

Large Language Models Powered Context-aware Motion Prediction in Autonomous Driving

Xiaoji Zheng, Lixiu Wu, Zhijie Yan, Yuanrong Tang, Hao Zhao
Chen Zhong[✉], Bokui Chen[✉], and Jiangtao Gong[✉]

Abstract—Motion prediction is among the most fundamental tasks in autonomous driving. Traditional methods of motion forecasting primarily encode vector information of maps and historical trajectory data of traffic participants, lacking a comprehensive understanding of overall traffic semantics, which in turn affects the performance of prediction tasks. In this paper, we utilized Large Language Models (LLMs) to enhance the global traffic context understanding for motion prediction tasks. We first conducted systematic prompt engineering, visualizing complex traffic environments and historical trajectory information of traffic participants into image prompts—Transportation Context Map (TC-Map), accompanied by corresponding text prompts. Through this approach, we obtained rich traffic context information from the LLM. By integrating this information into the motion prediction model, we demonstrate that such context can enhance the accuracy of motion predictions. Furthermore, considering the cost associated with LLMs, we propose a cost-effective deployment strategy: enhancing the accuracy of motion prediction tasks at scale with 0.7% LLM-augmented datasets. Our research offers valuable insights into enhancing the understanding of traffic scenes of LLMs and the motion prediction performance of autonomous driving. The source code is available at <https://github.com/AIR-DISCOVER/LLM-Augmented-MTR> and <https://aistudio.baidu.com/projectdetail/7809548>.

I. INTRODUCTION

Motion prediction is one of the most important tasks in the field of autonomous driving [1]–[3], which predicts the motion statuses of nearby agents by jointly considering nearby agents and road maps. This information will assist the decision module in making a more robust and safer driving decision. Therefore, this field already features a plethora of datasets and public competitions. For instance, since 2021, Waymo has been organizing competitions in motion prediction¹, attracting models and algorithms that have won championships, such as MTR [1], MTR++ [2], and MGTR [4]. There are three classes of motion prediction methods: goal-based methods [5], [6], direct-regression methods [7], [8], and methods that take the best of both [1], [2]. Although methods that take the best of both can balance results and computing resources greatly, they still encode information of agents and maps first and then decode to obtain final results. The methods above all lack a comprehensive understanding of overall traffic semantics.

The authors are with the Institute for AI Industry Research, Tsinghua University, Beijing, China. Corresponding Email: myheimu@gmail.com, chenbk@tsinghua.edu.cn, gongjiangtao@air.tsinghua.edu.cn

¹<https://waymo.com/open/challenges/2021/motion-prediction/>

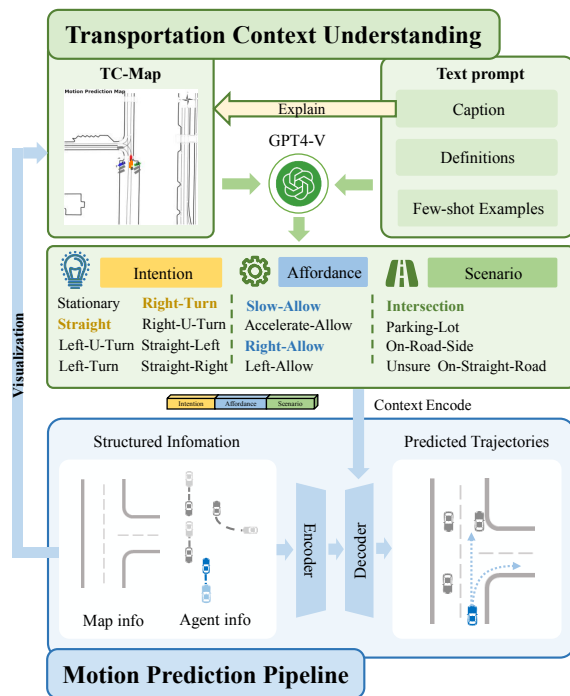


Fig. 1. Context-aware Motion Prediction Based on LLMs. We first visualize the structured information of one scenario in the motion prediction dataset, GPT4-V then understand the scenario via a visualized image and well-designed prompt. Finally, GPT4-V outputs transportation context information. This information will be used to augment traditional motion prediction algorithms.

Large Language Model (LLM) is popular in the field of autonomous driving, as not only the Language Model but also the World Model with common sense [9]–[11]. After the release of OpenAI’s GPT series of large language models, many works that use GPT to assist driving have been emerging [12]–[16]. These efforts leverage the inferential capabilities of large language models or through experimentation, aiding autonomous driving algorithms in better perception [17] and planning [12], [14], [15], and even facilitating end-to-end decision-making [13], [18]. However, teaching LLM to understand the interactive context among various transportation participants in diverse transportation scenarios is still a challenging task. Although many have applied vision-language models to traffic scenarios, their performance heavily relies on datasets annotated from a first-person perspective [19]–[21]. For abstract, information-overwhelming, and bird’s-eye view (BEV) traffic scenarios,

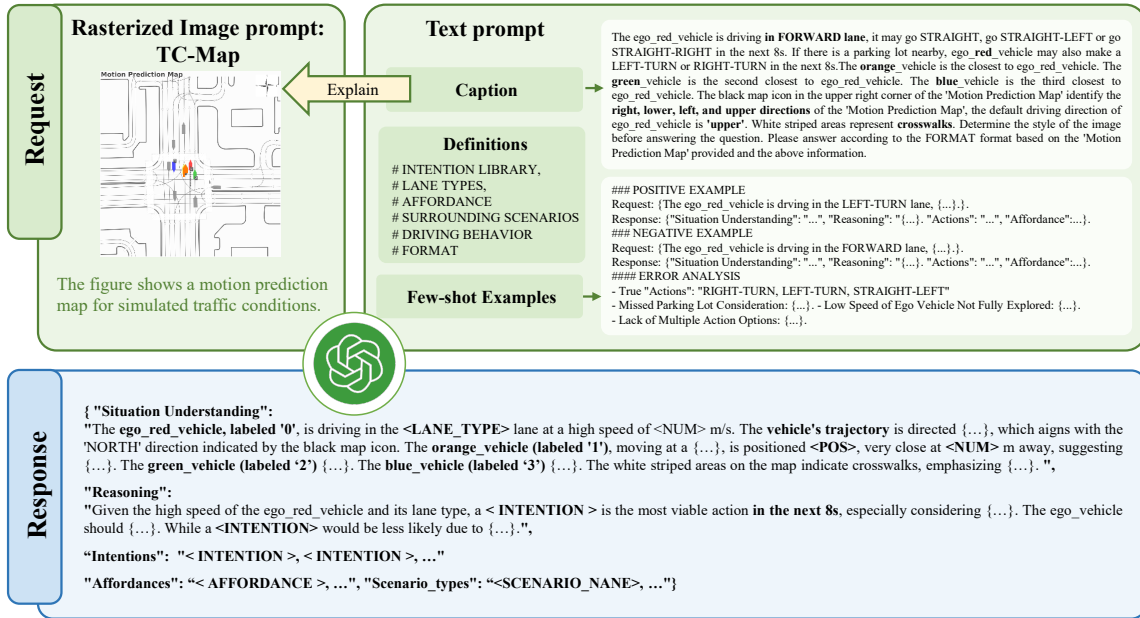


Fig. 2. Details of Transportation Context Generation Prompt.

there is a lack of relevant datasets and understanding methodologies.

In this paper, we propose a new method that enhances the traffic context understanding of motion prediction models using LLMs to make more accurate predictions. To enable an unfine-tuned LLM with vision (e.g. GPT4-V) to understand traffic context and output the necessary context information for motion prediction, we visualize the vector map data and historical data of traffic participants as a Transportation Context Map (TC-Map) to serve as an image prompt and propose a corresponding text prompt design. After successfully obtaining traffic context information generated by the LLM, we integrated this information into a classical motion prediction algorithm—MTR [1]. The results indicate that this context information effectively improves the accuracy of motion prediction. Additionally, considering the cost of LLMs, we also propose a cost-effective deployment strategy: by utilizing 0.7% LLM-enhanced datasets, we can empower motion prediction performance at scale.

Our contributions are as follows:

- We systematically designed and conducted prompt engineering to enable an unfine-tuned GPT4-V to comprehend complex traffic scenarios involving multiple traffic participants and to output context information such as intention, affordance, and scenario;
- We introduced a novel approach that combines the context information outputted by GPT4-V with the classical motion prediction pipeline [1], and we verified that this method enhances the effectiveness of motion prediction;
- We proposed and validated a deployment strategy based on a dataset with a limited amount of LLM-generative context, which reduced the deployment cost of this method.

II. METHOD

A. Get Transportation Context Information from GPT4-V

Translating images directly into motion poses is significantly challenging [22]. We leverage vision-enhanced LLMs to derive transportation context, identifying the agent of interest in the TC-Map as the ego-agent, categorized into vehicles (V), pedestrians (P), and cyclists (C). Illustrated in Fig. 2, the TC-Map designed based on the Waymo Open Motion Dataset (WOMD) [23] serves as the rasterized image prompt, and together with the text prompt forms the Transportation Context Generation Prompt (TCGP). The caption of the text prompt is obtained by filling in the template with code. Utilizing GPT4-V with this combined input facilitates the extraction of comprehensive context including situation understanding, reasoning, intentions of ego-agent, affordances [24], and scenario types. Additionally, six prompt design suggestions have been summarized for input prompts as shown in Fig. 3.

1) *TC-Map Design Suggestions*: Informed by prompt design suggestion 1 and 2. Suggestion 1 minimizes information overload. Drawing inspiration from the MTR Map Collection module [1], we introduced TC-Map Crop to focus the ego-agent's attention solely on relevant surrounding scenes. Initial GPT4-V tests revealed that dense data, like 11 frames of historical trajectory in JSON format, leading to unclear intentions. Thus, we customized TC-Map dimensions: 120m x 120m for vehicles, 80m x 80m for pedestrians, and 60m x 60m for cyclists, respectively, to filter out extraneous influences on intention discernment. Furthermore, to address the varying orientations of the ego-agent across different scenarios, suggestion 2 proposes normalizing the heading of the ego-agent to "North" with a common map guide icon to reduce confusion in direction identification.

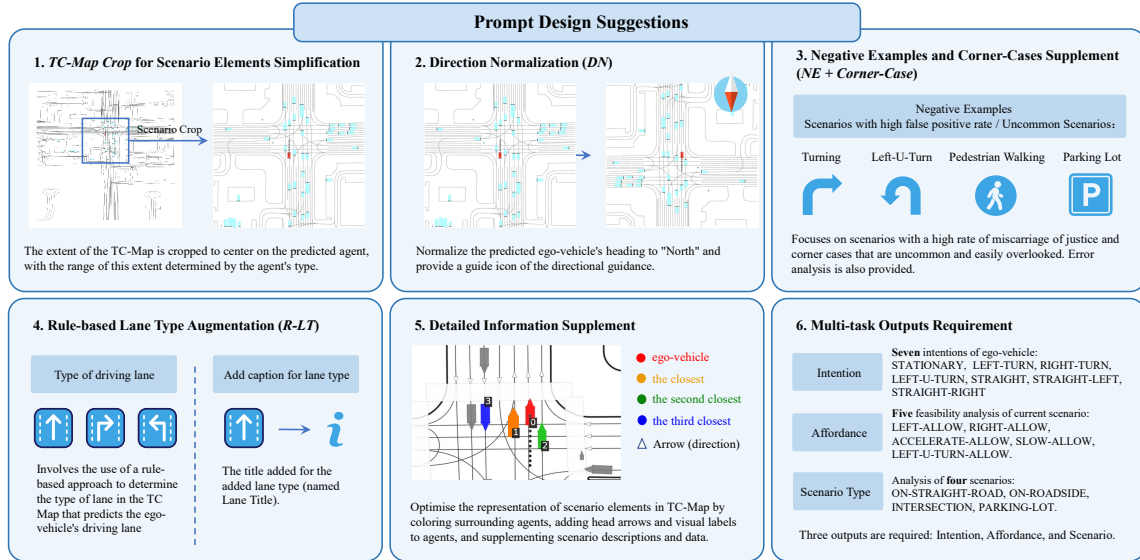


Fig. 3. Six Prompt Design Guideline for GPT4-V Understanding Motion Predication Context.

2) *Text Prompt Design Suggestions*: Addressing design suggestions 3 and 4, we introduce a structured approach to example creation. Suggestion 3 involves categorizing examples into positive and negative instances. Positive examples ensure a diverse and random generation of intentions. Negative examples are bad cases selected from GPT4-V's responses during the quantitative experiment of prompts, focusing on scenarios with high misjudgment rates and uncommon, easily overlooked corner cases. Additional error analysis is provided and placed closer to the end of the negative example. Suggestion 4 consists of two sentences, the first sentence contains the ego-agent's driving lane type in TC-Map determined by a rule-based algorithm, and the second sentence provides explanations for the added lane type. They aids GPT4-V in ruling out implausible intentions by assuming compliance with traffic regulations.

3) *Detailed Information Supplement*: Building on prompt design suggestion 5, we visualize more scenario elements using WOMD's comprehensive scenario details, including road edges and crosswalks. For agents, we optimize the visualization of different scenario elements in the TC-Map by adding indicative head arrows and visual labels. Moreover, this suggestion includes integrating dynamic elements such as the ego-agent's speed and the relative speeds, positions, and distances of nearby agents, enriching the context for more accurate motion prediction.

4) *Multi-task Outputs Requirement*: Matching to prompt design suggestion 6, we delineated seven intention categories for the ego-agent inspired by the official evaluation tool of Waymo [23]. *STRAIGHT-LEFT* and *STRAIGHT-RIGHT* anticipate potential lane changes, indicating the complexity of intentions. Predicting these intentions often requires GPT4-V to utilize rasterized images, the ego-agent's position, heading, and speed, the state of surrounding vehicles, and the direction of nearby lane lines predicting these two intentions is highly challenging. Additionally, our approach includes the creation

of five affordance types and four scenario types based on an understanding of comprehensive scenario analyses. These three types of information collectively provide rich, human-like transportation context information to enhance the accuracy of the motion prediction task.

B. Augment Motion Prediction Models via Transportation Context Information

1) *Motion Transformer*: We selected the Motion Transformer (MTR) [1], the state-of-the-art model in the WOMD Challenge 2022, as our foundational model. Leveraging the Transformer architecture [25], MTR encodes comprehensive scenario data—including agent movements and map features—before decoding this information to predict motion trajectories. This process underpins our model's ability to accurately anticipate future positions.

2) *Integrate with Transportation Context Information*: To enrich our motion prediction model with transportation context, we categorized context information into three types: intention, affordance, and scenario, each represented as natural language descriptors (e.g. [*Straight*, *Right-Turn*], [*Slow-Allow*, *Right-Allow*], and [*Intersection*] in Fig. 1). These descriptors are encoded as one-hot vectors, allowing for multiple active elements to reflect the non-exclusive nature of these descriptors. Notably, affordances and scenarios are directly encoded, while intentions receive weighted encoding to prioritize them.

More specifically, for affordance and scenario information generated by LLM, we simply encode the description words into one-hot vectors:

$$A = \text{OneHotEncode}(\text{AffordanceList}), \quad (1)$$

$$S = \text{OneHotEncode}(\text{ScenarioList}), \quad (2)$$

where $A \in \mathbb{R}^8$ and $S \in \mathbb{R}^4$ (see Fig. 1). LLM gives all possible intentions and the order of different description words stands for different possibilities. So, when we encode

TABLE I
EFFECTS OF DIFFERENT COMPONENTS IN TCGP

<i>Sugg 1</i>		<i>Sugg 2</i>		<i>Sugg 3</i>		<i>Sugg 4</i>		<i>Sugg 5</i>		ACC(<i>Ist-I</i>)	ACC	ACC(<i>S,S-L,S-R</i>)
TC-Map	Crop	DN	NE+Corner-Case	R-LT	TCGP-Opt	SPD-Add						
✗	✗	✗	✗	✗	✗	✗	0.3551	0.5140	0.5794			
✓	✗	✗	✗	✗	✗	✗	0.4579	0.5981	0.7383			
✓	✓	✗	✗	✗	✗	✗	0.4205	0.6635	0.7757			
✓	✓	✓	✗	✗	✗	✗	0.4112	0.6822	0.8037			
✓	✓	✓	✓	✓	✗	✗	0.5607	0.8317	0.9252			
✓	✓	✓	✓	✓	✓	✗	0.5514	0.8691	0.9252			
✓	✓	✓	✓	✓	✓	✓	0.5794	0.8691	0.9345			

Sugg: Prompt design suggestion. **TCGP-Opt**: Coloring agents surrounding the ego-agent in different colors to distinguish them easily, optimizing the visualization of scenario elements in the TC-Map and supplementing the corresponding scenario description in the text prompt. **SPD-Add**: Adding information about the speed of the ego-agent, the speeds of surrounding agents, and their relative positions and distances to the ego-agent in the text prompt. **ACC**: The accuracy of intentions output by GPT4-V. When GPT4-V outputs multiple intentions, if at least one intention matches the ground truth without additional clarification, it is considered correct. (*Ist-I*): Taking only the first intention from the list of intentions output by GPT4-V. (*S,S-L,S-R*) means consolidating the three intentions of *STRAIGHT*, *STRAIGHT-LEFT*, and *STRAIGHT-RIGHT* into a single *STRAIGHT* category.

intention information into a one-hot vector, we set different weights for different intentions in the list:

$$I = \text{WeightedOneHotEncode}(\text{IntentionList}). \quad (3)$$

The weight for i -th word in the list is:

$$\text{weight}_i = \text{length}(\text{IntentionInfoList}) - i, \quad (4)$$

where $I \in \mathbb{R}^5$ (see Fig. 1), and each intention's weight inversely correlates with its position in the list, emphasizing more probable intentions. So that the first word in the list has the maximal weight, and the last word holds the minimal weight.

To improve the performance of motion prediction and reduce the cost of computing resources, cluster-based intention points were provided in the process of decoding as prior knowledge [1]. Inspired by this, we also choose to integrate the transportation context information generated by LLM in the process of decoding.

The query content $Q \in \mathbb{R}^{K \times D}$ includes all scenario information in the data, where K stands for the number of trajectories that need to be predicted and D is the dimension of the feature. We first concatenate different types of context information and align the dimensions of traffic context information with query content, then integrate the context information generated by LLM via cross-attention after query content has been initialized.

$$TC = \text{MLP}([I, A, S]), \quad (5)$$

$$TC = \text{RepeatFirstDim}(TC, K), \quad (6)$$

$$Q_{tc} = \text{CrossAttn}(q = Q_0, k = TC, v = TC), \quad (7)$$

where $TC \in \mathbb{R}^D$ in Eq. 5 after multi-layer perceptron, and $TC \in \mathbb{R}^{K \times D}$ after repeating K times in Eq. 6. When conducting cross-attention, we use initialized query content Q_0 as query, and TC for both key and value in Eq. 7.

The following pipeline keeps all the same with MTR.




Fig. 4. Confusion Matrix of Intention Generated from GPT4-V

TABLE II
DATA AND INTENTION ACCURACY OF THREE EGO-AGENT TYPES

method	type	ACC(<i>Ist-I</i>)	ACC	ACC(<i>S,S-L,S-R</i>)
MTR	V	0.8097	0.9442	0.9635
	P	0.6906	0.9495	0.9564
	C	0.7226	0.8822	0.9321
	AVG	0.7644	0.9321	0.9532
GPT4-V	V	0.6224	0.8482	0.9080
	P	0.6222	0.8038	0.8133
	C	0.5838	0.7857	0.9030
	AVG	0.6145	0.8259	0.8865

III. EXPERIMENT

A. TCGP Generated Transportation Context Information

Data distribution. The data set used to generate transportation context information was randomly selected from the WOMD validation set, consisting of approximately 15,000 TC-Map samples distributed according to a vehicle (V) to pedestrian (P) to cyclist (C) ratio of 9:3:3. We create a confusion matrix for the results of intention in Tab. II as shown in Fig. 4. For simplification, the categories *STRAIGHT-LEFT* and *STRAIGHT-RIGHT* were amalgamated into a single *STRAIGHT* category. The matrix's rows denote the actual intentions, while its columns represent

the intentions as predicted by GPT4-V. Observing the main diagonal of the confusion matrix, it’s not hard to see that *STRAIGHT* has the highest prediction accuracy, followed by *LEFT-TURN* and *RIGHT-TURN* predictions. The comparatively lower accuracy observed for *STATIONARY* intentions stems from its correct identification only at a zero speed of the ego-agent, despite being labeled as *STATIONARY* at minimal speeds. Ego-agents intended for *LEFT-U-TURN* are frequently classified as *LEFT-TURN*, a consequence of the WOMB map’s left-turn lanes occasionally accommodating straight movements, thus increasing the likelihood of *LEFT-U-TURN* and *LEFT-TURN* being predicted as *STRAIGHT*. Similarly, *RIGHT-U-TURN* ego-agents are often mistaken as *STATIONARY*, and unexpectedly, they are frequently classified as *STATIONARY*. This anomaly primarily arises because intentions like *RIGHT-U-TURN* are most common among pedestrians, whose speed and walking direction can change abruptly, reflecting the challenge of precisely capturing the dynamic intentions of pedestrians. Despite GPT4-V’s intention prediction accuracy being lower than that of the MTR model, the integration of GPT4-V’s output still contributes valuable context for motion prediction, enhancing overall prediction accuracy, as evidenced in Tab. III.

Ablation Study. We perform ablation studies on TCGP to understand the effectiveness of each component of our prompt in Tab. I. The data comprising 107 TC-Maps samples for the ablation experiment were randomly sourced from the validation set of WOMB. We have the following observations: 1) TCGP achieves optimal performance when fully equipped with all its components, with every module playing a role in predicting intentions. 2) TC-Map Crop enables GPT4-V to allocate more attention to the areas that should be focused on when predicting the future intentions of the ego-agent, significantly enhancing the accuracy of intention prediction. 3) R-LT is crucial because it provides the type of lane in which the ego-agent is located, and the type of lane often determines the future direction of the ego-agent. These observations underscore the multifaceted contributions of TCGP’s components to the overall effectiveness of intention prediction, highlighting the importance of a holistic approach in optimizing predictive performance.

B. LLM Augmented Motion Prediction

Dataset and metrics. We evaluated our LLM-augmented motion prediction methods using the extensive WOMB dataset. Task requirements include predicting 6 future motion trajectories over 8 seconds, based on 1 second of historical data. The official WOMB evaluation tool calculates key metrics, notably the mean Average Precision (mAP), which serve as crucial benchmarks in the official leaderboard .

Data generation. To manage the high cost of generating transportation context information for WOMB’s 2 million agents (approximately \$0.1 per agent, totaling nearly \$200,000), we implemented a cost-effective strategy. Initially, LLM-generated transportation context for a small subset of WOMB was expanded to the remainder via a nearest neighbor algorithm, utilizing Euclidean distance for

similarity assessment (see Alg. 1). This approach efficiently scales LLM augmentation while maintaining data integrity.

Algorithm 1 Generate Scaled LLM Augmented Training Data with Minimal LLM Augmented Data

Input: T : the overall dataset without LLM augmented;
Encoder: scenario info encoder; *LLM*: large language model.

Output: T^* : the overall dataset with LLM augmented.

- 1: initial $T_1, T_2 \leftarrow \text{split}(T)$, where $|T_1| \gg |T_2|$
 - 2: $TC_1 \leftarrow []$, where TC_1 is Transportation Context info
 - 3: $TC_2 \leftarrow \text{LLM}(T_2)$
 - 4: $F_1 \leftarrow \text{Encoder}(T_1)$, where F_1 is feature vector list
 - 5: $F_2 \leftarrow \text{Encoder}(T_2)$
 - 6: **for** agent i in T_1 **do**
 - 7: $f_i \leftarrow F_1[i]$, where f_i is a feature vector
 - 8: $j \leftarrow \text{NearestNeighbor}(f_i, F_2)$
 - 9: $TC_1[i] = TC_2[j]$
 - 10: **end for**
 - 11: $T_1^* \leftarrow \text{Concatenate}(T_1, TC_1)$
 - 12: $T_2^* \leftarrow \text{Concatenate}(T_2, TC_2)$
 - 13: $T^* \leftarrow \text{Concatenate}(T_1^*, T_2^*)$
-

Performance comparison. We compared our LLM augmented motion prediction model with MTR on the validation set of WOMB. We trained both models on the 20% WOMB dataset. Our main results are in Table III. Our LLM-augmented model exhibited superior performance over the MTR model across all agent types (vehicles, pedestrians, cyclists) on the WOMB validation set. Notably, the average mAP saw an enhancement of 0.95%, underscoring the augmented model’s effectiveness.

TABLE III
 PERFORMANCE ON THE VALIDATION SET OF WOMB

method	type	mAP↑	minADE↓	minFDE↓	MR↓
MTR	V	0.3862	0.8257	1.6789	0.1809
	P	0.3587	0.3825	0.8088	0.0935
	C	0.2848	0.7985	1.6380	0.2212
	AVG	0.3432	0.6689	1.3752	0.1652
+LLM	V	0.3954	0.8147	1.6205	0.1751
	P	0.3754	0.3830	0.8070	0.0934
	C	0.2924	0.8102	1.6464	0.2306
	AVG	0.3527	0.6693	1.3580	0.1664

The best scores are expressed in **bold**.

Ablation Study. This study evaluates the contribution of each component within the transportation context generated by the LLM toward enhancing motion prediction accuracy. An ablation experiment was conducted by randomly selecting 5% of scenarios (approximately 24,000 scenarios) from the WOMB training set to train a model. The evaluation involved comparing the baseline model against the LLM-augmented model on the minimal part of the WOMB validation set with traffic context information generated by LLM. The results are presented in Table IV. I, A, and S are abbreviations of intention, affordance, and scenario, respectively.

V-mAP, P-mAP, and C-mAP mean mAP scores for vehicle, pedestrian, and cyclist, respectively. The results highlight the incremental value added by incorporating various types of transportation context information. Incorporation of all context types resulted in the highest mAP scores, marking a 1.67% increase over the baseline MTR model, thereby validating the efficacy of LLM-augmented context in enhancing motion prediction.

TABLE IV
ABLATION EXPERIMENT OF
LLM AUGMENTED MOTION PREDICTION

method	I	A	S	mAP \uparrow	V-mAP \uparrow	P-mAP \uparrow	C-mAP \uparrow
MTR	-	-	-	0.2970	0.3300	0.3507	0.2102
-	✓	-	-	0.3066	0.3377	0.3400	0.2422
-	-	✓	-	0.2995	0.3386	0.3277	0.2322
-	-	-	✓	0.3059	0.3341	0.3564	0.2272
-	✓	✓	-	0.3030	0.3294	0.3458	0.2337
-	✓	-	✓	0.3134	0.3318	0.3692	0.2392
-	-	✓	✓	0.3117	0.3389	0.3554	0.2409
+LLM	✓	✓	✓	0.3137	0.3385	0.3549	0.2476

The highest mAP score for each type of agent and the average mAP score is expressed in **bold**.

IV. CONCLUSIONS

In this paper, through systematic prompt engineering, we utilized the common sense and reasoning abilities of LLMs to extract human-like global context information from complex traffic scene motion predictions. By integrating these contexts into the traditional motion prediction pipeline, we enhanced the accuracy of trajectory prediction. We also proposed a cost-effective deployment strategy and verified its effectiveness.

Through meticulous prompt design, we achieved an impressive accuracy rate, underscoring the remarkable capacity of LLMs to comprehend complex and detailed transportation scenarios. To our knowledge, our work is pioneering in incorporating BEV-like TC-Maps into LLM prompts, offering a novel perspective for LLMs to interpret driving scenarios. Our experimental findings suggest that integrating transportation context information significantly enhances motion prediction capabilities. It is noteworthy that enabling LLMs to understand complex scene information remains challenging. Significant opportunities exist for exploring the integration of rich contextual information with motion prediction models. Our preliminary results suggest considerable potential for future development in this research area.

REFERENCES

- [1] S. Shi, L. Jiang, D. Dai, and B. Schiele, "Motion transformer with global intention localization and local movement refinement," *Advances in Neural Information Processing Systems*, vol. 35, pp. 6531–6543, 2022.
- [2] S. Shi, L. Jiang, D. Dai, and S. Bernt, "Mtr++: Multi-agent motion prediction with symmetric scene modeling and guided intention querying," 2024.
- [3] T. Zhao, Y. Xu, M. Monfort, W. Choi, C. Baker, Y. Zhao, Y. Wang, and Y. N. Wu, "Multi-agent tensor fusion for contextual trajectory prediction," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 12 126–12 134.
- [4] Y. Gan, H. Xiao, Y. Zhao, E. Zhang, Z. Huang, X. Ye, and L. Ge, "Mgtr: Multi-granular transformer for motion prediction with lidar," *arXiv preprint arXiv:2312.02409*, 2023.

- [5] J. Gu, C. Sun, and H. Zhao, "Densent: End-to-end trajectory prediction from dense goal sets," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 303–15 312.
- [6] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid, *et al.*, "Tnt: Target-driven trajectory prediction," in *Conference on Robot Learning*. PMLR, 2021, pp. 895–904.
- [7] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal, *et al.*, "Scene transformer: A unified architecture for predicting multiple agent trajectories," *arXiv preprint arXiv:2106.08417*, 2021.
- [8] B. Varadarajan, A. Hefny, A. Srivastava, K. S. Refaat, N. Nayakanti, A. Corrmann, K. Chen, B. Douillard, C. P. Lam, D. Anguelov, *et al.*, "Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 7814–7821.
- [9] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 824–24 837, 2022.
- [10] X. Wang, J. Wei, D. Schuurmans, Q. Le, E. Chi, S. Narang, A. Chowdhery, and D. Zhou, "Self-consistency improves chain of thought reasoning in language models," *arXiv preprint arXiv:2203.11171*, 2022.
- [11] S. Yao, D. Yu, J. Zhao, I. Shafran, T. L. Griffiths, Y. Cao, and K. Narasimhan, "Tree of thoughts: Deliberate problem solving with large language models," *arXiv preprint arXiv:2305.10601*, 2023.
- [12] J. Mao, Y. Qian, H. Zhao, and Y. Wang, "Gpt-driver: Learning to drive with gpt," *arXiv preprint arXiv:2310.01415*, 2023.
- [13] Z. Xu, Y. Zhang, E. Xie, Z. Zhao, Y. Guo, K. K. Wong, Z. Li, and H. Zhao, "Drivegpt4: Interpretable end-to-end autonomous driving via large language model," *arXiv preprint arXiv:2310.01412*, 2023.
- [14] W. Wang, J. Xie, C. Hu, H. Zou, J. Fan, W. Tong, Y. Wen, S. Wu, H. Deng, Z. Li, *et al.*, "Drivemlm: Aligning multi-modal large language models with behavioral planning states for autonomous driving," *arXiv preprint arXiv:2312.09245*, 2023.
- [15] J. Mao, J. Ye, Y. Qian, M. Pavone, and Y. Wang, "A language agent for autonomous driving," *arXiv preprint arXiv:2311.10813*, 2023.
- [16] Y. Jin, X. Shen, H. Peng, X. Liu, J. Qin, J. Li, J. Xie, P. Gao, G. Zhou, and J. Gong, "Surrealdriver: Designing generative driver agent simulation framework in urban contexts based on large language model," *arXiv preprint arXiv:2309.13193*, 2023.
- [17] 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," *arXiv preprint arXiv:2312.14150*, 2023.
- [18] T.-H. Wang, A. Maalouf, W. Xiao, Y. Ban, A. Amini, G. Rosman, S. Karaman, and D. Rus, "Drive anywhere: Generalizable end-to-end autonomous driving with multi-modal foundation models," *arXiv preprint arXiv:2310.17642*, 2023.
- [19] D. Wu, W. Han, T. Wang, Y. Liu, X. Zhang, and J. Shen, "Language prompt for autonomous driving," *arXiv preprint arXiv:2309.04379*, 2023.
- [20] 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," *arXiv preprint arXiv:2305.14836*, 2023.
- [21] M. Nie, R. Peng, C. Wang, X. Cai, J. Han, H. Xu, and L. Zhang, "Reason2drive: Towards interpretable and chain-based reasoning for autonomous driving," *arXiv preprint arXiv:2312.03661*, 2023.
- [22] K. Chitta, A. Prakash, and A. Geiger, "Neat: Neural attention fields for end-to-end autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 793–15 803.
- [23] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou, *et al.*, "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 9710–9719.
- [24] Y. Xu, X. Yang, L. Gong, H.-C. Lin, T.-Y. Wu, Y. Li, and N. Vasconcelos, "Explainable object-induced action decision for autonomous vehicles," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9523–9532.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.