

M³CAD: Towards Generic Cooperative Autonomous Driving Benchmark

Morui Zhu¹ Yongqi Zhu¹ Yihao Zhu¹ Qi Chen² Deyuan Qu² Song Fu¹ Qing Yang¹

¹University of North Texas ²Toyota InfoTech Labs

Abstract: We introduce M³CAD, a novel benchmark designed to advance research in generic cooperative autonomous driving. M³CAD comprises 204 sequences with 30k frames, spanning a diverse range of cooperative driving scenarios. Each sequence includes multiple vehicles and sensing modalities, e.g., LiDAR point clouds, RGB images, and GPS/IMU, supporting a variety of autonomous driving tasks, including object detection and tracking, mapping, motion forecasting, occupancy prediction, and path planning. This rich multimodal setup enables M³CAD to support both single-vehicle and multi-vehicle autonomous driving research, significantly broadening the scope of research in the field. To our knowledge, M³CAD is the most comprehensive benchmark specifically tailored for cooperative multi-task autonomous driving research. We evaluate the state-of-the-art end-to-end solution on M³CAD to establish baseline performance. To foster cooperative autonomous driving research, we also propose E2EC, a simple yet effective framework for cooperative driving solution that leverages inter-vehicle shared information for improved path planning. We release M³CAD, along with our baseline models and evaluation results, to support the development of robust cooperative autonomous driving systems. All resources will be made publicly available on our project webpage. <https://github.com/zhumorui/M3CAD>

Keywords: Cooperative Autonomous Driving, Benchmark, Multi-Task, End-to-End Planning

1 Introduction

Cooperative autonomous driving (CAD) refers to the paradigm where multiple autonomous vehicles communicate and coordinate with each other to enhance driving efficiency and safety. To advance research in this domain, there is an urgent need for a comprehensive benchmark that enables evaluation and comparison of CAD algorithms and solutions. The benchmark should incorporate the following key features. First, it needs to provide the scenarios involving multiple vehicles, focusing on the cases where they can collaborate with each other. Second, it must support the research for multiple cooperative driving tasks, e.g., cooperative perception, collaborative mapping, joint motion forecasting, and coordinated path planning. Third, it should incorporate diverse sensor types, e.g., LiDAR and cameras, as well as varying environmental settings, including rain and nighttime.

1.1 Limitations of Prior Works

Recent advancements in autonomous driving have demonstrated the effectiveness of end-to-end planning frameworks, e.g., UniAD [1], VAD [2], and Drive-VM [3]. Meanwhile, CAD has gained traction, with most efforts focusing on perception [4, 5]. Comprehensive study of cooperative autonomy, however, is still limited by several key factors. First, existing datasets such as KITTI [6] and NuScenes [7] are primarily designed for single-vehicle settings. They do not support research into multi-vehicle collaboration across various tasks. Second, almost all existing benchmarks provide ground truth based only on individual vehicle’s observations. This limits the ability to study how vehicles can build a shared, global understanding of the environment through data sharing, an

essential component for CAD. Third, the vehicle’s trajectory in most datasets, e.g., OpenV2V [8], often relies on simplified, straight-line trajectories during data collection, which fails to capture the complexity and diversity of real-world driving scenarios.

1.2 M³CAD

To address the above-mentioned limitations, we introduce a novel benchmark designed specifically to support research in **M**ulti-vehicle, **M**ulti-task, and **M**ulti-modality **C**ooperative **A**utonomous **D**riving (**M³CAD**). By leveraging the advanced rendering capabilities of the recently-released Unreal Engine 5 (UE5) in CARLA [9], we can simulate realistic multi-vehicle interactions within diverse driving scenarios. M³CAD comprises 204 sequences in total, offering over 30k frames and more than 267k annotated instances, along with the ground truth data for vehicle trajectories, mapping, and occupancy information. Each sequence includes labeled vehicles with their precise location and trajectory information at each timestamp, as well as the map and occupancy details. This comprehensive dataset supports a wide range of autonomous driving tasks, addressing the limitations in the single-vehicle end-to-end benchmarks (e.g., NuScenes [7]) and cooperative non-end-to-end datasets (e.g., OPV2V [8]).

1.3 Contributions

To the best of our knowledge, M³CAD is currently the most comprehensive benchmark for both *single-vehicle* and *cooperative autonomous driving research*, supporting multiple tasks like *object detection and tracking*, *mapping*, *motion forecasting*, *occupancy*, and *path planning*, while supporting more realistic vehicle movements and interactions in complex environments.

Comprehensiveness. M³CAD enables research on cooperative autonomous driving by supporting multiple vehicles collaborating across a range of tasks. While the ultimate goal of CAD is to achieve safe and efficient path planning, understanding how collaboration among vehicles would enhance these tasks is crucial. Currently, the lack of datasets supporting such research not only limits the community’s understanding of the challenges and limitations of CAD, but also hinders a deeper exploration of what information should be shared among vehicles. The interactions between various components of an autonomous driving system in a cooperative setting remain underexplored. Furthermore, M³CAD is versatile enough to also support single-vehicle autonomous driving tasks, making it a comprehensive benchmark for both single-agent and cooperative autonomous driving research.

Diverse Trajectory. Previous studies have pointed out that existing benchmarks often assume vehicles follow simple straight trajectories to simplify data collection and annotation, leading models to rely primarily on the ego vehicle’s state while neglecting the importance of environmental perception [10]. To better align with the requirements of real-world autonomous driving, the trajectories of ego vehicle in M³CAD is more complex and diverse. Thanks to the powerful rendering capabilities of the UE5 engine and the navigation capability provided by Inverted AI in complex environments (filled with obstacles, traffic lights, and dynamic objects), M³CAD generates realistic and diverse motion trajectories. This design significantly emphasizes the dependence of path planning task on environmental perception.

Flexibility. Alongside the M³CAD benchmark, we offer an easy-to-use tool that allows researchers to create custom driving scenarios in CARLA, while applying benchmarking solutions based on the real-world NuScenes dataset [7]. Most end-to-end autonomous driving solutions are evaluated against the NuScenes dataset which is inherently limited by its original design for single-vehicle scenarios. To support research in multi-vehicle autonomous driving, M³CAD allows researchers to customize driving scenarios, e.g., different weather conditions, various road layouts, and rare driving events, to align with their specific research goals. Such flexibility makes it easier for researchers to conduct targeted experiments and gain valuable insights from specific research problems or tasks.

Table 1: Detailed comparison of the M³CAD with existing benchmarks. CT: Cooperation Type, OD: Object Detection, OT: Object Tracking, MP: Mapping, MF: Motion Forecasting, OCC: Occupancy Prediction, PP: Path Planning.

Benchmarks	# Seq.	# Fr.	3D Label	Night/Rain	Straight	CT	Multiple Tasks					
							OD	OT	MP	MF	OCC	PP
KITTI [6]	22	15k	200k	X/X	—	X	✓	✓	X	X	X	X
NuScenes [7]	1k	40k	1.4M	✓/✓	69.8%	X	✓	✓	✓	✓	✓	✓
OPV2V [8]	73	11k	232k	X/X	77.7%	V2V	✓	✓	X	✓	X	X
V2V4Real [11]	43	20k	240k	X/X	—	V2V	✓	✓	X	✓	X	X
V2X-Seq [12]	95	15k	10.45k	✓/X	—	V2I	✓	✓	✓	✓	X	X
TUMTraf V2X [13]	10	2k	29.8k	✓/X	—	V2I	✓	✓	✓	✓	X	X
M ³ CAD (ours)	204	30k	267k	✓/✓	58.73%	V2V	✓	✓	✓	✓	✓	✓

2 Related Works

Benchmarks for Multi-Task Autonomous Driving. Multi-task autonomous driving framework jointly optimizes multiple modules for individual tasks while prioritizing the ultimate planning objective, benefiting from enhanced safety and interpretability [14]. To support multi-task research, the NuScenes benchmark [7] is proposed and is recognized as the most valuable resource in autonomous driving research, due to its comprehensive and high-quality data that encompasses a wide range of multi-task driving scenarios. Several end-to-end autonomous driving solutions have been evaluated using the NuScenes dataset, e.g., DiffStack [15], MP3 [16], P3 [17], ST-P3 [18], and VAD [2]. Among them, a notable example is UniAD [1], a unified framework that seamlessly integrates perception, prediction, and planning into a single architecture. While multi-task end-to-end autonomous driving frameworks have shown promising performance by jointly optimizing all modules, there remains a significant gap in exploring collaborative strategies among multiple vehicles for autonomous driving tasks.

Benchmarks for Cooperative Perception. Despite the progress in end-to-end autonomous driving, multi-vehicle collaboration has primarily focused on cooperative perception, leading this field underexplored. To support this direction, several benchmarks have been introduced, including OPV2V [8], V2V4Real [11], V2X-Seq [12], TUMTraf-V2X [13], and DAIR-V2X-C [19]. However, these datasets are largely centered around perception tasks (particularly object detection) which limits their utility for more comprehensive autonomous driving challenges such as object tracking, motion forecasting, and ultimately, path planning. Table 1 presents a comparative summary of existing benchmarks and our proposed M³CAD, highlighting differences in dataset scale, the number of 3D object annotations, the proportion of vehicles traveling in straight trajectories, and the range of supported autonomous driving tasks.

3 The Proposed M³CAD Benchmark

M³CAD is designed to serve as a versatile platform that facilitates various cooperative autonomous driving tasks. To achieve this goal, we deployed multiple vehicles in the *Town10HD* map in CARLA. Each vehicle is equipped with four 800 × 600-pixel (110° FOV) cameras, a 64-beam LiDAR, and GPS/IMU sensors to collect multi-modality sensor data during daytime, nighttime, and various weather conditions.

Annotation. We annotate the collected data based on the NuScenes format [7], enabling the research on multi-vehicle, multi-task and multi-sensor autonomous driving. To support research on multiple autonomous driving tasks, we provide in M³CAD comprehensive annotations for *object detection*, *tracking*, *motion prediction*, *occupancy forecasting*, and *path planning*. All objects located within the 102.4m × 102.4m bird’s-eye-view (BEV) range centered on the ego vehicle are annotated, including attributes: *location*, *speed*, *bounding box center* and *extent (length, width, height)*. To support effective collaboration between vehicles, M³CAD also provides transformation matrices that align the ego vehicle with any collaborating vehicles, enabling accurate information fusion.

For object detection and tracking tasks, only objects within the BEV range are used to emulate realistic constraints. In addition, we provide *motion histories* and *future trajectories* for dynamic

vehicles, as well as *binary BEV occupancy maps* from LiDAR point clouds for occupancy forecasting. Vehicle trajectories are logged for planning. Semantic map layers, including *drivable areas*, *lane dividers*, and *road dividers*, are aligned with the CARLA global coordinate system and structured following the nuScenes specification. To increase diversity and realism, data is collected under varied environmental conditions such as rain and nighttime.

Global Ground Truth. In the context of cooperative autonomous driving, we emphasize the importance of using *global ground truth*. *Global ground truth* refers to the information that is observable by all vehicles in the vicinity of the ego vehicle, including bounding boxes, class labels, and object IDs. It represents a complete, shared understanding of the environment, accessible to all vehicles within the cooperative system. In contrast, *local ground truth* is limited to what the ego vehicle can observe independently, without sharing or integrating data from other vehicles. Using only *local ground truth* can mislabel cooperative features during training, leading to performance degradation. This issue is prevalent in many existing solutions [20, 8, 21, 19, 22, 23, 24]. Therefore, we strongly advocate that *global ground truth* must be used to train cooperative solutions, i.e., M³CAD provides *global ground truth* to properly supervise model’s training on cooperative perception and planning.

Dataset Split. The M³CAD dataset is divided into training, validation, and test subsets using a 70/15/15 split. Table 2 provides a summary of the data statistics for each split.

Table 2: Statistics on multi-vehicle cooperation across the M³CAD dataset splits. For the training, validation, and test subsets, we report the average and standard deviation of the cooperation range between vehicles (measured in meters), the number of vehicles that can potentially cooperate with the ego vehicle, and the number of vehicles within the ego vehicle’s BEV range.

Datasets	# of Seq.	Cooperation Range (m)	# of Collaborating Vehicles per Frame	# of Vehicles in BEV
Training	143	29.95 ± 14.41	1.56 ± 0.68	8.81 ± 3.06
Validation	31	31.80 ± 14.11	1.30 ± 0.48	8.35 ± 3.99
Test	30	25.29 ± 16.30	1.28 ± 0.52	9.23 ± 3.11

Single-Vehicle Multiple Autonomous Driving Tasks. The rich data generated by these vehicles allows for the exploration of various tasks of autonomous driving. These tasks include object detection (OD), object tracking (OT), mapping (MP), motion forecasting (MF), occupancy prediction (OCC), and most importantly path planning (PP). Figure 1 shows the results of different tasks performed on the M³CAD dataset.

Multi-Vehicle Multiple Autonomous Driving Tasks. One of the prominent features of M³CAD is its support for *multi-vehicle collaboration*, enabling cooperative path planning to enhance the safety of autonomous driving vehicles. Although cooperative object detection has been extensively studied [20, 8, 24, 25, 23, 26], there remains a lack of comprehensive understanding of vehicle collaboration in other tasks, e.g., cooperative mapping, motion forecasting, and occupancy prediction, which could collectively enhance vehicle’s planning accuracy and driving safety. It’s important to note that not all components benefit equally from data sharing among vehicles, e.g., our experiment results indicate that cooperative motion forecasting can substantially influence the accuracy of path planning results.

As shown in Figure 2, with the assistance from the sender vehicle, ego vehicle can obtain a better perception of the surrounding environment, allowing it to follow a trajectory much closer to the ground truth (with a L2 error of 0.77m). Without M³CAD, it would be extremely difficult, if not impossible, to systematically evaluate the benefits and trade-offs of different tasks within cooperative autonomous driving. Last but not least, M³CAD can be easily extended to include data from additional sensor modalities, such as radar, depth cameras, or ultrasound sensors.

4 BEV-Oriented E2EC Framework

To support the cooperation among autonomous vehicles, we present a simple yet effective cooperative framework, namely End-to-End Cooperative (E2EC) framework, for conneted and autonomous




Figure 1: Illustrations of various autonomous driving tasks using the M³CAD dataset. (a), (b), (d), and (e) show sample images captured from the vehicle’s four cameras. (a) Demonstrates the path planning (PP) results, where the ego vehicle’s predicted trajectory is represented by a dotted line. (b) Shows object tracking (OT) and motion forecasting (MF) results where dotted lines represent predicted trajectories of other vehicles. (c) Presents object detection (OD) results in 3D space. (f) Depicts mapping (MP) and occupancy prediction (OCC) results.

driving, where the key is to support BEV feature fusion to realize collaborations among vehicles. As shown in Fig. 3, it supports end-to-end cooperative autonomous driving while enabling the study of individual cooperative tasks, making it a flexible tool for both research and practical deployment.

BEV Fusion Algorithms. Assuming vehicle v_j is selected for cooperation, it needs to send its BEV feature \mathcal{F}_j to the ego vehicle, along with its pose \mathcal{P}_j and location \mathcal{L}_j information. Based on the affine transformation function $T_j = \Psi(\mathcal{P}_j, \mathcal{L}_j)$, the transformation matrix T_j is obtained. Then, the ego vehicle transforms the received feature map to align with its perspective, resulting in a transformed feature map \mathcal{F}_j' . The transformed \mathcal{F}_j' will then be fused into the ego’s as follows:

$$\mathcal{F}^* = \Phi(\mathcal{F}_e \parallel \mathcal{F}_j') \in \mathbb{R}^{H \times W} \quad (1)$$

where \mathcal{F}_e is the ego vehicle’s BEV feature, H and W are the height and width of the feature map, and $\Phi(\cdot)$ denotes the fusion function. A variety of methods can be employed to implement Φ , including F-Cooper [20], Attentive Fusion [8], CoBEVT [24], V2VNet [25], Where2Comm [22], V2VAM [27], V2X-ViT [21], CoAlign [28], and SiCP [23]. Within the M³CAD benchmarks, several above-mentioned fusion strategies are provided, offering the flexibility of selecting among them or extending them with new fusion methods.

Cooperative Multiple Tasks. The proposed E2EC framework explicitly decomposes the autonomous driving pipeline into a series of modular tasks. This task-wise decomposition enables researchers to focus on solving specific sub-problems within the broader cooperative driving context. Each task in the pipeline serves a unique role and has value beyond direct path planning. For instance, collaborative mapping can benefit digital map providers, while multi-agent motion forecasting is vital for intelligent transportation systems and traffic management. The E2EC framework provides the flexibility to investigate individual tasks in isolation, while also supporting end-to-end integration to study inter-task dependencies.

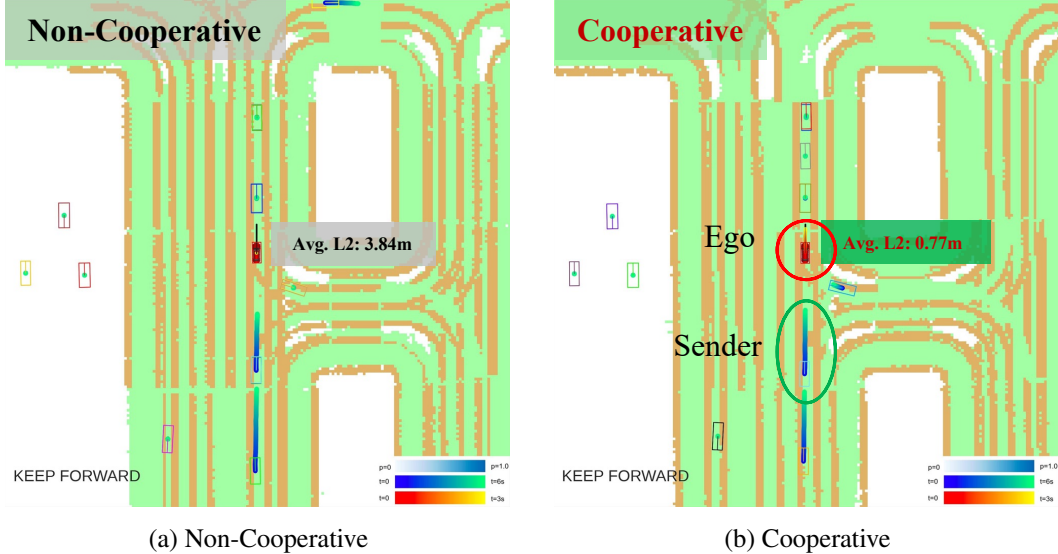


Figure 2: Qualitative comparison between cooperative and non-cooperative path planning. All visualizations are shown in the BEV, with each vehicle depicted as a uniquely colored box. The blue-green curves represent the predicted trajectories of vehicles over the next 6 seconds, showing only the top-1 predicted trajectories here. The red-yellow curves present the planned paths of the ego vehicle within a 3-second window, while black curves indicate the ground-truth trajectories. In the underlying semantic map, green areas indicate drivable regions and yellow dots mark lane dividers.

5 Experiments

We begin evaluating the M³CAD benchmark by assessing its support for single-vehicle and multi-task research, and then extend our analysis to the multi-vehicle and multi-task domain.

Performance of Tracking and Mapping. We evaluate UniAD [1], a widely used single-vehicle end-to-end autonomous driving model, across the NuScenes, M³CAD, and OPV2V* datasets ¹. To investigate whether and how cooperation between vehicles can enhance autonomous driving, we implement the proposed E2EC framework using the modules provided in UniAD. Table 3 shows single-vehicle (S.) and cooperative (C.) tracking and mapping performance, using the same metrics defined in UniAD.

Experiments show that UniAD can be effectively evaluated on both the NuScenes and M³CAD datasets, highlighting the generality and compatibility of the proposed benchmark. It also demonstrates that the end-to-end solution achieves better performance on M³CAD compared to its counterpart, the OPV2V* dataset which is also generated using the CARLA simulator. This performance gain may be largely due to M³CAD’s use of UE5 and high-definition maps, which enable the generation of more realistic and diverse data. Most importantly, enabling vehicles to share BEV features with one another can largely enhance the ego vehicle’s capabilities in tracking and mapping. Notably, *the recall metrics show the most significant gains*, suggesting that the cooperative approach enables better detection and tracking of objects compared to the non-cooperative baseline. This improvement is intuitive: by enabling cooperation between vehicles, the ego vehicle’s field of view and line of sight are extended, allowing it to perceive more objects in the environment.

Performance of Motion Forecasting, Occupancy Prediction and Path Planning. Table 4 presents the performance of UniAD in both non-cooperative and cooperative settings, evaluated across motion forecasting, occupancy prediction, and path planning tasks. Compared to its performance on NuScenes, *UniAD achieves consistently better results on M³CAD across nearly all tasks, with es-*

¹OPV2V* is an extension of the original OPV2V dataset, enhanced to support research on multi-task autonomous driving. For more details, please refer to the supplementary materials.

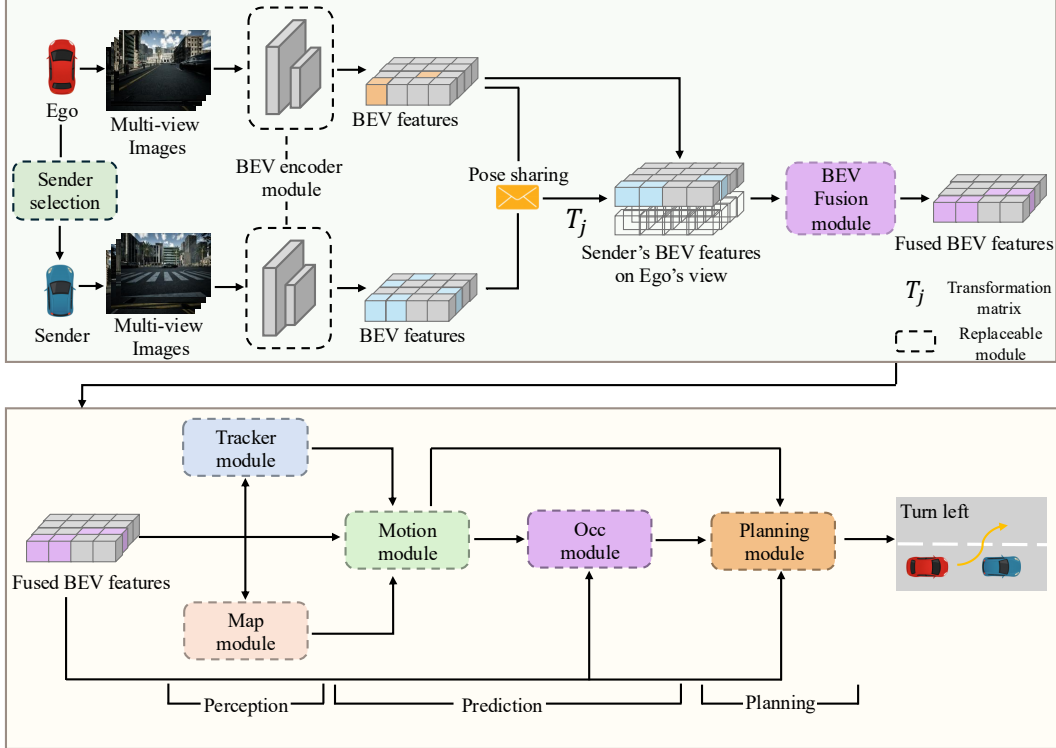


Figure 3: Multi-view images captured by vehicles are processed using a BEV encoder to extract BEV features. The ego vehicle then selects a sender vehicle, which shares its local BEV feature. Before processing the received features, the shared BEV is transformed into the ego vehicle’s perspective. This transformed feature is then fused with the ego’s feature. The final fused BEV is used to support various autonomous driving tasks. Importantly, each of the dashed modules in the system, e.g., sender selection, BEV encoder, BEV fusion, tracking, mapping, motion prediction, occupancy estimation, and planning, can be implemented using different methods or solutions.

Table 3: Comparison of UniAD’s first stage performance on different benchmarks. S.: Single-vehicle tasks, C.: Cooperative tasks.

Methods	Datasets	Object Detection and Tracking				Mapping (%)	
		mAP \uparrow	AMOTA \uparrow	AMOTP \downarrow	Recall \uparrow	IoU-Lane \uparrow	IoU-Road \uparrow
S.	NuScence	0.380	0.359	1.320	0.467	31.3	69.1
	OPV2V*	0.462	0.202	1.140	0.287	41.5	79.1
	M ³ CAD	0.457	0.211	0.662	0.475	49.6	94.0
C.	OPV2V	0.436	0.435	1.160	0.487	39.0	79.4
	M ³ CAD	0.785	0.774	0.579	0.846	50.7	95.3

pecially significant gains in motion forecasting. Specifically, M³CAD yields a 2.4 \times improvement in minADE (minimum Average Displacement Error), likely due to its inclusion of multi-vehicle interactions, which enhance the model’s understanding of object dynamics. These gains could also be driven by improvements in the mapping task (1.4 \times better IoU-Road), enabled by M³CAD’s richer HD maps. This additional context helps UniAD better capture road topology and constraints, resulting in more accurate and realistic trajectory predictions. All of these enhancements finally result in significantly improved path planning results, achieving an average L2 distance of just 0.31m between the predicted and ground truth trajectories, and an average collision probability of 0.01%.

5.1 Ablation Study

Importance of Global Ground Truth. To effectively train cooperative perception models, it is crucial to use global ground truth, as relying on local ground truth can lead to incomplete supervision

Table 4: Comparison of UniAD’s second stage performance on different benchmarks. S.:Single-vehicle tasks, C.: Cooperative tasks.

	Datasets	Motion Forecasting (m)			Occupancy Prediction (%)				Plan ($m, \%$)	
		ADE↓	FDE↓	MR↓	IoU-n↑	IoU-f↑	VPQ-n↑	VPQ-f↑	L2↓	Col.↓
S.	NuScence	0.71	1.02	0.151	63.4	40.2	54.7	33.5	1.03	0.31
	OPV2V*	0.42	0.56	0.001	71.1	39.6	61.4	34.7	2.10	0.41
	M ³ CAD	0.36	0.38	0.001	76.2	57.5	56.2	47.3	0.43	0.02
C.	OPV2V	0.48	0.64	0.0010	71.7	39.3	68.9	38.6	2.29	0.04
	M ³ CAD	0.28	0.30	0.0001	80.9	63.0	76.9	61.4	0.31	0.01

during training. Figure 4 highlights the impact of using only local ground truth on the OPV2V dataset, i.e., a noticeable drop in performance is observed when cooperative perception models are trained with incomplete ground truth, demonstrating the importance and necessity of global ground truth.




Figure 4: Comparison of local GT vs. global GT for the tracking task on OPV2V.

Importance of Perception Data. Since the ego vehicle in NuScenes primarily follows straight-line trajectories, this limits the diversity and complexity of its movement. Prior studies [10, 29] also confirmed that path planning outcomes from NuScenes are heavily influenced by the ego vehicle’s internal state, rather than the sensor data that capture the surrounding environment. We argue that sensor data plays a critical role in autonomous driving, especially for path planning. As shown in Table 5, Ego-MLP [10] performs comparably to UniAD on the NuScenes benchmark, but its performance drops significantly on the M³CAD benchmark. This suggests that when the ego vehicle exhibits more realistic and complex trajectories, the effectiveness of path planning results becomes significantly more dependent on the perception data.

Table 5: Open-loop planning performance on NuScenes and M³CAD benchmarks. † Results on NuScenes are reported from [10].

Benchmarks	Methods	Planning: Avg. L2 (m) ↓				Planning: Avg. Col. (%) ↓			
		1s	2s	3s	Avg	1s	2s	3s	Avg
NuScenes†	Ego-MLP	0.15	0.32	0.59	0.35	0.00	0.27	0.85	0.37
	UniAD	0.20	0.42	0.75	0.46	0.02	0.25	0.84	0.37
M ³ CAD	Ego-MLP	1.02	2.04	3.06	2.04	0.00	0.09	0.30	0.13
	UniAD	0.34	0.46	0.58	0.46	0.00	0.00	0.07	0.02

6 Limitations

The current dataset does not handle a large number of vehicles collaborating with each other. Due to the limitations of GPU memory, we cannot simulate hundreds of vehicles at once or gather their

sensor data. This means the dataset might not be as diverse as it could be and does not support testing large-scale collaboration between vehicles. This limitation is also evident in the fact that only Town10HD is supported by UE5, meaning that only a single digital map is used to generate the dataset. The current annotations only cover vehicles, without including other traffic participants. This may limit the research to vehicle detection and tracking, while tasks related to other road objects may not be supported. The baseline experiments focus on UniAD and provide reference results for single-vehicle and cooperative autonomous driving. Other state-of-the-art end-to-end autonomous driving solutions, e.g., VAD and VAD-V2, were not yet evaluated on M³CAD.

7 Conclusions

In this paper, we present M³CAD, a comprehensive benchmark designed for cooperative autonomous driving. M³CAD consists of 204 multimodal sequences totaling over 30,000 frames and supports a wide range of autonomous driving tasks. To the best of our knowledge, it is currently the most comprehensive benchmark specifically focused on multi-vehicle and multi-task scenarios in this field. We evaluate the state-of-the-art solution UniAD on M³CAD to benchmark its performance and provide a reference for future research. In addition, we introduce E2EC, a simple yet effective cooperative framework to support cooperation among vehicles on various tasks. We believe the M³CAD benchmark, together with the evaluations and proposed baseline, will help drive further research in cooperative autonomous driving and support its transition into real-world applications.

References

- [1] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17853–17862, 2023.
- [2] B. Jiang, S. Chen, Q. Xu, B. Liao, J. Chen, H. Zhou, Q. Zhang, W. Liu, C. Huang, and X. Wang. Vad: Vectorized scene representation for efficient autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8350, 2023.
- [3] Y. Wang, J. He, L. Fan, H. Li, Y. Chen, and Z. Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14749–14759, 2024.
- [4] S. Hong, Y. Liu, Z. Li, S. Li, and Y. He. Multi-agent collaborative perception via motion-aware robust communication network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15301–15310, 2024.
- [5] J. Zhang, K. Yang, Y. Wang, H. Wang, P. Sun, and L. Song. Ermvp: Communication-efficient and collaboration-robust multi-vehicle perception in challenging environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12575–12584, 2024.
- [6] Y. Liao, J. Xie, and A. Geiger. KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *arXiv preprint arXiv:2109.13410*, 2021.
- [7] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [8] R. Xu, H. Xiang, X. Xia, X. Han, J. Li, and J. Ma. Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2583–2589. IEEE, 2022.

- [9] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017.
- [10] Z. Li, Z. Yu, S. Lan, J. Li, J. Kautz, T. Lu, and J. M. Alvarez. Is ego status all you need for open-loop end-to-end autonomous driving? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14864–14873, 2024.
- [11] R. Xu, X. Xia, J. Li, H. Li, S. Zhang, Z. Tu, Z. Meng, H. Xiang, X. Dong, R. Song, et al. V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13712–13722, 2023.
- [12] H. Yu, W. Yang, H. Ruan, Z. Yang, Y. Tang, X. Gao, X. Hao, Y. Shi, Y. Pan, N. Sun, et al. V2x-seq: A large-scale sequential dataset for vehicle-infrastructure cooperative perception and forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5486–5495, 2023.
- [13] W. Zimmer, G. A. Wardana, S. Sritharan, X. Zhou, R. Song, and A. C. Knoll. Tumtraf v2x cooperative perception dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22668–22677, 2024.
- [14] L. Chen, P. Wu, K. Chitta, B. Jaeger, A. Geiger, and H. Li. End-to-end autonomous driving: Challenges and frontiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [15] P. Karkus, B. Ivanovic, S. Mannor, and M. Pavone. Diffstack: A differentiable and modular control stack for autonomous vehicles. In *Conference on robot learning*, pages 2170–2180. PMLR, 2023.
- [16] S. Casas, A. Sadat, and R. Urtasun. Mp3: A unified model to map, perceive, predict and plan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14403–14412, 2021.
- [17] A. Sadat, S. Casas, M. Ren, X. Wu, P. Dhawan, and R. Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16*, pages 414–430. Springer, 2020.
- [18] S. Hu, L. Chen, P. Wu, H. Li, J. Yan, and D. Tao. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *European Conference on Computer Vision*, pages 533–549. Springer, 2022.
- [19] T. U. Institute for AI Industry Research (AIR). Vehicle-infrastructure collaborative autonomous driving: Dair-v2x dataset, 2021.
- [20] Q. Chen, X. Ma, S. Tang, J. Guo, Q. Yang, and S. Fu. F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds. In *2019 ACM/IEEE Symposium on Edge Computing (SEC)*, page 88–100.
- [21] R. Xu, H. Xiang, Z. Tu, X. Xia, M.-H. Yang, and J. Ma. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. In *European conference on computer vision*, pages 107–124. Springer, 2022.
- [22] Y. Hu, S. Fang, Z. Lei, Y. Zhong, and S. Chen. Where2comm: Communication-efficient collaborative perception via spatial confidence maps. *Advances in neural information processing systems*, 35:4874–4886, 2022.

- [23] D. Qu, Q. Chen, T. Bai, H. Lu, H. Fan, H. Zhang, S. Fu, and Q. Yang. Sicip: Simultaneous individual and cooperative perception for 3d object detection in connected and automated vehicles. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8905–8912. IEEE, 2024.
- [24] R. Xu, Z. Tu, H. Xiang, W. Shao, B. Zhou, and J. Ma. Cobevt: Cooperative bird’s eye view semantic segmentation with sparse transformers. *arXiv preprint arXiv:2207.02202*, 2022.
- [25] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *European Conference on Computer Vision*, pages 605–621. Springer, 2020.
- [26] D. Qu, Q. Chen, Y. Zhu, Y. Zhu, S. S. Avedisov, S. Fu, and Q. Yang. Head: A bandwidth-efficient cooperative perception approach for heterogeneous connected and autonomous vehicles. *arXiv preprint arXiv:2408.15428*, 2024.
- [27] J. Li, R. Xu, X. Liu, J. Ma, Z. Chi, J. Ma, and H. Yu. Learning for vehicle-to-vehicle cooperative perception under lossy communication. *IEEE Transactions on Intelligent Vehicles*, 8(4): 2650–2660, 2023.
- [28] Y. Lu, Q. Li, B. Liu, M. Dianati, C. Feng, S. Chen, and Y. Wang. Robust collaborative 3d object detection in presence of pose errors. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4812–4818. IEEE, 2023.
- [29] J.-T. Zhai, Z. Feng, J. Du, Y. Mao, J.-J. Liu, Z. Tan, Y. Zhang, X. Ye, and J. Wang. Rethinking the open-loop evaluation of end-to-end autonomous driving in nuscenec, 2023. URL <https://arxiv.org/abs/2305.10430>.

Supplementary Materials

In this supplementary document, we provide additional details to support our main paper, “M³CAD: Towards Generic Cooperative Autonomous Driving Benchmark”. Section A includes more details of M³CAD benchmark, including the construction principle, the CARLA-based data acquisition pipeline, full sensor suite, coordinate systems, dataset splits, and global ground truth generations. Section B presents the details and hyperparameter settings of the E2EC framework implemented in our experiments. It includes the training details, e.g., training losses for perception, motion forecasting, occupancy prediction and planning. Section C covers the details on how to extend OPV2V dataset to make it support multi-task research. To obtain the OPV2V* datasets, we will discuss the map format conversion, coordinate frame transformations, and nuScenes-compatible formatting. Finally, Section D provides additional qualitative results across object detection, mapping, forecasting, occupancy and planning tasks.

A Details of M³CAD Benchmark

A.1 Data Acquisition.

We utilize the CARLA simulator integrated with UE5 as the rendering backend to simulate multi-vehicle traffic scenarios, collect sensor data, and generate corresponding ground truth annotations. The simulation is conducted within an upgraded version of the Town10HD map, featuring realistically remodeled vehicles and environments. As illustrated in Figure 5, each vehicle in the simulator is equipped with a comprehensive suite of sensors. The coordinate system of these sensors follows the standard right-handed convention, with the X-axis pointing forward, the Y-axis to the left, and the Z-axis upward.

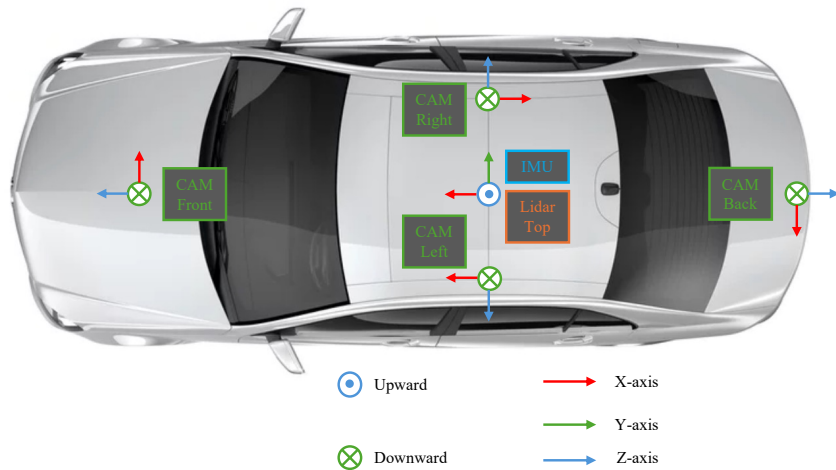


Figure 5: Sensor setup for each CAV in M³CAD. Each vehicle is equipped with four cameras (FRONT, LEFT, RIGHT, BACK), all with a 110° field of view and an image resolution of 800 × 600 pixels, with lens flare and motion blur disabled. A 64-beam LiDAR, mounted at a height of 1.9 m, operates at 10 Hz with a 100 m detection range. GPS/IMU measurements are sampled at 10 Hz, providing global position with an accuracy of approximately 20 mm and vehicle heading with an accuracy of approximately 2°.

In each simulation run, the environment includes 45 parked vehicles and approximately 30–40 NPC (non-playable characters) vehicles. To ensure stability and realism, only vehicle blueprints with *base_type* = “car” are used, excluding collision-prone models such as trailers or motorcycles. For scene diversity, vehicle colors and blueprint types are randomly assigned unless explicitly specified. As shown in Figure 6, we also provide data captured under different weather and lighting conditions. At each simulation timestamp, we record comprehensive ego vehicle information, including the ego pose in the CARLA world coordinate, speed (km/h), LiDAR sensor pose, and future planning




Figure 6: This set of images shows different weather conditions captured in CARLA, from the simplest (Sunny, at the top) to the most complex (Night + Rainy, at the bottom). The row labels indicate the weather conditions, and the column labels represent the camera view.

trajectories. The calibration data for all cameras, including both intrinsic parameters and extrinsic transformations to the LiDAR frame, are also provided to facilitate multi-sensor fusion and 3D perception.

A.2 Global Ground Truth

M³CAD provides *global ground truth* to properly supervise model’s training on cooperative perception and planning. As illustrated in Figure 7, the annotations span a $102.4\text{m} \times 102.4\text{m}$ BEV centered on the ego vehicle and include both visible and heavily occluded agents. This highlights the limitation of single-vehicle perception and the necessity of collaborative sensing in complex scenes.

A.3 Dataset Split and Other Tasks

Dataset Split. Our M³CAD dataset consists of 204 multi-modal sequences, and we apply the 70/15/15 split to divide the data into training, validation, and test sets. Specifically, 143 sequences are assigned to the training set, referred to as M³CAD_{Tra}, while the remaining 31 sequences are used for validation (M³CAD_{Val}) and 30 sequences for testing (M³CAD_{Tst}). In creating these splits, we make every effort to ensure that the distributions of the validation and test sets closely mirror each

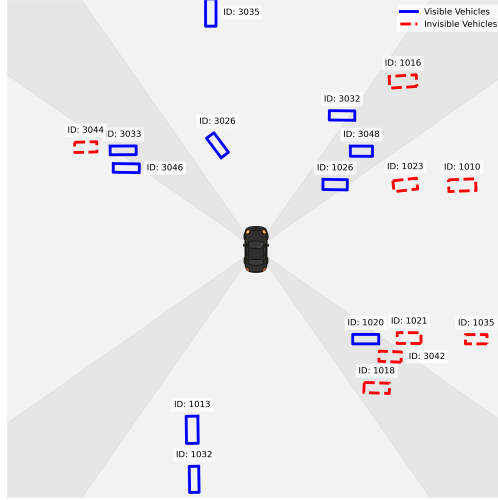


Figure 7: Global ground truth annotations in the BEV are centered on the ego vehicle, corresponding to frame 000290 of vehicle ID 37 in scene 2025_03_03_14_54_49_37. Blue boxes indicate vehicles that are visually identifiable from any ego vehicle, while red dashed boxes represent vehicles that are occluded due to heavy occlusion ($>95\%$). As the ego vehicle is equipped with four cameras, each having a 110° field of view, approximately 10° overlapping areas between adjacent views are shown as darker sectors.

other and collectively cover the majority of the Town10 in CARLA. Figure 8 presents a comparison of the density maps of car annotations for M^3CAD_{Tra} , M^3CAD_{Val} , and M^3CAD_{Tst} .




Figure 8: Polar log-scaled density maps of car annotations for the train, val, and test splits. The radial axis denotes the distance to the ego-vehicle in meters (shown up to $51.2m$), while the polar axis represents the yaw angle relative to the ego heading (0° pointing forward, angles increasing clockwise). Darker bins indicate a higher concentration of annotated car boxes within that angular-radial region.

B Details of BEV-Oriented E2EC Framework

B.1 Implementation of E2EC Framework

We implemented the E2EC framework using the modules provided by UniAD. Table 6 summarizes the main input and output of each module in our E2EC implementation. Starting from four synchronized RGB images, we first extract multi-scale image features with a ResNet-50 backbone and FPN neck. These features are then converted into bird’s-eye-view (BEV) features by the BEVFormer encoder. The SwapFusion encoder then combines two BEV features into one. Next, a Detection Transformer decoder outputs 3D detection results. In parallel, a SegDETR head generates a dense semantic map, and a lightweight CNN followed by a DETR up-sampler predicts instance-level oc-

cupancy over future time steps. Finally, MotionFormer predict vehicles’ future trajectories, and a Transformer decoder with an MLP produces a path planning results.

Table 6: E2EC architecture in the M3CAD benchmark. Shapes are specified by batch size B , number of vehicles $N = 2$, and BEV grid dimensions $256 \times 200 \times 200$. Here, P denotes the number of anchors, N_{agent} the number of agents, and T the number of time steps in the occupancy prediction.

Block	Framework	Main input(s)	Main output(s)
Image backbone & neck	ResNet-50 \rightarrow FPN	RGB images ($B \times 4 \times 3 \times H \times W$)	Multi-level feats: P2 $256 \times H/8 \times W/8$ P3 $256 \times H/16 \times W/16$ P4 $256 \times H/32 \times W/32$
BEV Encoder	Temporal self-attn, Spatial cross-attn	FPN feats, CAN-bus ($B \times 18$)	BEV feats ($B \times 256 \times 200 \times 200$)
Feature fusion	SwapFusionEncoder	Stacked BEVs ($B \times N \times 256 \times 200 \times 200$)	Fused BEV ($B \times 256 \times 200 \times 200$)
Perception & tracking	Detection Transformer Decoder, cls/reg heads	Fused BEV, Detection querie, Track queries	3D boxes ($B \times 900 \times 11$), Past trajectories ($B \times 900 \times 16$)
Mapping	SegDeformable Transform- er, mask heads	Fused BEV, Map queries	Thing masks ($B \times 300 \times 200 \times 200$), Stuff map ($B \times C_{\text{stuff}} \times 200 \times 200$)
Occupancy	Light CNN \rightarrow DETR \rightarrow CVT up-sampler	Fused BEV, Instance queries	Instance segment map ($B \times T \times 200 \times 200$)
Motion forecasting	MotionFormer	Fused BEV, Motion queries, Track queries, Map queries	Future trajectories ($B \times N_{\text{agent}} \times P \times 61$)
Path planning	Transformer Decoder \rightarrow MLP regressor	Fused BEV, Navigation token; Motion queries	Planning trajectory ($B \times 12 \times 2$)

B.2 Training Details

In the two-stage training process, Stage 1 involves jointly training the BEV encoding module, the BEV fusion module, and the perception modules (tracking and mapping). The total loss is defined as:

$$\mathcal{L}_{\text{stage1}} = \mathcal{L}_{\text{track}} + \mathcal{L}_{\text{map}} \quad (2)$$

In Stage 2, we freeze the BEV encoding module and train the remaining modules: the BEV fusion module, tracking, mapping, motion forecasting, occupancy prediction, and planning. The total loss is defined as:

$$\mathcal{L}_{\text{stage2}} = \mathcal{L}_{\text{track}} + \mathcal{L}_{\text{map}} + \mathcal{L}_{\text{motion}} + \mathcal{L}_{\text{occ}} + \mathcal{L}_{\text{planning}} \quad (3)$$

We train the network in two stages over a total of 10 epochs, using a batch size of 1 per GPU. AdamW with a weight decay of 0.01 is employed, along with a cosine-annealing learning rate schedule, preceded by a 500-iteration linear warm-up (with a warm-up ratio of 1/3). The parameters in the image backbone are updated with a $0.1 \times$ learning rate multiplier. All experiments are carried out on four NVIDIA RTX A6000 GPUs.

B.3 Sender Selection Criteria.

While prior works have focused extensively on effective feature fusion between ego and sender vehicles, little attention has been given to the impact of *sender vehicle selection*. Intuitively, vehicles that provide more complementary BEV features are likely to contribute more significantly to enhancing the ego vehicle’s perception of its surroundings. Therefore, identifying and selecting the most informative senders is crucial for improving cooperative perception.

In our implementation, we prioritize senders that observe the scene from different viewpoints to provide complementary information. For each frame, we first identify all surrounding vehicles within a $51.2m$ radius. Among these candidates, we calculate the relative yaw difference between each sender and the ego vehicle. The sender with the largest viewpoint difference is then selected, as it is more likely to observe occluded or otherwise unobserved areas in the ego’s field of view. This strategy enhances viewpoint diversity, ultimately improving the effectiveness of cooperative operation in complex scenarios. The E2EC framework supports customizable sender selection criteria, moving beyond the naive approach of randomly fusing BEV features from nearby vehicles. Possible selection criteria include the distance between the ego and sender vehicles, the size of the overlapping sensing region, or whether the sender is positioned along the predicted trajectory of the ego vehicle. This flexibility enables more strategic and effective cooperation among vehicles in diverse driving scenarios.

B.4 Transformation Matrix

To ensure consistency in the shared perception data, all received BEV features must be transformed into the ego vehicle’s perspective. This step is crucial to eliminate discrepancies in overlapping regions caused by differing viewpoints across vehicles. Specifically, a transformation matrix T_j is required to convert the BEV \mathcal{F}_j generated by any vehicle v_j into the ego vehicle’s coordinate system. This transformation is performed by sequentially mapping the BEV from v_j ’s local LiDAR coordinate system to its IMU frame, then to the global coordinate system, followed by mapping into the ego vehicle’s IMU frame, and finally to the ego’s LiDAR coordinate system. Within the E2EC framework and with the support from the M³CAD dataset, such transformation matrices can be readily obtained for any participating vehicle.

C OPV2V* Preparation

As the OPV2V does not support multiple autonomous driving tasks, e.g., mapping and path planning, it serves primarily as a benchmark for studying the cooperative perception of multiple vehicles. To add the multi-task feature, we developed a tool to extend the original OPV2V dataset by incorporating data and ground truth annotations for additional autonomous driving tasks, including object tracking, mapping, motion forecasting, occupancy grid prediction, and path planning. Based on using more accurate global ground truth, we constructed an extended dataset named OPV2V*, and it is designed to support multi-vehicle, multi-task research, enabling more comprehensive studies of cooperative autonomous driving systems.

C.1 Conversion to NuScenes Format

To ensure the compatibility with existing autonomous driving benchmarks, we converted the original OPV2V dataset into the nuScenes dataset format. This conversion facilitates consistent evaluation over various tasks, including detection, tracking, mapping, and planning. The converted dataset OPV2V* also supports direct comparison with other autonomous driving benchmarks using the original OPV2V dataset.

OPV2V* comprises 164 scenes and 25,561 frames of LiDAR point clouds, and 4 RGB camera images, collected from 7 vectorized maps (Town01, Town02, Town03, Town04, Town05, Town06, Town07, and Town10HD) in CARLA. The dataset includes a total of 675,984 ground truth 3D

object labels. To match the NuScenes format, the OPV2V dataset was re-indexed using consistent 0.5-second timestamp intervals. We also established the prev/next frame relationships among frames to support the tracking task. A key challenge in preparing the OPV2V* dataset was converting between different coordinate systems. OPV2V uses a left-handed coordinate system, while NuScenes uses a right-handed one. To make the data compatible, we carefully transformed all pose-related information, e.g., positions and orientations, to ensure all data aligns correctly. The sensor-to-sensor coordinate system transformations are illustrated in Figure 9.

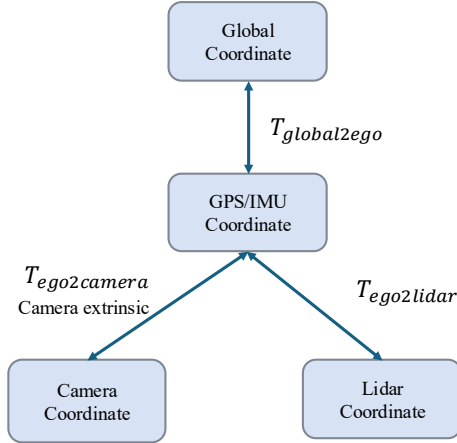


Figure 9: Transformations between different sensors’ coordinate systems where $T_{global2ego}$, $T_{ego2camera}$, and $T_{ego2lidar}$ represent the transformation matrix to convert coordinates from the global system to the ego vehicle’s system, from the ego vehicle’s system to the camera’s system, and from the ego vehicle’s system to the LiDAR’s system, respectively.

C.2 Map Format Conversion

Another challenge in preparing the OPV2V* dataset is converting the map data from the OpenDrive format (.xodr) into the vector map format (.json) supported by NuScenes. OpenDrive represents road networks with geometries, while the NuScenes vector map uses a layer-based representation, where each semantic element (like drivable areas or road segments) is represented in layers. These layers are designed as polygons in a JSON format with global coordinates. Each polygon is defined by a series of nodes, with each node representing a 2D coordinate in the global coordinate system.

To address the differences between the two formats, we developed a tool to convert the data. The tool extracts geometric and semantic features from the OpenDrive file, such as geometry, laneOffset, and laneSection. These features are then restructured into NuScenes vector map layers. For each OpenDrive element, we generate the corresponding polygon and assign it to the right category (e.g., drivable area, lane dividers, or road dividers) based on semantic mapping. Figure 10 visualizes this transformation process, demonstrating how a curved lane in OpenDrive is converted into a polygonal region in the NuScenes-style representation. Due to the large scale of the entire map, only a partial region of Town04 is visualized for clarity.

As OpenDrive represents road geometry using parametric equations, the basic unit it uses is a line segment. These segments can be made up of different geometric shapes like straight lines, arcs, and spirals. In contrast, the NuScenes vector map format uses nodes as the basic unit. Each semantic layer (like drivable area or road segment) is made up of polygons, with each polygon defined by an ordered list of nodes. To bridge the structural differences, we propose a two-stage conversion pipeline comprising **geometric discretization** and **semantic reconstruction**.

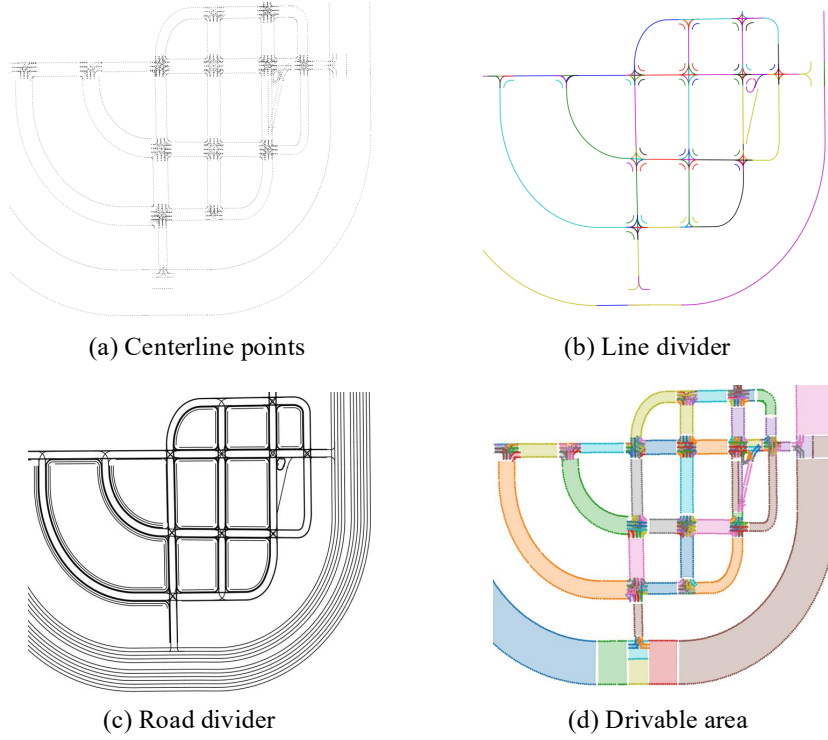


Figure 10: Visualization of the map Town04 conversion process from OpenDrive to NuScenes-style representation. (a) uniformly sampled centerline points; (b) generated lane dividers, where different colors represent segments corresponding to different road lanes; (c) generated road dividers; and (d) reconstructed drivable area represented as polygons, where each colored polygon corresponds to a separate road segment or structural element.

C.2.1 Geometric Discretization

Each road segment in OpenDrive map is preprocessed by uniformly sampling discrete points along its geometric shape. For line segments, which are defined by a starting point P_s , a length l , and a direction vector \mathbf{d} , we generate interpolated points with a fixed step size Δl , as shown in Equation (1). For arc segments, sampling is performed along the arc length s , and the corresponding 2D coordinates are computed using a parametric function $\mathcal{G}_{\text{arc}}(s)$, which is determined by the arc’s starting position, initial orientation, and curvature. The formulation is shown in Equation (2).

$$P_i = P_s + i \cdot \Delta l \cdot \mathbf{d}, \quad i = 0, 1, \dots, \left\lfloor \frac{l}{\Delta l} \right\rfloor \quad (1)$$

$$P_i = \mathcal{G}_{\text{arc}}(s_i), \quad s_i = i \cdot \Delta s, \quad i = 0, 1, \dots, \left\lfloor \frac{L}{\Delta s} \right\rfloor \quad (2)$$

Through this uniform sampling process, we obtain a centerline point set $\{C_i\}$, which serves as the geometric foundation for semantic reconstruction.

C.2.2 Semantic Reconstruction

Based on the discretized centerline point set $\{C_i\}$, we reconstruct NuScenes-style semantic polygons by incorporating lane attributes (e.g., `laneSection`). For each center point C_i , we compute the left and right lane boundaries according to the lane width, as shown in Equation (3) and Equation (4), where \mathbf{n}_i denotes the normal vector at C_i . The boundaries are then sampled using the same

method as in Section C.2.1. This process yields polygonal representations for semantic regions such as drivable areas, lane dividers, and road boundaries.

$$L_i = C_i + w(s_i) \cdot \mathbf{n}_i \quad (3)$$

$$R_i = C_i - w(s_i) \cdot \mathbf{n}_i \quad (4)$$

D More Qualitative Results

In this section, we present qualitative results to visually demonstrate the effectiveness of our benchmark. Figure 11 visually compares the performance of different autonomous driving tasks between the cooperative and non-cooperative solutions. The results highlight our benchmark’s ability to consistently deliver robust performance across all tasks, including object detection and tracking, mapping, motion forecasting, occupancy prediction, and path planning. Specifically, the cooperative method performs much better than the non-cooperative one in all tasks. Figure 12 shows the corresponding images captured by the four cameras on the ego vehicle. As we can see, the cooperative solution generally provides better motion forecasting results than the non-cooperative one.

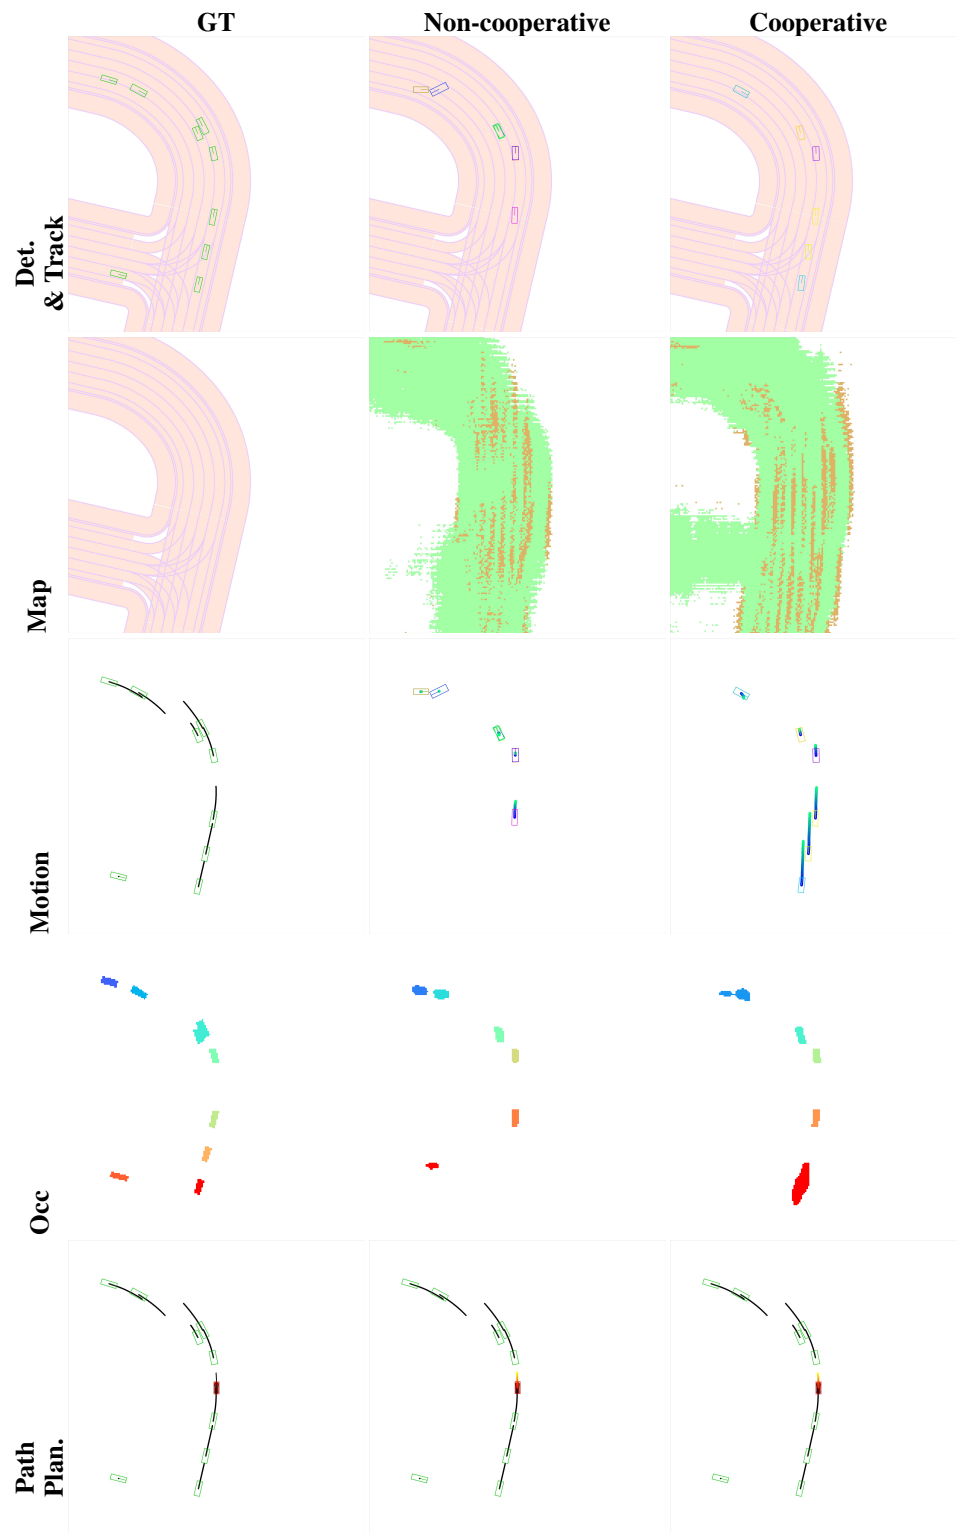


Figure 11: Qualitative comparison of the performance of multiple tasks under *ground truth* (GT), *non-cooperative*, and *cooperative*, corresponding to the frame 002002 of vehicle ID 60 in scene 2025_05_05_14_08_01_60



Figure 12: Rows list the camera view (front, back, left, right), columns compare the performance of single vehicle against the cooperative solution, showing the differences in object detection and motion forecasting results from different viewpoints.