Rethinking IoU-based Optimization for Single-stage 3D Object Detection

Hualian Sheng[1,2,3], Sijia Cai[2], Na Zhao[3*], Bing Deng[2], Jianqiang Huang[2], Xian-Sheng Hua[2], Min-Jian Zhao[1], and Gim Hee Lee[3]

[1] College of Information Science and Electronic Engineering, Zhejiang University
[2] Alibaba Cloud Computing Ltd.
[3] Department of Computer Science, National University of Singapore
{hlsheng,mjzhao}@zju.edu.cn, {nazhao, gimhee.lee}@comp.nus.edu.sg,
{stephen.csj, dengbing.db, jianqiang.hjq, xiansheng.hxs}@alibaba-inc.com

Abstract. Since Intersection-over-Union (IoU) based optimization maintains the consistency of the final IoU prediction metric and losses, it has been widely used in both regression and classification branches of single-stage 2D object detectors. Recently, several 3D object detection methods adopt IoU-based optimization and directly replace the 2D IoU with 3D IoU. However, such a direct computation in 3D is very costly due to the complex implementation and inefficient backward operations. Moreover, 3D IoU-based optimization is sub-optimal as it is sensitive to rotation and thus can cause training instability and detection performance deterioration. In this paper, we propose a novel Rotation-Decoupled IoU (RDIoU) method that can mitigate the rotation-sensitivity issue, and produce more efficient optimization objectives compared with 3D IoU during the training stage. Specifically, our RDIoU simplifies the complex interactions of regression parameters by decoupling the rotation variable as an independent term, yet preserving the geometry of 3D IoU. By incorporating RDIoU into both the regression and classification branches, the network is encouraged to learn more precise bounding boxes and concurrently overcome the misalignment issue between classification and regression. Extensive experiments on the benchmark KITTI and Waymo Open Dataset validate that our RDIoU method can bring substantial improvement for the single-stage 3D object detection. The code is available at https://github.com/hlsheng1/RDIoU.
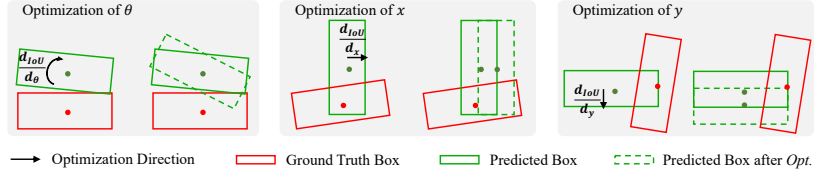
Keywords: 3D object detection; Single-stage; Rotation-Decoupled IoU

1 Introduction

Point cloud-based 3D object detection is an important task in robotics and autonomous driving. In contrast to the image-based object detection task, 3D object detection needs to predict 3D bounding boxes with higher degrees-of-freedom

---

[*] Corresponding author

Fig. 1. Illustration of direct 3D IoU optimization between the ground truth box (red) and predicted box (green) in Bird's Eye View (BEV), compared to the post-optimized bounding box (dotted green). We see that optimizing the 3D IoU over the bounding box parameters does not bring the predicted boxes closer to the ground truths.

and allocate confidence map in a larger search space. This causes difficulties in predicting accurate 3D bounding boxes with reliable confidence.

State-of-the-art 3D detection approaches can be generally divided into single-stage and two-stage methods. Two-stage methods [1,17,22,23,36] usually perform a second feature extraction or multi-scale feature aggregation based on the proposals generated by the region proposal network (RPN). As a result, the second-stage network can focus on the partial positive areas to avoid the sparsity of the whole point cloud. Furthermore, the second-stage network can utilize Intersection-over-Union (IoU) guided supervision to generate more reliable confidence predictions. In contrast to the coarse-to-fine pipeline in the two-stage methods, single-stage methods [9,10,31,36,37,38] adopt an end-to-end pipeline that designs one dense object detector to directly predict pixel-level object categories and bounding boxes. The end-to-end one-stage pipeline is much more elegant and effective than the two-stage pipeline. Nevertheless, the performance of single-stage methods is usually inferior to the two-stage methods. The performance gap is mainly caused by inaccurate regression precision and confidence assignment of the single-stage methods since they independently predict the 3D bounding box center, size and rotation with the Smooth-$\ell 1$ loss [6], which results in inconsistency between the loss and the final IoU-based evaluation metric.

To maintain the consistency between the loss and IoU-based evaluation metric, IoU-based optimization [11,12,20,40] has been widely studied and shown impressive performance improvements in image-based single-stage object detection. Inspired by this, several single-stage 3D object detection approaches [38,41] adopt the IoU-based optimization by replacing the 2D IoU with the 3D IoU. Despite their impressive performance gain in the axis-aligned cases, the computation of intersection area between two rotated 3D bounding boxes is much more complex than their 2D counterparts without rotation. A number of efforts [2,32,33,39] have been made to simplify the computation of rotated IoU intersection by approximations. For example, PIoU [2] only counts the pixels in the intersection area. However, none of these methods can be directly utilized to solve the **negative coupling effect of rotation on the 3D IoU**. As illustrated in Figure 1, optimizing the 3D IoU over the bounding box parameters can lead to further misalignment between the predicted and ground truth bounding boxes. This problem is caused by the coupling between the rotation and the

center and size parameters of the 3D bounding box. In contrast, the issue of negative coupling effect of rotation does not exist in the 2D IoU counterpart under common settings as the optimization over the center and size parameters directly minimizes the 2D IoU without rotation. To this end, we propose a new high-performance and efficient 3D bounding box optimization objective that satisfies the following three conditions: it must 1) be differentiable; 2) satisfy the consistency between the evaluation metric and optimization objective; 3) circumvent the negative coupling effect of rotation on the 3D IoU.

In this paper, we propose a Rotation-Decoupled IoU (RDIoU) method to model the interaction of two arbitrary-oriented 3D bounding boxes while stabilize the training process. Our main idea is to disentangle dependencies of the bounding box parameters by decoupling and handling the rotation individually at a loss level. Specifically, we utilize a 4-dimensional representation of the box to calculate the IoU-like criterion and remove complex rotation-edge variations. A fixed hyperparameter is introduced to further control the weights of rotation change in the computation of our RDIoU. Our proposed RDIoU formulation addresses the non-differentiable and instability issues caused by the rotation change while preserving the geometry of 3D IoU. By incorporating RDIoU into the regression and classification branches, we propose a RDIoU-guided DIoU loss [40] and a RDIoU-guided quality focal loss [11,12]. These two enhanced losses guide the optimization towards high-performance box localization and alleviate the misalignment between classification and regression, and thus significantly improve the single-stage 3D detection performance over other existing losses.

Based on our proposed RDIoU, we build a simple and elegant single-stage 3D detector. In the training phase, our RDIoU-based optimization guides the 3D CNN backbone to achieve better feature alignment and get more accurate box parameters without encountering the rotation-sensitive issues. Furthermore, our RDIoU-based single-stage 3D detector also enables us to effectively train the entire network in an end-to-end manner without the need of hyperparameter-sensitive, stage-wise training or time-consuming backward operations. Our main contributions in this paper are: 1) We propose the RDIoU-based optimization for 3D object detectors, which brings a more robust training process as compared to the 3D IoU-based optimization strategy. 2) We incorporate RDIoU into regression supervision and present the RDIoU-guided DIoU loss, which exhibits significantly better performance as compared to the 3D IoU-guided DIoU loss [38]. 3) We incorporate RDIoU into classification supervision and present the RDIoU-guided quality focal loss, which is able to help generate more reasonable confidence maps. 4) We conduct extensive experiments on two benchmark datasets (*i.e.,* KITTI and Waymo), and the promising results validate that our RDIoU can attain top performance with the commonly used backbone networks.
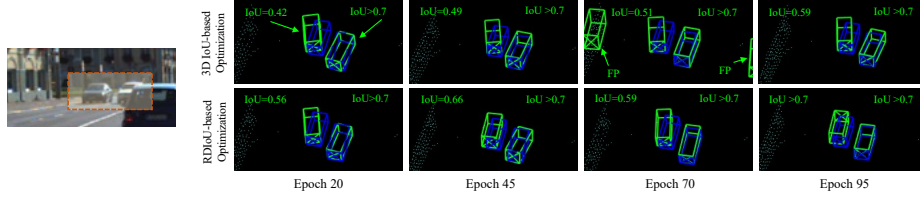
2   Related Work

Single-stage Detectors. Single-stage detectors predict the location and category from predefined anchor boxes or points over different spatial positions in

a single-shot manner such as SECOND [31] and PointPillar [10]. SECOND [31] is the first to apply 3D sparse convolution [7,8] to process the 3D voxel features, which are averaged by voxel-enclosed points. PointPillar [10] collapses the points in vertical pillars with a simplified PointNet [18,19], followed by a typical 2D CNN backbone. Point-GNN [25] proposes a graph neural network to learn better point features. 3DSSD [34] develops F-FPS as a supplement of D-FPS to build an anchor-free 3D object detector. SA-SSD [9] applies an auxiliary segment network to assist the voxel feature learning. CIA-SSD [37] utilizes an IoU-aware confidence rectification to improve the classification. SE-SSD [38] proposes a self-ensembling training schedule to improve the performance of a pretrained CIA-SSD model. These single-stage detectors are usually of high speed.

Two-stage Detectors. Two-stage detectors can be seen as the extension of single-stage detectors. These methods first generate high-quality proposals with categories based on RPN, and then refine each proposal's location and output the prediction confidence. PointRCNN [23] generates proposals based on segmented foreground objects and then refines the bounding box via a regression branch. Part A^2 [24] extends PointRCNN with an intra-object part supervision. PV-RCNN [22] first utilizes voxel-based backbone to generate high-quality proposals, and then devotes a point-based approach to aggregate multi-scale voxel features along with raw point features. Its variant Voxel-RCNN [4] proposes an accelerated point-based module to capture multi-scale voxel features without raw point features. LiDAR-RCNN [13] introduces a plug-and-play point-based module to refine the RPN proposals based on simple PointNet-like approach. CenterPoint [36] suggests an anchor-free method to improve the RPN detect head, and then learns more accurate bounding box regression based on BEV features within given proposals. Transformer architecture [28] recently appears in the 3D detection area. VoTr [16] improves the region proposal network by introducing Transformer into sparse convolution. CT3D [21] proposes a novel channel-wise Transformer for proposal refinement. These two-stage detectors usually exhibit high performance but suffer the problem of having too many hyperparameters and high latency, and thus limiting their usefulness in the industrial application. In comparison, it is imperative to improve the single-stage detector for the better trade-off between performance and latency.

IoU-based Optimization. IoU-based optimization has been effectively validated and implemented in 2D object detection without rotation, e.g. GIoU loss [20], DIoU/CIoU loss [40], quality focal loss [11,12], etc. These methods aim to directly optimize the final evaluation metric and have achieved promising performance. Recently, several 3D rotated object detection works study a direct extension of the 2D IoU to 3D domain and yield the 3D IoU loss [41] and ODIoU loss [38]. However, these 3D IoU-based losses usually need high resource cost such as huge GPU memory cost due to brute-force search or time-consuming back-propagation on the CPU. Moreover, it is sub-optimal to set the 3D IoU as the optimization objective as shown in Figure 1. Instead, we propose a more effective optimization objective, *i.e.* the RDIoU, to achieve both high-performance training and efficient back-propagation on the GPU.

Fig. 2. Qualitative comparison between the 3D IoU-based and RDIoU-based optimization strategies on the KITTI *train* set. The predicted and ground truth bounding boxes are shown in green and blue, respectively. The 3D IoU-based strategy is difficult to handle the hard sample, and the IoU is eventually stagnant under 0.7. In contrast, RDIoU-based strategy finally makes the prediction much closer to the ground truth.

3   Method

Our main contribution is the RDIoU, which is a better optimization objective than the existing 3D IoU in alleviating the negative coupling effect of the rotation during the training stage. In the following, we first discuss the details of our RDIoU and the comparison with 3D IoU. Subsequently, we insert our RDIoU into the regression and classification branches of existing 3D object detectors.
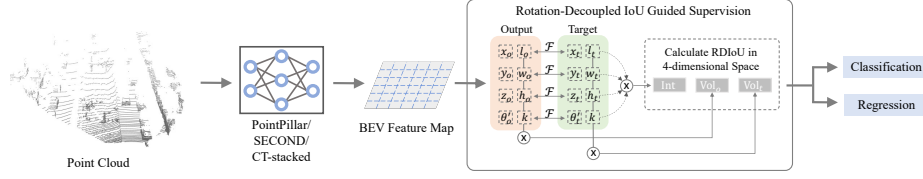
3.1   Our RDIoU

The optimization strategies for object detection generally fall into two types:

1. Non IoU-based: Each box parameter is individually optimized without considering their spatial connection. Typically, Smooth-$\ell 1$ loss [6] and focal loss [14] are used for regression and classification, respectively.
2. IoU-based: The box parameters are jointly optimized via maintaining consistency with the IoU evaluation metric. Typically, DIoU/CIoU loss [40] and quality focal loss [11] are used for regression and classification, respectively.

Compared to non IoU-based optimization, the IoU-based optimization strategy usually exhibits better performance due to the end-to-end metric learning. However, a naive application of this strategy on 3D object detection causes potential instability issue in the optimization due to the negative coupling effect of the rotation and thus affects the final convergence of the network, as shown in the first row of Figure 2. In view of this problem, we construct a more efficient optimization strategy that considers both the stability of the training process and the preservation of the geometry from the IoU on the bounding box. Figure 3 illustrates our proposed RDIoU method.

Let $(x_o, y_o, z_o, l_o, w_o, h_o, \theta_o)$ denotes the 3D bounding box parameters of the regression vector from the regression branch and $(x_t, y_t, z_t, l_t, w_t, h_t, \theta_t)$ denotes its corresponding target regression vector. (x, y, z) is the center coordinate, (l, w, h) is the size, and θ is the rotation of the bounding box. Furthermore, we use $(x_a, y_a, z_a, l_a, w_a, h_a, \theta_a)$ to denote the selected anchor box from the 3D object

Fig. 3. The illustration of RDIoU method. We advocate a Rotation-Decoupled IoU to model the intersection of two rotated 3D bounding boxes. This design is further set as the optimization target to guide the classification and regression learning, respectively.

detector and $(x_g, y_g, z_g, l_g, w_g, h_g, \theta_g)$ to denote the ground truth 3D bounding box. Following the previous works [31], the regression target can be encoded as:

$$x_t = \frac{x_g - x_a}{d}, y_t = \frac{y_g - y_a}{d}, z_t = \frac{z_g - z_a}{h_a}, \tag{1}$$

$$l_t = \frac{l_g}{l_a}, \quad w_t = \frac{w_g}{w_a}, \quad h_t = \frac{h_g}{h_a}, \quad \theta_t = \theta_g - \theta_a, \tag{2}$$

where $d = \sqrt{(l_a)^2 + (w_a)^2}$ is the diagonal of the base of anchor.

Since the negative coupling effect of rotation is the main source of inaccuracies in the 3D IoU-based optimization, we decouple the rotation into an independent dimension to avoid its complex coupling relationship with the other bounding box parameters. Specifically, the rotation is decoupled into the 4-*th* dimension with a fixed edge k. As a result, the center of the optimized bounding box is transformed from (x_t, y_t, z_t) to $(x_t, y_t, z_t, \theta_t)$, and thus we can easily calculate the IoU in this 4-dimensional space. Consequently, the RDIoU of two rotated 3D bounding boxes can be formulated with a regular IoU [20] as:

$$\text{RDIoU} = \text{Int}/(\text{Vol}_o + \text{Vol}_t - \text{Int}), \quad \text{where} \tag{3}$$

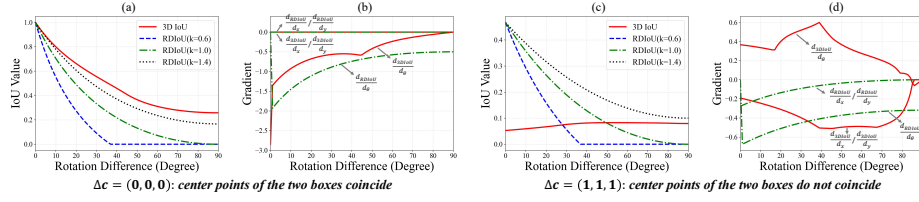$$\text{Int} = \mathcal{F}(x_o, x_t, l_o, l_t) * \mathcal{F}(y_o, y_t, w_o, w_t) * \mathcal{F}(z_o, z_t, h_o, h_t) * \mathcal{F}(\theta_{o'}, \theta_{t'}, k, k),$$
$$\text{Vol}_o = l_o * w_o * h_o * k, \quad \text{Vol}_t = l_t * w_t * h_t * k,$$

and $\quad \mathcal{F}(a_o, a_t, b_o, b_t) = \min(a_o + \frac{b_o}{2}, a_t + \frac{b_t}{2}) - \max(a_o - \frac{b_o}{2}, a_t - \frac{b_t}{2}).$

Following the common setting [31], we set $\theta_{o'} = \sin\theta_o \cos\theta_t$ and $\theta_{t'} = \cos\theta_o \sin\theta_t$. "Int" denotes the intersection volume between the two 4-dimensional bounding boxes. k is side length of the fourth dimension corresponding to center θ, which we empirically set as $k = 1$ to achieve the best performance.

Our proposed RDIoU inherits the geometry constraints from the 3D IoU through an IoU-like formulation in the 4-dimensional space. Every variable is differentiable with our RDIoU as the regression target, and thus can be trained with the regular gradient-based back-propagation on the GPU. Furthermore, the decoupling of the rotation as an independent variable in our RDIoU leads to the right optimization direction for each individual variable with the increasing IoU during optimization.

Fig. 4. Simulation experiments. We initialize two 3D boxes of size equal to the predefined anchor size, *i.e.*, (3.9,1.6,1.56). Here, Δc is the box center distance. The gradients of RDIoU are computed with $k = 1$.

3.2   More Comparison between RDIoU and 3D IoU

To further investigate the intrinsic difference, we simulate the values of 3D IoU and our RDIoU under increasing rotation differences. As shown in Figure 4(a), the values of RDIoU and 3D IoU are similar when the center points of the two boxes coincide. In contrast, as shown in Figure 4(c), the values become distinct when the center points of the two boxes do not coincide. Specifically, we can see that the 3D IoU yields unreasonable dynamics in Figure 4(c), where the computed IoU value increases even when the rotation difference becomes larger. We utilize the hyperparameter k to control the weight of rotation change. As can be seen from Figure 4(a) and (c), a smaller k results in a steeper curve. Furthermore, we plot the gradient changes of x, y and θ with increasing rotation difference in Figure 4(b) and (d) corresponding to Figure 4(a) and (c), respectively. From Figure 4(b), we can see that the gradients of 3D IoU and RDIoU are close when the center points of the two boxes coincide. However, as shown in Figure 4(d), the gradient over rotation of 3D IoU is obviously misdirected ($\frac{d_{3DIoU}}{d_{\theta}} > 0$) when the center points of the two boxes do not coincide. In this case, the gradient of 3D IoU changes drastically, while RDIoU still performs in a relatively smooth manner.

3.3   Incorporating RDIoU into Regression Supervision

We modify the DIoU loss [40] for 2D object detection to take our RDIoU for 3D object detection with rotation. Specifically, the modified regression loss aims to maximize the intersection volume, *i.e.* the RDIoU and minimize the normalized center distance in the 4-dimensional space. Let $\boldsymbol{c}_o = [x_o, y_o, z_o, \theta_{o'}]$ and $\boldsymbol{c}_t = [x_t, y_t, z_t, \theta_{t'}]$, the center distance penalty is:

$$\rho_c = \frac{\delta(\boldsymbol{c}_o, \boldsymbol{c}_t)}{\text{Diag}}, \tag{4}$$

where $\delta(\cdot)$ denotes the square of Euclidean distance and

$$\text{Diag} = \mathcal{G}(x_o, x_t, l_o, l_t) + \mathcal{G}(y_o, y_t, w_o, w_t) + \mathcal{G}(z_o, z_t, h_o, h_t) + \mathcal{G}(\theta_{o'}, \theta_{t'}, k, k),$$

where $\mathcal{G}(a_o, a_t, b_o, b_t) = \left(\max(a_o + \frac{b_o}{2}, a_t + \frac{b_t}{2}) - \min(a_o - \frac{b_o}{2}, a_t - \frac{b_t}{2})\right)^2$. As a result, the final RDIoU-guided DIoU loss for 3D bounding box regression is:

$$\mathcal{L}_{RL} = 1 - \text{RDIoU} + \rho_c. \tag{5}$$

Compared with the existing 3D variants of DIoU losses [38,41] that directly maximize the 3D intersection volume and minimize the 3D center distance, our RDIoU-guided DIoU loss has the advantages of more efficient learning and low resource cost. The rotation decoupling operator transforms the rotated 3D bounding box into a 4-dimensional latent space, which leads to stable optimization direction for each variable. Unlike [38,41] that need huge GPU memory cost or necessary pre-trained model, our RDIoU-guided DIoU loss can be easily implemented in GPU with a small cost. In the ablation studies, we show RDIoU-guided DIoU loss can surpass the 3D IoU-guided DIoU loss by a large margin.

3.4   Incorporating RDIoU into Classification Supervision

Solving the misalignment problem is also a crucial issue for object detection. Thus, we derive a RDIoU-guided quality focal loss to jointly supervise the classification and the bounding box quality estimation. The quality focal loss [11,12] has achieved good performance in 2D image detection, where the IoUs between predicted and ground truth boxes are set as the quality targets. For 3D detection with rotated 3D bounding boxes, an intuitive thought is to apply the quality focal loss with 3D IoU-guided quality supervision. However, 3D IoU can cause confusion when the higher prediction error corresponds to a larger 3D IoU (see Figure 1), and therefore it leads to bad dynamics during the optimization.

On the contrary, our RDIoU is less susceptible to the change of regression variables and thus it can be seen as a better measure of the regression quality. Consequently, we utilize our RDIoU as the quality estimation target. Similar to GFLV2 [11], we adopt joint representation targets with the quality focal loss for both classification and bounding box quality prediction. Specifically, the target for each sample can be expressed as $\boldsymbol{t} = \boldsymbol{r} \times e$, where $\boldsymbol{r} = [r_1, r_2, \ldots, r_m], r_i \in \{0, 1\}$ is the classification representation with m classes and $e \in [0, 1]$ is the scalar RDIoU representation. Taking the joint estimation as y, the RDIoU-guided quality focal loss is given by:

$$\mathcal{L}_{RQFL} = -\beta_1 |\text{RDIoU} - y|^{\beta_2} \big((1 - \text{RDIoU}) \log(1 - y) + \text{RDIoU} \log(y)\big), \quad (6)$$

where $\beta_1 = 0.25$ and $\beta_2 = 2$ follow the setting of focal loss. In this way, the predicted values of classification head model both the category probability and the corresponding bounding box confidence. Furthermore, we use the cross-entropy loss for direction classification ($\mathcal{L}_d$). Finally, the overall loss for training is:

$$\mathcal{L} = \mathcal{L}_{RQFL} + \gamma_1 \mathcal{L}_d + \gamma_2 \mathcal{L}_{RL}, \tag{7}$$

where γ_1 and γ_2 are hyperparameters to weigh the loss terms.

Backbone Network. Following the previous works [10,31], we voxelize the point cloud to produce a regular representation. In each voxel, we calculate the mean coordinates and intensities of the raw points as the voxel feature. Subsequently, the 3D backbone network gradually converts the initial voxel features into high-dimensional feature volumes, and the 3D features along Z-axis are stacked into 2D BEV feature maps. A 2D backbone network is designed to enhance the extraction of the BEV features for the final location and classification tasks. Motivated by the success of Transformer [28] in computer vision-related tasks, we adopt a stacked architecture that contains both the convolution block and the Transformer block to generate the final detection features. Specifically, the convolution block first extracts the BEV features while keeping the dimensions (both the number of channels and the size of feature map) unchanged. A convolution layer with stride two is then adopted to further reduce the spatial size. The Transformer block then processes the high-level features without changing the feature dimensions. Our designed conv-trans-stacked (CT-stacked) architecture shows better performance as compared to the pure convolution based network. The whole framework of our RDIoU is illustrated in Figure 3. The final feature map is $8\times$ downsampling of the input resolution. The classification, quality estimation, and regression are performed at pixel-level.

4   Experiments

In this section, we compare our RDIoU to other state-of-the-art methods on two popular benchmarks: KITTI [5] and Waymo Open Dataset [26]. Furthermore, we conduct extensive ablation studies to investigate the performance on different backbone networks and on each component of RDIoU to validate our design.

4.1   Datasets

KITTI. This dataset consists of 7,481 LiDAR samples for training and 7,518 LiDAR samples for testing. We further follow the common protocol [10,31] to split the original training data into 3,712 training samples and 3,769 validation samples for experimental studies.

Waymo Open Dataset. This dataset consists of 798 training sequences with 158,361 LiDAR samples, and 202 validation sequences with 40,077 LiDAR samples. Five LiDAR sensors are used for full 360-degree annotation instead of 90-degree as in KITTI. Currently, it is the largest dataset for autonomous driving.

4.2   Implementation Details

For KITTI, the raw point clouds are first clipped into $(0, 70.4)$m, $(-40, 40)m$, $(-3, 1)$m for X, Y, Z axis ranges with voxel size $(0.05, 0.05, 0.1)$m. For Waymo Open Dataset, the corresponding axis ranges are $(-75.2, 75.2)$m, $(-75.2, 75.2)$m, $(-2, 4)$m, and the voxel size is $(0.1, 0.1, 0.15)$m. We conduct all experiments based on the OpenPCDet [27] toolbox.

Type	Method	M	Car-3D (IoU=0.7)			Ped.-3D (IoU=0.5)			Cyc.-3D (IoU=0.5)		
			Easy	Mod.*	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
Two-stage	Part-A^2 [24]	✓	**89.56**	79.41	78.84	**65.69**	**60.05**	55.45	85.50	69.90	65.49
	STD [35]	✗	89.70	79.80	**79.30**	-	-	-	-	-	-
	PV-RCNN [22]	✓	89.35	83.69	78.70	63.12	54.84	51.78	**86.06**	69.48	64.50
	Voxel-RCNN [4]	✗	89.41	84.52	78.93	-	-	-	-	-	-
	VoTr-TSD [16]	✗	89.04	84.04	78.68	-	-	-	-	-	-
	CT3D [21]	✓	89.11	85.04	78.76	64.23	59.84	**55.76**	85.04	**71.71**	**68.05**
	CT3D [21]	✗	89.54	86.06	78.99	-	-	-	-	-	-
	BtcDet [30]	✗	-	**86.57**	-	-	-	-	-	-	-
Single-stage	VoxelNet [43]	✓	81.97	65.46	62.85	57.86	53.42	48.87	67.17	47.65	45.11
	SECOND [31]	✓	88.61	78.62	77.22	56.55	52.98	47.73	80.59	67.16	63.11
	PointPillar [10]	✓	86.46	77.28	74.65	57.75	52.29	47.91	80.06	62.69	59.71
	3DSSD [34]	✓	89.71	79.45	78.67	-	-	-	-	-	-
	SA-SSD [9]	✗	**90.15**	79.91	78.78	-	-	-	-	-	-
	CIA-SSD [37]	✗	90.04	78.91	78.80	-	-	-	-	-	-
	SE-SSD [38]	✗	-	85.71	-	-	-	-	-	-	-
	VoTr-SSD [16]	✗	87.86	78.27	76.93	-	-	-	-	-	-
	RDIoU (Ours)	✓	89.16	85.24	78.41	**63.26**	**57.47**	**52.53**	**83.32**	**68.39**	**63.63**
	RDIoU (Ours)	✗	89.76	**86.62**	**79.04**	-	-	-	-	-	-

Table 1. Performance comparisons with state-of-the-art methods on the KITTI *val* set with 11 recall positions. M :✓means training on three classes. M :✗means training only on car. Mod.* is the most important metric. The top-1 of two-stage and single-stage methods are bold, respectively.

The 3D backbone has four levels with feature dimensions $(16, 32, 64, 64)$, $(16, 32, 64, 128)$ for KITTI and Waymo Open Dataset, respectively. The 2D backbone contains two blocks, the first block is implemented by 5 CNN layers for KITTI and 6 CNN layers for Waymo Open Dataset to keep the same resolution with the output of 3D backbone. The second block is implemented by 1 CNN layer and 4 Transformer (using the setting of $4\times$ expansion in FFN layers and 4 attention heads) layers [15,28] with half the resolution. Finally, one fractionally-strided convolution layer is adopted to double the resolution.

All the evaluated models are trained from scratch in an end-to-end manner with the ADAM optimizer. The learning rate is decayed with a cosine annealing strategy, and the maximum is $3.5 \times e^{-4}$. We use a batch size of 32. For the hyperparameters of the overall loss, we set $\gamma_1 = 0.2$ and $\gamma_2 = 2$.

4.3   Results on Real-world Datasets

KITTI. KITTI is relatively small as compared to the Waymo Open Dataset. We follow the previous works to train models on *train* set, and report the results on *val* set. Furthermore, we report the detection results on KITTI *test* server by training the model with *train+val* set. All the evaluated models are reported in three difficulty levels (*i.e.*, *easy*, *moderate*, *hard*). Table 1 reports the results on the KITTI *val* set. For the most important 3D object detection metric on *moderate* level of car, our proposed RDIoU method surpasses the current best single-stage models CIA-SSD [37], SE-SSD [38], VoTr-SSD [16] with

Method	Reference	Stage	3D AP (IoU=0.7)		
			Easy	Moderate*	Hard
Part-A^2 [24]	TPAMI 2020	Two	87.81	78.49	73.51
STD [35]	ICCV 2019	Two	87.95	79.71	75.09
Point-GNN [25]	CVPR2020	Two	88.33	79.47	72.29
PV-RCNN [22]	CVPR 2020	Two	90.25	81.43	76.82
LiDAR-RCNN [13]	CVPR 2021	Two	85.97	74.21	69.18
VoTr-TSD [16]	ICCV 2021	Two	89.90	82.09	**79.14**
CT3D [21]	ICCV 2021	Two	87.83	81.77	77.16
BtcDet [30]	AAAI 2022	Two	**90.64**	**82.86**	78.09
VoxelNet [43]	CVPR 2018	Single	77.82	64.17	57.51
SECOND [31]	Sensors 2018	Single	83.34	72.55	65.82
PointPillar [10]	CVPR 2019	Single	82.58	74.31	68.99
3DSSD [34]	CVPR 2020	Single	88.36	79.57	74.55
SA-SSD [9]	CVPR 2020	Single	88.75	79.79	74.16
SE-SSD [38]	CVPR 2021	Single	**91.49**	**82.54**	77.15
RDIoU (Ours)	-	Single	90.65	82.30	**77.26**

Table 2. Performance comparisons with state-of-the-art methods for car detection on the KITTI *test* benchmark with 40 recall positions. The top-1 of two-stage and single-stage methods are bold, respectively.

+7.71%AP, +0.91%AP, +8.35%AP, respectively. Furthermore, our RDIoU even outperforms the best two-stage models VoTr-TSD [16], CT3D [21] BtcDet [30] with +2.58%AP, +0.56%AP, +0.05%AP, respectively. This superior performance strongly manifests the effectiveness of our proposed method.

For pedestrian and cyclist detection, we perform a three-class training for our RDIoU method and compare to the state-of-the-art methods with the same settings. It can be seen that our RDIoU leads a large margin compared with these single-stage methods (VoxelNet [43], SECOND [31] and PointPillar [10]). At the *moderate* level, our RDIoU outperforms SECOND with +4.49%AP and +1.23%AP on pedestrian and cyclist detection, respectively. This also proves that our RDIoU has the superior ability of detecting small objects. Note that the performances on small objects are marginally lower compared to the other two-stage methods such as Part-A^2 [24] and CT3D [21]. This is because single-stage methods usually assign more attention to the main objects (*i.e.,* car) compared to the two-stage methods which have the inherent advantage from the proposals.

Table 2 shows the car-3D detection results by submitting to the KITTI *test* server. Our RDIoU achieves 82.30%AP on the most important *moderate* level, surpassing the state-of-the-art methods of LiDAR-RCNN [13], VoTr-TSD [16], CT3D [21] by +8.09%AP, +0.21%AP, +0.53%AP, respectively. Note that these are all two-stage methods such that RDIoU has the absolute advantage of inference speed. Compared to SE-SSD [38] and BtcDet [30], our RDIoU is slightly lower by -0.24%AP and -0.56%AP, respectively. One possible reason is the mismatched data distributions between the KITTI *val* set and *test* set [13,24]. Moreover, SE-SSD requires extra data augmentation and complex self-ensembling

Method	Stage	3D AP/APH (IoU=0.7)				BEV AP/APH (IoU=0.7)			
		Overall	0-30m	30-50m	50m-Inf	Overall	0-30m	30-50m	50m-Inf
LEVEL_1									
MVF [42]	Two	62.9/-	86.3/-	60.0/-	36.0/-	80.4/-	93.6/-	79.2/-	63.1/-
PV-RCNN [22]	Two	70.3/69.7	91.9/91.3	69.2/68.5	42.2/41.3	80.0/82.1	97.4/96.7	83.0/82.0	65.0/63.2
Voxel-RCNN [4]	Two	75.6/-	92.5/-	74.1/-	53.2/-	88.2/-	97.6/-	87.3/-	77.7/-
LiDAR-RCNN [13]	Two	76.0/75.5	92.1/91.6	74.6/74.1	54.5/53.4	90.1/89.3	97.0/96.5	89.5/88.6	78.9/77.4
CenterPoint [36]	Two	76.7/76.2	-/-	-/-	-/-		-/-	-/-	-/-
VoTr-TSD [16]	Two	75.0/74.3	92.3/91.7	73.4/72.6	51.1/50.0	-/-	-/-	-/-	-/-
CT3D [21]	Two	76.3/-	92.5/-	75.1/-	**55.4/-**	**90.5/-**	**97.6/-**	88.1/-	78.9/-
BtcDet [30]	Two	**78.6/78.1**	**96.1/-**	**77.6/-**	54.5/-	-/-	-/-	-/-	-/-
PointPillar* [10]	Single	72.1/71.5	88.3/87.8	69.9/69.3	48.0/47.3	87.9/87.1	96.6/96.0	87.1/86.2	78.1/76.5
Pillar-OD [29]	Single	69.8/-	88.5/-	66.5/-	42.9/-	87.1/-	95.8/-	84.7/-	72.1/-
VoTr-SSD [16]	Single	69.0/68.4	88.2/87.6	66.7/66.1	42.1/41.4	-/-	-/-	-/-	-/-
RDIoU (Ours)	Single	**78.4/78.0**	**93.0/92.6**	**75.4/74.9**	**56.2/55.6**	**91.6/91.0**	**98.1/97.7**	**90.8/90.2**	**82.4/81.1**
LEVEL_2									
PV-RCNN [22]	Two	65.4/64.8	91.6/91.0	65.1/64.5	36.5/35.7	77.5/76.6	94.6/94.0	80.4/79.4	55.4/53.8
Voxel-RCNN [4]	Two	66.6/-	91.7/-	67.9/-	40.8/-	81.1/-	97.0/-	81.4/-	63.3/-
LiDAR-RCNN [13]	Two	68.3/67.9	91.3/90.9	68.5/68.0	42.4/41.8	81.7/81.0	94.3/93.9	**82.3/81.5**	65.8/64.5
CenterPoint [36]	Two	68.8/68.3	-/-	-/-	-/-	-/-	-/-	-/-	-/-
VoTr-TSD [16]	Two	65.9/65.3	-/-	-/-	-/-	-/-	-/-	-/-	-/-
CT3D [21]	Two	69.0/-	91.8/-	68.9/-	42.6/-	**81.7/-**	**97.1/-**	82.2/-	64.3/-
BtcDet [30]	Two	**70.1/69.6**	**96.0/-**	**70.1/-**	43.9/-	-/-	-/-	-/-	-/-
PointPillars* [10]	Single	63.6/63.1	87.4/86.9	62.9/62.3	37.2/36.7	81.3/80.4	94.0/93.5	81.7/80.8	65.5/64.1
VoTr-SSD [16]	Single	60.2/59.7	-/-	-/-	-/-	-/-	-/-	-/-	-/-
RDIoU (Ours)	Single	**69.5/69.1**	**92.3/91.9**	**69.3/68.9**	**43.7/43.1**	**83.1/82.5**	**97.5/97.1**	**85.2/84.6**	**68.3/67.2**

Table 3. Performance comparisons with state-of-the-art methods for the vehicle detection on the Waymo validation dataset. Here PointPillars* is implemented in mmdetection3D [3]. The top-1 of two-stage and single-stage methods are bold, respectively.

procedure while the two-stage method BtcDet requires a further box refinement module. Overall, the results on both *test* and *val* sets consistently reveal that RDIoU method is highly effective.

Waymo Open Dataset. All methods for comparison are evaluated with AP and average precision by heading (APH) at two difficulty levels defined in the official evaluation, where the *LEVEL_1* objects contain at least 5 points while the *LEVEL_2* objects contain 1∼5 inside points. The rotated IoU threshold is set to 0.7 for vehicle detection. We also report the detection results based on the different distances of the objects for adequate comparison. Table 3 reports the results on validation sequences of Waymo Open Dataset. It can be clearly seen that RDIoU achieves excellent performance on this large-scale, diverse, and challenging dataset. Our RDIoU beats almost all state-of-the-arts including the two-stage methods on all evaluation metrics by a large margin in a number of settings. Specifically, our RDIoU outperforms the latest state-of-the-art methods CenterPoint [36], VoTr-TSD [16], CT3D [21] with +1.7%, +3.4%, +2.1% on 3D AP of *LEVEL_1*, respectively, and achieves fairly close performance as compared to BtcDet [30] (only -0.2%AP performance drop). These inspired results further affirm the strong ability of our RDIoU method.

4.4   Ablation Studies

In this section, we detail the influence of each component that contributes to the final RDIoU design. All models are trained on the KITTI *train* set from scratch,

Method	3D$_{R11}$			3D$_{R40}$	FPS
	Easy	Mod.	Hard	Mod.	(Hz)
PointPillar [10]	87.08	77.74	76.24	79.88	33.8
PointPillar (+RDIoU)	**88.89**	**78.89**	**78.02**	**82.42**	33.8
SECOND [31]	88.78	78.74	77.51	82.85	30.5
SECOND (+RDIoU)	**89.24**	**86.10**	**78.60**	**85.80**	30.5
CT-stacked	88.93	78.91	77.63	83.01	26.6
CT-stacked (+RDIoU)	**89.76**	**86.62**	**79.04**	**86.20**	26.6

Table 4. Ablation study on different backbone networks. (+RDIoU) means replacing the classification and regression losses in baseline models with our proposed RDIoU.

Method	3D IoU guided	RDIoU guided	3D$_{R11}$			3D$_{R40}$
			Easy	Mod.	Hard	Mod.
Smooth$-\ell1$ loss [14]			89.19	78.86	77.53	83.14
IoU loss [41]	✓		88.80	81.94	77.67	83.60
		✓	**89.40**	**85.60**	**78.76**	**85.73**
CIoU loss [40]	✓		88.82	83.20	77.48	84.01
		✓	**89.43**	**86.21**	**78.90**	**85.98**
DIoU loss [40]	✓		88.83	83.59	77.93	84.20
		✓	**89.76**	**86.62**	**79.04**	**86.20**

Table 5. Ablation study on RDIoU-guided regression loss. All models utilize RDIoU-guided quality focal loss for classification.

and evaluated on the KITTI *val* set for fair comparison. We use CT-stacked backbone network as the default setting. The final performance is reported on 3D *moderate* level of car with 11 and 40 recall positions.

Effect of Different Backbone Networks. In Table 4, we plug RDIoU into another two commonly used 3D object detection backbone networks SECOND [31] and PointPillar [10]. It can be seen that RDIoU can help different backbone networks to achieve significant performance improvement. Specifically, plugging our RDIoU into PointPillar, SECOND and our CT-stacked backbone networks brings 3D$_{R40}$+2.54%, 3D$_{R40}$+2.95% and 3D$_{R40}$+3.19% improvements on *moderate* level, respectively. The contributing factor is that our RDIoU method can be integrated into any existing backbone networks to assist the model learning for better performance.

Effect of RDIoU-guided Regression Branch. As shown in Figure 5, although 3D IoU-based regression losses show better performance than the Smooth-$\ell1$ loss, our RDIoU-based regression losses can still further boost the performance by a large margin. Specifically, RDIoU-guided IoU loss [41], CIoU loss [40], DIoU loss [40] surpass the 3D IoU-guided losses with 3D$_{R40}$+2.13%, 3D$_{R40}$+1.97%, 3D$_{R40}$+2.00%, respectively. These gains in AP strongly prove the effectiveness of our proposed method, which encourages the network to learn more precise 3D bounding boxes. DIoU loss performs better than IoU loss due to the direct minimization for the normalized distance between predicted and ground truth

3D IoU guided QFL	RDIoU guided QFL	Easy	$3D_{R11}$ Mod.	Hard	$3D_{R40}$ Mod.
		89.16	84.78	78.15	85.12
✓		89.33	86.07	78.79	85.79
	✓	**89.76**	**86.62**	**79.04**	**86.20**

Table 6. Ablation study on RDIoU-guided classification loss. Traditional focal loss is adopted as the default setting. All models utilize RDIoU-guided DIoU loss.

bounding boxes. Note that the aspect ratio of the 3D real-world category is relatively stable, and thus CIoU loss does not bring performance improvement as compared to the DIoU loss.

Effect of RDIoU-guided Classification Branch. In Table 6, we investigate the influence of RDIoU-guided quality focal loss by replacing it with the traditional focal loss [14] and the 3D IoU-guided quality foal loss [12], respectively. As seen in 1^{st} and 2^{nd} rows of Table 6, the 3D IoU-guided quality focal loss performs better than the traditional focal loss due to the extra quality estimation for each predicted bounding box. 2^{nd} row and 3^{rd} row of Table 6 demonstrate that RDIoU-guided quality focal loss can also obviously exceed 3D IoU-guided quality focal loss ($3D_{R11}$+0.55%, $3D_{R40}$+0.41%). This significant improvement comes from the stable measurement of the predicted bounding boxes, which subsequently facilitates the generation of more reliable confidence maps.

5   Conclusion

This paper presents a novel RDIoU method to model the intersection of two arbitrary-oriented 3D bounding boxes for improving the single-stage point cloud detectors. RDIoU decouples the rotation variable as an independent term while preserving the geometry of 3D IoU. It exhibits robustness to the rotation change compared with the 3D IoU. Based on RDIoU, we propose the RDIoU-guided DIoU loss to enable stable optimization process and efficient back-propagation during training, and it achieves the best performance among the latest 3D regression losses. Moreover, we introduce the RDIoU-guided quality focal loss to address the misalignment problem between the regression results and the confidence maps, and it also exhibits much better than the 3D IoU-guided quality focal loss. The experimental results show that our RDIoU method can help the commonly used backbone networks to achieve state-of-the-art performance.

References

1. Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3d object detection network for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1907–1915 (2017)
2. Chen, Z., Chen, K., Lin, W., See, J., Yu, H., Ke, Y., Yang, C.: Piou loss: Towards accurate oriented object detection in complex environments. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 195–211 (2020)
3. Contributors, M.: MMDetection3D: OpenMMLab next-generation platform for general 3D object detection. https://github.com/open-mmlab/mmdetection3d (2020)
4. Deng, J., Shi, S., Li, P., Zhou, W., Zhang, Y., Li, H.: Voxel r-cnn: Towards high performance voxel-based 3d object detection. arXiv preprint arXiv:2012.15712 (2020)
5. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. The International Journal of Robotics Research **32**(11), 1231–1237 (2013)
6. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1440–1448 (2015)
7. Graham, B.: Sparse 3d convolutional neural networks. arXiv preprint arXiv:1505.02890 (2015)
8. Graham, B., van der Maaten, L.: Submanifold sparse convolutional networks. arXiv preprint arXiv:1706.01307 (2017)
9. He, C., Zeng, H., Huang, J., Hua, X.S., Zhang, L.: Structure aware single-stage 3d object detection from point cloud. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11873–11882 (2020)
10. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12697–12705 (2019)
11. Li, X., Wang, W., Hu, X., Li, J., Tang, J., Yang, J.: Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11632–11641 (2021)
12. Li, X., Wang, W., Wu, L., Chen, S., Hu, X., Li, J., Tang, J., Yang, J.: Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. Advances in Neural Information Processing Systems (NIPS) (2020)
13. Li, Z., Wang, F., Wang, N.: Lidar r-cnn: An efficient and universal 3d object detector. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7546–7555 (2021)
14. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2980–2988 (2017)
15. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021)
16. Mao, J., Xue, Y., Niu, M., Bai, H., Feng, J., Liang, X., Xu, H., Xu, C.: Voxel transformer for 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3164–3173 (2021)
17. Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum pointnets for 3d object detection from rgb-d data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 918–927 (2018)

18. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 652–660 (2017)
19. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. Advances in Neural Information Processing Systems (NIPS) **30**, 5099–5108 (2017)
20. Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 658–666 (2019)
21. Sheng, H., Cai, S., Liu, Y., Deng, B., Huang, J., Hua, X.S., Zhao, M.J.: Improving 3d object detection with channel-wise transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2743–2752 (2021)
22. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10529–10538 (2020)
23. Shi, S., Wang, X., Li, H.: Pointrcnn: 3d object proposal generation and detection from point cloud. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–779 (2019)
24. Shi, S., Wang, Z., Shi, J., Wang, X., Li, H.: From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2020)
25. Shi, W., Rajkumar, R.: Point-gnn: Graph neural network for 3d object detection in a point cloud. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1711–1719 (2020)
26. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2446–2454 (2020)
27. Team, O.D.: Openpcdet: An open-source toolbox for 3d object detection from point clouds. https://github.com/open-mmlab/OpenPCDet (2020)
28. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems (NIPS). pp. 5998–6008 (2017)
29. Wang, Y., Fathi, A., Kundu, A., Ross, D.A., Pantofaru, C., Funkhouser, T., Solomon, J.: Pillar-based object detection for autonomous driving. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 18–34 (2020)
30. Xu, Q., Zhong, Y., Neumann, U.: Behind the curtain: Learning occluded shapes for 3d object detection. arXiv preprint arXiv:2112.02205 (2021)
31. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors **18**(10), 3337 (2018)
32. Yang, X., Yan, J., Ming, Q., Wang, W., Zhang, X., Tian, Q.: Rethinking rotated object detection with gaussian wasserstein distance loss. arXiv preprint arXiv:2101.11952 (2021)
33. Yang, X., Yang, J., Yan, J., Zhang, Y., Zhang, T., Guo, Z., Sun, X., Fu, K.: Scrdet: Towards more robust detection for small, cluttered and rotated objects. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8232–8241 (2019)

34. Yang, Z., Sun, Y., Liu, S., Jia, J.: 3dssd: Point-based 3d single stage object detector. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11040–11048 (2020)
35. Yang, Z., Sun, Y., Liu, S., Shen, X., Jia, J.: Std: Sparse-to-dense 3d object detector for point cloud. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1951–1960 (2019)
36. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11784–11793 (2021)
37. Zheng, W., Tang, W., Chen, S., Jiang, L., Fu, C.W.: Cia-ssd: Confident iou-aware single-stage object detector from point cloud. arXiv preprint arXiv:2012.03015 (2020)
38. Zheng, W., Tang, W., Jiang, L., Fu, C.W.: Se-ssd: Self-ensembling single-stage object detector from point cloud. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14494–14503 (2021)
39. Zheng, Y., Zhang, D., Xie, S., Lu, J., Zhou, J.: Rotation-robust intersection over union for 3d object detection. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 464–480 (2020)
40. Zheng, Z., Wang, P., Liu, W., Li, J., Ye, R., Ren, D.: Distance-iou loss: Faster and better learning for bounding box regression. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). vol. 34, pp. 12993–13000 (2020)
41. Zhou, D., Fang, J., Song, X., Guan, C., Yin, J., Dai, Y., Yang, R.: Iou loss for 2d/3d object detection. In: International Conference on 3D Vision (3DV). pp. 85–94. IEEE (2019)
42. Zhou, Y., Sun, P., Zhang, Y., Anguelov, D., Gao, J., Ouyang, T., Guo, J., Ngiam, J., Vasudevan, V.: End-to-end multi-view fusion for 3d object detection in lidar point clouds. In: Conference on Robot Learning (CoRL). pp. 923–932 (2020)
43. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4490–4499 (2018)

A   Appendix

Some supplementary materials are provided to further validate the proposed approach. The supplementary materials include ablation studies on parameter k, more comparisons to existing ODIoU loss [38] and 3D IoU loss [41], inserting RDIoU into state-of-the-art 3D detectors, and visualization results, respectively.

Effect of parameter k. As shown in Table 7, we try different settings of the edge k to control the weight of rotation change. Here, we can see that k is an important factor to boost the final performance, and $k = 1.0$ is the best.

k	0.6	0.8	**1.0**	1.2	1.4	1.6	1.8
$3D_{R40}$	85.89	85.96	**86.20**	85.92	85.85	85.70	85.59

Table 7. Ablation study on parameter k, which reflects the sensitivity of RDIoU value to the rotation difference.

More comparison with 3D IoU [41] and ODIoU [38]. In Table 8, we provide more comparison results for 3D IoU and ODIoU with our proposed RDIoU. It can be seen that RDIoU outperforms both 3D IoU and ODIoU by large margins.

Method	$3D_{R11}$			$3D_{R40}$
	Easy	Mod.	Hard	Mod.
PointPillar[10]	87.08	77.74	76.24	79.88
PointPillar (+3D IoU)	86.91	77.93	76.83	80.12
PointPillar (+ODIoU)	87.15	78.29	77.08	80.58
PointPillar (+RDIoU)	**88.89**	**78.89**	**78.02**	**82.42**
SECOND[31]	88.78	78.74	77.51	82.85
SECOND (+3D IoU)	88.04	80.97	77.06	83.02
SECOND (+ODIoU)	88.69	82.82	77.41	83.88
SECOND (+RDIoU)	**89.24**	**86.10**	**78.60**	**85.80**
CT-stacked	88.93	78.91	77.63	83.01
CT-stacked (+3D IoU)	88.23	81.09	77.16	83.29
CT-stacked (+ODIoU)	88.70	82.89	77.53	84.01
CT-stacked (+RDIoU)	**89.76**	**86.62**	**79.04**	**86.20**

Table 8. Comparisons to 3D IoU and ODIoU.

Inserting RDIoU into state-of-the-art methods. We reproduce SOTA single-stage methods: CIA-SSD[37], SA-SSD[9], SE-SSD[38], and two-stage meth-

ods: PV-RCNN[22], Voxel-RCNN[4], CT3D[21] in Table 9. Our RDIoU can significantly improve their performance, especially for the single-stage detectors.

Type	Method	Easy	$3D_{R11}$ Mod.	Hard	$3D_{R40}$ Mod.
Single-stage	CIA-SSD[37]	**89.48**	78.54	77.35	81.93
	CIA-SSD (+RDIoU)	89.07	**85.44**	**78.55**	**85.23**
	SA-SSD[9]	89.26	79.28	78.35	82.65
	SA-SSD (+RDIoU)	**89.56**	**86.04**	**78.76**	**85.87**
	SE-SSD[38]	89.07	79.22	78.37	82.48
	SE-SSD (+RDIoU)	**89.24**	**85.98**	**78.60**	**85.24**
Two-stage	PV-RCNN[22]	89.31	84.49	78.78	84.93
	PV-RCNN (+RDIoU)	**89.47**	**86.21**	**79.01**	**85.96**
	Voxel-RCNN[4]	89.44	84.45	78.90	85.24
	Voxel-RCNN (+RDIoU)	**89.67**	**86.12**	**78.91**	**85.92**
	CT3D[21]	**89.54**	86.06	**78.99**	85.82
	CT3D (+RDIoU)	89.31	**86.27**	78.90	**85.94**

Table 9. Inserting our RDIoU into SOTA methods based on OpenPCDet[27].

Visualization. We provide some qualitative detection results on KITTI *test* set in Figure 5, it shows that RDIoU can produce high-quality 3D bounding boxes in diversified scenarios.

Fig. 5. Snapshots of qualitative results on KITTI test set. The output 3D bounding boxes are shown in green.