

TANet: Robust 3D Object Detection from Point Clouds with Triple Attention

Zhe Liu,¹ Xin Zhao,² Tengteng Huang,¹ Ruolan Hu,¹ Yu Zhou,¹ Xiang Bai^{1*}

¹Huazhong University of Science and Technology, Wuhan, China, 430074

²Institute of Automation, Chinese Academy of Sciences, Beijing, China, 100190

{m201772494, huangtengtng, huruolan, yuzhou, xbai}@hust.edu.cn, xzhao@nlpr.ia.ac.cn

Abstract

In this paper, we focus on exploring the robustness of the 3D object detection in point clouds, which has been rarely discussed in existing approaches. We observe two crucial phenomena: 1) the detection accuracy of the hard objects, e.g., Pedestrians, is unsatisfactory, 2) when adding additional noise points, the performance of existing approaches decreases rapidly. To alleviate these problems, a novel TANet is introduced in this paper, which mainly contains a Triple Attention (TA) module, and a Coarse-to-Fine Regression (CFR) module. By considering the channel-wise, point-wise and voxel-wise attention jointly, the TA module enhances the crucial information of the target while suppresses the unstable cloud points. Besides, the novel stacked TA further exploits the multi-level feature attention. In addition, the CFR module boosts the accuracy of localization without excessive computation cost. Experimental results on the validation set of KITTI dataset demonstrate that, in the challenging noisy cases, i.e., adding additional random noisy points around each object, the presented approach goes far beyond state-of-the-art approaches. Furthermore, for the 3D object detection task of the KITTI benchmark, our approach ranks the first place on Pedestrian class, by using the point clouds as the only input. The running speed is around 29 frames per second.

Introduction

3D object detection in point clouds has a large number of applications in real scenes, especially for autonomous driving and augmented reality. On the one hand, point clouds provide reliable geometric structure information and precise depth, while how to utilize such information effectively is an essential issue. On the other hand, point clouds are usually unordered, sparse, and unevenly distributed, which poses great challenges for accurate object detection.

In recent years, several approaches based on point clouds have been proposed in the 3D object detection community. PointRCNN (Shi, Wang, and Li 2019) directly operates on the raw point clouds, extracts the features by PointNets (Qi et al. 2017a; 2017b) and then estimates final results by two-stage detection networks. VoxelNet (Zhou and Tuzel 2018),

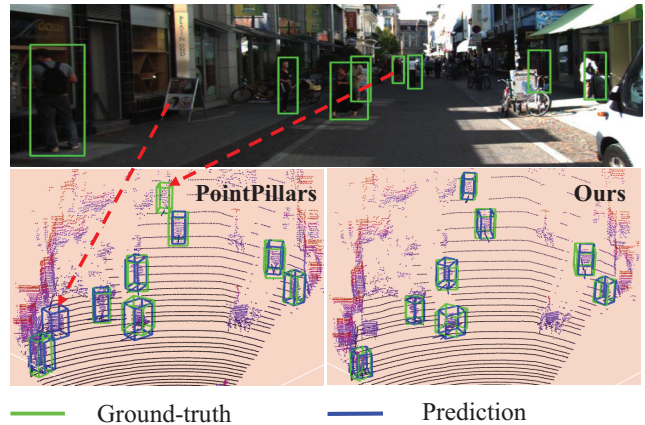


Figure 1: Detection results for Pedestrians. The first row shows the corresponding 2D image. The second row demonstrates the 3D detection results produced by PointPillars and our method, respectively. We highlight the missed and false detection in PointPillars with red arrows.

SECOND (Yan, Mao, and Li 2018) and PointPillars (Lang et al. 2019) convert the point clouds to the regular voxel grid and apply a series of convolutional operations to 3D object detection.

Although existing approaches have reported promising detection accuracy, the performance is still unsatisfactory in challenging cases, especially hard objects such as pedestrians. As shown in Fig. 1, PointPillars (Lang et al. 2018) misses a pedestrian and provides a false positive object. We reveal the intrinsic reasons from two aspects, 1) pedestrians have a smaller scale than cars, which makes fewer valid points scanned on them through LiDAR. 2) pedestrians frequently appear in a variety of scenes, so various background objects, such as trees, bushes, poles, *etc.*, might be close to the pedestrians, which results in enormous difficulty to recognize them correctly. Hence, detecting objects in the complex point clouds is still an extremely challenging task.

In this paper, we present a novel architecture named Triple Attention Network (TANet), as shown in Fig 2. The straightforward motivation is that, in the cases of severe noises, a

*Corresponding author

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

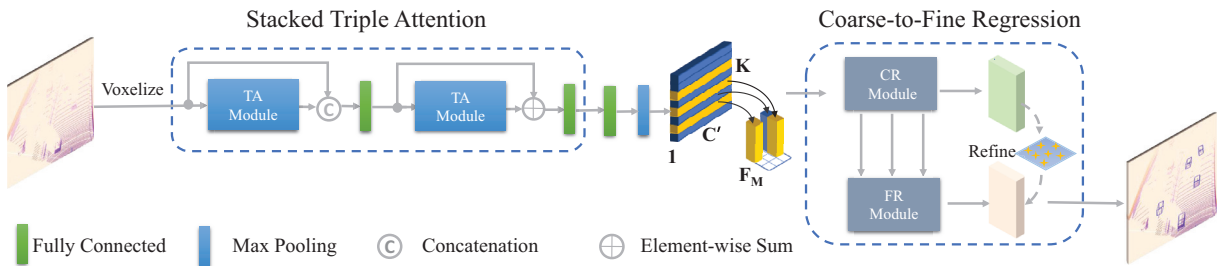


Figure 2: The full pipeline of TANet. First, we equally divide the point clouds into a voxel grid consisting of a set of voxels. Then, the stacked triple attention separately process each voxel to obtain a more discriminative representation. Subsequently, a compact feature representation for each voxel is extracted by aggregating the points inside it in a max-pooling manner. We arrange the voxel feature according to its original spatial position in the grid and thus lead to a feature representation for the voxel grid in the shape of $C' \times H \times W$. Finally, the coarse-to-fine regression is employed to generate the final 3D bounding boxes.

set of informative points can supply sufficient cues for the subsequent regression. In order to capture such informative cues, a TA module is introduced to enhance the discriminative points and suppress the unstable points. Specifically, the point-wise attention, channel-wise attention are learned respectively, and they are combined by element-wise multiplication. Besides, we also consider the voxel-wise attention, which represents the global attention of a voxel.

In the noisy cases, only applying a single regressor, e.g., one stage RPN, for 3D bounding box localization is unsatisfactory. To address this issue, we introduce an end-to-end trainable coarse-to-fine regression (CFR) mechanism. The coarse step provides a rough estimation of the object following (Zhou and Tuzel 2018; Lang et al. 2019). Then we present a novel Pyramid Sampling Aggregation (PSA) fusion approach, which supplies cross-layer feature maps. The refinement is implemented upon the fused cross-layer feature map to obtain the finer estimation.

Both the TA module and CFR mechanism are crucial for the robustness of 3D detectors, which is very important for the real scenario of automatic driving. Since not all the data in KITTI dataset is troubled by noises, in the experimental evaluation, we simulate the noisy cases by adding random points around each object. Extensive experiments demonstrate that our approach greatly outperforms the state-of-the-art approaches. Besides, our approach achieves the best results on the Pedestrian class of the KITTI benchmark, which further verifies the robustness of the presented detector.

In summary, the key contributions of the proposed method lie in:

- 1) We introduce a novel Triple Attention module, which takes the channel-wise, point-wise, and voxel-wise attention into consideration jointly, then the stacked operation is performed to obtain the multi-level feature attention, and hence the discriminative representation of the object is acquired.

- 2) We propose a novel coarse-to-fine regression, based on the coarse regression results, the fine regression is performed on the informative fused cross-layer feature map.

- 3) Our approach achieves convincing experimental results in the challenging noisy cases, and the quantitative comparisons on the KITTI benchmark illustrate that our method

achieves state-of-the-art performance and promising inference speed.

Related Work

With the rapid development of computer vision, much effort has been devoted to detect 3D objects from multi-view images, which can be roughly classified into two categories: the front view based approaches (Song and Chandraker 2015; Chen et al. 2016; Mousavian et al. 2017), and the bird’s eye view based approaches (Chen et al. 2017; Ku et al. 2018; Yang, Luo, and Urtasun 2018; Simon et al. 2018; Li, Zhang, and Xia 2016; Yang, Liang, and Urtasun 2018). However, it is difficult for these methods to localize the objects accurately due to the loss of depth information.

Recently, the research trend has gradually shifted from the RGB image to point cloud data. Detecting 3D objects based on the voxel grid has been widely concerned (Engelcke et al. 2017; Wang and Posner 2015). In these approaches, 3D point cloud space is divided into voxels equally, and only the non-empty voxels are encoded for computational efficiency. VoxelNet (Zhou and Tuzel 2018), SECOND (Yan, Mao, and Li 2018) and PointPillars (Lang et al. 2019) convert point clouds to a regular voxel grid and learn the representation of each voxel with the Voxel Feature Encoding (VFE) layer. Then, the 3D bounding boxes are computed by a region proposal network based on the learned voxel representation. In contrast to PointPillars, we focus on leveraging channel-wise, point-wise, and voxel-wise attention of point clouds to learn a more discriminative and robust representation for each voxel. To our best knowledge, the proposed method is the first one to design the attention mechanism suitable for the 3D object detection task. In addition, our method utilizes stacked Triple Attention (TA) modules to capture the multi-level feature attention.

In the 3D object detection task, the voxel grids based approaches, e.g. Voxelnet, PointPillars and SECOND, frequently adopt one-stage detection network (Liu et al. 2016; Fu et al. 2017; Shen et al. 2017), which can process more than 20 frames per second. In contrast, the raw point clouds based approaches, e.g., PointRCNN (Shi, Wang, and Li 2019), utilize two-stage architecture (Girshick 2015; Ren et

al. 2015; Lin et al. 2017a; He et al. 2017). Although these methods lead to better detection performance, they run at relatively slower speeds with less than 15 frames per second. Motivated by RefineDet (Zhang et al. 2018), we propose an end-to-end trainable coarse-to-fine regression mechanism. Our goal is to seek detection methods that can achieve a better trade-off better accuracy and efficiency.

3D Object Detection with TANet

In this section, we introduce TANet for 3D object detection based on voxels, which is an end-to-end trainable neural network. As shown in Fig. 2, our network architecture consists of two main parts: the Stacked Triple Attention and the Coarse-to-Fine Regression.

Before introducing the technical details, we present several basic definitions of 3D object detection. A point set in 3D space is defined as $\mathbf{P} = \{\mathbf{p}_i = [x_i, y_i, z_i, r_i]^T \in \mathbb{R}\}_{i=1,2,\dots,M}$, where x_i, y_i, z_i denote the coordinate values of each point along the axes X, Y, Z, r_i is the laser reflection intensity that can be treated as an extra feature, and M is the total number of points. Given an object in 3D space, it is represented by a 3D bounding box $(c_x, c_y, c_z, h, w, l, \theta)$, including its center c_x, c_y, c_z , size h, w, l , and orientation θ that indicates the heading angle around the up-axis.

Stacked Triple Attention

We suppose that a point cloud \mathbf{P} in the 3D space has the range of W^*, H^*, D^* along the X, Y, Z axes respectively. \mathbf{P} is equally divided into a specific voxel grid. Each voxel has the size of v_W^*, v_H^*, v_D^* . Consequently, the voxel grid is of size $W = W^*/v_W^*, H = H^*/v_H^*, D = D^*/v_D^*$. Note that D is always 1, since the voxel grid is not discretized along Z axis in the implementation.

As the point clouds are often sparse and unevenly distributed, a voxel has a variable number of points. Let N and C denote the maximum number of points of each voxel and the channel number of each point feature, respectively. A voxel grid \mathbf{V} that consists of K voxels can be defined as $\mathbf{V} = \{V^1, \dots, V^K\}$, where $V^k \in \mathbb{R}^{N \times C}$ indicates the k -th voxel of \mathbf{V} .

Point-wise Attention. Given a voxel V^k , we perform max-pooling to aggregate point features across the channel-wise dimensions, resulting in its point-wise responses $E^k \in \mathbb{R}^{N \times 1}$. To explore the spatial correlation of points, following the *Excitation* operation (Hu, Shen, and Sun 2018), two fully-connected layers are employed to encode the global responses, i.e.:

$$S^k = W_2 \delta(W_1 E^k), \quad (1)$$

where $W_1 \in \mathbb{R}^{r \times N}$, $W_2 \in \mathbb{R}^{N \times r}$ are the weight parameters of two fully-connected layers, respectively. δ is the ReLU activation function. $S^k \in \mathbb{R}^{N \times 1}$ is the point-wise attention of V^k . As shown in Fig. 3, the upper branch of the attention module is designed for describing the spatial correlation among the points inside each voxel.

Channel-wise Attention. Similar to the strategy for estimating the point-wise attention for a voxel, we compute the channel-wise attention with the middle branch of the attention module as shown in Fig. 3. A max-pooling operation

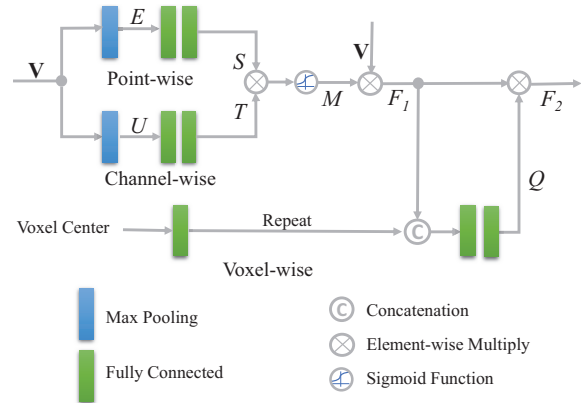


Figure 3: The architecture of TA module.

is performed to aggregate the channel features across their point-wise dimensions, which obtains the channel-wise responses of a voxel V^k to be $U^k \in \mathbb{R}^{1 \times C}$. Then, we compute $T^k = W_2 \delta(W_1 (U^k)^T)$. $W_1 \in \mathbb{R}^{r \times C}$ and $W_2 \in \mathbb{R}^{C \times r}$ represent the importance of feature channels for each voxel.

Given the k -th voxel V^k , we obtain the attention matrix $M^k \in \mathbb{R}^{N \times C}$ that combines spatial-wise attention S^k and channel-wise attention T^k through the element-wise multiply, i.e. :

$$M^k = \sigma(S^k \times T^k) \quad (2)$$

where σ denotes the *sigmoid* function, which is employed to normalize the values of the attention matrix to the range of $[0, 1]$. Thus, a feature representation $F_1^k = M^k \odot V^k \in \mathbb{R}^{N \times C}$ can be obtained, which properly weights the importance of all the points inside a voxel across the point-wise and channel-wise dimensions.

Voxel-wise Attention. The voxel-wise attention is further employed to judge the importance of the voxels. We first average the coordinates of all points inside each voxel as the voxel center, which can provide accurate location information. Then the voxel center is transformed into a higher dimension through a fully-connected layer, and it is combined with F_1^k in a concatenation fashion. Voxel-wise attention weight is defined as $Q = [q^1, \dots, q^k, \dots, q^K] \in \mathbb{R}^{K \times 1 \times 1}$, where q^k is obtained by compressing the point-wise and channel-wise dimensions to 1 via two fully-connected layers, respectively. Finally, a more robust and discriminative voxel feature is obtained by $F_2^k = q^k \cdot F_1^k$.

Through all the above operations, the feature representation F_2^k enhances the crucial features, which contributes significantly to our tasks while suppresses the irrelevant and noising features. For simplicity, we name the module integrating these three types of attention as Triple Attention (TA).

Stacked TA. As shown in Fig. 2, in our approach, two TA modules are stacked to exploit the multi-level feature attention. The first one directly operates on the original features of the point clouds, while the second one works on the higher dimensional features. For each TA module, we concatenate/sum its output with its input to fuse more feature

information. Then the higher-dimensional feature representation is obtained via a fully-connected layer. Finally, a max-pooling operation is used to aggregate all the point features of each voxel, which is treated the input of the CFR.

Coarse-to-Fine Regression

We employ a Coarse Regression (CR) module and a Fine Regression (FR) module for 3D box estimation. The details of these two modules are presented in the following.

The CR module uses a similar architecture with (Zhou and Tuzel 2018; Yan, Mao, and Li 2018; Lang et al. 2018). To be specific, as shown in the top module of Fig. 4, the output of Block1, Block2 and Block3 is denoted as B_1 , B_2 and B_3 , and the shapes of these blocks are $(C', H/2, W/2)$, $(2C', H/4, W/4)$ and $(4C', H/8, W/8)$, respectively. The CR module generates a feature map F_C with the size of $(6C', H/2, W/2)$, then follows the classification and regression branches, which provides the coarse boxes for FR module.

Based on the output of the CR module, the Pyramid Sampling Aggregation (PSA) module is leveraged to provide cross-layer feature maps, which is shown in the bottom module of Fig. 4. The high-level features supply larger receptive fields and richer semantic information, while the low-level features have a larger resolution. Thus, the cross-layer feature maps effectively capture multi-level information, leading to a more comprehensive and robust feature representation for objects. Specifically, a feature pyramid $\{B_1^1, B_1^2, B_1^3\}$ is achieved based upon B_1 , where B_1^1 is equivalent to B_1 . B_1^2 and B_1^3 are obtained by two down-sampling operations performing on B_1 , respectively. B_1^2 has the same size as B_2 , and B_1^3 has the same size as B_3 . Similarly, an up-sampling and a down-sampling are operated on B_2 to obtain $\{B_2^1, B_2^2, B_2^3\}$. In addition, $\{B_3^1, B_3^2, B_3^3\}$ is obtained by two up-sampling operations based on B_3 .

To make full use of the cross-layer features, we concatenate B_i^1, B_i^2 and B_i^3 for $i = 1, 2, 3$, respectively. Then a series of convolution operations followed by an up-sampling layer are executed, which results in the feature maps $UP = \{UP_1, UP_2, UP_3\}$, they have the same shape of $(2C' \times H/2 \times W/2)$.

In addition, the enriched feature maps from the PSA module are combined with the semantic information from CR module. Specifically, a 1×1 convolution is employed to transform F_C into F_B , F_B has the same dimension with feature maps in UP . Then each feature map in UP is combined with F_B by element-wise addition. A 3×3 convolution layer is performed on each fused feature map. F_R is obtained by concatenating the modulated features in UP , which serves as the feature map for Fine Regression. The regression branch of FR regards the coarse boxes of CR as the new anchors to regress the 3D bounding box, and perform the classification.

Loss Function

The multi-task loss function is employed for jointly optimizing the CR module and the FR module. The offsets of the bounding box regression between a prior anchor a and

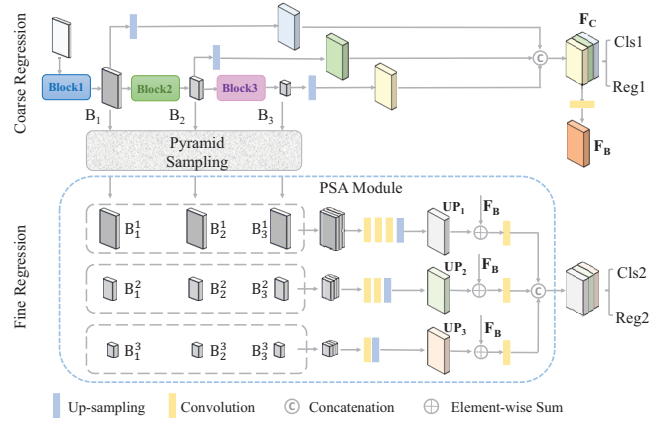


Figure 4: The architecture of Coarse-to-Fine Regression. The Pyramid Sampling indicates a series of down-sampling and up-sampling operations, which can be achieved via the pooling and the transposition convolution.

the ground-truth box g can be computed as:

$$\begin{aligned} \Delta_x^g &= \frac{x_g - x_a}{d_a}, \Delta_y^g = \frac{y_g - y_a}{d_a}, \Delta_z^g = \frac{z_g - z_a}{h_a} \\ \Delta_w^g &= \log\left(\frac{w_g}{w_a}\right), \Delta_l^g = \log\left(\frac{l_g}{l_a}\right), \Delta_h^g = \log\left(\frac{h_g}{h_a}\right) \\ \Delta_\theta^g &= \theta_g - \theta_a, \end{aligned} \quad (3)$$

where $d_a = \sqrt{(w_a)^2 + (l_a)^2}$. For simplicity, the residual vector $\Delta_g = (\Delta_x^g, \Delta_y^g, \Delta_z^g, \Delta_w^g, \Delta_l^g, \Delta_h^g, \Delta_\theta^g)$ is defined as the regression ground-truth. Similarly, $\Delta_p = (\Delta_x^p, \Delta_y^p, \Delta_z^p, \Delta_w^p, \Delta_l^p, \Delta_h^p, \Delta_\theta^p)$ indicates the offsets between the prior anchors and the predicted boxes. SmoothL1 (Girshick 2015) is used as our 3D bounding box regression loss L_{reg} . Besides, angle loss is employed for better restricting the orientation of 3D bounding box. Note that when the orientation angle of a 3D object is shifted $\pm\pi$ radians, it does not change the estimation of localization. Hence, the *sine* function is introduced to encode the loss of the orientation angle θ following (Yan, Mao, and Li 2018). Considering that the number of positive samples and negative samples is imbalanced, the Focal Loss (Lin et al. 2017b) is adopted as the classification loss L_{cls} . The superscript C and R represent the CR module and the FR module, respectively. It should be noted that the FR module leverages the coarse bounding boxes as the new anchor boxes, which is different from the CR module that utilizes the prior anchor boxes. The total loss function can be defined as:

$$\begin{aligned} L_{total} &= \alpha L_{cls}^C + \beta \frac{1}{N_{pos}^C} \sum L_{reg}^C(\Delta_p^C, \Delta_g^C) + \\ &\lambda \{ \alpha L_{cls}^R + \beta \frac{1}{N_{pos}^R} \sum L_{reg}^R(\Delta_p^R, \Delta_g^R) \} \end{aligned} \quad (4)$$

where N_{pos}^C and N_{pos}^R stand for the numbers of positive anchors in the CR module and FR module, respectively. α and β represent the balance weights for the classification loss and the regression loss, respectively. λ is used to balance the weight for the CR module and FR module.

Method	Number of noise points	Cars				Pedestrians				Cyclists			
		Easy	Moderate	Hard	3D mAP	Easy	Moderate	Hard	3D mAP	Easy	Moderate	Hard	3D mAP
PointRCNN (2019)	0	88.26	77.73	76.67	80.89	65.62	58.57	51.48	58.56	82.76	62.83	59.62	68.40
PointPillars (2019)	0	87.50	77.01	74.77	79.76	66.73	61.06	56.50	61.43	83.65	63.40	59.71	68.92
Ours	0	88.21	77.85	75.62	80.56	70.8	63.45	58.22	64.16	85.98	64.95	60.40	70.44
PointRCNN (2019)	20	88.24	76.95	74.73	79.97	62.00	56.17	49.52	55.90	81.55	61.98	57.20	66.91
PointPillars (2019)	20	87.21	76.74	74.54	79.50	64.44	59.02	55.00	59.49	82.66	62.52	58.23	67.80
Ours	20	88.17	77.68	75.31	80.39	69.98	62.70	57.65	63.44	85.55	64.06	60.03	69.88
PointRCNN (2019)	50	87.99	76.66	74.16	79.60	58.12	51.23	45.30	51.55	79.49	60.63	56.3	65.47
PointPillars (2019)	50	87.07	76.60	69.05	77.57	62.75	57.32	52.25	57.44	81.98	61.15	56.53	66.55
Ours	50	87.97	77.29	74.4	79.89	69.37	62.5	56.54	62.80	84.85	63.54	58.48	68.96
PointRCNN (2019)	100	87.56	75.98	69.34	77.63	55.57	48.35	42.88	48.93	76.77	56.66	52.92	62.12
PointPillars (2019)	100	86.62	76.06	68.91	77.20	60.31	55.17	49.65	55.04	80.97	58.02	54.6	64.53
Ours	100	87.52	76.64	73.86	79.34	67.30	60.77	54.45	60.84	84.53	61.64	57.44	67.87

Table 1: Performance comparison with PointRCNN and PointPillars for 3D object detection task on the KITTI validation set for *Cars*, *Pedestrians* and *Cyclists*. 3D mAP represents the mean average precision of each category.

Experiments

Experimental Dataset and Evaluation Metrics

All of the experiments are conducted on the KITTI dataset (Geiger, Lenz, and Urtasun), which contains 7481 training samples and 7518 test samples. Since the access to the ground truth for the test set is not available, we follow (Qi et al. 2018; Chen et al. 2017) to split the training samples into a training set consisting of 3712 samples and a validation set consisting of 3769 samples. When evaluating the performance on the test set, the training samples are re-split into the training and validation set according to the ratio of around 5:1. Our results are reported on both the KITTI validation set and test set for *Cars*, *Pedestrians* and *Cyclists* categories. For each category, three difficulty levels are involved (*Easy*, *Moderate* and *Hard*), which depend on the size, occlusion level and truncation of 3D objects.

The mean Average Precision (mAP) is utilized as our evaluation metric. For fair comparisons, we adopt the official evaluation protocol. Concretely, the IoU threshold is set to 0.7 for *Cars* and the IoU threshold to 0.5 for *Pedestrians* and *Cyclists*.

Implementation Details

For data augmentation, the points inside a ground-truth 3D bounding box along the Z-axis are rotated following the uniform distribution of $[-\pi/4, \pi/4]$ for orientation varieties. Besides, we further randomly flip the point clouds in 3D boxes along the X-axis. Random scaling in the range of $[0.95, 1.05]$ is also applied. Ground-truth boxes are randomly sampled and placed into raw samples to simulate the scenes crowded with multiple objects (Yan, Mao, and Li 2018).

Adaptive Moment Estimation (Adam) (Kingma and Ba 2015) is used for optimization with the learning rate of 0.0002. And our model is trained for about 200 epochs with a mini-batch size of 2. In all of our experiments, we randomly sample fixed N points for voxels containing more than N points. For the voxels containing less than N points, we simply pad them with zeros. In our settings, a large value of N is selected to be 100 for capturing sufficient cues to explore the spatial relationships. The dimension of the feature map for each voxel \mathbf{F}_M is 64 (e.g., $C' = 64$). All of our experiments are evaluated on a single Titan V GPU card. In our experiments, we set α , β , and λ to 1.0, 2.0 and 2.0 for total

loss, respectively. For more implementation details, please refer to our supplemental material.

Evaluation on KITTI dataset

In this part, our method is compared with the state-of-the-art approaches on the KITTI dataset using 1) noising point cloud data and 2) original point cloud data. For each task, three categories are involved (*Cars*, *Pedestrians* and *Cyclists*). The comprehensive experimental results are reported for the three categories under the three difficulty levels (*Easy*, *Moderate* and *Hard*).

Results on noising point cloud data. Extra challenge is introduced for detection by adding some noise points to each object. We think it is a relatively reasonable to sample these noise points that are closer to real scenes, since farther noising points will not interfere with the detection of a certain object. Specifically, the x , y and z coordinates of these noise points obey the uniform distribution of $[-l/2 + c_x, -3l + c_x] \cup [l/2 + c_x, 3l + c_x]$, $[-h/2 + c_y, -3h + c_y] \cup [h/2 + c_y, 3h + c_y]$ and $[-w/2 + c_z, -3w + c_z] \cup [w/2 + c_z, 3w + c_z]$, respectively. All models are trained with the official training data while tested on the noising validation point cloud data to evaluate their robustness for noises.

Quantitative results with state-of-the-art methods are presented in Table. 1. Although PointRCNN (Shi, Wang, and Li 2019) outperforms our method by 0.43% in terms of 3D mAP for *Cars*, our method shows its superior robustness for noises. With 100 noise points added, our method yields a 3D mAP of 79.34%, outperforming PointRCNN by 1.7%. For *Pedestrians*, our method achieves an improvement of 5.8% and 11.9% comparing with PointPillars (Lang et al. 2019) and PointRCNN. It can be observed that our method demonstrates great robustness for noises, especially for hard examples, e.g., *Pedestrians*, hard *Cyclists* and hard *Cars*.

On the whole, voxel-based methods (e.g., PointPillars and our framework) are more robust to noising points compared with the method based on the raw point clouds (e.g., PointRCNN). The main reason is that PointRCNN is a two-stage method that first optimizes the Region Proposal Network (RPN) separately and then optimizes the refinement network (i.e., RCNN) while fixing the parameters of RPN. For the contrast, our method is a coarse-to-fine detection network that can be end-to-end trainable, which is more robust

Method	Cars			Pedestrians			Cyclists			Modality	3D mAP (%)
	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard		
MV3D (2017)	71.09	62.35	55.12	-	-	-	-	-	-	LiDAR & Image	-
UberATG-ContFuse (2018)	82.54	66.22	64.04	-	-	-	-	-	-	LiDAR & Image	-
PC-CNN-V2 (2018)	84.33	73.80	64.83	-	-	-	-	-	-	LiDAR & Image	-
AVOD (2018)	73.59	65.78	58.38	38.28	31.51	26.98	60.11	44.90	38.80	LiDAR & Image	48.70
AVOD-FPN (2018)	81.94	71.88	66.38	50.80	42.81	40.88	64.00	52.18	46.61	LiDAR & Image	57.50
F-Pointnet (2018)	81.20	70.39	62.19	51.21	44.89	40.23	71.96	56.77	50.39	LiDAR & Image	58.80
MV3D (LiDAR) (2017)	66.77	52.73	51.31	-	-	-	-	-	-	LiDAR	-
Voxelnet (2018)	77.47	65.11	57.73	39.48	33.69	31.51	61.22	48.36	44.37	LiDAR	50.99
SECOND (2018)	83.13	73.66	66.20	51.07	42.56	37.29	70.51	53.85	46.90	LiDAR	58.35
PointPillars (2019)	79.05	74.99	68.30	52.08	43.53	41.49	75.78	59.07	52.92	LiDAR	60.80
PointRCNN (2019)	85.94	75.76	68.32	49.43	41.78	38.63	73.93	59.60	53.59	LiDAR	60.78
TANet (Ours)	83.81	75.38	67.66	54.92	46.67	42.42	73.84	59.86	53.46	LiDAR	62.00

Table 2: Performance comparison with previous approaches for 3D object detection task on the KITTI test split for *Cars*, *Pedestrians* and *Cyclists*. 3D mAP represents the mean average precision of all three categories on 3D object detection.

for interference feature.

Results on original point cloud data. The experimental results on the official KITTI test dataset are presented in Table. 2. Our method yields an 3D mAP of 62.00% over the three categories, outperforming the state-of-the-art methods PointPillars (Lang et al. 2019) and PointRCNN (Shi, Wang, and Li 2019) about 1.20% and 1.22%, respectively. In particular, for challenging objects (e.g. Pedestrians), our method achieves an improvement of 2.30% and 4.83% over PointPillars and PointRCNN, respectively. In addition, the visualization on the test set is provided in the supplemental material.

The visualization of TANet. We present the visualization of learned feature maps and predicted confidence scores feature produced by PointPillars (Lang et al. 2019) and TANet in Fig. 5. It can be observed from the visualized feature maps that TANet can focus on salient objects and ignore noising parts. Besides, compared to PointPillars, our method outputs higher confidence scores on the salient objects. For more challenging objects which PointPillars fails to detect, our method still obtains a satisfactory confidence score. This explains why our method performs better than PointPillars.

Running time. The average inference time of our method is 34.75ms, including 1) 8.0 ms for data pre-processing; 2) 12.42ms for voxel feature extraction; and 3) 14.33ms for detection and the post-processing procedure.

Ablation Studies

In this section, extensive ablation experiments are provided to analyze the effect of different components of TANet. All of our experimental results are conducted on the official validation set with 100 noise points for all the three categories. The official code of PointPillars (Lang et al. 2019) is reproduced as our baseline. Moreover, some experimental results can be seen on the official validation set without noise points on supplemental material.

Analysis of the attention mechanisms. Table. 3 presents extensive ablation studies of the proposed attention mechanisms. We remove the TA and Fine Regression(FR) module from our model as the baseline and achieves a 3D mAP of 65.59%. With only Point-wise Attention (PA) and only Channel-wise Attention (CA), the performance is boosted to 67.04% and 66.93%, respectively. And we name the parallel attention fusion mechanism for PA and CA as PACA.

As shown in Table. 3, specifically, when combining them together, PACA yields a 3D mAP of 67.38%, outperforming the baseline model by 1.8%. To verify the superiority of PACA better, we provide three alternatives (*Concat*, *PA-CA* and *CA-PA*), which combine the spatial attention and channel-wise attention in different ways. Concretely, *Concat* concatenates the outputs of these two kinds of attention along the channel-wise dimension. *PA-CA* (resp. *CA-PA*) represents cascading spatial attention (resp. channel-wise attention) and channel-wise attention (resp. spatial attention) sequentially. It can be obviously observed that PACA outperforms all these three combination mechanisms, suggesting PACA is of great importance for making rational use of both spatial and channel-wise information. It should be noted that all these different ways for attention combination bring significant improvements over the baseline. Moreover, based on the PACA, TA module takes the voxel-wise attention into consideration, which achieves the improvement with about 0.40%. It demonstrates the effectiveness and robustness of unifying the channel-wise attention point-wise attention and voxel-wise attention.

Effect of the PSA module. The effect of the PSA module is further explored. Our method also compares with the most representative single-shot refinement network RefineDet (Zhang et al. 2018) under two settings: with and without the TA module. Under the same setting, our method is consistently better than RefineDet. Without the TA module, it is noticed that the improvement of the PSA module is not so obvious. But surprisingly, with the TA module, the performance of the PSA achieves an evident improvement. It means that the PSA module has a fine complementarity with the TA module. The TA module can provide a robust and discriminate feature representations, and the PSA module can make full use of them to estimate 3D bounding boxes.

Conclusion

This paper proposes a novel TANet for 3D object detection in point clouds, especially for noising point clouds. The Triple Attention (TA) module and the Coarse-to-Fine Regression (CFR) module are the core parts of TANet. The former adaptively enhances crucial information of the objects and suppresses the interference points. The latter provides more accurate detection boxes without excessive computa-

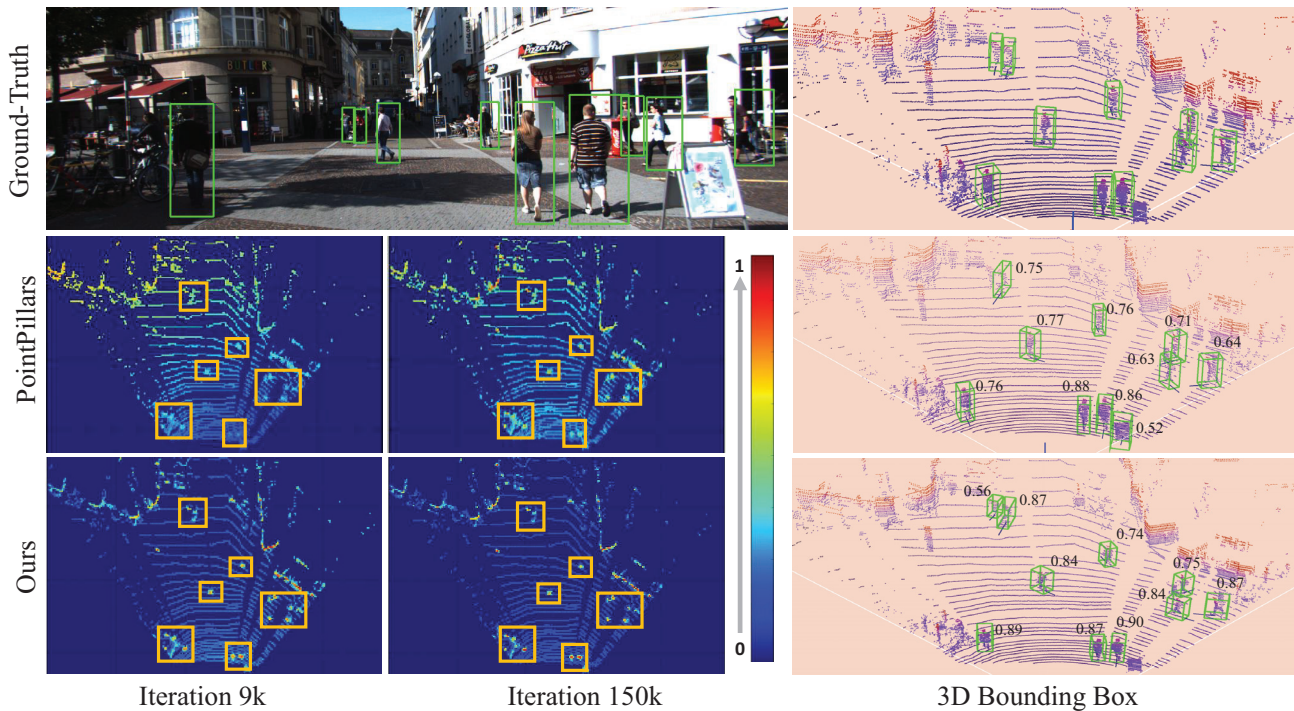


Figure 5: Visualization of learned feature map and predicted confidence score. The first row shows the ground-truth detection on 2D image and 3D point clouds. The second and third row illustrates the learned feature map and the predicted confidence score of PointPillars and our methods, respectively. We highlight some crucial areas for each feature map with a yellow rectangular box.

Method	Cars	Pedestrians	Cyclists	3D mAP
Baseline	77.20	55.04	64.53	65.59
CA	77.46	57.61	65.71	66.93
PA	77.58	57.59	65.97	67.04
PA-CA	77.90	56.65	66.10	66.88
CA-SA	77.82	56.92	66.04	66.93
Concat	77.87	57.32	66.17	67.12
PACA (Ours)	77.97	57.94	66.22	67.38
TA module (Ours)	78.33	58.43	66.61	67.79

Table 3: Ablation experiments on the effect of channel-wise attention, point-wise attention and voxel-wise attention, as well as different combination settings. All the experiments are conducted without FR module.

tion cost. Our method achieves state-of-the-art performance on the KITTI dataset. More importantly, in the more challenging cases that the additional noise points are added, extensive experiments further demonstrate the superior robustness of our method, which outperforms existing approaches by a large margin.

Acknowledgement

This work was supported by National Key R&D Program of China (No. 2018YFB1004600), to Dr. Xiang Bai by the National Program for Support of Top-notch Young Professionals and the Program for HUST Academic Frontier Youth Team 2017QYTD08. This work was also supported

Method	Cars	Pedestrians	Cyclists	3D mAP
Baseline	77.20	55.04	64.53	65.59
Baseline + RefineDet (2018)	77.27	55.86	65.23	66.12
Baseline + PSA	77.30	56.02	65.30	66.21
TA module	78.33	58.43	66.61	67.79
TA module + RefineDet (2018)	79.28	60.05	67.06	68.80
TA module + PSA	79.34	60.84	67.87	69.35

Table 4: Ablation experiments on the effect of the proposed PSA module.

by the National Natural Science Foundation of China No. 61703049 and Grant No. 61602485.

References

- Chen, X.; Kundu, K.; Zhang, Z.; Ma, H.; Fidler, S.; and Urtasun, R. 2016. Monocular 3d object detection for autonomous driving. In *CVPR*, 2147–2156.
- Chen, X.; Ma, H.; Wan, J.; Li, B.; and Xia, T. 2017. Multi-view 3d object detection network for autonomous driving. In *CVPR*, 1907–1915.
- Du, X.; Ang, M. H.; Karaman, S.; and Rus, D. 2018. A general pipeline for 3d detection of vehicles. In *ICRA*, 3194–3200.
- Engelcke, M.; Rao, D.; Wang, D. Z.; Tong, C. H.; and Posner, I. 2017. Vote3deep: Fast object detection in 3d point clouds using efficient convolutional neural networks. In *ICRA*, 1355–1361. IEEE.

- Fu, C.-Y.; Liu, W.; Ranga, A.; Tyagi, A.; and Berg, A. C. 2017. Dssd: Deconvolutional single shot detector. *ArXiv*.
- Geiger, A.; Lenz, P.; and Urtasun, R. Are we ready for autonomous driving? In *CVPR*, 3354–3361.
- Girshick, R. 2015. Fast r-cnn. In *ICCV*, 1440–1448.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*, 2961–2969.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *CVPR*, 7132–7141.
- Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. *ICLR*.
- Ku, J.; Mozifian, M.; Lee, J.; Harakeh, A.; and Waslander, S. L. 2018. Joint 3d proposal generation and object detection from view aggregation. In *IROS*, 1–8. IEEE.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2018. Pointpillars: Fast encoders for object detection from point clouds. *arXiv preprint arXiv:1812.05784*.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, 12697–12705.
- Li, B.; Zhang, T.; and Xia, T. 2016. Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916*.
- Liang, M.; Yang, B.; Wang, S.; and Urtasun, R. 2018. Deep continuous fusion for multi-sensor 3d object detection. In *ECCV*, 641–656.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *CVPR*, 2117–2125.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017b. Focal loss for dense object detection. In *ICCV*, 2980–2988.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *ECCV*, 21–37. Springer.
- Mousavian, A.; Anguelov, D.; Flynn, J.; and Kosecka, J. 2017. 3d bounding box estimation using deep learning and geometry. In *CVPR*, 7074–7082.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 5099–5108.
- Qi, C. R.; Liu, W.; Wu, C.; Su, H.; and Guibas, L. J. 2018. Frustum pointnets for 3d object detection from rgb-d data. In *CVPR*, 918–927.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 91–99.
- Shen, Z.; Liu, Z.; Li, J.; Jiang, Y.-G.; Chen, Y.; and Xue, X. 2017. Dsod: Learning deeply supervised object detectors from scratch. In *ICCV*, 1919–1927.
- Shi, S.; Wang, X.; and Li, H. 2019. Pointcnn: 3d object proposal generation and detection from point cloud. In *CVRR*, 770–779.
- Simon, M.; Milz, S.; Amende, K.; and Gross, H.-M. 2018. Complex-yolo: An euler-region-proposal for real-time 3d object detection on point clouds. In *ECCV*, 197–209. Springer.
- Song, S., and Chandraker, M. 2015. Joint sfm and detection cues for monocular 3d localization in road scenes. In *CVPR*, 3734–3742.
- Wang, D. Z., and Posner, I. 2015. Voting for voting in on-line point cloud object detection. In *Robotics: Science and Systems*, volume 1, 10–15607.
- Yan, Y.; Mao, Y.; and Li, B. 2018. Second: Sparsely embedded convolutional detection. *Sensors* 18(10):3337.
- Yang, B.; Liang, M.; and Urtasun, R. 2018. Hdnet: Exploiting hd maps for 3d object detection. In *Conference on Robot Learning*, 146–155.
- Yang, B.; Luo, W.; and Urtasun, R. 2018. Pixor: Real-time 3d object detection from point clouds. In *CVPR*, 7652–7660.
- Zhang, S.; Wen, L.; Bian, X.; Lei, Z.; and Li, S. Z. 2018. Single-shot refinement neural network for object detection. In *CVPR*, 4203–4212.
- Zhou, Y., and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, 4490–4499.