

# RSAR: Restricted State Angle Resolver and Rotated SAR Benchmark

Xin Zhang<sup>1</sup>, Xue Yang<sup>2</sup>, Yuxuan Li<sup>1</sup>, Jian Yang<sup>1</sup>, Ming-Ming Cheng<sup>1,3</sup>, Xiang Li<sup>1,3</sup> ✉

<sup>1</sup>VCIP, CS, Nankai University <sup>2</sup>Shanghai AI Laboratory <sup>3</sup>NKIARI, Shenzhen Futian

zhasion@mail.nankai.edu.cn, yangxue@pjlab.org.cn, yuxuan.li.17@ucl.ac.uk

{csjyang, cmm, xiang.li.implus}@nankai.edu.cn

## Abstract

Rotated object detection has made significant progress in the optical remote sensing. However, advancements in the Synthetic Aperture Radar (SAR) field are lagging behind, primarily due to the absence of a large-scale dataset. Annotating such a dataset is inefficient and costly. A promising solution is to employ a weakly supervised model (e.g., trained with available horizontal boxes only) to generate pseudo-rotated boxes for reference before manual calibration. Unfortunately, the existing weakly supervised models exhibit limited accuracy in predicting the object's angle. Previous works attempt to enhance angle prediction by using angle resolvers that decouple angles into cosine and sine encodings. In this work, we first reevaluate these resolvers from a unified perspective of dimension mapping and expose that they share the same shortcomings: these methods overlook the unit cycle constraint inherent in these encodings, easily leading to prediction biases. To address this issue, we propose the **Unit Cycle Resolver (UCR)**, which incorporates a unit circle constraint loss to improve angle prediction accuracy. Our approach can effectively improve the performance of existing state-of-the-art weakly supervised methods and even surpasses fully supervised models on existing optical benchmarks (i.e., DOTA-v1.0). With the aid of UCR, we further annotate and introduce **RSAR**, the largest multi-class rotated SAR object detection dataset to date. Extensive experiments on both RSAR and optical datasets demonstrate that our UCR enhances angle prediction accuracy. Our dataset and code can be found at: <https://github.com/zhasion/RSAR>.

## 1. Introduction

Rotated object detection [66], which provides more precise localization than horizontal object detection, is widely utilized in remote sensing [5, 17, 45, 46], 3D detection [27, 67], and scene text detection [21, 25, 30, 68]. Particularly

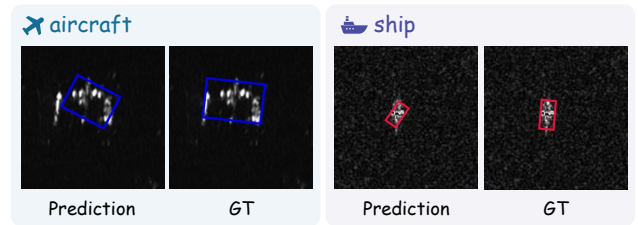


Figure 1. A comparison between prediction results from a weakly supervised model and the ground truth (GT). The weakly supervised model's accuracy in predicting the object angle requires further improvement.

in the field of remote sensing, the increased accessibility of optical satellite imagery [37, 45] facilitates numerous contributions to rotated object detection. Synthetic Aperture Radar (SAR) [36], a prominent remote sensing technology with all-weather imaging capabilities, attracts considerable attention in recent years. With the introduction of SAR datasets [39, 44, 64], a growing number of studies [4, 14, 20, 69] are focusing on SAR object detection.

Current research on SAR image interpretation primarily focuses on horizontal object detection, where performance tends to saturation [13, 14]. In contrast, progress in rotated SAR object detection is much slower. This is mainly due to the lack of a large-scale dataset specifically designed for rotated SAR object detection. Annotating such a dataset is both costly and inefficient, and a potential solution is to leverage a weakly supervised model. This weakly supervised model is trained with horizontal bounding boxes and generates pseudo-rotated detection boxes, which are then manually refined for efficient calibration. However, as shown in Fig. 1, accurately predicting the detected object's angle remains challenging for existing weakly supervised models. Therefore, improving angle prediction accuracy is crucial for enhancing the applicability of weakly supervised methods in SAR datasets.

The main challenge in accurate angle prediction lies in the angle boundary discontinuity problem [53]. To address this issue radically, previous related works [48, 58] utilize

✉ Corresponding Author.

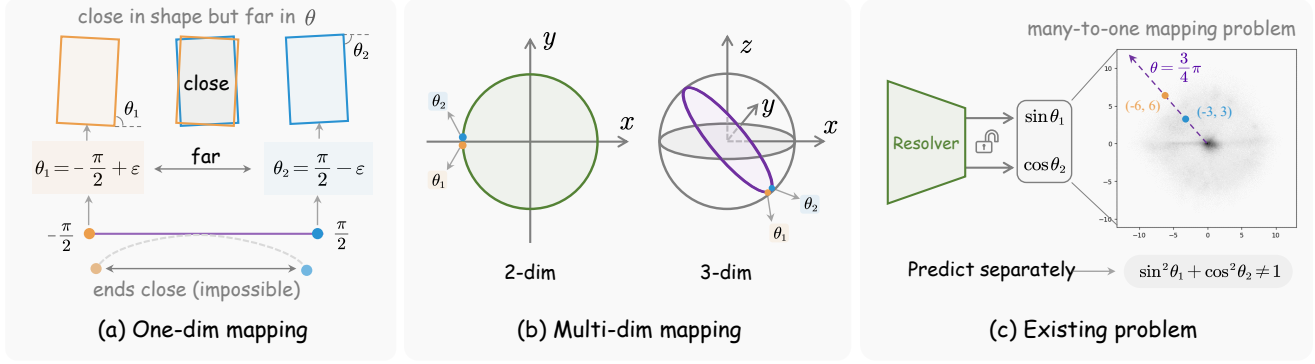


Figure 2. Diagram analysis of the angle boundary discontinuity problem from the unified perspective of *dimensional mapping*. (a) One-dimensional space leads to the problem of angle boundary discontinuity. (b) Mapping one-dimensional values to two-dimensional and three-dimensional spaces helps address the issue. (c) Existing methods overlook the unit circle constraint inherent in the angle encoding states, leading to a many-to-one problem that introduces biases in model optimization and predictions. Our unified perspective clarifies the angle boundary discontinuity issue and exposes the potential shortcomings of existing methods.

an angle resolver to decompose an angle value into sine and cosine components. Although their specific formulations vary, we reevaluate them from a unified and insightful perspective of *dimensional mapping*. All of these methods employ *dimensional mapping* to transform the discontinuous one-dimensional angle regression task into a continuous multi-dimensional encoding states regression task.

From this unified perspective, it is evident that previous approaches overlook the constraint that the angle encoding states must adhere to the unit circle conditions. Specifically, these methods [48, 58] predict multiple encoding states independently, which can lead to deviations from the unit circle constraint. This deviation complicates the optimization process and easily introduces biases in angle predictions. To address this, we propose an innovative restricted state angle resolver called Unit Cycle Resolver (UCR). It incorporates a unit cycle constraint loss to ensure the encoding states conform to the unit circle constraint. To justify its general effectiveness, we evaluate our method on the large-scale rotated object detection dataset DOTA-v1.0. Our method can effectively improve the performance of existing state-of-the-art weakly supervised methods and even achieves results comparable to that of fully supervised methods.

With the initial pseudo rotated box labels provided by the weakly supervised model [60] using our proposed UCR, we efficiently construct the RSAR dataset through manual calibration. RSAR is a comprehensive multi-class large-scale rotated SAR object detection dataset, which comprises 95,842 SAR images and 183,534 annotated instances across six typical SAR object categories. To our knowledge, RSAR is the largest rotated object detection dataset available in this field to date. Extensive experiments on RSAR and more optical datasets demonstrate that our UCR significantly enhances model performance in angle prediction.

Our contributions can be summarized as follows:

- We analyze existing angle resolver methods and their limitations from a unified perspective, proposing an innovative resolver called Unit Cycle Resolver (UCR). This resolver ensures that angle encodings adhere to inherent constraints by incorporating a unit circle constraint loss.
- We efficiently construct RSAR, a large-scale multi-class rotated SAR object detection dataset, with the help of the proposed UCR and existing weakly supervised model. To our best knowledge, RSAR is the largest rotated object detection dataset in the field of SAR to date.
- Experiments demonstrate that our method can significantly improve the performance of previous state-of-the-art approaches in weakly supervised tasks. It significantly improves the accuracy of angle prediction in both RSAR and optical datasets. Notably, the improved method achieves performance comparable to fully supervised methods in DOTA-v1.0.

## 2. Related Work

**SAR Object Detection:** Traditional SAR object detection methods primarily rely on manually designed features [35, 40], while deep neural networks are widely adopted in recent years [2, 3, 62]. The success of these networks largely depends on the availability of sufficient data. Early SAR object detection datasets mainly focus on ship objects. For instance, SSDD [64] introduces the first publicly available SAR ship detection dataset, and HRSID [44] is designed specifically for ship detection in high resolution. In recent years, larger datasets and more detection categories are proposed [19, 39, 43], including SARDet-100K [19], which provides a large-scale dataset for horizontal SAR object detection. Building on these datasets, various detection

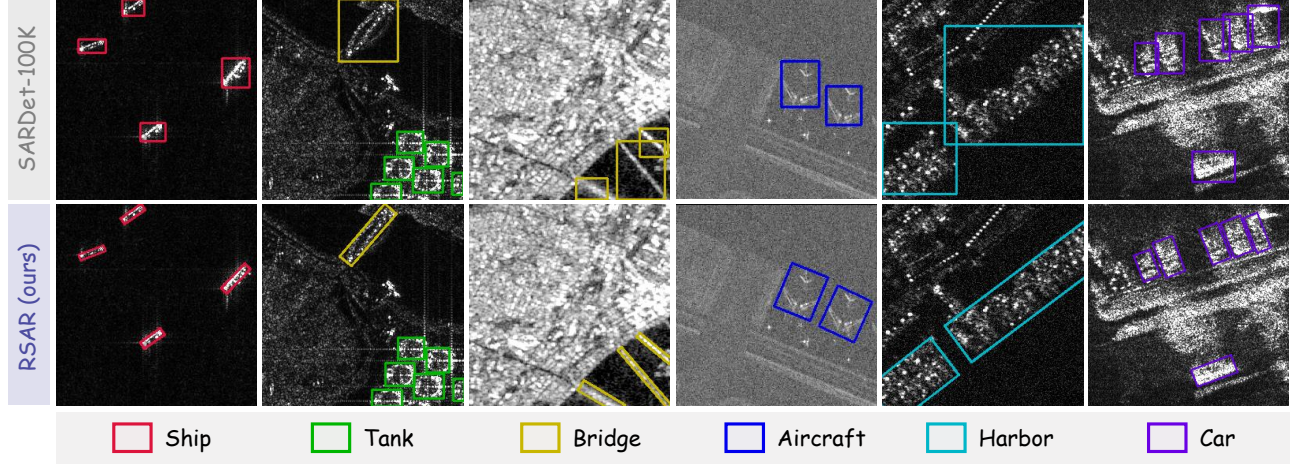


Figure 3. Visualization of images from RSAR and SARDet-100K. Rotated annotations in RSAR offer higher location accuracy compared to horizontal annotations in SARDet-100K. Rotated SAR object detection presents greater challenges than horizontal SAR object detection.

frameworks [3, 12, 14, 15, 63] are developed. However, current methods focus primarily on horizontal detection, where evaluation metrics reach saturation (*e.g.*, the evaluation metric of  $AP_{50}$  typically exceeding 90). In contrast, rotated object detection, which offers higher localization accuracy and presents more challenges, has progressed slowly. This delay is due to the limitations of single-category, small-scale datasets in this domain.

**Rotated Object Detection:** Existing models represent a detection bounding box using the format  $(cx, cy, w, h, \theta)$ . The key distinction between rotated object detection and horizontal object detection lies in the prediction of angles  $\theta$ . Currently, most mainstream approaches build upon horizontal detection frameworks [1, 10, 16, 32, 34] by incorporating angle regression predictions. These methods include single-stage frameworks like R<sup>3</sup>Det [52] and S<sup>2</sup>ANet [6], as well as two-stage frameworks such as ReDet [7] and SCRDet [50], along with DETR-based frameworks [57, 61, 71]. Weakly supervised tasks [29, 33, 55, 59, 60, 65] also play an important part in rotated object detection, particularly when rotated bounding box labels are unavailable. H2RBox [55] addresses this challenge by utilizing horizontal boxes as weak supervision signals for detecting rotated objects. H2RBox-v2 [60] further reduces the performance gap between weakly and fully supervised tasks by leveraging symmetry consistency information of the detected objects. However, the boundary discontinuity caused by angles  $\theta$  persists in most rotated object detection methods.

**Angle Boundary Discontinuity Problem:** Initially, using L1 loss to directly supervise the angle  $\theta$  for rotated object detection results in a sharp increase in angle loss at the boundaries of the defined angle range [53]. Subsequent research aims to address this issue. For instance, CSL [49] reformulates the angle regression task into a discrete an-

gle classification task, leading to various improved methods [41, 51]. However, this discrete approach introduces potential calculation errors. Later methods treat the detection bounding box as a Gaussian distribution, proposing GWD [53], KLD [54], and KFIOU [56] based on Gaussian joint optimization techniques. While these methods enhance model performance, they do not fully resolve the problem [48]. Recent approaches, such as PSC [58] and ACM [48], adopt angle encoding strategies to tackle this issue. Nonetheless, the relationship between these two methods remains unexplored, and the underlying constraint conditions are often overlooked.

### 3. Restricted State Angle Resolver

#### 3.1. Related Work of Angle Resolver

Inspired by the periodic ambiguity observed in absolute phase acquisition during optical measurements, PSC [58] introduces a phase-shifting coder technique to address the angle boundary discontinuity problem. The angle encoding states can be expressed as follows:

$$m_n = \cos\left(2\theta + \frac{2n\pi}{N_{\text{step}}}\right), \quad (1)$$

where  $n = 1, 2, \dots, N_{\text{step}}$  and  $N_{\text{step}}$  represents the total number of encodings. Here,  $\theta$  is the angle value in radians, and the range in the  $le_{90}$  notation is given by  $[-\pi/2, \pi/2)$ . The model predicts all encoding states  $M = \{m_1, \dots, m_n\}$  independently and decodes them accordingly [58].

Inspired by PSC, ACM [48] introduces a coding function based on the complex exponential function. The angle encoding states can be expressed as follows:

$$m = e^{j\omega\theta} = \cos(\omega\theta) + j \sin(\omega\theta), \quad (2)$$

where  $m$  denotes encoded value,  $j$  is the imaginary unit,  $\theta$  is the angle value, and  $\omega \in \mathbb{R}^+$  is the angular frequency.

### 3.2. Unified Perspective of Angle Resolver

The two methods utilize different forms of angle encoding, with ACM explicitly stating that these forms cannot be equivalent in any scenario. Nevertheless, both the issue of angle boundary discontinuity and the two methods can be viewed from a unified and insightful perspective of *dimensional mapping*. Table 1 shows the comparison between our unified perspective and the previous approach, which we will analyze in detail.

**One-dimensional mapping.** The core issue with angle boundary discontinuity arises from the challenge of making one-dimensional angle values equal at both ends of a defined range. To clarify, we assume the angle range is defined as  $\theta \in [-\pi/2, \pi/2)$  and  $\varepsilon$  represents a small positive value. As illustrated in Fig. 2, a rotated bounding box with an angle of  $\theta_1 = -\pi/2 + \varepsilon$  and another with an angle of  $\theta_2 = \pi/2 - \varepsilon$  nearly overlap in their shape. However, they exhibit a significant angle distance (*i.e.*,  $|\theta_1 - \theta_2|$ ) in one-dimensional linear space. This illustrates a limitation in one-dimensional space, where two angle values that are very far apart can represent rotated bounding boxes that are very close together.

**Two-dimensional mapping.** As illustrated in the left part of Fig. 2(b), the one-dimensional angle values can be mapped onto a circle in two-dimensional space, ensuring that the values at both ends of the angle range are equal. For simplicity, we assume the cycle is a unit circle with a radius of 1. Each point on the unit circle corresponds to a value within the defined one-dimensional angle range. This encoding mapping relationship can be represented by the following parametric equations:

$$m_1 = \cos(\theta), m_2 = \sin(\theta). \quad (3)$$

This is consistent with the form given in Eqn. (2) used in ACM. Mapping one-dimensional values to a circle in two-dimensional space already effectively resolves the issue of angle boundary discontinuity. However, this mapping relationship can be further extended into higher-dimensional space such as three-dimensional space.

**Three-dimensional mapping.** There are various ways to map one-dimensional value to a circle in three-dimensional space. For simplicity, we define the mapping range within the unit space (*i.e.*, each dimension has a value range of  $[-1, 1]$ ) and assume that the encoded values in each dimension follow the same distribution. The resulting constraint equations are as follows (see Appendix for detailed derivation):

$$\begin{cases} \sum_{i=1}^3 m_i^2 = \frac{3}{2} \\ \sum_{i=1}^3 m_i = 0 \end{cases}, \quad (4)$$

Unified Perspective ( <b>ours</b> )	Independent Perspective
one-dimension mapping	-
two-dimension mapping	ACM [48] (complex exponential)
three-dimension mapping	PSC [58] (absolute phase)

Table 1. A comparison between our unified perspective and the previous method’s independent perspective. Existing approaches (*e.g.*, ACM [48], PSC [58]) could be viewed as a case of our unified perspective.

where  $m_i$  represents the angle encoding in  $i_{th}$  dimension.

A valid solution to Eqn. (4) can be represented by Eqn. (1), as used in the PSC method. The right part of Fig. 2(b) provides a conceptual illustration of this case.

Although it is possible to map one-dimensional values onto a circular curve in higher-dimensional space, the constraints become progressively more complex as the number of dimensions increases. To simplify the calculations, we limit our focus to two-dimensional mapping and three-dimensional mapping.

### 3.3. Restricted State Angle Resolver

Through this analysis, we provide a clear explanation of the angle boundary discontinuity problem from the perspective of *dimensional mapping*, while also clarifying the mathematical principles behind existing methods. From this unified perspective, we reveal the limitation of previous approaches, where each angle encoding state is predicted independently (as shown in Fig. 2(c)). These predicted encoding states may not adhere to the inherent constraints of the unit circle. Without this constraint, a many-to-one mapping relationship exists between angle encoding states and angle values. Specifically, two angle encoding states that follow a linear scaling relationship could correspond to the same angle. This unrestricted solution space complicates the optimization of the model.

Building on the analysis above, we introduce a new angle resolver, termed the **Unit Circle Resolver (UCR)**. This approach facilitates *dimensional mapping* by satisfying the necessary constraints. We introduce the unit circle constraint loss to restrict the angle encoding state space of the network’s predictions. To simplify the calculations, we focus on two-dimensional mapping and three-dimensional mapping. We denote  $n$  as the dimension of mapping, and the form of this loss can be expressed as:

$$\mathcal{L}_{uc} = \left| \frac{n}{2} - \sum_{i=1}^n m_i^2 \right| + \sigma(n) \left| \sum_{i=1}^n m_i \right|, \quad (5)$$

where  $\sigma(n)$  is a piecewise function can be expressed as:

$$\sigma(n) = \begin{cases} 1, & n = 3 \\ 0, & n = 2 \end{cases}. \quad (6)$$

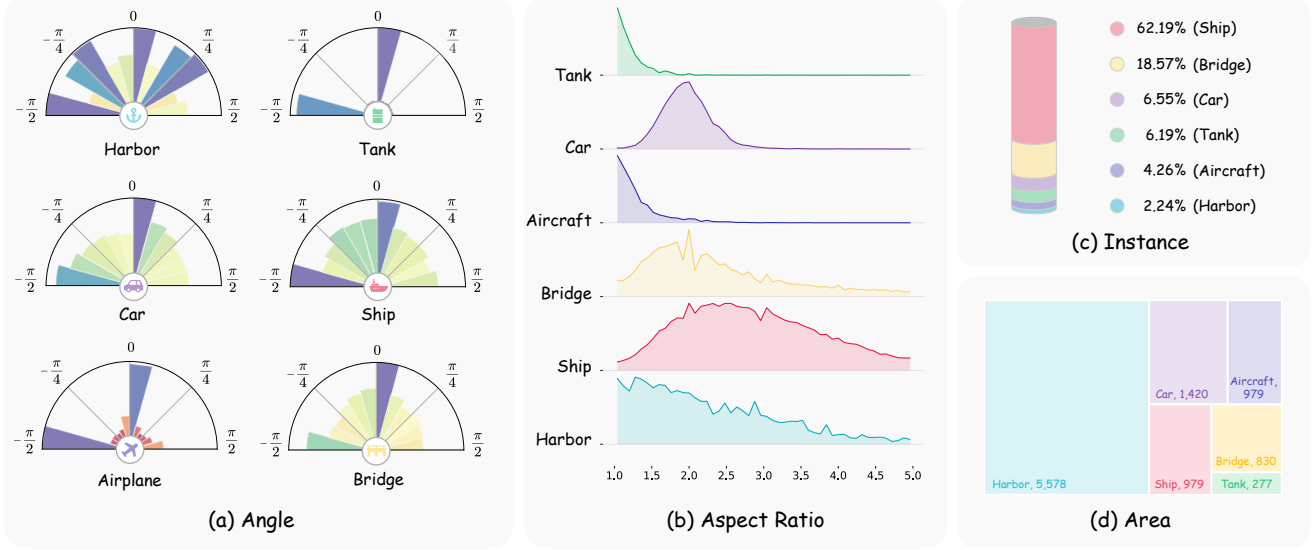


Figure 4. Statistical visualization of attributes for each category in RSAR. (a) The angle distribution of instances for each category (expressed in  $le_{90}$  angle notation). (b) The aspect ratio distribution of instances for each category. (c) The percentage of instances for each category. (d) The average instance pixel area for each category.

Since the model’s initial predictions for the angle encoding are relatively random, and the encoding values near the center of the unit circle exhibit significant fluctuations in the early stages, we introduce an invalid region to enhance training efficiency. Within this region, the angle encoding prediction is subject only to the constraint of Eqn. (5), without being influenced by the angle regression constraint. This region is defined as:

$$\sum_{i=1}^n m_i^2 < m_{invalid}. \quad (7)$$

For the rotated object detection tasks, the total loss can be written as:

$$\mathcal{L} = \mathcal{L}_{cls} + \lambda_{reg}\mathcal{L}_{reg} + \lambda_{uc}\mathcal{L}_{uc}, \quad (8)$$

where  $\mathcal{L}_{cls}$  is the loss of classification and  $\mathcal{L}_{reg}$  is the loss of bounding box positioning regression.  $\lambda_{reg}$  and  $\lambda_{uc}$  denote the loss weight coefficient of the corresponding loss.

Our UCR can be applied for critical angle prediction in the current weakly supervised model, H2RBox-v2 [60]. We leverage this model to generate rotated pseudo-labels, which assist in annotating a rotated SAR object dataset.

## 4. RSAR Dataset

SARDet-100K [19] consolidates 10 typical SAR datasets for standardization and unification. Building on this, we construct the rotated SAR object detection dataset, RSAR.

### 4.1. Dataset Annotation

**Data Cleaning:** Although SARDet-100K has processed the data, duplicate and unannotated images remain. Therefore, we first clean the original dataset before annotation. Specifically, we retain only one image with the maximum number of annotated objects in the case of repeated images and remove any images that lack annotations. This data-cleaning process helps prevent evaluation bias caused by dataset leaks.

**Data Annotation:** We use the horizontal bounding boxes from the dataset to train the weakly supervised model [60] using our UCR. This model is then employed to generate rotated bounding boxes as references for successive manual calibration. To ensure fairness in subsequent experiments, both the test and validation sets are fully manually annotated without reference from the weakly supervised model. It is worth noting that the orientation information for certain aircraft parts in the dataset was too ambiguous for accurate annotation, so we excluded this portion of the data.

**Unified Format:** We convert the annotations to DOTA format, which is easily compatible with mainstream detection frameworks.

### 4.2. Dataset Analysis

Our RSAR dataset comprises 95,842 images, 183,534 instances, and 6 typical object categories: **Ship, Tank, Bridge, Aircraft, Harbor, and Car**. The visualization of randomly sampled images from the dataset is shown in Fig. 3. It is evident that, compared to horizontal bounding box annotations, rotated bounding box annotations of-

Dataset	Categories	Images	Instances	RBB
HRSID [44]	1	5,604	16,969	
SAR-Ship [42]	1	39,729	50,885	
SAR-AIRcraft [43]	1*	4,368	16,463	
SIVED [24]	1	1,044	12,013	
SARDet-100K [19]	6	116,598	245,653	
RSDD-SAR [47]	1	7,000	10,263	✓
SSDD [64]	1	1,160	2,587	✓
SRSD [11]	1*	666	2,884	✓
OGSOD [39]	3	18,331	48,589	✓
DSSDD [9]	1	1,236	3,540	✓
<b>RSAR (ours)</b>	<b>6</b>	<b>95,842</b>	<b>183,534</b>	<b>✓</b>

Table 2. Comparison of existing SAR object detection datasets with RSAR. RBB indicates rotated bounding boxes are available, and \* denotes categories containing fine-grained subcategories. RSAR is the largest multi-class large-scale rotated SAR object detection dataset so far.

fer greater accuracy, while also introducing more challenges for the rotated SAR object detection task. To the best of our knowledge, RSAR is the first large-scale, multi-class rotated SAR object detection dataset.

To support the advancement of rotated SAR object detection, we analyze the annotation attributes for each category in the RSAR dataset. **(1) Angle:** We converted the angle values to the  $le_{90}$  notation ( $\theta \in [-\pi/2, \pi/2)$ ) and visualize the angle distribution for each category in Fig. 4(a). Notably, tanks are circular objects, so only horizontal annotations exist, consistent with the settings of previous datasets [45]. **(2) Aspect Ratio:** The aspect ratio of the objects also impacts detection accuracy. An aspect ratio close to 1 can introduce the square-like problem [53], while objects with an excessively high aspect ratio increase detection difficulty. Fig. 4(b) illustrates the distribution of the aspect ratio for each category. **(3) Instances:** Fig. 4(c) displays the proportion of instances for each category, with ships accounting for the largest share. It is evident that previous studies have predominantly focused on detecting ships. **(4) Area:** Fig. 4(d) presents the average pixel area size for each category.

### 4.3. Comparison with Existing Datasets

Table 2 presents the comparison between RSAR and other existing SAR object detection datasets. Previous datasets typically focus on a single object category, such as Ship (*e.g.*, HRSID, SAR-Ship) or Aircraft (*e.g.*, SAR-AIRcraft). However, single-category and single-scene detection data can introduce bias into specific models, undermining the evaluation of model generalization. In contrast, RSAR encompasses 6 categories of detection objects, making it the largest multi-category dataset among existing SAR object detection datasets. While the data cleaning process led to a slight reduction in the number of images and instances com-

pared to SARDet-100K, RSAR remains a COCO-level object dataset in terms of scale. In summary, RSAR is the first large-scale multi-class rotated SAR object detection dataset, setting it apart from existing datasets.

## 5. Experiments

**Optical Datasets: DOTA-v1.0 [45]** is a large optical aerial object detection dataset comprising 2,806 images, 15 object categories, and 188,282 instances. Following common practice [18, 60], we split the image into  $1,024 \times 1,024$  patches with an overlap of 200 pixels. All algorithms are trained on the training and validation sets, with evaluation conducted on the online test set. **HRSC [26]** is an optical remote sensing ship detection dataset consisting of 1,061 images containing 2,976 instances of ships. We utilize the pre-processing provided by MMRotate [70], where the images are scaled to  $800 \times 800$  for both training and testing.

**Evaluation Metric:** We use average precision (AP) as the evaluation metric widely adopted in this field. Following previous methods [61], we employ  $AP_{75}$  as the primary evaluation metric to more effectively evaluate the accuracy of angle prediction in rotated object detection. Additionally, we include mAP and  $AP_{50}$  as secondary evaluation metrics.

**Training Details:** To ensure a fair comparison, all methods are implemented using the MMRotate [70] framework, which is built on PyTorch [31]. All experiments are conducted on NVIDIA RTX 3090 GPUs, with 4 GPUs used for training on the RSAR dataset and 1 GPU for the other datasets. The DOTA-v1.0 and RSAR datasets are trained with a batch size of 2 for 12 epochs, while the HRSC dataset is trained with a batch size of 2 for 72 epochs. For weakly supervised tasks, we utilize horizontal bounding boxes as the supervision signal, while allowing the model to predict rotated bounding boxes. We adopt the previous state-of-the-art method, H2RBox-v2 [60], as the baseline model, which employs the FCOS [38] detector with a ResNet50 [8] backbone and an FPN [22] neck. All models are trained using AdamW optimizer [28], with an initial learning rate of  $1e-5$ .

### 5.1. Unit Cycle Resovler

Building on the previous HBox-supervised state-of-the-art model, H2RBox-v2, we evaluate the performance of our resolver against earlier methods across multiple datasets, including RSAR, DOTA-v1.0, and HRSC. The results, presented in Table 3, show significant improvements across all evaluation metrics compared to prior approaches. On the small-scale HRSC dataset, our method achieves the most pronounced improvement, with increases of 3.47% in  $AP_{75}$  and 2.79% in mAP metrics. On the large-scale DOTA-v1.0 optical dataset, our approach enhances the  $AP_{75}$  metric by 1.93% over previous methods. Remarkably, for the first time on the DOTA-v1.0 dataset, our method achieves weakly supervised performance that exceeds that of fully

Method	Resolver	DM	Supervised	RSAR			DOTA-v1.0 [45]			HRSC [26]		
				AP <sub>50</sub>	AP <sub>75</sub>	mAP	AP <sub>50</sub>	AP <sub>75</sub>	mAP	AP <sub>50</sub>	AP <sub>75</sub>	mAP
FCOS (R-50) [38]	-	-	RBB	66.66	31.45	34.22	71.44	41.76	41.81	89.26	77.47	63.21
H2RBox [55]	-	-	HBB	49.92	11.09	18.29	67.31	32.78	35.92	7.90	0.37	1.72
H2RBox-v2 [60]	ACM [48]	2	HBB	65.34	23.53	30.64	72.37	40.17	41.05	89.58	72.02	58.95
H2RBox-v2 [60]	PSC [58]	3	HBB	65.16	24.07	30.91	72.31	39.49	40.69	89.30	64.80	57.98
H2RBox-v2 [60]	UCR	2	HBB	<b>69.21</b>	<u>24.68</u>	<u>32.25</u>	<u>73.22</u>	<u>42.26</u>	<u>42.65</u>	<u>89.73</u>	<u>74.80</u>	<u>60.00</u>
H2RBox-v2 [60]	UCR	3	HBB	<u>68.33</u>	<b>26.17</b>	<b>32.64</b>	<b>73.99</b>	<b>42.10</b>	<b>43.10</b>	<b>89.74</b>	<b>75.49</b>	<b>61.74</b>

Table 3. The performance of our method and previous approaches on weakly supervised tasks across multiple datasets. DM refers to the dimension of mapping, RBB denotes the rotated bounding box, and HBB denotes the horizontal bounding box. Our method demonstrates significant improvements across all metrics compared to prior methods, which further validate the effectiveness of our unified analysis and the proposed UCR. The best score is in **bold** and the second-best is in underline.

Method	AP <sub>50</sub>	AP <sub>75</sub>	mAP
▼ <i>One-stage</i>			
RetinaNet [38]	57.67	22.72	27.65
R <sup>3</sup> Det [52]	63.94	25.02	30.50
S <sup>2</sup> ANet [6]	66.47	28.52	33.11
FCOS [23]	66.66	31.45	34.22
▼ <i>Two-stage</i>			
Faster RCNN [34]	63.18	24.88	30.46
O-RCNN [46]	64.82	32.69	33.62
ReDet [7]	64.71	32.84	34.30
RoI-Transformer [5]	66.95	32.65	35.02
▼ <i>DETR-based</i>			
Deformable DETR [71]	46.62	13.06	19.63
ARS-DETR [61]	61.14	28.97	31.56

Table 4. The fully supervised performance in various detectors on RSAR. All models use the default configuration in MMRotate.

$\lambda_{uc}$	RSAR			DOTA-v1.0 [45]		
	AP <sub>50</sub>	AP <sub>75</sub>	mAP	AP <sub>50</sub>	AP <sub>75</sub>	mAP
0	65.16	24.07	30.91	72.31	39.49	40.69
0.01	68.01	24.10	32.40	73.35	41.25	41.99
0.03	<b>68.45</b>	23.67	32.03	<b>73.99</b>	<b>42.10</b>	<b>43.10</b>
0.05	68.33	<b>26.17</b>	<b>32.64</b>	73.81	41.99	42.38
0.1	67.62	25.86	32.58	73.17	42.03	42.13
0.2	66.80	25.10	31.94	73.20	41.75	41.72

Table 5. Ablation experiments on different loss weights for unit cycle loss. The best performance on RSAR is achieved with a loss weight of 0.05, and the optimal loss weight for DOTA-v1.0 is 0.03.

supervised models using rotated bounding boxes. This demonstrates the effectiveness of our approach and establishes a strong baseline for rotated SAR object detection. On our RSAR dataset, we also observe an improvement of 2.10% in the AP<sub>75</sub> metric. Despite the challenges posed by complex environmental information associated with SAR images, our method effectively reduces the performance gap between weakly and fully supervised scenarios. These results further demonstrate that the weak supervision model in our method can effectively support dataset annotation.

unit cycle loss	invalid region	AP <sub>50</sub>	AP <sub>75</sub>	mAP
		72.37	40.17	41.05
	✓	72.68	41.57	42.09
✓		73.16	40.94	41.68
✓	✓	<b>73.22</b>	<b>42.26</b>	<b>42.65</b>

Table 6. Ablation experiments on various strategies in the two-dimensional mapping of UCR. Optimal results are achieved by combining the unit circle loss with the invalid areas.

## 5.2. Performance of RSAR in Various Detector

To better demonstrate the performance and characteristics of our RSAR dataset, we conduct experiments using mainstream one-stage, two-stage, and DETR-based rotated object detection frameworks. The experimental results are presented in Table 4. FCOS achieves the best performance among the one-stage frameworks, while RoI-Transformer outperforms others in the two-stage category. Overall, the two-stage frameworks deliver higher accuracy than other frameworks. Additionally, RSAR exhibits lower performance across various detectors when compared to the optical dataset DOTA-v1.0, highlighting the challenges of the rotated SAR object detection task.

## 5.3. Ablation Studies

**Loss weights:** Table 5 illustrates the impact of different loss weights for the unit cycle loss on model performance in weakly supervised tasks. The model achieves optimal performance on the RSAR dataset with a weight of 0.05, while the best results on the DOTA-v1.0 dataset are obtained with a weight of 0.03. Notably, compared to the baseline where our method is not applied ( $\lambda_{uc} = 0$ ), our approach significantly enhances both AP<sub>75</sub> and mAP metrics, even with a small loss weight. As the loss weight increases, the model’s performance begins to decline. This may be due to the fact that overly stringent constraints hinder the model’s ability to learn angles effectively. Overall, these results highlight the effectiveness and robustness of our method.

**Invalid region:** In our approach, we employ the unit circle loss to constrain the representation space of angle encod-

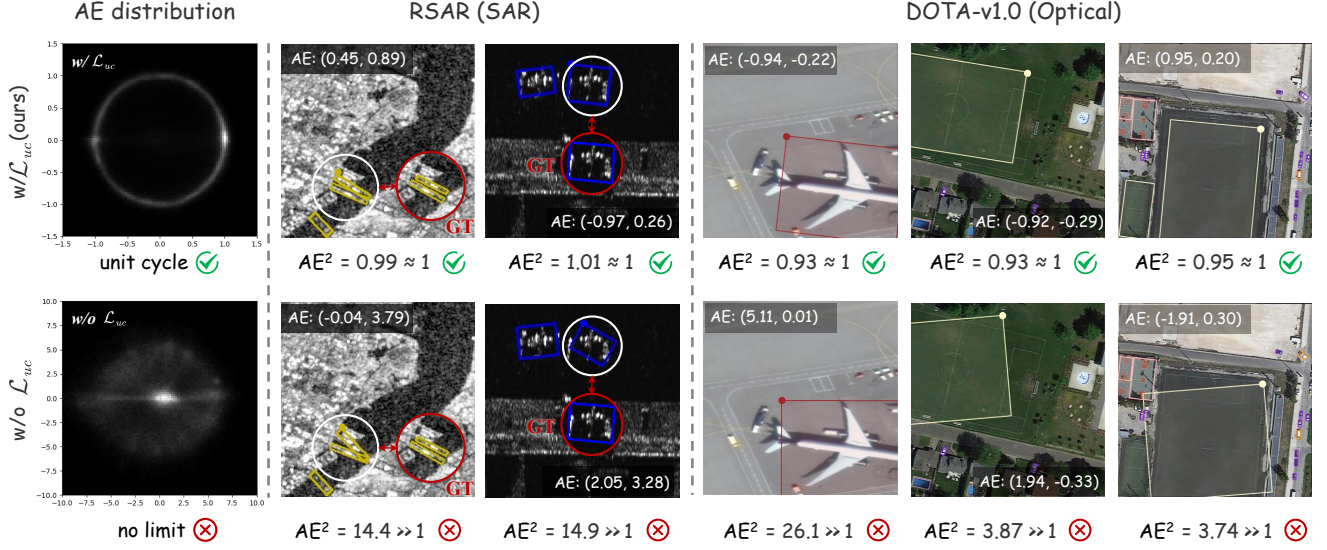


Figure 5. Comparison of visualized results on RSAR and DOTA-v1.0 in two-dimensional mapping. AE represents angle encoding, indicating the  $(\cos \theta, \sin \theta)$  of the model’s prediction, while  $AE^2$  denotes their sum of squares (*i.e.*,  $AE^2 = \sin^2 \theta + \cos^2 \theta$ ). We obtain the predicted values for all angle encodings of the bounding boxes on the test set and display their probability distribution statistics in the image on the left. The white regions in the probability distribution diagram correspond to areas where angle encoding has a higher probability of occurrence. Due to the unrestricted distribution in the angle encoding state space of the no-limit method, angle predictions may lack accuracy in certain scenarios. Our method significantly enhances the accuracy of angle prediction in the model.

ing and define the region near the center of the unit circle as an invalid region. Within this region, angle encoding is supervised solely by the unit circle loss and is exempt from angle regression supervision. Table 6 highlights the impact of these two strategies on model performance. The results show that simply defining the invalid region significantly improves performance, indicating that angle prediction in this region negatively affects the model and underscoring the importance of restricting the angle space. Overall, both strategies enhance model performance, with their combined use yielding even greater benefits.

**Visualization Result:** To better demonstrate how our method improves the model’s angle predictions, we compared the visualization results of our approach with the baseline method, as shown in Fig. 5. On the left, we present the distribution statistics of the model’s predicted angle encoding states on the test set, where white areas indicate regions with a high probability of occurrence for these encoding states. The baseline method, which lacks constraints, shows a randomly scattered distribution. In contrast, our method exhibits a well-defined unit circle distribution, indicating that our angle predictions are more regular. In the visualizations for both the RSAR and DOTA-v1.0 datasets, we observe that the baseline model tends to produce extreme predictions for the angle encoding states (*i.e.*,  $AE^2 \gg 1$ ), which negatively impacts the final prediction results. Our method provides more accurate predictions for

the angles of detection objects, as it adheres to the necessary constraints (*i.e.*,  $AE^2 \approx 1$ ).

## 6. Conclusion

In this paper, we address the issue of angle boundary discontinuity in rotated object detection from a unified perspective based on dimension mapping. Our analysis reveals that existing methods overlook the inherent unit cycle constraints in angle encoding, which leads to bias in angle predictions. To overcome this, we propose a unit circle resolver (UCR) that ensures angle encoding satisfies these constraint conditions, thereby improving the accuracy of angle prediction. We apply UCR in weakly supervised models to generate pseudo-rotated labels, and after manual calibration, we introduce RSAR—a large-scale, multi-class, rotated SAR object detection dataset. To our best knowledge, RSAR is the largest dataset in the field to date. Experimental results on both the RSAR and optical datasets demonstrate that our method significantly enhances angle prediction accuracy, even surpassing fully supervised models on DOTA-v1.0.

**Limitations and future work:** This paper primarily focuses on the improvement of weakly supervised models and the construction of a rotated SAR object detection dataset. However, designing a fully supervised detection model tailored for SAR detection using the RSAR dataset remains an important avenue for future research.

## Acknowledgement

This research was supported by the Young Scientists Fund of the National Natural Science Foundation of China (Grant No.62206134), the Fundamental Research Funds for the Central Universities 070-63233084, and the Tianjin Key Laboratory of Visual Computing and Intelligent Perception (VCIP). Computation is supported by the Supercomputing Center of Nankai University (NKSC). This work was supported by the National Science Fund of China under Grant No. 62361166670.

## References

- [1] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 2018. 3
- [2] Lifu Chen, Ru Luo, Jin Xing, Zhenhong Li, Zhihui Yuan, and Xingmin Cai. Geospatial transformer is what you need for aircraft detection in sar imagery. *TGRS*, 2022. 2
- [3] Zongyong Cui, Qi Li, Zongjie Cao, and Nengyuan Liu. Dense attention pyramid networks for multi-scale ship detection in sar images. *TGRS*, 2019. 2, 3
- [4] Yimian Dai, Minrui Zou, Yuxuan Li, Xiang Li, Kang Ni, and Jian Yang. Denodet: Attention as deformable multi-subspace feature denoising for target detection in sar images. *ArXiv*, 2024. 1
- [5] 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. 1, 7
- [6] Jiaming Han, Jian Ding, Jie Li, and Gui-Song Xia. Align deep features for oriented object detection. *TGRS*, 2021. 3, 7
- [7] Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotation-equivariant detector for aerial object detection. In *CVPR*, 2021. 3, 7
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [9] Yuxin Hu, Yini Li, and Zongxu Pan. A dual-polarimetric sar ship detection dataset and a memory-augmented autoencoder-based detection method. *Sensors*, 2021. 6
- [10] Ge-Peng Ji, Deng-Ping Fan, Yu-Cheng Chou, Dengxin Dai, Alexander Liniger, and Luc Van Gool. Deep gradient learning for efficient camouflaged object detection. *Machine Intelligence Research*, 2023. 3
- [11] Songlin Lei, Dongdong Lu, Xiaolan Qiu, and Chibiao Ding. Srsdd-v1. 0: A high-resolution sar rotation ship detection dataset. *Remote Sensing*, 2021. 6
- [12] Jianwei Li, Changwen Qu, and Jiaqi Shao. Ship detection in sar images based on an improved faster r-cnn. In *BIGSAR-DATA*, 2017. 3
- [13] Jianwei Li, Congan Xu, Hang Su, Long Gao, and Taoyang Wang. Deep learning for sar ship detection: Past, present and future. *Remote Sensing*, 2022. 1
- [14] Ke Li, Di Wang, Zhangyuan Hu, Wenxuan Zhu, Shaofeng Li, and Quan Wang. Unleashing channel potential: Space-frequency selection convolution for sar object detection. In *CVPR*, 2024. 1, 3
- [15] Weijie Li, Wei Yang, Tianpeng Liu, Yuenan Hou, Yuxuan Li, Zhen Liu, Yongxiang Liu, and Li Liu. Predicting gradient is better: Exploring self-supervised learning for sar atr with a joint-embedding predictive architecture. *ISPRS*, 2024. 3
- [16] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *NeurIPS*, 2020. 3
- [17] 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. 1
- [18] Yuxuan Li, Xiang Li, Yimian Dai, Qibin Hou, Li Liu, Yongxiang Liu, Ming-Ming Cheng, and Jian Yang. Lsknet: A foundation lightweight backbone for remote sensing. *IJCV*, 2024. 6
- [19] Yuxuan Li, Xiang Li, Weijie Li, Qibin Hou, Li Liu, Ming-Ming Cheng, and Jian Yang. Sardet-100k: Towards open-source benchmark and toolkit for large-scale sar object detection. *NeurIPS*, 2024. 2, 5, 6
- [20] Yuxuan Li, Xiang Li, Yunheng Li, Yicheng Zhang, Yimian Dai, Qibin Hou, Ming-Ming Cheng, and Jian Yang. Sm3det: A unified model for multi-modal remote sensing object detection. *ArXiv*, 2024. 1
- [21] Minghui Liao, Zhen Zhu, Baoguang Shi, Gui-song Xia, and Xiang Bai. Rotation-sensitive regression for oriented scene text detection. In *CVPR*, 2018. 1
- [22] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 6
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, and Kaiming He. Focal loss for dense object detection. In *ICCV*, 2017. 7
- [24] Xin Lin, Bo Zhang, Fan Wu, Chao Wang, Yali Yang, and Huiqin Chen. Sived: A sar image dataset for vehicle detection based on rotatable bounding box. *Remote Sensing*, 2023. 6
- [25] Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. Fots: Fast oriented text spotting with a unified network. In *CVPR*, 2018. 1
- [26] 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*, 2017. 6, 7
- [27] Zongdai Liu, Dingfu Zhou, Feixiang Lu, Jin Fang, and Liangjun Zhang. Autosshape: Real-time shape-aware monocular 3d object detection. In *ICCV*, 2021. 1
- [28] I Loshchilov. Decoupled weight decay regularization. *ArXiv*, 2017. 6
- [29] Junwei Luo, Xue Yang, Yi Yu, Qingyun Li, Junchi Yan, and Yansheng Li. Pointobb: Learning oriented object detection via single point supervision. In *CVPR*, 2024. 3
- [30] Jianqi Ma, Weiyuan Shao, Hao Ye, Li Wang, Hong Wang, Yingbin Zheng, and Xiangyang Xue. Arbitrary-oriented scene text detection via rotation proposals. *IEEE transactions on multimedia*, 2018. 1
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 2019. 6

- [32] J Redmon. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 3
- [33] Botao Ren, Xue Yang, Yi Yu, Junwei Luo, and Zhidong Deng. Pointobb-v2: Towards simpler, faster, and stronger single point supervised oriented object detection. *ArXiv*, 2024. 3
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *TPAMI*, 2016. 3, 7
- [35] Frank C Robey, Daniel R Fuhrmann, Edward J Kelly, and Ramon Nitzberg. A cfar adaptive matched filter detector. *IEEE TAES*, 1992. 2
- [36] Guang-Cai Sun, Yanbin Liu, Jixiang Xiang, Wenkang Liu, Mengdao Xing, and Jianlai Chen. Spaceborne synthetic aperture radar imaging algorithms: An overview. *IEEE GRSM*, 2021. 1
- [37] 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*, 2022. 1
- [38] Zhi Tian, Xiangxiang Chu, Xiaoming Wang, Xiaolin Wei, and Chunhua Shen. Fully convolutional one-stage 3d object detection on lidar range images. *NeurIPS*, 2022. 6, 7
- [39] Chao Wang, Rui Ruan, Zhicheng Zhao, Chenglong Li, and Jin Tang. Category-oriented localization distillation for sar object detection and a unified benchmark. *TGRS*, 2023. 1, 2, 6
- [40] Haipeng Wang, Feng Xu, and Shanshan Chen. Saliency detector for sar images based on pattern recurrence. *IEEE J-STARS*, 2016. 2
- [41] Jian Wang, Fan Li, and Haixia Bi. Gaussian focal loss: Learning distribution polarized angle prediction for rotated object detection in aerial images. *TGRS*, 2022. 3
- [42] Yuanyuan Wang, Chao Wang, Hong Zhang, Yingbo Dong, and Sisi Wei. A sar dataset of ship detection for deep learning under complex backgrounds. *Remote Sensing*, 2019. 6
- [43] Zhirui Wang, Yuzhuo Kang, Xuan Zeng, Yuelei Wang, Ting Zhang, and Xian Sun. Sar-aircraft-1.0: High-resolution sar aircraft detection and recognition dataset. *Journal of Radars*, 2023. 2, 6
- [44] Shunjun Wei, Xiangfeng Zeng, Qizhe Qu, Mou Wang, Hao Su, and Jun Shi. Hrsid: A high-resolution sar images dataset for ship detection and instance segmentation. *IEEE Access*, 2020. 1, 2, 6
- [45] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dots: A large-scale dataset for object detection in aerial images. In *CVPR*, 2018. 1, 6, 7
- [46] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *ICCV*, 2021. 1, 7
- [47] Congan Xu, Hang Su, Jianwei Li, Yu Liu, Libo Yao, Long Gao, Wenjun Yan, and Taoyang Wang. Rsdd-sar: Rotated ship detection dataset in sar images. *Journal of Radars*, 2022. 6
- [48] Hang Xu, Xinyuan Liu, Haonan Xu, Yike Ma, Zunjie Zhu, Chenggang Yan, and Feng Dai. Rethinking boundary discontinuity problem for oriented object detection. In *CVPR*, 2024. 1, 2, 3, 4, 7
- [49] Xue Yang and Junchi Yan. On the arbitrary-oriented object detection: Classification based approaches revisited. *IJCV*, 2022. 3
- [50] Xue Yang, Jirui Yang, Junchi Yan, Yue Zhang, Tengfei Zhang, Zhi Guo, Xian Sun, and Kun Fu. Srdet: Towards more robust detection for small, cluttered and rotated objects. In *ICCV*, 2019. 3
- [51] Xue Yang, Liping Hou, Yue Zhou, Wentao Wang, and Junchi Yan. Dense label encoding for boundary discontinuity free rotation detection. In *CVPR*, 2021. 3
- [52] Xue Yang, Junchi Yan, Ziming Feng, and Tao He. R3det: Refined single-stage detector with feature refinement for rotating object. In *AAAI*, 2021. 3, 7
- [53] 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. 1, 3, 6
- [54] 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. 3
- [55] Xue Yang, Gefan Zhang, Wentong Li, Xuehui Wang, Yue Zhou, and Junchi Yan. H2rbox: Horizontal box annotation is all you need for oriented object detection. *ICLR*, 2023. 3, 7
- [56] Xue Yang, Yue Zhou, Gefan Zhang, Jirui Yang, Wentao Wang, Junchi Yan, Xiaopeng Zhang, and Qi Tian. The kfiou loss for rotated object detection. *ICLR*, 2023. 3
- [57] Zhongyu Yang, Chen Shen, Wei Shao, Tengfei Xing, Runbo Hu, Pengfei Xu, Hua Chai, and Ruini Xue. Ldtr: Transformer-based lane detection with anchor-chain representation. *Computational Visual Media*, 2024. 3
- [58] Yi Yu and Feipeng Da. Phase-shifting coder: Predicting accurate orientation in oriented object detection. In *CVPR*, 2023. 1, 2, 3, 4, 7
- [59] Yi Yu, Xue Yang, Qingyun Li, Feipeng Da, Jifeng Dai, Yu Qiao, and Junchi Yan. Point2rbox: Combine knowledge from synthetic visual patterns for end-to-end oriented object detection with single point supervision. In *CVPR*, 2024. 3
- [60] Yi Yu, Xue Yang, Qingyun Li, Yue Zhou, Feipeng Da, and Junchi Yan. H2rbox-v2: Incorporating symmetry for boosting horizontal box supervised oriented object detection. *NeurIPS*, 2024. 2, 3, 5, 6, 7
- [61] Ying Zeng, Yushi Chen, Xue Yang, Qingyun Li, and Junchi Yan. Ars-detr: Aspect ratio-sensitive detection transformer for aerial oriented object detection. *TGRS*, 2024. 3, 6, 7
- [62] Jinsong Zhang, Mengdao Xing, and Yiyuan Xie. Fec: A feature fusion framework for sar target recognition based on electromagnetic scattering features and deep cnn features. *TGRS*, 2020. 2
- [63] Tianwen Zhang, Xiaoling Zhang, and Xiao Ke. Quad-fpn: A novel quad feature pyramid network for sar ship detection. *Remote Sensing*, 2021. 3

- [64] Tianwen Zhang, Xiaoling Zhang, Jianwei Li, Xiaowo Xu, Baoyou Wang, Xu Zhan, Yanqin Xu, Xiao Ke, Tianjiao Zeng, Hao Su, et al. Sar ship detection dataset (ssdd): Official release and comprehensive data analysis. *Remote Sensing*, 2021. [1](#), [2](#), [6](#)
- [65] Wenqing Zhao and Lijiao Xu. Weakly supervised target detection based on spatial attention. *Visual Intelligence*, 2024. [3](#)
- [66] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object detection with deep learning: A review. *TNNLS*, 2019. [1](#)
- [67] Yu Zheng, Danyang Zhang, Sinan Xie, Jiwen Lu, and Jie Zhou. Rotation-robust intersection over union for 3d object detection. In *ECCV*, 2020. [1](#)
- [68] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: an efficient and accurate scene text detector. In *CVPR*, 2017. [1](#)
- [69] Yue Zhou, Xue Jiang, Guozheng Xu, Xue Yang, Xingzhao Liu, and Zhou Li. Pvt-sar: An arbitrarily oriented sar ship detector with pyramid vision transformer. *IEEE J-STARS*, 2022. [1](#)
- [70] 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. [6](#)
- [71] 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. [3](#), [7](#)