

# LIGA-Stereo: Learning LiDAR Geometry Aware Representations for Stereo-based 3D Detector

Xiaoyang Guo   Shaoshuai Shi   Xiaogang Wang   Hongsheng Li  
CUHK-SenseTime Joint Laboratory, The Chinese University of Hong Kong  
{xyguo, ssshi, xgwang, hsl1}@ee.cuhk.edu.hk

## Abstract

*Stereo-based 3D detection aims at detecting 3D objects from stereo images, which provides a low-cost solution for 3D perception. However, its performance is still inferior compared with LiDAR-based detection algorithms. To detect and localize accurate 3D bounding boxes, LiDAR-based detectors encode high-level representations from LiDAR point clouds, such as accurate object boundaries and surface normal directions. In contrast, high-level features learned by stereo-based detectors are easily affected by the erroneous depth estimation due to the limitation of stereo matching. To solve the problem, we propose LIGA-Stereo (LiDAR Geometry Aware Stereo Detector) to learn stereo-based 3D detectors under the guidance of high-level geometry-aware representations of LiDAR-based detection models. In addition, we found existing voxel-based stereo detectors failed to learn semantic features effectively from indirect 3D supervisions. We attach an auxiliary 2D detection head to provide direct 2D semantic supervisions. Experiment results show that the above two strategies improved the geometric and semantic representation capabilities. Compared with the state-of-the-art stereo detector, our method has improved the 3D detection performance of cars, pedestrians, cyclists by 10.44%, 5.69%, 5.97% mAP respectively on the official KITTI benchmark. The gap between stereo-based and LiDAR-based 3D detectors is further narrowed. The code is available at <https://xy-guo.github.io/liga/>.*

## 1. Introduction

In recent years, LiDAR-based 3D detection [57, 28, 46, 45, 61, 60] has achieved increasing performance and stability in autonomous driving and robotics. However, the high cost of LiDAR sensors has limited its applications in low-cost products. Stereo matching [3, 18, 26, 35] is the most common depth sensing technique using only cameras. Compared with LiDAR sensors, stereo cameras are at a




Figure 1. Our method utilizes superior geometry-aware features from LiDAR-based 3D detection models to guide the training of stereo-based 3D detectors.

much lower cost and higher resolutions, which makes it a suitable alternative solution for 3D perception. 3D detection from stereo images aims at detecting objects using estimated depth maps [39, 52, 64] or implicit 3D geometry representations [30, 10, 53]. However, the performance of existing stereo-based 3D detection algorithms is still inferior compared with LiDAR-based algorithms.

LiDAR-based detection algorithms take raw point cloud as inputs and then encode the 3D geometry information into intermediate and high-level feature representations. To detect and localize accurate 3D bounding boxes, the model must learn robust local features about object boundaries and surface normal directions, which are essential for predicting accurate bounding box size and orientation. The features learned by LiDAR-based detectors provide robust high-level summarization of accurate 3D geometry structures. In comparison, due to the limitation of stereo matching, the inaccurately estimated depth or implicit 3D representation have difficulties in encoding accurate 3D geometry of objects, especially for distant ones. In addition, the target box supervisions only provide object-level supervisions (location, size, and orientation).

This inspires us to utilize superior LiDAR detection models to guide the training of stereo detection model via imitating the geometry-aware representations encoded by the LiDAR model. Comparing with traditional knowledge distillation [24] for recognition tasks, we did not take the final erroneous classification and regression predictions from

the LiDAR model as “soft” targets, which we found benefits little for training stereo detection networks. The erroneous regression targets would constrain the upper-bound accuracy of bounding box regression. Instead, we force our model to align intermediate features with those of LiDAR models, which encode high-level geometry representations of the scene. The features from LiDAR models could provide powerful and discriminative high-level geometry-aware features, such as surface normal directions and boundary locations. On the other hand, the LiDAR features can provide extra regularization to alleviate the overfitting problem caused by erroneous stereo predictions.

Besides learning better geometry features, we further explore how to learn better semantic features for boosting the 3D detection performance. Instead of learning semantic features from indirect 3D supervisions, we propose to attach an auxiliary multi-scale 2D detection head on the 2D semantic features, which could directly guide the learning of 2D semantic features. Our baseline model, Deep Stereo Geometry Network (DSGN) [10], failed to benefit from extra semantic features effectively according to their ablation studies. We argue that the network provides erroneous semantic supervisions from indirect 3D supervisions because of depth estimation errors, while our proposed direct guidance could greatly benefit 3D detection performance from better learning of 2D semantic features. Experiment results show that the performance is further improved, especially for classes with few samples like *cyclist*.

The contributions can be summarized as follows. 1) We propose to utilize features from superior LiDAR-based detection models to guide the training of stereo-based 3D detection model. LiDAR features encode compact 3D geometry representations of the scene to guide and regularize the stereo features. 2) By attaching an auxiliary 2D detection head to provide direct 2D supervisions, our model significantly improves the learning efficiency for semantic features, which further improves the recall rate especially for rare categories. 4) On the official KITTI 3D detection benchmark, our proposed method surpasses state-of-the-art models by 10.44%, 5.69%, and 5.97% mAP on the *car*, *pedestrian* and *cyclist* classes respectively.

## 2. Related Work

**Stereo Matching.** Mayer *et al.* [35] proposed the first deep stereo algorithm DispNet, which regresses disparity map from feature-based correlation cost volume. DispNet is then extended using multi-stage refinement [37, 31] and auxiliary semantic features [48, 59]. State-of-the-art stereo models construct feature-based cost volume by concatenating left-right 2D features for all disparity candidates and then apply 3D aggregation network to predict the disparity probability distribution [26, 3, 18, 68]. State-of-the-art stereo detection networks [52, 64, 10, 53] also estimate depth by

similar network structures. Zhang *et al.* [66] proposed novel cost volume aggregation strategies to improve the computational efficiency. Yin *et al.* [63] accelerated stereo matching by hierarchically estimating local disparity distributions at each scale and composing them together to form the final match density. Xu *et al.* [55] proposed a sparse points based intra-scale cost aggregation to alleviate the edge-fattening issue. Cheng *et al.* [11] employed neural architecture search (NAS) to automatically search the optimal network structure for stereo matching.

**LiDAR-based 3D Detection.** By leveraging the more accurate depth information captured by LiDAR sensors, 3D detection approaches [71, 40, 58, 46, 28, 47, 45] with LiDAR point clouds generally achieve better performance than image-based approaches. To learn effective features from irregular and sparse point clouds, most existing approaches adopt the voxelization operation to transfer point clouds to regular grids, where the 3D space is first divided into regular 3D voxels [71, 47] or bird-view 2D grids [58, 28] to be processed by convolutions for detection head. Yan *et al.* [57] proposed to utilize sparse convolutions [17] for efficient feature learning from sparse voxels. Du *et al.* [13] presented a feature imitation strategy to learn better perceptual features from synthesized conceptual scenes. Inspired by them, we propose to imitate the much informative feature maps from LiDAR models for better guidance beyond 3D box annotations.

**Stereo-based 3D Detection.** Stereo-based 3D detection algorithms can be roughly divided into three types:

1) *2D-based methods* [7, 54, 30, 39, 65, 49, 38] first detect 2D bounding box proposals and then regress instance-wise 3D boxes. Stereo-RCNN [30] extended Faster R-CNN [43] for stereo-inputs to associate left and right images. Disp R-CNN [49] and ZoomNet [65] incorporated extra instance segmentation mask and part location map to improve detection quality. However, the final performance is limited by the recall of 2D detection algorithms, and 3D geometry information is not fully utilized.

2) *Pseudo-LiDAR* [52, 64, 41, 29, 15] based 3D detection first estimate depth maps and then detect 3D bounding boxes using existing LiDAR-based algorithms. Pseudo-LiDAR++ [64] adapted stereo cost volume to depth cost volume for direct depth estimation. Qian *et al.* [41] made Pseudo-LiDAR pipeline end-to-end trainable. However, these models only take geometry information into consideration, which is lack of complementary semantic features.

3) *Volume-based methods* construct 3D anchor space [42] or detect from 3D stereo volume [10, 53]. DSGN [10] directly constructs differentiable volumetric representation, which encodes implicit 3D geometry structure of the scene, for one-stage stereo-based 3D detection. PLUME [53] directly constructs geometry volume in 3D space for acceleration. Our work takes DSGN [10] as the baseline model.

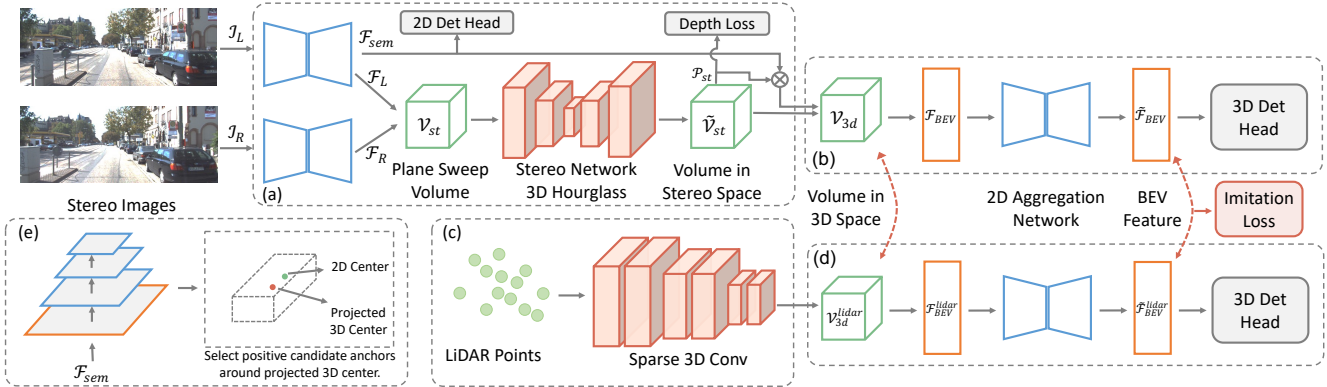


Figure 2. The framework of our stereo-based 3D detection algorithm. (a) The stereo matching network which takes a stereo image pair as inputs and outputs a feature volume in 3D space. (b) and (d) share the same structure, which detects 3D objects from dense / sparse 3D feature volumes, respectively. (c) The teacher LiDAR-based detection network [57], which processes raw LiDAR point clouds using sparse 3D convolution layers. (e) Our auxiliary 2D detection head for direct semantic supervisions.

We solve several key problems existed in [10] and surpass state-of-the-art models by a large margin.

**Knowledge Distillation.** Distillation is first proposed by Hinton *et al.* [24] for model compression by supervising student networks with “softened labels” from predictions of large teacher networks. Going further from “softened labels”, knowledge from intermediate layers provide richer information from the teacher [44, 23, 62, 25, 27]. Recently, knowledge distillation has been successfully applied to detection [4, 51, 5, 56, 14] and semantic segmentation [36, 33, 22, 9]. The knowledge can also be transferred across modalities [19, 1, 2, 69]. However, the feature imitation from LiDAR-based to stereo-based 3D detectors is first explored in this paper.

### 3. Our Approach

In this work, to improve the performance of our model, we developed two strategies to learn better geometric and semantic features respectively. Due to the limitation of stereo matching, stereo-based detectors are fragile to the erroneous depth estimation, especially for low-texture surfaces, blurry boundaries and occlusion areas. In comparison, the features learned by LiDAR-based detectors provide robust high-level geometry-aware representations (accurate boundaries and surface normal directions). To minimize the gap between LiDAR-based and stereo-based detectors, we propose to utilize LiDAR models to guide the training of stereo-based detectors for better geometric supervision. In addition, we employ auxiliary 2D semantic supervisions to improve the learning efficiency for semantic features.

In the following section, we will first revisit the baseline model, Deep Stereo Geometry Network (DSGN) [10], in Sec. 3.1 to make the paper self-contained. In Sec. 3.2, we describe the proposed LiDAR imitation strategy for encod-

ing better geometry representations. In Sec. 3.3, we discuss why the baseline model is of low efficiency for learning semantic features and propose the corresponding solution. The training losses are specified in Sec. 3.4.

#### 3.1. Revisit of Deep Stereo Geometry Network

In this paper, we utilize Deep Stereo Geometry Network (DSGN) [10] as our baseline model, which directly detects objects using implicit volumetric representations.

**Volume in Stereo Space.** Given a left-right image pair  $(\mathcal{I}_L, \mathcal{I}_R)$  and their features  $(\mathcal{F}_L, \mathcal{F}_R)$ , a plane-sweep volume  $\mathcal{V}_{st}$  is constructed by concatenating left features and corresponding right features for every candidate depth level,

$$\mathcal{V}_{st}(u, v, w) = \text{concat} \left[ \mathcal{F}_L(u, v), \mathcal{F}_R\left(u - \frac{fL}{d(w)s_{\mathcal{F}}}, v\right) \right], \quad (1)$$

in which  $(u, v)$  is the current pixel coordinate.  $w=0, 1, \dots$  is the candidate depth index, and  $d(w)=w \cdot v_d + z_{min}$  is the function to calculate its corresponding depth, where  $v_d$  is the depth interval, and  $z_{min}$  is the minimum depth of the detection area.  $f$ ,  $L$ , and  $s_{\mathcal{F}}$  are the camera focal length, the baseline of the stereo camera pair, and the stride of the feature map, respectively. After filtering  $\mathcal{V}_{st}$  with a 3D cost volume aggregation network, we obtain an aggregated stereo volume  $\mathcal{V}_{st}$  and a depth distribution volume  $\mathcal{P}_{st}$ .  $\mathcal{P}_{st}(u, v, \cdot)$  represents the depth probability distribution of pixel  $(u, v)$  over all discrete depth levels  $d(w)$ .

**Volume in 3D Space.** In order to convert the feature volume from stereo space to normal 3D space, the 3D detection area is divided into small voxels of the same size. For each voxel, we project its center  $(x, y, z)$  back into the stereo space using the feature intrinsics  $\mathbf{K}_{\mathcal{F}}$  to obtain its reprojected pixel coordinate  $(u, v)$  and depth index

$d^{-1}(z)=(z-z_{min})/v_d$ . The volume in the 3D space is then defined as the concatenation of resampled stereo volume and semantic features masked by depth probability,

$$\mathcal{V}_{3d}(x, y, z) = \text{concat} \left[ \mathcal{V}_{st}(u, v, d^{-1}(z)), \mathcal{F}_{sem}(u, v) \cdot \mathcal{P}_{st}(u, v, d^{-1}(z)) \right], \quad (2)$$

in which  $\mathcal{V}_{st}$  and  $\mathcal{F}_{sem}$  provide geometric and semantic features respectively. Note that we ignore the trilinear and bilinear resampling operators for simplicity.

**Feature in BEV Space and Detection Head.** The 3D volume  $\mathcal{V}_{3d}$  is then rearranged into 2D bird’s eye view (BEV) feature  $\mathcal{F}_{BEV}$  by merging the channel dimension and the height dimension as [10]. A 2D aggregation network and detection heads are attached to  $\mathcal{F}_{BEV}$  to generate the aggregated BEV feature  $\tilde{\mathcal{F}}_{BEV}$  and predict the final 3D bounding boxes, respectively. The training loss is divided into two parts, depth regression loss and 3D detection loss,

$$\mathcal{L}_{bst} = \mathcal{L}_{depth} + \mathcal{L}_{det}. \quad (3)$$

We have made several modifications to the above baseline to improve both performance and efficiency. Please see details in Sec. 3.4 and Sec. 1 of the supplementary materials.

### 3.2. Learning LiDAR Geometry-aware Representations

LiDAR-based detection models [71, 40, 58, 46, 28, 47, 45] take raw point clouds as inputs, which are then encoded into high-level features (such as accurate boundaries and surface normal directions) for accurate bounding box localization. For stereo-based models, pure depth loss and detection loss could not well make the model learn such features for occlusion areas, non-textured areas and distant objects due to erroneous depth representations.

Inspired by the above observation, we propose to guide the training of stereo-based detectors using the high-level geometry-aware features from LiDAR models. In this paper, we utilize SECOND [57] as our LiDAR “teacher”. To make the features as consistent as possible between two models, we employ 2D aggregation network and detection head of the same structure, as shown in Fig. 2 (see details in the supplementary materials).

Due to the huge differences between the backbone of the LiDAR model and the stereo model, we only focus on how to perform feature imitation for high-level features  $\mathbb{F}_{im} \subseteq \{\mathcal{V}_{3d}, \mathcal{F}_{BEV}, \tilde{\mathcal{F}}_{BEV}\}$ , which are of the same shape between the two models. Our imitation loss aims at minimizing the feature distances between stereo feature  $\mathcal{F}_{im} \in \mathbb{F}_{im}$  and its corresponding LiDAR feature  $\mathcal{F}_{im}^{lidar}$ ,

$$\mathcal{L}_{im} = \sum_{\mathcal{F}_{im} \in \mathbb{F}_{im}} \frac{1}{N_{pos}} \left\| M_{fg} M_{sp} \left( g(\mathcal{F}_{im}) - \frac{\mathcal{F}_{im}^{lidar}}{\mathbb{E}[|\mathcal{F}_{im}^{lidar}|]} \right) \right\|_2^2, \quad (4)$$




Figure 3. Illustration of incorrect semantic supervisions caused by depth estimation errors (drawn in bird’s eye view). Green and red squares represent positive and negative anchors, and orange dash lines represent estimated surface with large errors. Values of  $\mathcal{P}_{st}$  are large around estimated surfaces, thus semantic features of foreground objects are falsely resampled around the wrong surfaces. Therefore, the supervisions for semantic features are inaccurate.

where  $g$  is a single convolution layer with kernel size 1 followed by an optional ReLU layer (depending on whether ReLU is applied to  $\mathcal{F}_{im}^{lidar}$ ).  $M_{fg}$  is the foreground mask, which is 1 inside any ground-truth object box, and 0 otherwise.  $M_{sp}$  is the sparse mask of  $\mathcal{F}_{im}^{lidar}$ , which is 1 for non-empty indices and 0 otherwise.  $N_{pos} = \sum M_{fg} M_{sp}$  is the normalization factor. Note that the LiDAR feature  $\mathcal{F}_{im}^{lidar}$  is normalized by the channel-wise expectation of non-zero absolute values of  $\mathcal{F}_{im}^{lidar}$  to make sure the scale stability of the L2 loss.

Experiment results show that the detection performance benefit little from imitating whole feature maps, thus  $M_{fg}$  is essential to make the imitation loss focus on foreground objects. We found that  $\mathcal{F}_{BEV}$  and  $\mathcal{V}_{3d}$  provide the most effective supervisions for training our stereo detection network. Please check the ablation studies in Sec. 4.3.1.

### 3.3. Improving Semantic Features by Direct 2D Supervisions

From Eq. (2) we can see before resampling semantic feature into 3D space, it is first multiplied by the depth probability from  $\mathcal{P}_{st}$  (which can be seen as the estimation of 3D occupancy mask. Please see how we supervise  $\mathcal{P}_{st}$  in Eq. (6)). In this way, semantic features are only resampled near the estimated surface (orange dashed lines in Fig. 3). However, when there exist large errors in the estimated depth values, the semantic features will be resampled to the wrong positions as illustrated in Fig. 3. As a result, the resampled semantic features deviate from the ground-truth positions and are then assigned with negative anchors (red squares in Fig. 3), and there is no resampled semantic feature around the positive anchors (green squares in Fig. 3). Therefore, the supervision signals for the semantic features are incorrect in this case, which causes the low learning efficiency of semantic features.

To solve the problem, we add an auxiliary 2D detection head to provide direct supervisions for learning semantic

Sensor	Method	AP <sub>3D</sub> (IoU=0.7)			AP <sub>BEV</sub> (IoU=0.7)			AP <sub>3D</sub> (IoU=0.5)			AP <sub>BEV</sub> (IoU=0.5)		
		Easy	Mod	Hard	Easy	Mod	Hard	Easy	Mod	Hard	Easy	Mod	Hard
LiDAR	MV3D (LiDAR) [8]	71.29	62.68	56.56	86.55	78.10	76.67	-	-	-	-	-	-
	PL++: P-RCNN + SL* [64]	75.1	63.8	57.4	88.2	76.9	73.4	-	-	-	-	-	-
	SECOND [57]	87.43	76.48	69.10	89.96	87.07	79.66	-	-	-	-	-	-
	Point-RCNN [46]	88.88	78.63	77.38	-	-	-	-	-	-	-	-	-
	SECOND (our teacher)	88.82	78.57	77.40	89.93	87.75	86.67	98.12	90.17	89.64	98.16	90.20	89.71
Stereo	3DOP [6]	6.55	5.07	4.10	12.63	9.49	7.59	46.04	34.63	30.09	55.04	41.25	34.55
	TLNet [42]	18.15	14.26	13.72	29.22	21.88	18.83	59.51	43.71	37.99	62.46	45.99	41.92
	Stereo-RCNN [30]	54.11	36.69	31.07	68.50	48.30	41.47	85.84	66.28	57.24	87.13	74.11	58.93
	PL: F-PointNet [52]	59.4	39.8	33.5	72.8	51.8	33.5	89.5	75.5	66.3	89.8	77.6	68.2
	PL++: P-RCNN [64]	67.9	50.1	45.3	82.0	64.0	57.3	89.7	78.6	75.1	89.8	83.8	77.5
	OC-Stereo [39]	64.07	48.34	40.39	77.66	65.95	51.20	89.65	80.03	70.34	90.01	80.63	71.06
	Disp-RCNN [49]	64.29	47.73	40.11	77.63	64.38	50.68	90.47	79.76	69.71	90.67	80.45	71.03
	DSGN [10]	72.31	54.27	47.71	83.24	63.91	57.83	-	-	-	-	-	-
	CG-Stereo [29]	76.17	57.82	54.63	87.31	68.69	65.80	90.58	87.01	79.76	97.04	88.58	80.34
	PLUME-Large [53]	-	-	-	84.7	71.1	65.1	-	-	-	91.3	86.6	81.6
	LIGA-Stereo (Ours)	<b>84.92</b>	<b>67.06</b>	<b>63.80</b>	<b>89.35</b>	<b>77.26</b>	<b>69.05</b>	<b>97.06</b>	<b>89.97</b>	<b>87.94</b>	<b>97.22</b>	<b>90.27</b>	<b>88.36</b>

Table 1. Car detection results on the KITTI *validation* set. The results are evaluated using the original KITTI metric with 11 recall values for fair comparison. If not specified, all results in other tables are evaluated using 40 recall positions. \* utilizes 4-beam LiDAR to refine stereo depth estimation.

Sensor	Method	Car AP <sub>3D</sub>			Car AP <sub>BEV</sub>			Car AP <sub>2D</sub>		
		Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
LiDAR	MV3D [8]	74.97	63.63	54.00	86.62	78.93	69.80	-	-	-
	SECOND [57]	83.34	72.55	65.82	89.39	83.77	78.59	-	-	-
	PointPillars [28]	82.58	74.31	68.99	90.07	86.56	82.81	94.00	91.19	88.17
	Point-RCNN [46]	86.96	75.64	70.70	92.13	87.39	82.72	95.92	91.90	87.11
	PV-RCNN [45]	90.25	81.43	76.82	94.98	90.65	86.14	98.17	94.70	92.04
	PL++ (SDN+GDC)* [64]	68.38	54.88	49.16	84.61	73.80	65.59	94.95	85.15	77.78
Stereo	Stereo R-CNN [30]	47.58	30.23	23.72	61.92	41.31	33.42	93.98	85.98	71.25
	PL: AVOD [52]	54.53	34.05	28.25	67.30	45.00	38.40	85.40	67.79	58.50
	PL++: P-RCNN [64]	61.11	42.43	36.99	78.31	58.01	51.25	94.46	82.90	75.45
	ZoomNet [65]	55.98	38.64	30.97	72.94	54.91	44.14	94.22	83.92	69.00
	OC-Stereo [39]	55.15	37.60	30.25	68.89	51.47	42.97	87.39	74.60	62.56
	Disp-RCNN [49]	68.21	45.78	37.73	79.76	58.62	47.73	93.45	82.64	70.45
	DSGN [10]	73.50	52.18	45.14	82.90	65.05	56.60	95.53	86.43	78.75
	CG-Stereo [29]	74.39	53.58	46.50	85.29	66.44	58.95	96.31	90.38	82.80
	CDN [15]	74.52	54.22	46.36	83.32	66.24	57.65	95.85	87.19	79.43
	PLUME-Middle [53]	-	-	-	83.0	66.3	56.7	-	-	-
	LIGA-Stereo (Ours)	<b>81.39</b>	<b>64.66</b>	<b>57.22</b>	<b>88.15</b>	<b>76.78</b>	<b>67.40</b>	<b>96.43</b>	<b>93.82</b>	<b>86.19</b>

Table 2. Car detection results on the KITTI *test* set (official KITTI leaderboard). \* utilizes 4-beam LiDAR to refine stereo depth estimation.

features. Instead of creating feature pyramids (FPN) using multi-level features like [32, 67], we take only a single feature map  $\mathcal{F}_{sem}$  to construct feature pyramids as shown in Fig. 2(e), which forms an information “bottleneck” to enforce all the semantic features to be encoded into  $\mathcal{F}_{sem}$ . Five consecutive convolution layers with a stride of 2 are attached to  $\mathcal{F}_{sem}$  to construct a multi-level feature pyramid, which are then connected to an ATSS [67] head for 2D detection. Since we find that the dilated convolutions and spatial pyramid pooling (SPP) [20] have already produced highly semantic features with large receptive fields, we ignore the top-down path of FPN for simplicity. Please see detailed network structures in Sec. 1 of the supplementary materials.

To ensure the semantic alignment between 2D and 3D features, 2D detection head should predict high scores

around re-projected 3D object centers. We made a small modification to the positive sample assignment algorithm of ATSS [67]. For each ground-truth bounding box  $g$ , we select  $k$  candidate anchors from each scale if their centers are closest to the re-projected 3D object centers, instead of 2D bounding box centers as in [67]. The candidate anchors are then filtered by the dynamic IoU threshold as in ATSS [67] to assign the final positive samples.

### 3.4. Modifications to Baseline and Training Losses

Note that we made several important modifications to DSGN to obtain a faster and more robust baseline model. 1) Decrease the number of channels and layers to reduce memory consumption and computation cost. 2) Utilize SECOND [57] detection head for 3D detection. 3) Replace the smooth-L1 depth regression loss with the uni-modal depth

Method	Pedestrian AP <sub>3D</sub>			Pedestrian AP <sub>BEV</sub>			Cyclist AP <sub>3D</sub>			Cyclist AP <sub>BEV</sub>		
	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
OC-Stereo [39]	24.48	17.58	15.60	29.79	20.80	18.62	29.40	16.63	14.72	32.47	19.23	17.11
DSGN [10]	20.53	15.55	14.15	26.61	20.75	18.86	27.76	18.17	16.21	31.23	21.04	18.93
Disp-RCNN [49]	37.12	25.80	22.04	40.21	28.34	24.46	40.05	24.40	21.12	44.19	27.04	23.58
CG-Stereo [29]	33.22	24.31	20.95	39.24	29.56	25.87	47.40	30.89	27.73	55.33	36.25	32.17
LIGA-Stereo (Ours)	<b>40.46</b>	<b>30.00</b>	<b>27.07</b>	<b>44.71</b>	<b>34.13</b>	<b>30.42</b>	<b>54.44</b>	<b>36.86</b>	<b>32.06</b>	<b>58.95</b>	<b>40.60</b>	<b>35.27</b>

Table 3. Pedestrian and cyclist detection results on KITTI *test* set (official KITTI leaderboard).

loss [63] in Eq. (6) based on Kullback-Leibler divergence. 4) Use a combination of L1 loss and auxiliary rotated 3D IoU loss [70] for better bounding box regression. 5) Attach a small U-Net to the 2D backbone to encode full-resolution feature maps for stereo volume construction. Please check Sec. 1 & 2.1 of the supplementary material for detailed network structure and performance analysis for these modifications.

The new overall loss of our model is formulated as

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{depth} + \mathcal{L}_{det} + \lambda_{im}\mathcal{L}_{im} + \lambda_{2d}\mathcal{L}_{2d}, \\ \mathcal{L}_{det} &= \mathcal{L}_{cls} + \lambda_{reg}^{L1}\mathcal{L}_{reg}^{L1} + \lambda_{reg}^{IoU}\mathcal{L}_{reg}^{IoU} + \lambda_{dir}\mathcal{L}_{dir}, \end{aligned} \quad (5)$$

where  $\mathcal{L}_{cls}$ ,  $\mathcal{L}_{reg}^{L1}$ , and  $\mathcal{L}_{dir}$  are the same classification loss, box regression loss, and direction classification loss as in SECOND [57].  $\mathcal{L}_{reg}^{IoU}$  is the average rotated IoU loss [70] between 3D box predictions and ground-truth bounding boxes. The uni-modal depth loss  $\mathcal{L}_{depth}$  is formulated as

$$\begin{aligned} \mathcal{L}_{depth} &= \frac{1}{N_{gt}} \sum_{u,v} \sum_w \left[ \right. \\ &\left. - \max \left( 1 - \frac{|d^* - d(w)|}{v_d}, 0 \right) \log \mathcal{P}_{st}(u, v, w) \right], \end{aligned} \quad (6)$$

in which  $d^*$  is the ground-truth depth. Note that the loss is only applied to the pixels  $(u, v)$  with valid LiDAR depths, and  $N_{gt}$  denotes the number of valid pixels. Compared with the original L1 loss, the loss in Eq. (6) provides more concentrated supervisions to the target distribution.

## 4. Experiments

### 4.1. Implementation Details

**Dataset.** We evaluate our method on the challenging KITTI 3D object detection benchmark [16]. There are 7,481 training and 7,518 testing stereo image pairs with synchronized LiDAR point clouds in the dataset. Following previous papers [46, 10], the training images are split into training set with 3,712 images and validation set with 3,769 images. There are three object classes in KITTI dataset, *car*, *pedestrian*, and *cyclist*, and each class is divided into three difficulty levels, *easy*, *moderate*, and *hard*.

**Evaluation Metric.** We performed detailed evaluation and analysis for 2D, BEV (bird-view) and 3D detection per-

formance. All the performance results are measured using IoU-based criteria to compute mean averaged precision (mAP) over 40 recall values, except for the results in Table 1, which are evaluated using the old metric with 11 recall values for fair comparison with previous methods.

**Network Structure.** The main structure of our stereo detector follows that of DSGN [10], which consists of 2D feature extraction network (left side of Fig. 2(a)), stereo aggregation network (right side of Fig. 2(a)), and BEV feature aggregation network (Fig. 2(b)). As described in Sec. 3.4, we made several important modifications to [10] to reduce computation cost and memory consumption. Compared with [10], we reduced the numbers of blocks of *conv2* to *conv5* from {3, 6, 12, 4} to {3, 4, 6, 3}, which is the same as ResNet-34 [21]. In addition, we append a small U-Net on the top of the 2D backbone to upsample the SPP feature back into full resolution to provide high-resolution features for stereo matching. The number of channels of the stereo aggregation network is halved from 64 to 32, and the 3D hourglass module for the 3D geometry volume  $\mathcal{V}_{3d}$  in [10] is removed. For the 3D detection head, we follow the open source implementation of SECOND [57] in OpenPCDet [50] to replace the original FCOS head in [10], which we found having inferior performance.

For the LiDAR “teacher”, we employ SECOND [57] because of its simplicity and similarity with our stereo detector. To make the teacher model and the student model as consistent as possible, we modify the stride of the last sparse convolution layer from 2 to 1, to obtain 1/4 down-sampled features to match the feature map size of our stereo detector. Please refer to Sec. 1 of the supplementary materials for detailed network structures of our stereo detector and the LiDAR “teacher”.

**Training Details.** Our stereo detector is trained using AdamW [34] optimizer, with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ . The batch size is fixed to 8. We train the networks with 8 NVIDIA TITAN Xp GPUs, with 1 training sample on each GPU. The model is first trained for 50 epochs using a base learning rate of 0.001, and then for 10 more epochs with a reduced learning rate of 0.0001. The weight decay is set to 0.0001. We apply only horizontal flipping as augmentations. Instead of training individual networks for different classes as [10], we employ only a single model to train on all three classes of the KITTI dataset. To improve training

#	trick	IM	2D	pt	Car AP <sub>3D-IoU0.7</sub>			Car AP <sub>3D-IoU0.5</sub>			Ped AP <sub>3D-IoU0.5</sub>			Ped AP <sub>3D-IoU0.25</sub>			Cyclist AP <sub>3D-IoU0.5</sub>			Cyclist AP <sub>3D-IoU0.25</sub>		
					Easy	Mod	Hard	Easy	Mod	Hard	Easy	Mod	Hard	Easy	Mod	Hard	Easy	Mod	Hard	Easy	Mod	Hard
					91.89	81.08	78.36	99.36	94.35	93.63	72.30	64.79	57.82	86.83	82.03	76.02	85.17	65.81	62.00	89.31	69.41	66.07
					72.31	52.29	47.12	94.73	82.33	77.17	36.84	31.42	27.55	67.38	60.39	53.88	35.39	23.16	22.29	46.28	30.58	29.17
a.	✗				76.44	56.73	49.52	96.25	86.92	79.46	24.39	19.10	16.03	73.12	60.43	52.90	29.12	16.27	15.14	38.81	23.36	21.49
b.	✓				82.15	62.74	57.34	98.85	90.21	84.81	41.18	33.72	29.19	76.19	66.04	57.98	45.63	26.77	24.97	51.50	32.39	30.48
c.	✓	✓			84.99	65.67	60.16	98.84	92.47	87.32	<b>47.82</b>	<b>38.89</b>	<b>33.66</b>	<b>82.78</b>	<b>73.65</b>	<b>64.96</b>	57.54	33.70	31.23	74.39	48.44	44.93
d.	✓		✓		84.45	63.32	57.65	96.23	89.94	84.55	40.72	34.23	29.12	76.13	66.36	58.99	52.65	30.26	28.14	58.59	37.96	35.42
e.	✓	✓	✓		85.26	65.64	60.12	<b>98.98</b>	92.50	87.33	43.97	36.95	31.82	77.57	69.14	62.13	54.59	34.07	31.42	65.19	42.96	40.28
f.	✓			✓	81.05	63.07	56.50	98.91	92.22	87.04	40.29	34.13	29.31	77.60	68.54	61.42	42.77	26.13	24.26	52.11	34.63	32.78
g.	✓	✓	✓	✓	<b>86.84</b>	<b>67.71</b>	<b>62.02</b>	98.87	<b>93.21</b>	<b>87.97</b>	45.54	37.80	32.09	81.87	72.18	64.10	<b>60.00</b>	<b>37.31</b>	<b>34.25</b>	<b>79.89</b>	<b>52.17</b>	<b>48.09</b>

Table 4. Ablation studies on the KITTI *validation* set. The *tricks* include full-resolution features for stereo volume construction, auxiliary 3D IoU loss, and the uni-modal depth loss as described in Sec. 3.4. *IM* denotes imitating geometry-aware representations of LiDAR models in Sec. 3.2. *2D* denotes auxiliary direct 2D supervision in Sec. 3.3. *pt* means ImageNet [12] pre-trained weights, which helps learn better semantic features. (We only load pre-trained weights for layer *conv1-3* due to structure differences). \* The results for DSGN [10] is obtained by evaluating the officially released checkpoints using the new evaluation metrics with 40 recall values for fair comparison.

stability, the weight normalizers for all the loss terms are averaged across all GPUs to avoid unstable gradients. For the hyper-parameters of different losses,  $\lambda_{reg}^{L1}=0.5$ ,  $\lambda_{reg}^{IoU}=1.0$ ,  $\lambda_{dir}=0.2$ ,  $\lambda_{2d}=1.0$ ,  $\lambda_{im}=1.0$ . The input size is fixed to  $320 \times 1248$ , we crop the upper part of input images which does not contain any object to reduce memory consumption. The overall training time is about 12 hours.

For the KITTI dataset, the detection area is set to  $[2m, 59.6m]$  for the  $Z$  (depth) axis,  $[-30m, 30m]$  for the  $X$  axis, and  $[-1m, 3m]$  for the  $Y$  axis (in camera coordinate system). The voxel size of  $\mathcal{V}_{3d}$  in our stereo model is  $(0.2m, 0.2m, 0.2m)$ . The input voxel size of the LiDAR detector is set to  $(0.05m, 0.1m, 0.05m)$ . The spatial resolutions of BEV features of both models are  $0.2m \times 0.2m$ .

## 4.2. Main Results

In this section, we provide detailed comparison with state-of-the-art 3D detectors on the KITTI dataset (see Table 2, 3, 4). Our model is only trained with the KITTI dataset, without any pre-training on Scene Flow datasets.

As shown in Table 1, our method surpasses the state-of-the-art method *CG-Stereo* by 8.75%, 9.24%, 9.17% in 3D mAP (IoU=0.7) for *easy*, *moderate* and *hard* difficulty levels of *car* class. Our model even achieves almost the same performance as our LiDAR teacher (see *SECOND (teacher)* item in Table 1) based on 0.5-IoU evaluation metrics. The gap between our model and the LiDAR teacher is only 0.2% mAP for 3D mAP with 0.5 IoU threshold, which proves that stereo-based 3D detection has the potential to replace LiDAR devices as a low-cost solution.

We also submitted our results to the official evaluation benchmark for evaluating our performance on the *test* set of the KITTI dataset. For BEV performance, we surpasses *PLUME-Middle* [53] by 10.48% mAP. For 3D detection performance, we surpasses *CDN* [15] by 10.44% mAP. Compared with *PL++ (SDN+GDC)*, which is Pseudo-LiDAR++ [64] using 4-beam LiDAR for refinement, we

even achieves much better results. Our method significantly reduces the gap between LiDAR-based methods and stereo-based 3D detection methods. For *pedestrian* and *cyclist* categories, our method exceeds *CG-stereo* by 5.69% and 5.97% mAP for 3D detection respectively. Please see Sec. 3 of the supplementary material for visualization results.

**Computation cost and memory usage.** Compared with DSGN [10], our model has much less layers and number of channels. As shown in Table 5, the memory consumption is greatly decreased from 29GB to 10GB for training and 6GB to 4.9GB for testing. For the running time, although we used NVIDIA TITAN Xp, which is expected to be half the speed of NVIDIA V100, our model still decreased half of the running time compared with DSGN.

## 4.3. Ablation Study

In this section, we conduct ablation experiments to validate the contribution of each component of our model on the KITTI validation set. The results are summarized in Table 4. The first row *SECOND (teacher)* is the performance of our LiDAR “teacher”. The model (a.) is our re-implementation of DSGN [10], but with less channels and using *SECOND* [57] detection head (due to memory limitation and simplicity), which achieves superior performance on cars compared with the original paper, but inferior results for pedestrians and cyclists because we train only a single model for all three classes instead of individual models. We applied a series of important *tricks* to model (a.) to obtain a stronger baseline model (b.), including full-resolution features for stereo volume construction, auxiliary 3D IoU loss for bounding box regression, and the uni-modal depth loss in Eq. (6). Please refer to Sec. 2 of the supplementary materials for details. By now, we have obtained a stronger baseline model with 61.35%, 34.11%, 24.04% 3D mAPs on the three classes of the KITTI validation set. In the following parts, we will prove the effectiveness of the proposed LiDAR feature imitation strategy in Sec. 4.3.1 and the auxil-

iary direct 2D supervisions in Sec. 4.3.2.

### 4.3.1 Ablation study of LiDAR feature imitation

**Influence of LiDAR feature imitation.** We first study whether imitating LiDAR features can benefit 3D detection and which layer is the optimal one. From Table 6 we can see that each of the three feature layers can provide meaningful geometry-aware guidance, among which  $\mathcal{V}_{3d}$  and  $\tilde{\mathcal{F}}_{BEV}$  produce the most effective guidance. We also tried different combinations of imitation layers but found no further gain. However, when combining  $\mathcal{V}_{3d}$  and  $\tilde{\mathcal{F}}_{BEV}$  as the imitation features, we found the model gives more stable results. As a result, we choose  $\mathbb{F}_{im} = \{\mathcal{V}_{3d}, \tilde{\mathcal{F}}_{BEV}\}$  as our final choice.

**Necessity of foreground mask.** We then study the necessity of the foreground mask  $\mathcal{M}_{fg}$ . If  $\mathcal{M}_{fg}$  is removed, the imitation loss would be applied for both foreground and background features with an extreme class imbalance ratio. As expected, after removing  $\mathcal{M}_{fg}$ , there is almost no improvement in our model (w/o  $\mathcal{M}_{fg}$ ) compared with the *baseline* model as shown in Table 6.

**Imitation weight.** As shown in the last section of Table 6, we tested different weights  $\lambda_{im}$  for the LiDAR imitation loss term and found 1.0 is the optimal choice.

Therefore, the optimal choice is to imitate foreground features of the 3D volume  $\mathcal{V}_{3d}$  and the aggregated BEV features  $\tilde{\mathcal{F}}_{BEV}$ . In Table 4, by comparing model (b.) and (c.), the feature imitation strategy improves 3D detection performance by 2.9%, 5.2% and 6.9% on *car*, *pedestrian* and *cyclist* classes of the KITTI dataset. Similar conclusions can be drawn by comparing model (d.) and (e.).

### 4.3.2 The effectiveness of direct 2D supervisions

Due to the inefficiency of the 3D supervision for learning semantic features, we appended an auxiliary 2D detection head to provide direct 2D semantic supervisions for the semantic feature  $\mathcal{F}_{sem}$ . By comparing the model (d.) and the baseline model (b.) in Table 4, the direct 2D supervision improves 3D AP of cars from 62.74% to 63.32%, pedestrians from 33.72% to 34.23%, and cyclists from 26.77% to 30.26%. The category with few data, *cyclist*, benefits more from this strategy.

We also evaluated our semantic features by training 2D detection only, but found poor performance on the KITTI dataset, the 2D APs are only 88%, 52% and 36% for cars, pedestrians and cyclists. We owe the poor results to the lack of data and pre-trained weights. With ImageNet pre-trained weights, the 2D APs are improved to 92%, 54% and 41%. To learn more discriminative 2D semantic features, we also performed experiments with ImageNet pre-trained weights to initialize the model. The two proposed strategies can consistently improve 3D detection performance with pre-trained weights (see model (f.) and (g.) in Table 4). Our

Method	Inference Time (s)	GPU	Memory (GB)	AP <sub>3D</sub> (%) IoU 0.7 / 0.5
DSGN [10]	0.67	V100	29 / 6.0	53.95 / 79.10
Ours	0.35	TITAN Xp	10 / 4.9	68.47 / 92.85

Table 5. Time and memory consumption comparison. The *memory* column shows both training and testing memory usage.

Imitation Settings	Car AP <sub>3D</sub>		
	Easy	Moderate	Hard
baseline	82.15	62.74	57.34
$\mathbb{F}_{im} = \{\mathcal{V}_{3d}\}$	84.50	<b>65.83</b>	<b>60.35</b>
$\mathbb{F}_{im} = \{\mathcal{F}_{BEV}\}$	84.60	65.16	58.04
$\mathbb{F}_{im} = \{\tilde{\mathcal{F}}_{BEV}\}$	84.83	65.10	59.42
$\mathbb{F}_{im} = \{\mathcal{V}_{3d}, \tilde{\mathcal{F}}_{BEV}\}$	<b>84.99</b>	65.67	60.16
$\mathbb{F}_{im} = \{\mathcal{V}_{3d}, \mathcal{F}_{BEV}, \tilde{\mathcal{F}}_{BEV}\}$	82.95	63.49	58.17
w/o $\mathcal{M}_{fg}$	81.13	62.01	54.83
$\lambda_{im}=0.2$	82.61	64.15	58.23
$\lambda_{im}=0.5$	85.14	65.44	59.85
$\lambda_{im}=1.0$	<b>84.99</b>	<b>65.67</b>	<b>60.16</b>
$\lambda_{im}=2.0$	84.47	65.11	59.84

Table 6. Ablation studies for the LiDAR imitation loss. The optimal settings are to imitate  $\{\mathcal{V}_{3d}, \tilde{\mathcal{F}}_{BEV}\}$  with foreground mask  $\mathcal{M}_{fg}$ , and the loss coefficient is set to 1.0.

final model (model (g.)) is able to learn both high-quality geometric features and semantic features, which further reduces the gap between LiDAR-based and stereo-based 3D detection algorithms. Note that due to the structure differences between our backbone and ResNet-34, we only load pre-trained weights for the first three blocks from *conv1* to *conv3*. For the detailed 2D detection results, please refer to Sec. 2 of the supplementary materials.

## 5. Conclusion

In this paper, we propose to learn stereo-based 3D detectors under the guidance of high-level geometry-aware LiDAR features and direct semantic supervisions, which successfully improved the geometric and semantic capabilities. Our model surpasses the state-of-the-art algorithms over 10.44% mAP on the official KITTI 3D detection benchmark, which is closing the gap between stereo-based and LiDAR-based 3D detection algorithms. However, stereo-based 3D detectors still suffer from occlusions, non-textured area and distant objects. By utilizing more advanced “teachers” and more robust stereo algorithms, we expect this problem to be solved step by step in the future.

## Acknowledgement

This work is supported in part by the General Research Fund through the Research Grants Council of Hong Kong under Grants (Nos. 14204021, 14208417, 14207319, 14202217, 14203118, 14208619, ), in part by Research Impact Fund Grant No. R5001-18, in part by CUHK Strategic Fund.



## References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Asr is all you need: Cross-modal distillation for lip reading. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2143–2147. IEEE, 2020.
- [2] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [3] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *CVPR*, pages 5410–5418, 2018.
- [4] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 742–751, 2017.
- [5] Rui Chen, Haizhou Ai, Chong Shang, Long Chen, and Zijie Zhuang. Learning lightweight pedestrian detector with hierarchical knowledge distillation. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 1645–1649. IEEE, 2019.
- [6] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Andrew G Berneshawi, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals for accurate object class detection. In *Advances in Neural Information Processing Systems*, pages 424–432, 2015.
- [7] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals using stereo imagery for accurate object class detection. volume 40, pages 1259–1272. IEEE, 2017.
- [8] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *CVPR*, 2017.
- [9] Yuhua Chen, Wen Li, and Luc Van Gool. Road: Reality oriented adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7892–7901, 2018.
- [10] Yilun Chen, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Dsgn: Deep stereo geometry network for 3d object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [11] Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Tom Drummond, Hongdong Li, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. *arXiv preprint arXiv:2010.13501*, 2020.
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [13] Liang Du, Xiaoqing Ye, Xiao Tan, Jianfeng Feng, Zhenbo Xu, Errui Ding, and Shilei Wen. Associate-3ddet: perceptual-to-conceptual association for 3d point cloud object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13329–13338, 2020.
- [14] Heitor Feliz, Walber M Rodrigues, David Macêdo, Francisco Simões, Adriano LI Oliveira, Veronica Teichrieb, and Cleber Zanchettin. Squeezed deep 6dof object detection using knowledge distillation. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2020.
- [15] Divyansh Garg, Yan Wang, Bharath Hariharan, Mark Campbell, Kilian Q. Weinberger, and Wei-Lun Chao. Wasserstein distances for stereo disparity estimation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [16] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012.
- [17] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. *CVPR*, 2018.
- [18] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *CVPR*, 2019.
- [19] Saurabh Gupta, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2827–2836, 2016.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1904–1916, 2015.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [22] Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, and Youliang Yan. Knowledge adaptation for efficient semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 578–587, 2019.
- [23] Byeongho Heo, Jeessoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1921–1930, 2019.
- [24] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [25] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv preprint arXiv:1707.01219*, 2017.
- [26] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *ICCV*, pages 66–75, 2017.
- [27] Jangho Kim, Seonguk Park, and Nojun Kwak. Paraphrasing complex network: Network compression via factor transfer. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

- [28] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019.
- [29] Chengyao Li, Jason Ku, and Steven L Waslander. Confidence guided stereo 3d object detection with split depth estimation. *IROS*, 2020.
- [30] Peiliang Li, Xiaozhi Chen, and Shaojie Shen. Stereo r-cnn based 3d object detection for autonomous driving. In *CVPR*, pages 7644–7652, 2019.
- [31] Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Wei Chen, Linbo Qiao, Li Zhou, and Jianfeng Zhang. Learning for disparity estimation through feature constancy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2811–2820, 2018.
- [32] Tsung-Yi Lin, Piotr Dollár, Ross B Girshick, Kaiming He, Bharath Hariharan, and Serge J Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [33] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2604–2613, 2019.
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [35] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, pages 4040–4048, 2016.
- [36] Ravi Teja Mullapudi, Steven Chen, Keyi Zhang, Deva Ramanan, and Kayvon Fatahalian. Online model distillation for efficient video inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3573–3582, 2019.
- [37] Jiahao Pang, Wenxiu Sun, Jimmy SJ Ren, Chengxi Yang, and Qiong Yan. Cascade residual learning: A two-stage convolutional neural network for stereo matching. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 887–895, 2017.
- [38] Wanli Peng, Hao Pan, He Liu, and Yi Sun. Ida-3d: Instance-depth-aware 3d object detection from stereo vision for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13015–13024, 2020.
- [39] Alex D. Pon, Jason Ku, Chengyao Li, and Steven L. Waslander. Object-centric stereo matching for 3d object detection. In *2020 IEEE International Conference on Robotics and Automation, ICRA 2020, Paris, France, May 31 - August 31, 2020*, pages 8383–8389. IEEE, 2020.
- [40] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018.
- [41] Rui Qian, Divyansh Garg, Yan Wang, Yurong You, Serge Belongie, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. End-to-end pseudo-lidar for image-based 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5881–5890, 2020.
- [42] Zengyi Qin, Jinglu Wang, and Yan Lu. Triangulation learning network: from monocular to stereo 3d object detection. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [43] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
- [44] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitt-nets: Hints for thin deep nets. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [45] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10529–10538, 2020.
- [46] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Point-rcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–779, 2019.
- [47] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. 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.
- [48] Xiao Song, Xu Zhao, Hanwen Hu, and Liangji Fang. Edgestereo: A context integrated residual pyramid network for stereo matching. In *Asian Conference on Computer Vision*, 2018.
- [49] Jiaming Sun, Linghao Chen, Yiming Xie, Siyu Zhang, Qin-hong Jiang, Xiaowei Zhou, and Hujun Bao. Disp r-cnn: Stereo 3d object detection via shape prior guided instance disparity estimation. In *CVPR*, 2020.
- [50] OpenPCDet Development Team. Openpcdet: An open-source toolbox for 3d object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020.
- [51] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4933–4942, 2019.
- [52] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *CVPR*, pages 8445–8453, 2019.
- [53] Yan Wang, Bin Yang, Rui Hu, Ming Liang, and Raquel Urtasun. Plume: Efficient 3d object detection from stereo images. *arXiv preprint arXiv:2101.06594*, 2021.
- [54] Bin Xu and Zhenzhong Chen. Multi-level fusion based 3d object detection from monocular images. In *Proceedings of*

- the IEEE conference on computer vision and pattern recognition*, pages 2345–2353, 2018.
- [55] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1959–1968, 2020.
- [56] Jiaolong Xu, Yiming Nie, Peng Wang, and Antonio M López. Training a binary weight object detector by knowledge transfer for autonomous driving. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2379–2384. IEEE, 2019.
- [57] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018.
- [58] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018.
- [59] Guorun Yang, Hengshuang Zhao, Jianping Shi, Zhidong Deng, and Jiaya Jia. Segstereo: Exploiting semantic information for disparity estimation. In *ECCV*, pages 636–651, 2018.
- [60] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11040–11048, 2020.
- [61] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *ICCV*, 2019.
- [62] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4133–4141, 2017.
- [63] Zhichao Yin, Trevor Darrell, and Fisher Yu. Hierarchical discrete distribution decomposition for match density estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6044–6053, 2019.
- [64] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. In *ICLR*, 2020.
- [65] X. Ye X. Tan W. Yang S. Wen E. Ding A. Meng L. Huang Z. Xu, W. Zhang. Zoomnet: Part-aware adaptive zooming neural network for 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [66] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *CVPR*, pages 185–194, 2019.
- [67] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9759–9768, 2020.
- [68] Youmin Zhang, Yimin Chen, Xiao Bai, Suihanjin Yu, Kun Yu, Zhiwei Li, and Kuiyuan Yang. Adaptive unimodal cost volume filtering for deep stereo matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12926–12934, 2020.
- [69] Long Zhao, Xi Peng, Yuxiao Chen, Mubbasir Kapadia, and Dimitris N Metaxas. Knowledge as priors: Cross-modal knowledge generalization for datasets without superior knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6528–6537, 2020.
- [70] Dingfu Zhou, Jin Fang, Xibin Song, Chenye Guan, Junbo Yin, Yuchao Dai, and Ruigang Yang. Iou loss for 2d/3d object detection. In *2019 International Conference on 3D Vision (3DV)*, pages 85–94. IEEE, 2019.
- [71] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4490–4499, 2018.