

V2V4Real: A Real-world Large-scale Dataset for Vehicle-to-Vehicle Cooperative Perception

Runsheng Xu¹, Xin Xia^{1*}, Jinlong Li², Hanzhao Li¹, Shuo Zhang¹, Zhengzhong Tu³, Zonglin Meng¹, Hao Xiang¹, Xiaoyu Dong⁴, Rui Song^{5,6}, Hongkai Yu², Bolei Zhou¹, Jiaqi Ma¹

University of California, Los Angeles¹ Cleveland State University² University of Texas at Austin³
 Northwestern University⁴ Technical University of Munich⁵ Fraunhofer Institute⁶

Abstract

Modern perception systems of autonomous vehicles are known to be sensitive to occlusions and lack the capability of long perceiving range. It has been one of the key bottlenecks that prevents Level 5 autonomy. Recent research has demonstrated that the Vehicle-to-Vehicle (V2V) cooperative perception system has great potential to revolutionize the autonomous driving industry. However, the lack of a real-world dataset hinders the progress of this field. To facilitate the development of cooperative perception, we present V2V4Real, the first large-scale real-world multi-modal dataset for V2V perception. The data is collected by two vehicles equipped with multi-modal sensors driving together through diverse scenarios. Our V2V4Real dataset covers a driving area of 410 km, comprising 20K LiDAR frames, 40K RGB frames, 240K annotated 3D bounding boxes for 5 classes, and HDMaps that cover all the driving routes. V2V4Real introduces three perception tasks, including cooperative 3D object detection, cooperative 3D object tracking, and Sim2Real domain adaptation for cooperative perception. We provide comprehensive benchmarks of recent cooperative perception algorithms on three tasks. The V2V4Real dataset can be found at research.seas.ucla.edu/mobility-lab/v2v4real/.

1. Introduction

Perception is critical in autonomous driving (AV) for accurate navigation and safe planning. The recent development of deep learning brings significant breakthroughs in various perception tasks such as 3D object detection [22, 35, 42], object tracking [43, 56], and semantic segmentation [47, 57]. However, single-vehicle vision systems still suffer from many real-world challenges, such as occlusions and short-range perceiving capability [15, 40, 49], which

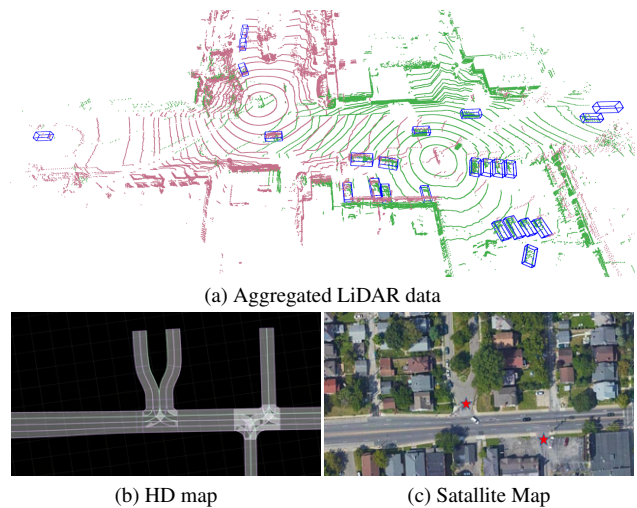


Figure 1. A data frame sampled from V2V4Real: (a) aggregated LiDAR data, (b) HD map, and (c) satellite map to indicate the collective position. More qualitative examples of V2V4Real can be found in the supplementary materials.

can cause catastrophic accidents. The shortcomings stem mainly from the limited field-of-view of the individual vehicle, leading to an incomplete understanding of the surrounding traffic.

A growing interest and recent advancement in cooperative perception systems have enabled a new paradigm that can potentially overcome the limitation of single-vehicle perception. By leveraging vehicle-to-vehicle (V2V) technologies, multiple connected and automated vehicles (CAVs) can communicate and share captured sensor information simultaneously. As shown in a complex intersection in Fig. 1, for example, the ego vehicle (red LiDAR) struggles to perceive the upcoming objects located across the way due to occlusions. Incorporating the LiDAR features from the nearby CAV (green scans) can largely broaden the sensing range of the vehicle and make it even see across the occluded corner.

Despite the great promise, however, it remains chal-

*Corresponding Author, email address: x35xia@ucla.edu

Dataset	Year	Real/ Sim	V2X	Size (km)	RGB images	LiDAR	Maps	3D boxes	Classes	Locations
Kitti [14]	2012	Real	No	-	15k	15k	No	200k	8	Karlsruhe
nuScenes [2]	2019	Real	No	33	1.4M	400k	Yes	1.4M	23	Boston, SG
Argo [3]	2019	Real	No	290	107k	22k	Yes	993k	15	2x USA
Waymo Open [36]	2019	Real	No	-	1M	200k	Yes	12M	4	3x USA
OPV2V [50]	2022	Sim	V2V	-	44k	11k	Yes	230k	1	CARLA
V2X-Sim [20]	2022	Sim	V2V&I	-	60K	10k	Yes	26.6k	1	CARLA
V2XSet [49]	2022	Sim	V2V&I	-	44K	11k	Yes	230k	1	CARLA
DAIR-V2X [54]	2022	Real	V2I	20	39K	39K	No	464K	10	Beijing, CN
V2V4Real (ours)	2022	Real	V2V	410	40K	20K	Yes	240K	5	Ohio, USA

Table 1. Comparison of the proposed dataset and existing representative autonomous driving datasets.

lenging to validate V2V perception in real-world scenarios due to the lack of public benchmarks. Most of the existing V2V datasets, including OPV2V [50], V2X-Sim [20], and V2XSet [49], rely on open-source simulators like CARLA [11] to generate synthetic road scenes and traffic dynamics with simulated connected vehicles. However, it is well known that there exists a clear domain gap between synthetic data and real-world data, as the traffic behavior and sensor rendering in simulators are often not realistic enough [27, 37]. Hence, models trained on these benchmarks may not generalize well to realistic driving situations.

To further advance innovative research on V2V cooperative perception, we present a large-scale multimodal and multitask V2V autonomous driving dataset, which covers 410 km road and contains 20K LiDAR frames with more than 240K 3D bounding box annotations. Compared to the only existing real-world cooperative dataset DAIR-V2X [54], our proposed V2V4Real dataset shows several strengths: (1) DAIR-V2X focuses on Vehicle-to-Infrastructure (V2I) applications without supporting V2V perception. Compared to V2I, V2V does not require the pre-installed sensors restricted in a certain area, which is more flexible and scalable. Our dataset fills the gap by focusing on the important V2V cooperation. (2) V2V4Real includes four diverse road types, including intersection, highway entrance ramp, highway straight road, and city straight road, covering broader driving areas and greater mileage. (3) We also provide high-definition (HD) maps that can be used for road topology prediction and semantic bird’s-eye-view (BEV) map understanding. (4) We construct several benchmarks that can train and evaluate recent autonomous perception algorithms, including 3D object detection, object tracking, and Sim2Real domain adaption, while DAIR-V2X only has a single track. 5) We have provided 8 state-of-the-art cooperative perception algorithms for benchmarking, whereas DAIR-V2X only implements 3 baseline methods. Unlike DAIR-V2X, which can be only accessed within China¹, we will make all the data, bench-

marks, and models publically available across the globe. Our contributions can be summarized as follows:

- We build the V2V4Real, a large real-world dataset dedicated to V2V cooperative autonomous perception. All the frames are captured by multi-modal sensor readings from real-world diverse scenarios in Columbus, Ohio, in the USA.
- We provide more than 240K annotated 3D bounding boxes for 5 vehicle classes, as well as corresponding HDMaps along the driving routes, which enables us to train and test cooperative perception models in real-world scenarios.
- We introduce three cooperative perception tasks, including 3D object detection, object tracking, and Sim2Real, providing comprehensive benchmarks with several SOTA models. The results show the effectiveness of V2V cooperation in multiple tasks.

2. Related Work

2.1. Autonomous Driving Datasets.

Public datasets have contributed to the rapid progress of autonomous driving technologies in recent years. Tab. 1 summarizes the recent autonomous driving datasets. The earlier datasets mainly focus on 2D annotations (boxes, masks) for RGB camera images, such as Cityscapes [8], Synthia [34], BDD100K [53], to name a few. However, achieving human-level autonomous driving requires accurate perception and localization in the 3D real world, whereas learning the range or depth information from pure 2D images is an ill-posed problem.

To enable robust perception in 3D or map-view, multimodal datasets that typically involve not only camera images but also range data such as Radar or LiDAR sensors have been developed [2, 14, 36]. KITTI [14] was a pioneering dataset that provides multimodal sensor readings, including front-facing stereo camera and LiDAR for 22 sequences, annotated with 200k 3D boxes and tasks of 3D object detection, tracking, stereo, and optical flow. Subsequently, NuScenes [2] and Waymo Open dataset [36] is

¹<https://thudair.baai.ac.cn/index>

the most recent multimodal datasets providing an orders-of-magnitude larger number of scenes (over 1K), with 1.4M and 993K annotated 3D boxes, respectively. Despite remarkable progress, those datasets only aim at developing single-vehicle driving capability, which has been demonstrated to have limited ability to handle severe occlusions as well as long-range perception [40, 45, 46, 50].

The recent development of V2V technologies has made it possible for vehicles to communicate and fuse multimodal features collaboratively, thus yielding a much broader perception range beyond the limit of single-view methods. OPV2V [50] builds the first-of-a-kind 3D cooperative detection dataset using CARLA and OpenCDA co-simulation. V2XSet [49] and V2X-Sim [20] further explore the viability of vehicle-to-everything (V2X) perception using synthesized data generated from CARLA simulator [11]. Unlike the above-simulated datasets, DAIR-V2X is the first real-world dataset for cooperative detection. However, DAIR-V2X only concentrates on V2I cooperation, neglecting the important V2V application, which can be more flexible and more likely to be scalable. As V2V and V2I perception has major differences, i.e., V2V perception needs to deal with more diverse traffic scenarios and occlusions [49], a real-world dataset for V2V perception is needed. Furthermore, DAIR-V2X only spans limited road types (i.e., only intersections) and constrained driving route length (only 20km).

2.2. 3D Detection

3D object detection plays a critical role in the success of autonomous driving. Based on available sensor modality, 3D detection has roughly three categories. (1) **Camera-based** detection denotes approaches that detect 3D objects from a single or multiple RGB images [16, 32, 33, 35, 42]. For instance, ImVoxelNet [35] builds a 3D volume in 3D world space and samples multi-view features to obtain the voxel representation. DETR3D [42] models 3D objects using queries to index into extracted 2D multi-camera features, which directly estimate 3D bounding boxes in 3D spaces. Especially, an additional detection head and attention modules can further improve small object detection accuracy for RGB images [23] (2) **LiDAR-based** detection typically converts LiDAR points into voxels or pillars, resulting in 3D voxel-based [51, 58] or 2D pillar-based methods [18, 52]. Since 3D voxels are usually expensive to process, PointPillars [18] propose to compress all the voxels along the z -axis into a single pillar, then predicting 3D boxes in the bird’s-eye-view space. Benefiting from its fast processing and real-time performance, many recent 3D object detection models follow this pillar-based approach [12, 41]. (3) **Camera-LiDAR fusion** presents a recent trend in 3D detection that fuses information from both image and LiDAR points. One of the key challenges in multimodal fusion is how to align the image features with point

clouds. Some methods [29, 39] use a two-step framework, e.g., first detect the object in 2D images, then use the obtained information to further process point clouds; more recent works [21, 28] develop end-to-end fusion pipelines and leverage cross-attention [38] to perform feature alignment.

2.3. V2V/V2X Cooperative Perception

Due to the intrinsic limitation of camera/LiDAR devices, occlusions and long-distance perception are extremely challenging for single-vehicle systems, which can potentially cause catastrophic consequences in complex traffic environments [50]. Cooperative systems, on the other hand, can unlock the possibility of multi-vehicle detection that tackles the limitation of single-vehicle perception. Among these, V2V (Vehicle-to-Vehicle) approaches center on collaborations between vehicles, while V2X (Vehicle-to-Everything) involves correspondence between vehicles and infrastructure. V2V/V2X cooperative perception can be roughly divided into three categories: (1) Early Fusion [5] where raw data is shared among CAVs, and the ego vehicle makes predictions based on the aggregated raw data, (2) Late Fusion [31] where detection outputs (e.g., 3D bounding boxes, confidence scores) are shared, then fused to a ‘consensus’ prediction, and (3) Intermediate Fusion [4, 26, 40, 50] where intermediate representations are extracted based on each agent’s observation and then shared with CAVs.

Recent methods typically choose the intermediate neural features computed from each agent’s sensor data as the transmitted features, which achieves the best trade-off between accuracy and bandwidth requirements. For instance, V2VNet [40] adopted graph neural networks to fuse intermediate features. F-Cooper [4] employed max-pooling fusion to aggregate shared Voxel features. Cooperant [10] used Point Transformer [55] to deliver point features and conduct experiments under AustoCastSim [30]. CoBEVT [47] proposed local-global sparse attention that captures complex spatial interactions across views and agents to improve the performance of cooperative BEV map segmentation. AttFuse [50] proposed an agent-wise self-attention module to fuse the received intermediate features. V2X-ViT [49] presented a unified vision transformer for multi-agent multi-scale perception and achieves robust performance under GPS error and communication delay.

3. V2V4Real Dataset

To expedite the development of V2V Cooperative Perception for autonomous driving, we propose **V2V4Real**, the real-world, large-scale, multi-modal dataset with diverse driving scenarios. This dataset is annotated with both 3D bounding boxes and HDMaps for the research of multi-vehicle cooperative perception. In this section, we first detail the setup of data collection (Sec. 3.1), and then describe

the data annotation approach (Sec. 3.2), and finally analyze the data statistics (Sec. 3.3).



Figure 2. **The information of the collection vehicles.**a) The Tesla vehicle. b) The Ford Fusion vehicle. c) The sensor setup for both vehicles. Note that the photo of Tesla is taken from the rear camera of Ford, and that of Ford is taken from the front camera of Tesla.

3.1. Data Acquisition

Sensor Setup. We collect the V2V4Real via two experimental connected automated vehicles including a Tesla vehicle (Fig. 2a) and a Ford Fusion vehicle (Fig. 2b) retrofitted by Transportation Research Center (TRC) company and AutonomouStuff (AStuff) Company respectively. Both vehicles are equipped with a Velodyne VLP-32 LiDAR sensor, two mono cameras (front and rear), and GPS/IMU integration systems. The sensor layout configuration can be found in Fig. 2c, and the detailed parameters are listed in Table. 2.

Driving Route. The two vehicles drive simultaneously in Columbus, Ohio, and their distance is maintained within 150 meters to ensure overlap between their views. To enrich the diversity of sensor-view combinations, we vary the relative poses of the two vehicles across different scenarios (see Sec. 3.3 for details). We collect driving logs for three days that cover 347 km of highway road and 63 km of city road. The driving routes are visualized in Fig. 3, wherein the red route is on day 1 (freeway with one to five lanes), the yellow is on day 2 (city road, one to two lanes), and the green route is on day 3 (highway, two to four lanes).

Data Collection. We collect 19 hours of driving data of 310K frames. We manually select the most representative 67 scenarios, each 10-20 seconds long. We sample the frames at 10Hz, resulting in a total of 20K frames of LiDAR point cloud and 40K frames of RGB images. For each scene, we ensure that the asynchronizations between two vehicles’ sensor systems are less than 50 ms. All the sce-

Sensors	Details
2x Camera	RGB, Tesla: 1280×720 , Ford: 2064×1544 32 channels, 1.2 M points per second,
1x LiDAR	200 m capturing range, -25° to 15° vertical FOV, ± 3 cm error, 10Hz
GPS & IMU	Tesla: RT3000, Ford: Novatel SPAN E1

Table 2. Sensor specifications for each vehicle.

narios are aligned with maps containing drivable regions, road boundaries, as well as dash lines.

3.2. Data Annotation




Figure 3. Driving routes of our two collection vehicles. Different colors represent the routes collected on different days.

Coordinate System. Our dataset includes four different coordinate systems: the LiDAR coordinate system for Tesla and Ford Fusion, the HDmap coordinate, and the earth-earth, fixed-coordinate (ECEF). We annotate the 3D bounding boxes separately based on each vehicle’s LiDAR coordinate system such that each vehicle’s sensor data alone can also be treated as single-agent detection tasks. We utilize the positional information provided by GPS on the two vehicles to initialize the relative pose of the two vehicles for each frame. The origin of the HDMap aligns with the initial frame of Tesla for each driving route.

3D Bounding boxes annotation. We employ SusTech-Point [19], a powerful opensource labeling tool, to annotate 3D bounding boxes for the collected LiDAR data. We hire two groups of professional annotators. One group is responsible for the initial labeling, and the other further refines the annotations. There are five object classes in total, including cars, vans, pickup trucks, semi-truck, and buses. For each object, we annotate its 7-degree-of-freedom 3D bounding box containing x, y, z for the centroid position and l, w, h, yaw for the bounding box extent and yaw angles. We also record each object’s driving state (*i.e.* dynamic or parking). To facilitate downstream applications such as tracking and behavior prediction, we assign consistent id and size for the same object in different timestamps.

Since the bounding boxes are annotated separately for the two collection vehicles, an object in the Tesla’s frame could have the same id as a different object in Ford Fusion’s frame. To avoid such issues, all the object ids in Tesla are labeled between 0 – 1000, while ids in Ford Fusion range




Figure 4. The distribution of the relative poses between the two collection vehicles.




Figure 5. The distribution of vehicle types in collected dataset.

from 1001 – 2000. Moreover, identical objects could have different ids in the annotation files of the two collection vehicles. To solve this issue, we transform the objects from different coordinates to a unified coordinate system and calculate the BEV IoU between all objects. For the objects that have IoU larger than a certain threshold, we assign them the same object id and unify their bounding box sizes.

Map Annotation. The HD map generation pipeline refers to generating a global point cloud map and vector map. To generate the point cloud map, we fuse a sequence of point cloud frames together. More specifically, we first preprocess each LiDAR frame by removing the dynamic objects while keeping the static elements. Then, a Normal Transformation Distribution scan matching algorithm is applied to compute the relative transformation between two consecutive LiDAR frames. The LiDAR odometry can then be constructed by taking the transformation. However, the noise imbued in the LiDAR data can lead to accumulated errors in the estimated transformation matrix as the frame index increases. Therefore, we compensate for these errors by further integrating the translation and heading information provided by the on-vehicle GPS/IMU system and applying Kalman filter [7]. Finally, all the points in different frames are transformed onto the map coordinate to form a global point cloud map. The aggregated point cloud maps will be imported to RoadRunner [9] to produce the vector




Figure 6. **Left:** Number of LiDAR points (e -based log scale) within ground truth bounding boxes with respect to radial distance from the ego vehicle. **Right:** Bounding box size distributions.

maps. The road is drawn and inferred from the intensity information visualized by distinct colors in Roadrunner. We then output the OpenDRIVE (Xodr) maps and convert them to lanelet maps [1] as the final format.

3.3. Data Analysis

Fig. 4 reveals the distribution of relative poses between the two collection vehicles across all scenarios. It can be observed that the two vehicles have a variety of relative poses, generating diverse view combinations of scenes. As Fig. 5 describes, most of the objects in V2V4Real belong to the Car class, while Pickup Truck ranks second. The number of Vans and Semi-Trucks are similar, while Bus has the least quantities. Fig. 6 shows the LiDAR points density distribution inside different objects bounding boxes and the bounding boxes’ size distribution. As we may see in the left figure, when there is only one vehicle (Tesla) scanning the environment, the number of LiDAR points within bounding boxes drops dramatically as the radial distance increases. Enhanced by the shared visual information from the other vehicle (Ford Fusion), the LiDAR point density of each object increases significantly and still retains at a high level even when the distance reaches 100 m. This validates the great benefits that cooperative perception can bring to the system. As the right figure reveals, the annotated objects have diverse bounding box sizes, with lengths ranging from 2.5 m to 23 m, widths ranging from 1.5 m to 4.5 m, and heights ranging from 1 m to 4.5 m, demonstrating the diversity of our data.

4. Tasks

Our dataset supports multiple cooperative perception tasks, including detection, tracking, prediction, localization, etc. In this paper, we focus on cooperative detection, tracking, and Sim2Real transfer learning tasks.

4.1. Cooperative 3D Object Detection

Scope. The V2V4Real detection task requires users to leverage multiple LiDAR views from different vehicles to perform 3D object detection on the ego vehicle. Compared to the single-vehicle detection task, cooperative detection has several domain-specific challenges:

- **GPS error:** There exists unavoidable error in the relative pose of the collaborators [24], which can produce global misalignments when transforming the data into a unified coordinate system.
- **Asynchronicity:** The sensor measurements of collaborators are usually not well-synchronized, which is caused by the asynchrony of the distinct sensor systems as well as the communication delay during the data transmission process [49].
- **Bandwidth limitation:** Typical V2V communication technologies require restricted bandwidth, which limits the transmitted data size [31, 40, 49]. Therefore, cooperative detection algorithms must consider the trade-off between accuracy and bandwidth requirements.

The major mission of this track is to design efficient cooperative detection methods to handle the above challenges.

Groundtruth. During training or testing, one of the two collection vehicles will be selected as the ego vehicle, and the other will transform its annotated bounding boxes to the ego’s coordinate. In this way, the groundtruth is defined in a unified (the ego) coordinate system. Note that in the training phase, the ego vehicle is randomly picked, while during testing, we fix Tesla as ego. Due to asynchronicity and localization errors, the bounding boxes from two vehicles corresponding to the same object have some offsets. In such a case, we select the one annotated in the ego vehicle as the groundtruth.

Evaluation. The evaluation range in x and y direction are $[-100, 100]$ m and $[-40, 40]$ m with respect to the ego vehicle. Similar to DAIR-V2X [54], we categorize different vehicle types as the same class and focus only on vehicle detection. We use the Average Precision (AP) at Intersection-over-Union (IoU) 0.5 and 0.7 as the metric to evaluate the performance of vehicle detection. To assess the transmission cost, Average MegaByte (AM) is employed, which represents the transmitted data size specified by the algorithm. Following [49, 54], we evaluate all the models under two settings: 1) *Sync* setting, under which the data transmission is regarded as instantaneous, whereas the asynchrony is only induced by the distinct cycles of the sensor systems. 2) *Async* setting, where we consider the data transmission delay as 100 ms. We simulate such communication delay by retrieving the LiDAR data from the previous timestamp from the non-ego vehicle.

Benchmarking methods. We evaluate most commonly adopted fusion strategies as Fig. 7 demonstrated for coop-

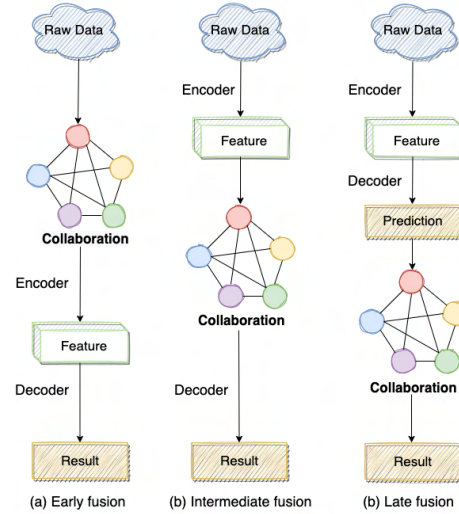


Figure 7. **The three different fusion strategies:** (a) Early Fusion, (b) Intermediate Fusion, and (c) Late Fusion.

erative perception with state-of-the-art methods in the domain. In total, four fusion strategies are considered:

- *No Fusion:* Only ego vehicle’s point cloud is used for visual reasoning. This strategy serves as the baseline.
- *Late Fusion:* Each vehicle detects 3D objects utilizing its own sensor observations and delivers the predictions to others. Then the receiver applies Non-maximum suppression to produce the final outputs.
- *Early Fusion:* The vehicles will directly transmit the raw point clouds to other collaborators and the ego vehicle will aggregate all the point clouds to its own coordinate frame, which preserves complete information but requires large bandwidths.
- *Intermediate Fusion:* The collaborators will first project their LiDAR to the ego vehicle’s coordinate system and then extract intermediate features using a neural feature extractor. Afterward, the encoded features are compressed and broadcasted to the ego vehicle for cooperative feature fusion. We benchmark a number of leading intermediate methods, including AttFuse [50], F-Cooper [4], V2VNet [40], V2X-Vit [49], and CoBEVT [47] (see Sec. 2.3 for detail descriptions). Similar to previous works [47, 49, 50], we train a simple auto-encoder to compress the intermediate features by $32\times$ to save bandwidth and decompress them to the original size on the ego side.

4.2. Object Tracking

Scope. In this track, we study whether and how object tracking models can obtain benefits from the cooperative system. There are two major approaches to tracking algorithms: joint detection and tracking and tracking by detection. In this paper, we focus on the second class.

Method	Sync (AP@IoU=0.5/0.7)				Async (AP@IoU=0.5/0.7)				AM (MB)
	Overall	0-30m	30-50m	50-100m	Overall	0-30m	30-50m	50-100m	
No Fusion	39.8/22.0	69.2/42.6	29.3/14.4	4.8/1.6	39.8/22.0	69.2/42.6	29.3/14.4	4.8/1.6	0
Late Fusion	55.0/26.7	73.5/36.8	43.7/22.2	36.2/17.3	50.2/22.4	70.7/34.2	41.0/19.8	26.1/7.8	0.003
Early Fusion	59.7/32.1	76.1/46.3	42.5/20.8	47.6/ 21.1	52.1/25.8	74.6/43.6	34.5/16.3	30.2/ 9.5	0.96
F-Cooper [4]	60.7/31.8	80.8/46.9	45.6/23.6	32.8/13.4	53.6/26.7	79.0/44.1	38.7/19.5	18.1/6.0	0.20
V2VNet [40]	64.5/34.3	80.6/51.4	52.6/26.6	42.6/14.6	56.4/28.5	78.6/48.0	44.2/21.5	25.6/6.9	0.20
AttFuse [50]	64.7/33.6	79.8/44.1	53.1/29.3	43.6/19.3	57.7/27.5	78.6/41.4	45.5/23.8	27.2/9.0	0.20
V2X-ViT [49]	64.9/ 36.9	82.0/ 55.3	51.7/26.6	43.2/16.2	55.9/29.3	79.7/ 50.4	43.3/21.1	24.9/7.0	0.20
CoBEVT [47]	66.5/36.0	82.3/51.1	52.1/28.2	49.1/19.5	58.6/29.7	80.3/48.3	44.7/22.8	30.5/8.7	0.20

Table 3. Cooperative 3D object detection benchmark.

Evaluation. We employ the same evaluation metrics in [2, 43] for object tracking, including 1) Multi Object Tracking Accuracy (MOTA), 2) Mostly Tracked Trajectories (MT), 3) Mostly Lost Trajectories (ML), 4) Average Multiobject Tracking Accuracy (AMOTA), 5) Average Multiobject Tracking Precision (AMOTP), and 6) scaled Average Multiobject Tracking Accuracy (sAMOTA). Specifically, the AMOTA and AMOTP average MOTA and MOTP across all recall thresholds, which takes into account the prediction confidence, compared to traditional MOTA and MOTP metrics. sAMOTA is proposed by [2] to guarantee a more linear span over the entire $[0, 1]$ range significantly difficult tracking tasks.

Baselines tracker. We implement AB3Dmot tracker [43] as our baseline tracker. Given the detection results from the cooperative detection models, AB3Dmot combines the 3D Kalman Filter with Birth and Death Memory technique to achieve an efficient and robust tracking performance.

4.3. Sim2Real Domain Adaptation

Scope. Data labeling is time-consuming and expensive for the perception system [44]. When it comes to cooperative perception, the cost can dramatically expand as the labelers need to annotate multiple sensor views, which is impossible to scale up. A potential solution is to employ infinite and inexpensive simulation data. However, it is known that there is a significant domain gap between simulated and real-world data distributions. Therefore, this track investigates how to utilize domain adaptation methods to reduce domain discrepancy in the cooperative 3D detection task.

Training. We define the target domain as the V2V4Real dataset and the source domain as a large-scale open simulated OPV2V dataset [50]. The training data consists of two parts: the OPV2V training set with provided annotations, and V2V4Real training set’s LiDAR point cloud without access to the labels. Participants should leverage domain adaptation algorithms to enable the cooperative detection models to generate domain-invariant features.

Evaluation. The evaluation will be conducted on the test set of V2V4Real dataset under the *Sync* setting, and the assessment protocol is the same as the cooperative 3D object

detection track.

Evaluated methods. The baseline method is to train the detection models on OPV2V and directly test on V2V4Real without any domain adaptation. To demonstrate the effectiveness of domain adaptation, we implement a similar method as in [6], which applies two domain classifiers for feature-level and object-level adaption and utilizes gradient reverse layer (GRL) [13] to backpropagate the gradient to assist the model for generating domain-invariant features.

5. Experiments

5.1. Implementation Details

The dataset is split into the train/validation/test set with 14,210/2,000/3,986 frames, respectively, for all three tasks. All the detection models employ PointPillar [18] as the backbone to extract 2D features from the point cloud. we train all models with 60 epochs, a batch size of 4 per GPU (RTX3090), a learning rate of 0.001, and we decay the learning rate with a cosine annealing [25]. Early stopping is used to find the best epoch. We also add normal point cloud data augmentations for all experiments, including scaling, rotation, and flip [18]. We employ AdamW [17] with a weight decay of 1×10^{-2} to optimize our models. For the tracking task, we take the previous 3 frames together with the current frame as the inputs.

5.2. 3D LiDAR Object Detection

Tab. 3 demonstrates the quantitative comparison between various cooperative 3D detection models on our V2V4Real dataset. We can observe that:

- Compared to the single-vehicle perception baseline, all cooperative perception methods can significantly boost performance by at least 15.2% in terms of overall AP at IoU 0.5. Furthermore, the accuracy of all evaluation ranges is improved, whereas long-range detection has the most benefits with a minimum of 28.0% and 11.8% gain for AP@0.5 and AP@0.7, respectively.
- Under both *Sync* and *Async* settings, intermediate fusion methods achieve the best trade-off between accu-

Method	AMOTA(\uparrow)	AMOTP(\uparrow)	sAMOTA(\uparrow)	MOTA(\uparrow)	MT(\uparrow)	ML(\downarrow)
No Fusion	16.08	41.60	53.84	43.46	29.41	60.18
Late Fusion	29.28	51.08	71.05	59.89	45.25	31.22
Early Fusion	26.19	48.15	67.34	60.87	40.95	32.13
F-Cooper [4]	23.29	43.11	65.63	58.34	35.75	38.91
AttFuse [50]	28.64	50.48	73.21	63.03	46.38	28.05
V2VNet [40]	30.48	54.28	75.53	64.85	48.19	27.83
V2X-ViT [49]	30.85	54.32	74.01	64.82	45.93	26.47
CoBEVT [47]	32.12	55.61	77.65	63.75	47.29	30.32

Table 4. **Cooperative Tracking benchmark.** All numbers represent percentages.

racy and transmission cost. Among all the intermediate fusion methods, CoBEVT has the best performance in terms of AP@0.5, 1.6% higher than the second best model V2X-ViT, 6.8% higher than *Early Fusion*, and 11.5% higher than *Late Fusion* in the *Sync* setting.

- Except for *No Fusion*, all other methods’ AP dropped significantly when the communication delay was introduced. For instance, CoBEVT, V2X-ViT, and V2VNet drops 6.3%, 7.6%, and 5.8% at AP@0.7, respectively. This observation highlighted the importance of robustness to the asynchrony for cooperative perception methods.

5.3. 3D Object Tracking

Tab. 4 shows the benchmark results for cooperative tracking. It can be seen that when AB3Dmot combines with cooperative detection, the performance is dramatically better than the single-vehicle tracking method. Similar to the cooperative detection track, CoBEVT [47] achieves the best performance in most of the evaluation metrics, including AMOTA (16.04% higher than baseline), sAMOTA (23.81% higher than baseline), and AMOTP (14.01% better than baseline).

5.4. Sim2Real Domain Adaptation

As Tab. 5 reveals, there exist serious domain gaps between the simulated dataset OPV2V and our real-world dataset V2V4Real. Without any domain adaptation, only seeing the simulated data will decrease the accuracy of the detection models by 42.2%, 37.1%, 41.3%, 37.5%, 33.9 for AttFuse, F-Cooper, V2VNet, V2X-ViT, and CoBEVT. Applying the domain adaption technique alleviates the performance drop by an average of 7.46%. Furthermore, the strongest model, CoBEVT, can reach 40.2% after employing the domain adaptation, which is higher than the *No Fusion* baseline method that uses real-world data for training.

6. Conclusion

We present V2V4Real, a large-scale real-world dataset that covers up to 410 km driving areas, contains 20K LiDAR frames, 40K RGB images, and are annotated with

Method	AP@IoU=0.5	AP drop
AttFuse [50]	22.5	42.2
AttFuse w/ D.A.	23.4 (+0.9)	41.3
F-Cooper [4]	23.6	37.1
F-Cooper w/ D.A.	37.3 (+13.7)	23.4
V2VNet [40]	23.2	41.3
V2VNet w/ D.A.	26.3 (+3.1)	38.2
V2X-ViT [49]	27.4	37.5
V2X-ViT w/ D.A.	39.5 (+12.1)	25.4
CoBEVT [47]	32.6	33.9
CoBEVT w/ D.A.	40.2 (+7.6)	26.3

Table 5. **Domain Adaptation benchmark.** The number in the bracket indicates the precision gain when using domain adaptation. AP drop refers to the precision gap compared to directly training on the V2V4Real dataset.

240K bounding boxes as well as HDMaps, to promote V2V cooperative perception research. We further introduce three V2V perception benchmarks involving 3D object detection, object tracking, and Sim2Real domain adaptation, which opens up the possibility for future task development. V2V4Real will be made fully available to the public to accelerate the progress of this new field. We plan to release the benchmarks and baseline models for HDMap learning tasks and camera images in the next version.

Broader impact. Although the proposed benchmark covers various driving scenes for V2V perception, there may still exist extremely challenging scenarios that do not appear in our training set. In such cases, the models should be trained more carefully in order not to hinder generalization abilities. Out-of-distribution detection is also an important topic that has not been investigated within the scope of this paper. These issues should be taken care of by future related research for robust and safe autonomous perception.

7. Acknowledgement

The project belongs to OpenCDA ecosystem [48] and is funded in part by the Federal Highway Administration project and California RIMI Program. Special thanks go to Transportation Research Center Inc for their collaboration in experimental data collection and processing.

References

- [1] Philipp Bender, Julius Ziegler, and Christoph Stiller. Lanelets: Efficient map representation for autonomous driving. In *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pages 420–425. IEEE, 2014. 5
- [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 multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2, 7
- [3] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019. 2
- [4] Qi Chen, Xu Ma, Sihai Tang, Jingda Guo, Qing Yang, and Song Fu. 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*, pages 88–100, 2019. 3, 6, 7, 8
- [5] Qi Chen, Sihai Tang, Qing Yang, and Song Fu. Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 514–524. IEEE, 2019. 3
- [6] Yuhua Chen, Wen Li, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Domain adaptive faster r-cnn for object detection in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3339–3348, 2018. 7
- [7] Charles K Chui, Guanrong Chen, et al. *Kalman filtering*. Springer, 2017. 5
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 2
- [9] Valter Crescenzi, Giansalvatore Mecca, Paolo Merialdo, et al. Roadrunner: Towards automatic data extraction from large web sites. In *VLDB*, volume 1, pages 109–118, 2001. 5
- [10] Jiaxun Cui, Hang Qiu, Dian Chen, Peter Stone, and Yuke Zhu. Coopernaut: end-to-end driving with cooperative perception for networked vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17252–17262, 2022. 3
- [11] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 2, 3
- [12] Lue Fan, Ziqi Pang, Tianyuan Zhang, Yu-Xiong Wang, Hang Zhao, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Embracing single stride 3d object detector with sparse transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8458–8468, 2022. 3
- [13] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 7
- [14] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 2
- [15] Yushan Han, Hui Zhang, Huifang Li, Yi Jin, Congyan Lang, and Yidong Li. Collaborative perception in autonomous driving: Methods, datasets and challenges. *arXiv preprint arXiv:2301.06262*, 2023. 1
- [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. 3
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [18] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. 3, 7
- [19] E Li, Shuaijun Wang, Chengyang Li, Dachuan Li, Xiangbin Wu, and Qi Hao. Sustech points: A portable 3d point cloud interactive annotation platform system. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 1108–1115. IEEE, 2020. 4
- [20] Yiming Li, Ziyang An, Zixun Wang, Yiqi Zhong, Siheng Chen, and Chen Feng. V2x-sim: A virtual collaborative perception dataset for autonomous driving. *arXiv preprint arXiv:2202.08449*, 2022. 2, 3
- [21] Yingwei Li, Adams Wei Yu, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Yifeng Lu, Denny Zhou, Quoc V Le, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17182–17191, 2022. 3
- [22] Yiming Li, Zhiding Yu, Christopher Choy, Chaowei Xiao, Jose M Alvarez, Sanja Fidler, Chen Feng, and Anima Anandkumar. Voxformer: Sparse voxel transformer for camera-based 3d semantic scene completion. *arXiv preprint arXiv:2302.12251*, 2023. 1
- [23] Wei Liu, Karoll Quijano, and Melba M Crawford. Yolov5-tassel: detecting tassels in rgb uav imagery with improved yolov5 based on transfer learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:8085–8094, 2022. 3
- [24] Wei Liu, Xin Xia, Lu Xiong, Yishi Lu, Letian Gao, and Zhuoping Yu. Automated vehicle sideslip angle estimation considering signal measurement characteristic. *IEEE Sensors Journal*, 21(19):21675–21687, 2021. 6
- [25] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 7

- [26] Yifan Lu, Quanhao Li, Baoan Liu, Mehrdad Dianati, Chen Feng, Siheng Chen, and Yanfeng Wang. Robust collaborative 3d object detection in presence of pose errors. *arXiv preprint arXiv:2211.07214*, 2022. 3
- [27] Sivabalan Manivasagam, Shenlong Wang, Kelvin Wong, Wenyuan Zeng, Mikita Sazanovich, Shuhan Tan, Bin Yang, Wei-Chiu Ma, and Raquel Urtasun. Lidarsim: Realistic lidar simulation by leveraging the real world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11167–11176, 2020. 2
- [28] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7077–7087, 2021. 3
- [29] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018. 3
- [30] Hang Qiu, Pohan Huang, Nam Asavisanu, Xiaochen Liu, Konstantinos Psounis, and Ramesh Govindan. Autocast: Scalable infrastructure-less cooperative perception for distributed collaborative driving. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '22, 2022. 3
- [31] Zaydoun Yahya Rawashdeh and Zheng Wang. 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)*, pages 3961–3966. IEEE, 2018. 3, 6
- [32] Cody Reading, Ali Harakeh, Julia Chae, and Steven L Waslander. Categorical depth distribution network for monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8555–8564, 2021. 3
- [33] Thomas Roddick, Alex Kendall, and Roberto Cipolla. Orthographic feature transform for monocular 3d object detection. *arXiv preprint arXiv:1811.08188*, 2018. 3
- [34] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*, pages 3234–3243, 2016. 2
- [35] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2397–2406, 2022. 1, 3
- [36] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 2
- [37] Simon Suo, Sebastian Regalado, Sergio Casas, and Raquel Urtasun. Trafficsim: Learning to simulate realistic multi-agent behaviors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10400–10409, 2021. 2
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [39] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4604–4612, 2020. 3
- [40] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *European Conference on Computer Vision*, pages 605–621. Springer, 2020. 1, 3, 6, 7, 8
- [41] Yue Wang, Alireza Fathi, Abhijit Kundu, David A Ross, Caroline Pantofaru, Tom Funkhouser, and Justin Solomon. Pillar-based object detection for autonomous driving. In *European Conference on Computer Vision*, pages 18–34. Springer, 2020. 3
- [42] 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*, pages 180–191. PMLR, 2022. 1, 3
- [43] Xinshuo Weng, Jianren Wang, David Held, and Kris Kitani. 3d multi-object tracking: A baseline and new evaluation metrics. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10359–10366. IEEE, 2020. 1, 7
- [44] Hao Xiang, Runsheng Xu, Xin Xia, Zhaoliang Zheng, Bolei Zhou, and Jiaqi Ma. V2xp-asg: Generating adversarial scenes for vehicle-to-everything perception. *arXiv preprint arXiv:2209.13679*, 2022. 7
- [45] Runsheng Xu, Weizhe Chen, Hao Xiang, Lantao Liu, and Jiaqi Ma. Model-agnostic multi-agent perception framework. *arXiv e-prints*, pages arXiv–2203, 2022. 3
- [46] Runsheng Xu, Jinlong Li, Xiaoyu Dong, Hongkai Yu, and Jiaqi Ma. Bridging the domain gap for multi-agent perception. *arXiv preprint arXiv:2210.08451*, 2022. 3
- [47] Runsheng Xu, Zhengzhong Tu, Hao Xiang, Wei Shao, Bolei Zhou, and Jiaqi Ma. Cobevt: Cooperative bird’s eye view semantic segmentation with sparse transformers. *arXiv preprint arXiv:2207.02202*, 2022. 1, 3, 6, 7, 8
- [48] Runsheng Xu, Hao Xiang, Xu Han, Xin Xia, Zonglin Meng, Chia-Ju Chen, Camila Correa-Jullian, and Jiaqi Ma. The openda open-source ecosystem for cooperative driving automation research. *IEEE Transactions on Intelligent Vehicles*, pages 1–13, 2023. 8
- [49] Runsheng Xu, Hao Xiang, Zhengzhong Tu, Xin Xia, Ming-Hsuan Yang, and Jiaqi Ma. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIX*, pages 107–124. Springer, 2022. 1, 2, 3, 6, 7, 8
- [50] Runsheng Xu, Hao Xiang, Xin Xia, Xu Han, Jinlong Li, and Jiaqi Ma. Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communica-

- tion. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2583–2589. IEEE, 2022. 2, 3, 6, 7, 8
- [51] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 3
- [52] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018. 3
- [53] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. 2
- [54] Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, 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*, pages 21361–21370, 2022. 2, 6
- [55] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021. 3
- [56] Zelin Zhao, Ze Wu, Yueqing Zhuang, Boxun Li, and Jiaya Jia. Tracking objects as pixel-wise distributions. In *European Conference on Computer Vision*, pages 76–94. Springer, 2022. 1
- [57] Brady Zhou and Philipp Krähenbühl. Cross-view transformers for real-time map-view semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13760–13769, 2022. 1
- [58] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018. 3