

# V2V-GoT: Vehicle-to-Vehicle Cooperative Autonomous Driving with Multimodal Large Language Models and Graph-of-Thoughts

Hsu-kuang Chiu<sup>1,2</sup> Ryo Hachiuma<sup>1</sup> Chien-Yi Wang<sup>1</sup> Yu-Chiang Frank Wang<sup>1</sup>  
 Min-Hung Chen<sup>1\*</sup> Stephen F. Smith<sup>2\*</sup>

**Abstract**—Current state-of-the-art autonomous vehicles could face safety-critical situations when their local sensors are occluded by large nearby objects on the road. Vehicle-to-vehicle (V2V) cooperative autonomous driving has been proposed as a means of addressing this problem, and one recently introduced framework for cooperative autonomous driving has further adopted an approach that incorporates a Multimodal Large Language Model (MLLM) to integrate cooperative perception and planning processes. However, despite the potential benefit of applying graph-of-thoughts reasoning to the MLLM, this idea has not been considered by previous cooperative autonomous driving research. In this paper, we propose a novel graph-of-thoughts framework specifically designed for MLLM-based cooperative autonomous driving. Our graph-of-thoughts includes our proposed novel ideas of occlusion-aware perception and planning-aware prediction. We curate the V2V-GoT-QA dataset and develop the V2V-GoT model for training and testing the cooperative driving graph-of-thoughts. Our experimental results show that our method outperforms other baselines in cooperative perception, prediction, and planning tasks. Our code and dataset are released to facilitate open-source research at <https://eddyhkchiu.github.io/v2vgot.github.io/>.

## I. INTRODUCTION

Today’s autonomous vehicles rely mainly on mounted cameras or LiDAR sensors to perceive the world, understand the dynamic surrounding scenes, and take driving decisions over time. Inherently such reliance on the vehicle’s local sensors can be limiting, particularly in situations where vehicles and other potential obstacles are occluded by other large nearby objects, such as buses or trucks. To mitigate this safety issue, recent research has proposed mechanisms for V2V-enabled cooperative perception [1]–[8]. In a cooperative perception framework, each Connected Autonomous Vehicle (CAV) can share its individual perception information with others and improve overall detection and tracking accuracy.

How to best realize cooperative perception remains a question for research. One recent trend in autonomous driving research is using Large Language Models (LLMs) and their variants to build end-to-end driving systems that digest the raw sensor input and then generate the planning results. Such LLM-based autonomous driving models [9]–[11] are motivated by the belief that LLM-based approaches could benefit from the reasoning and generalization capabilities that large-scale pre-trained data provides. However, most LLM-

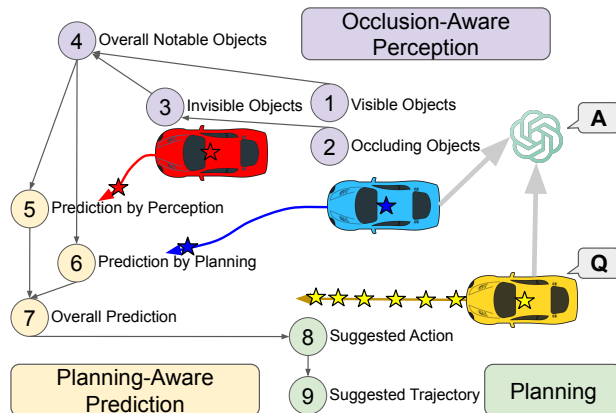


Fig. 1: Illustration of our proposed graph-of-thoughts reasoning framework for cooperative autonomous driving. All Connected Autonomous Vehicles (CAVs) share their perception features with the Multimodal Large Language Model (MLLM), as illustrated by the grey arrows. Any CAV can ask the MLLM to provide a suggested future trajectory or answer perception or prediction questions. The MLLM fuses the perception features from all CAVs and performs inference by following the graph-of-thoughts. If two QA nodes are connected by a directed edge in the graph, as illustrated by black arrows, the answer of the parent node QA is used as the input context of the child node QA. Other colored curved arrows illustrate the predicted or suggested future trajectories. Color stars represent current locations of objects, predicted or suggested future waypoints.

based driving research focuses on individual autonomous vehicles without cooperative perception or planning.

Some recent pioneering research works [12]–[15] have started to explore LLM-based cooperative autonomous driving. However, the V2V cooperation in prior works [13], [14] only involves the LLM as a pure language-based negotiator to resolve planning conflicts, which has not taken advantage of the multimodal understanding ability of MLLMs. In contrast, previous work V2V-LLM [12] adopts MLLMs to fuse perception features from multiple CAVs and provide answers to perception and planning questions. Although V2V-LLM [12] shows promising results, it has not considered the additional benefit of applying chain-of-thoughts or graph-of-thoughts reasoning on MLLMs.

In this paper, we propose a new graph-of-thoughts reasoning framework designed for MLLM-based cooperative autonomous driving, as illustrated in Figure 1. We follow the

<sup>1</sup>NVIDIA, <sup>2</sup>Carnegie Mellon University, \* equally advising.

The authors thank Boyi Li, Zhiding Yu, Boris Ivanovic, and Marco Pavone from NVIDIA for valuable discussions and feedback.

This research was funded by NVIDIA, the CMU Safety21 University Transportation Center, and CMU Robotics Institute.

problem setting from V2V-LLM [12], where all CAVs share their perception features with a MLLM. Any CAV can ask the MLLM to answer a driving question. The MLLM fuses the perception features from all CAVs and performs inference to provide the answer. Unlike previous work [9], [12], we propose a novel graph-of-thoughts reasoning framework specifically designed for cooperative autonomous driving scenarios, as illustrated in Figure 1 and 2.

To verify the effectiveness of our proposed idea, we first curate the V2V-GoT-QA dataset, built upon the V2V4Real [16] cooperative driving dataset. Our V2V-GoT-QA includes 9 different types of QA samples, as illustrated in Figure 2. Our proposed **occlusion-aware perception** questions (Q1 - Q4) consider visible, occluding, and invisible objects. Our proposed **planning-aware prediction** questions (Q5 - Q7) include prediction by perception features and prediction by other CAVs' current planned future trajectories. Our planning questions (Q8 - Q9) provide the suggested action settings and suggested waypoints of future trajectories to avoid potential collisions. Different types of QA are connected by a directed edge in our graph-of-thoughts, as shown in Figure 1. The answer of the parent node QA is used as the input context of the child node QA.

In addition to the newly curated dataset, we also develop our baseline model V2V-GoT, as shown in Figure 3. Unlike the prior work V2V-LLM [12] and DriveLM [9] that only take the perception features at the current timestep as visual input, our V2V-GoT includes the perception features from both current and previous timesteps to better capture the temporal dynamics of the surrounding objects.

In our experiments, we compare the final planning performance of our proposed V2V-GoT and other baseline methods from the prior work V2V-LLM [12], including *no fusion*, *early fusion*, and *intermediate fusion*. Experimental results show that our model outperforms all other baselines in the planning tasks, achieving the lowest collision rates and L2 errors. Furthermore, our ablation study also verifies that the newly proposed QA types in our graph-of-thoughts, namely occlusion-aware perception QAs (Q1 - Q4) and planning-aware prediction QAs (Q5 - Q7), are helpful in achieving better cooperative perception, prediction, and planning performance for MLLM-based cooperative autonomous driving.

Our contribution can be summarized as follows.

- We present a novel graph-of-thoughts reasoning framework for MLLM-based cooperative autonomous driving. The approach includes QA types designed specifically for cooperative driving, including **occlusion-aware perception** and **planning-aware prediction**.
- We curate the V2V-GoT-QA dataset and develop the V2V-GoT model to establish the benchmark and baseline method for future comparative research on MLLM-based cooperative autonomous driving with graph-of-thoughts reasoning.
- Experimental results demonstrate the effectiveness of V2V-GoT in improving overall perception, prediction, and planning performance for cooperative autonomous driving.

## II. RELATED WORK

### A. Vehicle-to-Vehicle Cooperative Perception

V2V cooperative perception algorithms are proposed to improve the overall detection and tracking accuracy in cooperative driving scenarios by sharing perception information between CAVs. Pioneering works F-cooper [17] and V2VNet [18] propose the intermediate fusion approach that shares perception feature maps, achieving a good balance between performance and communication cost. V2V4Real [16] is the first real cooperative driving dataset available worldwide with detection and tracking benchmarks that evaluate the real-world performance of numerous abundant cooperative detection [1]–[3] and cooperative tracking algorithms [5], [19]. However, this group of research focuses on perception and has not explored the planning part of autonomous driving. In contrast, our work covers both cooperative perception and planning.

### B. LLM-based Autonomous Driving

A more recent autonomous driving research trend is applying an LLM to build end-to-end autonomous driving models due to its promising reasoning ability. Pioneering work GPT-driver [20], [21] encodes the state of the ego vehicle and its object detection results into text and allows LLM to identify notable objects and suggest a future trajectory. Other works [22], [23] apply MLLMs to perform scene understanding from the image input. More advanced work, DriveLM [9], and LingoQA [24] develop MLLM-based end-to-end models to further generate the suggested action or future trajectory as output. DriveLM [9] also considers graph-of-thoughts but only for a single autonomous vehicle without V2V cooperation. This group of research focuses only on individual autonomous vehicles without cooperative perception or planning. In contrast, our work develops MLLMs-based models for V2V cooperative driving.

### C. LLM-based Cooperative Autonomous Driving

A few very recent works [12]–[15] start exploring LLM-based cooperative autonomous driving. CoDrivingLLM [14] uses an LLM as a pure language-based conflict coordinator in simulated cooperative driving scenarios in an intersection. CoLMDriver [13] uses a Vision Language Model (VLM) as the initial planner of an individual autonomous vehicle without cooperation, while the V2V cooperation module still uses an LLM as a negotiator based on pure language in simulated environments. Neither prior work has taken advantage of the multimodal understanding ability of MLLMs in their V2V cooperation modules nor verified their methods in real-world cooperative driving datasets. Another pioneering work V2V-LLM [12] proposes an MLLM-based cooperative driving model and shows promising results in real-world datasets. However, the chain-of-thoughts or graph-of-thoughts reasoning capabilities have not been explored in this prior work. In contrast, our work proposes a novel graph-of-thoughts reasoning framework, including the special QAs designed for cooperative driving, such as occlusion-aware perception and planning-aware prediction.

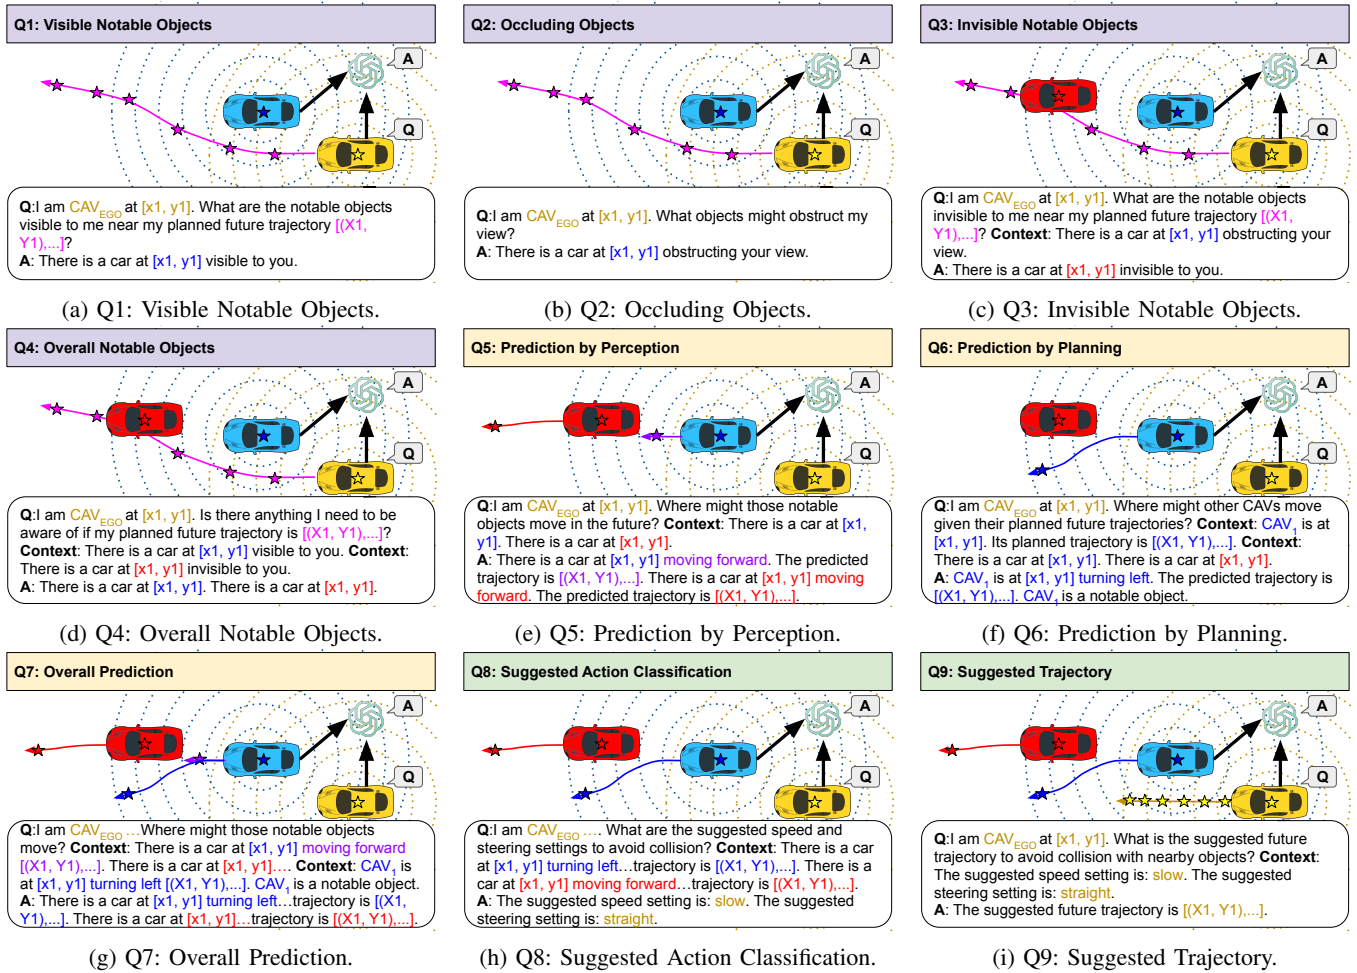


Fig. 2: Illustration of V2V-GoT-QA's 9 types of QA pairs: Perception (Q1 - Q4), Prediction (Q5 - Q7), and Planning (Q8 - Q9). The black arrows pointing at the MLLM indicate the perception data from CAVs. Other colored arrows represent predicted or suggested future trajectories.

### III. V2V-GOT-QA DATASET

#### A. Problem Setting

We build on the MLLM-based cooperative autonomous driving problem setting from V2V-LLM [12]. The MLLM fuses the perception features from all CAVs and performs inference to answer a question. Unlike the prior work V2V-LLM [12], we propose a novel graph-of-thoughts reasoning framework, including the ideas of **occlusion-aware perception** and **planning-aware prediction**. We design 9 types of question-answer pairs in our V2V-GoT-QA dataset, including perception, prediction, and planning, as illustrated in Figure 2. Different types of QAs are connected by directed edges in the graph-of-thoughts, as shown in Figure 1 and 4a. The answer of the parent node QA is used as the input context of the child node QA.

#### B. Dataset Details

Following V2V-LLM [12], we use V2V4Real [16] as the base dataset and curate various QA pairs. V2V4Real [16] has 7105 training frames and 1993 testing frames. For each frame that is 3 seconds before the end of each driving sequence in

V2V4Real [16], we create one QA pair for each CAV per QA type. Overall, the resulting V2V-GoT-QA dataset has 110610 training and 31014 testing QA pairs.

#### C. Dataset Curation

We use the ground-truth bounding box annotations, each CAV and object's future trajectory, and their geometric relationships in V2V4Real [16] to curate the QA pairs of V2V-GoT-QA dataset. In our proposed graph-of-thoughts structure, the answer of the parent node QA is used as the input context of the child node QA. For training data, we use the ground-truth answer to generate the input context of the child node QA. During inference, we use the model inference output of the parent node QA to generate the input context of the child node QA.

The QA types included in the V2V-GoT-QA dataset span tasks relating to Perception (Q1 - Q4), Prediction (Q5 - Q7), and Planning (Q8 - Q9). Details of each QA design rationale, how they are connected, dataset curation, and evaluation metrics are described in the following.

#### D. Perception QAs

The main goal of the perception QAs is to identify notable objects near ego CAV's current planned future trajectory. We propose the idea of **occlusion-aware perception** to the graph-of-thoughts reasoning and divide this task into subtasks of identifying notable objects that are visible (Figure 2a) and invisible (Figure 2c) to the ego CAV separately and then merge the results. This divide-and-conquer design can potentially make it easier for the model to learn which part of the perception feature maps to focus on in order to find the notable objects.

**Q1. Visible Notable Objects (Figure 2a):** We ask the MLLM to identify notable objects that are visible to the ego CAV and near its current planned future trajectory. Six waypoints from a planned future trajectory in the next 3 seconds are used as the reference trajectory in the question. To generate the ground-truth answer in the training data, at most 3 ground-truth objects within 10 meters of the reference trajectory that are detected by the ego CAV's individual 3D object detector are used as the ground-truth answer. We use the F1 score as the evaluation metric.

**Q2. Occluding Objects (Figure 2b):** Before identifying notable objects invisible to the ego CAV, locating the occluding objects first may provide useful information. To generate the ground-truth answer in the training data, at most 3 ground-truth objects closest to the ego CAV that are not occluded by other objects are used as the answer. We use the F1 score as the evaluation metric.

**Q3. Invisible Notable Objects (Figure 2c):** We ask the MLLM to identify notable objects that are invisible to the ego CAV and near its current planned future trajectory. In addition to the question, we also provide the input context from the answer of Q2. Occluding Objects (Figure 2b). Such context could be useful for the model to identify notable objects that are occluded from the ego CAV's field of view by focusing on the occluded region but in other CAVs' perception feature maps. To generate the ground-truth answer in the training data, at most 3 ground-truth objects within 10 meters of the reference trajectory that are not detected by the ego CAV's individual 3D object detector are used as the answer. We use the F1 score as the evaluation metric.

**Q4. Overall Notable Objects (Figure 2d):** This question simply takes the answers from Q1. Visible Notable Objects (Figure 2a) and Q3. Invisible Notable Objects (Figure 2c) as the input context and merges them to generate the final notable object identification answer, which are then used as the input context for the subsequent prediction questions. We use the F1 score as the evaluation metric.

#### E. Prediction QAs

Predicting the future trajectories of notable objects is critical for the final planning task to suggest a future trajectory for the ego CAV that can avoid potential collisions. We propose the idea of **planning-aware prediction** to the graph-of-thoughts reasoning and divide the prediction task into two parts: prediction by perception (Figure 2e) and prediction by

planning (Figure 2f), and then merge their results to generate the final prediction output.

**Q5. Prediction by Perception (Figure 2e):** We ask the MLLM to predict the future trajectories of the notable objects and classify their movement into one of the following 4 categories: *moving forward*, *turning left*, *turning right*, and *staying at the same location*. In addition to the question, the MLLM also takes the answer of Q4. Overall Notable Objects (Figure 2d) as context. With this context, the MLLM knows where the notable objects are at the current timestep. This prediction process relies mainly on the perception features in the current and previous timesteps from all CAVs. We use the ground-truth future trajectories of the notable objects and perform motion classification with heuristic threshold values on the trajectories as the ground-truth answers. We use L2 error as the evaluation metric.

**Q6. Prediction by Planning (Figure 2f):** If a notable object suddenly accelerates, decelerates, or changes directions, it may be difficult to predict its future trajectory by only observing its current and past locations and motions. However, if a notable object is also a CAV, it can directly share its planned future trajectory with other CAVs. That planning information could be used as a more accurate predicted future trajectory to help other CAVs' prediction and planning procedures. In Q6, we include the current planned future trajectories of other CAVs and the answer of Q4. Overall Notable Objects (Figure 2d) as the input context in the question. To generate the answer, we perform motion classification in the same manner as for Q5. We further identify whether other CAVs are notable objects or not in Q6's answer. To measure performance, we calculate the binary classification accuracy in determining whether the model can correctly identify if other CAVs are notable objects or not.

**Q7. Overall Prediction (Figure 2g):** This question takes the answers from Q5. Prediction by Perception (Figure 2e) and Q6. Prediction by Planning (Figure 2f) as the input context and merges them to generate the final prediction answer. If the answer of Q6 indicates that another CAV is a notable object, the model is expected to learn to use that CAV's planned future trajectory as its predicted future trajectory. The merged prediction output is then used as the input context for subsequent planning tasks. To measure performance, we compute L2 errors on the prediction output.

#### F. Planning QAs

We first classify the suggested action's speed and steering settings, and then use them as the input context to provide the suggested future trajectory.

**Q8. Suggested Action Classification (Figure 2h):** We use the prediction results from the answer of Q7. Overall Prediction (Figure 2g) as the input context and ask the MLLM to provide the suggested speed and steering settings. We adopt a similar approach from DriveLM [9]. The speed setting can be one of the following 5 categories: *fast*, *moderate*, *slow*, *very slow*, and *stop*. The steering setting can be one of the following 5 categories: *left*, *slightly left*, *straight*,

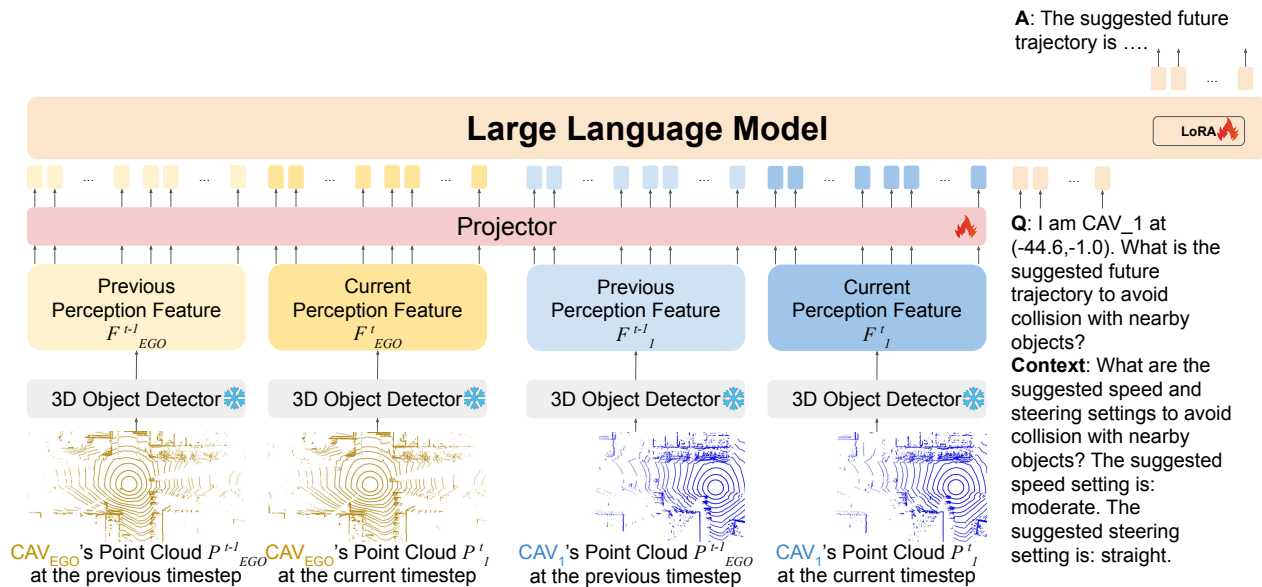


Fig. 3: Model architecture of V2V-GoT.

*slightly right*, and *right*. We classify the speed and steering settings by applying heuristic threshold values on the average difference between consecutive waypoint coordinates. To measure performance, we calculate the L1 errors between the output and ground-truth answer indices on the suggested speed and steering settings. For example, if the output speed and steering settings of a sample are *fast* and *slightly left*, while the ground-truth answer settings are *very slow* and *right*, the L1 error of this sample is  $|0 - 3| + |1 - 4| = 6$ .

**Q9. Suggested Trajectory (Figure 2i):** We use the suggested speed and steering setting from the answer of Q8. Suggested Action Classification (Figure 2h) as context and ask the LLM to provide the 6 waypoints of the suggested future trajectory for the next 3 seconds that avoids collisions. We use the ground-truth future trajectory as the ground-truth answer. To measure performance, we calculate the L2 errors and collision rates.

#### IV. V2V-GoT MODEL

##### A. Architecture

We build our V2V-GoT model with LLaVA-v1.5-7b [25] as the base MLLM architecture, as shown in Figure 3. Unlike the original LLaVA [25] that uses an image encoder, we apply a LiDAR-based 3D object detector, PointPillars [26], to extract perception features from each individual CAV's point cloud. Unlike V2V-LLM [12]'s model that only uses the perception features at the current timestep, our model uses the perception features at the current and previous timesteps from all CAVs as the input of the project layers in the MLLM to generate the visual tokens. The MLLM takes the visual tokens and the language tokens from the question and the context as input and generates the final answer in the natural language format.

##### B. Training and Inference

We follow the similar training settings and hyperparameters from V2V-LLM [12] and LLaVA [25]. During training,

we only train the projector layers and the LoRA [27] parts of the model and freeze the remaining parts. Our model is trained with NVIDIA H100-80GB GPUs and a batch size of 32. We use Adam optimizer with a starting learning rate  $2e^{-5}$  and train our model for 10 epochs on our V2V-GoT-QA training dataset. During testing inference, we follow our proposed graph-of-thoughts illustrated in Figure 1 and 4a. If two QA nodes are connected by a directed edge in the graph-of-thoughts, the inference model output of the parent node QA is used as the input context of the child node QA.

#### V. EXPERIMENT

##### A. Baseline Methods

We adopt V2V-LLM [12] and its baseline methods with different fusion approaches, such as *no fusion*, *early fusion*, and *intermediate fusion*, to compare with our proposed V2V-GoT. In *no fusion*, the MLLM only takes a single CAV's perception features as visual input and answers the driving-related question. This baseline shows the performance in driving scenarios without cooperative perception. In *early fusion*, the union of the point cloud of both CAVs is used to extract perception features, which are then used as the visual input of the MLLM. This approach usually requires high communication costs and is thus impractical for actual deployment. In *intermediate fusion*, such as CoBEVT [3], V2X-ViT [2], and AttFuse [1], cooperative 3D object detectors are used to extract perception features as MLLM's visual input. These approaches usually achieve a good balance of performance and communication cost. V2V-LLM [12] was proposed as a new fusion method that uses MLLM to fuse the scene-level and object-level features from multiple CAVs and achieves the best performance in V2V-QA [12]'s perception and planning tasks.

Our proposed V2V-GoT-QA dataset and the prior V2V-QA [12] dataset use the same base dataset V2V4Real [16] and have the same final cooperative planning tasks: our Q9. Suggested Trajectory (Figure 2i) and V2V-QA [12]'s Q5.

TABLE I: Testing performance of V2V-GoT in the planning task of V2V-GoT-QA dataset, in comparison with baseline methods, which are adopted from V2V-LLM [12]. To have a fair comparison to our proposed V2V-GoT, we modified baseline methods to also take perception features at the current and the previous timesteps as visual input. L2: L2 distance error. CR: Collision rate. Comm: Communication cost. In each column, the **best** results are in boldface, and the second-best results are in underline.

Method	L2 (m) ↓				CR (%) ↓				Comm(MB) ↓
	1s	2s	3s	Average	1s	2s	3s	Average	
<i>No Fusion</i>	3.47	5.79	8.26	5.84	1.48	4.24	7.72	4.48	<b>0</b>
<i>Early Fusion</i>	3.48	5.61	7.82	5.63	1.16	3.51	5.66	3.44	1.9208
<i>Intermediate Fusion</i>									
AttFuse [1]	3.65	6.21	8.75	6.20	1.19	4.41	6.38	3.99	<u>0.4008</u>
V2X-ViT [2]	3.46	5.80	8.19	5.81	1.45	4.24	6.59	4.09	<u>0.4008</u>
CoBEVT [3]	3.38	5.42	7.46	5.42	1.31	4.41	5.75	3.82	<u>0.4008</u>
<i>LLM Fusion</i>									
V2V-LLM [12]	2.90	4.91	6.98	4.93	0.75	2.87	4.93	2.85	0.4068
V2V-GoT (Ours)	<b>1.65</b>	<b>2.63</b>	<b>3.59</b>	<b>2.62</b>	<b>0.12</b>	<b>1.92</b>	<b>3.45</b>	<b>1.83</b>	0.4068

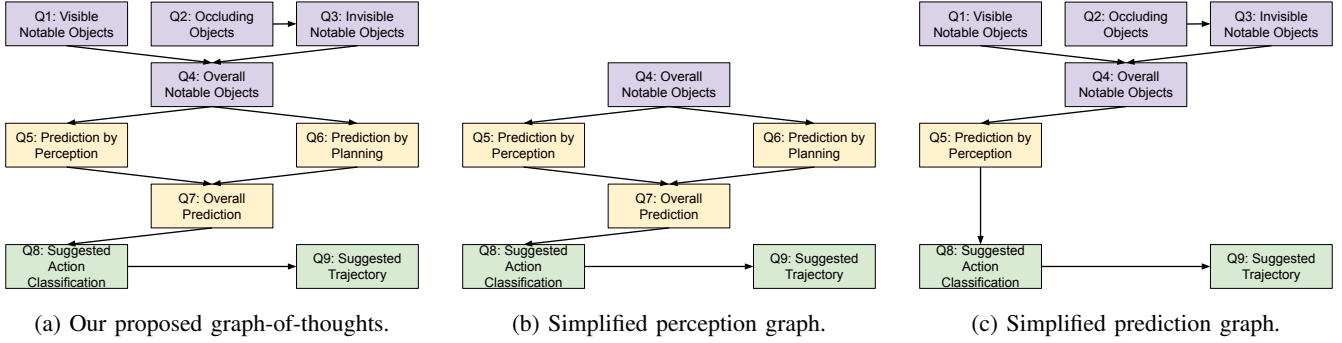


Fig. 4: Different graph-of-thoughts structures for cooperative autonomous driving. The QA types include Perception (Q1 - Q4), Prediction (Q5 - Q7), and Planning (Q8 - Q9). If two nodes are connected by a directed edge, the answer of the parent node QA is used as the input context of the child node QA.

TABLE II: Testing performance of V2V-GoT in all question-answering tasks of V2V-GoT-QA dataset, in comparison with different inference graphs. F1: F1 score. L2: L2 distance error. L1: L1 distance error. CR: Collision rate. Comm: Communication cost. In each column, the **best** results are in boldface.

Method	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9		Comm(MB) ↓
	F1 ↑	F1 ↑	F1 ↑	F1 ↑	L2 (m) ↓	Accuracy	L2 (m) ↓	L1 ↓	L2 (m) ↓	CR (%) ↓	
Simplified Perception	-	-	-	58.0	8.18	82.0	7.66	0.0958	2.84	1.89	0.4068
Simplified Prediction	52.5	30.1	44.0	<b>60.8</b>	<b>8.05</b>	-	-	0.0900	3.24	2.17	0.4068
V2V-GoT (Ours)	52.5	30.1	44.0	<b>60.8</b>	<b>8.05</b>	<b>87.4</b>	<b>7.61</b>	<b>0.0876</b>	<b>2.62</b>	<b>1.83</b>	0.4068

Therefore, we can directly compare our method with the baseline methods on the final planning performance. Note that prior work V2V-LLM [12] only takes the perception features at the current timestep as the visual input. To have a fair comparison with our V2V-GoT, we modify their baseline models to also take the perception features at both of the current and previous timesteps as visual input, following our model architecture shown in Figure 3.

### B. Quantitative Results

Table I shows the testing performance of V2V-GoT in the planning task of V2V-GoT-QA in comparison with baseline methods. Our newly proposed V2V-GoT is seen to achieve the best final planning performance with the lowest L2 errors and collision rates compared to all baselines with different fusion approaches. In particular, the results indicate that the introduction of graph-of-thoughts reasoning designed for cooperative autonomous driving indeed boosts the performance in cooperative autonomous driving planning tasks.

### C. Ablation Study

To further verify the impacts of our newly proposed ideas of **occlusion-aware perception** and **planning-aware prediction** in our graph-of-thoughts reasoning framework, as illustrated in Figures 1 and 4a, we perform an ablation study on two additional graph structures: a *simplified perception graph* shown in Figure 4b and a *simplified prediction graph* shown in Figure 4c. The ablation results are given in Table II.

1) *Simplified Perception Graph*: To verify the impact of **occlusion-aware perception**, we first create additional training data samples of Q4. Overall Notable Objects (Figure 2d) but without context. Then we train another model with newly generated data samples and the existing training dataset, and perform testing inference following the *simplified perception graph* Figure 4b. From Table II, we can see that the *simplified perception graph* results in worse perception performance in Q4 and worse subsequent prediction and planning performance in Q7 and Q9, compared to our proposed V2V-GoT. This result indicates the importance of our occlusion-aware

perception design in our final proposed graph-of-thoughts.

2) *Simplified Prediction Graph*: To verify the impact of **planning-aware prediction**, we use the same V2V-GoT model but run testing inference on the *simplified prediction graph* (Figure 4c). From Table II, we can see that this *simplified prediction graph* approach also results in poorer planning performance than our proposed V2V-GoT, due to the poorer prediction result, indicating the importance of our planning-aware prediction design. Furthermore, note that use of either simplified graph individually already offers performance improvement over the baseline V2V-LLM [12] (compare respective CR(%) values in Tables I and II).

## VI. COMMUNICATION COST

The communication cost of V2V-GoT is the same as the cost of the prior work V2V-LLM [12]. Although V2V-LLM [12] only uses the perception features at the current timestep and V2V-GoT uses the features at the current and the previous timesteps, the same perception features only need to be transferred from a CAV to the MLLM once. The MLLM can save and reuse the perception features received in the current timestep for multiple QAs of the graph-of-thoughts in the current and the next timesteps. The texts of the intermediate questions and answers do not need to be transferred between CAVs and the MLLM unless requested. The MLLM can perform the inference with the graph-of-thoughts reasoning and send the final planning answer to the CAVs in the end. Therefore, the overall communication cost of V2V-GoT is the same as prior work V2V-LLM [12].

## VII. QUALITATIVE RESULTS

Figures 5, 6, and 7 show the qualitative testing results of our proposed V2V-GoT in Q4. Overall Notable Objects (Figure 2d), Q7. Overall Prediction (Figure 2g), and Q9. Suggested Trajectory (Figure 2i), respectively. The input context information of each question comes from the model inference output of the associated parent question(s). In general, we can observe that our method generates perception, prediction, and planning output results close to ground-truth answers. More qualitative results can be found in the accompanying video.

## VIII. CONCLUSION

In this work, we propose a novel graph-of-thoughts reasoning framework for MLLM-based cooperative autonomous driving. Our proposed graph-of-thoughts includes occlusion-aware perception and planning-aware prediction, which are designed to take advantage of V2V information sharing in cooperative driving scenarios and the multimodal understanding ability of MLLMs. To verify the effectiveness of our proposed ideas, we curate the V2V-GoT-QA dataset and develop the V2V-GoT model. Our experimental results show that our proposed method outperforms all other baseline methods. Moreover, the ablation study further indicates that our innovative designs in the graph-of-thoughts improve overall cooperative perception, prediction, and planning performance. To facilitate open-source research, we have publicly released our V2V-GoT-QA dataset and V2V-GoT code.

### Q4: All Notable Objects



**Q:** I am **CAV\_EGO** at **(0,0,0.0)**. Is there anything I need to be aware of if my planned future trajectory is **[(5.5,0.2),(10.9,0.3),(16.4,0.4),(21.9,0.6),(27.4,0.8),(33.0,0.9)]**?

**Context:** There is no notable object visible to you. There is a car at (33.3,-1.8) invisible to you. There is a car at (39.1,-2.0) invisible to you.

**V2V-GoT:** There is a car at **(33.3,-1.8)**. There is a car at **(39.1,-2.0)**.


**GT:** There is a car at **(35.4,-1.3)**. There is a car at **(40.2,-1.7)**.

Fig. 5: Qualitative testing sample result of V2V-GoT on Q4. Overall Notable Objects (Figure 2d). The context information is from the testing inference output of parent question Q3. Invisible Notable Objects (Figure 2c) and Q1. Visible Notable Objects (Figure 2a). **Magenta x**: current location of the asking CAV. **Magenta curve**: reference trajectory in the question. **Yellow x**: model output. **Green o**: ground-truth answer.

## REFERENCES

- [1] 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 *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [2] 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 (ECCV)*, 2022.
- [3] R. Xu, Z. Tu, H. Xiang, W. Shao, B. Zhou, and J. Ma, "Cobevt: Cooperative bird's eye view semantic segmentation with sparse transformers," in *Conference on Robot Learning (CoRL)*, 2022.
- [4] H.-k. Chiu and S. F. Smith, "Selective communication for cooperative perception in end-to-end autonomous driving," in *IEEE International Conference on Robotics and Automation (ICRA) Workshop*, 2023.
- [5] H.-k. Chiu, C.-Y. Wang, M.-H. Chen, and S. F. Smith, "Probabilistic 3d multi-object cooperative tracking for autonomous driving via differentiable multi-sensor kalman filter," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [6] H. Xiang, R. Xu, and J. Ma, "Hm-vit: Hetero-modal vehicle-to-vehicle cooperative perception with vision transformer," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [7] J. Cui, H. Qiu, D. Chen, P. Stone, and Y. Zhu, "Coopernaut: End-to-end driving with cooperative perception for networked vehicles," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [8] Z. Zhou, S. Z. Zhao, T. Cai, Z. Huang, B. Zhou, and J. Ma, "Turbotrain: Towards efficient and balanced multi-task learning for multi-agent perception and prediction," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025.
- [9] C. Sima, K. Renz, K. Chitta, L. Chen, H. Zhang, C. Xie, P. Luo, A. Geiger, and H. Li, "Drivelm: Driving with graph visual question answering," in *European Conference on Computer Vision (ECCV)*, 2024.
- [10] S. Wang, Z. Yu, X. Jiang, S. Lan, M. Shi, N. Chang, J. Kautz, Y. Li, and J. M. Alvarez, "OmniDrive: A holistic vision-language dataset for autonomous driving with counterfactual reasoning," in *CVPR*, 2025.

## Q7: Overall Prediction



**Q:** I am CAV\_1 at (-50.9,-1.4). Where might those notable objects move in the future?

**Context:** There is a car at (0.0,0.0) moving forward. The predicted future trajectory is [(50.5,1.9)]. There is a car at (-17.0,3.1) moving forward. The predicted future trajectory is [(43.9,5.1)]. There is a car at (-19.8,-7.3) moving forward. The predicted future trajectory is [(-8.5,-12.2)]. There is a car at (-49.2,-1.3) moving forward. The predicted future trajectory CAV\_EGO is at (0.0,0.0) moving forward. Its planned future trajectory is [(38.8,1.3)]. CAV\_EGO is a notable object.


**V2V-GoT:** There is a car at (0.0,0.0) moving forward. The predicted future trajectory is [(38.8,1.3)]. There is a car at (-17.0,3.1) moving forward. The predicted future trajectory is [(43.9,5.1)]. There is a car at (-19.8,-7.3) moving forward. The predicted future trajectory is [(-8.5,-12.2)].

**GT:** There is a car at (0.0,0.0) moving forward. The predicted future trajectory is [(38.8,1.3)]. There is a car at (-14.4,3.0) moving forward. The predicted future trajectory is [(42.1,4.7)]. There is a car at (-16.3,-5.2) moving forward. The predicted future trajectory is [(-17.1,-4.3)].

Fig. 6: Qualitative testing sample result of V2V-GoT on Q7. Overall Prediction (Figure 2g). The context information is from the testing inference output of parent question Q5. Prediction by Perception (Figure 2e) and Q6. Prediction by Planning (Figure 2f). Magenta  $\times$ : current location of the asking CAV. Blue line and  $\times$ : model output of the predicted future trajectories, starting, and ending waypoints of notable objects. Green line and  $\circ$ : ground-truth answer.

- [11] R. Tian, B. Li, X. Weng, Y. Chen, E. Schmerling, Y. Wang, B. Ivanovic, and M. Pavone, "Tokenize the world into object-level knowledge to address long-tail events in autonomous driving," in *Conference on Robot Learning (CoRL)*, 2024.
- [12] H.-k. Chiu, R. Hachiuma, C.-Y. Wang, S. F. Smith, Y.-C. F. Wang, and M.-H. Chen, "V2v-llm: Vehicle-to-vehicle cooperative autonomous driving with multi-modal large language models," <https://arxiv.org/abs/2502.09980>, 2025.
- [13] C. Liu, G. Liu, Z. Wang, J. Yang, and S. Chen, "Colmdriver: Llm-based negotiation benefits cooperative autonomous driving," *arXiv preprint arXiv:2503.08683*, 2025.
- [14] S. Fang, J. Liu, M. Ding, Y. Cui, C. Lv, P. Hang, and J. Sun, "Towards interactive and learnable cooperative driving automation: a large language model-driven decision-making framework," *IEEE Transactions on Vehicular Technology*, pp. 1–12, 2025.
- [15] X. Gao, Y. Wu, R. Wang, C. Liu, Y. Zhou, and Z. Tu, "Langcoop: Collaborative driving with language," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2025, pp. 4235–4246.
- [16] R. Xu, X. Xia, J. Li, H. Li, S. Zhang, Z. Tu, Z. Meng, H. Xiang, X. Dong, R. Song, H. Yu, B. Zhou, and J. Ma, "V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [17] Q. Chen, X. Ma, S. Tang, J. Guo, Q. Yang, and S. Fu, "F-cooper:

## Q9: Suggested Trajectory



**Q:** I am CAV\_1 at (-44.6,-1.0). What is the suggested future trajectory to avoid collision with nearby objects?

**Context:** What are the suggested speed and steering settings to avoid collision with nearby objects? The suggested speed setting is: moderate. The suggested steering setting is: straight.

**V2V-GoT:** The suggested future trajectory is [(-37.9,-0.8),(-31.3,-0.5),(-24.0,-0.2),(-17.8,0.1),(-10.8,0.3),(-3.9,0.5)].

**GT:** The suggested future trajectory is [(-36.4,-0.8),(-28.3,-0.6),(-20.2,-0.3),(-13.1,-0.2),(-5.8,0.0),(0.3,0.2)].

Fig. 7: Qualitative testing sample result of V2V-GoT on Q9. Suggested Trajectory (Figure 2i). The context information is from the testing inference output of parent question Q8. Suggested Action Classification (Figure 2h). Magenta  $\times$ : current location of the asking CAV. Blue curve and  $\times$ : model output of the suggested future trajectory and the ending waypoint. Green curve and  $\circ$ : ground-truth answer.

Feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds," in *ACM/IEEE Symposium on Edge Computing (SEC)*, 2019.

- [18] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, J. Tu, and R. Urtasun, "V2vnet: Vehicle-to-vehicle communication for joint perception and prediction," in *European Conference on Computer Vision (ECCV)*, 2020.
- [19] H. Su, S. Arakawa, and M. Murata, "Cooperative 3d multi-object tracking for connected and automated vehicles with complementary data association," in *2024 IEEE Intelligent Vehicles Symposium (IV)*, 2024, pp. 285–291.
- [20] J. Mao, Y. Qian, J. Ye, H. Zhao, and Y. Wang, "Gpt-driver: Learning to drive with gpt," in *Advances in Neural Information Processing Systems (NeurIPS) Workshop (Foundation Models for Decision Making)*, 2023.
- [21] J. Mao, J. Ye, Y. Qian, M. Pavone, and Y. Wang, "A language agent for autonomous driving," in *Conference On Language Modeling (COLM)*, 2024.
- [22] T. Qian, J. Chen, L. Zhuo, Y. Jiao, and Y.-G. Jiang, "Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2024.
- [23] D. Wu, W. Han, T. Wang, Y. Liu, X. Zhang, and J. Shen, "Language prompt for autonomous driving," *arXiv preprint*, 2023.
- [24] A.-M. Marcu, L. Chen, J. Hünermann, A. Karnsund, B. Hanotte, P. Chidananda, S. Nair, V. Badrinarayanan, A. Kendall, J. Shotton, et al., "Lingoqa: Visual question answering for autonomous driving," in *European Conference on Computer Vision (ECCV)*, 2024.
- [25] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [26] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [27] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *International Conference on Learning Representations (ICLR)*, 2022.