

TricubeNet: 2D Kernel-Based Object Representation for Weakly-Occluded Oriented Object Detection

Beomyoung Kim^{1†} Janghyeon Lee^{2†} Sihaeng Lee^{2†} Doyeon Kim³ Junmo Kim³
NAVER CLOVA¹ LG AI Research² KAIST³

Abstract

We present a novel approach for oriented object detection, named *TricubeNet*, which localizes oriented objects using visual cues (*i.e.*, heatmap) instead of oriented box offsets regression. We represent each object as a 2D Tricube kernel and extract bounding boxes using simple image-processing algorithms. Our approach is able to (1) obtain well-arranged boxes from visual cues, (2) solve the angle discontinuity problem, and (3) can save computational complexity due to our anchor-free modeling. To further boost the performance, we propose some effective techniques for size-invariant loss, reducing false detections, extracting rotation-invariant features, and heatmap refinement. To demonstrate the effectiveness of our *TricubeNet*, we experiment on various tasks for weakly-occluded oriented object detection: detection in an aerial image, densely packed object image, and text image. The extensive experimental results show that our *TricubeNet* is quite effective for oriented object detection. Code is available at <https://github.com/qjadudl1994/TricubeNet>.

1. Introduction

Object detection is one of the fundamental computer vision tasks, and deep learning-based methods [32, 4, 36] have shown remarkable performance. However, existing detectors often focus on detecting a horizontal bounding box that is not sufficient in some cases. First, for densely arranged oriented objects, an intersection-over-union (IOU) between adjacent horizontal bounding boxes tends to be large, and some of these boxes will be filtered out by non-maximum suppression (NMS). Second, since the horizontal bounding box can contain many redundant areas, it is not suitable for real-world applications that require tighter and more accurate boxes, such as aerial images and scene text images. To detect the object in a more accurate form, oriented object detection has attracted much attention recently.

[†]This work was done when the author worked at KAIST.

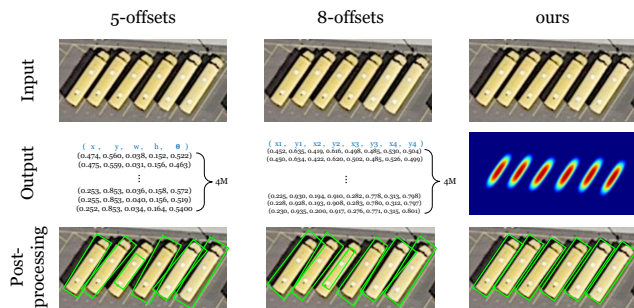


Figure 1. Existing anchor-based methods that regress five or eight offsets require a huge number of anchor boxes (about 4M) and infer a bit scatty boxes. We represent each object as a 2D Tricube kernel and extract bounding boxes using simple image-processing algorithms. Our approach, named *TricubeNet*, does not require the anchor boxes and can obtain well-arranged boxes.

Most oriented object detectors [26, 7, 40] adopt Faster R-CNN [32] or RetinaNet [21] as their baseline model and additionally infer an angle of the object. They adopt anchor augmentation strategy and regress oriented box offsets in form of 5-offsets (x, y, w, h, θ) or 8-offsets ($x_1, y_1, x_2, y_2, x_3, y_3, x_4, y_4$). They have reigned on the throne with state-of-the-art performance, however, some limitations remain. (1) Regressing the box offsets might have trouble in obtaining well-arranged boxes of densely arranged oriented objects as in Figure 1; (2) Regressing the angle offset causes angle discontinuity problem; the angle discontinuity on the boundary leads to the loss fluctuation during training; (3) They require huge computational complexity due to the anchor augmentation and a heavy IoU calculation for the oriented box. For example, when they define anchor boxes with three scales, five aspect ratios, and six angles and adopt FPN [20] architecture with 800×800 input resolution, they require about 4M total anchor boxes ($3 \times 5 \times 6 = 90$ anchor boxes per a pixel location). Some might argue that anchor-free approaches such as [48, 28] can reduce the computational complexity, however, they also suffer from the angle discontinuity problem due to their angle regression.

In this paper, we introduce a novel approach for the oriented object detection, named **TricubeNet**. We localize

oriented objects using *visual cues* (*i.e.*, heatmap) instead of the box offset regression. As shown in Figure 1, we represent each object as a 2D Tricube kernel whose shape visually describes the width, height, and angle of the object, and then extract bounding boxes using simple image-processing algorithms. Our approach is able to (1) obtain well-arranged oriented boxes from visual cues of arranged objects as shown in Figure 1, (2) solve the angle discontinuity problem by taking away the angle regression, (3) save computational complexity due to our anchor-free modeling, and (4) is a simple one-stage anchor-free detector.

Furthermore, for the competitive result, we should handle some challenging factors of the oriented object detection: various shapes and sizes of objects, densely arranged objects, a huge number of objects, false detections, and complexity of the background. To deal with these challenging factors, we propose some techniques. The first is a *Size-Weight Mask* (SWM). The pixel-wise mean squared error (MSE) loss causes a size-imbalance problem, that is, a small object tends to be given small loss, weakening the detection of small objects. To make a size-invariant loss function, we propose the SWM. Second, to give a balanced loss between foreground and background pixels and reduce false-positive detections at once, we introduce a *False-Positive Example Mining* (FPEM) technique. Third, we propose a *Multi-Angle Convolution* (MAC) module to extract a rotation-invariant feature for the oriented object. Last, we design a repetitive refinement stage to refine the output heatmap and call this technique *heatmap cascade refinement*.

We verify the effectiveness of the TricubeNet in various tasks: oriented object detection in the aerial image (DOTA [39]), the densely packed object image (SKU110K-R [10, 28]), and the scene text image (MSRA-TD500 [45], ICDAR 2015 [16]). We target detection of weakly-occluded oriented objects and choose the above highly practical tasks in real-world applications. The experimental results show that TricubeNet is quite effective to detect the oriented object with a simple anchor-free one-stage process.

In summary, our contributions are as follows:

- We propose a novel oriented object detector, TricubeNet, which localizes oriented objects using visual cues (*i.e.*, heatmap) instead of the box offset regression.
- Our approach can obtain well-arranged oriented boxes, solve the angle discontinuity problem, and save computational complexity by eliminating anchor boxes.
- We propose some techniques (*i.e.*, SWM, FPEM, MAC, and heatmap cascade refinement) to properly detect the oriented object and boost the performance.
- We verify the effectiveness of our TricubeNet from extensive experimental results on various tasks.

2. Related Work

Two-stage object detectors consist of two processes: extract object region candidates and crop the region of interest (ROI) of each object and predict the class and bounding box offsets of the object. R-CNN [9] uses a selective search [37] method to generate the bounding box candidates and feeds them to the classifier. SPP [12] and Fast R-CNN [8] crop the ROIs from feature maps and feed them to the classifier. Faster R-CNN [32] generates bounding box candidates from a region proposal network (RPN), which allows training in an end-to-end manner. Recently, Cascade R-CNN [4] designs iterative refinement steps for the high-quality bounding box prediction. Although the two-stage object detectors achieve state-of-the-art performance, they require a high computational complexity.

One-stage object detectors classify and regress the bounding boxes at once using anchor boxes. SSD [23] densely produces the bounding boxes from multi-level feature maps using various sizes of anchor boxes and removes overlapping bounding boxes using NMS post-processing. RetinaNet [21] uses a focal loss to alleviate the class-imbalance problem between positive and negative anchor boxes.

Anchor-free object detectors have recently been proposed and eliminate the anchor box in the network design. CornerNet [18] is a keypoint-based anchor-free approach that represents objects as pairs of corner keypoints and groups them. CenterNet [48] represents objects as center points with width and height regression. The anchor-free detectors are simple and efficient but report a lower performance than the two-stage detectors.

Oriented object detectors often adopt the object detectors and additionally regress the angle for oriented objects. Adopting the Faster-RCNN [32] as a baseline detector, RRPNet [26] exploits rotated anchor boxes and changes the IoU calculation for rotated boxes, RoI Transformer [7] extracts geometry-robust pooled features, and Gliding vertex [40] employs a simple object representation method which glides the vertex of the horizontal bounding box on each corresponding side. Adopting the RetinaNet [21] as a baseline detector, RSDet [29] proposes an eight-parameter regression model using a rotation sensitivity error (RSE) for handling the angle discontinuity problem and R³Det [41] is a fast one-stage detector with a feature refinement module to handle the feature misalignment problem. SCRDet [44] proposes an IoU loss to alleviate the angle discontinuity problem. These approaches adopt an anchor augmentation strategy with a huge amount of box candidates to obtain oriented boxes. To reduce computational complexity by eliminating anchor box, DRN [28] adopts CenterNet [48] as their baseline detector taking anchor-free modeling and additionally regresses the angle offset with a proposed feature selection module and dynamic refinement head to dynamically refine the network. Note that DRN represents objects as




Figure 2. Overview of TricubeNet. It produces one channel heatmap per category where each oriented object is represented as a 2D Tricube kernel. The backbone network consists of a fully convolutional encoder–decoder architecture. H and W are the height and width of the image, respectively; C is the number of categories; R is the downsampling rate. From the heatmap cascade refinement, we progressively refine the output heatmap (H) and extract bounding boxes from the lastly refined heatmap (H_3) using simple image-processing algorithms.




Figure 3. (a): 2D Tricube kernel, (b): 2D Gaussian kernel, (c): object as 2D Tricube kernel, and (d): object as 2D Gaussian kernel.

center points and obtains other box offsets using regression, suffering from the angle discontinuity problem. Meanwhile, we represent the whole object region as a 2D Tricube kernel and obtain a bounding box using image-processing algorithms, detouring the angle discontinuity problem.

The text detectors widely have adopted a segmentation-based anchor-free approach by utilizing the fact that there is no occlusion in the text. PixelLink [6] performs the instance segmentation using a pixel-wise eight-neighbors link prediction. CRAFT [2] is a character-level detection approach instead of word-level detection by representing each character as a 2D Gaussian kernel.

3. Method

3.1. 2D Tricube kernel

Recent key-point detection approaches [18, 48, 27] represent each key-point as a 2D Gaussian kernel and show remarkable performances. Motivated by these approaches, we assume that a 2D kernel function is a good candidate for representing the key features and is an easy-to-learn form for deep neural networks. For the oriented object detection, however, the 2D Gaussian kernel is not a suitable choice because its circular form often fails to represent the angle of an object. For example, in Figure 3(d), the 2D Gaussian kernel for the rotated square-form object cannot represent the angle of the object. For suitable modeling for the oriented object detection, we choose a 2D Tricube kernel as our object representation method. Unlike the 2D Gaussian kernel that has a circular-form distribution, the 2D Tricube




Figure 4. Illustration of multi-angle convolution (MAC) module which extracts rotation-invariant features.

kernel has a rectangular-form distribution as in Figure 3(a) and is defined as $(1 - |x|^3)^\gamma \cdot (1 - |y|^3)^\gamma$ where γ is set to 7 in all experiments. The 2D Tricube kernel can properly represent the angle of the object as shown in Figure 3(c).

3.2. Architecture

As illustrated in Figure 2, TricubeNet consists of a fully convolutional encoder–decoder architecture. Let $I \in \mathbb{R}^{H \times W \times 3}$ be an input image, where H is height and W is width. The network predicts the output heatmap $H \in \mathbb{R}^{\frac{H}{R} \times \frac{W}{R} \times C}$, where R is the downsampling rate and C is the number of categories. We adopt Hourglass-104 [27] network as our backbone network. We apply a bilinear upsampling layer on the output heatmap for higher pixel precision; the downsampling rate R is set to 2 in all experiments.

Multi-Angle Convolution (MAC) Module. In our heatmap-based detection framework, extracting rotation-invariant features is important to detect orient objects. However, a convolutional neural network (CNN) has a limitation in extracting a rotation-invariant feature. To this end, we design a multi-angle convolution (MAC) module that is light and effective to extract the rotation-invariant features. The MAC module is illustrated in Figure 4.

First, we define n kinds of angles, θ_i , $i \in (1, \dots, n)$, where the range is $[0, \frac{\pi}{2})$. Next, when the input feature map X consists of K channels, we divide the channel of the X by n by applying a convolutional layer with a kernel size of 1 and channel of $\frac{K}{n}$. We denote the divided feature map as

Algorithm 1: Rotation Convolution

Input: Feature maps $X \in \mathbb{R}^{H \times W \times K}$,
rotation angle θ

Output: Rotation-invariant feature maps :
 $\hat{X} \in \mathbb{R}^{H \times W \times K}$

$\hat{X} \leftarrow \text{Rescale}(X, 1/(\sin \theta + \cos \theta))$

$\hat{X} \leftarrow \text{Rotate}(\hat{X}, \theta)$

$\hat{X} \leftarrow \text{Conv}3 \times 3(\hat{X})$

$\hat{X} \leftarrow \text{Rotate}(\hat{X}, 2\pi - \theta)$

$\hat{X} \leftarrow \text{Rescale}(\hat{X}, \sin \theta + \cos \theta)$

return \hat{X}

\mathcal{F}_i . Then, we apply a rotation convolutional layer with an angle of θ_i and kernel size 3 on \mathcal{F}_i . Last, we concatenate the outputs of the rotation convolution with n different angles applied. Formally, the output of the MAC module \hat{X} , which is the rotation invariant feature map, is defined as:

$$\hat{X} = \text{concat}([\text{RCConv}(\mathcal{F}_i, \theta_i), \dots, \text{RCConv}(\mathcal{F}_n, \theta_n)]), \quad (1)$$

where *concat* is the concatenate operation and *RCConv* is the rotation convolutional layer with an angle of θ . In all experiments, we set $n = 4$ and $(\theta_1, \theta_2, \theta_3, \theta_4) = (0, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3})$.

A recent study, DRN [28], proposed a rotation convolutional layer, which rotates the kernel of CNN to extract the rotation-invariant feature. However, rotating the kernel of CNN requires heavy computations and additional parameters. For the efficient implementation of the rotation convolutional layer, we approximate the rotating of the kernel by rotating feature maps and then applying a normal CNN.

The detailed process of our rotation convolutional layer is described in Algorithm 1. First, we re-scale the feature map with a scale factor of $1/(\sin \theta + \cos \theta)$ to keep the original shape after rotating and then rotate the feature map by θ . Then, we apply a normal convolutional layer with kernel size 3 on the rotated feature map and return the output feature map back to the original angle and scale.

Heatmap Cascade Refinement. The two-stage detectors, such as Cascade R-CNN [4], take advantage of a progressive cascade refinement of the predicted bounding box; it repeatedly crops the region of the predicted bounding box and predicts refined bounding box. In contrast, the cascade refinement step is less studied on the heatmap-based anchor-free approaches. Here, in our study, we propose a heatmap cascade refinement, which progressively refines the output heatmap in a pixel-wise manner, as illustrated in Figure 2. Specifically, when the last feature map our backbone network is denoted as X , the r -th refined feature map Y_r and r -th refined heatmap H_r are defined as:




Figure 5. Illustration of the ground truth (GT) heatmap and size-weight mask generation procedure.

$$\begin{aligned} Y_1 &= \text{MAC}(X), H_1 = f_c(Y_1), \\ Y_2 &= \text{MAC}(X + Y_1), H_2 = f_c(Y_2), \\ &\vdots \\ Y_r &= \text{MAC}(X + Y_{r-1}), H_r = f_c(Y_r), \end{aligned} \quad (2)$$

where *MAC* is the multi-angle convolution module and *f_c* is a convolutional layer with a kernel size of 1 and a channel of C . Through the heatmap cascade refinement, we progressively refine the output heatmap by taking the rotation-invariant feature map of the previous refinement step.

3.3. From bounding boxes to heatmap

Here, we describe how to generate a ground truth heatmap. The whole generation process is illustrated in Figure 5. First, we create a square 2D Tricube kernel and normalize it to a value between zero to one. Then, we reshape the kernel to the same size as the ground truth oriented bounding box. Last, we insert this reshaped kernel into the heatmap. If two Tricube kernels overlap, we take an element-wise maximum operation. The generated ground truth heatmap is denoted as $\hat{H} \in \mathbb{R}^{\frac{H}{R} \times \frac{W}{R} \times C}$.

3.4. Objective

We optimize the TricubeNet using pixel-wise mean-squared error (MSE) objective function between H and \hat{H} . However, giving the pixel-wise MSE loss function has two problems. First is a size-imbalance problem. Since we assign the 2D Tricube kernel according to the size of each object, large objects tend to have large losses and small objects tend to have little losses; this weakens the detection of small objects. Second is a class-imbalance problem. Most of the pixels in the heatmap are background pixels, therefore, it interferes with focusing on the foreground pixels. To alleviate the above problems, we propose a size-weight mask (SWM) and false-positive example mining (FP EM). **Size-Weight Mask (SWM).** For the size-invariant loss function, we weight the MSE loss according to the size of the object using the size-weight mask (SWM); a large object is given a low weight and a small object is given a high weight. The SWM is denoted as $M \in \mathbb{R}^{\frac{H}{R} \times \frac{W}{R} \times C}$. For each object, foreground pixels in the SWM contain a weight that is inversely proportional to the size of the ob-

ject. Also, background pixels contain a weight of 1. This mask is generated in the same manner as the ground truth heatmap (Section 3.3) and illustrated in Figure 5. When we denote the size of the i -th object as S_i , the set of pixels of i -th object as \mathcal{P}_i^{obj} , and the number of objects in an image as N , the SWM for i -th object, M_i , is defined as:

$$S = \sum_{i=1}^N S_i, \quad M_i(p) = \begin{cases} \frac{S}{S_i \times N} & p \in \mathcal{P}_i^{obj} \\ 1 & \text{otherwise.} \end{cases} \quad (3)$$

False-Positive Example Mining (FPEM). To alleviate the class-imbalance problem, we adopt Online Hard Example Mining (OHEM) [34] that is an effective sampling method to make more focus on the foreground objects. However, focusing too much on the foreground objects yields a lot of false-positive detections. For balanced sampling while reducing false-positives, we propose false-positive exemplar mining (FPEM). Specifically, we extract the positive and false-positive pixels during training in an online manner:

$$\begin{aligned} \mathcal{P}^{pos} &= \{p \mid \hat{H}(p) > 0\}, \\ \mathcal{P}^{fp} &= \{p \mid \hat{H}(p) = 0 \text{ and } H(p) > 0\}, \end{aligned} \quad (4)$$

where \mathcal{P}^{pos} and \mathcal{P}^{fp} is the set of positive and false-positive pixels, respectively. Then we sample the ratio of $|\mathcal{P}^{pos}|$ and $|\mathcal{P}^{fp}|$ to be 1:3. Here, the order of \mathcal{P}^{fp} is sorted according to the distance between \hat{H} and H and then sampled in the top-k manner.

The final objective function L is defined as,

$$L = \frac{1}{S} \sum_{p \in \mathcal{P}} M(p) \times (H(p) - \hat{H}(p))^2, \quad (5)$$

where \mathcal{P} is the set of sampled pixels from FPEM.

3.5. From heatmap to oriented bounding boxes

We apply simple post-processing algorithms to extract the oriented bounding boxes from the heatmap, as illustrated in Figure 6. First, we obtain a binary heatmap $S \in [0, 1]^{\frac{H}{R} \times \frac{W}{R} \times C}$ from the output heatmap H with a threshold τ . By increasing the τ , we can separate the weakly-occluded objects. Second, through the connected component labeling (CCL) [13] algorithm, we allocate an ID of each Tricube kernel. Third, we obtain the rotated rectangular box with a minimum area containing each Tricube kernel using OpenCV [3] functions. Last, since we drop some part of the Tricube kernel through the threshold τ , we scale up the kernel back to its original size by multiplying a scale factor s . For the choice of τ and s , we adaptively set the s according to the τ utilizing the distribution of the Tricube kernel, *i.e.*, $y = (1 - |x|^3)^\gamma$. Note that the distribution of the Tricube kernel is maintained even if the input resolution or the object size is changed, therefore, we can adaptively set $s = \frac{1}{1 - (1 - |\tau|^3)^\gamma}$, where γ is set to 7 as described in 3.1.




Figure 6. From heatmap to oriented bounding boxes. First, (a) from the heatmap, (b) we obtain binary heatmap using threshold τ . (c) We label each kernel using the connected component labeling (CCL) algorithm and (d) extract the contour points and find the minimum-area rectangle for each kernel.

4. Experiments

4.1. Dataset

To demonstrate our effectiveness, we experiment on various practical datasets of weakly-occluded objects: aerial image, densely packed object image, and scene text image.

DOTA [39] is the largest dataset for oriented object detection in aerial images; it consists of 2806 images with sizes ranging from 800×800 to 4000×4000 . The number of training, validation, and testing images are 1411, 458, and 937, respectively. It contains 15 categories of objects and 188,282 instances with a wide variety of scales, orientations, and shapes. Following other conventions [7, 28], we resize the image and crop a series of 1024×1024 patches from the original images with a stride of 824. For training, we resize the images at three scales (0.75, 1.0, and 1.5). For single-scale testing and multi-scale testing, we resize the images at one scale (1.0) and three scales (0.75, 1.0, and 1.5), respectively. The performance on the test set is measured on the official DOTA evaluation server¹.

SKU110K-R [28] is an extended dataset of SKU110K [10] for the densely packed oriented object detection. It contains thousands of supermarket store images with various scales from 1840×1840 to 4320×4320 , viewing angles, lighting conditions, and noise levels. Each image contains an average of 154 tightly packed objects, up to 718 objects. The rotation data augmentation with six angles (-45° , -30° , -15° , 15° , 30° , and 45°) is performed to original images. After the augmentation, the number of training, validation, and testing images are 57,533, 4,116, and 20,587, respectively.

MSRA-TD500 [45] is a dataset for multi-lingual, long, and oriented text detection in both indoors and outdoors natural images. The images contain English and Chinese scripts and each text is labeled by a rotated rectangle. It consists of 300 training images and 200 testing images.

ICDAR2015 [16] is proposed in Robust Reading Competition for incidental scene text detection. There are 1000 training images and 500 testing images. Each text is annotated as word level with a quadrangle of four vertexes.

¹<https://captain-whu.github.io/DOTA>

Table 1. Experiment results on DOTA dataset. MS and Flip denote the multi-scale and flip test time augmentation, respectively. PL-Plane, BD-Baseball Diamond, BR-Bridge, GTF-Ground Field Track, SV-Small Vehicle, LV-Large Vehicle, SH-Ship, TC-Tennis Court, BC-Basketball Court, ST-Storage Tank, SBF-Soccer Ball Field, RA-Roundabout, HA-Harbor, SP-Swimming Pool, and HC-Helicopter.

Method	Backbone	MS	Flip	PL	BD	BR	GTF	SV	LV	SH	TC	BC	ST	SBF	RA	HA	SP	HC	mAP
two-stage																			
R-DFPN [42]	ResNet-101			80.92	65.82	33.77	58.94	55.77	50.94	54.78	90.33	66.34	68.66	48.73	51.76	55.10	51.32	35.88	57.94
RRPN [26]	ResNet-101			80.94	65.75	35.34	67.44	59.92	50.91	55.81	90.67	66.92	72.39	55.06	52.23	55.14	53.35	48.22	61.01
R ² CNN [15]	ResNet-101			88.52	71.20	31.66	59.30	51.85	56.19	57.25	90.81	72.84	67.38	56.69	52.84	53.08	51.94	53.58	60.67
ICN [1]	ResNet-101			81.36	74.30	47.70	70.32	64.89	67.82	69.98	90.76	79.06	78.20	53.64	62.90	67.02	64.17	50.23	68.16
RoI Trans [7]	ResNet-101	✓		88.64	78.52	43.44	75.92	68.81	73.68	83.59	90.74	77.27	81.46	58.39	53.54	62.83	58.93	47.67	69.56
SCRDet [44]	ResNet-101			89.41	78.83	50.02	65.59	69.96	57.63	72.26	90.73	81.41	84.39	52.76	63.62	62.01	67.62	61.16	69.83
SCRDet [44]	ResNet-101	✓		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
Gliding vertex [40]	ResNet-101	✓		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
one-stage																			
FR-O [39]	ResNet-101			79.42	77.13	17.70	64.05	53.30	38.02	37.16	89.41	69.64	59.28	50.30	52.91	47.89	47.40	46.30	54.13
RetinaNet [21]	ResNet-50	✓		88.87	74.46	40.11	58.03	63.10	50.61	63.63	90.89	77.91	76.38	48.26	55.85	50.67	60.23	34.23	62.22
R ³ Det [41]	ResNet-101	✓		89.54	81.99	48.46	62.52	70.48	74.29	77.54	90.80	81.39	83.54	61.97	59.82	65.44	67.46	60.05	71.69
R ³ Det [41]	ResNet-152	✓		89.49	81.17	50.53	66.10	70.92	78.66	78.21	90.81	85.26	84.23	61.81	63.77	68.16	69.83	67.17	73.74
RSDet [29]	ResNet-101	✓		89.80	82.90	48.60	65.20	69.50	70.10	70.20	90.50	85.60	83.40	62.50	63.90	65.60	67.20	68.00	72.20
RSDet [29]	ResNet-152	✓		90.10	82.00	53.80	68.50	70.20	78.70	73.60	91.20	87.10	84.70	64.30	68.20	66.10	69.30	63.70	74.10
anchor-free																			
CenterNet [48]	Hourglass-104			89.02	69.71	37.62	63.42	65.23	63.74	77.28	90.51	79.24	77.93	44.83	54.64	55.93	61.11	45.71	65.04
CenterNet [48]	Hourglass-104	✓		89.56	79.83	43.80	66.54	65.58	66.09	83.11	90.72	83.72	84.30	55.62	58.71	62.48	68.33	50.77	69.95
DRN [28]	Hourglass-104			88.91	80.22	43.52	63.35	73.48	70.69	84.94	90.14	83.85	84.11	50.12	58.41	67.62	68.60	52.50	70.70
DRN [28]	Hourglass-104	✓		89.45	83.16	48.98	62.24	70.63	74.25	83.99	90.73	84.60	85.35	55.76	60.79	71.56	68.82	63.92	72.95
DRN [28]	Hourglass-104	✓	✓	89.71	82.34	47.22	64.10	76.22	74.43	85.84	90.57	86.18	84.89	57.65	61.93	69.30	69.63	58.48	73.23
TricubeNet (ours)	Hourglass-104			87.51	73.62	43.21	63.67	76.97	72.97	84.36	89.21	83.59	84.60	47.29	61.77	73.36	68.74	69.40	72.17
TricubeNet (ours)	Hourglass-104	✓		88.28	80.46	47.32	70.09	76.97	72.97	84.52	90.73	83.87	84.60	56.92	62.91	73.36	68.74	71.63	74.22
TricubeNet (ours)	Hourglass-104	✓	✓	88.75	82.12	49.24	72.98	77.64	74.53	84.65	90.81	86.02	85.38	58.69	63.59	73.82	69.67	71.08	75.26

4.2. Training details

For SKU110K-R, MSRA-TD500, and ICDAR2015 datasets, we set the input resolution to 800×800 and apply random cropping, random rotating, color jittering for data augmentation. For DOTA dataset, we set the input resolution to 1024×1024 and apply random rotating and color jittering data augmentation. Adam [17] is used as the optimizer and the learning rate is set to $2.5e-4$ for all datasets. For DOTA dataset, we train the model with 140 epochs and drop the learning rate by a factor 10 at 90 and 120 epochs. For SKU110K-R dataset, we train the model with 20 epochs without the learning rate decay. Following existing text detection approaches [6, 5], we pre-train the TricubeNet on SynthText [11] dataset with 3 epochs and finetuned on MSRA-TD500 and ICDAR2015 datasets. For MSRA-TD500 and ICDAR2015 datasets, we train the model with 180 epochs and drop the learning rate by a factor 10 at 120 and 160 epochs.

4.3. Testing details

We extract oriented bounding boxes using the post-processing algorithm described in Section 3.5. For DOTA dataset, we apply two kinds of test time augmentation: multi-scale testing and flip augmentation. For multi-scale testing, three scales (0.75, 1.0, and 1.5) and Soft-NMS for oriented bounding boxes are used. For the flip augmentation, we average output heatmaps. For the evaluation metric of DOTA, MSRA-TD500, and ICDAR2015, we adopt the mean average precision (mAP) of the 0.5 polygon IoU threshold. For the SKU110K-R dataset, we adopt COCO-style [22] evaluation method for oriented boxes: mAP at

Table 2. Evaluation results on SKU110K-R using the COCO-style metric.

Method	mAP	AP ₅₀	AP ₇₅	AR ₃₀₀
YOLOv3-Rotate [31]	49.1	-	51.1	58.2
CenterNet-4point [48]	34.3	-	19.6	42.2
CenterNet [48]	54.7	-	61.1	62.2
DRN [28]	55.9	-	63.1	63.3
TricubeNet (ours)	57.7	94.7	65.2	62.5

IoU=0.5:0.05:0.95, average precision AP₇₅ at IoU of 0.75, and average recall AR₃₀₀ at IoU=0.5:0.05:0.95.

4.4. State-of-the-art comparisons

We compare the performance of TricubeNet with the state-of-the-art methods on the test set of each dataset: Table 1 for DOTA, Table 2 for SKU110K-R, Table 3 for MSRA-TD500 and ICDAR2015. For the comparison, we apply two-steps heatmap cascade refinement. The anchor-based two-stage and one-stage detectors require a huge number of anchor boxes and multiple loss functions for the box offsets regression; it demands huge computational complexity and careful hyperparameter tuning. Meanwhile, anchor-free detectors can alleviate the above problems concerned with the anchor box by eliminating the anchor box in the network design, but their performance is slightly inferior to that of anchor-based detectors. However, our TricubeNet achieves highly competitive performance in all datasets despite its anchor-free setting. Especially, our outstanding performance in small vehicles (SV) shows that TricubeNet can accurately detect small objects without anchor box tuning, and we can obtain well-arranged oriented

Table 3. Evaluation on MSRA-TD500 and ICDAR2015 testset. * denotes for multi-scale test. R, P, and H denote recall, precision, and H-mean respectively.

Method	MSRA-TD500			ICDAR 2015		
	R	P	H	R	P	H
two-stage						
Wang <i>et al.</i> [38]	82.1	85.2	83.6	86.0	89.2	87.6
Gliding vertex [40]	84.3	88.8	86.5	-	-	-
one-stage						
SegLink [33]	70.0	86.0	77.0	76.8	73.1	75.0
RRD* [19]	73.0	87.0	79.0	80.0	88.0	83.8
Lyu <i>et al.</i> * [25]	76.2	87.6	81.5	79.7	89.5	84.3
Direct* [14]	81.0	91.0	86.0	80.0	85.0	82.0
anchor-free						
Zhang <i>et al.</i> [46]	67.0	83.0	74.0	43.0	71.0	54.0
EAST* [49]	67.4	87.3	76.1	78.3	83.3	80.7
TextSnake [24]	73.9	83.2	78.3	80.4	84.9	82.6
PixelLink* [6]	73.2	83.0	77.8	82.0	85.5	83.7
CRAFT [2]	78.2	88.2	82.9	84.3	89.8	86.9
TricubeNet (ours)	80.8	90.4	85.3	75.2	90.1	82.0

Table 4. The effect of each kernel representing an object.

kernel	Tricube	Gaussian	Effective Rect	Binary Rect
mAP	75.26	72.12	69.31	58.52

boxes as shown in Figure 7(a). In addition, the sufficiently high performance of 94.7% AP_{50} on SKU110K-R shows that TricubeNet is effective in detecting densely packed objects. When applying the same test augmentation, Gliding vertex [40], one of the two-stage detectors, shows the state-of-the-art performance (75.0% *v.s.* 74.2%) as in Table 1, we argue that our efficiency is higher than them because we exploit the advantages of the one-stage anchor-free setting and simply solve the loss discontinuity problem.

4.5. Ablation study

To analyze the effectiveness of each component of TricubeNet, we conduct an ablation study. For the evaluation, we experiment on DOTA validation set following all parameter settings as given in Section 4.2.

First, we investigate which 2D kernel is more effective in representing the object. For the comparison, we choose one of four different types of 2D kernels: 2D Tricube kernel, 2D Gaussian kernel, binary rectangle, effective rectangle. The effective rectangle is a shrunk binary rectangle, which is employed in FSAF [50], and we use a binary rectangle shrunk by 40% as the effective rectangle. Table 4 shows the DOTA validation score of each kernel. The effective rectangle achieves higher performance than the binary rectangle because the binary rectangle often fails to separate weakly-occluded objects. Although the Gaussian kernel has a limitation in representing the angle of the object as discussed in Section 3.1, the Gaussian kernel shows a 2.81% higher performance than the effective rectangle; this result demon-

Table 5. Ablation study for the proposed techniques: size-weight mask (SWM), false-positive example mining (FP EM), multi-angle convolution (MAC) module, one-step heatmap refinement (Cascade 1), and two-step heatmap refinement (Cascade 2).

	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L	Params	GFlops
Baseline	30.4	61.6	25.2	16.3	32.4	37.6	188M	1015G
+ SWM	31.1	64.0	25.6	19.2	33.0	36.0	+0M	+0G
+ FP EM	32.3	66.8	25.2	20.5	35.4	37.6	+0M	+0G
+ MAC	33.6	67.7	28.7	20.2	35.9	39.8	+0.4M	+28G
+ Cascade 1	35.3	69.6	30.3	24.1	37.2	40.1	+1.2M	+85G
+ Cascade 2	35.6	70.1	30.9	24.6	38.5	40.3	+2.1M	+143G

Table 6. The analysis of each method for extracting rotation-invariant features.

	AP	AP_{50}	AP_{75}	Params	GFlops
CNN	32.3	66.8	25.2	+0.00M	+00G
DCN [51]	32.9	67.0	28.1	+0.55M	+40G
RCL [28]	33.4	67.6	28.5	+0.49M	+34G
MAC (ours)	33.6	67.7	28.7	+0.43M	+28G

strates the advantage of the kernel-based object representation. Since the Tricube kernel can solve the limitation of the Gaussian kernel while taking advantage of the kernel-based object representation, the Tricube kernel is the most suitable choice for the oriented object and achieves the highest performance among the four types of kernels.

Second, we evaluate the effect of proposed techniques, *i.e.*, size-weight mask (SWM), false-positive example mining (FP EM), multi-angle convolution (MAC) module, and heatmap cascade refinement. For a more precise analysis, we measure the performance using COCO-style evaluation metric as in Table 5. AP_S , AP_M , and AP_L denote the AP when the object size is smaller than 32×32 , between 32×32 and 96×96 , and larger than 96×96 , respectively. In addition, we report the number of parameters and GFlops. We set the baseline model as the hourglass-104 network trained using MSE loss without any proposed techniques and evaluate the effect of each technique. When applying both SWM and FP EM, AP and AP_S are considerably improved by 1.9% and 4.2%, respectively; it clearly shows that SWM and FP EM help in solving size-imbalance and class-imbalance problems. The MAC module, which is designed for extracting the rotation-invariant feature, yields 1.3% performance improvement. One-step heatmap refinement further improves the performance by 2.3%. However, two-step heatmap refinement only yields 0.3% improvement; we conclude that two-step heatmap refinement is enough to produce a high-quality heatmap.

Last, in Table 6, we evaluate our MAC module by comparing with a conventional convolution (CNN), deformable convolution [51] (DCN), and rotation convolution layer [28] (RCL). Compared to CNN and DCN, our MAC module can properly extract rotation-invariant features, so the performance is improved by 1.3% and 0.7%, respectively. Since both MAC and RCL are designed to extract rotation-

invariant features, our MAC module shows a similar performance improvement to the RCL. However, the RCL requires heavier computations and about 60% more GPU memory usage than our MAC during the training. From the above result, we prove that the proposed MAC module is more effective and efficient in our TricubeNet, which is a fully-heatmap-based approach.


4.6. Qualitative Results

We collect some qualitative results on each dataset: Figure 7(a) for DOTA, Figure 7(b) for MSRA-TD500, Figure 7(c) for ICDAR 2015, and Figure 7(d) for SKU110K-R. From the results, we can notice that our TricubeNet handles the challenging factors of oriented object detection. Specifically, as in Figure 7(a), TricubeNet can properly detect objects of various sizes within the multi-category classification problem; as in Figure 7(c), TricubeNet can accurately detect the rotated text despite the complex background; as in Figure 7(d), TricubeNet is outstanding in the detection of a huge number of densely arranged oriented objects. Furthermore, we can identify that the consistency of the rotated bounding boxes detecting the arranged objects is very high.

5. Conclusion and Future Work

We present a novel approach for oriented object detection, named TricubeNet. Our main concept is that we localize objects using visual cues (*i.e.*, heatmap) instead of box offsets regression. We represent each object as a 2D Tricube kernel and extract bounding boxes using simple image-processing algorithms. Unlike anchor-based approaches, TricubeNet is able to obtain well-arranged oriented boxes from visual cues, solve the angle discontinuity problem by taking away the angle regression, and save the computational complexity due to our anchor-free modeling. Additionally, for the better fit in oriented object detection, we propose some effective techniques: SWM for the size-invariant loss function; FPEM for balancing between foreground and background pixels and reducing false-positive detections; MAC module to extract the rotation-invariant features; heatmap cascade refinement for progressively refine the output heatmap. The extensive experimental results on various datasets show that our TricubeNet is considerably effective and efficient for oriented object detection.

For future work, our failure cases are should be addressed. TricubeNet is quite effective to detect weakly-occluded oriented objects, however, it has trouble detecting heavily occluded objects. Also, the extracted box offsets from our post-processing algorithm are integer type, which weakens the precise detection of tiny objects. Addressing these limitations would give a great improvement.




(a) DOTA



(b) MSRA-TD500



(c) ICDAR 2015



(d) SKU110K-R

Figure 7. Qualitative results of TricubeNet.

Acknowledgement

This research was supported in part by Autonomous Driving Center, R&D Division, Hyundai Motor Company when the authors worked at KAIST.

References

- [1] Seyed Majid Azimi, Eleonora Vig, Reza Bahmanyar, Marco Körner, and Peter Reinartz. Towards multi-class object detection in unconstrained remote sensing imagery. In *Asian Conference on Computer Vision*, pages 150–165. Springer, 2018. 6
- [2] Youngmin Baek, Bado Lee, Dongyoon Han, Sangdoon Yun, and Hwalsuk Lee. Character region awareness for text detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9365–9374, 2019. 3, 7
- [3] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000. 5
- [4] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018. 1, 2, 4
- [5] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee. Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 502–511, 2019. 6
- [6] Dan Deng, Haifeng Liu, Xuelong Li, and Deng Cai. Pixellink: Detecting scene text via instance segmentation. *arXiv preprint arXiv:1801.01315*, 2018. 3, 6, 7
- [7] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2849–2858, 2019. 1, 2, 5, 6
- [8] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 2
- [9] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 2
- [10] Eran Goldman, Roei Herzig, Aviv Eisenschtat, Jacob Goldberger, and Tal Hassner. Precise detection in densely packed scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5227–5236, 2019. 2, 5
- [11] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2315–2324, 2016. 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916, 2015. 2
- [13] Lifeng He, Yuyan Chao, Kenji Suzuki, and Kesheng Wu. Fast connected-component labeling. *Pattern recognition*, 42(9):1977–1987, 2009. 5
- [14] Wenhao He, Xu-Yao Zhang, Fei Yin, and Cheng-Lin Liu. Multi-oriented and multi-lingual scene text detection with direct regression. *IEEE Transactions on Image Processing*, 27(11):5406–5419, 2018. 7
- [15] Yingying Jiang, Xiangyu Zhu, Xiaobing Wang, Shuli Yang, Wei Li, Hua Wang, Pei Fu, and Zhenbo Luo. R2cnn: rotational region cnn for orientation robust scene text detection. *arXiv preprint arXiv:1706.09579*, 2017. 6
- [16] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pages 1156–1160. IEEE, 2015. 2, 5
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [18] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 734–750, 2018. 2, 3
- [19] Minghui Liao, Zhen Zhu, Baoguang Shi, Gui-song Xia, and Xiang Bai. Rotation-sensitive regression for oriented scene text detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5909–5918, 2018. 7
- [20] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 1
- [21] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 1, 2, 6
- [22] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [23] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 2
- [24] Shangbang Long, Jiaqiang Ruan, Wenjie Zhang, Xin He, Wenhao Wu, and Cong Yao. Textsnake: A flexible representation for detecting text of arbitrary shapes. In *Proceedings of the European conference on computer vision (ECCV)*, pages 20–36, 2018. 7
- [25] Pengyuan Lyu, Cong Yao, Wenhao Wu, Shuicheng Yan, and Xiang Bai. Multi-oriented scene text detection via corner localization and region segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7553–7563, 2018. 7
- [26] 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*, 20(11):3111–3122, 2018. 1, 2, 6

- [27] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. [1](#), [2](#), [6](#), [7](#)
- [28] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11207–11216, 2020. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [29] Wen Qian, Xue Yang, Silong Peng, Yue Guo, and Chijun Yan. Learning modulated loss for rotated object detection. *arXiv preprint arXiv:1911.08299*, 2019. [2](#), [6](#)
- [30] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [31] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. [6](#)
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. [1](#), [2](#)
- [33] Baoguang Shi, Xiang Bai, and Serge Belongie. Detecting oriented text in natural images by linking segments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2550–2558, 2017. [7](#)
- [34] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 761–769, 2016. [5](#)
- [35] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5693–5703, 2019.
- [36] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10781–10790, 2020. [1](#)
- [37] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013. [2](#)
- [38] Xiaobing Wang, Yingying Jiang, Zhenbo Luo, Cheng-Lin Liu, Hyunsoo Choi, and Sungjin Kim. Arbitrary shape scene text detection with adaptive text region representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6449–6458, 2019. [7](#)
- [39] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3974–3983, 2018. [2](#), [5](#), [6](#)
- [40] Yongchao Xu, Mingtao Fu, Qimeng Wang, Yukang Wang, Kai Chen, Gui-Song Xia, and Xiang Bai. Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. [1](#), [2](#), [6](#), [7](#)
- [41] Xue Yang, Qingqing Liu, Junchi Yan, and Ang Li. R3det: Refined single-stage detector with feature refinement for rotating object. *arXiv preprint arXiv:1908.05612*, 2019. [2](#), [6](#)
- [42] Xue Yang, Hao Sun, Kun Fu, Jirui Yang, Xian Sun, Menglong Yan, and Zhi Guo. Automatic ship detection in remote sensing images from google earth of complex scenes based on multiscale rotation dense feature pyramid networks. *Remote Sensing*, 10(1):132, 2018. [6](#)
- [43] Xue Yang, Hao Sun, Xian Sun, Menglong Yan, Zhi Guo, and Kun Fu. Position detection and direction prediction for arbitrary-oriented ships via multitask rotation region convolutional neural network. *IEEE Access*, 6:50839–50849, 2018.
- [44] 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 *Proceedings of the IEEE International Conference on Computer Vision*, pages 8232–8241, 2019. [2](#), [6](#)
- [45] Cong Yao, Xiang Bai, Wenyu Liu, Yi Ma, and Zhuowen Tu. Detecting texts of arbitrary orientations in natural images. In *2012 IEEE conference on computer vision and pattern recognition*, pages 1083–1090. IEEE, 2012. [2](#), [5](#)
- [46] Zheng Zhang, Chengquan Zhang, Wei Shen, Cong Yao, Wenyu Liu, and Xiang Bai. Multi-oriented text detection with fully convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4159–4167, 2016. [7](#)
- [47] Lin Zhou, Haoran Wei, Hao Li, Wenzhe Zhao, Yi Zhang, and Yue Zhang. Arbitrary-oriented object detection in remote sensing images based on polar coordinates. *IEEE Access*, 2020.
- [48] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. [1](#), [2](#), [3](#), [6](#)
- [49] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: an efficient and accurate scene text detector. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 5551–5560, 2017. [7](#)
- [50] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 840–849, 2019. [7](#)
- [51] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9308–9316, 2019. [7](#)