V2X-Sim: Multi-Agent Collaborative Perception Dataset and Benchmark for Autonomous Driving

Yiming Li[®], Student Member, IEEE, Dekun Ma, Ziyan An, Zixun Wang, Student Member, IEEE, Yiqi Zhong, Siheng Chen[®], and Chen Feng[®], Member, IEEE

Abstract—Vehicle-to-everything (V2X) communication techniques enable the collaboration between vehicles and many other entities in the neighboring environment, which could fundamentally improve the perception system for autonomous driving. However, the lack of a public dataset significantly restricts the research progress of collaborative perception. To fill this gap, we present V2X-Sim, a comprehensive simulated multi-agent perception dataset for V2X-aided autonomous driving. V2X-Sim provides: (1) multi-agent sensor recordings from the road-side unit (RSU) and multiple vehicles that enable collaborative perception, (2) multi-modality sensor streams that facilitate multi-modality perception, and (3) diverse ground truths that support various perception tasks. Meanwhile, we build an open-source testbed and provide a benchmark for the state-of-the-art collaborative perception algorithms on three tasks, including detection, tracking and segmentation. V2X-Sim seeks to stimulate collaborative perception research for autonomous driving before realistic datasets become widely available.

Index Terms—Deep learning for visual perception, multi-robot systems, data sets for robotic vision.

I. INTRODUCTION

PERCEPTION is a fundamental capability for autonomous vehicles, which allows them to represent, identify, and interpret sensory input for understanding the complex surroundings. In literature, single-vehicle perception has been intensively studied thanks to the well-established driving datasets [1]–[3], and researchers have proposed various algorithms to deal with different downstream tasks [4]–[6].

Despite recent advances in single-vehicle perception, the individual viewpoint often results in degraded perception in

Manuscript received 24 February 2022; accepted 30 June 2022. Date of publication 21 July 2022; date of current version 23 August 2022. This letter was recommended for publication by Associate Editor I. Gilitschenski and Editor C. C. Lerma upon evaluation of the reviewers' comments. This work was supported in part by the NSF CPS Program under Grant CMMI-1932187 and CNS-2121391, in part by the National Natural Science Foundation of China under Grant 62171276, in part by the Science and Technology Commission of Shanghai Municipal under Grant 21511100900, and in part by CALT under Grant 2021-01. (Corresponding authors: Siheng Chen; Chen Feng.)

Yiming Li, Dekun Ma, Ziyan An, Zixun Wang, and Chen Feng are with the New York University, Brooklyn, NY 11201 USA (e-mail: yimingli9702 @gmail.com; dm4524@nyu.edu; annieziyan1222@gmail.com; craddywang@gmail.com; cfeng@nyu.edu).

Yiqi Zhong is with the University of Southern California, Los Angeles 91103 USA (e-mail: yiqizhon@usc.edu).

Siheng Chen is with the Cooperative Medianet Innovation Center, Shanghai Jiao Tong University and Shanghai AI Laboratory, Shanghai 200240, China (e-mail: sihengc@sjtu.edu.cn).

Digital Object Identifier 10.1109/LRA.2022.3192802

Our dataset and code are available at https://ai4ce.github.io/V2X-Sim/

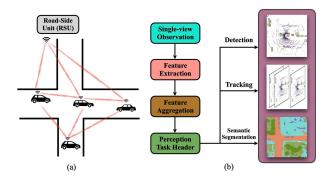


Fig. 1. (a) Intersection for vehicle-to-everything (V2X) communication. (b) Workflow of multi-agent collaborative perception with intermediate-/feature-based strategy. We benchmark collaborative object detection, multi-object tracking, and semantic segmentation in the bird's eye view (BEV).

long-range or occluded areas. A promising solution to this problem is through vehicle-to-everything (V2X) [7], a cutting-edge communication technology that enables dialogue between a vehicle and other entities, including vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I). With the aid of V2X communication, we are able to upgrade single-vehicle perception to collaborative perception, which introduces more viewpoints to help autonomous vehicles see further, better and even see through occlusion, thereby fundamentally enhancing the capability of perception.

Collaborative perception naturally draws on communication and perception. Its development requires expertise from both communities. Recently, the communication community has made enormous efforts to promote such a study [8]–[10]; however, only a few works have been proposed from the perspective of perception [11]–[15]. One major reason for this is the lack of well-designed and organized collaborative perception datasets. Due to the immaturity of V2X and the cost of simultaneously operating multiple autonomous vehicles, it is very expensive and laborious to build such a real dataset for research communities. Therefore, we synthesize a comprehensive and publicly available dataset, named as V2X-Sim, to advance the study of collaborative perception for V2X-communication-aided autonomous driving.

To generate V2X-Sim, we employ SUMO [16], a micro-traffic simulation, to produce numerically-realistic traffic flow, and CARLA [17], a widely-used open-source simulator for autonomous driving research, to retrieve well-synchronized sensor streams from multiple vehicles as well as the road-side unit (RSU). Meanwhile, multi-modality sensor streams of different entities are recorded to enable cross-modality perception. In addition, diverse annotations including bounding boxes, vehicle

2377-3766 © 2022 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

TABLE I

COMPARISON OF COLLABORATIVE PERCEPTION DATASETS FOR AUTONOMOUS DRIVING. THERE ARE NO PUBLIC DATASETS WHICH SUPPORT BOTH V2V AND V2I RESEARCH: MULTI-AGENT DATA ARE EITHER GENERATED BY SIMULATORS [12], [20] OR CREATED BY SELECTING CONSECUTIVE FRAMES FROM SINGLE-AGENT REAL DATASETS [11], [21], [22]. SEVERAL WORKS COLLECT DATA FROM MULTIPLE INFRASTRUCTURE SENSORS: [23] IN SIMULATION, [24], [25] IN REAL WORLD. OUR DATASET IS THE FIRST PUBLIC MULTI-AGENT MULTI-MODALITY DATASET WHICH SUPPORTS DIFFERENT COLLABORATIVE PERCEPTION TASKS.

Dataset	Scenario	Source		Sensor RGB Depth LiDAR IMU/GPS				Tasks S Det. Track. Seg.		
V2V-Sim [12]	V2V	LiDARsim [26]	x	х	/	х	1	/	х	Х
Cooper [11]	V2V	KITTI [2]	X	Х	/	1	1	1	Х	Х
MFSL [21]	V2V	KITTI [2]	1	Х	X	1	×	Х	/	Х
VANETs [22]	V2V	KITTI [2]	1	Х	/	1	1	Х	Х	Х
OPV2V [20]	V2V	OpenCDA [27] & CARLA [17]	1	Х	/	1	1	/	Х	1
CODD [28]	V2V	CARLA [17]	×	Х	/	X	1	/	Х	/
Cooper (inf) [23]	V2I	CARLA [17]	/	/	Х	X	1	/	Х	1
WIBAM [24]	V2I	Real-world	1	Х	Х	Х	1	/	х	1
DAIR-V2X [25]	V2I	Real-world	1	×	1	1	1	Х	Х	1
V2X-Sim (Ours)	V2V&V2I	CARLA [17] & SUMO [16]	/	/	1	1	1	1	1	1

trajectories, and semantic labels are provided to facilitate various downstream tasks. To better serve multi-agent, multi-modality, and multi-task perception research for autonomous driving, we further provide a benchmark for three crucial perception tasks (collaborative detection, tracking, and segmentation) on the proposed dataset using the state-of-the-art collaboration strategies [12], [13], [18], [19]. In summary, our contributions are two-fold:

- We propose V2X-Sim, a comprehensive V2X perception dataset for autonomous driving, to support multi-agent multi-modality multi-task perception research.
- We create an open-source testbed for collaborative perception methods, and provide a benchmark on three tasks to encourage more research in this field.

II. RELATED WORK

Autonomous driving dataset: Since the pioneering dataset KITTI [2] was released, the autonomous driving community has been trying to increase the dataset comprehensiveness in terms of driving scenarios, sensor modalities, and data annotations. Regarding driving scenarios, current datasets cover crowded urban scenes [29], adverse weather conditions [30], night scenes [31], and multiple cities [1] to enrich the data distribution. As for sensor modalities, nuScenes [1] collects data with Radar, RGB cameras, and LiDAR in a 360° viewpoint; WoodScape [32] captures data with fisheye cameras. Regarding data annotations, semantic labels in both images [33]–[36] and point clouds [37], [38] are provided to enable semantic segmentation; 2D/3D box trajectories are offered [39], [40] to facilitate tracking and prediction. In summary, most real datasets emphasize the data comprehensiveness in single-vehicle situations, but ignore the multi-vehicle scenarios.

V2X system and dataset: By sharing information with other vehicles or the RSU, V2X mitigates the shortcomings of single-vehicle perception and planning such as the limited sensing range and frequent occlusion [7]. Previous research [41] has developed an enhanced cooperative microscopic traffic model in V2X scenarios, and studied the effect of V2X in traffic disturbance scenarios. [42] proposes a multi-modal cooperative perception system that provides see-through, lifted-seat, satellite and all-around views to drivers. More recently, COOPER [11] investigates raw-data level collaborative perception to improve the detection capability for autonomous driving. V2VNet [12]

proposes intermediate-feature level collaboration to promote the vehicle's perception and prediction capability. Several works use multiple infrastructure sensors to jointly perceive the environment and employ output-level collaboration with vehicleto-infrastructure communication [23], [24]. As for the datasets, [11], [21], [22] simulate the V2V scenarios with different frames from KITTI [2]. Yet, these datasets are unrealistic multi-agent datasets for the measurements are not captured by different viewpoints. Some other works use a platoon strategy for data capture [43], [44], but they are biased because the observations were highly correlated with each other. The most relevant work is V2V-Sim [12], which is based on a high-quality LiDAR simulator [26]. Unfortunately, V2V-Sim does not include the V2I scenario and is not publicly available. Moreover, OPV2V [20] and CODD [28] only support the detection task in the V2V scenario. Existing collaborative perception datasets are summarized in Table I: V2X-Sim¹ is currently the most comprehensive one with multi-agent multi-modality sensory streams in both V2V and V2I scenarios, and can support various downstream tasks such as multi-agent collaborative detection, tracking, and semantic segmentation.

III. V2X-SIM DATASET

V2X-Sim could enable more research on the collaboration strategy among vehicles to achieve a more robust perception. This could fundamentally benefit autonomous driving, intelligent transportation systems, smart cities, *etc.*.

A. Sensor Suite on Vehicles and RSU

Multi-modality data is essential for robust perception. To ensure the comprehensiveness of our dataset, we equip each vehicle with a sensor suite composed of RGB cameras, LiDAR, GPS, IMU, and RSU with RGB cameras and LiDAR.

Sensor configuration: On both the vehicle and RSU, the camera and LiDAR cover 360° horizontally to enable full-view perception. Specifically, each vehicle carries six RGB cameras following the nuScenes configuration [1]; the RSU is equipped with four RGB cameras toward four directions at the crossroad. We employ depth camera, semantic segmentation camera, semantic LiDAR, and BEV semantic segmentation camera in CARLA to obtain the corresponding ground-truth for each RGB camera. Note that the BEV semantic segmentation camera is based on orthogonal projection while the ego-vehicle semantic segmentation camera uses perspective projection. Table II summarizes the detailed sensor specification.

Sensor layout and coordinate system: The overall sensor layout and coordinate system is shown in Fig. 2. The BEV semantic camera shares the same x, y position with LiDAR yet is placed higher to ensure a certain size of field of view. Note that we invert the y-axis in CARLA and use a right-hand coordinate system following nuScenes [1].

Diverse annotations: To assist downstream tasks including detection, tracking and semantic segmentation, we provide various annotations such as 3D bounding boxes, pixel-wise and point-wise semantic labels. Each box is defined by the location of its center in x, y, z coordinates, and its width, length, and height. In total, there are 23 categories such as the pedestrian, building, ground, etc. In addition, precise depth values are provided for depth estimation.

¹This work extends the LiDAR-based V2V data in our previous work [13] with more modalities, scenarios and downstream tasks.

 $\label{eq:table ii} \textbf{TABLE II} \\ \textbf{Sensor Suite of Vehicle (V) and Intersection (I)}.$

Sensor	Description
V: 6 × RGB camera I: 4 × RGB camera	Each vehicle is equipped with 6 cameras. Each camera has a FoV of 70°, except for the back camera that has a FoV of 110°. Each RSU has 4 cameras looking diagonally downward at 35° with a 70° FoV. The image size is 1600×900.
V: 6 × Depth cam- era I: 4 × Depth camera	Each vehicle has 6 depth cameras with the same setting as RGB cameras. Each RSU has 4 depth cameras with the same setting as its RGB cameras.
V: 6 × Semantic camera I: 4 × Se- mantic camera	
V&I: 1 × BEV se- mantic camera	Each vehicle and RSU has one BEV semantic camera at the top, looking downward. Both the raw images (semantic tags encoded in the red channel) and the converted colored images are provided. The image size is 900×900.
V&I: 1 × LiDAR and Semantic Li- DAR	We attach one LiDAR and one semantic LiDAR on top of the ego vehicle and the intersection center. Specs: 32 channels, 70m max range, 250,000 points per second, 20 Hz rotation frequency.

B. CARLA-SUMO Co-Simulation

We consider it a realistic V2X scenario when multiple vehicles with their own routes are simultaneously located in the same intersection. Each intersection is also equipped with one RSU with sensing capability. Regarding the traffic simulation, there are several non-public simulators which can explicitly generate data tailored for collaborative perception such as scenes with occlusion, and sensor range limitations [45], [46]. Yet in this work, we use open-source CARLA-SUMO co-simulation for traffic flow simulation and data recording. Vehicles are spawned in CARLA via SUMO to roam around in the town with random routes. Hundreds of vehicles are spawned in different towns (Town03, Town04 and Town05 that have crossroads as well as)T-junctions in both the crowded downtown and suburb highway). We record several log files in different towns. Afterwards, at different junctions, we read out 100 scenes from the log files. Each scene has a 20-second duration, and we choose M(M =2, 3, 4, 5) vehicles as well as one RSU in each scene as intelligent agents to share information. Example scenarios are shown in Fig. 3(a).

C. Downstream Tasks

Our dataset not only supports single-agent perception tasks such as 3D object detection, tracking, image-/point-cloud-based semantic segmentation, depth estimation, but also enables collaborative perception such as collaborative 3D object detection, tracking, and collaborative BEV semantic segmentation in urban driving scenes. We provide a benchmark for collaborative perception algorithms.

D. Dataset Statistics

Annotation statistics: We provide statistics on the annotations and objects to highlight the inclusiveness and diversity of our dataset. In Fig. 3(b) we analyze the size of the cars' bounding boxes within a 70 m range from ego vehicles in each frame. The great variation of car sizes indicates that our scenes contain a diverse set of car makes and models that well includes most of the common real-world vehicles. Fig. 3(c) shows the annotation count in each frame for vehicles within 0-30 m, 30-50 m, and 50-70 m ranges from each ego vehicle. It suggests that our dataset features both crowded scenes (up to 100 annotations within 50-70 m from the ego vehicle) and less crowded driving scenarios (as low as 10 annotations within 30 m from the ego vehicle). Fig. 3(d) contains statistics on the number of LiDAR points per annotation for single-agent and multi-agent scenarios

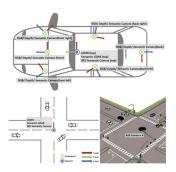


Fig. 2. Sensor layout and coordinate systems.

respectively. The number of total LiDAR points of each object annotation increases when there are more than one agents observing the same object. Specifically, for a single agent, there are 183.83 points in each annotation on average, but the number goes up to 875.59 points per annotation for multiple agents.

Scene features: We analyze the distance between each two ego vehicles for every frame, as shown in Fig. 3(e). An overwhelming percentage of the ego vehicles are presented within 20-30 meters from each other, suggesting they are geographically closely connected. The speed of cars within 70 m from ego vehicles are shown in Fig. 3(f). Given the fact that our scenes are selected near intersections, we notice that a major fraction of vehicles are slower than 10 km/h. However, the maximum speed is as high as 90+ km/h. Fig. 3(g) shows the percentage of annotations observed by a certain number of ego vehicles, up to 5. Over 60% of the annotations are observed by at least two ego vehicles.

IV. COLLABORATIVE PERCEPTION BENCHMARK

We benchmark three crucial perception tasks in autonomous driving within the collaboration setting: detection, tracking, and semantic segmentation. The three tasks have been extensively studied since they generate essential perception knowledge for autonomous vehicles to make safer decisions. For performance evaluation, we follow the same evaluation protocol of the three tasks in the single-agent scenario, except that we utilize the information shared by other agents while the single-agent perception do not have access to such information. We report the results in two scenarios: (1) V2V only, and (2) V2X (V2V + V2I).

Dataset format and split: Our V2X-Sim dataset follows the same storage format of nuScenes [1], an authoritative multimodality single-agent autonomous driving dataset. nuScenes collected real-world single-agent driving data to promote the single-agent autonomous driving research; while we simulate multi-agent scenarios to facilitate the next-generation V2Xaided collaborative perception technology. Each scene in our dataset contains a 20 s traffic flow at a certain intersection of three CARLA towns, and the multi-modality multi-agent sensory streams are recorded at 5 Hz, meaning each scene is composed of 100 frames. We generate 100 scenes with a total of 10,000 frames, and in each scene, 2-5 vehicles are selected as the collaboration agents. We use 8,000/1,000/1,000 frames for training/validation/testing respectively, and we ensure that there is no overlap in terms of the intersections across training/validation/testing set. Each frame has data sampled from multiple agents (vehicles and RSU). There are 37,200 samples in the training set, 5,000 samples in the validation set, and 5,000 samples in the test set. The split is shared across tasks.

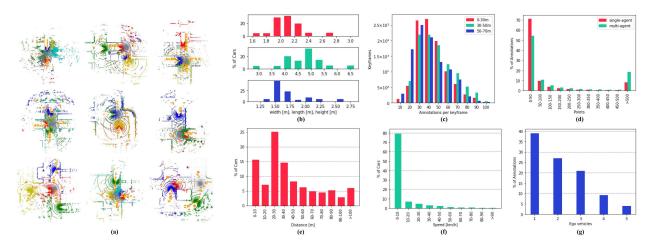


Fig. 3. (a) Visualizations of the bird's eye view point cloud from different scenes. Gray denotes the point cloud captured by the RSU LiDAR. Each color (except for gray) represents a vehicle, and the orange boxes denote the vehicles in the scene. (b) Statistics for car bounding box sizes. (c) Counts for annotations per keyframe where the annotated vehicles are presented within 0-30 m, 30-50 m, and 50-70 m of the ego vehicles. (d) Counts for LiDAR points per annotation. (e) Statistics of the distance between every two ego vehicles for all frames. (f) Speed of cars located within 70 m from ego vehicles. (g) Percentage of annotated vehicles observed by 1-5 ego vehicles.

Implementation details: Bird's-eye-view (BEV) is a widely used and powerful representation in autonomous driving because it can describe the surrounding objects and overall context via a compact top-down 2D map [47]. Therefore, BEV-based representation is adopted in all three tasks: we use a 3D voxel grid to represent the 3D world, employ binary representation, and assign each voxel a positive label if the voxel contains point cloud data. Since the generated 3D voxel grid can be considered as a pseudo-image whose height dimension is the channel dimension, we can perform the efficient 2D convolution instead of the heavy 3D convolution. Specifically, we crop the points located in the region of $[-32, 32] \times [-32, 32] \times [-3, 2]$ meters for vehicles defined in the ego-vehicle Euclidean coordinate and $[-32, 32] \times [-32, 32] \times [-8, -3]$ meters for RSU. The width and length of each voxel are 0.25 m, and the height is 0.4 m, meaning the generated BEV-based pseudo-image has a dimension of $256 \times 256 \times 13$ ($W \times L \times H$). Note that the models in all the tasks consume the 3D BEV map and generate perception results in 2D BEV.

Benchmark models: We aim to benchmark collaborative perception strategies rather than the well-studied individual perception methods. We consider early/intermediate/late/no collaboration models for the benchmark. The intermediate models, including DiscoNet [13], V2VNet [12], When2com [18], and Who2com [19], are based on the communication of the intermediate features of the neural network. The methods in our benchmark are as follows:

- Lower-bound: The single-agent perception model without collaboration which processes a single-view point cloud is considered as the lower-bound.
- *Co-lower-bound:* Collaborative lower-bound fuses the output from different single-agent perception models.
- *Upper-bound:* The early collaboration model which transmits raw point cloud data is the upper-bound.
- *DiscoNet* [13]: DiscoNet uses a directed collaboration graph with matrix-valued edge weight to adaptively highlight the informative spatial regions and reject the noisy regions of the messages sent by the partners. After adaptive message fusion, the updated features will be transmitted to the output head for perception.

- *V2VNet* [12]: V2VNet uses a pose-aware graph neural network to propagate agents' information, and employs a convolutional gated recurrent unit based module to aggregate other agents' information. After several rounds of neural message passing, the updated features are fed into the output head to generate perception results.
- When2com [18]: When2com employs attention-based mechanism for communication group construction: the partners with satisfactory correlation scores will be selected as the collaborators. After the attention-score-based weighted fusion, the updated features will be fed into the output head for perceptions. The model with pose information is marked by *.
- Who2com [19]: Who2com shares a similar idea with When2com, yet it uses handshake mechanism to select the collaborator: the partner with the highest score will be selected as the collaborator. The model with pose information is marked by *.

We implement a 3D perception pipeline that can be integrated with all of the communication methods mentioned above. Since the source codes of V2VNet is not publicly available, we reimplement the V2VNet in PyTorch according to its pseudocode. For when2com/who2com, we borrow their communication modules from its official code. All of the intermediate collaboration modules use the same architecture and conduct collaboration at the same intermediate feature layer. Moreover, all of the methods are trained with the same setting to ensure that the performance gain comes from the collaboration instead of irrelevant techniques.

A. Collaborative Object Detection in BEV

Problem definition: As the most crucial perception task in autonomous driving, 3D object detection aims to recognize and localize the objects in 3D scenes given a single frame, with the following tracking, prediction and planning modules all heavily relying on the detections. The models consume voxelized point cloud and output BEV bounding boxes.

Backbone and evaluation: We use a classic anchor-based detector composed of a convolutional encoder, a convolutional

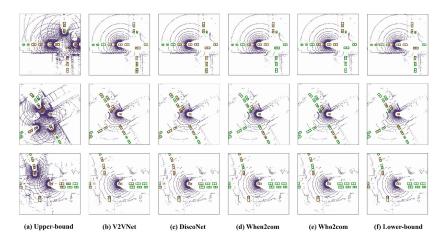


Fig. 4. Visualizations of BEV detection on V2X-Sim. Red and green boxes are the predictions and ground-truths respectively.

TABLE III
QUANTITATIVE RESULTS OF COLLABORATIVE BEV DETECTION ON V2X-SIM

Method	APo w/o RSU	@ IoU=0.5 w/ RSU Δ	APo w/o RSU	@ IoU=0.7 Jw/ RSU Δ
Lower-bound	49.90	46.96 ↓ 2.94	44.21	42.33 ↓ 1.88
Co-lower-bound	43.99	42.98 \ \ 1.01	39.10	38.26 ↓ 0.84
When2com [18]	44.02	46.39 ↑ 2.37	39.89	40.32 ↑ 0.43
When2com* [18]	45.35	48.28 ↑ 2.93	3 40.45	41.43 ↑ 0.68
Who2com [19]	44.02	46.39 ↑ 2.37	39.89	40.32 ↑ 0.43
Who2com* [19]	45.35	48.28 ↑ 2.93	3 40.45	41.13 ↑ 0.68
V2VNet [12]	68.35	72.08 ↑ 3.73	62.83	65.85 ↑ 3.02
DiscoNet [13]	69.03	72.87 ↑ 3.84	63.44	66.40 ↑ 4.81
Upper-bound	70.43	77.08 ↑ 6.65	67.04	72.57 \(\phi\) 5.53

RSU Denotes the Road-Side Unit. AP Denotes the Average Precision. Δ is the Absolute Gain in AP Introduced by RSU

decoder, and an output header for classification and regression [48]. Regarding the loss function, we use the binary crossentropy loss to supervise the box classification and the smooth L_1 loss to supervise the box regression, following [48]. We employ the generic BEV detection evaluation metric: Average Precision (AP) at Intersection-over-Union (IoU) thresholds of 0.5 and 0.7. We target the vehicle detection and report the results on the test set

Quantitative results: Table III demonstrates the quantitative comparisons on AP (@IoU = 0.5/0.7). We find that: (1) the upper-bound performs best amongst all methods, and it improves lower-bound significantly by 41.1% and 51.6% in terms of AP@0.5 and AP@0.7 in the scenario of V2V only, validating the necessity of collaboration; (2) V2V and V2I jointly can generally improve the perception over V2V only with more viewpoints, e.g., adding V2I can bring an improvement of 9.4% for upper-bound, and 5.5% for V2VNet in terms of AP@0.5; (3) among the intermediate models, DiscoNet achieves the best performance via the well-designed distilled collaboration graph, V2VNet achieves the second best performance via the powerful neural message passing, and When2com/Who2com only achieve comparable performance with lower-bound since the attention-mechanism is not suitable in point-cloud-based collaborative perception: the agents usually need complementary information rather than a highly similar one; (4) the late collaboration model (co-lower-bound) hurts the detection

performance because of introducing extra false positives from other vehicles.

Qualitative results: The qualitative results are shown in Fig. 4. We see that the collaboration can fundamentally mitigate the problems of long-range perception and occlusion.

B. Collaborative Multi-Object Tracking in BEV

Problem definition: Different from detection, multi-object tracking requires the generation of temporally consistent perception results. Multi-object tracking in BEV is to use bounding boxes, object categories, and object identities to track different objects within a temporal sequence.

Evaluation metrics: We mainly utilize HOTA (Higher Order Tracking Accuracy) [49] to evaluate our BEV tracking performance. HOTA can evaluate detection, association, and localization performance via a single unified metric. In addition, the classic multi-object tracking accuracy (MOTA) and multi-object tracking precision (MOTP) [50] are also employed. MOTA can measure detection errors and association errors. MOTP solely measures localization accuracy.

Baseline tracker: We implement SORT [51] as our baseline tracker. Given the detection results, SORT will combine the Kalman Filter and Hungarian algorithm to achieve an accurate and efficient tracking performance.

Quantitative results: Quantitative comparisons of BEV tracking are shown in Table IV. Similar to BEV detection, upper-bound achieves the best performance in terms of MOTA and HOTA. Meanwhile, adding V2I can improve MOTA largely yet cannot help too much in MOTP. Co-lower-bound shows good performance in localization accuracy (MOTP). A more advanced tracker is required to exploit the collaboration for filling the performance gap.

C. Collaborative Semantic Segmentation in BEV

Problem definition: We aim to conduct semantic segmentation in BEV using only geometry point cloud. In the collaborative perception scenarios, there are measurements collected by multiple agents with distinct viewpoints. Therefore, there are more information in the scene, facilitating the semantic scene understanding.

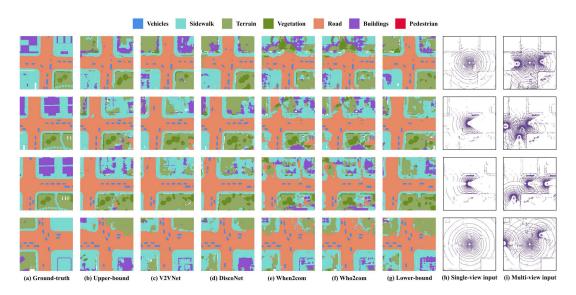


Fig. 5. Visualizations of collaborative BEV semantic segmentation.

TABLE IV

QUANTITATIVE RESULTS OF BEV TRACKING ON V2X-SIM. MOTA: MULTIPLE OBJECT TRACKING ACCURACY. MOTP: MULTIPLE OBJECT TRACKING PRECISION. HOTA: HIGHER ORDER TRACKING ACCURACY. DETA: DETECTION ACCURACY. ASSA: ASSOCIATION ACCURACY. DETRE: DETECTION RECALL. DETPR: DETECTION PRECISION. ASSRE: ASSOCIATION RECALL. ASSPR: ASSOCIATION PRECISION. Loca: Localization Accuracy. The Number to the Left of () DENOTES THE PERFORMANCE IN V2V SOLELY. THE NUMBER IN () REPRESENTS THE PERFORMANCE GAIN BY ADDING V2I

Method	MOTA	MOTP	HOTA	DetA	AssA	DetRe	DetPr	AssRe	AssPr	LocA
Lower-bound	35.72 (\\$3.87)	84.16 (\\$0.74)	34.27 (\1.68)	33.64 (\\$3.24)	36.18 (\\$0.06)	35.07 (\\$3.54)	82.49 (†0.96)	46.70 (†0.23)	58.72 (†0.10)	86.43 (†0.38)
Co-lower-bound	21.53 (†0.58)	85.76 (†0.15)	39.16 (\\$0.71)	41.14 (\\$0.93)	38.18 (\(\psi 0.62\))	59.54 (\$\dagger\$2.52)	54.68 (†0.79)	50.92 (\$\dagger\$0.65)	55.78 (†0.84)	87.64 (†0.38)
When2com [18]	29.48 (†2.45)	86.10 (\$\dagge 2.79)	30.94 (†1.01)	27.90 (†2.04)	35.33 (†0.06)	28.67 (†2.58)	86.11 (\14.81)	46.30 (\(\psi 0.15 \)	59.20 (\\$\\$0.36)	87.98 (\$1.98)
When2com* [18]	30.17 (†1.43)	84.95 (\1.44)	31.34 (\(\frac{1}{2}\)0.43)	29.11 (†1.05)	35.42 (\(\phi\)0.21)	30.28 (†1.32)	83.81 (\(\phi\)0.29)	46.65 (\\$\tau 0.29)	58.61 (†0.18)	86.14 (†0.17)
Who2com [19]	29.48 (†2.46)	86.10 (\(\psi 2.79\)	30.94 (†1.01)	27.90 (†2.04)	35.33 (\(\psi\)0.06)	28.67 (†2.58)	86.11 (\14.81)	46.30 (\\$\d\ 0.15)	59.20 (\\$\\$0.36)	87.98 (\$1.98)
Who2com* [19]	30.17 (†1.43)	84.95 (\1.44)	31.34 (†0.43)	29.11 (†1.06)	35.42 (\(\dagger)0.21\)	30.28 (†1.33)	83.81 (†0.29)	46.65 (\\$\d\ 0.29)	58.61 (†0.18)	86.14 (†0.17)
V2VNet [12]	55.29 (†2.29)	85.21 (\\$\tau 0.53)	43.68 (†0.91)	50.71 (†1.93)	38.76 (†0.24)	53.40 (†2.51)	84.45 (\1.07)	50.22 (†0.53)	58.50 (\\$\tau0.07)	87.22 (†0.38)
DiscoNet [13]	56.69 (†2.26)	86.23 (\$\dagge 0.41)	44.76 (†1.09)	52.41 (†2.18)	39.25 (†1.11)	54.87 (†2.58)	86.29 (\$\dagger\$0.95)	50.86 (†1.02)	58.94 (\$\dagge 0.15)	88.07 (†0.34)
Upper-bound	58.00 (†3.92)	85.61 (†0.25)	44.83 (†4.24)	52.94 (†4.24)	38.95 (\$\dagger\$0.75)	55.07 (†4.68)	86.54 (\$\dagger\$0.30)	50.35 (\$\\$0.86)	58.71 (†0.15)	87.48 (†0.06)

TABLE V QUANTITATIVE RESULTS OF BEV SEGMENTATION ON V2X-SIM. THE NUMBER TO THE LEFT OF () DENOTES THE PERFORMANCE IN V2V SOLELY

Method	Vehicle	Sidewalk	Terrain	Road	Building	Pedestrian	Vegetation	mIoU
Lower-bound	45.93 (†2.22)	42.39 (\12.75)	47.03 (†0.20)	65.76 (\1.27)	25.38 (\1.89)	20.59 (\$\dagger3.09)	35.83 (\(\gamma 0.66\))	36.64 (\\$\tau 0.87)
Co-lower-bound	47.67 (†2.43)	48.79 (\$1.41)	50.92 (†0.85)	70.00 (\$\\$0.65)	25.26 (†0.17)	10.78 (\1.77)	39.46 (†2.69)	38.38 (†0.46)
When2com [18]	48.43 (\(\phi0.03\))	33.06 (†1.38)	36.89 (†1.76)	57.74 (†1.56)	29.20 (†1.18)	20.37 (†0.57)	39.17 (\psi 0.01)	34.49 (†0.88)
When2com* [18]	47.74 (†1.23)	33.60 (\\$0.40)	35.81 (\(\frac{1}{1}.05\))	56.75 (\(\phi\)0.48)	26.11 (\\$0.92)	19.16 (\(\frac{1}{2}\)0.04)	39.64 (\12.55)	33.81 (\(\psi 0.47 \)
Who2com [19]	48.40 (†0.06)	32.76 (†1.68)	36.04 (†2.61)	57.51 (†1.79)	29.17 (†1.21)	20.36 (†0.58)	39.08 (\(\phi\)0.08)	34.31 (†1.06)
Who2com* [19]	47.74 (†1.23)	33.60 (\\$0.40)	35.81 (†1.05)	56.75 (†0.48)	26.11 (\\$0.92)	19.16 (†0.04)	39.64 (\12.55)	33.81 (\\$0.47)
V2VNet [12]	58.42 (†3.09)	48.33 (\$\\$3.87)	48.51 (\1.59)	70.02 (†0.46)	28.58 (†5.18)	21.99 (†0.57)	41.42 (†0.35)	41.11 (†0.74)
DiscoNet [13]	56.66 (†1.19)	46.98 (\$1.74)	50.22 (\$\psi 1.05)	68.62 (\$\dagger\$0.25)	27.36 (†5.58)	22.02 (†0.82)	42.50 (†0.95)	40.84 (†0.53)
Upper-bound	64.09 (†5.34)	41.34 (†2.42)	48.20 (†0.74)	67.05 (†2.04)	29.07 (†0.74)	31.54 (†3.15)	45.04 (†0.70)	42.29 (†1.98)

The Number in () Represents the Performance Gain by Adding V2I

Baseline segmentation method and evaluation metrics: We follow the backbone architecture as well as the loss function of U-Net [52] in our baseline method. The input is a BEV-based voxelized point cloud, and the output is BEV semantic segmentation. We label and predict seven categories as listed in Table V, while the remaining is unlabeled. In our benchmark, we evaluate the segmentation performance using mean Intersection-over-Union (mIoU).

Quantitative results: As illustrated in Table V, we find that: (1) V2VNet, DiscoNet, and upper-bound achieve comparable performance in terms of terrain and road categories; (2) the attention-based methods (i.e., when2com, who2com) performs

worse because the attention-based mechanisms try to find the collaboration partners with high correlation scores. Whereas, in 3D perception, the collaborators with complementary information should be prioritized during the collaboration. Such nonalignment can make it quite hard for the attention model to learn; (3) there is a large gap between lower-bound and upper-bound regarding the two safety-critical categories: vehicle (45.93% v.s. 64.09%) and pedestrian (20.59% v.s. 31.54%), proving the values of collaboration; (4) employing V2V and V2I jointly can generally enhance the vehicle segmentation over using V2V solely; (5) co-lower-bound performs better than the lower-bound.

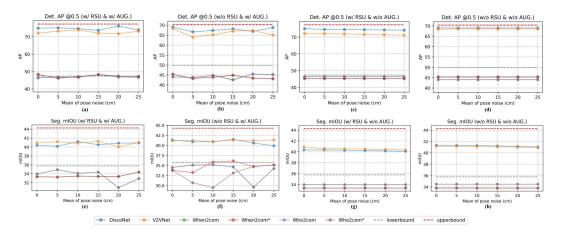


Fig. 6. Experimental results of a robustness test conducted under various magnitudes of pose noise. AUG. means augmentation which adds pose noise during training.

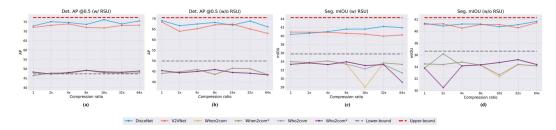


Fig. 7. Experimental results of a bandwidth test conducted under various compression ratios.

Qualitative results: Fig. 5 shows the semantics segmentation results. The results of upper-bound restore the semantic information with rich point cloud data. Intermediate-based collaboration strategies V2VNet and DiscoNet can achieve satisfactory performance yet When2com and Who2com hurt the performance compared to the lower-bound.

D. Discussions on Pose Noise and Compression Ratio

We further examine the robustness of different intermediate models against realistic pose noise (Gaussian noise with a mean of 0.05 m-0.25 m and a standard deviation of 0.02 cm), as shown in Fig. 6. We can see that the models perform comparably with or without pose noise in the training phase, and all the intermediate methods have shown stable performance against the pose noise. The reason is that the intermediate feature map has a relatively low spatial resolution (each grid in the feature map denotes a coverage of $2m \times 2m$), thus is not vulnerable to noisy pose. Meanwhile, we employ an 1×1 autoencoder to further compress the feature channel of the transmitted feature map. We test the models with different compression ratios, and we found that the jointly learned 1×1 autoencoder can even improve the performance a little bit, and most intermediate models achieve comparable performance at different levels of compression, as shown in Fig. 7.

V. CONCLUSION

We propose V2X-Sim dataset based on CARLA-SUMO co-simulation, in order to enable multi-agent collaborative perception research in autonomous driving. Our dataset provides

multi-agent multi-modality sensor streams captured by the vehicles and road-side unit (RSU) in realistic traffic flows. Diverse annotations are provided to support a variety of 3D perception tasks. In addition, we benchmark several state-of-the-art collaborative perception methods in collaborative BEV detection, tracking, and semantic segmentation tasks. Future works include the simulation of latency issues as well as the development of novel evaluation metrics in collaborative perception tasks. We believe our work can inspire many relevant research areas including but not limited to autonomous driving, computer vision, multi-robot system, communication engineering, and machine learning.

ACKNOWLEDGMENT

The authors would like to thank anonymous reviewers for their helpful suggestions, and NYU high performance computing (HPC) for the support.

REFERENCES

- H. Caesar et al., "NuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11618–11628.
- [2] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the KITTI vision benchmark suite," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2012, pp. 3354–3361.
- [3] P. Sun et al., "Scalability in perception for autonomous driving: Waymo open dataset," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2443–2451.
- [4] E. Arnold, O. Y. Al-Jarrah, M. Dianati, S. Fallah, D. Oxtoby, and A. Mouzakitis, "A survey on 3D object detection methods for autonomous driving applications," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 10, pp. 3782–3795, Oct. 2019.

- [5] S. M. Marvasti-Zadeh, L. Cheng, H. Ghanei-Yakhdan, and S. Kasaei, "Deep learning for visual tracking: A comprehensive survey," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 5, pp. 3943–3968, May 2022.
- [6] S. Minaee, Y. Boykov, F. Porikli, A. Plaza, N. Kehtarnavaz, and D. Terzopoulos, "Image segmentation using deep learning: A survey," *IEEE Trans. Pattern Anal. Mach. Intelli.*, vol. 44, no. 7, pp. 3523–3542, Jul. 2022.
- [7] Z. MacHardy, A. Khan, K. Obana, and S. Iwashina, "V2X access technologies: Regulation, research, and remaining challenges," *IEEE Commun. Surv. Tut.*, vol. 20, no. 3, pp. 1858–1877, Jul.—Sep. 2018.
- [8] M. Muhammad and G. A. Safdar, "Survey on existing authentication issues for cellular-assisted V2X communication," Veh. Commun., vol. 12, pp. 50–65, 2018.
- [9] M. Hasan, S. Mohan, T. Shimizu, and H. Lu, "Securing vehicle-to-everything (V2X) communication platforms," *IEEE Trans. Intell. Veh.*, vol. 5, no. 4, pp. 693–713, Dec. 2020.
- [10] V. Mannoni, V. Berg, S. Sesia, and E. Perraud, "A comparison of the V2X communication systems: Its-g5 and C-V2X," in *Proc. IEEE 89th Veh. Technol. Conf.*, 2019, pp. 1–5.
- [11] Q. Chen, S. Tang, Q. Yang, and S. Fu, "Cooper: Cooperative perception for connected autonomous vehicles based on 3D point clouds," in *Proc. IEEE 39th Int. Conf. Distrib. Comput. Syst.*, 2019, pp. 514–524.
- [12] T.-H. Wang et al., "V2VNet: Vehicle-to-vehicle communication for joint perception and prediction," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 605–621.
- [13] Y. Li, S. Ren, P. Wu, S. Chen, C. Feng, and W. Zhang, "Learning distilled collaboration graph for multi-agent perception," in *Proc. Neural Inf. Process. Syst.*, 2021, pp. 29541–29552.
- [14] Y. Yuan and M. Sester, "Comap: A synthetic dataset for collective multiagent perception of autonomous driving," *Int. Arch. Photogrammetry*, *Remote Sens. Spatial Inf. Sci.*, vol. 43, pp. 255–263, 2021.
- [15] Y. Yuan, H. Cheng, and M. Sester, "Keypoints-based deep feature fusion for cooperative vehicle detection of autonomous driving," *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 3054–3061, Apr. 2022.
- [16] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker, "Recent development and applications of SUMO-simulation of urban mobility," *Int. J. Adv. Syst. Meas.*, vol. 5, no. 3/4, pp. 128–138, 2012.
- [17] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proc. 1st Annu. Conf. Robot Learn.*, 2017, pp. 1–16.
- [18] Y.-C. Liu, J. Tian, N. Glaser, and Z. Kira, "When2com: Multi-agent perception via communication graph grouping," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 4105–4114.
- [19] Y.-C. Liu, J. Tian, C.-Y. Ma, N. Glaser, C.-W. Kuo, and Z. Kira, "Who2com: Collaborative perception via learnable handshake communication," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 6876–6883.
- [20] R. Xu, H. Xiang, X. Xia, X. Han, J. Liu, and J. Ma, "OPV2V: An open benchmark dataset and fusion pipeline for perception with vehicle-tovehicle communication," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2022, pp. 2583–2589.
- [21] Z. Xiao, Z. Mo, K. Jiang, and D. Yang, "Multimedia fusion at semantic level in vehicle cooperactive perception," in *Proc. IEEE Int. Conf. Multi*media Expo Workshops, 2018, pp. 1–6.
- [22] Y. Maalej, S. Sorour, A. Abdel-Rahim, and M. Guizani, "Vanets meet autonomous vehicles: A multimodal 3D environment learning approach," in *Proc. IEEE Glob. Commun. Conf.*, 2017, pp. 1–6.
- [23] E. Arnold, M. Dianati, R. de Temple, and S. Fallah, "Cooperative perception for 3D object detection in driving scenarios using infrastructure sensors," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 3, pp. 1852–1864, Mar. 2022.
- [24] M. Howe, I. Reid, and J. Mackenzie, "Weakly supervised training of monocular 3D object detectors using wide baseline multi-view traffic camera data," in *Proc. Brit. Mach. Vis.*, 2021, p. 394.
- [25] H. Yu et al., "DAIR-V2X: A large-scale dataset for vehicle-infrastructure cooperative 3D object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 21361–21370.
- [26] S. Manivasagam et al., "LiDArsim: Realistic LiDAR simulation by leveraging the real world," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11164–11173.
- [27] R. Xu, Y. Guo, X. Han, X. Xia, H. Xiang, and J. Ma, "OpenCDA: An open cooperative driving automation framework integrated with co-simulation," in *Proc. IEEE Int. Intell. Transp. Syst. Conf.*, 2021, pp. 1155–1162.
- [28] E. Arnold, S. Mozaffari, and M. Dianati, "Fast and robust registration of partially overlapping point clouds," *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 1502–1509, Apr. 2022.

- [29] A. Patil, S. Malla, H. Gang, and Y.-T. Chen, "The H3D dataset for full-surround 3D multi-object detection and tracking in crowded urban scenes," in *Proc. Int. Conf. Robot. Automat.*, 2019, pp. 9552–9557.
- [30] M. Pitropov et al., "Canadian adverse driving conditions dataset," Int. J. Robot. Res., vol. 40, pp. 681–690, 2021.
- [31] Q.-H. Pham et al., "A*3D dataset: Towards autonomous driving in challenging environments," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 2267–2273.
- [32] S. Yogamani et al., "Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 9307–9317.
- [33] X. Huang et al., "The apolloscape dataset for autonomous driving," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops, 2018, pp. 1067–10676.
- [34] M. Cordts et al., "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3213–3223.
- [35] G. Ros, L. Sellart, J. Materzynska, D. Vázquez, and A. M. López, "The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3234–3243.
- [36] G. Neuhold, T. Ollmann, S. R. Buló, and P. Kontschieder, "The mapillary vistas dataset for semantic understanding of street scenes," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5000–5009.
- [37] J. Behley et al., "SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 9296–9306.
- [38] Q. Hu, B. Yang, S. Khalid, W. Xiao, A. Trigoni, and A. Markham, "To-wards semantic segmentation of urban-scale 3D point clouds: A dataset, benchmarks and challenges," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 4977–4987.
- [39] M.-F. Chang et al., "Argoverse: 3D tracking and forecasting with rich maps," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8740–8749.
- [40] S. Ettinger et al., "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9710–9719.
- [41] D. Jia and D. Ngoduy, "Enhanced cooperative car-following traffic model with the combination of V2V and V2I communication," *Transp. Res. Part B-methodological*, vol. 90, pp. 172–191, 2016.
- [42] S.-W. Kim et al., "Multivehicle cooperative driving using cooperative perception: Design and experimental validation," *IEEE Trans. Intell. Transp. Syst.*, vol. 16, no. 2, pp. 663–680, Apr. 2015.
- [43] Z. Y. Rawashdeh and Z. Wang, "Collaborative automated driving: A machine learning-based method to enhance the accuracy of shared information," in *Proc. Int. Conf. Intell. Transp. Syst.*, 2018, pp. 3961–3966.
- [44] Q. Chen et al., "DSRC and radar object matching for cooperative driver assistance systems," in *Proc. IEEE Intell. Veh. Symp.*, 2015, pp. 1348–1354.
- [45] S. Suo, S. Regalado, S. Casas, and R. Urtasun, "Trafficsim: Learning to simulate realistic multi-agent behaviors," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 10400–10409.
- [46] S. Tan et al., "Scenegen: Learning to generate realistic traffic scenes," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 892–901.
- [47] P. Wu, S. Chen, and D. Metaxas, "Motionnet: Joint perception and motion prediction for autonomous driving based on bird's eye view maps," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 11382–11392.
- [48] W. Luo, B. Yang, and R. Urtasun, "Fast and furious: Real time end-to-end 3D detection, tracking and motion forecasting with a single convolutional net," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 3569–3577.
- [49] J. Luiten et al., "HOTA: A higher order metric for evaluating multi-object tracking," Int. J. Comput. Vis., vol. 129, no. 2, pp. 548–578, Oct. 2020.
- [50] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: The clear mot metrics," *EURASIP J. Image Video Process.*, pp. 1–10, 2008, Art. no. 246309.
- [51] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, "Simple online and realtime tracking," in *Proc. IEEE Int. Conf. Image Process.*, 2016, pp. 3464–3468.
- [52] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Med. Image Comput. Comput.-Assist. Interv.*, 2015, pp. 234–241.