

# DID-M3D: Decoupling Instance Depth for Monocular 3D Object Detection

Liang Peng<sup>1,2</sup>, Xiaopei Wu<sup>1</sup>, Zheng Yang<sup>2</sup>, Haifeng Liu<sup>1</sup>, and Deng Cai<sup>1,2</sup> ✉

<sup>1</sup> State Key Lab of CAD&CG, Zhejiang University, China

{pengliang, wuxiaopei, haifengliu}@zju.edu.cn   dengcai@cad.zju.edu.cn

<sup>2</sup> Fabu Inc., Hangzhou, China

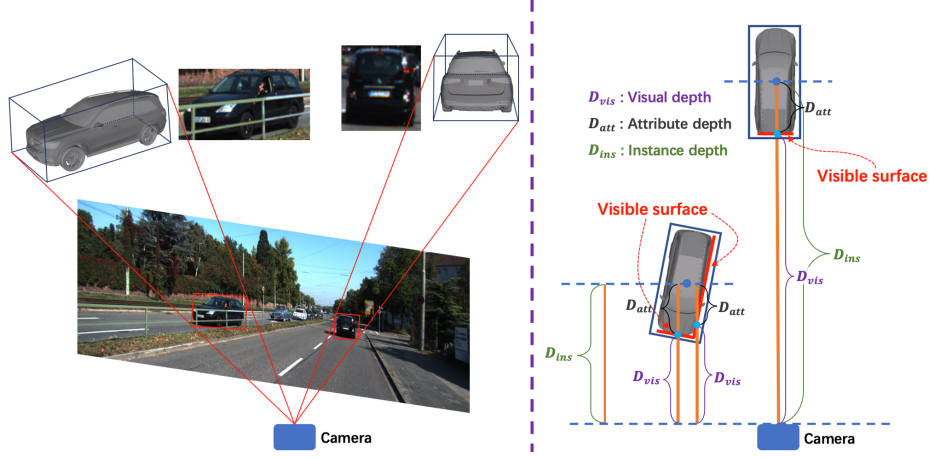
yangzheng@fabu.ai

**Abstract.** Monocular 3D detection has drawn much attention from the community due to its low cost and setup simplicity. It takes an RGB image as input and predicts 3D boxes in the 3D space. The most challenging sub-task lies in the instance depth estimation. Previous works usually use a direct estimation method. However, in this paper we point out that the instance depth on the RGB image is non-intuitive. It is coupled by visual depth clues and instance attribute clues, making it hard to be directly learned in the network. Therefore, we propose to reformulate the instance depth to the combination of the instance visual surface depth (**visual depth**) and the instance attribute depth (**attribute depth**). The visual depth is related to objects' appearances and positions on the image. By contrast, the attribute depth relies on objects' inherent attributes, which are invariant to the object affine transformation on the image. Correspondingly, we decouple the 3D location uncertainty into visual depth uncertainty and attribute depth uncertainty. By combining different types of depths and associated uncertainties, we can obtain the final instance depth. Furthermore, data augmentation in monocular 3D detection is usually limited due to the physical nature, hindering the boost of performance. Based on the proposed instance depth disentanglement strategy, we can alleviate this problem. Evaluated on KITTI, our method achieves new state-of-the-art results, and extensive ablation studies validate the effectiveness of each component in our method. The codes are released at <https://github.com/SPengLiang/DID-M3D>.

**Keywords:** monocular 3D detection, instance depth estimation.

## 1 Introduction

Monocular 3D object detection is an important topic in the self-driving and computer vision community. It is popular due to its low price and configuration simplicity. Rapid improvements [6,5,28,36,60,25] have been conducted in recent years. A well-known challenge in this task lies in instance depth estimation, which is the bottleneck towards boosting the performance since the depth information is lost after the camera projection process.



**Fig. 1.** We decouple the instance depth into visual depths and attribute depths due to the coupled nature of instance depth. Please refer to the text for more details. Best viewed in color with zoom in.

Many previous works [2,12,39] directly regress the instance depth. This manner does not consider the ambiguity brought by the instance depth itself. As shown in Figure 1, for the right object, its instance depth is the sum of car tail depth and half-length of the car, where the car length is ambiguous since both car’s left and right sides are invisible. For the left object, except for the intuitive visible surface depth, the instance depth further depends on the car dimension and orientation. We can observe that the instance depth is non-intuitive. It requires the network to additionally learn instance inherent attributes on the instance depth head. Previous direct estimation and mediate optimization methods do not fully consider this coupled nature. Thus they lead to suboptimal performance on the instance depth estimation, showing less accurate results.

Based on the analysis above, in this paper we propose to decouple the instance depth to **instance visual surface depth (visual depth)** and **instance attribute depth (attribute depth)**. We illustrate some examples in Figure 1. For each point (or small patch) on the object, the visual depth denotes the absolute depth towards the agent (car/robot) camera, and we define the attribute depth as the relative depth offset from the point (or small patch) to the object’s 3D center. This decoupled manner encourages the network to learn different feature patterns of the instance depth. Visual depth on monocular imagery depends on objects’ appearances and positions [11] on the image, which is affine-sensitive. By contrast, attribute depth highly relies on object inherent attributes (*e.g.*, dimensions and orientations) of the object. It focuses on features inside the RoI, which is affine-invariant. (See Section 4.1 and 4.2 for detailed discussion). Thus the attribute depth is independent of the visual depth, and decoupled instance

depth allows us to use separate heads to extract different types of features for different types of depths.

Specifically, for an object image patch, we divide it into  $m \times n$  grids. Each grid can denote a small area on the object, with being assigned a visual depth and the corresponding attribute depth. Considering the occlusion and 3D location uncertainty, we use the uncertainty to denote the confidence of each depth prediction. At inference, every object can produce  $m \times n$  instance depth predictions, thus we take advantage of them and associated uncertainties to adaptively obtain the final instance depth and confidence.

Furthermore, prior works usually are limited by the diversity of data augmentation, due to the complexity of keeping alignment between 2D and 3D objects when enforcing affine transformation in a 2D image. Based on the decoupled instance depth, we show that our method can effectively perform data augmentation, including the way using affine transformation. It is achieved by the affine-sensitive and affine-invariant nature of the visual depth and attribute depth, respectively (See Section 4.3). To demonstrate the effectiveness of our method, we perform experiments on the widely used KITTI dataset. The results suggest that our method outperforms prior works with a significant margin.

In summary, our contributions are listed as follows:

1. We point out the coupled nature of instance depth. Due to the entangled features, the previous way of directly predicting instance depth is suboptimal. Therefore, we propose to decouple the instance depth into attribute depths and visual depths, which are independently predicted.
2. We present two types of uncertainties to represent depth estimation confidence. Based on this, we propose to adaptively aggregate different types of depths into the final instance depth and correspondingly obtain the 3D localization confidence.
3. With the help of the proposed attribute depth and visual depth, we alleviate the limitation of using affine transformation in data augmentation for monocular 3D detection.
4. Evaluated on KITTI benchmark, our method sets the new state of the art (SOTA). Extensive ablation studies demonstrate the effectiveness of each component in our method.

## 2 Related Work

### 2.1 LiDAR-based 3D Object Detection

LiDAR-based methods utilize precise LiDAR point clouds to achieve high performance. According to the representations usage, they can be categorized into point-based, voxel-based, hybrid, and range-view-based methods. Point-based methods [37,42,51] directly use raw point clouds to preserve fine-grained structures of objects. However, they usually suffer from high computational costs. Voxel-based methods [58,50,30,56,53] voxelize the unordered point clouds into regular grids so that the CNNs can be easily applied. These methods are more

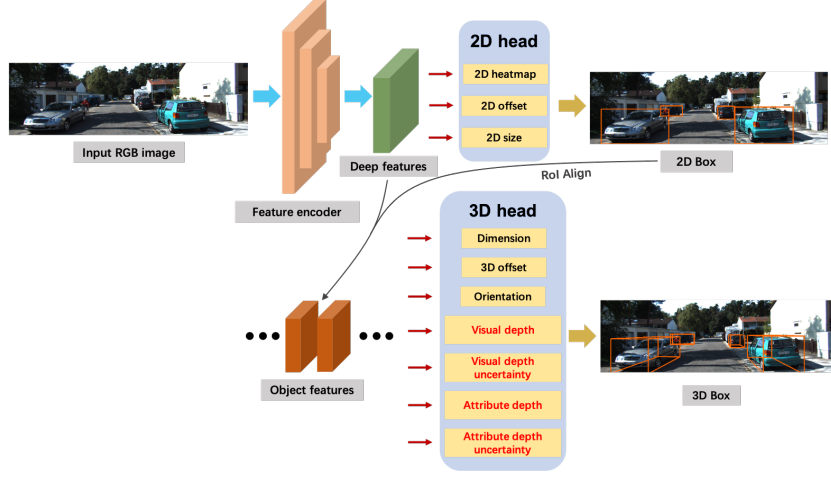
efficient, but voxelization introduces quantization errors, resulting in information loss. To explore advantages of different representations, some hybrid methods [8,52,43,41,32] are proposed. They validate that combining point-based and voxel-based methods can achieve a better trade-off between accuracy and efficiency. Range-view-based methods [1,19,4,13] organize point clouds in range view, which is a compact representation of point clouds. These methods are also computationally efficient but are under-explored.

## 2.2 Monocular 3D Object Detection

Due to the low cost and setup simplicity, monocular 3D object detection is popular in recent years. Previous monocular works can be roughly divided into image-only based and depth-map based methods. The pioneer method [6] integrates different types of information such as segmentation and scene priors for performing 3D detection. To learn spatial features, OFTNet [40] projects 2D image features to the 3D space. M3D-RPN [2] attempts to extract depth-aware features by using different convolution kernels on different image columns. Kinematic3D [3] uses multi-frames to capture the temporal information by employing a tracking module. GrooMeD-NMS [18] develops a trainable NMS-step to boost the final performance. At the same period of time, many works fully take advantage of the scene and geometry priors [59,44,48,23]. Due to the ill-posed nature of monocular imagery, some works resort to using the dense depth estimation [49,12,26,35,46,39,47]. With the help of estimated depth maps, RoI-10D [29] use CAD models to augment training samples. Pseudo-LiDAR [49,27] based methods are also popular. They convert estimated depth maps to point clouds, then well-designed LiDAR 3D detectors can be directly employed. CaDDN [39] predicts a categorical depth distribution, to precisely project depth-aware features to 3D space. In sum, benefiting from rapid developments of deep learning technologies, monocular 3D detection has conducted remarkable progress.

## 2.3 Estimation of Instance Depth

Most monocular works directly predict the instance depth. There are also some works that use auxiliary information to help the instance depth estimation in the post-processing. They usually take advantage of the geometry constraints and scene priors. The early work Deep3DBox [31] regresses object dimensions and orientations, the remaining 3D box parameters including the instance depth are estimated by 2D-3D box geometry constraints. This indirect way has poor performance because it does not fully use the supervisions. To use geometric projection constraints, RTM3D [21] predicts nine perspective key-points (center and corners) of a 3D bounding box in the image space. Then the initially estimated instance depth can be optimized by minimizing projection errors. KM3D [20] follows this line and integrates this optimization into a end-to-end training process. More recently, MonoFlex [55] also predicts the nine perspective key-points. In addition to the directly predicted instance depth, it uses the projection heights in pair key-points, using geometric relationships to produce new instance depths.



**Fig. 2.** Network framework. The overall design follows GUPNet [25]. The estimated 2D boxes are used to extract specific features for each object, followed by 3D box heads, which predict required 3D box parameters. The red parts in 3D heads denote the newly-proposed components. They are used to decouple the instance depth.

MonoFlex develops an ensemble strategy to obtain the final instance depth. Differing from previous works, GUPNet [25] proposes an uncertainty propagation method for the instance depth. It uses estimated object 3D dimensions and 2D height to obtain initial instance depth, with additionally predicting a depth bias to refine the instance depth. GUPNet mainly focuses on tackling the error amplification problem in the geometry projection process. MonoRCNN [44] also introduces a distance decomposition based on the 2D height and 3D height. Such methods use geometric or auxiliary information to refine the estimated instance depth, while they do not fully use the coupled nature of the instance depth.

### 3 Overview and Framework

**Preliminaries** Monocular 3D detection takes an image captured by an RGB camera as input, predicting amodal 3D bounding boxes of objects in 3D space. These 3D boxes are parameterized by the 3D center location  $(x, y, z)$ , dimension  $(h, w, l)$ , and the orientation  $(\theta)$ . Please note, in the self-driving scenario, the orientation usually refers to the yaw angle, and the roll and pitch angles are zeros by default. Also, the ego-car/robot has been calibrated.

In this paragraph, we provide an overview and describe the framework. The overall framework is shown in Figure 2. First, the network takes an RGB image  $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$  as the input. After feature encoding, we have deep features  $\mathbf{F} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$ , where  $C$  is the channel number. Second, deep features  $\mathbf{F}$  are fed into three 2D detection heads, namely, 2D heatmap  $\mathbf{H} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times B}$ , 2D offset  $\mathbf{O}_{2d} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 2}$ , and 2D size  $\mathbf{S}_{2d} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times 2}$ , where  $B$  is the number of categories. By

combining such 2D head predictions, we can achieve 2D box predictions. Then, with 2D box estimates, single object features are obtained from deep features  $\mathbf{F}$  by RoI Align. We have object features  $\mathbf{F}_{obj} \in \mathbb{R}^{n \times 7 \times 7 \times C}$ , where  $7 \times 7$  is the RoI Align size and  $n$  refers to the number of RoIs. Finally, these object features  $\mathbf{F}_{obj}$  are fed into 3D detection heads to produce 3D parameters. Therefore, we have 3D box dimension  $\mathbf{S}_{3d} \in \mathbb{R}^{n \times 3}$ , 3D center projection offset  $\mathbf{O}_{3d} \in \mathbb{R}^{n \times 2}$ , orientation  $\Theta \in \mathbb{R}^{n \times k \times 2}$  (we follow the multi-bin design [31] where  $k$  is the bin number), visual depth  $\mathbf{D}_{vis} \in \mathbb{R}^{n \times 7 \times 7}$ , visual depth uncertainty  $\mathbf{U}_{vis} \in \mathbb{R}^{n \times 7 \times 7}$ , attribute depth  $\mathbf{D}_{att} \in \mathbb{R}^{n \times 7 \times 7}$ , and attribute depth uncertainty  $\mathbf{U}_{att} \in \mathbb{R}^{n \times 7 \times 7}$ . Using the parameters above, we can achieve the final 3D box predictions. In the following sections, we will detail the proposed method.

## 4 Decoupled Instance Depth

We divide the RoI image patch into  $7 \times 7$  grids, assigning each grid a visual depth value and an attribute depth value. We provide the ablation on the grid size for visual and attribute depth in experiments (See Section 5.4 and Table 4). In the following, we first detail the two types of depths, followed by the decoupled-depth based data augmentation, then introduce the way of obtaining the final instance depth, and finally describe loss functions.


### 4.1 Visual Depth

The visual depth denotes the physical depth of the object surface on the small RoI image grid. For each grid, we define the visual depth as the average pixel-wise depth within the grid. If the grid is  $1 \times 1$  pixel, the visual depth is equal to the pixel-wise depth. Given that a pixel denotes the quantified surface of the object, we can regard visual depths as the general extension of pixel-wise depths.

The visual depth in monocular imagery has an important property. For a monocular-based system, visual depth highly relies on the object’s 2D box size (the faraway object appears small on the images and vice versa) and the position on the image (lower  $v$  coordinates under image coordinate system indicate larger depths) [11]. Therefore, if we perform an affine transformation to the image, the visual depth should be correspondingly transformed, where the depth value should be scaled. We call this nature the affine-sensitive.

### 4.2 Attribute Depth

The attribute depth refers to the depth offset from the visual surface to the object’s 3D center. We call it attribute depth because it is more likely related to the object’s inherent attributes. For example, when the car orientation is parallel to  $z$ -axis (depth direction) in 3D space, the attribute depth of the car tail is the car’s half-length. By contrast, the attribute depth is car’s half-width if the orientation is parallel to  $x$ -axis. We can see that the attribute depth depends on the object semantics and its inherent attributes. In contrast to the affine-sensitive



**Fig. 3.** Affine transformation based data augmentation. We do not change the object’s inherent attributes, *i.e.*, attribute depths, dimensions, and observation angles. The visual depth is scaled according to the 2D height scale factor. The 3D center projection is transformed together with the image affine transformation.

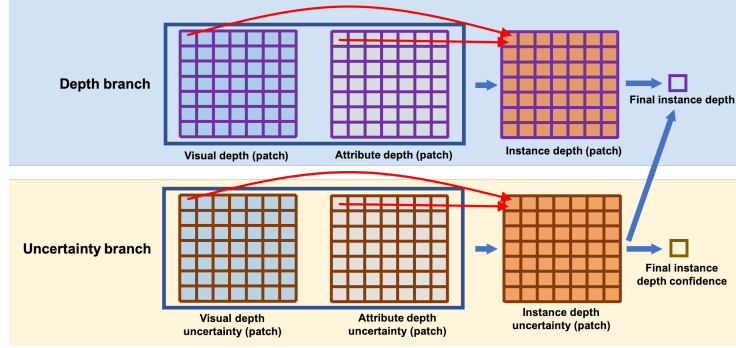
nature of visual depth, attribute depth is invariant to any affine transformation because object inherent characteristics will not change. We call this nature the affine-invariant.

As described above, we use two separate heads to estimate the visual depth and attribute depth, respectively. The disentanglement of instance depth has several advantages: (1) The object depth is decoupled in a reasonable and intuitive manner, thus we can more comprehensively and precisely represent the object; (2) The network is allowed to extract different types of features for different types of depths, which facilitates the learning; (3) Benefitting from the decoupled depth, our method can effectively perform affine transformation based data augmentation, which is usually limited in previous works.

### 4.3 Data Augmentation

In monocular 3D detection, many previous works are limited by data augmentation. Most of them only employ photometric distortion and flipping transformation. Data augmentation using affine transformation is hard to be adopted because the transformed instance depth is agnostic. Based on the decoupled depth, we point out that our method can alleviate this problem.

We illustrate an example in Figure 3. Specifically, we add the random cropping and scaling strategy [57] in the data augmentation. The 3D center projection point on the image follows the same affine transformation process of the image. The visual depth is scaled by the scale factor along  $y$ -axis on the image because



**Fig. 4.** Depth flow for an object. We use the visual depth, attribute depth, and the associated uncertainty to obtain the final instance depth.

$d = \frac{f \cdot h_{3d}}{h_{2d}}$ , where  $f, h_{3d}, h_{2d}$  is the focal length, object 3D height and 2D height, respectively. Conversely, the attribute depth keeps the same due to its affine-invariant nature. We do not directly scale the instance depth, as this manner will implicitly damage the attribute depth. Similarly, other inherent attributes of objects, *i.e.*, the observation angle and the dimension, keep the same as original values. We empirically show that the data augmentation works well. We provide the ablations in experiments (See Section 5.4 and Table 3).

#### 4.4 Depth Uncertainty and Aggregation

The 2D classification score cannot fully express the confidence in monocular 3D detection because of the difficulty in 3D localization. Previous works [25,45] use the instance depth confidence or 3D IoU loss, integrating with 2D detection confidence to represent the final 3D detection confidence. Given that we have decoupled the instance depth into visual depth and attribute depth, we can further decouple the instance depth uncertainty. Only when an object has low visual uncertainty and low attribute depth uncertainty simultaneously, the instance depth can have high confidence.

Inspired by [16,25], we assume every depth prediction is a Laplace distribution. Specifically, for each visual depth  $d_{vis}$  in  $\mathbf{D}_{vis} \in \mathbb{R}^{n \times 7 \times 7}$  and the corresponding uncertainty  $u_{vis}$  in  $\mathbf{U}_{vis} \in \mathbb{R}^{n \times 7 \times 7}$ , they follow the Laplace distribution  $L(d_{vis}, u_{vis})$ . Similarly, the attribute depth distribution is  $L(d_{att}, u_{att})$ , where  $d_{att}$  in  $\mathbf{D}_{att} \in \mathbb{R}^{n \times 7 \times 7}$  and  $u_{att}$  in  $\mathbf{U}_{att} \in \mathbb{R}^{n \times 7 \times 7}$ . Therefore, the instance depth distribution derived by associated visual and attribute depth is  $L(\tilde{d}_{ins}, \tilde{u}_{ins})$ , where  $\tilde{d}_{ins} = d_{vis} + d_{att}$  and  $\tilde{u}_{ins} = \sqrt{u_{vis}^2 + u_{att}^2}$ . Then we use  $\tilde{\mathbf{D}}_{ins(patch)} \in \mathbb{R}^{n \times 7 \times 7}$  and  $\tilde{\mathbf{U}}_{ins(patch)} \in \mathbb{R}^{n \times 7 \times 7}$  to denote the instance depth and the uncertainty on the RoI patch. We illustrate the depth flow process in Figure 4.

To obtain the final instance depth, we first convert the uncertainty to probability [16,25], which can be written as  $\mathbf{P}_{ins(patch)} = \exp(-\tilde{\mathbf{U}}_{ins(patch)})$ , where



$\mathbf{P}_{ins(patch)} \in \mathbb{R}^{n \times 7 \times 7}$ . Then we aggregate the instance depth on the patch to the final instance depth. For the  $i^{th}$  object ( $i = 1, \dots, n$ ), we have:

$$d_{ins} = \sum \frac{\tilde{\mathbf{D}}_{ins(patch)_i} \mathbf{P}_{ins(patch)_i}}{\sum \mathbf{P}_{ins(patch)_i}} \quad (1)$$

The corresponding final instance depth confidence is:

$$p_{ins} = \sum \left( \frac{\mathbf{P}_{ins(patch)_i}}{\sum \mathbf{P}_{ins(patch)_i}} \mathbf{P}_{ins(patch)_i} \right) \quad (2)$$

Therefore, the final 3D detection confidence is  $p = p_{2d} \cdot p_{ins}$ , where  $p_{2d}$  is the 2D detection confidence.

#### 4.5 Loss Functions

**2D detection part:** As shown in Figure 2, for the 2D object detection part, we follow the design in CenterNet [57]. The 2D heatmap  $\mathbf{H}$  aims to indicate the rough object center on the image. The size is  $\frac{H}{4} \times \frac{W}{4} \times B$ , where  $H, W$  is the input image size and  $B$  is the number of categories. The 2D offset  $O_{2d}$  refers to the residual towards rough 2D centers, and the 2D size  $S_{2d}$  denotes the 2D box height and width. Following CenterNet [57], we have loss functions  $\mathcal{L}_H, \mathcal{L}_{O_{2d}}, \mathcal{L}_{S_{2d}}$ , respectively.

**3D detection part:** For the 3D object dimension, we follow the typical transformation and loss design [2]  $\mathcal{L}_{S_{3d}}$ . For the orientation, the network predicts the observation angle and uses the multi-bin loss [31]  $\mathcal{L}_\Theta$ . Also, we use the 3D center projection on the image plane and the instance depth to recover the object’s 3D location. For the 3D center projection, we achieve it by predicting the 3D projection offset to the 2D center. The loss function is:  $\mathcal{L}_{O_{3d}} = \text{Smooth}L_1(O_{3d}, O_{3d}^*)$ . We use  $*$  to denote corresponding labels. As mentioned above, we decouple the instance depth into visual depth and attribute depth. The visual depth labels are obtained by projecting LiDAR points onto the image and the attribute depth labels are obtained by subtracting instance depth labels with visual depth labels. Combing with the uncertainty [16,9], the visual depth loss is:  $\mathcal{L}_{D_{vis}} = \frac{\sqrt{2}}{u_{vis}} \|d_{vis} - d_{vis}^*\| + \log(u_{vis})$ , where  $u_{vis}$  is the uncertainty. Similarly, we have attribute depth loss  $\mathcal{L}_{D_{att}}$  and instance depth loss  $\mathcal{L}_{D_{ins}}$ . Among these loss terms, the losses concerning instance depth ( $\mathcal{L}_{D_{vis}}, \mathcal{L}_{D_{att}}$ , and  $\mathcal{L}_{D_{ins}}$ ) play the most important role since they matter objects’ localization in the 3D space. We empirically set 1.0 for weights of all loss terms, and the overall loss is:

$$\mathcal{L} = \mathcal{L}_H + \mathcal{L}_{O_{2d}} + \mathcal{L}_{S_{2d}} + \mathcal{L}_{S_{3d}} + \mathcal{L}_\Theta + \mathcal{L}_{O_{3d}} + \mathcal{L}_{D_{vis}} + \mathcal{L}_{D_{att}} + \mathcal{L}_{D_{ins}} \quad (3)$$

## 5 Experiments

### 5.1 Implementation Details

We conduct experiments on 2 NVIDIA RTX 3080Ti GPUs with batch size 16. We use the PyTorch framework [34]. We train the network with 200 epochs and

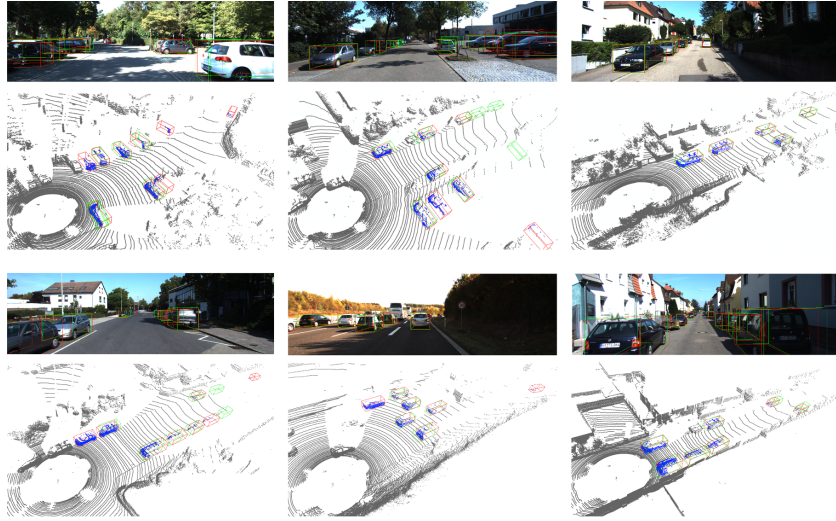
**Table 1.** Comparisons on KITTI testing set. The **red** refers to the highest result and **blue** is the second-highest result. Our method achieves state-of-the-art results. Note that DD3D[33] uses a large private dataset (DDAD15M), which includes 15M frames.

Approaches	Venue	Runtime	AP <sub>BEV</sub> (IoU=0.7)  <sub>R40</sub>			AP <sub>3D</sub> (IoU=0.7)  <sub>R40</sub>		
			Easy	Moderate	Hard	Easy	Moderate	Hard
MonoGRNet [38]	<i>AAAI19</i>	400ms	18.19	11.17	8.73	15.74	9.61	4.25
ROI-10D [29]	<i>CVPR19</i>	200ms	9.78	4.91	3.74	4.32	2.02	1.46
MonoPSR [17]	<i>CVPR19</i>	200ms	18.33	12.58	9.91	10.76	7.25	5.85
M3D-RPN [2]	<i>ICCV19</i>	160ms	21.02	13.67	10.23	14.76	9.71	7.42
AM3D [27]	<i>ICCV19</i>	400ms	25.03	17.32	14.91	16.50	10.74	9.52
MonoPair [9]	<i>CVPR20</i>	60ms	19.28	14.83	12.89	13.04	9.99	8.65
D4LCN [12]	<i>CVPR20</i>	200ms	22.51	16.02	12.55	16.65	11.72	9.51
RTM3D [21]	<i>ECCV20</i>	40ms	19.17	14.20	11.99	14.41	10.34	8.77
PatchNet [26]	<i>ECCV20</i>	400ms	22.97	16.86	14.97	15.68	11.12	10.17
Kinematic3D [3]	<i>ECCV20</i>	120ms	26.69	17.52	13.10	19.07	12.72	9.17
Neighbor-Vote [10]	<i>MM21</i>	100ms	27.39	18.65	16.54	15.57	9.90	8.89
Ground-Aware [23]	<i>RAL21</i>	50ms	29.81	17.98	13.08	21.65	13.25	9.91
MonoRUn [5]	<i>CVPR21</i>	70ms	27.94	17.34	15.24	19.65	12.30	10.58
DDMP-3D [46]	<i>CVPR21</i>	180ms	28.08	17.89	13.44	19.71	12.78	9.80
Monodle [28]	<i>CVPR21</i>	40ms	24.79	18.89	16.00	17.23	12.26	10.29
CaDDN [39]	<i>CVPR21</i>	630ms	27.94	18.91	17.19	19.17	13.41	11.46
GrooMeD-NMS [18]	<i>CVPR21</i>	120ms	26.19	18.27	14.05	18.10	12.32	9.65
MonoEF [59]	<i>CVPR21</i>	30ms	29.03	19.70	17.26	21.29	13.87	11.71
MonoFlex [55]	<i>CVPR21</i>	35ms	28.23	19.75	16.89	19.94	13.89	12.07
MonoRCNN [44]	<i>ICCV21</i>	70ms	25.48	18.11	14.10	18.36	12.65	10.03
AutoShape [24]	<i>ICCV21</i>	40ms	30.66	20.08	15.95	22.47	14.17	11.36
GUPNet [25]	<i>ICCV21</i>	34ms	30.29	21.19	18.20	22.26	15.02	13.12
PCT [47]	<i>NeurIPS21</i>	45ms	29.65	19.03	15.92	21.00	13.37	11.31
MonoCon [22]	<i>AAAI22</i>	26ms	<b>31.12</b>	<b>22.10</b>	<b>19.00</b>	<b>22.50</b>	<b>16.46</b>	<b>13.95</b>
DD3D [33]	<i>ICCV21</i>	-	<i>32.35</i>	<i>23.41</i>	<i>20.42</i>	<i>23.19</i>	<i>16.87</i>	<i>14.36</i>
<b>DID-M3D (ours)</b>	<i>ECCV22</i>	40ms	<b>32.95</b>	<b>22.76</b>	<b>19.83</b>	<b>24.40</b>	<b>16.29</b>	<b>13.75</b>
Improvements			+1.83	+0.66	+0.83	+1.90	-0.17	-0.20

employ the Hierarchical Task Learning (HTL) [25] training strategy. The Adam optimizer is used with the initial learning rate  $1e-5$ . The learning rate increases to  $1e-3$  in the first 5 epochs by employing the linear warm-up strategy and decays in epoch 90 and 120 with rate 0.1. We set  $k$  12 in the multi-bin orientation  $\Theta$ . For the backbone and head, we follow the design in [25,54]. Inspired by CaDDN [39], we project LiDAR point clouds onto the image frame to create sparse depth maps and then depth completion [15] is performed to generate depth values at each pixel in the image. Considering the space limitation, we provide more experimental results and discussion in the supplementary material.

## 5.2 Dataset and Metrics

Following the commonly adopted setup in previous works [2,12,18,24], we perform experiments on the widely used KITTI [14] 3D detection dataset. The KITTI dataset provides 7,481 training samples and 7,518 testing samples, where training sample labels are publicly available and the labels of testing samples keep secret in the KITTI website, which are only used for online evaluation and ranking. To conduct ablations, previous work further divides the 7,481 samples into a new training set with 3,712 samples and a *val* set with 3,769 samples. This



**Fig. 5.** Qualitative results on KITTI *val* set. **Red:** ground-truth 3D boxes; **Green:** our predictions. We can observe that the model conducts accurate 3D box predictions. Best viewed in color with zoom in.

data split [7] is widely adopted by most previous works. Additionally, KITTI divides objects into the *easy*, *moderate*, and *hard* level according to the object 2D box height (related to the depth), occlusion and truncation levels. For evaluation metrics, we use the suggested  $AP_{40}$  metric [45] under the two core tasks, *i.e.*, 3D and bird’s-eye-view (BEV) detection.

### 5.3 Performance on KITTI Benchmark

We compare our method (DID-M3D) with other methods in KITTI test set, which is the official benchmark for monocular 3D object detection. The results are shown in Table 1. We can see that our method achieves a new state-of-the-art. For example, compared to GUPNet [25] (*ICCV21*), we boost the performance from 21.19/15.02 to 22.26/16.29 under the moderate setting. As for PCT [47] (*NeurIPS21*), we exceed it with 3.23/2.92 AP under the moderate setting, which is a significant improvement. When compared to the recent method MonoCon[22] (*AAAI22*), our method still shows better performance on all BEV metrics and a 3D metric. Also, the runtime of our method is comparable to other real-time methods. Such results validate the superiority of our method. Additionally, to demonstrate the generalizability on other categories, we perform experiments on cyclist and pedestrian categories. As shown in Table 5, our method brings obvious improvements to the baseline (without employing the proposed components). The results suggest that our method works well for other categories.

Moreover, we provide qualitative results on the RGB image and 3D space in Figure 5. We can observe that for most simple cases (*e.g.*, close objects without

occlusion and truncation), the model predictions are quite precise. However, for the heavily occluded, truncated, or faraway objects, the orientation or instance depth is less accurate. This is a common dilemma for most monocular works due to the limited information in monocular imagery. In the supplementary material, we will provide more experimental results and make detailed discussions on failure cases.

#### 5.4 Ablation Study

To investigate the impact of each component in our method, we conduct detailed ablation studies on KITTI. Following the common practice in previous works, all ablation studies are evaluated in the *val* set on the car category.

**Decoupled Instance Depth.** We report the results in Table 2. Experiment (a) is the baseline using the direct instance depth prediction. To make fair comparisons, for the baseline, we also employ the grid design (Experiment (b)). Similar to our method, it means that the network also produce  $7 \times 7$  instance depth predictions for every object, which are all supervised in training and averaged at inference. Then, we decouple the instance depth into visual depths and attribute depths (Experiment (b)  $\rightarrow$  (c)), this simple modification improves the accuracy significantly. This result indicates the network performs suboptimally due to the coupled nature of instance depth, demonstrating our viewpoint. From Experiment (c)  $\rightarrow$  (d,e), we can see that the depth uncertainty brings improvements, because the uncertainty stabilizes the training of depth, benefiting the network learning. When simultaneously enforcing both two types of uncertainty, the performance is further boosted. Please note, the decoupled instance depth is the precondition of decoupled uncertainty. Given that the two types of depth uncertainty are achieved, we can obtain the final instance depth uncertainty (Experiment (f)  $\rightarrow$  (g)). This can be regarded as the 3D location confidence. It is used to combine with the original 2D detection confidence, which brings obvious improvements. Finally, we can use the decoupled depth and corresponding uncertainties to adaptively obtain the final instance depth (Experiment (h)), while previous experiments use the average value on the patch. We can see that this design enhances the performance. In summary, by using the decoupled depth strategy, we improve the baseline performance from 16.79/11.24 to 22.76/16.12 (Experiment (b)  $\rightarrow$  (h)). It is an impressive result. Overall, the ablations validate the effectiveness of our method.

**Affine Transformation Based Data Augmentation.** In this paragraph we aim to understand the effect of affine transformation based data augmentation. The comparisons are shown in Table 3. We can see that the method obviously benefits from affine-based data augmentation. Note that the proper depth transformation is very important. When enforcing affine-based data augmentation, the visual depth should be scaled while the attribute depth should not be changed due to their affine-sensitive and affine-invariant nature, respectively.

**Table 2.** Ablation for decoupled instance depth. “Dec.”: decoupled; “ID.”: instance depth; “ $u_{vis}$ ”: visual depth uncertainty; “ $u_{att}$ ”: attribute depth uncertainty; “Conf.”: confidence; “AA.”: adaptive aggregation.

Experiments	Grid	Dec. ID.	$u_{vis}$	$u_{att}$	ID. Conf.	ID. AA.	AP <sub>BEV</sub> /AP <sub>3D</sub> (IoU=0.7)  <sub>R40</sub>		
							Easy	Moderate	Hard
(a)							19.84/14.07	16.25/11.20	14.13/9.97
(b)	✓						21.01/14.67	16.79/11.24	15.41/10.18
(c)	✓	✓					22.98/16.95	18.72/13.24	16.57/11.23
(d)	✓	✓	✓				25.13/16.85	19.84/13.45	17.19/12.03
(e)	✓	✓		✓			24.83/17.29	19.66/13.50	17.06/11.48
(f)	✓	✓	✓	✓			25.23/18.14	20.06/13.91	17.63/12.52
(g)	✓	✓	✓	✓	✓		29.34/21.51	21.53/15.57	18.55/12.84
(h)	✓	✓	✓	✓	✓	✓	<b>31.10/22.98</b>	<b>22.76/16.12</b>	<b>19.50/14.03</b>

If we change the attribute depth without scaling visual depth, the detector even performs worse than the one without affine-based data augmentation ( $AP_{3D}$  downgrades from 12.76 to 12.65). It is because this manner misleads the network in the training with incorrect depth targets. After revising visual depth, the network can benefit from the augmented training samples, boosting the performance from 19.05/12.76 to 21.74/15.48 AP under the moderate setting. We can see that the improper visual depth can result in larger impacts compared to the improper attribute depth on the final performance, as the visual depth has a larger value range. Finally, we obtain the best performance when employing the proper visual depth and attribute depth transformation strategy.

**Table 3.** Ablation for affine transformation based data augmentation. “Aff. Aug.” in the table denotes the affine-based data augmentation.

w/ Aff. Aug.	Scaled $d_{vis}$	Scaled $d_{att}$	AP <sub>BEV</sub> /AP <sub>3D</sub> (IoU=0.7)  <sub>R40</sub>		
			Easy	Moderate	Hard
			25.63/17.61	19.05/12.76	16.34/11.21
✓			26.97/18.98	19.80/14.33	17.71/12.00
✓		✓	22.23/15.27	18.98/12.65	16.52/10.75
✓	✓	✓	28.67/21.43	21.74/15.48	18.76/13.47
✓	✓		<b>31.10/22.98</b>	<b>22.76/16.12</b>	<b>19.50/14.03</b>

**Grid Size for Visual and Attribute Depth.** As described in Section 4, we divide the RoI image patch into  $m \times m$  grids, where each grid has a visual depth and an attribute depth. This paragraph investigates the impact brought by the grid size. When increasing grid size  $m$ , visual depths and attribute depths are becoming fine-grained. This tendency makes visual depth more intuitive, which is close to the pixel-wise depth. However, the fine-grained grid will lead to sub-optimal performance in terms of learning object attributes since the attributes focus on the overall object. It indicates that there exists a trade-off. Therefore,

**Table 4.** Ablation for the grid size on visual depth and attribute depth.

Grid size	AP <sub>BEV</sub> /AP <sub>3D</sub> (IoU=0.7)  <sub>R40</sub>		
	Easy	Moderate	Hard
1 × 1	26.11/19.39	18.76/13.01	16.00/11.34
3 × 3	27.03/19.73	19.91/14.33	18.10/12.25
5 × 5	29.20/21.36	21.53/15.13	18.61/12.53
9 × 9	30.21/21.78	22.28/14.98	18.93/12.47
13 × 13	29.88/21.67	21.97/15.29	18.96/12.72
19 × 19	28.20/20.36	21.53/15.13	18.61/12.53
<b>7 × 7</b>	<b>31.10/22.98</b>	<b>22.76/16.12</b>	<b>19.50/14.03</b>

we perform ablations on the grid size  $m$ , as shown in Table 4. We achieve the best performance when  $m$  is set to 7.

**Table 5.** Comparisons on pedestrian and cyclist categories on KITTI *val* set under IoU criterion 0.5. Our method brings obvious improvements to the baseline.

Approaches	Pedestrian, AP <sub>BEV</sub> /AP <sub>3D</sub>   <sub>R40</sub>			Cyclist, AP <sub>BEV</sub> /AP <sub>3D</sub>   <sub>R40</sub>		
	Easy	Moderate	Hard	Easy	Moderate	Hard
Baseline	5.97/5.13	4.75/3.88	3.87/3.05	5.11/4.27	2.68/2.46	2.50/2.17
Baseline+Ours	<b>8.86/7.27</b>	<b>7.01/5.87</b>	<b>5.46/4.89</b>	<b>6.13/5.54</b>	<b>3.09/2.59</b>	<b>2.67/2.49</b>

## 6 Conclusion

In this paper, we point out that the instance depth is coupled by visual depth clues and object inherent attributes. Its entangled nature makes it hard to be precisely estimated with the previous direct method. Therefore, we propose to decouple the instance depth into visual depths and attribute depths. This manner allows the network to learn different types of features for instance depth. At the inference stage, the instance depth is obtained by aggregating visual depth, attribute depth, and associated uncertainties. Using the decoupled depth, we can effectively perform affine transformation based data augmentation on the image, which is usually limited in previous works. Finally, extensive experiments demonstrate the effectiveness of our method.

## Acknowledgments

This work was supported in part by The National Key Research and Development Program of China (Grant Nos: 2018AAA0101400), in part by The National Nature Science Foundation of China (Grant Nos: 62036009, U1909203, 61936006, 61973271), in part by Innovation Capability Support Program of Shaanxi (Program No. 2021TD-05).

## References

1. Bewley, A., Sun, P., Mensink, T., Anguelov, D., Sminchisescu, C.: Range conditioned dilated convolutions for scale invariant 3d object detection. arXiv preprint arXiv:2005.09927 (2020)
2. Brazil, G., Liu, X.: M3d-rpn: Monocular 3d region proposal network for object detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 9287–9296 (2019)
3. Brazil, G., Pons-Moll, G., Liu, X., Schiele, B.: Kinematic 3d object detection in monocular video. In: European Conference on Computer Vision. pp. 135–152. Springer (2020)
4. Chai, Y., Sun, P., Ngiam, J., Wang, W., Caine, B., Vasudevan, V., Zhang, X., Anguelov, D.: To the point: Efficient 3d object detection in the range image with graph convolution kernels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16000–16009 (2021)
5. Chen, H., Huang, Y., Tian, W., Gao, Z., Xiong, L.: Monorun: Monocular 3d object detection by reconstruction and uncertainty propagation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10379–10388 (2021)
6. Chen, X., Kundu, K., Zhang, Z., Ma, H., Fidler, S., Urtasun, R.: Monocular 3d object detection for autonomous driving. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2147–2156 (2016)
7. Chen, X., Kundu, K., Zhu, Y., Ma, H., Fidler, S., Urtasun, R.: 3d object proposals using stereo imagery for accurate object class detection. IEEE transactions on pattern analysis and machine intelligence **40**(5), 1259–1272 (2017)
8. Chen, Y., Liu, S., Shen, X., Jia, J.: Fast point R-CNN. In: ICCV (2019)
9. Chen, Y., Tai, L., Sun, K., Li, M.: Monopair: Monocular 3d object detection using pairwise spatial relationships. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12093–12102 (2020)
10. Chu, X., Deng, J., Li, Y., Yuan, Z., Zhang, Y., Ji, J., Zhang, Y.: Neighbor-vote: Improving monocular 3d object detection through neighbor distance voting. arXiv preprint arXiv:2107.02493 (2021)
11. Dijk, T.v., Croon, G.d.: How do neural networks see depth in single images? In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2183–2191 (2019)
12. Ding, M., Huo, Y., Yi, H., Wang, Z., Shi, J., Lu, Z., Luo, P.: Learning depth-guided convolutions for monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11672–11681 (2020)
13. Fan, L., Xiong, X., Wang, F., Wang, N., Zhang, Z.: Rangedet: In defense of range view for lidar-based 3d object detection. arXiv preprint arXiv:2103.10039 (2021)
14. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 3354–3361. IEEE (2012)
15. Hu, M., Wang, S., Li, B., Ning, S., Fan, L., Gong, X.: Penet: Towards precise and efficient image guided depth completion. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 13656–13662. IEEE (2021)
16. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? Advances in neural information processing systems **30** (2017)
17. Ku, J., Pon, A.D., Waslander, S.L.: Monocular 3d object detection leveraging accurate proposals and shape reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11867–11876 (2019)

18. Kumar, A., Brazil, G., Liu, X.: Groomed-nms: Grouped mathematically differentiable nms for monocular 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8973–8983 (2021)
19. Li, B., Zhang, T., Xia, T.: Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916* (2016)
20. Li, P., Zhao, H.: Monocular 3d detection with geometric constraint embedding and semi-supervised training. *IEEE Robotics and Automation Letters* **6**(3), 5565–5572 (2021)
21. Li, P., Zhao, H., Liu, P., Cao, F.: Rtm3d: Real-time monocular 3d detection from object keypoints for autonomous driving. *arXiv preprint arXiv:2001.03343* (2020)
22. Liu, X., Xue, N., Wu, T.: Learning auxiliary monocular contexts helps monocular 3d object detection. *arXiv preprint arXiv:2112.04628* (2021)
23. Liu, Y., Yixuan, Y., Liu, M.: Ground-aware monocular 3d object detection for autonomous driving. *IEEE Robotics and Automation Letters* **6**(2), 919–926 (2021)
24. Liu, Z., Zhou, D., Lu, F., Fang, J., Zhang, L.: Autoshape: Real-time shape-aware monocular 3d object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 15641–15650 (2021)
25. Lu, Y., Ma, X., Yang, L., Zhang, T., Liu, Y., Chu, Q., Yan, J., Ouyang, W.: Geometry uncertainty projection network for monocular 3d object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3111–3121 (2021)
26. Ma, X., Liu, S., Xia, Z., Zhang, H., Zeng, X., Ouyang, W.: Rethinking pseudo-lidar representation. *arXiv preprint arXiv:2008.04582* (2020)
27. Ma, X., Wang, Z., Li, H., Zhang, P., Ouyang, W., Fan, X.: Accurate monocular 3d object detection via color-embedded 3d reconstruction for autonomous driving. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 6851–6860 (2019)
28. Ma, X., Zhang, Y., Xu, D., Zhou, D., Yi, S., Li, H., Ouyang, W.: Delving into localization errors for monocular 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4721–4730 (2021)
29. Manhardt, F., Kehl, W., Gaidon, A.: Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2069–2078 (2019)
30. Mao, J., Xue, Y., Niu, M., Bai, H., Feng, J., Liang, X., Xu, H., Xu, C.: Voxel transformer for 3d object detection. *arXiv preprint arXiv:2109.02497* (2021)
31. Mousavian, A., Anguelov, D., Flynn, J., Kosecka, J.: 3d bounding box estimation using deep learning and geometry. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7074–7082 (2017)
32. Noh, J., Lee, S., Ham, B.: Hvpr: Hybrid voxel-point representation for single-stage 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14605–14614 (2021)
33. Park, D., Ambrus, R., Guizilini, V., Li, J., Gaidon, A.: Is pseudo-lidar needed for monocular 3d object detection? In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3142–3152 (2021)
34. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: *Advances in neural information processing systems*. pp. 8026–8037 (2019)
35. Peng, L., Liu, F., Yan, S., He, X., Cai, D.: Ocm3d: Object-centric monocular 3d object detection. *arXiv preprint arXiv:2104.06041* (2021)



36. Peng, L., Liu, F., Yu, Z., Yan, S., Deng, D., Yang, Z., Liu, H., Cai, D.: Lidar point cloud guided monocular 3d object detection. *arXiv preprint arXiv:2104.09035* (2021)
37. 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 conference on computer vision and pattern recognition*. pp. 918–927 (2018)
38. Qin, Z., Wang, J., Lu, Y.: Monogrnnet: A geometric reasoning network for monocular 3d object localization. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 8851–8858 (2019)
39. Reading, C., Harakeh, A., Chae, J., Waslander, S.L.: Categorical depth distribution network for monocular 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8555–8564 (2021)
40. Roddick, T., Kendall, A., Cipolla, R.: Orthographic feature transform for monocular 3d object detection. *arXiv preprint arXiv:1811.08188* (2018)
41. 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: *CVPR*. pp. 10529–10538 (2020)
42. Shi, S., Wang, X., Li, H.: PointRCNN: 3D object proposal generation and detection from point cloud. In: *CVPR*. pp. 770–779 (2019)
43. 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* (2020)
44. Shi, X., Ye, Q., Chen, X., Chen, C., Chen, Z., Kim, T.K.: Geometry-based distance decomposition for monocular 3d object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 15172–15181 (2021)
45. Simonelli, A., Bulò, S.R., Porzi, L., López-Antequera, M., Kotschieder, P.: Disentangling monocular 3d object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1991–1999 (2019)
46. Wang, L., Du, L., Ye, X., Fu, Y., Guo, G., Xue, X., Feng, J., Zhang, L.: Depth-conditioned dynamic message propagation for monocular 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 454–463 (2021)
47. Wang, L., Zhang, L., Zhu, Y., Zhang, Z., He, T., Li, M., Xue, X.: Progressive coordinate transforms for monocular 3d object detection. *Advances in Neural Information Processing Systems* **34** (2021)
48. Wang, T., Xinge, Z., Pang, J., Lin, D.: Probabilistic and geometric depth: Detecting objects in perspective. In: *Conference on Robot Learning*. pp. 1475–1485. PMLR (2022)
49. Wang, Y., Chao, W.L., Garg, D., Hariharan, B., Campbell, M., Weinberger, K.Q.: Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8445–8453 (2019)
50. Yan, Y., Mao, Y., Li, B.: SECOND: Sparsely embedded convolutional detection. *Sensors* **18**(10), 3337 (2018)
51. Yang, Z., Sun, Y., Liu, S., Jia, J.: 3DSSD: Point-based 3D single stage object detector. In: *CVPR*. pp. 11040–11048 (2020)
52. Yang, Z., Sun, Y., Liu, S., Shen, X., Jia, J.: STD: Sparse-to-dense 3D object detector for point cloud. In: *ICCV*. pp. 1951–1960 (2019)
53. 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 (June 2021)

54. Yu, F., Wang, D., Shelhamer, E., Darrell, T.: Deep layer aggregation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2403–2412 (2018)
55. Zhang, Y., Lu, J., Zhou, J.: Objects are different: Flexible monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3289–3298 (2021)
56. 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. pp. 14494–14503 (2021)
57. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019)
58. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4490–4499 (2018)
59. Zhou, Y., He, Y., Zhu, H., Wang, C., Li, H., Jiang, Q.: Monocular 3d object detection: An extrinsic parameter free approach. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7556–7566 (2021)
60. Zou, Z., Ye, X., Du, L., Cheng, X., Tan, X., Zhang, L., Feng, J., Xue, X., Ding, E.: The devil is in the task: Exploiting reciprocal appearance-localization features for monocular 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2713–2722 (2021)