miao, Zhiqiang Zhou, Junjie Song, Yunpeng Dong, and Xue Yang. Task interleaving and orientation estimation for high-precision oriented object detection in aerial images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 196:241–255, 2023.
- 75. Xue Yang, Gefan Zhang, Xiaojiang Yang, Yue Zhou, Wentao Wang, Jin Tang, Tao He, and Junchi Yan. Detecting rotated objects as gaussian distributions and its 3-d generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4335–4354, 2022.
- 76. Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *CVPR*, 2022.
- 77. Ding Jia, Yuhui Yuan, Haodi He, Xiaopei Wu, Haojun Yu, Weihong Lin, Lei Sun, Chao Zhang, and Han Hu. Detrs with hybrid matching. In *CVPR*, 2023.
- Tianhe Ren, Shilong Liu, Feng Li, Hao Zhang, Ailing Zeng, Jie Yang, Xingyu Liao, Ding Jia, Hongyang Li, He Cao, et al. detrex: Benchmarking detection transformers. arXiv, 2023.