

Densely Constrained Depth Estimator for Monocular 3D Object Detection

Yingyan Li^{1,2,4,5}, Yuntao Chen, Jiawei He^{1,2,4}, and Zhaoxiang Zhang^{1,2,3,4,5}

¹ Institute of Automation, Chinese Academy of Sciences (CASIA)

² University of Chinese Academy of Sciences (UCAS)

³ Centre for Artificial Intelligence and Robotics, HKISI-CAS

⁴ National Laboratory of Pattern Recognition (NLPR)

⁵ School of Future Technology, UCAS

{liyinyan2021,hejiawei2019,zhaoxiang.zhang}@ia.ac.cn,
chenyuntao08@gmail.com

Abstract. Estimating accurate 3D locations of objects from monocular images is a challenging problem because of lacking depth. Previous work shows that utilizing the object’s keypoint projection constraints to estimate multiple depth candidates boosts the detection performance. However, the existing methods can only utilize vertical edges as projection constraints for depth estimation. So these methods only use a small number of projection constraints and produce insufficient depth candidates, leading to inaccurate depth estimation. In this paper, we propose a method that utilizes *dense* projection constraints from edges of any direction. In this way, we employ much more projection constraints and produce considerable depth candidates. Besides, we present a graph matching weighting module to merge the depth candidates. The proposed method *DCD* (Densely Constrained Detector) achieves state-of-the-art performance on the KITTI and WOD benchmarks. Code is released at <https://github.com/BraveGroup/DCD>.

Keywords: Monocular 3D object detection, dense geometric constraint, message passing, graph matching

1 Introduction

Monocular 3D detection [17,45,7,51] has become popular because images are large in number, easy to obtain, and have dense information. Nevertheless, the lack of depth information in monocular images is a fatal problem for 3D detection. Some methods [2,22] use deep neural networks to regress the 3D bounding boxes directly, but it is challenging to estimate the 3D locations of the objects from 2D images. Another line of work [44,8,30] employs a pre-trained depth estimator. However, training the depth estimator is separated from the detection part, requiring a large amount of additional data. In addition, some works [17,51,23] use geometric constraints, i.e., regresses the 2D/3D edges, and then estimates the object’s depth from the 2D-3D edge projection constraints.

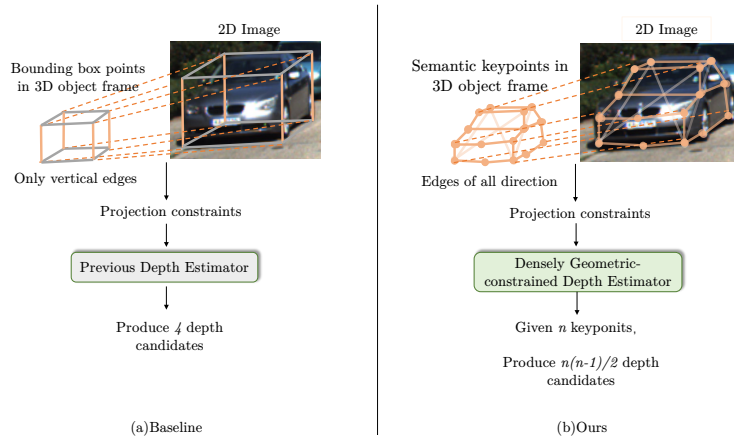


Fig. 1. The object’s depth is estimated by 2D-3D edge projection constraints. This figure compares the involved edges in the object’s depth estimation between (a) previous work and (b) ours. The previous work only deals with vertical edges. Our work is able to handle the edges of any direction.

These works employ 3D shape prior information and exhibit state-of-the-art performance, which is worthy of future research.

A problem of the previous work is that their geometric constraints are insufficient. Specifically, some existing methods [51,25,50] estimate the height of the 2D bounding box and the 3D bounding box, and then generate the depth candidates of an object from 2D-3D height projection constraints. The final depth is produced by weighting all the depth candidates. As Fig. 1 shows, this formulation is only suitable for the vertical edges, which means they only utilize a tiny amount of constraints and 3D prior, leading to inaccurate depth estimations.

Some of the depth candidates are of low quality, so weighting is needed. However, the previous work’s weighting methods are suboptimal. Since the final depth is derived from the weighted average of depth candidates, the weight should reflect the quality of each depth candidate. Existing methods [23] use a branch to regress the weight of each depth candidate directly, and this branch is paralleled to the keypoints regression branch. So the weighting branch does not know each keypoint’s quality. Some work predicts the uncertainty of each depth to measure the quality of the depth and use the uncertainty to weight [25,51]. However, they obtain the uncertainty of each depth candidate independently, and they do not supervise the weight explicitly.

To address the problem of insufficient geometric constraints, we propose a *Densely Geometric-constrained Depth Estimator (DGDE)*. *DGDE* can estimate depth candidates from projection constraints provided by edges of any direction, no more limited to the vertical edges. This estimator allows better use of the shape information of the object. In addition, training the neural network with

abundant 2D-3D projection constraints helps the neural network understand the mapping relationship from the 2D plane to the 3D space.

To weight the depth candidates properly, we propose a new depth candidates weighting module that employs graph matching, named *Graph Matching Weighting module*. We construct *complete graphs* based on 2D and 3D semantic keypoints. In a 2D keypoint graph, the 2D keypoint coordinates are placed on the vertices, and an edge represents a pair of 2D keypoints. The 3D keypoint graph is constructed in the same way. We then match the 2D edges and 3D edges and produce the matching scores. The 2D-3D edge matching score is used as the weight of the corresponding depth candidate. These weights are explicitly supervisable. Moreover, the information of the entire 2d/3d edges is used to generate each 2d-3d edge matching score.

In summary, our main contributions are:

1. We propose a *Dense Geometric-constrained Depth Estimator (DGDE)*. Different from the previous methods, *DGDE* estimates depth candidates utilizing projection constraints of edges of any direction. Therefore, considerable 2D-3D projection constraints are used, producing considerable depth candidates. We produce high-quality final depth based on these candidates.
2. We propose an effective and interpretable *Graph Matching Weighting module (GMW)*. We construct the 2D/3D graph from 2D/3D keypoints respectively. Then we regard the graph matching score of the 2D-3D edge as the weight of the corresponding depth candidate. This strategy utilizes all the keypoints' information and produces explicitly supervised weights.
3. We localize each object more accurately by weighting the estimated depth candidates with corresponding matching scores. Our *Densely Constrained Detector (DCD)* achieves state-of-the-art performance on the KITTI and Waymo Open Dataset (WOD) benchmarks.

2 Related Work

Monocular 3D Object Detection. Monocular 3D object detection [16,5,7,13] becomes more and more popularity because monocular images are easy to obtain. The existing methods can be divided into two categories: single-center-point-based and multi-keypoints-based.

Single-center-point-based methods [22,54,52] use the object's center point to represent an object. In detail, M3D-RPN [2] proposes a depth-aware convolutional layer with the estimated depth. MonoPair [7] discovers that the relationships between nearby objects are useful for optimizing the final results. Although single-center-point-based methods are simple and fast, location regression is unstable because only one center point is utilized. Therefore, the multi-keypoints-based methods [23,51,17] are drawing more and more attention recently.

Multi-keypoints-based methods predict multiple keypoints for an object. More keypoints provide more projection constraints. The projection constraints are useful for training the neural network because constraints build the mapping relationship from the 2D image plane to the 3D space. Deep MANTA [5] defines

4 wireframes as templates for matching cars, while 3D-RCNN [16] proposes a render-and-compare loss. Deep3DBox [29] utilizes 2D bounding boxes as constraints to refine 3D bounding boxes. KM3D [17] localizes the objects utilizing eight bounding boxes points projection. MonoJSG [20] constructs an adaptive cost volume with semantic features to model the depth error. AutoShape [23], highly relevant to this paper, regresses 2D and 3D semantic keypoints and weighs them by predicted scores from a parallel branch. MonoDDE [19] is a concurrent work. It uses more geometric constraints than MonoFlex [51]. However, it only uses geometric constraints based on the object’s center and bounding box points, while we use dense geometric constraints derived from semantic keypoints. In this paper, we propose a novel edge-based depth estimator. The estimator produces depth candidates by projection constraints, consisting of edges of any direction. We weight each edge by graph matching to derive the final depth of the object.

Graph Matching and Message Passing. Graph matching [37] is defined as matching vertices and edges between two graphs. This problem is formulated as a Quadratic Assignment Problem (QAP) [3] originally. It has been widely applied in different tasks, such as multiple object tracking [14], semantic keypoint matching [53] and point cloud registration [10]. In the deep learning era, graph matching has become a differentiable module. The spectral relaxation [49], quadratic relaxation [14] and Lagrange decomposition [32] are common in use. However, one simple yet effective way to implement the graph matching layer is using the Sinkhorn [4,33,42] algorithm, which is used in this paper. This paper utilizes a graph matching module to achieve the message passing between keypoints. This means that we calculate the weighting score not only from the keypoint itself but also by taking other keypoints’ regression quality into consideration, which is a kind of message passing between keypoints. All the keypoints on the object help us judge each keypoint’s importance from the perspective of the whole object. Message passing [12] is a popular design, e.g., in graph neural network [1] and transformer [39]. The message passing module learns to aggregate features between nodes in the graph and brings the global information to each node. In the object detection task, Relation Networks [15] is a pioneer work utilizing message passing between proposals. Message passing has also been used for 3D vision. Message passing between frames [46], voxels [9] and points [34] is designed for LiDAR-based 3D object detection. However, for monocular 3D object detection, message passing is seldom considered. In recent work, PGD [43] conducts message passing in the geometric relation graph. But it does not consider the message passing within the object.

3 Methodology

The overview of our framework is in Fig. 2. We employ a single-stage detector [51] to detect the object’s 3D attributes from the monocular image. We propose the *Densely Geometric-constrained Depth Estimator (DGDE)*, which can calculate the depth from any direction’ 2D-3D edge. The *DGDE* can effectively utilize the semantic keypoints of the object and produce many depth candidates. Be-

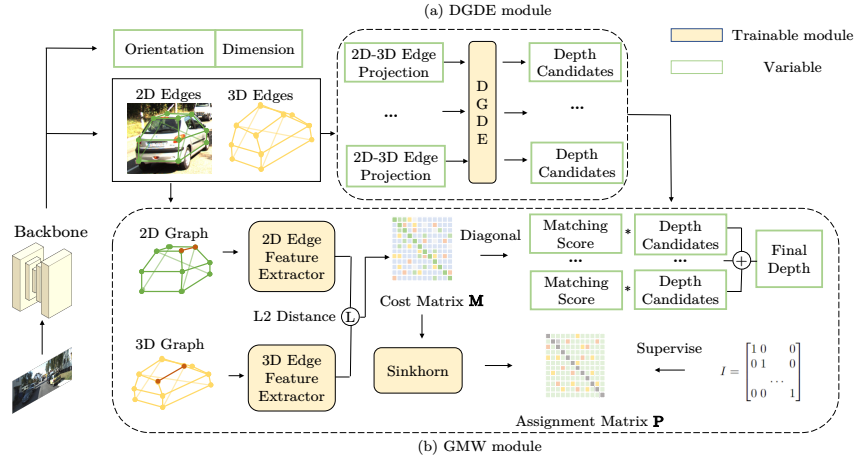


Fig. 2. Overview of our framework. (a) We propose a method *Densely Geometric-constrained Depth Estimator* (DGDE). DGDE is able to estimate the object’s depth candidates from 2D-3D projection constraints of edges of any direction. (b) *Graph Matching Weighting module* (GMW) obtains the weights of estimated depth candidates by Graph Matching. A robust depth is derived from combing the multiply depth candidates with corresponding weights.

sides, we utilize the regressed 2D edges, 3D edges, and orientation as the input for our 2D-3D edge Graph Matching network. Our *Graph Matching Weighting module* (GMW) matches each 2D-3D edge and produces a matching score. By combining the multiple depths with their corresponding matching scores, we can finally generate a robust depth for the object.

3.1 Geometric-based 3D Detection Definition

The geometric-based monocular 3D object detection estimates the object’s location by 2D-3D projection constraints. Specifically, the network predicts the object’s dimension (h, w, l) , rotation r_y , since autonomous driving datasets generally assume that the ground is flat. Assuming an object has n semantic keypoints, we regress the i -th ($i = 1, 2, \dots, n$) keypoint’s 2D coordinate (u^i, v^i) in image coordinate and 3D coordinate (x_o^i, y_o^i, z_o^i) in object frame. The object frame’s coordinate origin is the object’s center point. The i -th 2D-3D keypoint projection constraint is established from $(u^i, v^i, x_o^i, y_o^i, z_o^i, r_y)$. Given n semantic 2D-3D keypoint projection constraints, it is an overdetermined problem for solving 3D object location (x_c, y_c, z_c) , which is the translation vector for transforming the points from the object frame into the camera frame. The method of generating semantic keypoints of each object is adapted from [23]. We establish a few car models by PCA and refine the models by the 3D points segmented from the points cloud and 2D masks. After we obtain the keypoints, we can use our *DGDE* to estimate the object’s depth from the keypoint projection constraints.

3.2 Densely Geometric-constrained Depth Estimation

While previous depth estimation methods [51] only take vertical edges into account, our *DGDE* can handle edges of any direction. Therefore, we are able to utilize much more constraints to estimate the depth of each depth candidate.

Next, we will show the details of estimating dense depth candidates of an object from 2D-3D keypoint projection constraints. The solution is based on the keypoint's projection relationship from 3D space to the 2D image. The i -th ($i = 1, 2, \dots, n$) keypoint's 3D coordinate (x_o^i, y_o^i, z_o^i) is defined in the object frame and is projected on a 2D image plane by the equation:

$$s_i[u_i, v_i, 1]^T = \mathbf{K}[\mathbf{R}|\mathbf{t}][x_o^i, y_o^i, z_o^i, 1]^T, \quad (1)$$

where s_i is the i -th keypoint's depth, \mathbf{K} is the camera intrinsic matrix and $\mathbf{K}, \mathbf{R}, \mathbf{t}$ is represented as:

$$\mathbf{K} = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix}, \mathbf{R} = \begin{bmatrix} \cos r_y & 0 & \sin r_y \\ 0 & 1 & 0 \\ -\sin r_y & 0 & \cos r_y \end{bmatrix}, \mathbf{t} = [x_c, y_c, z_c]^t. \quad (2)$$

By Eq. (1) and Eq. (2), the equation of i -th keypoint's projection constraint is denoted as:

$$\begin{cases} s_i = z_c - x_o^i \sin r_y + z_o^i \cos r_y, \\ \tilde{u}_i(z_c - x_o^i \sin r_y + z_o^i \cos r_y) = x_c + x_o^i \cos r_y + z_o^i \sin r_y, \\ \tilde{v}_i(z_c - x_o^i \sin r_y + z_o^i \cos r_y) = y_c + y_o^i, \end{cases} \quad (3)$$

where $\tilde{u}_i = \frac{u_i - c_x}{f_x}$, $\tilde{v}_i = \frac{v_i - c_y}{f_y}$. Intuitively, $(z_c - x_o^i \sin r_y + z_o^i \cos r_y)$ means an object's i -th 3D keypoint's Z coordinate (i.e., depth) in the camera coordinate. $(x_c + x_o^i \cos r_y + z_o^i \sin r_y)$ means 3D keypoint's X coordinate while $(y_c + y_o^i)$ means its Y coordinate.

Similarly, the j -th ($j = 1, 2, \dots, n$) projection constraint is denoted as:

$$\begin{cases} s_j = z_c - x_o^j \sin r_y + z_o^j \cos r_y, \\ \tilde{u}_j(z_c - x_o^j \sin r_y + z_o^j \cos r_y) = x_c + x_o^j \cos r_y + z_o^j \sin r_y, \\ \tilde{v}_j(z_c - x_o^j \sin r_y + z_o^j \cos r_y) = y_c + y_o^j. \end{cases} \quad (4)$$

From Eq. (3) and Eq. (4), we can densely obtain the z_c from the i -th, j -th, $i \neq j$ keypoint (i.e., $edge_{ij}$) projection constraints as:

$$z_c^{ij} = \begin{cases} \frac{l_i - l_j}{\tilde{u}_i - \tilde{u}_j}, \\ \frac{h_i - h_j}{\tilde{v}_i - \tilde{v}_j}, \end{cases} \quad (5)$$

where $l_i = x_o^i \cos(r_y) + z_o^i \sin(r_y) + u^i(x_o^i \sin(r_y) - z_o^i \cos(r_y))$ and $h_i = y_o^i + v^i(x_o^i \sin(r_y) - z_o^i \cos(r_y))$. This equation reveals that depth can be calculated by

the projection constraints of an edge of any direction. Given z_c , we can estimate x_c, y_c from Eq. (3) as $x_c^i = u_i z_c - l_i, y_c^i = v_i z_c - h_i$.

We generate $m = n(n-1)/2$ depth candidates given n keypoints. It is inevitable to meet some low-quality depth candidates in such a large number of depths. Therefore, an appropriate weighting method is necessary to ensemble these depth candidates.

3.3 Depth Weighting by Graph Matching

As we estimate the depth candidate $z_c^{ij}(i, j = 1, \dots, n)$ for the object o from *DGDE*, the final depth z_c of the object can be weighted from these depth estimations according to the estimation quality $w_{i,j}$, as

$$z_c = \sum_{i < j} w_{i,j} z_c^{ij}. \quad (7)$$

In this section, we propose a new weighting method, called *Graph Matching Weighting module (GMW)*.

Graph Construction and Edge Feature extraction. We construct 2D keypoint graph $\mathcal{G}_{2d} = (\mathcal{V}_{2d}, \mathcal{E}_{2d})$ and 3D keypoint graph $\mathcal{G}_{3d} = (\mathcal{V}_{3d}, \mathcal{E}_{3d})$. In \mathcal{G}_{2d} , each vertex $i \in \mathcal{V}_{2d}$ denotes a predicted keypoint ${}^{(2d)}\mathbf{p}^i = [u^i, v^i]$ in image coordinate and the edge ${}^{(2d)}e^{i,j} \in \mathcal{E}_{2d}$ denotes the pair of ${}^{(2d)}\mathbf{p}^i$ and ${}^{(2d)}\mathbf{p}^j$, the edge feature ${}^{(2d)}\mathbf{f}^{i,j}$ is extracted from the 2D coordinate ${}^{(2d)}\mathbf{p}^i$ and ${}^{(2d)}\mathbf{p}^j$. The 3D keypoint graph is almost similar to the 2D keypoint graph. The only difference is that the vertex $i' \in \mathcal{V}_{3d}$ denotes the 3D coordinate ${}^{(3d)}\mathbf{p}^{i'} = [x_o^{i'}, y_o^{i'}, z_o^{i'}]$.

Following [47], the 2D and 3D edge feature extractor are as

$${}^{(2d)}\mathbf{f}_k^{i,j} = \text{ReLU}_k(\text{BN}_k(\text{CN}_k({}^{(2d)}\text{FC}_k({}^{(2d)}\mathbf{f}_{k-1}^{i,j})))), \quad (8)$$

$${}^{(3d)}\mathbf{f}_k^{i',j'} = \text{ReLU}_k(\text{BN}_k(\text{CN}_k({}^{(3d)}\text{FC}_k({}^{(3d)}\mathbf{f}_{k-1}^{i',j'})))), \quad (9)$$

where $k \in \{1, \dots, K\}$ denotes the index of layers, and FC, CN, BN, ReLU denote fully-connected layer, Context Normalization [47], Batch Normalization, and ReLU, respectively. It is worth mentioning that Context Normalization extracts the global information of all edges. The input of the edge feature extractor is $\mathbf{f}_0^{i,j} = [\mathbf{p}^i, \mathbf{p}^j]$, where $[\cdot]$ denotes the concatenation of the vectors. The output of edge feature extractor ${}^{(2d)}\mathbf{f}^{i,j}$ and ${}^{(3d)}\mathbf{f}^{i',j'}$ should be L2-normalized to $[0, 1]$.

Graph matching layer. Given the extracted 2D and 3D edge features, the Cost Matrix $\mathbf{M} \in \mathbb{R}^{m \times m}$ is calculated from the L2 distance between each 2D edge feature ${}^{(2d)}\mathbf{f}^{i,j}$ on the edge s and 3D edge feature ${}^{(3d)}\mathbf{f}^{i',j'}$ on the edge t :

$$\mathbf{M}_{s,t} = \text{L2}({}^{(2d)}\mathbf{f}^{i,j}, {}^{(3d)}\mathbf{f}^{i',j'}), (s, t \in \{1, \dots, m\}), \quad (10)$$

where m denotes the number of edges. Then we take \mathbf{M} as the input of declarative Sinkhorn layer[4] to gain the Assignment Matrix \mathbf{P} . The Sinkhorn layer iteratively optimizes \mathbf{P} by minimizing the objective function:

$$\mathcal{F}(\mathbf{M}, \mathbf{P}) = \sum_{s=1}^m \sum_{t=1}^m (\mathbf{M}_{s,t} \mathbf{P}_{s,t} + \alpha \mathbf{P}_{s,t} (\log \mathbf{P}_{s,t} - 1)), \quad (11)$$

where $\mathbf{P}_{s,t}(\log \mathbf{P}_{s,t} - 1)$ is a regularization term and α is the coefficient. $\mathbf{P} \in U(\mathbf{a}, \mathbf{b})$ as:

$$U(\mathbf{a}, \mathbf{b}) = \{\mathbf{X} \in \mathbb{R}_+^{m \times m} | \mathbf{X}\mathbf{1}^m = \mathbf{1}^m, \mathbf{X}^T\mathbf{1}^m = \mathbf{1}^m\}, \quad (12)$$

where $\mathbf{1}^m$ is an m -dimensional vector with all values to be 1. Note that, Sinkhorn is a differentiable graph matching solver that can make the whole pipeline learnable. When calculating the final depth z_c according to Eq. (7), the weight $\mathbf{w} = \text{Softmax}(1/\text{Diag}(\mathbf{M}))$, where $\text{Diag}(\mathbf{M})$ means the vector consisting of diagonal elements of matrix \mathbf{M} . Intuitively, it means that we take the similarity of the 2D and 3D edges with the same semantic label as the prediction quality.

Loss function. We design regression loss \mathcal{L}_m^r to supervise the final weighted depth z_c , and classification loss \mathcal{L}_m^c to supervise the assignment matrix \mathbf{P} of graph matching output to be an identity matrix. Specifically, \mathcal{L}_m^c is the Binary Entropy Loss (BCE), \mathcal{L}_m^r is an L1 loss:

$$\mathcal{L}_m^c = \sum_{s=1}^m \sum_{t=1}^m \text{BCE}(\mathbf{P}_{s,t}, \mathbf{P}_{s,t}^*), \mathcal{L}_m^r = \left| \sum_{i < j} w_{i,j} z_c^{ij} - z_c^* \right|, \quad (13)$$

where $\mathbf{P}^* = \mathbf{I}$ is the ground truth assignment matrix, z_c^* is the ground truth depth of the object. The final matching loss is $\mathcal{L}_m = \mathcal{L}_m^c + \beta \mathcal{L}_m^r$, where β is a hyper-parameter.

4 Experiments

4.1 Setup

Dataset. We evaluate our method on the KITTI [11] and Waymo Open Dataset v1.2 (WOD) [38]. The KITTI [11] dataset is collected from Europe Streets. It consists of 7481 images for training and 7518 images for testing. We divide the training data into a train set (3712 images) and a validation set (3769 images) as in [55]. Waymo Open Dataset (WOD) [38] has 798 training sequences, 202 validation sequences, and 150 test sequences. The dataset contains images captured by 5 high-resolution cameras in complex environments and is much more challenging than KITTI [11] dataset. We only use the images from the FRONT camera for training and evaluation. The training set has 158,081 images and the validation set has 39,848 images.

Evaluation Metrics. For the KITTI dataset, we compare our methods with previous methods on the *test* set using $AP_{3D|R_{40}}$ result from the test server. In ablation studies, the results on *val* set are reported. For the WOD, we focus on the category of vehicle. Following the official evaluation criteria, we compare the performance with the state-of-the-art methods using average precision (AP) and average precision weighted by heading (APH) metrics. The results are shown in Table 3. The objects are classified into two difficulty levels (LEVEL_1, LEVEL_2) according to the object’s points’ number under the LiDAR sensor.

Table 1. The result on the KITTI test server compared with other public methods in recent years.

Methods	Reference	Category	$AP_{3D R40 IoU@0.7}$			$AP_{BEV R40 IoU@0.7}$		
			Easy	Mod.	Hard	Easy	Mod.	Hard
PatchNet [26]	ECCV20	Pretrained Depth	15.68	11.12	10.17	22.97	16.86	14.97
D4LCN [8]	CVPR20		16.65	11.72	9.51	22.51	16.02	12.55
DDMP-3D [40]	CVPR21		19.71	12.78	9.80	28.08	17.89	13.44
CaDDN [31]	CVPR21	LiDAR Auxiliary	19.17	13.41	11.46	27.94	18.91	17.19
RTM3D [18]	ECCV20	Directly Regress	14.41	10.34	8.77	19.17	14.20	11.99
Movi3D [36]	ECCV20		15.19	10.90	9.26	22.76	17.03	14.85
Ground-Aware [21]	RAL21		21.65	13.25	9.91	29.81	17.98	13.08
MonoDLE [28]	CVPR21		17.23	12.26	10.29	24.79	18.89	16.00
MonoRCNN [35]	ICCV21		18.36	12.65	10.03	25.48	18.11	14.10
MonoEF [56]	CVPR21		21.29	13.87	11.71	29.03	19.70	17.26
MonoRUn [6]	CVPR21	Geometric-based	19.65	12.30	10.58	27.94	17.34	15.24
AutoShape [23]	ICCV21		22.47	14.17	11.36	30.66	20.08	15.59
GUPNet [25]	ICCV21		22.20	15.02	13.12	30.29	21.19	18.20
MonoFlex(Baseline) [51]	CVPR21		19.94	13.89	12.07	28.23	19.75	16.89
DCD(Ours)	ECCV22		23.81	15.90	13.21	32.55	21.50	18.25

4.2 Implementation Details

Detection Framework. We apply the 3D object detection framework following [51], which uses DLA-34 [48] as backbone. We use MultiBin loss [29] for rotation. L1 loss is adopted to estimate dimension, 2D/3D keypoints and depth.

Keypoints. The source of keypoints is discussed in Sec. 3.1. We use 73 keypoints in total consisting of the following parts: (1) 63 semantic keypoints; (2) 8 bounding box corners and the top center and the bottom center of the 3D bounding box. There are 2628 unique keypoint pairs that can be generated from 73 keypoints, so we can obtain 2628 depth estimations at most for each object. For robustness, we select 1500 depth estimations as the final candidates for weighting. The details of the selection strategy are in Sec. 4.4.

Training and Inference. For the KITTI dataset, all the input images are padded into 1280×384 . We train the model using AdamW [24] optimizer with an initial learning rate of $3e-4$ for 100 epochs. The learning rate decays by $10 \times$ at 80 and 90 epochs. We train the model on 2 RTX2080Ti GPUs and the batch size is 8. We train the weighting network (i.e., matching network) separately. The weighting network employs the AdamW optimizer with learning rate $1e-4$ and weight decay $1e-5$. We first train using the classification loss in the weighting network for 50 epochs and add the regression loss for another 50 epochs. During inference, only monocular images are needed. For the WOD, the input size of the images is 1920×1280 . We ignore objects whose 2D bounding box’s width or height is less than 20 pixels. We train our detection model for 20 epochs with 8 RTX2080Ti GPUs. The batch size is 8 and the learning rate is set to $8e-5$, decayed $10 \times$ at the 18-th epoch. The rest of the experiment settings are the same as KITTI.



Fig. 3. The qualitative results on the KITTI *val* set. Rather than representing an object as a bounding box, we utilize semantic keypoints to represent an object. The red boxes represent the ground truth, while the green boxes represent the prediction.

4.3 Comparison with State-of-the-arts Methods

Table 1 shows the comparison with other state-of-the-art methods on the KITTI [11] test set. Our method (*DCD*) achieves state-of-the-art performance on both AP_{3D} and AP_{BEV} . We surpass the directly-regress-depth-based method and pretrain-depth-estimator-based methods by a large margin. It reveals that geometric constraints are critical to accurately locating objects. Compared with other geometric-based methods, we still have a significant improvement. We outperform our baseline MonoFlex [51] by **2.01** in AP_{3D} Mod. level, which reveals the importance of sufficient geometric constraints for monocular 3D detection. We also surpass other geometric-based methods such as AutoShape [23] and GUPNet [25] thanks to the dense geometric constraints and effective weighting method. The result of Pedestrian and Cyclist is in Table 2. Using only bounding box corners as input, we can still observe an improvement over MonoFlex [51]. It shows that our method can handle the problem that some objects are hard to obtain semantic keypoints (without CAD models or non-rigid). The qualitative result is in Fig. 3.

We also achieve state-of-the-art performance on WOD [38] as Table 3 shows. We surpass the previous state-of-the-art methods such as CaDDN [31] and PCT [41]. We also re-implement MonoFlex [51] on WOD [38] for a fair comparison. Compared with the baseline method, we improve the AP IoU@0.7 by **0.87**.

4.4 Ablation Studies

Keypoints enable better depth estimation. We utilize multiple keypoints rather than bounding boxes to represent an object. The keypoints can accurately reflect the object’s outline, which provides meaningful shape prior and abundant

Table 2. The $AP_{3D|R40}$ results for Pedestrian and Cyclist on KITTI *test* set. We use the bounding box corners as the input of *DGDE*.

Method	Pedestrian			Cyclist		
	Easy	Mod.	Hard	Easy	Mod.	Hard
M3D-RPN [2]	4.92	3.48	2.94	0.94	0.65	0.47
MonoPair [7]	10.02	6.68	5.53	3.79	2.12	1.83
MonoFlex [51]	9.43	6.31	5.26	4.17	2.35	2.04
AutoShape [23]	5.46	3.74	3.03	5.99	3.06	2.70
Ours	10.37	6.73	6.28	4.72	2.74	2.41

Table 3. The result on WOD [38] *val* set. *Italics*: These methods utilize the whole *train* set, while the others use 1/3 amount of images in *train* set. ‡: M3D-RPN is re-implemented by [31]. †: PatchNet is re-implemented by [41]. *: MonoFlex is our baseline and re-implemented ourselves.

Difficulty	Method	3D AP (IoU@0.7)	3D APH (IoU@0.7)	3D AP (IoU@0.5)	3D APH (IoU@0.5)
LEVEL_1/LEVEL_2	M3D-RPN‡ [2]	0.35/0.33	0.34/0.33	3.79/3.61	3.63/3.46
	PatchNet† [27]	0.39/0.38	0.37/0.36	2.92/2.42	2.74/2.28
	PCT [41]	0.89/0.66	0.88/0.66	4.20/4.03	4.15/3.99
	CaDDN [31]	5.03/4.49	4.99/4.45	17.54/16.51	17.31/16.28
	<i>MonoFlex</i> * [51]	11.70/10.96	11.64/10.90	32.26/30.31	32.06/30.12
	<i>DCD(Ours)</i>	12.57/11.78	12.50/11.72	33.44/31.43	33.24/31.25

information for depth estimation. To show the benefits gained from dense geometric constraints, we visualize the predicted depths’ error on KITTI *val* set as Fig. 4 shows.

The effectiveness of DGDE. To discover the inner workings of multiple 2D-3D keypoints projection constraints, we apply the DGDE on the state-of-the-art method MonoFlex [51]. MonoFlex predicts eight bounding box corners and two top-down center points. In addition to the ten keypoints, we add another regression branch on MonoFlex to predict 63 semantic keypoints. Basically, they are sampled from the model surface representing the rough skeletons.

In Table 4, the improvements from (d) to (e) are significant (+1.30 on the Hard level) even with only 10 keypoints. With more keypoints ((c) and (f)), DGDE achieves holistic improvements on both the uncertainty-based weighting method and our matching-based method.

The more keypoints, the better performances. Our depth estimator can produce depth candidates by edge projection constraints of arbitrary directions. To fully realize the potential of our depth estimator, we use all the extra 63 semantic keypoints. Thus, it is easy to generate numerous edges and obtain considerable depth candidates (2628). With such a large number of depth candidates, it is more likely to generate an accurate and robust final depth. In Table 4, the model (f) with all keypoints outperforms our baseline by a large margin of 2.31 AP on Easy level and 1.51 AP on Mod. level.

The effectiveness of Graph Matching Weighting module (GMW). As the number of projection constraints increases, it is urgent to weigh the con-

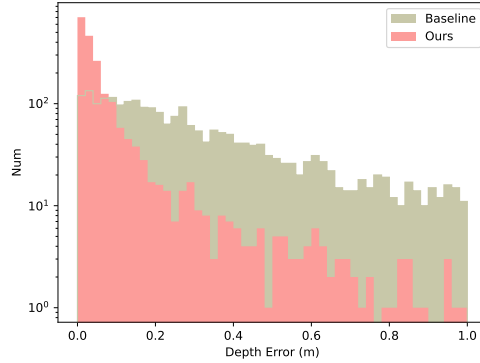


Fig. 4. This figure is the histogram of the depth error on the KITTI *val* set. X-axis represents the distance from the predicted final depth to the GT depth, and Y-axis represents the number of depths of objects with log scale. As the figure shows, our method can estimate depth much more accurately than baseline.

straints appropriately since some are of low quality. To this end, we apply *GMW* and compare it with uncertainty-based methods. The uncertainty-based methods ((a), (b) and (c)) estimate uncertainty independently for each edge. Thus they are not capable of exploiting the global instance information of all edge projection constraints. Different from that, *GMW* is able to exploit the global information and achieve much better results. For example, model (f) (w/ *GMW*) surpasses model (c) (w/o *GMW*) by **1.55** AP on the Hard level, where the model is provided with global information to deal with severe occlusions.

Table 4. Quantitative results using the state-of-the-art method MonoFlex [51] as baseline. This table shows the effectiveness of *DGDE* and *GMW*. The Sec. 4.4 explains the strategy of choosing 1500 depth candidates.

	Weighting Method	<i>DGDE</i>	#Keypoints	#Depth Candidates	$AP_{3D R40 IoU@0.7}$		
					Easy	Mod	Hard
(a)	Uncertainty (Baseline) [51]		10	5	21.63	15.87	13.38
(b)	Uncertainty	✓	10	45	21.72	16.09	13.35
(c)	Uncertainty	✓	73	1500	22.84	16.53	13.77
(d)	<i>GMW</i>		10	5	22.58	16.14	13.63
(e)	<i>GMW</i>	✓	10	45	23.30	16.91	14.93
(f)	<i>GMW</i>	✓	73	1500	23.94	17.38	15.32

Study of supervision priority in *GMW*. In Table 5, we find that enabling depth regression supervision at the beginning detracts the performance. There is a straightforward explanation: when the match is incorrect, supervising the weighted depth will make the gradients noisy.

The Number of Edges. The numerical calculation of the depth by Eq. (6) is very unstable when the denominator is too close to zero. We made a histogram for the analysis. As shown in Fig. 5, the vast majority of these small-denominator depths are of poor quality. For this reason, we use a mask to ignore the depth candidates with extremely small denominators. This ablation study shows how the number of selected depth candidates influences the model performance. As Table 6 shows, the $AP|_{R40}$ increases when the number of selected depths is from 50 to 1500, peaks when the number of depths is 1500 and decreases afterward.

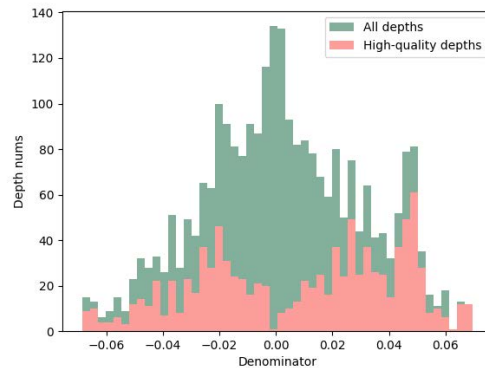


Fig. 5. The horizontal axis represents the value of the denominator of Eq. (6) for an object. The heights of green bins represent the number of depth candidates computed by *DGDE*, while the heights of red bins represent the number of high-quality depth candidates whose distance from ground truth depth is less than 0.5m.

Table 5. Ablation of supervision priority of *GMW* module.

Reg loss	Cls loss	Reg loss start	$AP_{3D R40 IoU@0.7}$		
			Easy	Mod	Hard
	✓	-	23.38	17.03	15.01
✓	✓	0 epoch	22.93	16.83	14.72
✓	✓	50 epochs	23.94	17.38	15.32

Table 6. Ablation of the different number of edges.

#Edges	$AP_{3D R40 IoU@0.7}$		
	Easy	Mod	Hard
50	23.35	16.93	15.01
500	23.58	17.11	15.13
1500	23.94	17.38	15.32
2628	23.37	16.98	14.98

4.5 Discussion about DCD and AutoShape

Although both of our method and AutoShape [23] utilize multiple keypoints to estimate the object’s location, there are three critical differences:

- AutoShape directly uses all 2D-3D *keypoints* projection constraints to solve the object depth. Our method solves a depth candidate from each *edge* constraint. Thus, our edge constraints are not only in a larger number but also in higher order than the keypoint constraints.
- AutoShape generates keypoint weights independently without explicit interaction between keypoints. Our method uses a learnable graph matching module to model the edge constraints, so we produce each depth’s weight based on all the edge constraints, leading to better weighting.

We re-implement AutoShape’s depth estimation and weighting method on our baseline and the experiment result is in Table 7.

Table 7. We re-implement the AutoShape’s depth estimation and weighting regression methods on our baseline. The combined depth estimator combines all keypoint projection constraints as input and produces one depth as output.

Method	#Keypoints	Depth Estimator	#Depth Candidates	$AP_{3D R40 IoU@0.7}$		
				Easy	Mod	Hard
Autoshape [23]	73	Combined	1	22.37	16.48	14.58
<i>DCD</i> (Ours)	73	<i>DGDE</i>	1500	23.94	17.38	15.32

5 Conclusion

This paper proposes a method that can densely calculate an object’s depth from 2D-3D projection constraints of edges of any direction. Therefore, we can obtain $n(n-1)/2$ depths for an object with n keypoints. Moreover, we propose a novel framework that can generate reliable weights for each depth by matching the 2D-3D edges. We finally produce a robust depth by combing each depth candidate with its weight. The experiments show the effectiveness of our method, where we outperform all the existing methods in the KITTI and WOD benchmarks.

Acknowledgements

This work was supported in part by the Major Project for New Generation of AI (No.2018AAA0100400), the National Natural Science Foundation of China (No. 61836014, No. U21B2042, No. 62072457, No. 62006231). Also, our sincere and hearty appreciations go to Lue Fan, who polishes our paper and offers many valuable suggestions.

References

1. Battaglia, P.W., Hamrick, J.B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., et al.: Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261 (2018) [4](#)
2. Brazil, G., Liu, X.: M3d-rpn: Monocular 3d region proposal network for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9287–9296 (2019) [1](#), [3](#), [11](#)
3. Burkard, R.E., Çela, E., Pardalos, P.M., Pitsoulis, L.S.: The Quadratic Assignment Problem, pp. 1713–1809. Springer US, Boston, MA (1998) [4](#)
4. Campbell, D., Liu, L., Gould, S.: Solving the blind perspective-n-point problem end-to-end with robust differentiable geometric optimization. In: European Conference on Computer Vision. pp. 244–261. Springer (2020) [4](#), [7](#)
5. Chabot, F., Chaouch, M., Rabarisoa, J., Teuliere, C., Chateau, T.: Deep manta: A coarse-to-fine many-task network for joint 2d and 3d vehicle analysis from monocular image. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2040–2049 (2017) [3](#)
6. 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) [9](#)
7. 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) [1](#), [3](#), [11](#)
8. 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 Workshops. pp. 1000–1001 (2020) [1](#), [9](#)
9. Fan, L., Pang, Z., Zhang, T., Wang, Y.X., Zhao, H., Wang, F., Wang, N., Zhang, Z.: Embracing single stride 3d object detector with sparse transformer. arXiv preprint arXiv:2112.06375 (2021) [4](#)
10. Fu, K., Liu, S., Luo, X., Wang, M.: Robust point cloud registration framework based on deep graph matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8893–8902 (2021) [4](#)
11. 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) [8](#), [10](#)
12. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. In: International conference on machine learning. pp. 1263–1272. PMLR (2017) [4](#)
13. Grabner, A., Roth, P.M., Lepetit, V.: 3d pose estimation and 3d model retrieval for objects in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3022–3031 (2018) [3](#)
14. He, J., Huang, Z., Wang, N., Zhang, Z.: Learnable graph matching: Incorporating graph partitioning with deep feature learning for multiple object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5299–5309 (June 2021) [4](#)
15. Hu, H., Gu, J., Zhang, Z., Dai, J., Wei, Y.: Relation networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3588–3597 (2018) [4](#)

16. Kundu, A., Li, Y., Rehg, J.M.: 3d-rnn: Instance-level 3d object reconstruction via render-and-compare. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3559–3568 (2018) [3](#), [4](#)
17. 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) [1](#), [3](#), [4](#)
18. 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* **2** (2020) [9](#)
19. Li, Z., Qu, Z., Zhou, Y., Liu, J., Wang, H., Jiang, L.: Diversity matters: Fully exploiting depth clues for reliable monocular 3d object detection. *CoRR abs/2205.09373* (2022) [4](#)
20. Lian, Q., Li, P., Chen, X.: Monojsg: Joint semantic and geometric cost volume for monocular 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1070–1079 (2022) [4](#)
21. 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) [9](#)
22. Liu, Z., Wu, Z., Tóth, R.: Smoke: single-stage monocular 3d object detection via keypoint estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 996–997 (2020) [1](#), [3](#)
23. 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) [1](#), [2](#), [3](#), [4](#), [5](#), [9](#), [10](#), [11](#), [13](#), [14](#)
24. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017) [9](#)
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) [2](#), [9](#), [10](#)
26. Ma, X., Liu, S., Xia, Z., Zhang, H., Zeng, X., Ouyang, W.: Rethinking pseudo-lidar representation. In: European Conference on Computer Vision. pp. 311–327. Springer (2020) [9](#)
27. Ma, X., Liu, S., Xia, Z., Zhang, H., Zeng, X., Ouyang, W.: Rethinking pseudo-lidar representation. In: European Conference on Computer Vision. pp. 311–327. Springer (2020) [11](#)
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: CVPR. pp. 4721–4730 (2021) [9](#)
29. 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) [4](#), [9](#)
30. 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) [1](#)
31. Reading, C., Harakeh, A., Chae, J., Waslander, S.L.: Categorical depth distribution network for monocular 3d object detection. *arXiv preprint arXiv:2103.01100* (2021) [9](#), [10](#), [11](#)

32. Rolínek, M., Swoboda, P., Zietlow, D., Paulus, A., Musil, V., Martius, G.: Deep graph matching via blackbox differentiation of combinatorial solvers. In: European Conference on Computer Vision. pp. 407–424. Springer (2020) [4](#)
33. Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4938–4947 (2020) [4](#)
34. 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. pp. 2743–2752 (2021) [4](#)
35. Shi, X., Ye, Q., Chen, X., Chen, C., Chen, Z., Kim, T.K.: Geometry-based distance decomposition for monocular 3d object detection. arXiv preprint arXiv:2104.03775 (2021) [9](#)
36. Simonelli, A., Bulo, S.R., Porzi, L., Ricci, E., Kotschieder, P.: Towards generalization across depth for monocular 3d object detection. In: ECCV. pp. 767–782. Springer (2020) [9](#)
37. Sun, H., Zhou, W., Fei, M.: A survey on graph matching in computer vision. In: Zheng, Q., Zheng, X., Zhao, X., Yan, W., Zhang, N., Wang, L. (eds.) 13th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics, CISP-BMEI 2020, Chengdu, China, October 17-19, 2020. pp. 225–230. IEEE (2020) [4](#)
38. Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., Vasudevan, V., Han, W., Ngiam, J., Zhao, H., Timofeev, A., Ettinger, S., Krivokon, M., Gao, A., Joshi, A., Zhang, Y., Shlens, J., Chen, Z., Anguelov, D.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) [8](#), [10](#), [11](#)
39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017) [4](#)
40. 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: CVPR. pp. 454–463 (2021) [9](#)
41. 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) [10](#), [11](#)
42. Wang, R., Yan, J., Yang, X.: Learning combinatorial embedding networks for deep graph matching. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3056–3065 (2019) [4](#)
43. 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) [4](#)
44. Weng, X., Kitani, K.: Monocular 3d object detection with pseudo-lidar point cloud. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019) [1](#)
45. Yan, C., Salman, E.: Mono3d: Open source cell library for monolithic 3-d integrated circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers* **65**(3), 1075–1085 (2017) [1](#)
46. Yang, Z., Zhou, Y., Chen, Z., Ngiam, J.: 3d-man: 3d multi-frame attention network for object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1863–1872 (2021) [4](#)

47. Yi, K.M., Trulls, E., Ono, Y., Lepetit, V., Salzmann, M., Fua, P.: Learning to find good correspondences. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2666–2674 (2018) [7](#)
48. 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) [9](#)
49. Zanfir, A., Sminchisescu, C.: Deep learning of graph matching. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2684–2693 (2018) [4](#)
50. Zhang, Y., Ma, X., Yi, S., Hou, J., Wang, Z., Ouyang, W., Xu, D.: Learning geometry-guided depth via projective modeling for monocular 3d object detection. arXiv preprint arXiv:2107.13931 (2021) [2](#)
51. Zhang, Y., Lu, J., Zhou, J.: Objects are different: Flexible monocular 3d object detection. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021. pp. 3289–3298. Computer Vision Foundation / IEEE (2021) [1](#), [2](#), [3](#), [4](#), [6](#), [9](#), [10](#), [11](#), [12](#)
52. Zhou, D., Song, X., Dai, Y., Yin, J., Lu, F., Liao, M., Fang, J., Zhang, L.: Iafa: Instance-aware feature aggregation for 3d object detection from a single image. In: Proceedings of the Asian Conference on Computer Vision (2020) [3](#)
53. Zhou, F., De la Torre, F.: Factorized graph matching. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 127–134. IEEE (2012) [4](#)
54. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019) [3](#)
55. Zhou, Y., Tuzel, O.: Voxynet: 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) [8](#)
56. 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) [9](#)