

Poly Kernel Inception Network for Remote Sensing Detection

Xinhao Cai^{1*}, Qiuxia Lai^{2*}, Yuwei Wang^{1*}, Wenguan Wang^{3†}, Zeren Sun¹, Yazhou Yao^{1†}

¹ Nanjing University of Science and Technology ² Communication University of China ³ Zhejiang University

<https://github.com/NUST-Machine-Intelligence-Laboratory/PKINet>

Abstract

Object detection in remote sensing images (RSIs) often suffers from several increasing challenges, including the large variation in object scales and the diverse-ranging context. Prior methods tried to address these challenges by expanding the spatial receptive field of the backbone, either through large-kernel convolution or dilated convolution. However, the former typically introduces considerable background noise, while the latter risks generating overly sparse feature representations. In this paper, we introduce the Poly Kernel Inception Network (PKINet) to handle the above challenges. PKINet employs multi-scale convolution kernels without dilation to extract object features of varying scales and capture local context. In addition, a Context Anchor Attention (CAA) module is introduced in parallel to capture long-range contextual information. These two components work jointly to advance the performance of PKINet on four challenging remote sensing detection benchmarks, namely DOTA-v1.0, DOTA-v1.5, HRSC2016, and DIOR-R.

1. Introduction

Object detection in remote sensing images (RSIs) has gained substantial attention in recent years [11, 56, 64]. This task is dedicated to discerning the presence of specific objects within RSIs and subsequently ascertaining their categories and precise locations. In contrast to generic object detection that typically produces horizontal bounding boxes, remote sensing object detection aims to generate bounding boxes that align accurately with the orientation of the objects. Consequently, numerous prior efforts have been dedicated to developing various oriented bounding box (OBB) detectors [10, 20, 31, 65, 67, 71] and improving the angle prediction accuracy for OBBs [68, 70, 72–74]. Nevertheless, the unique characteristics of RSIs remain relatively under-explored when it comes to improving the feature extraction for object detection.

RSIs, including aerial and satellite images, are typi-

*Equal contribution.

†Corresponding author.

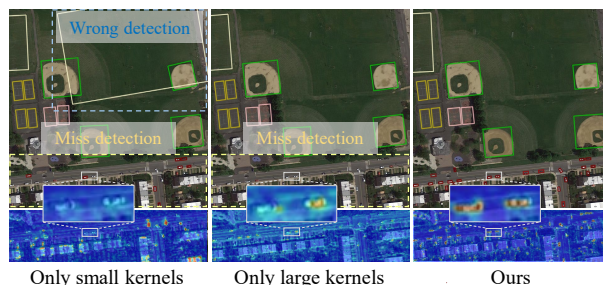
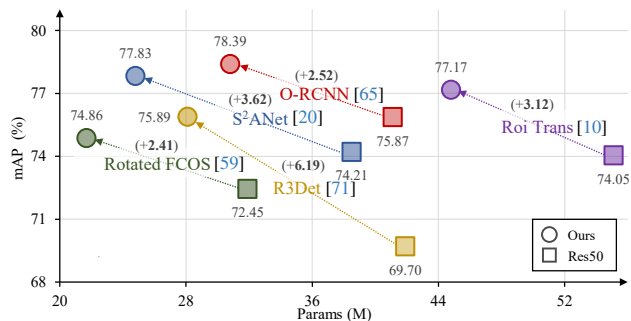


Figure 1. **Top:** Our approach yields solid performance gains over various remote sensing detectors [10, 20, 59, 65, 71] with fewer parameters on DOTA-v1.0 [64]. **Bottom:** Networks with small kernels miss long-range context in large object detection, whereas those with large kernels introduce noise for small objects. Our multi-scale convolution, however, handles scale variations well.

cally acquired from a bird’s-eye perspective, offering high-resolution views of the Earth’s surface. Consequently, objects depicted in RSIs exhibit a wide range of scales, spanning from expansive ones like soccer fields to relatively diminutive entities such as vehicles. Furthermore, the accurate recognition of these objects relies not solely on their appearances, but also on their contextual information, *i.e.*, the surrounding environment in which they are situated. To address the large variation in the object scales, some methods employ explicit data augmentation techniques [2, 54, 82] to improve the robustness of the features against scale variations. Some resort to multi-scale feature integration [37, 81] or pyramidal feature hierarchy [33, 61] to extract features rich in scale information. Nevertheless, a limitation of these methods is that the receptive fields for objects of varying scales remain identical, thereby failing to provide sufficient contextual information for larger objects.

Recently, LSKNet [32] proposes to selectively enlarge the spatial receptive field for larger objects to capture more scene context information. This is achieved by incorporating large-kernel convolutions [12, 18, 38, 43] and dilated convolutions into the backbone network. However, it is noteworthy that the use of large-kernel convolutions may introduce a significant amount of background noise, which could be detrimental to the accurate detection of small objects. On the other hand, dilated convolutions, though effective at enlarging the receptive field, might inadvertently overlook fine-grained details within that field, potentially resulting in overly sparse feature representations.

To address the challenges posed by the large variation in object scales and the diverse-ranging context within RSIs, in this paper, we present a powerful and lightweight feature extraction backbone network named Poly Kernel Inception Network (PKINet) for remote sensing object detection. Unlike previous methods that rely on large-kernel or dilated convolutions to expand the receptive field, PKINet arranges multiple depth-wise convolution kernels of different sizes without dilation in parallel, and extracts dense texture features across varying receptive fields. These texture features are adaptively fused along the channel dimension, enabling the collection of local contextual information. To further encompass long-range contextual information, we introduce a Context Anchor Attention (CAA) mechanism, which leverages global average pooling and 1D strip convolutions to capture the relationships between distant pixels and enhances the features within the central region. The two components work jointly to facilitate the extraction of adaptive features with both local and global contextual information, thereby improving the performance of remote sensing object detection.

To the best of our knowledge, PKINet represents the pioneering effort in exploring the application of inception-style convolutions and global context attention in remote sensing object detection, aiming to effectively tackle the challenges posed by the considerable variations in object scale and contextual diversity. Extensive experiments on widely used remote sensing benchmarks DOTA-v1.0 [64], DOTA-v1.5 [64], HRSC2016 [41], and DIOR-R [3] demonstrate the effectiveness of our method. In addition to its exceptional feature extraction capabilities, our model is lightweight compared with previous methods thanks to the strategic use of depth-wise and 1D convolutions.

2. Related Work

The challenges faced by remote sensing object detection primarily stem from objects with arbitrary orientations and substantial scale variations [3, 11, 40, 56, 64, 75]. The majority of previous methods have focused on oriented bounding box (OBB) detection. Nonetheless, an emerging trend is to design effective feature extraction backbones tailored

to the characteristics of remote sensing images (RSIs).

OBB for Remote Sensing Object Detection. To address the challenge of arbitrary orientations of the objects in RSIs, one research direction focuses on *developing specialized OBB detectors*. This includes introducing feature refinement techniques into the detector neck [69, 71], extracting the rotated region of interest (RoI) [10, 65], designing specific detection heads for the OBBs [21, 26, 48], *etc.* Though improving over general horizontal bounding box (HBB) detectors, these methods often suffer from issues like boundary discontinuity due to their relatively inflexible object representations obtained by augmenting horizontal object representations with additional angle parameters. To mitigate the aforementioned issues, another line of research has been dedicated to *developing new object representations* for detecting OBBs [15, 31, 62, 67, 70, 76]. For example, Xu *et al.* [67] propose to describe a multi-oriented object by adding four gliding offset variables to classical HBB representation. Li *et al.* [31] characterize oriented objects using a set of points to achieve more accurate orientation estimation. Some others [4, 27, 72, 73] utilize Gaussian distributions to model the OBBs for object detection and design new loss functions [51] to guide the learning process.

Although these methods are promising in addressing challenges related to arbitrary orientations, they typically rely on standard backbones for feature extraction, which often overlook the unique characteristics of RSIs that are essential for object detection, *e.g.*, the large object scale variations and the diverse contextual information. In contrast, we propose a feature extraction backbone to deal with the challenges posed by the large object scale variations.

Feature Extraction for Remote Sensing Object Detection. To better handle the unique challenges such as large object scale variations in RSIs, certain methods emphasize the *extraction of multi-scale features* through approaches like data augmentation [2, 54, 82], multi-scale feature integration [39, 61, 81, 83], feature pyramid network (FPN) enhancement [16, 25, 35, 80], or multi-scale anchor generation [19, 24, 52]. Recently, there has been a noteworthy development in the *design of feature extraction backbones* specifically for remote sensing object detection. Some [21, 50] focus on extracting features suitable for objects of varying orientations with equivalent receptive fields. Some [32] enlarge the spatial receptive field for larger objects using large kernels [12, 38, 43], which inevitably introduces background noise for smaller objects. Some [8, 17, 79] adopt multi-scale convolution kernels in order to address challenges across various fields, yet research in remote sensing detection remains scarce.

Similar to [32], we propose a new feature extraction backbone PKINet to address the challenges posed by the large variation in object scales and diverse context in RSIs. There are **two key differences** between the two methods.

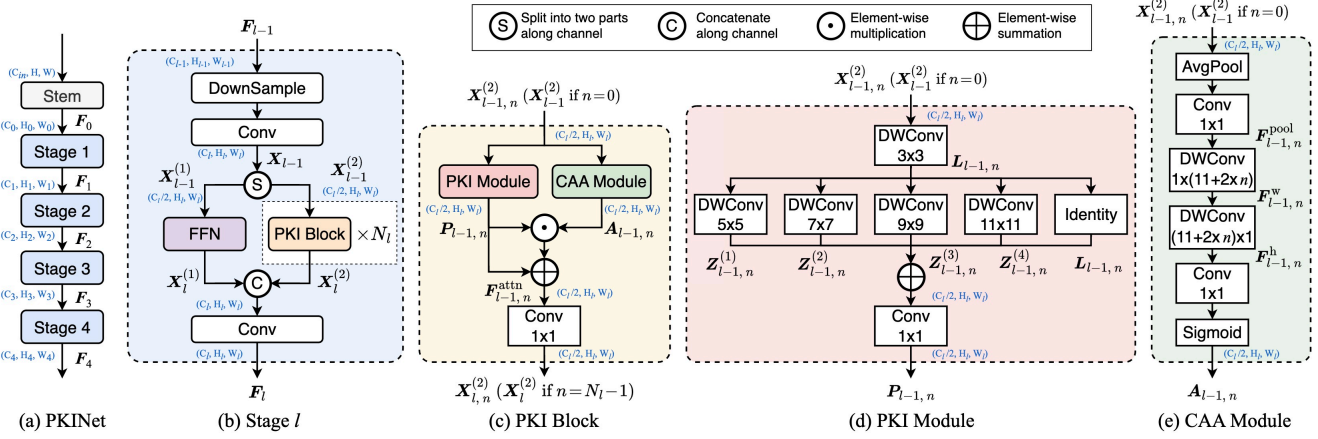


Figure 2. **PKINet overview.** (a) **PKINet** consists of four stages, where the spatial resolution of the l -th stage output is $(C_l \times H_l \times W_l)$. Each (b) **Stage** (§3.1) implies a cross-stage partial (CSP) structure, where the input is split in half along the channel dimension and fed to a Feed-Forward Network (FFN) and a sequence of N_l PKI Blocks, respectively. Each (c) **PKINet Block** contains a (d) **PKI Module** (§3.2) and a (e) **CAA Module** (§3.3). Here, $n = 0, \dots, N_l - 1$ means that the PKI/CAA Module is in the n -th PKI Block of the l -th stage.

Firstly, instead of relying on large-kernel or dilated convolutions to expand the receptive field, PKINet utilizes inception-style depth-wise convolution without dilation to extract multi-scale texture features across varying receptive fields. Secondly, our method incorporates a Context Anchor Attention (CAA) mechanism to capture the long-range contextual information. The two components collaborate to facilitate the extraction of adaptive features with both local and global contextual information, thereby improving the performance of remote sensing object detection.

3. Methodology

As shown in Fig. 2(a), our PKINet is a feature extraction backbone similar to VGG [55] and ResNet [22], which consists of four stages. Each stage (§3.1) implies a cross-stage partial (CSP) structure [60], where the stage input is split and fed into two paths. One path is a simple Feed-Forward Network (FFN). The other path consists of a sequence of PKI Blocks, and each PKI Block contains a PKI Module (§3.2) and a CAA Module (§3.3). The outputs of the two paths are concatenated to yield the output of the stage. PKINet can be incorporated with various oriented object detectors such as Oriented RCNN [65] to produce the final object detection results for RSIs.

3.1. PKI Stage

There are four stages arranged sequentially in PKINet. The input and output of stage l are $F_{l-1} \in \mathbb{R}^{C_{l-1} \times H_{l-1} \times W_{l-1}}$ and $F_l \in \mathbb{R}^{C_l \times H_l \times W_l}$, respectively. The structure of stage l is shown in Fig. 2(b), which implies a cross-stage partial (CSP) structure [60]. Specifically, the stage input F_{l-1} after initial processing is split in half along the channel dimension and fed into two paths:

$$\begin{aligned} X_{l-1} &= \text{Conv}_{3 \times 3}(\text{DS}(F_{l-1})) \in \mathbb{R}^{C_l \times H_l \times W_l}, \\ X_{l-1}^{(1)} &= X_{l-1}[: \frac{1}{2}C_l, \dots], X_{l-1}^{(2)} = X_{l-1}[\frac{1}{2}C_l :, \dots], \end{aligned} \quad (1)$$

where DS denotes the downsampling operation. One path is a simple Feed-Forward Network (FFN), which takes in $X_{l-1}^{(1)} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$ and then output $X_l^{(1)} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$. The other path consists of a sequence of N_l PKI Blocks, which processes $X_{l-1}^{(2)} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$ and yields $X_l^{(2)} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$. As shown in Fig. 2 (c), PKI Block contains a PKI Module and a CAA Module, which will be detailed in §3.2 and §3.3, respectively. The final output of stage l is:

$$F_l = \text{Conv}_{1 \times 1}(\text{Concat}(X_l^{(1)}, X_l^{(2)})) \in \mathbb{R}^{C_l \times H_l \times W_l}, \quad (2)$$

where Concat refers to the concatenation operation.

3.2. PKI Module

A PKINet Block consists of a PKI Module and a CAA Module. In this section, we look into the details of PKI Module. We will introduce CAA Module in §3.3.

As discussed in §1, different from general object detection, remote sensing object detection aims to locate and recognize objects of varying sizes within a single image. To address the challenges related to large variations in object scales, we introduce *PKI Module* to capture multi-scale texture features. As showcased in Fig. 2(d), PKI Module is an inception-style module [57, 77] that comprises a small-kernel convolution to grasp local information, followed by a set of parallel depth-wise convolutions to capture contextual information across multiple scales. Formally, PKI Module within the n -th PKI Block of the l -th stage can be represented mathematically as follows:

$$\begin{aligned} L_{l-1, n} &= \text{Conv}_{k_s \times k_s}(X_{l-1, n}^{(2)}), n = 0, \dots, N_l - 1, \\ Z_{l-1, n}^{(m)} &= \text{DWConv}_{k^{(m)} \times k^{(m)}}(L_{l-1, n}), m = 1, \dots, 4. \end{aligned} \quad (3)$$

Here, $L_{l-1, n} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$ is the local feature extracted by the $k_s \times k_s$ convolution, and $Z_{l-1, n}^{(m)} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$ is the context feature extracted by the m -th $k^{(m)} \times k^{(m)}$ depth-wise convolution (DWConv). In our experiment, we set

$k_s = 3$ and $k^{(m)} = (m + 1) \times 2 + 1$. For $n = 0$, we have $\mathbf{X}_{l-1, n}^{(2)} = \mathbf{X}_{l-1}^{(2)}$. Note that our PKI Module does not use dilated convolution, thereby preventing the extraction of overly sparse feature representations.

Then, the local and contextual features are fused by a convolution of size 1×1 , characterizing the interrelations among various channels:

$$\mathbf{P}_{l-1, n} = \text{Conv}_{1 \times 1}(\mathbf{L}_{l-1, n} + \sum_{m=1}^4 \mathbf{Z}_{l-1, n}^{(m)}), \quad (4)$$

where $\mathbf{P}_{l-1, n} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$ represents the output feature. The 1×1 convolution serves as a channel fusion mechanism to integrate features with varying receptive field sizes. In this way, our PKI Module could capture a broad spectrum of contextual information without compromising the integrity of local texture features.

3.3. Context Anchor Attention (CAA)

As discussed above, the inception-style PKI Module in the PKI Block focuses on extracting multi-scale local contextual information. To capture long-range contextual information, inspired by [32, 58], we further integrated a Context Anchor Attention (CAA) module into the PKI Block. CAA aims to grasp contextual interdependencies among distant pixels while augmenting the central features concurrently. An illustration of CAA is presented in Fig. 2(e). Take the CAA Module in the n -th PKI Block of the l -th stage as an example, we adopt average pooling followed by a 1×1 convolution to obtain the local region feature:

$$\mathbf{F}_{l-1, n}^{\text{pool}} = \text{Conv}_{1 \times 1}(\mathcal{P}_{\text{avg}}(\mathbf{X}_{l-1, n}^{(2)})), \quad n=0, \dots, N_l-1, \quad (5)$$

where \mathcal{P}_{avg} represents average pooling operation. For $n = 0$, we have $\mathbf{X}_{l-1, n}^{(2)} = \mathbf{X}_{l-1}^{(2)}$. Then, we apply two depth-wise strip convolutions as an approximation to a standard large-kernel depth-wise convolution:

$$\begin{aligned} \mathbf{F}_{l-1, n}^{\text{w}} &= \text{DWConv}_{1 \times k_b}(\mathbf{F}_{l-1, n}^{\text{pool}}), \\ \mathbf{F}_{l-1, n}^{\text{h}} &= \text{DWConv}_{k_b \times 1}(\mathbf{F}_{l-1, n}^{\text{w}}). \end{aligned} \quad (6)$$

We opt for depth-wise strip convolutions based on two primary considerations. First, strip convolution is lightweight. Compared to a conventional $k_b \times k_b$ 2D depth-wise convolution, we can achieve a similar effect with a couple of 1D depth-wise kernels with a parameter reduction of $k_b/2$. Second, strip convolution can facilitate the identification and extraction of features for objects with slender shapes, such as bridges. To increase the receptive field of CAA Module as the PKI Block it belongs to goes deeper, we set $k_b = 11 + 2 \times l$, *i.e.*, we calculate the kernel size k_b as the function of the PKI Block depth n . Such a design enhances the ability of PKINet to establish the relationship between long-range pixels, and would not significantly increase the computational cost thanks to the strip depth-wise design.

Finally, our CAA Module produces an attention weight

	$\frac{H_l}{H} \times \frac{W_l}{W}$	Layer Specification		PKINet	
				T	S
Stem	$\frac{1}{2} \times \frac{1}{2}$	Down-samp.	Kernel Size	3×3 , stride 2	
Stage 1	$\frac{1}{4} \times \frac{1}{4}$	Down-sampling	Kernel Size	3×3 , stride 2	
			Embed. Dim	32	64
		PKI Block	Kernel Size	3×3 to 11×11	
			#Block (N_1)	4	
Stage 2	$\frac{1}{8} \times \frac{1}{8}$	Down-sampling	Kernel Size	3×3 , stride 2	
			Embed. Dim	64	128
		PKI Block	Kernel Size	3×3 to 11×11	
			#Block (N_2)	14	12
Stage 3	$\frac{1}{16} \times \frac{1}{16}$	Down-sampling	Kernel Size	3×3 , stride 2	
			Embed. Dim	128	256
		PKI Block	Kernel Size	3×3 to 11×11	
			#Block (N_3)	22	20
Stage 4	$\frac{1}{32} \times \frac{1}{32}$	Down-sampling	Kernel Size	3×3 , stride 2	
			Embed. Dim	256	512
		PKI Block	Kernel Size	3×3 to 11×11	
			#Block (N_4)	4	
Parameters (M)				4.13	13.69
FLOPs (G)				22.70	70.20

Table 1. **Configurations of two variants of PKINet.** Here, ‘‘T’’ denotes ‘‘Tiny’’, and ‘‘S’’ denotes ‘‘Small’’. See §3.4 for details.

$\mathbf{A}_{l-1, n} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$, which is further used to enhance the output of PKI Module (*cf.* Eq. (4)):

$$\begin{aligned} \mathbf{A}_{l-1, n} &= \text{Sigmoid}(\text{Conv}_{1 \times 1}(\mathbf{F}_{l-1, n}^{\text{h}})), \\ \mathbf{F}_{l-1, n}^{\text{attn}} &= (\mathbf{A}_{l-1, n} \odot \mathbf{P}_{l-1, n}) \oplus \mathbf{P}_{l-1, n}. \end{aligned} \quad (7)$$

Here, Sigmoid function ensures that the attention map $\mathbf{A}_{l-1, n}$ is in range $(0, 1)$, \odot denotes the element-wise multiplication, \oplus denotes the element-wise summation, and $\mathbf{F}_{l-1, n}^{\text{attn}} \in \mathbb{R}^{\frac{1}{2}C_l \times H_l \times W_l}$ is the enhanced feature. The output of the n -th PKI Block in the l -th stage is obtained by:

$$\mathbf{X}_{l, n}^{(2)} = \text{Conv}_{1 \times 1}(\mathbf{F}_{l-1, n}^{\text{attn}}). \quad (8)$$

For $n = N_l - 1$, we have $\mathbf{X}_l^{(2)} = \mathbf{X}_{l, n}^{(2)}$, *i.e.*, we denote the output of the last PKI Block as $\mathbf{X}_l^{(2)}$.

3.4. Implementation Details

In this paper, we present two variants of the proposed backbone, namely **PKINet-T** and **PKINet-S**, where ‘‘T’’ stands for ‘‘Tiny’’, and ‘‘S’’ stands for ‘‘Small’’. The Stem structure consists of three 3×3 convolution layers with strides $(2, 1, 1)$, respectively. For both PKINet-T and PKINet-S, $H_l = H/2^{(l+1)}$, $W_l = W/2^{(l+1)}$ for $l = 0, \dots, 4$, and H, W are the height and width of the input, respectively. For PKINet-T, $C_0 = 32$, $C_l = 2^{l-1} \times C_0$ for $l = 1, \dots, 4$, and the number of PKI Blocks of the four stages are $(4, 14, 22, 4)$, respectively. For PKINet-S, $C_0 = 64$, $C_l = 2^{l-1} \times C_0$ for $l = 1, \dots, 4$, and the number of PKI Blocks of the four stages are $(4, 12, 20, 4)$, respectively. Please note that although PKINet-T comprises more PKI Blocks compared to PKINet-S, it contains significantly fewer parameters owing to the halved channel number in the intermediate features. The detailed configurations of the two variants of PKINet are listed in Table 1.

Method	Backbone	#P ↓	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	mAP ↑
DETR-based																		
AO ² -DETR [7]	ResNet-50 [22]	40.8M	87.99	79.46	45.74	66.64	78.90	73.90	73.30	90.40	80.55	85.89	55.19	63.62	51.83	70.15	60.04	70.91
O ² -DETR [47]	ResNet-50 [22]	-	86.01	75.92	46.02	66.65	79.70	79.93	89.17	90.44	81.19	76.00	56.91	62.45	64.22	65.80	58.96	72.15
ARS-DETR [78]	ResNet-50 [22]	41.6M	86.61	77.26	48.84	66.76	78.38	78.96	87.40	90.61	82.76	82.19	54.02	62.61	72.64	72.80	64.96	73.79
One-stage																		
SASM [26]	ResNet-50 [22]	36.6M	86.42	78.97	52.47	69.84	77.30	75.99	86.72	90.89	82.63	85.66	60.13	68.25	73.98	72.22	62.37	74.92
R3Det-GWD [72]	ResNet-50 [22]	41.9M	88.82	82.94	55.63	72.75	78.52	83.10	87.46	90.21	86.36	85.44	64.70	61.41	73.46	76.94	57.38	76.34
R3Det-KLD [73]	ResNet-50 [22]	41.9M	88.90	84.17	55.80	69.35	78.72	84.08	87.00	89.75	84.32	85.73	64.74	61.80	76.62	78.49	70.89	77.36
O-RepPoints [31]	ResNet-50 [22]	36.6M	87.02	83.17	54.13	71.16	80.18	78.40	87.28	90.90	85.97	86.25	59.90	70.49	73.53	72.27	58.97	75.97
Rotated FCOS [59]	ResNet-50 [22]	31.9M	88.52	77.54	47.06	63.78	80.42	80.50	87.34	90.39	77.83	84.13	55.45	65.84	66.02	72.77	49.17	72.45
	PKINet-S	21.7M	88.56	82.89	47.96	58.20	81.09	83.09	88.23	90.88	84.57	85.81	57.98	66.26	75.12	80.93	51.39	74.86
	ResNet-50 [22]	41.9M	89.00	75.60	46.64	67.09	76.18	73.40	79.02	90.88	78.62	84.88	59.00	61.16	63.65	62.39	37.94	69.70
R3Det [71]	ARC-R50 [50]	65.2M	89.49	78.04	46.36	68.89	77.45	72.87	82.76	90.90	83.07	84.89	58.72	68.61	64.75	68.39	49.67	72.32
	PKINet-S	28.1M	89.63	82.40	49.77	71.72	79.95	81.39	87.79	90.90	84.20	86.09	61.08	66.55	73.06	73.85	59.95	75.89
	ResNet-50 [22]	38.5M	89.11	82.84	48.37	71.11	78.11	78.39	87.25	90.83	84.90	85.64	60.36	62.60	65.26	69.13	57.94	74.12
S ² ANet [20]	ARC-R50 [50]	71.8M	89.28	78.77	53.00	72.44	79.81	77.84	86.81	90.88	84.27	86.20	60.74	68.97	66.35	71.25	65.77	75.49
	PKINet-S	24.8M	89.67	84.16	51.94	71.89	80.81	83.47	88.29	90.80	87.01	86.94	65.02	69.53	75.83	80.20	61.85	77.83
Two-stage																		
SCRDet [69]	ResNet-50 [22]	41.9M	89.98	80.65	52.09	68.36	68.36	60.32	72.41	90.85	87.94	86.86	65.02	66.68	66.25	68.24	65.21	72.61
G.V. [67]	ResNet-50 [22]	41.1M	89.64	85.00	52.26	77.34	73.01	73.14	86.82	90.74	79.02	86.81	59.55	70.91	72.94	70.86	57.32	75.02
CenterMap [44]	ResNet-50 [22]	41.1M	89.02	80.56	49.41	61.98	77.99	74.19	83.74	89.44	78.01	83.52	47.64	65.93	63.68	67.07	61.59	71.59
ReDet [21]	ResNet-50 [22]	31.6M	88.79	82.64	53.97	74.00	78.13	84.06	88.04	90.89	87.78	85.75	61.76	60.39	75.96	68.07	63.59	76.25
Roi Trans [10]	ResNet-50 [22]	55.1M	89.01	77.48	51.64	72.07	74.43	77.55	87.76	90.81	79.71	85.27	58.36	64.11	76.50	71.99	54.06	74.05
	PKINet-S	44.8M	89.33	85.59	55.75	74.69	74.69	79.13	88.05	90.90	87.43	86.90	61.67	64.25	77.77	75.38	66.08	77.17
Rotated Faster R-CNN [53]	ResNet-50 [22]	41.1M	89.40	81.81	47.28	67.44	73.96	73.12	85.03	90.90	85.15	84.90	56.60	64.77	64.70	70.28	62.22	73.17
	ARC-R50 [50]	74.4M	89.49	82.11	51.02	70.38	79.07	75.06	86.18	90.91	84.23	86.41	56.10	69.42	65.87	71.90	63.47	74.77
	PKINet-S	30.8M	89.33	85.27	52.34	73.03	73.72	75.60	86.97	90.88	86.52	87.30	64.23	64.20	75.63	80.31	61.47	76.45
O-RCNN [65]	ResNet-50 [22]	41.1M	89.46	82.12	54.78	70.86	78.93	83.00	88.20	90.90	87.50	84.68	63.97	67.69	74.94	68.84	52.28	75.87
	ARC-R50 [50]	74.4M	89.40	82.48	55.33	73.88	79.37	84.05	88.06	90.90	86.44	84.83	63.63	70.32	74.29	71.91	65.43	77.35
	LSKNet-S [32]	31.0M	89.66	85.52	57.72	75.70	74.95	78.69	88.24	90.88	86.79	86.38	66.92	63.77	77.77	74.47	64.82	77.49
	PKINet-S	30.8M	89.72	84.20	55.81	77.63	80.25	84.45	88.12	90.88	87.57	86.07	66.86	70.23	77.47	73.62	62.94	78.39

Table 2. **Experimental results on DOTA-v1.0 dataset [64]** under single-scale training and testing setting. PKINet-S is pretrained on ImageNet-1K [9] for 300 epochs similar to the compared methods [10, 65, 71]. See §4.2 for details.

4. Experiment

4.1. Experimental Setup

Datasets. We conduct extensive experiments on four popular remote sensing object detection datasets:

- **DOTA-v1.0 [64]** is a large-scale dataset for remote sensing detection which contains 2806 images, 188,282 instances, and 15 categories with a large variety of orientations and scales. The dataset is comprised of 1,411, 458, and 937 images for `train`, `val`, and `test`, respectively.
- **DOTA-v1.5 [64]** is a more challenging dataset based on DOTA-v1.0 which is released for DOAI Challenge 2019. This iteration includes the addition of a novel category named `Container Crane (CC)` and a substantial increase in the number of minuscule instances that are less than 10 pixels, containing 403,318 instances in total.
- **HRSC2016 [41]** is a remote sensing dataset for ship detection that contains 1061 aerial images whose size ranges from 300×300 and 1500×900 . The images splits into 436/181/444 for `train/val/test`.
- **DIOR-R [3]** provides OBB annotations based on remote sensing dataset DIOR [30] dataset. It contains 23,463 im-

Backbone	#Params ↓	#FLOPs ↓	mAP ↑
ResNet-18 [22]	11.2M	38.1G	74.20
PKINet-T (ours)	4.1M	22.7G	77.87
ResNet-50 [22]	23.3M	86.1G	75.87
PKINet-S (ours)	13.7M	70.2G	78.39

Table 3. **Comparison with ResNet [22] backbone on DOTA-v1.0 dataset [64]**. Params and FLOPs are computed for backbones only. All the backbones are pretrained on ImageNet-1K [9] for 300 epochs and built within Oriented RCNN [65]. See §4.2 for details.

ages with the size of 800×800 and 192,518 annotations.

Training. Our training process contains ImageNet [9] pre-train and remote sensing object detector training. For ImageNet pre-train, our PKINet is trained on the ImageNet-1K under the MMPretrain [6] toolbox. In the main experiment, we train it for 300 epochs for higher performance like previous works [32, 50, 65, 71]. In the process of pre-train, we adapt the AdamW [29] optimizer with a momentum of 0.9 and a weight decay of 0.05. Cosine schedule [45] and warm-up strategy are employed to adjust the learning rate. We use 8 GPUs with a batch size of 1024 for pertaining. For remote sensing object detector training, experiments are conducted on MMRotate [84] framework. To compare with

Method	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	CC	mAP \uparrow
RetinaNet-O [36]	71.43	77.64	42.12	64.65	44.53	56.79	73.31	90.84	76.02	59.96	46.95	69.24	59.65	64.52	48.06	0.83	59.16
FR-O [53]	71.89	74.47	44.45	59.87	51.28	68.98	79.37	90.78	77.38	67.50	47.75	69.72	61.22	65.28	60.47	1.54	62.00
Mask R-CNN [23]	76.84	73.51	49.90	57.80	51.31	71.34	79.75	90.46	74.21	66.07	46.21	70.61	63.07	64.46	57.81	9.42	62.67
HTC [1]	77.80	73.67	51.40	63.99	51.54	73.31	80.31	90.48	75.12	67.34	48.51	70.63	64.84	64.48	55.87	5.15	63.40
ReDet [21]	79.20	82.81	51.92	71.41	52.38	75.73	80.92	90.83	75.81	68.64	49.29	72.03	73.36	70.55	63.33	11.53	66.86
DCFL [66]	-	-	-	-	56.72	-	80.87	-	-	75.65	-	-	-	-	-	-	67.37
LSKNet-S [32]	72.05	84.94	55.41	74.93	52.42	77.45	81.17	90.85	79.44	69.00	62.10	73.72	77.49	75.29	55.81	42.19	70.26
PKINet-S (ours)	80.31	85.00	55.61	74.38	52.41	76.85	88.38	90.87	79.04	68.78	67.47	72.45	76.24	74.53	64.07	37.13	71.47

Table 4. **Experimental results on DOTA-v1.5 dataset** [64] compared with state-of-the-art methods with single-scale training and testing. PKINet-S backbone is pretrained on ImageNet-1K [9] for 300 epochs, as the compared methods [21, 66]. PKINet-S is built within the framework of Oriented RCNN [65]. See §4.2 for details.

other methods, we use `trainval` sets of these benchmarks and their `test` sets for testing. Following the settings of previous methods [21, 65, 71, 78], we crop original images into 1024×1024 patches with overlaps of 200 for DOTA-v1.0 and DOTA-v1.5 datasets. For HRSC2016 and DIOR-R datasets, the input size is set as 800×800 . Models are trained with 30 epochs, 30 epochs, 60 epochs, and 36 epochs for DOTA-v1.0, DOTA-v1.5, HRSC2016, and DIOR-R. We employ AdamW [29] optimizer with a weight decay of 0.05. The initial learning rate is set to 0.0002. All flops reported are calculated when the input image size is 1024×1024 . To prevent over-fitting, images undergo random resizing and flipping during training following previous methods [21, 65, 71, 78]. Five-run average mAP of our method are reported for HRSC2016 and DIOR-R.

Testing. The image resolution at the testing stage remains consistent with the training stage. For the sake of fairness, we do not apply any test-time data augmentation.

Evaluation Metric. The mean average precision (mAP) and the Average Precision at 0.5 threshold (AP_{50}) are reported.

Reproducibility. Our algorithm is implemented in PyTorch. We use eight NVIDIA RTX 4090 GPUs for ImageNet pre-training and four NVIDIA Tesla V100 GPUs for downstream training and testing.

4.2. Quantitative Results

Performance on DOTA-v1.0 [64]. To begin with, we make a comparison with ResNet [22] built within the framework of Oriented RCNN [65] on DOTA-v1.0 in Table 3. PKINet-T outperforms by **3.67%** using only **36.7%** of the parameters and **59.6%** of the compute needed by ResNet-18. PKINet-S also excels, improving by **2.52%** with just **58.8%** of the parameters and **81.53%** of the compute of ResNet-50.

Our PKINet backbone, when paired with multiple detection architectures shown in Table 2, consistently outperforms ResNet-50 and more networks designed for the remote sensing detection task (*i.e.*, ARC [50] and LSKNet [32]). For one-stage architectures, our backbone is able to bring **2.41%/6.19%/3.71%** mAP improvement compared with ResNet-50 for Rotated FCOS [59], R3Det [71], and S²ANet [20] respectively. Even when integrated with the classical S²ANet [20], our method surpasses the previous

Method	#Params \downarrow	mAP (07) \uparrow	mAP (12) \uparrow
DRN [49]	-	-	92.70
GWD [72]	47.4M	89.85	97.37
Rot Trans. [10]	55.1M	86.20	-
Gliding Vertex [67]	41.1M	88.20	-
CenterMap [44]	41.1M	-	92.80
AOPG [3]	-	90.34	96.22
R3Det [71]	41.9M	89.26	96.01
S ² ANet [20]	38.6M	90.17	95.01
ReDet [21]	31.6M	90.46	97.63
O-RCNN [65]	41.1M	90.50	97.60
O-RepPoints [31]	36.6M	90.38	97.26
LSKNet [32]	31.0M	90.65	98.46
PKINet-S (ours)	30.8M	90.70	98.54

Table 5. **Experimental results on HRSC2016 dataset** [41]. PKINet-S is pretrained on ImageNet-1K [9] for 300 epochs which is consistent with previous methods [21, 32, 46] and built within the framework of Oriented RCNN [65]. mAP (07/12): VOC 2007 [13]/2012 [14] metrics. See §4.2 for details.

Method	RetinaNet-O [36]	FR-OBB [53]	RT [10]	LSKNet-S [32]
mAP \uparrow	57.55	59.54	63.87	65.90
Method	GGHL [28]	Oriented Rep [65]	DCFL [66]	PKINet-S (ours)
mAP \uparrow	66.48	66.71	66.80	67.03

Table 6. **Experimental results on DIOR-R dataset** [30]. Following previous methods [28, 65, 66], PKINet-S is pretrained on ImageNet-1K [9] for 300 epochs and built within the framework of Oriented RCNN [65]. See §4.2 for details.

approaches, achieving a performance of **77.83%**. For two-stage architectures, PKINet also achieves remarkable gain (**3.12%/3.28%/2.23%**). When equipped with the advanced detector Oriented RCNN [65], the performance reaches the superior **78.39%** with a remarkable performance improvement for small categories compared with the previous best method LSKNet [32] (**5.3%/5.76%** for SV/LV). For the RA category that needs more contextual information, PKINet also achieves **6.46%** enhancement compared with LSKNet.

Performance on DOTA-v1.5 [64]. As shown in Table 4, our approach achieves outstanding performance on the more challenging dataset DOTA-v1.5 with minuscule instances, evidencing its efficacy and generalization ability to small objects. Our PKINet outperforms the former state-of-the-art methods, achieving an improvement of **1.21%**.

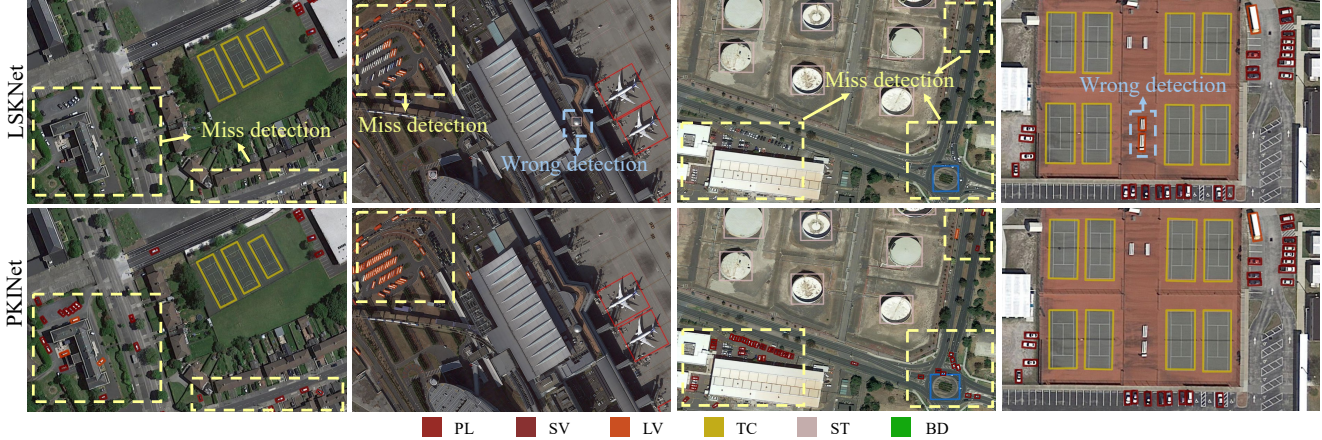


Figure 3. Visual results on DOTA-v1.0 dataset [64]. Top: LSKNet [32]; Bottom: our PKINet. See §4.3 for details.

Method	ResNet-18 [22]	PVT-T [63]	ConvNeXt-N [43]	PKINet-S (ours)
#Params ↓	11.2	12.9	15.6	13.7
mAP ↑	34.0	36.7	41.3	43.4

Table 7. Experimental results on COCO 2017 dataset [34]. All models are pretrained on ImageNet-1K [9] for 300 epochs and are based on Mask R-CNN [23]. See §4.2 for details.

Performance on HRSC2016 [41]. Our PKINet-S surpasses 12 leading methods on the HRSC2016 dataset with fewer parameters, as illustrated in Table 5. The slight edge over LSKNet [32] mainly stems from HRSC2016 merging 31 subclasses into a single ‘ship’ category for training and testing. This protocol doesn’t fully showcase our method’s strengths in managing inter-class object size variation.

Performance on DIOR-R [3]. We present comparison results on DIOR-R, as shown in Table 6. We achieve the best performance with 67.03%.

Performance on COCO 2017 [34]. To assess the versatility of PKINet as a general framework adaptable to various forms of bounding boxes, we evaluate our method on the widely-used general detection benchmark COCO. As can be seen in Table 7, PKINet outperforms several famous backbones with similar parameters, thereby further affirming the efficacy of our method as a general-purpose backbone that is not confined to RSIs.

4.3. Qualitative Results

Fig. 3 depicts representative visual results on DOTA [64]. As seen, compared to the previous best-performing method LSKNet [32] which merely relies on large kernels, our PKINet demonstrates a strong ability to adjust to significant size variations of target objects in a scene, ensuring detection of larger items (*e.g.*, PL, TC, ST, and BD) while retaining focus on smaller ones (*e.g.*, SV and LV).

4.4. Diagnostic Experiments

To gain more insights into PKINet, a set of ablative studies on DOTA-v1.0 is conducted with Oriented RCNN [65] as

the detector. All the backbones mentioned in this section are trained on ImageNet-1K [9] for 100 epochs for efficiency.

Multi-scale Kernel Design. First, the critical multi-scale kernel design in PKINet (*cf.* §3.2) is investigated in Table 8a. It demonstrates that using only small 3×3 kernels yields poor performance due to limited texture information extraction. Then a multi-scale kernel structure is adopted whose kernel size ranges from 3×3 to 11×11 with a stride of 2. Under this setting, the model shows the best performance. Next, a stride of 4 when the kernel size increases is tested and its performance is sub-optimal. Further trials with only large kernels led to increased computation but decreased performance, dropping by 0.49% and 0.84%, indicating that large kernels may introduce background noise and bring a performance drop (*cf.* §1).

Then, we investigate the kernel number in multi-scale kernel design, detailed in §3.2. As Table 8b shows, with only two kernels (only 3×3 and 5×5 kernels are reserved), the network can’t capture long-range pixel relationships. As the number of kernels rises, network performance improves, achieving optimal outcomes with five kernels.

Kernel Dilations. Then, we examine the effect of dilations in our PKI module (*cf.* §3.2). As displayed in Table 8d, there is a performance degradation (-1.09%) despite the increase in the receptive field compared to no kernel dilations. As we further increase the degree of dilation, a further drop in performance occurs. This proves that merely applying dilation to expand the receptive field does not work.

Context Anchor Attention. Next, the effectiveness of CAA module (*cf.* §3.3) is proved. To start with, CAA is applied with different kernel sizes to check the impact in Table 8f. The three kernel sizes in the first column represent the size in average pooling and two strip convolutions. As can be seen, smaller kernels fail to capture long-range dependencies, reducing performance, while larger kernels improve this by including more context. Our expansive kernel size strategy that increases the kernel size of strip convo-

Kernel Design	#Params ↓	#FLOPs ↓	mAP ↑	Kernel Number	#Params ↓	#FLOPs ↓	mAP ↑	Stage Apply	#Params ↓	#FLOPs ↓	mAP ↑
(3, 3, 3, 3, 3)	12.62M	62.40G	76.94	2	12.56M	61.95G	75.76	None	12.03M	61.72G	77.13
(3, 5, 7, 9, 11)	13.69M	70.20G	78.16	3	12.78M	63.57G	76.07	1	12.19M	64.04G	77.35
(3, 5, 9, 13, 17)	14.99M	79.57G	78.07	4	13.13M	66.24G	77.53	2	12.31M	65.45G	77.48
(11, 11, 11, 11, 11)	15.13M	80.61G	77.67	5	13.69M	70.20G	78.16	3	12.97M	66.59G	77.72
(15, 15, 15, 15, 15)	17.44M	92.45G	77.32	6	14.35M	75.26G	78.10	ALL	13.69M	70.20G	78.16

(a) multi-scale kernel design

(b) kernel number

(c) location for implementing CAA

Kernel Design	#Params ↓	#FLOPs ↓	mAP ↑
(3, 3, 3)	13.50M	68.95G	77.52
(5, 5, 5)	13.52M	69.08G	77.71
(5, 7, 7)	13.54M	69.21G	77.76
(7, 11, 11)	13.58M	69.47G	77.89
Expansive	13.69M	70.20G	78.16

(f) kernel size in CAA

Kernel Dilations	Max RF	mAP ↑	CSP	Blocks	#Params ↓	#FLOPs ↓	mAP ↑
(1, 1, 1, 1, 1)	13	78.16	✓	(4, 12, 20, 4)	13.69M	70.20G	78.16
(2, 2, 2, 2, 2)	24	77.07	×	(4, 12, 20, 4)	42.59M	182.07G	-
(3, 3, 3, 3, 3)	36	76.95	×	(2, 2, 4, 2)	17.30M	58.60G	77.83

(d) kernel dilations

(e) cross-stage partial structure

Table 8. **A set of ablative studies on DOTA-v1.0** [64]. The adopted network designs are marked in red. All the networks are pretrained on ImageNet-1K [9] for 100 epochs and built with the framework of Oriented RCNN [65]. See §4.4 for details.

lutions as the blocks deepen achieves the best performance.

After that, since there are four stages in our PKINet, how the implementing location affects the final performance is investigated. As revealed in Table 8c, CAA module (*c.f.* §3.3) can bring performance improvement when implemented at any stage. Consequently, when deploying CAA module at all stages, the performance gain reaches **1.03%**. **Cross-Stage Partial Structure.** Table 8e further explores the impact of the Cross-Stage Partial (CSP) structure. Eliminating CSP leads to exponential increases in both parameters and computational costs (by **211%** and **159%**, respectively). Reducing stage blocks from (4, 12, 20, 4) to (2, 2, 4, 2) allows models without CSP structure to achieve a similar parameter count as the former one, but results in sub-optimal performance due to fewer blocks.

4.5. Analysis

To measure the model’s detection sensitivity concerning the sizes of different categories, we utilize Pearson Correlation Coefficient (PCC) [5] to quantify the linear correlation between the average bounding box area per category and the average detection score per category of DOTA-v1.0 [64].

First, we calculate the average area of all the annotations for the k -th category, denoted as S_k . The average area for all categories \bar{S} is calculated as $\bar{S} = \frac{1}{K} \sum_{k=1}^K S_k$, where K is the number of categories. The mean scores for each category Q_k and for all categories \bar{Q} are computed in a similar manner. Second, we calculate the covariance between the category-wise average areas $\{S_k\}_{k=1}^K$ and the category-wise average scores $\{Q_k\}_{k=1}^K$ as $D = \frac{1}{(K-1)} \sum_{k=1}^K (S_k - \bar{S}) \times (Q_k - \bar{Q})$. Finally, PCC is computed as:

$$r = D / \sigma_S \sigma_Q. \quad (9)$$

Here, σ_S and σ_Q are the standard deviations of the category-wise average areas $\{S_k\}_{k=1}^K$ and the category-wise average scores $\{Q_k\}_{k=1}^K$, respectively. A PCC absolute value $|r|$ close to 0 suggests a minimal linear correlation, indicating that the model’s detection performance is rarely influ-

Methods	mAP ↑	$ r $ ↓
S ² ANet [20]	74.13	0.23
O-RCNN [65]	75.87	0.22
ARC [26]	77.35	0.29
LSKNet [32]	77.49	0.24
PKINet-S (ours)	78.39	0.19

Table 9. **Comparison of mAP and PCC (r) on DOTA-v1.0** dataset [64]. See §4.5 for details.

enced by the size of the object. As illustrated in Table 9, our PKINet achieves both the highest mAP and the lowest PCC absolute value $|r|$, indicating that PKINet is the least sensitive to size variations across different categories.

5. Discussion and Conclusion

In this paper, we propose Poly Kernel Inception Network (PKINet) for remote sensing object detection, which aims at tackling the challenges posed by considerable variations in object scale and contextual diversity in remote sensing images. PKINet employs parallel depth-wise convolution kernels of various sizes to capture dense texture features effectively across different scales. A Context Anchor Attention mechanism is also introduced to capture long-range contextual information further. We experimentally show that PKINet achieves state-of-the-art performance on four famous remote sensing benchmark datasets.

Limitations and Future Work. While both PKINet-T and PKINet-S have demonstrated superior detection performance over previous methods, limitations in our computational resources have restricted PKINet from scaling up the model capacity to achieve its maximal potential. Similar studies on model scalability have received substantial interest in general object detection, as highlighted in Swin Transformer [42] and ConvNeXt [43]. We leave further investigation into the scalability of PKINet for future research.

Acknowledgement. This work was supported by the National Natural Science Foundation of China (No. 62102182, 62372405, 62202227, 62302217, 62306292), CCF-Tencent Open Fund.

References

- [1] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *CVPR*, pages 4974–4983, 2019. 6
- [2] Yukang Chen, Peizhen Zhang, Zeming Li, Yanwei Li, Xiangyu Zhang, Gaofeng Meng, Shiming Xiang, Jian Sun, and Jiaya Jia. Stitcher: Feedback-driven data provider for object detection. *arXiv preprint arXiv:2004.12432*, 2(7):12, 2020. 1, 2
- [3] Gong Cheng, Jiabao Wang, Ke Li, Xingxing Xie, Chunbo Lang, Yanqing Yao, and Junwei Han. Anchor-free oriented proposal generator for object detection. *IEEE TGRS*, 60:1–11, 2022. 2, 5, 6, 7
- [4] Gong Cheng, Yanqing Yao, Shengyang Li, Ke Li, Xingxing Xie, Jiabao Wang, Xiwen Yao, and Junwei Han. Dual-aligned oriented detector. *IEEE TGRS*, 60:1–11, 2022. 2
- [5] Israel Cohen, Yiteng Huang, Jingdong Chen, Jacob Benesty, Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. Pearson correlation coefficient. *Noise Reduction in Speech Applications*, pages 1–4, 2009. 8
- [6] MMPreTrain Contributors. Openmmlab’s pre-training toolbox and benchmark. <https://github.com/open-mmlab/mmpretrain>, 2023. 5
- [7] Linhui Dai, Hong Liu, Hao Tang, Zhiwei Wu, and Pinhao Song. Ao2-detr: Arbitrary-oriented object detection transformer. *IEEE TCSVT*, 2022. 5
- [8] Yingpeng Dai, Chenglin Li, Xiaohang Su, Hongxian Liu, and Jiehao Li. Multi-scale depthwise separable convolution for semantic segmentation in street–road scenes. *Remote Sensing*, 15(10):2649, 2023. 2
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 5, 6, 7, 8
- [10] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *CVPR*, pages 2849–2858, 2019. 1, 2, 5, 6
- [11] 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 TPAMI*, 44(11):7778–7796, 2021. 1, 2
- [12] Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *CVPR*, pages 11963–11975, 2022. 2
- [13] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>. 6
- [14] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 6
- [15] Kun Fu, Zhonghan Chang, Yue Zhang, and Xian Sun. Point-based estimator for arbitrary-oriented object detection in aerial images. *IEEE TGRS*, 59(5):4370–4387, 2020. 2
- [16] Haoyuan Guo, Xi Yang, Nannan Wang, Bin Song, and Xinbo Gao. A rotational libra r-cnn method for ship detection. *IEEE TGRS*, 58(8):5772–5781, 2020. 2
- [17] Meng-Hao Guo, Cheng-Ze Lu, Qibin Hou, Zhengning Liu, Ming-Ming Cheng, and Shi-Min Hu. Segnext: Rethinking convolutional attention design for semantic segmentation. *NeurIPS*, 35:1140–1156, 2022. 2
- [18] Meng-Hao Guo, Cheng-Ze Lu, Zheng-Ning Liu, Ming-Ming Cheng, and Shi-Min Hu. Visual attention network. *CVM*, 9(4):733–752, 2023. 2
- [19] Wei Guo, Wen Yang, Haijian Zhang, and Guang Hua. Geospatial object detection in high resolution satellite images based on multi-scale convolutional neural network. *Remote Sensing*, 10(1):131, 2018. 2
- [20] Jiaming Han, Jian Ding, Jie Li, and Gui-Song Xia. Align deep features for oriented object detection. *IEEE TGRS*, 60:1–11, 2021. 1, 5, 6, 8
- [21] Jiaming Han, Jian Ding, Nan Xue, and Gui-Song Xia. Redet: A rotation-equivariant detector for aerial object detection. In *CVPR*, pages 2786–2795, 2021. 2, 5, 6
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 3, 5, 6, 7
- [23] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017. 6, 7
- [24] Jie-Bo Hou, Xiaobin Zhu, and Xu-Cheng Yin. Self-adaptive aspect ratio anchor for oriented object detection in remote sensing images. *Remote Sensing*, 13(7):1318, 2021. 2
- [25] Liping Hou, Ke Lu, and Jian Xue. Refined one-stage oriented object detection method for remote sensing images. *IEEE TIP*, 31:1545–1558, 2022. 2
- [26] Liping Hou, Ke Lu, Jian Xue, and Yuqiu Li. Shape-adaptive selection and measurement for oriented object detection. In *AAAI*, pages 923–932, 2022. 2, 5, 8
- [27] Liping Hou, Ke Lu, Xue Yang, Yuqiu Li, and Jian Xue. G-rep: Gaussian representation for arbitrary-oriented object detection. *Remote Sensing*, 15(3):757, 2023. 2
- [28] Zhanchao Huang, Wei Li, Xiang-Gen Xia, and Ran Tao. A general gaussian heatmap label assignment for arbitrary-oriented object detection. *IEEE TIP*, 31:1895–1910, 2022. 6
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5, 6
- [30] Ke Li, Gang Wan, Gong Cheng, Liqiu Meng, and Junwei Han. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS J. P&RS*, 159:296–307, 2020. 5, 6
- [31] Wentong Li, Yijie Chen, Kaixuan Hu, and Jianke Zhu. Oriented reppoints for aerial object detection. In *CVPR*, pages 1829–1838, 2022. 1, 2, 5, 6
- [32] Yuxuan Li, Qibin Hou, Zhaohui Zheng, Ming-Ming Cheng, Jian Yang, and Xiang Li. Large selective kernel network

- for remote sensing object detection. In *ICCV*, pages 16794–16805, 2023. 2, 4, 5, 6, 7, 8
- [33] Xi Liang, Jing Zhang, Li Zhuo, Yuzhao Li, and Qi Tian. Small object detection in unmanned aerial vehicle images using feature fusion and scaling-based single shot detector with spatial context analysis. *IEEE TCSVT*, 30(6):1758–1770, 2019. 1
- [34] 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*, pages 740–755, 2014. 7
- [35] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 2
- [36] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 6
- [37] Zhao Lin, Kefeng Ji, Xiangguang Leng, and Gangyao Kuang. Squeeze and excitation rank faster r-cnn for ship detection in sar images. *IEEE GRSL*, 16(5):751–755, 2018. 1
- [38] Shiwei Liu, Tianlong Chen, Xiaohan Chen, Xuxi Chen, Qiao Xiao, Boqian Wu, Mykola Pechenizkiy, Decebal Mocanu, and Zhangyang Wang. More convnets in the 2020s: Scaling up kernels beyond 51x51 using sparsity. *arXiv preprint arXiv:2207.03620*, 2022. 2
- [39] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, pages 21–37, 2016. 2
- [40] Zikun Liu, Hongzhen Wang, Lubin Weng, and Yiping Yang. Ship rotated bounding box space for ship extraction from high-resolution optical satellite images with complex backgrounds. *IEEE GRSL*, 13(8):1074–1078, 2016. 2
- [41] Zikun Liu, Liu Yuan, Lubin Weng, and Yiping Yang. A high resolution optical satellite image dataset for ship recognition and some new baselines. In *ICPRAM*, pages 324–331, 2017. 2, 5, 6, 7
- [42] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 8
- [43] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, pages 11976–11986, 2022. 2, 7, 8
- [44] Yang Long, Gui-Song Xia, Shengyang Li, Wen Yang, Michael Ying Yang, Xiao Xiang Zhu, Liangpei Zhang, and Deren Li. On creating benchmark dataset for aerial image interpretation: Reviews, guidances, and million-aid. *IEEE J-STARS*, 14:4205–4230, 2021. 5, 6
- [45] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 5
- [46] Chengqi Lyu, Wenwei Zhang, Haiyan Huang, Yue Zhou, Yudong Wang, Yanyi Liu, Shilong Zhang, and Kai Chen. Rtmnet: An empirical study of designing real-time object detectors. *arXiv preprint arXiv:2212.07784*, 2022. 6
- [47] Teli Ma, Mingyuan Mao, Honghui Zheng, Peng Gao, Xiaodi Wang, Shumin Han, Errui Ding, Baochang Zhang, and David Doermann. Oriented object detection with transformer. *arXiv preprint arXiv:2106.03146*, 2021. 5
- [48] Qi Ming, Zhiqiang Zhou, Lingjuan Miao, Hongwei Zhang, and Linhao Li. Dynamic anchor learning for arbitrary-oriented object detection. In *AAAI*, pages 2355–2363, 2021. 2
- [49] Xingjia Pan, Yuqiang Ren, Kekai Sheng, Weiming Dong, Haolei Yuan, Xiaowei Guo, Chongyang Ma, and Changsheng Xu. Dynamic refinement network for oriented and densely packed object detection. In *CVPR*, pages 11207–11216, 2020. 6
- [50] Yifan Pu, Yiru Wang, Zhuofan Xia, Yizeng Han, Yulin Wang, Weihao Gan, Zidong Wang, Shiji Song, and Gao Huang. Adaptive rotated convolution for rotated object detection. In *ICCV*, pages 6589–6600, 2023. 2, 5, 6
- [51] Wen Qian, Xue Yang, Silong Peng, Junchi Yan, and Yue Guo. Learning modulated loss for rotated object detection. In *AAAI*, pages 2458–2466, 2021. 2
- [52] Heqian Qiu, Hongliang Li, Qingbo Wu, Fanman Meng, King Ngi Ngan, and Hengcan Shi. A2rmnet: Adaptively aspect ratio multi-scale network for object detection in remote sensing images. *Remote Sensing*, 11(13):1594, 2019. 2
- [53] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 28, 2015. 5, 6
- [54] Pourya Shamsolmoali, Masoumeh Zareapoor, Jocelyn Chanussot, Huiyu Zhou, and Jie Yang. Rotation equivariant feature image pyramid network for object detection in optical remote sensing imagery. *IEEE TGRS*, 60:1–14, 2021. 1, 2
- [55] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [56] 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 J. P&RS*, 184:116–130, 2022. 1, 2
- [57] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. 3
- [58] Yehui Tang, Kai Han, Jianyuan Guo, Chang Xu, Chao Xu, and Yunhe Wang. Ghostnetv2: enhance cheap operation with long-range attention. *NeurIPS*, 35:9969–9982, 2022. 4
- [59] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. 1, 5, 6
- [60] Chien-Yao Wang, Hong-Yuan Mark Liao, Yueh-Hua Wu, Ping-Yang Chen, Jun-Wei Hsieh, and I-Hau Yeh. Cspnet: A new backbone that can enhance learning capability of cnn. In *CVPRW*, pages 390–391, 2020. 3
- [61] Guanqun Wang, Yin Zhuang, He Chen, Xiang Liu, Tong Zhang, Lianlin Li, Shan Dong, and Qianbo Sang. Fsod-net: Full-scale object detection from optical remote sensing imagery. *IEEE TGRS*, 60:1–18, 2022. 1, 2

- [62] Jinwang Wang, Jian Ding, Haowen Guo, Wensheng Cheng, Ting Pan, and Wen Yang. Mask obb: A semantic attention-based mask oriented bounding box representation for multi-category object detection in aerial images. *Remote Sensing*, 11(24):2930, 2019. [2](#)
- [63] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, pages 568–578, 2021. [7](#)
- [64] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. DOTA: A large-scale dataset for object detection in aerial images. In *CVPR*, pages 3974–3983, 2018. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [65] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *ICCV*, pages 3520–3529, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [66] Chang Xu, Jian Ding, Jinwang Wang, Wen Yang, Huai Yu, Lei Yu, and Gui-Song Xia. Dynamic coarse-to-fine learning for oriented tiny object detection. In *CVPR*, pages 7318–7328, 2023. [6](#)
- [67] Yongchao Xu, Mingtao Fu, Qimeng Wang, Yukang Wang, Kai Chen, Gui-Song Xia, and Xiang Bai. Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE TPAMI*, 43(4):1452–1459, 2020. [1](#), [2](#), [5](#), [6](#)
- [68] Xue Yang and Junchi Yan. Arbitrary-oriented object detection with circular smooth label. In *ECCV*, pages 677–694, 2020. [1](#)
- [69] Xue Yang, Jirui Yang, Junchi Yan, Yue Zhang, Tengfei Zhang, Zhi Guo, Xian Sun, and Kun Fu. ScrDet: Towards more robust detection for small, cluttered and rotated objects. In *ICCV*, pages 8232–8241, 2019. [2](#), [5](#)
- [70] Xue Yang, Liping Hou, Yue Zhou, Wentao Wang, and Junchi Yan. Dense label encoding for boundary discontinuity free rotation detection. In *CVPR*, pages 15819–15829, 2021. [1](#), [2](#)
- [71] Xue Yang, Junchi Yan, Ziming Feng, and Tao He. R3det: Refined single-stage detector with feature refinement for rotating object. In *AAAI*, pages 3163–3171, 2021. [1](#), [2](#), [5](#), [6](#)
- [72] Xue Yang, Junchi Yan, Qi Ming, Wentao Wang, Xiaopeng Zhang, and Qi Tian. Rethinking rotated object detection with gaussian wasserstein distance loss. In *ICML*, pages 11830–11841, 2021. [1](#), [2](#), [5](#), [6](#)
- [73] 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*, 34:18381–18394, 2021. [2](#), [5](#)
- [74] Xue Yang, Yue Zhou, Gefan Zhang, Jirui Yang, Wentao Wang, Junchi Yan, Xiaopeng Zhang, and Qi Tian. The kfiou loss for rotated object detection. *arXiv preprint arXiv:2201.12558*, 2022. [1](#)
- [75] Yazhou Yao, Tao Chen, Hanbo Bi, Xinhao Cai, Gensheng Pei, Guoye Yang, Zhiyuan Yan, Xian Sun, Xing Xu, and Hai Zhang. Automated object recognition in high-resolution optical remote sensing imagery. *NSR*, 10(6):nwad122, 2023. [2](#)
- [76] Jingru Yi, Pengxiang Wu, Bo Liu, Qiaoying Huang, Hui Qu, and Dimitris Metaxas. Oriented object detection in aerial images with box boundary-aware vectors. pages 2150–2159, 2021. [2](#)
- [77] Weihao Yu, Pan Zhou, Shuicheng Yan, and Xinchao Wang. Inceptionnext: When inception meets convnext. *arXiv preprint arXiv:2303.16900*, 2023. [3](#)
- [78] Ying Zeng, Xue Yang, Qingyun Li, Yushi Chen, and Junchi Yan. Ars-detr: Aspect ratio sensitive oriented object detection with transformer. *arXiv preprint arXiv:2303.04989*, 2023. [5](#), [6](#)
- [79] Ke Zhang, Inuwa Mamuda Bello, Yu Su, Jingyu Wang, and Ibrahim Maryam. Multiscale depthwise separable convolution based network for high-resolution image segmentation. *IJRS*, 43(18):6624–6643, 2022. [2](#)
- [80] Wenhua Zhang, Licheng Jiao, Yuxuan Li, Zhongjian Huang, and Haoran Wang. Laplacian feature pyramid network for object detection in vhr optical remote sensing images. *IEEE TGRS*, 60:1–14, 2021. [2](#)
- [81] Yuanlin Zhang, Yuan Yuan, Yachuang Feng, and Xiaoqiang Lu. Hierarchical and robust convolutional neural network for very high-resolution remote sensing object detection. *IEEE TGRS*, 57(8):5535–5548, 2019. [1](#), [2](#)
- [82] Wei Zhao, Wenping Ma, Licheng Jiao, Puhua Chen, Shuyuan Yang, and Biao Hou. Multi-scale image block-level f-cnn for remote sensing images object detection. *IEEE Access*, 7:43607–43621, 2019. [1](#), [2](#)
- [83] Zhuo Zheng, Yanfei Zhong, Ailong Ma, Xiaobing Han, Ji Zhao, Yanfei Liu, and Liangpei Zhang. Hynet: Hyper-scale object detection network framework for multiple spatial resolution remote sensing imagery. *ISPRS J. P&RS*, 166:1–14, 2020. [2](#)
- [84] Yue Zhou, Xue Yang, Gefan Zhang, Jiabao Wang, Yanyi Liu, Liping Hou, Xue Jiang, Xingzhao Liu, Junchi Yan, Chengqi Lyu, Wenwei Zhang, and Kai Chen. Mmrotate: A rotated object detection benchmark using pytorch. In *ACM MM*, 2022. [5](#)