

---

# BEVFusion: A Simple and Robust LiDAR-Camera Fusion Framework

---

Tingting Liang<sup>1\*</sup> Hongwei Xie<sup>2\*</sup> Kaicheng Yu<sup>2\*</sup> Zhongyu Xia<sup>1</sup> Zhiwei Lin<sup>1</sup>  
Yongtao Wang<sup>1§</sup> Tao Tang<sup>2&3</sup> Bing Wang<sup>2</sup> Zhi Tang<sup>1</sup>

<sup>1</sup> Wangxuan Institute of Computer Technology, Peking University, China

<sup>2</sup> DAMO Academy, Alibaba Group, China

<sup>3</sup> Shenzhen Campus of Sun Yat-sen University, China

{tingtingliang, xiazhongyu, zhiweilin, wyt, tangzhi}@pku.edu.cn  
{hongwei.xie.90, kaicheng.yu.yt trent.tangtao, bluecewang6}@gmail.com

## Abstract

Fusing the camera and LiDAR information has become a de-facto standard for 3D object detection tasks. Current methods rely on point clouds from the LiDAR sensor as queries to leverage the feature from the image space. However, people discovered that this underlying assumption makes the current fusion framework infeasible to produce any prediction when there is a LiDAR malfunction, regardless of minor or major. This fundamentally limits the deployment capability to realistic autonomous driving scenarios. In contrast, we propose a surprisingly simple yet novel fusion framework, dubbed BEVFusion, whose camera stream does not depend on the input of LiDAR data, thus addressing the downside of previous methods. We empirically show that our framework surpasses the state-of-the-art methods under the normal training settings. Under the robustness training settings that simulate various LiDAR malfunctions, our framework significantly surpasses the state-of-the-art methods by 15.7% to 28.9% mAP. To the best of our knowledge, we are the first to handle realistic LiDAR malfunction and can be deployed to realistic scenarios without any post-processing procedure. The code is available at <https://github.com/ADLab-AutoDrive/BEVFusion>.

## 1 Introduction

Vision-based perception tasks, like detecting bounding boxes in 3D space, have been a critical aspect of fully autonomous driving tasks [56, 42, 57, 41]. Among all the sensors of a traditional vision on-vehicle perception systems, LiDAR and camera are usually the two most critical sensors that provide accurate point cloud and image features of a surrounding world. In the early stage of perception systems, people design separate deep models for each sensor [37, 38, 59, 16, 53], and fuse the information via post-processing approaches [32]. Note that, people discover that bird's eye view (BEV) has been a de-facto standard for autonomous driving scenarios as, generally speaking, car cannot fly [20, 23, 39, 16, 54, 33]. However, it is often difficult to regress 3D bounding boxes on pure image inputs due to the lack of depth information, and similarly, it is difficult to classify objects on point clouds when LiDAR does not receive enough points.

Recently, people have designed LiDAR-camera fusion deep networks to better leverage information from both modalities. Specifically, the majority of works can be summarized as follow: i) given one or a few points of the LiDAR point cloud, LiDAR to world transformation matrix and the essential matrix (camera to world); ii) people transform the LiDAR points [43, 46, 47, 46, 17, 60] or proposals

---

<sup>§</sup>Corresponding Author.

<sup>\*</sup>Equal Contribution.

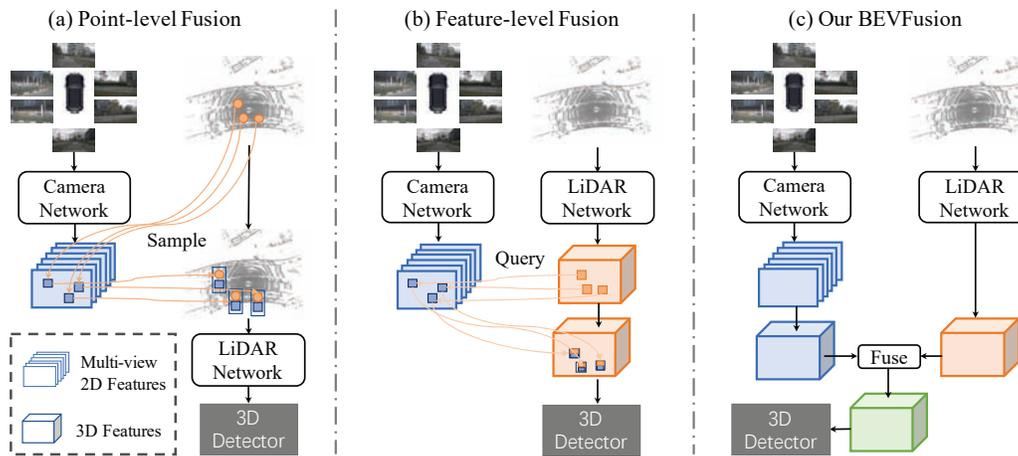


Figure 1: **Comparison of our framework with previous LiDAR-camera fusion methods.** Previous fusion methods can be broadly categorized into (a) point-level fusion mechanism [43, 46, 47, 46, 17, 60] that project image features onto raw point clouds, and (b) feature-level fusion mechanism [4, 61, 1, 21] that projects LiDAR feature or proposals on each view image separately to extract RGB information. (c) In contrast, we propose a novel yet surprisingly simple framework that disentangles the camera network from LiDAR inputs.

[4, 61, 1, 21] into camera world and use them as queries, to select corresponding image features. This line of work constitutes the state-of-the-art methods of 3D BEV perception.

However, one underlying assumption that people overlooked is, that as one needs to generate image queries from LiDAR points, the current LiDAR-camera fusion methods intrinsically depend on the raw point cloud of the LiDAR sensor, as shown in Fig. 1. In the realistic world, people discover that if the LiDAR sensor input is missing, for example, LiDAR points reflection rate is low due to object texture, a system glitch of internal data transfer, or even the field of view of the LiDAR sensor cannot reach 360 degrees due to hardware limitations [62], current fusion methods fail to produce meaningful results<sup>1</sup>. This fundamentally hinders the applicability of this line of work in the realistic autonomous driving system.

We argue the ideal framework for LiDAR-camera fusion should be, that each model for a single modality should not fail regardless of the existence of the other modality, yet having both modalities will further boost the perception accuracy. To this end, we propose a surprisingly simple yet effective framework that disentangles the LiDAR-camera fusion dependency of the current methods, dubbed BEVFusion. Specifically, as in Fig. 1 (c), our framework has two independent streams that encode the raw inputs from the camera and LiDAR sensors into features within the same BEV space. We then design a simple module to fuse these BEV-level features after these two streams, so that the final feature can be passed into modern task prediction head architecture [20, 59, 1].

As our framework is a general approach, we can incorporate current single modality BEV models for camera and LiDAR into our framework. We moderately adopt Lift-Splat-Shoot [34] as our camera stream, which projects multi-view image features to the 3D ego-car coordinate features to generate the camera BEV feature. Similarly, for the LiDAR stream, we select three popular models, two voxel-based ones and a pillar-based one [59, 1, 20] to encode the LiDAR feature into the BEV space.

On the nuScenes dataset, our simple framework shows great generalization ability. Following the same training settings [20, 59, 1], BEVFusion improves PointPillars and CenterPoint by 18.4% and 7.1% in mean average precision (mAP) respectively, and achieves a superior performance of 69.2% mAP comparing to 68.9% mAP of TransFusion [1], which is considered as state-of-the-art. Under the robust setting by randomly dropping the LiDAR points inside object bounding boxes with a probability of 0.5, we propose a novel augmentation technique and show that our framework surpasses

<sup>1</sup>See [62] and Sec. 4.4 for more details.

all baselines significantly by a margin of 15.7% ~28.9% mAP and demonstrate the robustness of our approach.

Our contribution can be summarized as follow: i) we identify an overlooked limitation of current LiDAR-camera fusion methods, which is the dependency of LiDAR inputs; ii) we propose a simple yet novel framework to disentangle LiDAR camera modality into two independent streams that can generalize to multiple modern architectures; iii) we surpass the state-of-the-art fusion methods under both normal and robust settings.

## 2 Related Works

Here, we categorize the 3D detection methods broadly based on their input modality.

**Camera-only.** In the autonomous driving domain, detecting 3D objects with only camera-input has been heavily investigated in recent years thanks to the KITTI benchmark [11]. Since there is only one front camera in KITTI, most of the methods have been developed to address monocular 3D detection [29, 40, 28, 19, 63, 66, 39, 49, 48]. With the development of autonomous driving datasets that have more sensors, like nuScenes[2] and Waymo [44], there exists a trend of developing methods [50, 51, 53] that take multi-view images as input and found to be significantly superior to monocular methods. However, voxel processing is often accompanied by high computation.

As in common autonomous driving datasets, Lift-Splat-Shoot (LSS) [34] uses depth estimation network to extract the implied depth information of multi-perspective images and transform camera feature maps into 3D Ego-car coordinate. Methods [39, 16, 54] are also inspired by LSS [34] and refer to the LiDAR for the supervision on depth prediction. A similar idea can also be found in BEVDet [16, 15], the state-of-the-art method in multi-view 3D object detection. MonoDistill [7] and LiGA Stereo [12] improve performance by unify LiDAR information to a camera branch.

**LiDAR-only.** LiDAR methods initially lie in two categories based on their feature modality: i) point-based methods that directly operate on the raw LiDAR point clouds [38, 37, 36, 42, 57, 22]; and ii) transforming the original point clouds into a Euclidean feature space, such as 3D voxels [65] and feature pillar [20, 52, 59]. Recently, people started to exploit these two feature modalities in a single model to increase the representation power [4, 64, 41]. Another line of work is to exploit the benefit of the bird's eye view plane [20, 10, 45].

**LiDAR-camera fusion.** As the features produced by LiDAR and camera contain complementary information in general, people started to develop methods that can be jointly optimized on both modalities and soon become a de-facto standard in 3D detection. As in Fig. 1, these methods can be divided into two categories depending on their fusion mechanism, (a) point-level fusion where one queries the image features via the raw LiDAR points and then concatenates them back as additional point features [17, 43, 47, 60]; (b) feature-level fusion where one firstly projects the LiDAR points into a feature space [61] or generates proposals [1], queries the associated camera features then concatenates back to the feature space [24, 6]. The latter constitutes the state-of-the-art methods in 3D detection, specifically, TransFusion [1] uses the bounding box prediction of LiDAR features as a proposal to query the image feature, then adapts a Transformer-like architecture to fuse the information back to LiDAR features. Similarly, DeepFusion [21] projects LiDAR features on each view image as queries and then leverages cross-attention for two modalities.

An overlooked assumption of the current fusion mechanism is they heavily rely on the LiDAR point clouds, in fact, if the LiDAR input is missing, these methods will inevitably fail. This will hinder the deployment of such algorithms in realistic settings. In contrast, our BEVFusion is a surprisingly simple yet effective fusion framework that fundamentally overcomes this issue by disentangling the camera branch from the LiDAR point clouds as shown in Fig. 1(c). In addition, concurrent works [27, 58] also address this problem and propose effective LiDAR-camera 3D perception models.

**Other modalities.** There exist other works to leverage other modalities, such as fusing camera-radar by feature map concatenation [3, 18, 31, 30]. While interesting, these methods are beyond the scope of our work. Despite a concurrent work [5] aims to fuse multi-modalities information in a single network, its design is limited to one specific detection head [53] while our framework can be generalized to arbitrary architectures.

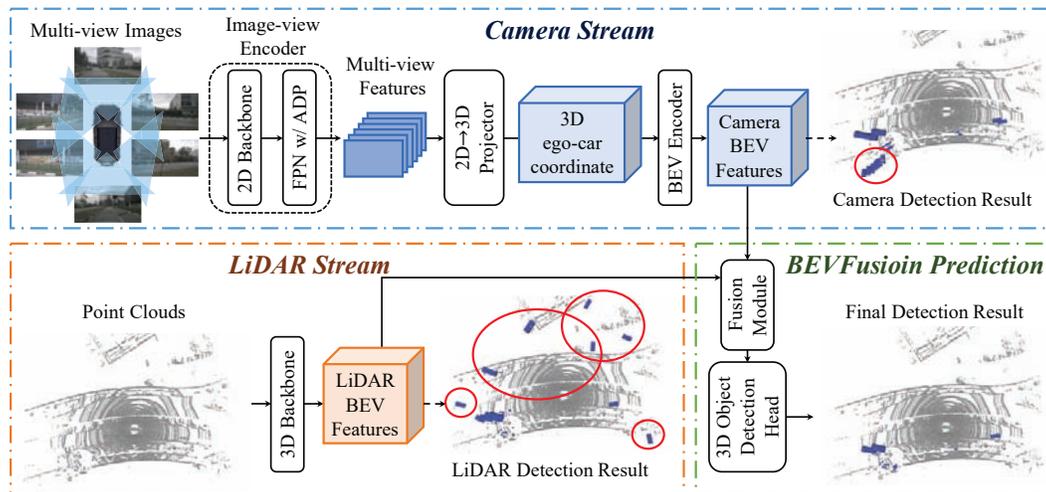


Figure 2: **An overview of BEVFusion framework.** With point clouds and multi-view image inputs, two streams separately extract features and transform them into the same BEV space: i) the camera-view features are projected to the 3D ego-car coordinate features to generate camera BEV feature; ii) 3D backbone extracts LiDAR BEV features from point clouds. Then, a fusion module integrates the BEV features from two modalities. Finally, a task-specific head is built upon the fused BEV feature and predicts the target values of 3D objects. In detection result figures, blue boxes are predicted bounding boxes, while red circled ones are the false positive predictions.

### 3 BEVFusion: A General Framework for LiDAR-Camera Fusion

As shown in Fig. 2, we present our proposed framework, BEVFusion, for the 3D object detection in detail. As our fundamental contribution is disentangling the camera network from the LiDAR features, we first introduce the detailed architecture of the camera and LiDAR stream, then present a dynamic fusion module to incorporate features from these modalities.

#### 3.1 Camera stream architecture: From multi-view images to BEV space

As our framework has the capability to incorporate any camera streams, we begin with a popular approach, Lift-Splat-Shoot (LSS) [35]. As the LSS is originally proposed for BEV semantic segmentation instead of 3D detection, we find out that directly using the LSS architecture has inferior performance, hence we moderately adapt the LSS to improve the performance (see Sec. 4.5 for ablation study). In Fig. 2 (top), we detail the design of our camera stream in the aspect of an image-view encoder that encodes raw images into deep features, a view projector module that transforms these features into 3D ego-car coordinate, and an encoder that finally encodes the features into the bird's eye view (BEV) space.

**Image-view Encoder** aims to encode the input images into semantic information-rich deep features. It consists of a 2D backbone for basic feature extraction and a neck module for scale variate object representation. Different from LSS [34] which uses the convolutional neural network ResNet [13] as the backbone network, we use the more representative one, CB-Swin-Tiny [25] as the backbone. Following [34], we use a standard Feature Pyramid Network (FPN) [26] on top of the backbone to exploit the features from multi-scale resolutions. To better align these features, we first propose a simple feature *Adaptive Module* (ADP) to refine the upsampled features. Specifically, we apply an adaptive average pooling and a  $1 \times 1$  convolution for each upsampled feature before concatenating. See Appendix Sec. A for the detailed module architecture.

**View Projector Module.** As the image features are still in 2D image coordinate, we design a view projector module to transform them into 3D ego-car coordinate. We apply  $2D \rightarrow 3D$  view projection proposed in [34] to construct the Camera BEV feature. The adopted view projector takes the image-view feature as input and densely predicts the depth through a classification manner. Then, according

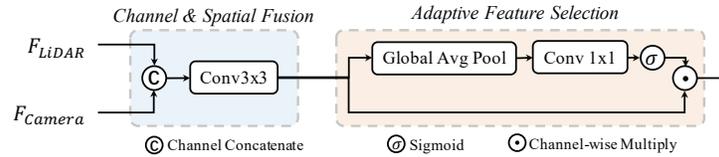


Figure 3: **Dynamic Fusion Module.**

to camera extrinsic parameters and the predicted image depth, we can derive the image-view features to render in the predefined point cloud and obtain a pseudo voxel  $V \in R^{X \times Y \times Z \times C}$ .

**BEV Encoder Module.** To further encode the voxel feature  $V \in R^{X \times Y \times Z \times C}$  into the BEV space feature ( $\mathbf{F}_{\text{Camera}} \in R^{X \times Y \times C_{\text{Camera}}}$ ), we design a simple encoder module. Instead of applying pooling operation or stacking 3D convolutions with stride 2 to compress  $z$  dimension, we adopt the *Spatial to Channel (S2C) operation* [54] to transform  $V$  from 4D tensor to 3D tensor  $V \in R^{X \times Y \times (ZC)}$  via reshaping to preserve semantic information and reduce cost. We then use four  $3 \times 3$  convolution layers to gradually reduce the channel dimension into  $C_{\text{Camera}}$  and extract high-level semantic information. Different from LSS [34] which extracts high-level features based on downsampled low-resolution features, our encoder directly processes full-resolution Camera BEV features to preserve the spatial information.

### 3.2 LiDAR stream architecture: From point clouds to BEV space

Similarly, our framework can incorporate any network that transforms LiDAR points into BEV features,  $\mathbf{F}_{\text{LiDAR}} \in R^{X \times Y \times C_{\text{LiDAR}}}$ , as our LiDAR streams. A common approach is to learn a parameterized voxelization [65] of the raw points to reduce the  $Z$ -dimension and then leverage networks consisting of sparse 3D convolution [56] to efficiently produce the feature in the BEV space. In practice, we adopt three popular methods, PointPillars [20], CenterPoint [59] and TransFusion [1] as our LiDAR stream to showcase the generalization ability of our framework.

### 3.3 Dynamic fusion module

To effectively fuse the BEV features from both camera ( $\mathbf{F}_{\text{Camera}} \in R^{X \times Y \times C_{\text{Camera}}}$ ) and LiDAR ( $\mathbf{F}_{\text{LiDAR}} \in R^{X \times Y \times C_{\text{LiDAR}}}$ ) sensors, we propose a dynamic fusion module in Fig. 3. Given two features under the same space dimension, an intuitive idea is to concatenate them and fuse them with learnable static weights. Inspired by Squeeze-and-Excitation mechanism [14], we apply a simple channel attention module to select important fused features. Our fusion module can be formulated as:

$$\mathbf{F}_{\text{fused}} = f_{\text{adaptive}}(f_{\text{static}}([\mathbf{F}_{\text{Camera}}, \mathbf{F}_{\text{LiDAR}}])), \quad (1)$$

where  $[\cdot, \cdot]$  denotes the concatenation operation along the channel dimension.  $f_{\text{static}}$  is a static channel and spatial fusion function implemented by a  $3 \times 3$  convolution layer to reduce the channel dimension of concatenated feature into  $C_{\text{LiDAR}}$ . With input feature  $\mathbf{F} \in R^{X \times Y \times C_{\text{LiDAR}}}$ ,  $f_{\text{adaptive}}$  is formulated as:

$$f_{\text{adaptive}}(\mathbf{F}) = \sigma(\mathbf{W}f_{\text{avg}}(\mathbf{F})) \cdot \mathbf{F}, \quad (2)$$

where  $\mathbf{W}$  denotes linear transform matrix (e.g.,  $1 \times 1$  convolution),  $f_{\text{avg}}$  denotes the global average pooling and  $\sigma$  denotes sigmoid function.

### 3.4 Detection head

As the final feature of our framework is in BEV space, we can leverage the popular detection head modules from earlier works. This is further evidence of the generalization ability of our framework. In essence, we compare our framework on top of three popular detection head categories, anchor-based [20], anchor-free-based [59], and transform-based [1].

## 4 Experiments

In this section, we present our experimental settings and the performance of BEVFusion to demonstrate the effectiveness, strong generalization ability, and robustness of the proposed framework.

Table 1: **Generalization ability of BEVFusion.** We validate the effectiveness of our fusion framework on nuScenes validation set, compared to single modality streams over three popular methods [20, 59, 1]. Note that each method here defines the structure of the LiDAR stream and associated detection head while the camera stream remains the same as in Sec. 3.1.

Modality		PointPillars		CenterPoint		TransFusion-L	
Camera	LiDAR	mAP	NDS	mAP	NDS	mAP	NDS
✓		22.9	31.1	27.1	32.1	22.7	26.1
	✓	35.1	49.8	57.1	65.4	64.9	69.9
✓	✓	53.5	60.4	64.2	68.0	67.9	71.0

#### 4.1 Experimental settings

**Dataset.** We conduct comprehensive experiments on a large-scale autonomous-driving dataset for 3D detection, nuScenes [2]. Each frame contains six cameras with surrounding views and one point cloud from LiDAR. There are up to 1.4 million annotated 3D bounding boxes for 10 classes. We use nuScenes detection score (NDS) and mean average precision (mAP) as evaluation metrics. See Appendix B.1 for more details.

**Implementation details.** We implement our network in PyTorch using the open-sourced MMDetection3D [8]. We conduct BEVFusion with CB-Swin-Tiny [25] as 2D backbone for image-view encoder. PointPillars [20], CenterPoint [59], and TransFusion-L [1] are chosen as our LiDAR stream and 3D detection head. We set the image size to  $448 \times 800$  and the voxel size following the official settings of the LiDAR stream [20, 59, 1]. Our training consists of two stages: i) We first train the LiDAR stream and camera stream with multi-view image input and LiDAR point clouds input, respectively. Specifically, we train both streams following their LiDAR official settings in MMDetection3D [8]; ii) We then train BEVFusion for another 9 epochs that inherit weights from two trained streams. Note that no data augmentation (i.e., flipping, rotation, or CBGS [67]) is applied when multi-view image input is involved. In this version of paper, we additionally train the fusion phase with BEV-space data augmentation following [16, 27] when comparing with the state-of-the-art methods. In particular, we add the BEV-space augmentation implemented by [27]<sup>2</sup> for a better result. During testing, we follow the settings of LiDAR-only detectors [20, 59, 1] in MMDetection3D [8] without any extra post-processing. See Appendix Sec. B.1 for the detailed hyper-parameters and settings.

#### 4.2 Generalization ability

To demonstrate the generalization ability of our framework, we adapt three popular LiDAR-only detector as our LiDAR stream and detection head, PointPillars [20], CenterPoint [59] and TransFusion-L [1], as described in Sec. 3. If not specified, all experimental settings follow their original papers. In Table 1, we present the results of training two single modality streams, followed by jointly optimized. Empirical results show that our BEVFusion framework can significantly boost the performance of these LiDAR-only methods. Despite the limited performance of the camera stream, our fusion scheme improves PointPillars by 18.4% mAP and 10.6% NDS, and CenterPoint and TransFusion-L by a margin of 3.0% ~ 7.1% mAP. This evidences that our framework can generalize to multiple LiDAR backbone networks.

As our method relies on a two-stage training scheme, we nonetheless report the performance of a single stream in the bottom part of Table 1. We observe that the LiDAR stream constantly surpasses the camera stream by a significant margin. We ascribe this to the LiDAR point clouds providing robust local features about object boundaries and surface normal directions, which are essential for accurate bounding box prediction.

#### 4.3 Comparing with the state-of-the-art methods

Here, we use TransFusion-L as our LiDAR stream and present the results on the test set of nuScenes in Table 2. Without any test time augmentation or model ensemble, our BEVFusion surpasses all previous LiDAR-camera fusion methods and achieves the state-of-the-art performance of 69.2% mAP

<sup>2</sup><https://github.com/mit-han-lab/bevfusion>

Table 2: Results on the nuScenes validation (top) and test (bottom) set.

Method	Modality	mAP	NDS	Car	Truck	C.V.	Bus	Trailer	Barrier	Motor.	Bike	Ped.	T.C.
FUTR3D [5]	LC	64.2	68.0	86.3	61.5	26.0	71.9	42.1	64.4	73.6	63.3	82.6	70.1
BEVFusion	LC	67.9	71.0	88.6	65.0	28.1	75.4	41.4	72.2	76.7	65.8	88.7	76.9
BEVFusion*	LC	69.6	72.1	89.1	66.7	30.9	77.7	42.6	73.5	79.0	67.5	89.4	79.3
PointPillars[20]	L	30.5	45.3	68.4	23.0	4.1	28.2	23.4	38.9	27.4	1.1	59.7	30.8
CBGS[67]	L	52.8	63.3	81.1	48.5	10.5	54.9	42.9	65.7	51.5	22.3	80.1	70.9
CenterPoint[59] <sup>†</sup>	L	60.3	67.3	85.2	53.5	20.0	63.6	56.0	71.1	59.5	30.7	84.6	78.4
TransFusion-L [1]	L	65.5	70.2	86.2	56.7	28.2	66.3	58.8	78.2	68.3	44.2	86.1	82.0
PointPainting[46]	LC	46.4	58.1	77.9	35.8	15.8	36.2	37.3	60.2	41.5	24.1	73.3	62.4
3D-CVF[61]	LC	52.7	62.3	83.0	45.0	15.9	48.8	49.6	65.9	51.2	30.4	74.2	62.9
PointAugmenting[47] <sup>†</sup>	LC	66.8	71.0	87.5	57.3	28.0	65.2	60.7	72.6	74.3	50.9	87.9	83.6
MVP[60]	LC	66.4	70.5	86.8	58.5	26.1	67.4	57.3	74.8	70.0	49.3	89.1	85.0
FusionPainting[55]	LC	68.1	71.6	87.1	60.8	30.0	68.5	61.7	71.8	74.7	53.5	88.3	85.0
TransFusion[1]	LC	68.9	71.7	87.1	60.0	33.1	68.3	60.8	78.1	73.6	52.9	88.4	<b>86.7</b>
BEVFusion (Ours)	LC	<b>69.2</b>	<b>71.8</b>	<b>88.1</b>	<b>60.9</b>	<b>34.4</b>	<b>69.3</b>	<b>62.1</b>	<b>78.2</b>	72.2	52.2	<b>89.2</b>	85.2
BEVFusion (Ours)*	LC	<b>71.3</b>	<b>73.3</b>	<b>88.5</b>	<b>63.1</b>	<b>38.1</b>	<b>72.0</b>	<b>64.7</b>	<b>78.3</b>	<b>75.2</b>	<b>56.5</b>	<b>90.0</b>	<b>86.5</b>

<sup>†</sup> These methods exploit double-flip during the test time. The best and second best results are marked in red and blue.

Notion of class: Construction vehicle (C.V.), pedestrian (Ped.), traffic cone (T.C.). Notion of modality: Camera (C), LiDAR (L).

\* These methods exploit BEV-space data augmentation during training.

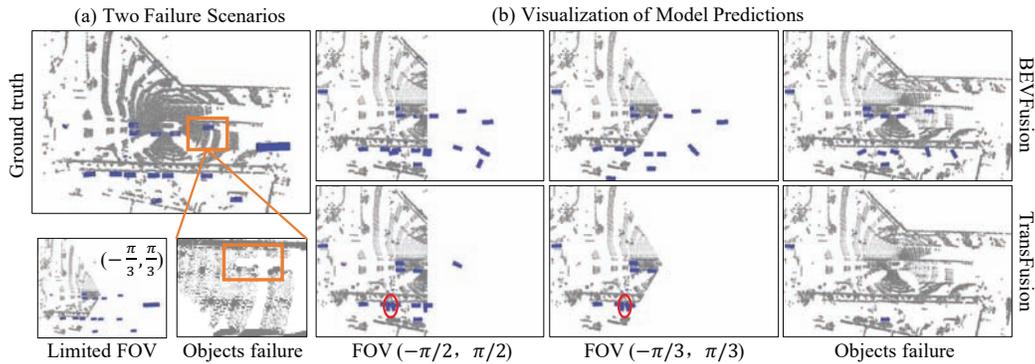


Figure 4: **Visualization of predictions under robustness setting.** (a) We visualize the point clouds under the BEV perspective of two settings, limited field-of-view (FOV) and LiDAR fails to receive object reflection points, where the orange box indicates the object points are dropped. Blue boxes are bounding boxes and red-circled boxes are false-positive predictions. (b) We show the predictions of the state-of-the-art method, TransFusion, and ours under three settings. Obviously, the current fusion approaches fail inevitably when the LiDAR input is missing, while our framework can leverage the camera stream to recover these objects.

comparing to the 68.9% mAP of TransFusion[1]. Note that we do not conduct data augmentation when multi-view image input is involved, while data augmentation plays a critical part in other cutting-edge methods. It is worth noticing that the original TransFusion[1] is a two-stage detector, whose model consists of two independent detection heads. By contrast, our BEVFusion with TransFusion-L as LiDAR backbone only contains one detection head, yet still outperforms the two-stage baseline by 0.3% mAP. As the only difference between our framework and TransFusion is the fusion mechanism, we ascribe this performance gain to a comprehensive exploration of the multi-modality modeling power of BEVFusion.

#### 4.4 Robustness experiments

Here, we demonstrate the robustness of our method against all previous baseline methods on two settings, LiDAR and camera malfunctioning. See [62] for more details.

Table 3: **Results on robustness setting of limited LiDAR field-of-view.** Our method significantly boosts the performance of LiDAR-only methods over all settings. Note that, compared to the TransFusion with camera fusion, our method still achieves over 15.3% mAP and 6.6% NDS improvement, showcasing the robustness of our approach.

FOV	Metrics	PointPillars		CenterPoint		TransFusion		LC
		LiDAR	↑BEVFusion	LiDAR	↑BEVFusion	LiDAR	↑BEVFusion	
$(-\pi/2, \pi/2)$	mAP	12.4	36.8 (+24.4)	23.6	45.5 (+21.9)	27.8	46.4 (+18.6)	31.1
	NDS	37.1	45.8 (+8.7)	48.0	54.9 (+6.9)	50.5	55.8 (+5.3)	49.2
$(-\pi/3, \pi/3)$	mAP	8.4	33.5 (+25.1)	15.9	40.9 (+25.0)	19.0	41.5 (+22.5)	21.0
	NDS	34.3	42.1 (+7.8)	43.5	49.9 (+6.4)	45.3	50.8 (+5.5)	41.2

Table 4: **Results on robustness setting of object failure cases.** Here, we report the results of baseline and our method that trained on the nuScenes dataset with and without the proposed robustness augmentation (Aug.). All settings are the same as in Table 3.

Aug.	Metrics	Pointpillars		Centerpoint		Transfusion		LC
		LiDAR	↑BEVFusion	LiDAR	↑BEVFusion	LiDAR	↑BEVFusion	
	mAP	12.7	34.3 (+21.6)	31.3	40.2 (+8.9)	34.6	40.8 (+6.2)	38.1
	NDS	36.6	49.1 (+12.5)	50.7	54.3 (+3.6)	53.6	56.0 (+2.4)	55.4
✓	mAP	-	41.6 (+28.9)	-	54.0 (+22.7)	-	50.3 (+15.7)	37.2
✓	NDS	-	51.9 (+15.3)	-	61.6 (+10.9)	-	57.6 (+4.0)	51.1

#### 4.4.1 Robustness experiments against LiDAR Malfunctions

To validate the robustness of our framework, we evaluate detectors under two LiDAR malfunctions: i) when the LiDAR sensor is damaged or the LiDAR scan range is restricted, i.e., semi-solid lidars; ii) when objects cannot reflect LiDAR points. We provide visualization of these two failure scenarios in Fig. 4 (a) and evaluate detectors on nuScenes validation set.

**Data augmentation for robustness.** We propose two data augmentation strategies for above two scenarios: i) we simulate the LiDAR sensor failure situation by setting the points with limited Field-of-View (FOV) in range  $(-\pi/3, \pi/3)$ ,  $(-\pi/2, \pi/2)$ . ii) To simulate the object failure, we use a dropping strategy where each frame has a 0.5 chance of dropping objects, and each object has a 0.5 chance of dropping the LiDAR points inside it. Below, we finetune detectors with these two data augmentation strategies, respectively.

**LiDAR sensor failure.** The nuScenes dataset provides a Field-of-View (FOV) range of  $(-\pi, \pi)$  for LiDAR point clouds. To simulate the LiDAR sensor failure situation, we adopt the first aforementioned robust augmentation strategy in Table 3. Obviously, the detector performance degrades as the LiDAR FOV decreases. However, when we fuse camera stream, with the presence of corruptions, the BEVFusion models, in general, are much more robust than their LiDAR-only counterparts, as shown in Fig. 4 (b). Notably, for PointPillars, the mAP increases by 24.4% and 25.1% when LiDAR FOV in  $(-\pi/2, \pi/2)$ ,  $(-\pi/3, \pi/3)$ , respectively. As for TransFusion-L, BEVFusion improves its LiDAR stream by a large margin of over 18.6% mAP and 5.3% NDS. The vanilla LiDAR-camera fusion approach [1] proposed in TransFusion (denoted as LC in Table 3 and Table 4) heavily relies on LiDAR data, and the gain is limited to less than 3.3% mAP while NDS is decreased. The results reveal that fusing our camera stream during training and inference compensates for the lack of LiDAR sensors to a substantial extent.

**LiDAR fails to receive object reflection points.** Here exist common scenarios when LiDAR fails to receive points from the object. For example, on rainy days, the reflection rate of some common objects is below the threshold of LiDAR hence causing the issue of object failure [62]. To simulate such a scenario, we adopt the second aforementioned robust augmentation strategy on the validation set. As shown in Table 4, when we directly evaluate detectors trained without robustness augmentation, BEVFusion shows higher accuracy than the LiDAR-only stream and vanilla LiDAR-camera fusion approach in TransFusion. When we finetune detectors on the robust augmented training set, BEVFusion largely improves PointPillars, CenterPoint, and TransFusion-L by 28.9%,

Table 5: **Results on robustness setting of camera failure cases.** F denotes front camera.

Approach	Clean		Missing F		Preserve F		Stuck	
	mAP	NDS	mAP	NDS	mAP	NDS	mAP	NDS
DETR3D[53]	34.9	43.4	25.8	39.2	3.3	20.5	17.3	32.3
PointAugmenting[47]	46.9	55.6	42.4	53.0	31.6	46.5	42.1	52.8
MVX-Net[43]	61.0	66.1	47.8	59.4	17.5	41.7	48.3	58.8
TransFusion[1]	66.9	70.9	65.3	70.1	64.4	69.3	65.9	70.2
BEVFusion	67.9	71.0	65.9	70.7	65.1	69.9	66.2	70.3

Table 6: **Ablating the camera stream.**

BE	ADP	LB	mAP↑	NDS↑
			13.9	24.5
✓			17.9	27.0
✓	✓		18.0	27.1
✓	✓	✓	22.9	31.1

BE: our simple BEV Encoder. ADP: adaptive module in FPN. LB: a large backbone.

Table 7: **Ablating Dynamic Fusion Module.**

CSF	AFS	Pointpillars		Centerpoint		Transfusion	
		mAP↑	NDS↑	mAP↑	NDS↑	mAP↑	NDS↑
		35.1	49.8	57.1	65.4	64.9	69.9
✓		51.6	57.4	63.0	67.4	67.3	70.5
✓	✓	53.5	60.4	64.2	68.0	67.9	71.0

Dynamic Fusion Module consists of CSF and AFS. CSF: channel& spatial fusion (Fig. 3(left)). AFS: adaptive feature selection (Fig. 3(right)).

22.7%, and 15.7% mAP. Specifically, the vanilla LiDAR-camera fusion method in TransFusion has a gain of only 2.6% mAP, which is smaller than the performance before finetuning, we hypothesize the reason is that the lack of foreground LiDAR points brings wrong supervision during training on the augmented dataset. The results reveal that fusing our camera stream during training and inference largely compensates for the lack of object LiDAR points to a substantial extent. We provide visualization in Fig. 4 (b).

#### 4.4.2 Robustness against camera malfunctions

We further validate the robustness of our framework against camera malfunctions under three scenarios in [62]: i) front camera is missing while others are preserved; ii) all cameras are missing except for the front camera; iii) 50% of the camera frames are stuck. As shown in Table 5, BEVFusion still outperforms camera-only [53] and other LiDAR-camera fusion methods [47, 43, 1] under above scenarios. The results demonstrate the robustness of BEVFusion against camera malfunctions.

#### 4.5 Ablation

Here, we ablate our design choice of the camera stream and the dynamic fusion module.

**Components for camera stream.** We conduct ablation experiments to validate the contribution of each component of our camera stream using different components in Table 6. The naive baseline with ResNet50 and feature pyramid network as multi-view image encoder, ResNet18 as BEV encoder following LSS [34], PointPillars [20] as detection head only obtains 13.9% mAP and 24.5% NDS. As shown in Table 6, there are several observations: (i) when we replace the ResNet18 BEV encoder with our simple BEV encoder module, the mAP and NDS are improved by 4.0% and 2.5%. (2) Adding the adaptive feature alignment module in FPN helps improve the detection results by 0.1%. (3) Concerning a larger 2D backbone, i.e., CB-Swin-Tiny, the gains are 4.9% mAP and 4.0% NDS. The camera stream equipped with PointPillars finally achieves 22.9% mAP and 31.1% NDS, showing the effectiveness of our design for the camera stream.

**Dynamic fusion module.** To illustrate the performance of our fusing strategy for two modalities, we conduct ablation experiments on three different 3D detectors, PointPillars, CenterPoint, and TransFusion and the LiDAR stream as baseline. As shown in 7, with a simple channel& spatial fusion (left part in Fig. 3), BEVFusion greatly improves its LiDAR stream by 16.5% (35.1% → 51.6%) mAP for PointPillars, 5.9% (57.1% → 63.0%) mAP for CenterPoint, and 2.4% (64.9% → 67.3%) mAP for TransFusion. When adaptive feature selection (right part in Fig. 3) is adopted, the mAP can be further improved by 1.9%, 1.2%, and 0.6% mAP for PointPillars, CenterPoint, and TransFusion,

respectively. The results demonstrate the necessity of fusing camera and LiDAR BEV features and effectiveness of our Dynamic Fusion Module in selecting important fused features.

## **5 Conclusion**

In this paper, we introduce BEVFusion, a surprisingly simple yet unique LiDAR-camera fusion framework that disentangles the LiDAR-camera fusion dependency of previous methods. Our framework comprises two separate streams that encode raw camera and LiDAR sensor inputs into features in the same BEV space, followed by a simple module to fuse these features such that they can be passed into modern task prediction head architectures. The extensive experiments demonstrate the strong robustness and generalization ability of our framework against the various camera and LiDAR malfunctions. We hope that our work will inspire further investigation of robust multi-modality fusion for the autonomous driving task.

## **Broader Impacts Statement and Limitations**

This paper studies robust LiDAR-camera fusion for 3D object detection. Since the detection explored in this paper is for generic objects and does not pertain to specific human recognition, so we do not see potential privacy-related issues. However, the biased-to-training-data model may pose safety threats when applied in practice. The research may inspire follow-up studies or extensions, with potential applications in autonomous driving tasks. While our study adopts a simple camera stream as the baseline, we also encourage the community to expand the architecture, e.g, with temporal multi-view camera input and alignment between intermediate LiDAR and camera features. We leave the extension of our method towards building such systems for future work.

## **Acknowledgment**

We thank the anonymous reviewers for their careful reading of our manuscript and their many insightful comments and suggestions.

## References

- [1] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. *arXiv preprint arXiv:2203.11496*, 2022.
- [2] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [3] Simon Chadwick, Will Maddern, and Paul Newman. Distant vehicle detection using radar and vision. In *International Conference on Robotics and Automation (ICRA)*, 2019.
- [4] Xiaozhi Chen, Huimin Ma, Jixiang Wan, B. Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [5] Xuanyao Chen, Tianyuan Zhang, Yue Wang, Yilun Wang, and Hang Zhao. Futr3d: A unified sensor fusion framework for 3d detection. *arXiv preprint arXiv:2203.10642*, 2022.
- [6] Yukang Chen, Yanwei Li, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Focal sparse convolutional networks for 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [7] Zhiyu Chong, Xinzhu Ma, Hong Zhang, Yuxin Yue, Haojie Li, Zhihui Wang, and Wanli Ouyang. Monodistill: Learning spatial features for monocular 3d object detection. *arXiv preprint arXiv:2201.10830*, 2022.
- [8] MMDetection3D Contributors. Mmdetection3d: Open-mmlab next-generation platform for general 3d object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020.
- [9] M. Everingham, L. Gool, Christopher K. I. Williams, J. Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal on Computer Vision (IJCV)*, 2009.
- [10] Lue Fan, Xuan Xiong, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Rangedet: In defense of range view for lidar-based 3d object detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [11] Andreas Geiger, Philip Lenz, and R. Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [12] Xiaoyang Guo, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Liga-stereo: Learning lidar geometry aware representations for stereo-based 3d detector. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [15] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022.
- [16] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021.
- [17] Tengting Huang, Zhe Liu, Xiwu Chen, and X. Bai. EPNet: Enhancing point features with image semantics for 3d object detection. In *European Conference on Computer Vision (ECCV)*, 2020.
- [18] Vijay John and Seiichi Mita. Rvnet: deep sensor fusion of monocular camera and radar for image-based obstacle detection in challenging environments. In *Pacific-Rim Symposium on Image and Video Technology (PSIVT)*, 2019.
- [19] Abhinav Kumar, Garrick Brazil, and Xiaoming Liu. Groomed-nms: Grouped mathematically differentiable nms for monocular 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [20] Alex H. Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. PointPillars: Fast encoders for object detection from point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [21] Yingwei Li, Adams Wei Yu, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Bo Wu, Yifeng Lu, Denny Zhou, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. *arXiv preprint arXiv:2203.08195*, 2022.
- [22] Zhichao Li, Feng Wang, and Naiyan Wang. Lidar r-cnn: An efficient and universal 3d object detector. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [23] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Qiao Yu, and Jifeng Dai. Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. *arXiv preprint arXiv:2203.17270*, 2022.

- [24] Ming Liang, Binh Yang, Shenlong Wang, and R. Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *European Conference on Computer Vision (ECCV)*, 2018.
- [25] Tingting Liang, Xiaojie Chu, Yudong Liu, Yongtao Wang, Zhi Tang, Wei Chu, Jingdong Chen, and Haibing Ling. Cbnet: A composite backbone network architecture for object detection. *arXiv preprint arXiv:2107.00420*, 2021.
- [26] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [27] Zhijian Liu, Haotian Tang, Alexander Amini, Xingyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. *arXiv preprint arXiv:2205.13542*, 2022.
- [28] Zongdai Liu, Dingfu Zhou, Feixiang Lu, Jin Fang, and Liangjun Zhang. Autoshape: Real-time shape-aware monocular 3d object detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [29] Yan Lu, Xinzhu Ma, Lei Yang, Tianzhu Zhang, Yating Liu, Qi Chu, Junjie Yan, and Wanli Ouyang. Geometry uncertainty projection network for monocular 3d object detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [30] Ramin Nabati and Hairong Qi. Centerfusion: Center-based radar and camera fusion for 3d object detection. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [31] Felix Nobis, Maximilian Geisslinger, Markus Weber, Johannes Betz, and Markus Lienkamp. A deep learning-based radar and camera sensor fusion architecture for object detection. In *Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, 2019.
- [32] Su Pang, Daniel Morris, and Hayder Radha. Clocs: Camera-lidar object candidates fusion for 3d object detection. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [33] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [34] Jonah Philion and S. Fidler. Lift, Splat, Shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision (ECCV)*, 2020.
- [35] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision (ECCV)*, 2020.
- [36] C. Qi, W. Liu, Chenxia Wu, Hao Su, and L. Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [37] C. Qi, Hao Su, Kaichun Mo, and L. Guibas. PointNet: Deep learning on point sets for 3d classification and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [38] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Neural Information Processing Systems (NeurIPS)*, 2017.
- [39] Cody Reading, Ali Harakeh, Julia Chae, and Steven L Waslander. Categorical depth distribution network for monocular 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [40] Thomas Roddick, Alex Kendall, and R. Cipolla. Orthographic feature transform for monocular 3d object detection. In *British Machine Vision Conference (BMVC)*, 2019.
- [41] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [42] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN: 3d object proposal generation and detection from point cloud. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [43] Vishwanath A Sindagi, Yin Zhou, and Oncel Tuzel. Mvx-net: Multimodal voxelnet for 3d object detection. In *International Conference on Robotics and Automation (ICRA)*, 2019.
- [44] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, P. Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, S. Ettinger, Maxim Krivokon, A. Gao, Aditya Joshi, Y. Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [45] Pei Sun, Weiyue Wang, Yuning Chai, Gamaleldin F. Elsayed, Alex Bewley, Xiao Zhang, Cristian Sminchisescu, and Drago Anguelov. RSN: Range sparse net for efficient, accurate lidar 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

- [46] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [47] Chunwei Wang, Chao Ma, Ming Zhu, and Xiaokang Yang. PointAugmenting: Cross-modal augmentation for 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [48] Li Wang, Liang Du, Xiaoqing Ye, Yanwei Fu, Guodong Guo, Xiangyang Xue, Jianfeng Feng, and Li Zhang. Depth-conditioned dynamic message propagation for monocular 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [49] Li Wang, Li Zhang, Yi Zhu, Zhi Zhang, Tong He, Mu Li, and Xiangyang Xue. Progressive coordinate transforms for monocular 3d object detection. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- [50] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *Conference on Robot Learning (CoRL)*, 2022.
- [51] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [52] Yue Wang, Alireza Fathi, Abhijit Kundu, David A. Ross, Caroline Pantofaru, Thomas A. Funkhouser, and Justin M. Solomon. Pillar-based object detection for autonomous driving. In *European Conference on Computer Vision (ECCV)*, 2020.
- [53] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning (CoRL)*, 2022.
- [54] Enze Xie, Zhiding Yu, Daquan Zhou, Jonah Philion, Anima Anandkumar, Sanja Fidler, Ping Luo, and Jose M. Alvarez. M<sup>2</sup>bev: Multi-camera joint 3d detection and segmentation with unified birds-eye view representation. *arXiv preprint arXiv:2204.05088*, 2022.
- [55] Shaoqing Xu, Dingfu Zhou, Jin Fang, Junbo Yin, Bin Zhou, and Liangjun Zhang. FusionPainting: Multimodal fusion with adaptive attention for 3d object detection. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2021.
- [56] Yan Yan, Yuxing Mao, and B. Li. SECOND: Sparsely embedded convolutional detection. *Sensors*, 2018.
- [57] Zetong Yang, Y. Sun, Shu Liu, and Jiaya Jia. 3DSSD: Point-based 3d single stage object detector. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [58] Zeyu Yang, Jiaqi Chen, Zhenwei Miao, Wei Li, Xiatian Zhu, and Li Zhang. Deepinteraction: 3d object detection via modality interaction. In *NeurIPS*, 2022.
- [59] Tianwei Yin, Xingyi Zhou, and Philipp Krährenbühl. Center-based 3d object detection and tracking. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [60] Tianwei Yin, Xingyi Zhou, and Philipp Krährenbühl. Multimodal virtual point 3d detection. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- [61] Jin Hyeok Yoo, Yeocheol Kim, Ji Song Kim, and J. Choi. 3D-CVF: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. In *European Conference on Computer Vision (ECCV)*, 2020.
- [62] Kaicheng Yu, Tang Tao, Hongwei Xie, Zhiwei Lin, Zhongwei Wu, Zhongyu Xia, Tingting Liang, Haiyang Sun, Jiong Deng, Dayang Hao, et al. Benchmarking the robustness of lidar-camera fusion for 3d object detection. *arXiv preprint arXiv:2205.14951*, 2022.
- [63] Yunpeng Zhang, Jiwen Lu, and Jie Zhou. Objects are different: Flexible monocular 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [64] Yin Zhou, Pei Sun, Y. Zhang, Dragomir Anguelov, J. Gao, Tom Y. Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. MVF: End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *Conference on Robot Learning (CoRL)*, 2019.
- [65] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-end learning for point cloud based 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [66] Yunsong Zhou, Yuan He, Hongzi Zhu, Cheng Wang, Hongyang Li, and Qinhong Jiang. Monocular 3d object detection: An extrinsic parameter free approach. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [67] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*, 2019.

## Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? **[Yes]**
- Did you include the license to the code and datasets? **[No]** The code and the data are proprietary.
- Did you include the license to the code and datasets? **[N/A]**

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]** See Sec. 1.
  - (b) Did you describe the limitations of your work? **[Yes]** See Sec. 4.2 and Broader Impacts Statement and Limitations.
  - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** See Broader Impacts Statement and Limitations.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]** Personally identifiable information is blurred for privacy protection.
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? **[N/A]**
  - (b) Did you include complete proofs of all theoretical results? **[N/A]**
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]**
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See Sec. 4.1 and Appendix Sec. B.1.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[N/A]** Our experiments are stable with multiple runs. All results are observed via a fixed seed.
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[Yes]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? **[Yes]** See Sec. 4.1.
  - (b) Did you mention the license of the assets? **[No]** The data used in our paper are publicly released.
  - (c) Did you include any new assets either in the supplemental material or as a URL? **[Yes]** See Appendix.
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[Yes]** The data used in our paper are publicly released.
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[Yes]** The data used in our work does not contain personally identifiable information or offensive content.
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**