

# DOLPHINS: Dataset for Collaborative Perception enabled Harmonious and Interconnected Self-driving\*

Ruiqing Mao<sup>1</sup>[0000–0001–7169–3922], Jingyu Guo<sup>1</sup>[0000–0002–1040–8259], Yukuan Jia<sup>1</sup>[0000–0002–2378–9433], Yuxuan Sun<sup>1,2</sup>[0000–0001–9720–6993], Sheng Zhou<sup>1</sup>[0000–0003–0651–0071], and Zhisheng Niu<sup>1</sup>[0000–0003–0420–2024]

<sup>1</sup> Department of Electronic Engineering, Tsinghua University, Beijing 100084, China  
 {mrq20, guojy18, jyk20}@mails.tsinghua.edu.cn  
 {sheng.zhou, niuzhs}@tsinghua.edu.cn

<sup>2</sup> School of Electronic and Information Engineering, Beijing Jiaotong University, Beijing 100044, China  
 yxsun@bjtu.edu.cn

**Abstract.** Vehicle-to-Everything (V2X) network has enabled collaborative perception in autonomous driving, which is a promising solution to the fundamental defect of stand-alone intelligence including blind zones and long-range perception. However, the lack of datasets has severely blocked the development of collaborative perception algorithms. In this work, we release DOLPHINS: Dataset for cOLlaborative Perception enabled Harmonious and INTERconnected Self-driving, as a new simulated *large-scale various-scenario multi-view multi-modality* autonomous driving dataset, which provides a ground-breaking benchmark platform for interconnected autonomous driving. DOLPHINS outperforms current datasets in six dimensions: temporally-aligned images and point clouds from both vehicles and Road Side Units (RSUs) enabling both *Vehicle-to-Vehicle (V2V)* and *Vehicle-to-Infrastructure (V2I)* based collaborative perception; 6 typical scenarios with dynamic weather conditions make the most *various* interconnected autonomous driving dataset; meticulously selected viewpoints providing *full coverage* of the key areas and every object; 42376 frames and 292549 objects, as well as the corresponding 3D annotations, geo-positions, and calibrations, compose the *largest* dataset for collaborative perception; Full-HD images and 64-line LiDARs construct *high-resolution* data with sufficient details; well-organized APIs and open-source codes ensure the *extensibility* of DOLPHINS. We also construct a benchmark of 2D detection, 3D detection, and multi-view collaborative perception tasks on DOLPHINS. The experiment results show that the raw-level fusion scheme through V2X communication can help to improve the precision as well as to reduce the necessity of expensive LiDAR equipment on vehicles when RSUs exist, which may accelerate the

\* This work is sponsored in part by the Nature Science Foundation of China (No. 62022049, No. 61871254), and by the project of Tsinghua University-Toyota Joint Research Center for AI Technology of Automated Vehicle.

popularity of interconnected self-driving vehicles. DOLPHINS dataset and related codes are now available on [www.dolphins-dataset.net](http://www.dolphins-dataset.net).

**Keywords:** Collaborative perception · Interconnected self-driving · Dataset

## 1 Introduction

One major bottleneck of achieving ultra-reliability in autonomous driving is the fundamental defect of stand-alone intelligence due to the single perception viewpoint. As illustrated in Fig. 1(a), the autonomous vehicle could not detect the pedestrians in its blind zone caused by the truck, which may lead to a severe accident. Great efforts have been put into single-vehicle multi-view object detection with multiple heterogeneous sensors [6, 8] or homogeneous sensors [20, 35], but the intrinsic limitation of stand-alone intelligence still exists.

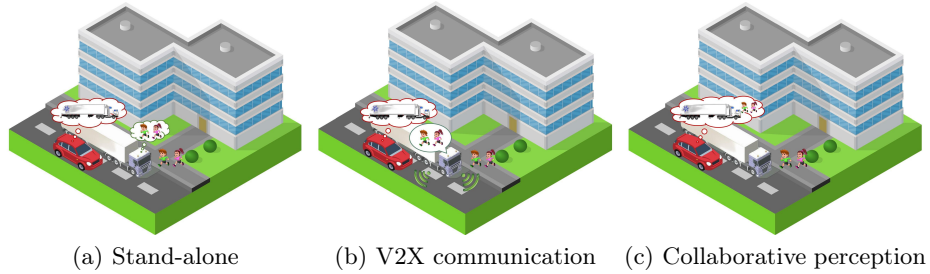


Fig. 1: An illustration of the advantages of collaborative perception over stand-alone intelligence.

Thanks to the Vehicle-to-Everything (V2X) network [42], *interconnected autonomous driving* is a highly-anticipated solution to occlusions, and thus enables advanced autonomous driving capabilities in complex scenarios such as intersections and overtaking. A vehicle can exchange the local sensor data with other terminals (as shown in Fig. 1(b)), including other vehicles and Road Side Units (RSUs), and then perform the object detection by fusing data from multiple viewpoints. The shared sensor data might contain information about the object in the blind zones of the ego vehicle, potentially enhancing the perception reliability [40] as in Fig. 1(c). This procedure is named as *collaborative perception*, which can be categorized into three levels: raw-level (early fusion, e.g. [5]), feature-level (middle fusion, e.g. [4, 34]), and object-level (late fusion, e.g. [17]).

However, the lack of large-scale datasets for collaborative autonomous driving has been seriously restricting the research of collaborative perception algorithms. Traditional datasets focus on a single viewpoint, i.e., the ego vehicle. In the past decade, KITTI [13], nuScenes [1], and Waymo Open [33] have successfully accelerated the development of stand-alone self-driving algorithms with a huge

amount of multi-modality data. But all of the information is collected from the ego vehicle view. Unfortunately, the most challenging but the greatest beneficial issue is the large parallax due to strong perspective changes between different terminals, i.e., aux vehicles and RSUs, as illustrated in Fig. 2. The large parallax leads to various occlusion relationships between objects, which may help the terminals to fulfill the blind zones, but also put forward the matching of the same object from different perspectives. Recently, some pioneer works have concentrated on datasets with multiple viewpoints, such as OPV2V [38], V2X-Sim [21], and DAIR-V2X [41]. Nevertheless, either data from aux vehicles (Vehicle-to-Vehicle, V2V) and RSUs (Vehicle-to-Infrastructure, V2I) are not provided simultaneously, or only an intersection scenario is considered. A more comprehensive dataset is required to fully support the development of V2X-based collaborative autonomous driving algorithms.

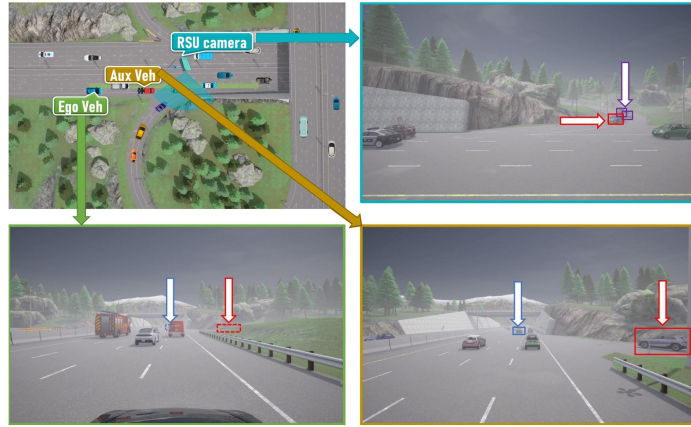


Fig. 2: An example of multi-view object detection in DOLPHINS dataset. There is a right merging lane in front of the ego vehicle. Because of the occlusion, the ego vehicle can hardly detect the purple vehicle (*red box*) on the branch and the police car (*blue box*). The auxiliary vehicle is in front of the ego vehicle, which can see both object vehicles distinctly. Additionally, the RSU can detect another two vehicles (*purple box*) on the branch.

To meet the demands, we present **DOLPHINS**, a new **D**ataset for **c**OLlaborative **P**erception enabled **H**armonious and **I**nterconnected **S**elf-driving. We use the CARLA simulator [9] to complete this work, which can provide us with realistic environment modeling and real-time simulations of the dynamics and sensors of various vehicles. Fig. 3 briefly demonstrates the advantages of DOLPHINS in six dimensions.

**V2X** DOLPHINS contains temporally-aligned images and point clouds from both aux vehicles and RSUs simultaneously, which provides a universal out-

of-the-box benchmark platform for the development and verification of V2V and V2I enabled collaborative perception without extra generation of data.

**Variety** DOLPHINS includes 6 typical autonomous driving scenarios, which is second only to real-world single-vehicle datasets [13, 1]. Our dataset includes urban intersections, T-junctions, steep ramps, highways on-ramps, and mountain roads, as well as dynamic weather conditions. Different scenarios raise different challenges to autonomous driving, such as dense traffic, ramp occlusions, and lane merging. More detailed information on traffic scenarios is presented in Sec. 3.1.

**Viewpoints** Considering the actual driving situation, 3 different viewpoints are meticulously set for each scenario, including both RSUs and vehicles. The data collected from viewpoints can achieve full coverage of key areas in each scenario as illustrated in Fig. 4. More specific locations of each viewpoint are illustrated in Fig. 5.

**Scale** In total, temporally-aligned images and point clouds are recorded over 42376 frames from each viewpoint, which is much larger than any other dataset for collaborative perception. 3D information of 292549 objects is annotated in KITTI format for ease of use, along with the geo-positions and calibrations. Statistical analysis of objects is provided in Sec. 3.4.

**Resolution** DOLPHINS furnishes high-resolution images and point clouds to maintain sufficient details. Full-HD ( $1920 \times 1080$ ) cameras and 64-line LiDARs equipped on both vehicles and RSUs, which are both among the highest quality in all datasets. Detailed descriptions of sensors are stated in Sec. 3.2.

**Extensibility** We also release the related codes of DOLPHINS, which contains the well-organized API to help researchers to generate additional data on demand, which makes DOLPHINS easily extensible and highly flexible.

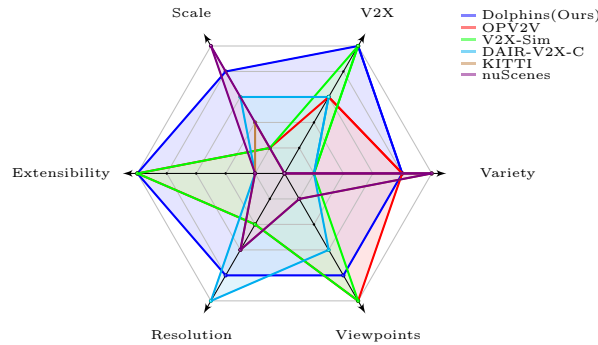


Fig. 3: A comparison with 3 brand new collaborative perception datasets: OPV2V [38], V2X-Sim [21], and DAIR-V2X-C [41], as well as 2 well-known single-vehicle autonomous driving datasets: KITTI [13] and nuScenes [1]. A detailed comparison is provided in Sec. 2.

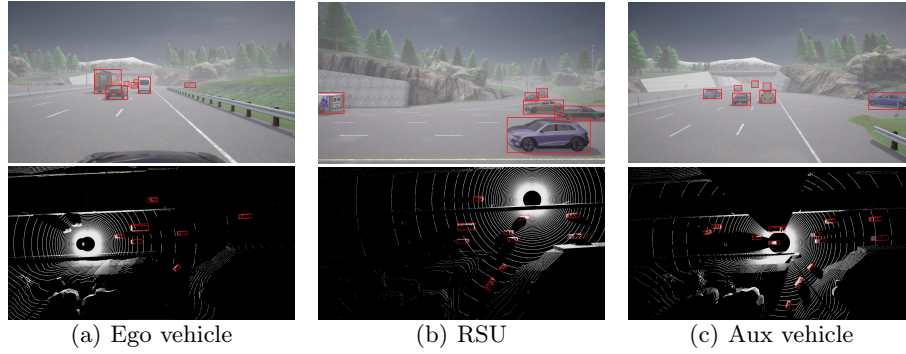


Fig. 4: An illustration of temporary-aligned images and point clouds from three viewpoints. The position of each viewpoint is demonstrated in Fig. 2.

We also conduct a comprehensive benchmark of state-of-the-art algorithms on DOLPHINS. Three typical tasks are considered: 2D object detection, 3D object detection, and multi-view collaborative perception. Other tasks, such as tracking, are also supported in DOLPHINS but not exhibited here. Besides, we construct two raw-level fusion schemes: the point clouds from the ego vehicle and the other viewpoint, and the image from the ego vehicle and point clouds from the RSU. The results of the raw-level fusion algorithms reveal the dual character of interconnected self-driving: enhancing the precision with more information or reducing the cost of sensors on the self-driving vehicles within the same precision.

As a new *large-scale various-scenario multi-view multi-modality* dataset, we hope this work brings a new platform to discover the potential benefits of connected intelligence. Our main contributions are summarized as:

- i. release DOLPHINS dataset with different scenarios, multiple viewpoints, and multi-modal sensors, aiming to inspire the research of collaborative autonomous driving;
- ii. provide open source codes for on-demand generation of data;
- iii. benchmark several state-of-the-art methods in 2D object detection, 3D object detection, and multi-view collaborative perception, illustrating the possibility of solving blind zones caused by occlusions as well as cutting the cost of self-driving vehicles by V2V and V2I communication.

## 2 Related works

There are many relative research areas, such as object detection, collaborative perception, and autonomous driving dataset. Due to the space limitation, some representative works which inspire us are introduced here, and the differences with our proposed dataset are highlighted.

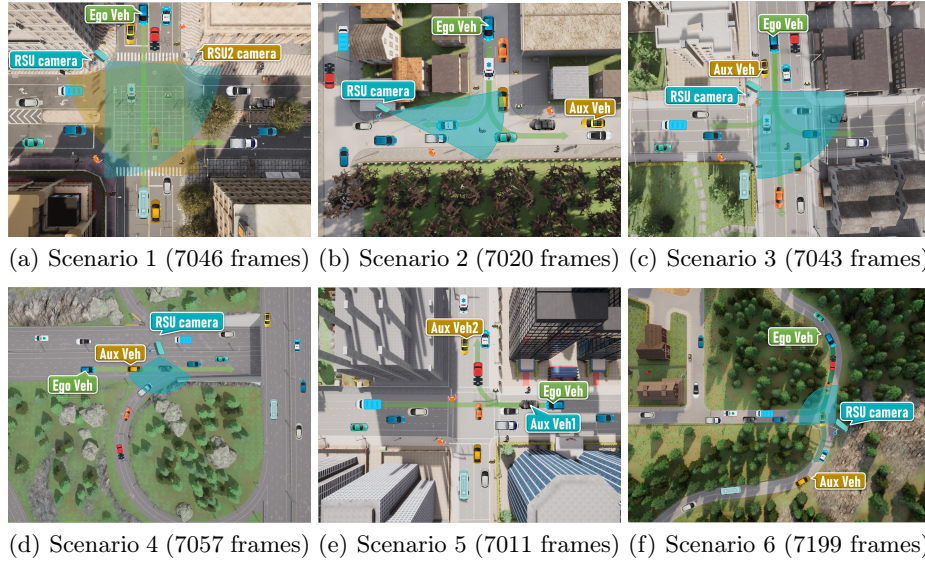


Fig. 5: All ego vehicles are driving along a pre-defined route (*green arrows*), while each RSU camera is settled with a fixed direction and range (*blue or brown sector mark*). We also mark positions where the ego vehicle or possible auxiliary vehicles are initialized. Among all scenarios, (a) and (e) are two intersection scenarios; (b) is the scenario of a T-junction with moderate rain; (c) is also a crossroads while the ego vehicle is on a steep ramp; (d) is a scenario existing a right merging lane on the expressway, and the weather is foggy; (f) is the scenario of a mountain road. All scenarios have plenty of occlusion situations.

**Object detection** is one of the most important tasks in autonomous driving. Typically, there are two kinds of object detectors, distinguished by whether to generate region proposals before the object detection and bounding box regression. R-CNN family [15, 14, 30, 16] is the representative of two-stage detectors, which exhibits epoch-making performance. On the other hand, the single-stage detectors, such as SSD [23] and YOLO [27–29], focus on the inference time and perform significantly faster than the two-stage competitors. Recently, CenterNet [10] and CornerNet [19] propose a new detection method without anchor generation. They directly predict the key points per-pixel, which makes the detection pipeline much simpler. DETR [2] firstly brings transformer architecture into object detection tasks.

**Collaborative perception** is a growing topic in the intelligent transportation society. Due to the 3D information provided by point clouds, LiDAR-based data fusion and object detection have been widely discussed. [5] proposes a raw-level fusion on point clouds with a deep network for detection. [26, 36] aim to use deep neural networks to enhance the perception outputs for sharing. V2VNet [34] con-



siders the feature-level fusion and uses a compressor to reduce the size of original point clouds, which is vital in bandwidth-limited V2X communications. V2V-ViT [37] introduces vision transformer to conquer the noises from communication. [25] and [24] consider the pure image fusion on feature-level and conduct a re-identification task during the detection. The latter uses Graph Neural Network (GNN) for perception and clustering.

**Autonomous driving datasets** are the key to evaluating the performance of detection methods. The commonly used KITTI [13] and nuScenes [1] only contain data from the ego vehicle. Pasadena Multi-view ReID dataset is proposed in [25], which contains data from different viewpoints of a single object. However, the objects are only street trees, which is not enough for autonomous driving. OPV2V [38] uses CARLA simulator [9] to produce multi-view autonomous driving data, but it only considers V2V communication. V2X-Sim [21] is also a CARLA-based simulated dataset. The first version of V2X-Sim only contains point clouds from different vehicles, which can only be applied for V2V communication. The second version of V2X-Sim contains both RGB images and the infrastructure viewpoints. Nevertheless, it still only considers the intersections scenario, and the BEV Lidar on the infrastructure is not realistic. By late February 2022, a new real-world connected autonomous driving dataset DAIR-V2X [41] is released. It consists of images and point clouds from one vehicle and one RSU, and contains both high-ways and intersections. However, DAIR-V2X is not capable of V2V data fusion or any other scenarios with more than two terminals. Our proposed dataset is generated by the CARLA simulator with six different scenarios and reasonable settings of RSUs and aux vehicles. Besides, with the related codes (which will also be released with the dataset), researchers can add any type and any number of sensors at any location as needed. A comparison to the above datasets is provided in Table 1.

Table 1: A detailed comparison between datasets. For DAIR-V2X, we choose DAIR-V2X-C since only this part is captured synchronously by both vehicles and infrastructure sensors.

Dataset	Year	V2X	Scenarios	Viewpoints	Frames	Extensibility	Resolution
KITTI	2012	none	-	1	15 k	×	1382×512 64 lines
nuScenes	2019	none	-	1	1.4 M	×	1600×1200 32 lines
OPV2V	2021	V2V	6	2-7 (avg. 3)	11.5 k	✓	800×600 64 lines
V2X-Sim	2022	V2V+V2I	1	2-5	10 k	✓	1600×900 32 lines
DAIR-V2X-C	2022	V2I	1	2	39 k	×	1920×1080 I: 300 lines; V: 40 lines
DOLPHINS (Ours)	2022	V2V+V2I	6	3	42 k	✓	1920×1080 64 lines

### 3 DOLPHINS Dataset

#### 3.1 Settings of traffic scenarios

We select six typical autonomous driving scenarios and several common types of weather from the preset scenarios of the CARLA simulator (as shown in Fig. 5). In each scenario, we set three units (RSU or vehicles) to collect both images and point cloud information. The first unit is attached to the vehicle we drive, namely, the ego vehicle, which provides us with the main viewpoint. In each simulation round, we initialized it at a specific location. The other two units will also be set up at appropriate positions. They are set on the RSUs or the auxiliary vehicles selected from the scenario and initialized at a specially designated point with a stochastic vehicle model. We initially set 20-30 vehicles as well as 10-15 pedestrians within a specific area around the ego vehicle, which is 100-150 meters in length and 100 meters in width. The initial locations are randomly selected from the preset locations provided by the CARLA simulator, which guarantees that no collisions will happen.

In each scenario, our ego vehicle chooses a specific route. At the same time, we collect the information of all sensors synchronously every 0.5 seconds in the simulation environment, i.e., at the rate of 2 fps. After the vehicle passes through the specific scenario, we wind up the current simulation round, reinitialize the scenario and start a new one. During each round, except for our ego vehicle and the possible auxiliary vehicle, all other traffic participants appear in a reasonable position randomly at the beginning and choose their route by themselves freely.

#### 3.2 Settings of sensors

We equip each unit with a LiDAR and an RGB camera, whose parameters are listed in Table 2. For the convenience of calibration between different sensors, we install both camera and LiDAR on the same point. The position of sensors on the vehicle is illustrated in the supplementary material.

#### 3.3 Extra data and calibrations

For each scenario, We divide our data into the training set and the test set at the ratio of 8:2. Each set contains the original pictures taken by the camera, the point cloud information generated by LiDAR, and the ground truth labels, and the calibration files. The labels include the following information: (i) 2D bounding box of the object in the image, (ii) 3D object dimensions and location, (iii) the value of *alpha* and *rotation\_y* which are defined in the KITTI Vision Benchmark [13]. Except for the above data, we further introduce two extra pieces of information in DOLPHINS: the locations of key vehicles and the context-aware labels. These two kinds of data are essential for collaborative autonomous driving. The geo-positions of vehicles can greatly help to align the perceptual information from different perspectives through coordinate transformations. Actually, to the best of our knowledge, all the published multi-view



Table 2: Parameters of sensors on different units

Sensor type	Parameter attributes	RSU	Vehicle
RGB Camera	Horizontal field of view in degrees	90	90
	Resolution	$1920 \times 1080$	$1920 \times 1080$
	Height in meters	4	$0.3 + h_{veh}^1$
LiDAR	Number of lasers	64	64
	Maximum distance to measure in meters	200	200
	Points generated by all lasers per second	$2.56 \times 10^6$	$2.56 \times 10^6$
	LiDAR rotation frequency	20	20
	Angle in degrees of the highest laser	0	2
	Angle in degrees of the lowest laser	-40	-24.8
	General proportion of points that are randomly dropped	0.1	0.1

1.  $h_{veh}$  denotes the height of the ego vehicle

collaborative perception algorithms are based on the locations of each vehicle, no matter image-based [24, 25] or LiDAR-based [5, 4, 34]. Besides, the interconnected autonomous vehicles can have wider perception fields with the help of other transportation participants and the RSUs, which means they can detect invisible objects. Most of the datasets only provide the labels of those who are in the view angle of sensors, which is not enough for the vehicles to make safe and timely decisions. We provide the labels of all traffic participants within 100 meters in front of or behind the ego unit, as well as 40 meters in the left and right side directions.

### 3.4 Data analysis

To further analyze the data components of the dataset, we calculate the number of cars and pedestrians in each scenario both in the training dataset and the test dataset (as illustrated in Table 3). What’s more, we categorize each object into three detection difficulty levels based on the number of laser points reflected by it in the point clouds. Easy objects reflect more than 16 points, as well as hard objects have no visible point, and the remaining objects are defined as moderate ones. In other words, the difficulty level actually indicates the occlusion level of each object. Since it is unlikely for us to manually annotate the occlusion level, such kind of definition is a suitable and convenient approximation. From the statistical analysis, it turns out that there is no pedestrian in scenarios 4 and 6, i.e., on high-way and mountain roads, which is self-consistent. Scenario 1 contains the most cars and pedestrians, as it is a crowded intersection. Scenario 2 is a T-junction, which has fewer directions for vehicles to travel. Scenario 3 is a steep ramp, which will be the hardest scenario along with Scenario 6, because of the severe occlusions caused by height difference.

Table 3: Statistical analysis of objects in DOLPHINS training and test dataset

Scenario	Training Dataset						Test Dataset					
	Car			Pedestrians			Car			Pedestrians		
	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
1	27548	4423	1090	12370	2117	349	7048	1096	296	3079	579	96
2	15428	1290	567	5641	2281	314	3895	330	155	1481	579	79
3	14365	4029	4291	3003	3462	584	3631	1068	1049	789	889	150
4	34012	11771	4089	0	0	0	8497	2937	1053	0	0	0
5	31648	6201	1993	4797	9734	1476	7918	1578	440	1161	2446	394
6	14035	2203	8531	0	0	0	3531	513	2150	0	0	0
Total	137036	29917	20561	25811	17594	2723	34520	7522	5143	6510	4493	719

## 4 Benchmarks

In this section, we provide benchmarks of three typical tasks on our proposed DOLPHINS dataset: 2D object detection, 3D object detection, and multi-view collaborative perception. For each task, we implement several classical algorithms.

### 4.1 Metrics

We first aggregate the training datasets of six scenarios altogether. The composed dataset will contain various background characteristics and occlusion relationships, which help the model to have better generalization ability. The training dataset is split for training and validation at the ratio of 5:3, then the performance of each detector is examined on the test dataset. Similar to KITTI [13], we use Average Precision (AP) at Intersection-over-Union (IoU) threshold of 0.7 and 0.5 to illustrate the goodness of detectors on cars, as well as IoU of 0.5 and 0.25 for the pedestrians since the pedestrians are much smaller than cars. The degree of difficulty is cumulative in the test, that is, the ground truths of easy objects are also considered in moderate and hard tests.

### 4.2 Experiment details

We use MMDetection [3] and MMDetection3D [7] to construct the training and test pipeline. As for 2D object detection tasks, we finetune the COCO [22] pre-trained models on our dataset. We also provide the GPU memory consumption and the inference speed to illustrate the differences between different methods, where the experiment is set with a batch size equal to 1. All the experiments are performed on 8 RTX 3090 GPUs.

### 4.3 2D object detection

As mentioned in Sec. 2, there are four typical detection paradigms: anchor-based two-stage detectors, anchor-based one-stage detectors, anchor-free one-stage detectors, and the vision transformer. In this part, we select Faster R-CNN [30]

as the representative of anchor-based two-stage detectors, YOLOv3 [29] for the anchor-based one-stage detectors, YOLOX [12] and TOOD [11] for the anchor-free one-stage detectors, and DETR [2] for the vision transformer. Specifically, we set the backbone network of Faster-RCNN and YOLOv3 to be Resnet-50, so that the size of these networks is close to each other. The experiment results are illustrated in Table 4 (left) and Table 5. It shows that all the detectors can have good knowledge of different scenarios. However, the modern anchor-free detectors can significantly speed up the entire inference procedure without loss of precision. One abnormal result is the surprising rise of AP in hard and moderate tasks compared with easy tasks, especially in the pedestrian detection. A reasonable explanation is the large proportion of moderate and hard objects due to the characteristics of different scenarios. For example, in Scenario 6, which is a mountain road, nearly half of the objects are severely occluded due to the undulating planes. Thus, the detectors tend to propose much more candidate objects to match those hard objects, which leads to low AP in easy tasks because of the false positives. It is proved by the high recall scores in easy tasks. The same thing happens in pedestrian detection, where the pedestrians are smaller and thus more likely to be hard ones. However, low AP is not equal to poor performance. On the contrary, meeting the ultra-reliability demands of self-driving, a higher recall rate is much more meaningful than the AP, which can alert the vehicles to the potential dangers in blind zones. More detailed analysis can be found in the supplementary material.

Table 4: 2D and 3D object detection analysis on speed and cost

2D Method	Inference speed (fps)	Memory usage (MB)	3D Method	Inference speed (fps)	Memory usage (MB)
Faster-RCNN	35.6	2513	SECOND	45.7	2433
YOLOv3	50.7	2285	PointPillars	36.8	3483
YOLOX-S	58.1	2001	PV-RCNN	13.1	2899
YOLOX-L	36.5	2233	MVX-Net	11.0	11321
TOOD	26.8	2247			
DETR	26.3	2419			

#### 4.4 3D object detection

As for the 3D object detection tasks, different modals of sensors lead to different detector architectures. We choose SECOND [39], PointPillars [18], and PV-RCNN [31] as SOTA LiDAR-based methods in this part. What’s more, the multi-modal detectors can combine the segmentation information from images and the depth information from LiDARs, which is an advantage to the detection of small objects which reflect few points, e.g. pedestrians. We also test MVX-Net [32] on our multi-modality dataset. The experiment results are illustrated in Table 4 (right) and Table 6. The results show that Scenarios 3 and 6 are the corner cases where the AP is significantly lower than in other scenarios. Due to the steep ramp, the LiDAR on the ego vehicle is hard to detect the opposite vehicles and pedestrians, which is the fundamental defect of stand-alone intelligence.

Table 5: 2D object detection results on DOLPHINS

Scenario	Method	Car AP@IoU=0.7			Car AP@IoU=0.5			Pedestrian AP@IoU=0.5			Pedestrian AP@IoU=0.25		
		Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
1	Faster R-CNN	89.18	80.73	80.62	90.07	89.72	88.76	87.12	90.77	90.82	87.12	90.79	90.88
	YOLOv3	84.76	79.36	79.19	86.93	89.90	89.79	85.22	90.41	90.54	85.40	90.66	90.79
	YOLOX-S	76.21	69.18	68.31	87.57	84.79	82.44	87.19	90.21	89.54	87.46	90.63	90.40
	YOLOX-L	84.72	79.58	78.79	88.95	88.90	87.42	87.76	90.66	90.70	87.82	90.73	90.81
	TOOD	88.30	80.04	79.97	89.76	90.07	89.71	89.02	90.83	90.87	89.03	90.83	90.88
2	DETR	82.29	77.76	75.61	89.40	88.43	86.79	87.16	89.44	89.19	87.77	90.17	90.14
	Faster R-CNN	90.50	90.31	90.19	90.82	90.81	90.80	87.36	90.49	90.18	87.36	90.66	90.60
	YOLOv3	89.33	89.46	89.48	90.17	90.27	90.47	75.96	90.05	90.11	75.96	90.47	90.59
	YOLOX-S	86.46	79.98	79.53	90.20	89.85	89.54	86.99	86.70	80.76	87.24	89.66	87.92
	YOLOX-L	89.48	86.87	80.85	90.61	90.47	90.25	86.29	89.51	87.65	86.29	89.98	89.72
3	TOOD	90.39	90.20	90.10	90.85	90.84	90.82	88.12	90.75	90.53	93.67	90.80	90.76
	DETR	89.90	88.52	87.30	90.60	90.53	90.47	86.80	89.63	87.12	89.83	90.51	90.12
	Faster R-CNN	85.23	88.55	80.25	86.54	90.18	81.51	75.89	89.77	90.67	75.89	89.80	90.76
	YOLOv3	84.64	78.65	71.03	87.46	88.97	80.94	44.31	88.77	89.23	44.63	89.80	90.74
	YOLOX-S	85.95	76.67	66.72	89.42	87.02	77.75	82.54	87.99	81.18	86.54	89.91	85.83
4	YOLOX-L	88.55	80.41	71.25	89.64	89.70	80.39	86.55	90.39	89.65	89.02	90.43	90.37
	TOOD	88.25	80.47	71.50	89.62	90.30	81.26	85.60	90.15	90.69	85.60	90.18	90.83
	DETR	87.94	85.00	79.18	89.17	88.71	86.74	85.73	87.49	85.87	85.73	88.69	87.56
	Faster R-CNN	89.33	81.21	81.09	89.40	89.68	88.70	N/A	N/A	N/A	N/A	N/A	N/A
	YOLOv3	82.43	78.22	77.66	83.95	89.21	89.50	N/A	N/A	N/A	N/A	N/A	N/A
5	YOLOX-S	88.97	76.37	70.29	89.65	86.24	84.43	N/A	N/A	N/A	N/A	N/A	N/A
	YOLOX-L	89.96	80.81	79.58	90.07	89.20	88.25	N/A	N/A	N/A	N/A	N/A	N/A
	TOOD	90.08	81.19	80.22	90.22	90.04	81.63	N/A	N/A	N/A	N/A	N/A	N/A
	DETR	88.17	76.00	74.13	89.44	86.50	86.00	N/A	N/A	N/A	N/A	N/A	N/A
	Faster R-CNN	89.70	80.86	81.04	90.25	89.34	88.82	68.07	80.10	81.06	71.49	80.17	81.41
6	YOLOv3	83.57	78.58	78.80	85.45	89.64	89.77	38.03	87.20	86.48	40.79	89.25	90.52
	YOLOX-S	79.04	75.71	70.72	88.53	86.92	85.07	65.48	79.95	77.82	65.54	85.47	79.99
	YOLOX-L	87.88	80.14	80.18	89.20	89.53	88.75	64.05	87.73	81.18	64.05	88.61	87.24
	TOOD	89.12	85.12	80.73	89.87	89.92	89.96	77.25	88.18	81.41	78.71	89.16	88.43
	DETR	84.26	78.03	75.95	89.58	88.35	87.00	74.59	81.21	77.55	77.04	86.34	83.82
7	Faster R-CNN	77.97	79.60	90.10	78.09	79.74	90.55	N/A	N/A	N/A	N/A	N/A	N/A
	YOLOv3	79.67	79.68	88.85	80.00	80.12	90.35	N/A	N/A	N/A	N/A	N/A	N/A
	YOLOX-S	81.63	79.53	79.73	82.34	81.28	88.61	N/A	N/A	N/A	N/A	N/A	N/A
	YOLOX-L	76.59	76.87	87.35	76.85	77.32	90.07	N/A	N/A	N/A	N/A	N/A	N/A
	TOOD	82.54	81.62	90.63	82.65	81.77	89.93	N/A	N/A	N/A	N/A	N/A	N/A
8	DETR	85.09	83.27	81.81	85.47	84.03	89.89	N/A	N/A	N/A	N/A	N/A	N/A

What's more, PV-RCNN [31] gains significantly better performance at the cost of taking nearly four times as long as SECOND [39]. MVX-Net [32] is inferior to those pure LiDAR-based methods, but it achieves surprising performance in pedestrians, which means the rich segmentation of information from images is profitable for the detection of small objects.

#### 4.5 Multi-view collaborative perception

Based on the information to exchange, collaborative perception can be categorized into three levels: raw-level (early fusion), feature-level (middle fusion), and object-level (late fusion). Due to the 3D information provided by point clouds, LiDAR-based data fusion and object detection have been widely discussed. We realize a raw-level fusion algorithm based on DOLPHINS LiDAR data through the superposition of point clouds from different perspectives. However, not all

Table 6: 3D object detection results on DOLPHINS

Scenario	Method	Car			Car			Pedestrian			Pedestrian		
		AP@IoU=0.7			AP@IoU=0.5			AP@IoU=0.5			AP@IoU=0.25		
		Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
1	SECOND	95.65	90.37	87.36	98.79	96.05	92.97	74.17	70.08	68.19	96.78	95.86	93.47
	PointPillar	96.63	92.09	88.83	98.55	96.19	93.13	70.12	65.41	63.56	95.57	94.38	91.98
	PV-RCNN	98.14	93.87	90.69	98.90	96.18	92.98	83.67	80.36	78.37	97.31	96.50	94.09
	MVX-Net	89.25	84.3	83.99	89.62	89.51	87.02	91.36	88.93	86.49	99.61	99.55	97.06
2	SECOND	96.33	91.32	88.77	98.48	97.50	95.56	66.71	59.35	57.27	95.07	91.00	87.94
	PointPillar	97.39	93.06	91.70	98.23	97.33	95.48	59.65	53.33	51.35	92.52	87.58	84.49
	PV-RCNN	98.57	94.76	91.15	99.12	97.44	94.33	78.74	71.59	68.94	97.37	93.91	90.60
	MVX-Net	93.77	91.04	88.63	96.33	95.91	93.72	89.12	82.22	79.78	99.37	98.72	94.19
3	SECOND	80.30	67.16	54.96	85.51	75.08	64.29	49.87	29.31	27.05	92.73	60.26	55.88
	PointPillar	78.94	68.29	56.78	85.67	75.75	66.75	37.49	22.04	20.00	82.64	52.55	48.25
	PV-RCNN	85.93	73.85	60.59	87.54	77.18	64.30	63.95	37.46	34.53	90.93	58.96	54.32
	MVX-Net	68.96	58.47	48.73	71.80	61.67	56.25	71.96	41.36	37.10	93.99	56.62	53.85
4	SECOND	97.81	92.11	84.47	99.33	97.35	90.50	N/A	N/A	N/A	N/A	N/A	N/A
	PointPillar	98.07	94.00	86.52	98.79	97.57	91.01	N/A	N/A	N/A	N/A	N/A	N/A
	PV-RCNN	99.37	95.54	87.70	99.50	97.78	89.98	N/A	N/A	N/A	N/A	N/A	N/A
	MVX-Net	91.76	86.39	83.73	91.96	89.19	86.55	N/A	N/A	N/A	N/A	N/A	N/A
5	SECOND	96.49	91.41	87.44	98.68	96.36	92.71	75.33	65.03	58.90	97.37	94.29	87.52
	PointPillar	97.45	92.92	89.30	98.81	96.52	92.94	71.36	62.07	56.27	97.32	93.59	85.70
	PV-RCNN	98.57	94.39	90.49	99.23	97.21	93.15	90.47	79.07	71.64	98.94	95.45	87.29
	MVX-Net	91.69	86.67	84.17	94.45	91.94	89.39	84.87	75.83	68.91	99.52	99.36	91.81
6	SECOND	90.53	82.60	56.05	97.54	82.11	68.15	N/A	N/A	N/A	N/A	N/A	N/A
	PointPillar	89.31	82.30	57.32	97.44	92.66	70.62	N/A	N/A	N/A	N/A	N/A	N/A
	PV-RCNN	95.95	89.29	62.05	98.23	93.95	69.28	N/A	N/A	N/A	N/A	N/A	N/A
	MVX-Net	87.53	75.29	52.60	90.76	80.55	57.97	N/A	N/A	N/A	N/A	N/A	N/A

the LiDAR-based 3D detection algorithms can be adapted to raw-level fusion schemes. Since many detectors use voxels to represent the point clouds of a district, the height of voxels is limited to reduce the computation complexity. The limitation will not be violated when the cooperators are on the same horizontal plane, as in [5] and [38]. However, when the data are from RSUs or from vehicles on a mountain road (as in Scenario 3 and 6 in Fig. 5), the height of the aggregated point clouds will be too large to tackle through traditional voxel processing. In our experiment settings, PointPillars [18] is the only algorithm to be compatible with the raw-level fusion scheme.

What's more, we also extend the MVX-Net to the collaborative autonomous driving scenarios. With the help of the point clouds from the LiDARs on the RSUs, whose locations are usually much higher, the ego vehicle can have a wider view with fewer occlusions. In addition, a single LiDAR on the RSU could free all the nearby autonomous vehicles from the necessity of equipping expensive LiDARs by sharing its point clouds through the V2I network, which brings great benefits to the realization of Level-5 autonomous driving. In this work, we use the point clouds from the RSU (or the aux vehicle 1 in Scenario 5) instead of the ego vehicle by transforming the coordinates.

Table 7 illustrates the multi-view collaborative perception on PointPillars [18] and MVX-Net [32], and the corresponding AP difference compared with stand-alone detection. It turns out that as for the superposition of raw point clouds, the

Table 7: Multi-view collaborative perception results on DOLPHINS

Scenario	Method	Car AP@IoU=0.7			Car AP@IoU=0.5			Pedestrian AP@IoU=0.5			Pedestrian AP@IoU=0.25		
		Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
1	PointPillar CP	97.13	95.19	94.57	97.63	96.40	95.89	72.73	70.60	70.23	94.01	93.69	93.31
	Difference	0.52%	3.37%	6.46%	-0.93%	0.22%	2.96%	3.72%	7.93%	10.49%	-1.63%	-0.73%	1.45%
	MVX-Net CP	89.34	84.30	84.02	89.64	89.49	87.00	90.77	86.08	85.92	99.61	99.53	97.04
2	Difference	0.10%	0.00%	0.04%	0.02%	-0.02%	-0.02%	-0.65%	-3.20%	-0.66%	0.00%	-0.02%	-0.02%
	PointPillar CP	97.97	97.03	96.42	98.52	97.89	97.55	57.51	51.89	50.92	91.94	88.52	86.54
	Difference	0.60%	4.27%	5.15%	0.30%	0.58%	2.17%	-3.59%	-2.70%	-0.84%	-0.63%	1.07%	2.43%
3	MVX-Net CP	93.75	90.99	88.47	96.41	96.19	93.81	88.52	81.72	79.25	99.29	98.76	94.13
	Difference	-0.02%	-0.05%	-0.18%	0.08%	0.29%	0.10%	-0.67%	-0.61%	-0.66%	-0.08%	0.04%	-0.06%
	PointPillar CP	81.48	72.73	66.30	86.88	77.69	72.33	32.62	19.89	18.47	73.24	48.12	44.80
4	Difference	3.22%	6.50%	16.77%	1.41%	2.56%	8.36%	-12.99%	-9.75%	-7.65%	-11.37%	-8.43%	-7.15%
	MVX-Net CP	69.04	58.48	48.77	71.61	61.60	54.18	70.82	39.76	37.43	91.87	56.60	51.76
	Difference	0.12%	0.02%	0.08%	-0.26%	-0.11%	-3.68%	-1.58%	-3.87%	0.89%	-2.26%	-0.04%	-3.88%
5	PointPillar CP	97.60	96.22	94.40	97.93	97.00	95.74	N/A	N/A	N/A	N/A	N/A	N/A
	Difference	-0.48%	2.36%	9.11%	-0.87%	-0.58%	5.20%	N/A	N/A	N/A	N/A	N/A	N/A
	MVX-Net CP	91.70	86.38	83.69	91.92	89.22	86.58	N/A	N/A	N/A	N/A	N/A	N/A
6	Difference	-0.07%	-0.01%	-0.05%	-0.04%	0.03%	0.03%	N/A	N/A	N/A	N/A	N/A	N/A
	PointPillar CP	96.38	94.24	92.77	96.87	95.82	94.38	65.50	61.53	58.69	93.00	91.62	87.82
	Difference	-1.10%	1.42%	3.89%	-1.96%	-0.73%	1.55%	-8.21%	-0.87%	4.30%	-4.44%	-2.10%	2.47%
7	MVX-Net CP	91.63	86.57	84.06	94.49	91.94	89.40	82.56	72.02	67.19	99.78	96.93	89.44
	Difference	-0.07%	-0.12%	-0.13%	0.04%	0.00%	0.01%	-2.72%	-5.02%	-2.50%	0.26%	-2.45%	-2.58%
	PointPillar CP	94.46	91.43	79.34	97.77	96.63	87.17	N/A	N/A	N/A	N/A	N/A	N/A
8	Difference	5.77%	11.09%	38.42%	0.34%	4.28%	23.44%	N/A	N/A	N/A	N/A	N/A	N/A
	MVX-Net CP	87.10	74.85	52.34	90.45	80.43	59.52	N/A	N/A	N/A	N/A	N/A	N/A
	Difference	-0.49%	-0.58%	-0.49%	-0.34%	-0.15%	2.67%	N/A	N/A	N/A	N/A	N/A	N/A

ego vehicle can gain plentiful benefits from the richer information directly from another perspective. Under those circumstances with severe occlusions such as Scenario 3 and 6 and for those hard objects, the cooperative perception-based PointPillars [18] achieves up to 38.42% increment in AP. However, the extra noise also infects the detection of small objects, which is also discussed in the supplementary material. On the other hand, as for the MVX-Net with the local camera and RSU LiDAR, the performance is nearly the same as the one with stand-alone sensors. It shows the opportunity to enable high-level autonomous driving on cheap, LiDAR-free vehicles through the sensors on infrastructures.

## 5 Conclusions

In this paper, we present a new large-scale various-scenario multi-view multi-modality autonomous driving dataset, DOLPHINS, to facilitate the research on collaborative perception-enabled connected autonomous driving. All the data are temporally-aligned and generated from three viewpoints, including both vehicles and RSUs, in six typical driving scenarios, along with the annotations, calibrations, and the geo-positions. What's more, we benchmark several SOTA algorithms on traditional 2D/3D object detection and brand-new collaborative perception tasks. The experiment results suggest that not only the extra data from V2X communication can eliminate the occlusions, but also the RSUs at appropriate locations can provide equivalent point clouds to the nearby vehicles, which can greatly reduce the prime cost of self-driving cars. In the future, we are going to further extend the number of infrastructures and aux vehicles, and construct more realistic maps of the downtown.

## References

1. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11621–11631 (2020)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *ECCV* (2020)
3. Chen, K., Wang, J., Pang, J., Cao, Y., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Xu, J., Zhang, Z., Cheng, D., Zhu, C., Cheng, T., Zhao, Q., Li, B., Lu, X., Zhu, R., Wu, Y., Dai, J., Wang, J., Shi, J., Ouyang, W., Loy, C.C., Lin, D.: MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155* (2019)
4. Chen, Q., Ma, X., Tang, S., Guo, J., Yang, Q., Fu, S.: F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds. In: *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*. p. 88–100. SEC '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3318216.3363300>, <https://doi.org/10.1145/3318216.3363300>
5. Chen, Q., Tang, S., Yang, Q., Fu, S.: Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds. In: *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. pp. 514–524. IEEE (2019)
6. Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3d object detection network for autonomous driving. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. pp. 1907–1915 (2017)
7. Contributors, M.: MMDetection3D: OpenMMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d> (2020)
8. Deng, J., Czarnecki, K.: Mlod: A multi-view 3d object detection based on robust feature fusion method. In: *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. pp. 279–284. IEEE (2019)
9. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An open urban driving simulator. In: *Proceedings of the 1st Annual Conference on Robot Learning*. pp. 1–16 (2017)
10. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet: Keypoint triplets for object detection. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6569–6578 (2019)
11. Feng, C., Zhong, Y., Gao, Y., Scott, M.R., Huang, W.: Tood: Task-aligned one-stage object detection. In: *ICCV* (2021)
12. Ge, Z., Liu, S., Wang, F., Li, Z., Sun, J.: YOLOX: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430* (2021)
13. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2012)
14. Girshick, R.: Fast r-cnn. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (December 2015)
15. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 580–587 (2014)



16. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
17. Kim, S.W., Liu, W., Ang, M.H., Frazzoli, E., Rus, D.: The impact of cooperative perception on decision making and planning of autonomous vehicles. *IEEE Intelligent Transportation Systems Magazine* **7**(3), 39–50 (2015)
18. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 12697–12705 (2019)
19. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: Proceedings of the European conference on computer vision (ECCV). pp. 734–750 (2018)
20. Li, P., Chen, X., Shen, S.: Stereo r-cnn based 3d object detection for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7644–7652 (2019)
21. Li, Y., An, Z., Wang, Z., Zhong, Y., Chen, S., Feng, C.: V2x-sim: A virtual collaborative perception dataset for autonomous driving (2022). <https://doi.org/10.48550/ARXIV.2202.08449>, <https://arxiv.org/abs/2202.08449>
22. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft coco: Common objects in context (2014). <https://doi.org/10.48550/ARXIV.1405.0312>, <https://arxiv.org/abs/1405.0312>
23. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: European conference on computer vision. pp. 21–37. Springer (2016)
24. Nassar, A.S., D’aronco, S., Lefèvre, S., Wegner, J.D.: Geograph: Graph-based multi-view object detection with geometric cues end-to-end. In: European Conference on Computer Vision. pp. 488–504. Springer (2020)
25. Nassar, A.S., Lefèvre, S., Wegner, J.D.: Simultaneous multi-view instance detection with learned geometric soft-constraints. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6559–6568 (2019)
26. Rawashdeh, Z.Y., Wang, Z.: Collaborative automated driving: A machine learning-based method to enhance the accuracy of shared information. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC). pp. 3961–3966. IEEE (2018)
27. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 779–788 (2016)
28. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7263–7271 (2017)
29. Redmon, J., Farhadi, A.: YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018)
30. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
31. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10529–10538 (2020)
32. Sindagi, V.A., Zhou, Y., Tuzel, O.: Mvx-net: Multimodal voxelnet for 3d object detection. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 7276–7282. IEEE (2019)

33. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2446–2454 (2020)
34. Wang, T.H., Manivasagam, S., Liang, M., Yang, B., Zeng, W., Urtasun, R.: V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In: European Conference on Computer Vision. pp. 605–621. Springer (2020)
35. Wang, Y., Guizilini, V.C., Zhang, T., Wang, Y., Zhao, H., Solomon, J.: Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In: Conference on Robot Learning. pp. 180–191. PMLR (2022)
36. Xiao, Z., Mo, Z., Jiang, K., Yang, D.: Multimedia fusion at semantic level in vehicle cooperative perception. In: 2018 IEEE International Conference on Multimedia & Expo Workshops (ICMEW). pp. 1–6. IEEE (2018)
37. Xu, R., Xiang, H., Tu, Z., Xia, X., Yang, M.H., Ma, J.: V2x-vit: Vehicle-to-everything cooperative perception with vision transformer (2022). <https://doi.org/10.48550/ARXIV.2203.10638>, <https://arxiv.org/abs/2203.10638>
38. Xu, R., Xiang, H., Xia, X., Han, X., Li, J., Ma, J.: Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication (2021). <https://doi.org/10.48550/ARXIV.2109.07644>, <https://arxiv.org/abs/2109.07644>
39. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors (2018)
40. Yang, Q., Fu, S., Wang, H., Fang, H.: Machine-learning-enabled cooperative perception for connected autonomous vehicles: Challenges and opportunities. IEEE Network **35**(3), 96–101 (2021)
41. Yu, H., Luo, Y., Shu, M., Huo, Y., Yang, Z., Shi, Y., Guo, Z., Li, H., Hu, X., Yuan, J., et al.: Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21361–21370 (2022)
42. Zeadally, S., Javed, M.A., Hamida, E.B.: Vehicular communications for its: standardization and challenges. IEEE Communications Standards Magazine **4**(1), 11–17 (2020)