

Prompting Multi-Modal Tokens to Enhance End-to-End Autonomous Driving Imitation Learning with LLMs

Yiqun Duan^{1,*}, Qiang Zhang^{2,*}, Renjing Xu^{2,†}

Abstract—The utilization of Large Language Models (LLMs) within the realm of reinforcement learning, particularly as planners, has garnered a significant degree of attention in recent scholarly literature. However, a substantial proportion of existing research predominantly focuses on planning models for robotics that transmute the outputs derived from perception models into linguistic forms, thus adopting a ‘pure-language’ strategy. In this research, we propose a hybrid End-to-End learning framework for autonomous driving by combining basic driving imitation learning with LLMs based on multi-modality prompt tokens. Instead of simply converting perception results from the separated train model into pure language input, our novelty lies in two aspects. 1) The end-to-end integration of visual and LiDAR sensory input into learnable multi-modality tokens, thereby intrinsically alleviating description bias by separated pre-trained perception models. 2) Instead of directly letting LLMs drive, this paper explores a hybrid setting of letting LLMs help the driving model correct mistakes and complicated scenarios. The results of our experiments suggest that the proposed methodology can attain driving scores of 49.21%, coupled with an impressive route completion rate of 91.34% in the offline evaluation conducted via CARLA. These performance metrics are comparable to the most advanced driving models.

I. INTRODUCTION

Recent advancements in AI systems for autonomous driving can be classified into two main categories: Pipeline Formation and End-to-End (E2E) Formation. Pipeline formation [1], [2], [3], [4], [5], [6], [7], [8], involves manually decomposing the drive into sequential modules. These modules consist of tasks such as localization [9], [10], scene reconstruction [11], [12], planning [13], [14], and control [15], [16] according to certain rules. However, these separately learned rules struggle to cover long-tail scenarios, which are crucial in driving situations. Thus, researchers are aiming to make the system learn using an end-to-end formation through reinforcement or imitation learning, to better mimic human-like decision-making and bring it a step closer to handling these important scenarios effectively.

E2E driving approaches [17], [18], [19], [20], [21] employ state-to-action imitation or reinforcement learning strategies on intermediate feature representation states. This enables agents to behave appropriately in various driving contexts. The introduction of multi-modal sensory fusion techniques

[22], [23], [24], [25] have significantly improved the perception state representation. These techniques make use of both visual and LiDAR data [26], [27], [28], [24], [25], [29], presenting it in a Bird’s Eye View (BEV) format or as a flat tensor for the decoding process. ST-P3 [17] approach and Transfuser [22] reach SOTA performance with the multi-modality imitation learning.

However, these methods emphasize on state representation and pay less attention to the decoder. Generally, the learning system predicts waypoints for the next n time steps based on the encoded feature, with an additional controller translating waypoints into actual controls. However, merely predicting waypoints falls short in the face of the complex demands of autonomous driving. This limitation hinders the practical application of end-to-end driving systems. Interfuser [30] and TCP [31] respectively improve the decoder by introducing safety constraints on the controller and directly adding a branch to predict safety control signals. UniAD [32] try to unify pipelines with an end-to-end training framework and reach SOTA performance by introducing more immediate supervision such as occupancy and lanes. Especially, Think-Twice [33] proposes to involve the decoder in an anthropomorphized approach that enables interaction between the encoder and decoder following ‘driving after thinking twice’.

However, it’s hard to manually list all possible driving rules or components comprehensively. In that case, letting the autonomous system learn semantically could benefit the end-to-end driving tasks. ADAPT [34] has given a first exploration of modeling the driving task as an image captioning model, which simultaneously predicts driving actions and driving explanation as a caption using visual transformers. More recent works have investigated using Large Language Models (LLMs) for wider autonomous/robotic systems. PaLM-E [35] utilized pre-trained LLM to complete multiple embodied tasks, including sequential robotic manipulation planning, visual question answering, and captioning. Mini-GPT4 [36] and LLaVA [37] introduce practical approaches for visual context tuning for LLMs. However, end-to-end autonomous driving has not yet been adequately explored. Moreover, these methods concentrate on the language model in an auto-regressive way. How to incorporate imitation learning and reinforcement learning in autonomous driving scenarios remain uninvestigated.

In this paper, we propose to model end-to-end autonomous driving as a multi-modality language model. For state encoding, we propose a two-branch framework to encode both visual and LiDAR features into joint feature token representations. We prompt the multi-modality visual tokens

¹Yiqun Duan is with the HAI Centre, Australia Artificial Intelligence Institute, School of Computer Science, University of Technology Sydney, 2007, Ultimo, Australia yiqun.duan@student.uts.edu.au

²Qiang Zhang and Renjing Xu are with the School of Computer Science at The Hong Kong University of Science and Technology (Guangzhou) qzhang749@connect.hkust-gz.edu.cn, renjingxu@hkust-gz.edu.cn

* are equal contributors, † is the corresponding author

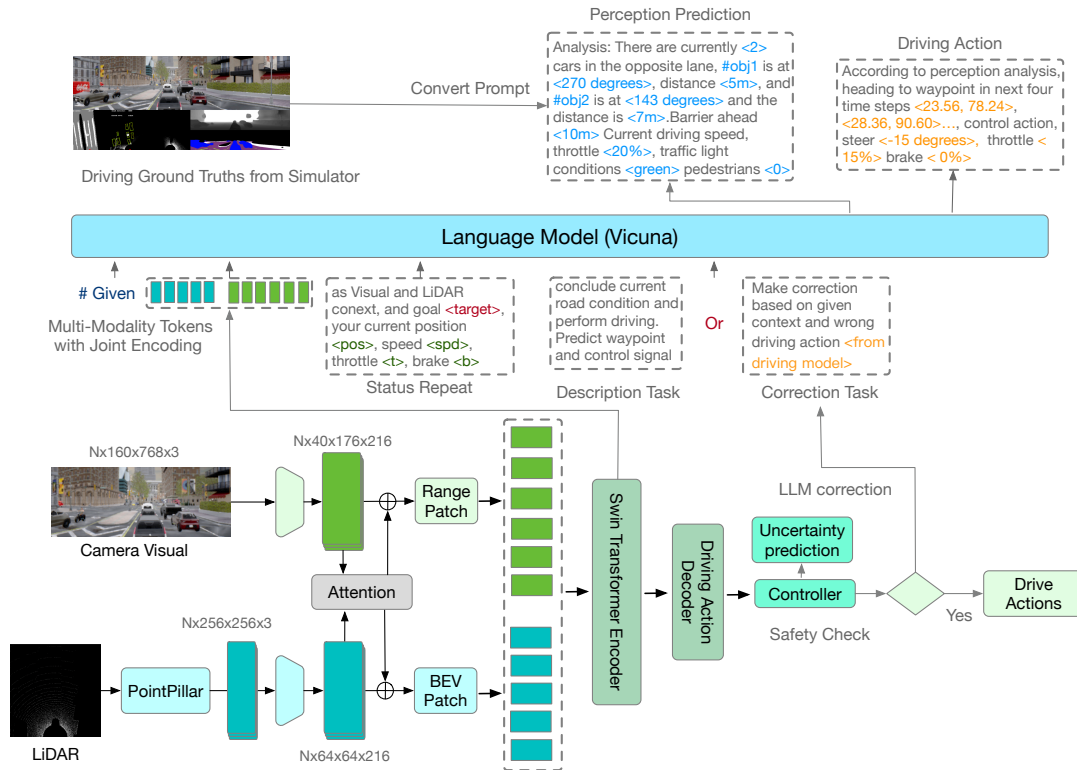


Fig. 1. Overview of utilizing language model for autonomous driving. The system takes the camera and LiDAR input and extracts shallow feature maps with two branches. Then a joint Swin transformer encodes perception input into joint token representation. The prompt is constructed by sequentially concatenating multi-modality joint tokens, status repeat tokens, and driving task tokens. The task prompts consist of two kinds 1) directly predicting the driving actions and 2) driving action correction given the driving output. The controller model also predicts an uncertainty score by an MLP layer, which decides whether to ask GPT to correction or drive by itself. The model is trained by auto-regressively predicting perception description and the driving action. Driving actions are executed by a final controller.

into driving languages, which contain driving descriptions followed by driving actions, not only waypoints but also control signals. The driving-like language formation could help the model learn comprehensive driving skills not merely waypoints, not merely control signals, but also the reasoning before taking the action. Despite the auto-regression loss, we design a subtle mechanism to incorporate a language model with driving metrics reward-guided reinforcement tuning. Considering the randomness of the large language generation model, 1) the generated results may exceed the normal safety threshold or do not contain waypoints or driving actions. 2) the generated waypoints and the control signal may conflict with each other. To enhance driving safety, we design a mechanism of re-quiring the driving decision with enhanced prompt to enable model to rethink the driving decision.

We have performed comprehensive experiments using the off-board driving simulator CARLA [38], a platform capable of providing the requisite driving information necessary for generating prompted language supervision. The findings of our experiments suggest that LLMs can attain a level of driving performance comparable to the current state-of-the-art models in this domain. The incorporation of language models into driving systems with reinforcement learning enhancement offers clear advantages. When driving actions are prompted through language, it enables model learning that extends beyond the mere prediction of waypoints. It

also allows for the extraction of the underlying driving logic that is inherently contained within the language itself. This methodology serves as a step towards the development of a more personalized driving model, thereby bringing us closer to the goal of a more human-like autonomous driving system. The contribution lies threefold.

- A novel framework incorporates multi-modality perception as joint token representations into language prompts for end-to-end autonomous driving.
- A simple yet effective prompting approach on unified observation, current states, trajectory, and control actions in continuous prompts.
- Incorporating reward-guided reinforcement learning supervision on language prompts in autonomous driving scenarios.

II. RELATED WORKS

Imitation Learning for End-to-End Driving: The evolution of End-to-End autonomous driving has typically followed a *destructured, then structured* trajectory. Early explorations, such as ALVINN [39], and DAVE [40], modeled a simple projection relationship between the input from a single-view camera and the steering wheel or accelerator angle. The prevalent approach at this stage was to eliminate human prior knowledge, thereby offering the network a more direct learning target to improve the fitting capability for

fully supervised behavior cloning. However, sensory input was limited in this early era, leading to the introduction of multiple imitation learning techniques [41], [42], [43] to enhance the data quality for behavior cloning. Subsequent developments saw the expansion of camera sensor usage [44] from a single view to multiple views to enhance planning stability. In more recent works, the reintroduction of prior knowledge has been explored to augment driving performance. WOR [45] and MARL[46] suggest that incorporating prior trained knowledge from related vision tasks such as detection [47], [48], [49] could enhance model performance. NEAT [20] introduced an attention module between different camera views to boost feature quality. ST-P3 [17], VAD [50], and UniAD [32] amalgamate human prior module design from pipeline methods to increase driving rationality.

Reinforcement Learning for LLMs: For language models, relying solely on supervised learning to generate accurate responses is inadequate as it's difficult to determine whether the model truly understands the answer to a question. The RL framework motivates a model to produce correct answers in a natural manner without resorting to fabrication because the model receives poor scores for incorrect answers. Reinforcement Learning from Human Feedback (RLHF) [51] is a critical technology in the current wave of large language models (LLMs). Three of OpenAI's most prominent LLMs, namely InstructGPT [52], ChatGPT, and GPT4, rely on RLHF as their core technology. Although PaLM-E integrated the 540 billion-parameter PaLM [53] LLM and the 22 billion-parameter Vision Transformer [54] to create an embodied multimodal language model, its approach on the autonomous driving task lacked sufficient detail. Haomo's DriveGPT ¹ for the autonomous driving scenario proposes to convert BEV information into drive language and learning language tokens. However, it is not solely end-to-end as it does not learn directly with raw tokens from sensors such as cameras and LiDAR.

Multi-Modality Fusion For Driving Tasks: The limited capabilities of a monotonous camera sensor have led to the exploration of other sensory modalities, such as LiDAR, Radar, and Openmap, to enhance driving stability [55], [56], [57], [58]. Most previous approaches to modality fusion have been specifically designed for certain perception tasks such as 3D object detection [24], [25], depth estimation [29], [59], and instance motion forecasting [60], [61]. BEVFusion [25], for instance, fuses LiDAR with image features by using completely independent encoders with geometric projection into the same Bird's-eye-view (BEV) feature space at a late stage. However, Transfuser [22] observed that pure geometrical fusion representation impairs the performance of comprehensive and complex urban E2E autonomous driving.

III. METHODOLOGY

A. Multi-Modality Joint Token Encoder

In order to provide a more suitable driving context for GPT-based models, we propose to fuse Visual-LiDAR per-

ception input into a joint token representation. The architecture is described in Fig. 1, which contains two stages, *early fusion* and *late fusion*. This two-stage network is designed as follows:

Early Fusion: In the first stage, two distinct CNN branches are utilized to extract shallow features from a monotonic image and LiDAR inputs, respectively. The image branch concatenates three front view camera inputs, each with a field of view of 60 degrees, into a single monotonic view, subsequently reshaped to the shape $3 \times 160 \times 704$. On the other hand, the LiDAR branch processes raw LiDAR input with PointPillar [55] to create a bird's eye view (BEV) feature with shape $33 \times 256 \times 256$. As these lower-level features retain strong geometric relations, the independent encoder can extract tight local feature representation with fewer distractions. During this stage, we apply cross-modality self-attention to enhance geometrical feature fusion.

The feature map \mathbf{F}_{im} , \mathbf{F}_{Li} representing the image branch and LiDAR branch are reshaped by flattening [62], [22] and concatenating into sequence $\mathbf{F}_{in} = \text{cat}(\mathbf{F}_{im}, \mathbf{F}_{Li})$. The transformer layer is applied on $\mathbf{F}_{in} \in \mathbb{R}^{N^* \times C}$ to perform multi-head attention between each token in N^* dimension as defined in Eq. 1.

$$\begin{aligned} \mathbf{Q}_F &= \mathbf{F}_{in}W_Q, \mathbf{K}_F = \mathbf{F}_{in}W_K, \mathbf{V}_F = \mathbf{F}_{in}W_V, \\ \mathbf{F}_{Li} &= m_1(\mathbf{A}) + \mathbf{F}_{in}, \mathbf{F}_{im} = m_2(\mathbf{A}) + \mathbf{F}_{in}, \\ \mathbf{A} &= \text{softmax}\left(\frac{\mathbf{Q}_F\mathbf{K}_F^T}{\sqrt{D}}\right)\mathbf{V}_F \end{aligned} \quad (1)$$

where m_1 and m_2 denote two different MLP layers that project the joint attention into proper shape and are directly added to the original feature map after normalization.

Late Fusion: After low-level feature extraction with attention, we intend to align features from both modalities into a unified semantic token space, where we treat each $16 \times 16 \times C$ feature segment to be a semantic *word* as shown in Fig. 1. To maintain the spatial relationships across these tokens, we use the cross-attention mechanism.




Fig. 2. Illustration of joint token representation, the perception token is aligned with segmentation embedding and positional embedding to distinguish from normal word token.

A segment embedding is added with positional embedding $\mathbf{PE} + \mathbf{SE}$ to indicate the position information is from range view or BEV feature space. The feature maps from different modalities are tokenized and concatenated into unified tokens $\mathbf{F}_s \in \mathbb{R}^{N^* \times C}$ with patch size 16×16 by adding $\mathbf{PE} + \mathbf{SE}$ to the token. A 4-layer transformer encoder is applied on the unified token sequence \mathbf{F}_s by setting query $Q(\mathbf{F}_s)$, key $K(\mathbf{F}_s)$, and value $V(\mathbf{F}_s)$. As the late fusion is performed

¹<https://drivegpt.haomoi.com/>

by share transformer encoding, this could introduce a better joint semantic feature representation between especially for the language model.

B. Prompt Construction

In the endeavor to integrate Large Language Models (LLMs) within the domain of autonomous driving, a salient challenge lies in the effective design of language prompts and the translation of supervisory signals into language constructs. A key factor in this process entails the creation of prompts that adeptly incorporate multimodal tokens, while also considering the vehicle’s instantaneous state, the ambient environment, and short-term as well as long-term objectives. Instead of letting the LLMs directly drive, the basic design logic is to let LLMs help the basic driving model correct driving actions. This setting alleviates the LLMs generating actions or goals not exist in real scenarios.

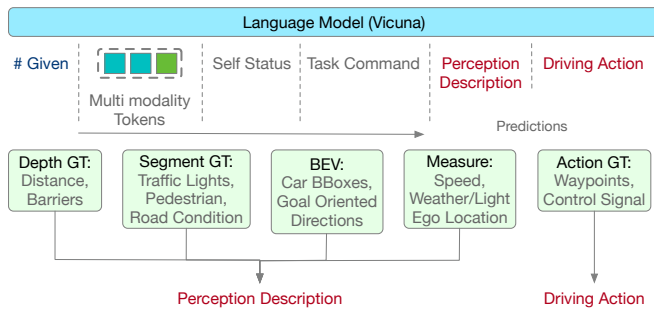


Fig. 3. Illustration of the prompt construction. Both traditional perception supervision and driving actions are converted into language descriptions through a descriptor.

The given information during driving is first contacted sequentially, 1) multi-modality tokens 2) self-status of the car 3) driving task command. Here, the self-status of the car denotes the current speed, throttle, brake, and current position and driving task commands denote the instruction (perform driving/solve drive conflicts). After that, supervisions are appended sequentially as language guidance considering two aspects 1) perception description and 2) driving action as shown in Fig. 3. Perception description guides the language model to predict perception results which are mostly used as auxiliary heads supervision in previous works [22], [63], [50], [32], such as segmentation, depth prediction, BEV prediction, and detection. Driving action predicts both the waypoints in the next 4 time steps and control actions. A notable advantage of auto-regressive models over their traditional counterparts is their inherent ability to learn driving logic, facilitated by their intrinsic capacity to assimilate the complex, underlying linguistic logic. Towards this end, we have proposed a comprehensive prompt design strategy, which is depicted above.

Tasks: The task prompts are designed to have three modes.

1) Given the multi-modality tokens directly generate perception observation output, and driving actions. Normally the driving language output is not used for driving considering the driving cost, yet the normal learning parallel with the




Fig. 4. Illustration of the re-query mechanism when there is disagreement.

driving model is crucial for establishing driving logic. 2) Re-query the LLM if the driving modal’s language output contradicts to the safety controller. 3) Given the multi-modality tokens and driving output, correct the driving error by LLMs.

C. Re-Query Mechanism

A remarkable problem of auto-regressive models hinders the performance is the uncertainty of the language prediction. Since in Section III-B, we propose to predict comprehensive driving actions and waypoints together in sequential order, it is observed in our experiment that the waypoints prediction might be conflicted with the control actions as illustrated in Fig. 4. Here the driving actions and waypoints are extracted through predefined regular expressions. The waypoints are converted to the control signal by the traditional pid controller used in previous works [22], [63]. If conflicts between the predicted action and the exceed the pre-defined threshold the system will trigger a re-query mechanism. At this stage, the task commands part in the prompt (Fig. 1) will be changed from the primary driving commands to query with the previous disagreement. If without an agreement, we will select the waypoints leaded action for final driving control. This mechanism could help the language model to ‘think twice’ when dealing with unclear or hard situations.

D. Reinforcement Guided Tuning

Pure auto-regressive learning is not sufficient to deal with the regression trajectories. Yet, under our design, the waypoints are purely learned in auto-regressive with MLE loss. In that case, we add a reinforcement-guided regression loss to boost the model with the prediction accuracy. Different from InstructGPT [52] that need human feedback to apply reinforcement fine-tuning, our scenarios have automated generated guidance. As the goal is to imitate the expert driving trajectories which is accessible in our task setting, we can assume the description of the expert trajectories as y_w . Since the GPT-based model only could put supervision in token wise yet, while the whole sequence is meaningful for regression, we introduce Proximal Policy Optimization Algorithms [64] (PPO) with masks to apply the supervision. The tuning loss L_{rl} is presented as a reward in policy gradient as defined below:

$$\mathcal{L}_{rl} = LLM(\mathbf{x}^{T_i} | \mathbf{x}^{1:T_i-1})R(\mathbf{x}^{1:T_i}, mask), \quad (2)$$

where mask denotes whether this is a parameter token. If the token is not a parameter token the reward will be always 0. \mathbf{x}^{T_i} denotes the predicted token at time step T_i and $R(\mathbf{x}^{1:T_i})$ denotes the reward function for waypoint prediction

and safety metrics. Because there is no sense in letting discriminators only judge on an incomplete sequence, Monte Carlo search (MC search) is applied to the sample and average unknown $|x| - T_i$ tokens (where $|x|$ is the total length).

$$R(x^{1:T_i}, mask) = \begin{cases} \frac{1}{N} \sum_{n=1}^N \|\mathbf{w}_t - \mathbf{w}_t^{gt}\|_1 \\ \mathcal{L}_{safety}(x^{1:T_i}), & x^{1:T_i} \in safety \end{cases} \quad (3)$$

where $\mathbf{x}_n^{1:|x|}$ denotes n^{th} sampled sequence by MC search, and \mathcal{L}_{safety} is the safety metrics used in InterFuser [30]. The training loss we use for training is the combination of normal language loss and reward loss, $L = L_{MLE} + \lambda L_{rl}$.

IV. EXPERIMENT

A. Experimental Setting

Simulation Environment: This paper performs the end-to-end driving task under the simulator. We use the **CARLA Simulator 0.9.14**², a high-performance driving simulator under urban scenario as our simulation environment. The driving is performed in a variety of areas under various weather and lighting conditions. The agent is expected to perform driving given predefined routes under complicated traffic and weather conditions, where the routes are a sequence of sparse goal locations with GPS coordinates.

Data Collection: We use the simulator to collect 1000 routes in 8 official town maps with an average length of 400m by using CARLA rule-based expert auto-agent [65] with 228k training pairs.

Evaluation Benchmark: This paper is evaluated by both online benchmark and offline benchmarks. Although the CARLA simulator provides an official evaluation leaderboard, the use time is restricted to only 200 hours per month yet K80 GPUs are not feasible to deploy large LLMs, which makes the official leaderboard unsuitable for ablation studies or obtaining detailed statistics involving We conduct our ablation comparison based on the Longset6 Benchmark proposed by Transfuser [22].

Evaluation Metrics: For both online evaluation and offline evaluation, we follow the official evaluation metrics to calculate three main metrics, *Route Completion (RC)*, *Infraction Score (IS)*, and *Driving Score (DS)*. The RC score is the percentage of route distance completed. Given R_i as the completion by the agent in route i , RC is calculated by averaging the completion rate $RC = \frac{1}{N} \sum_i R_i$. The IS is calculated by accumulating the infractions P_i incurred by the agent during completing the routes.

B. Implementation Details

Architectual: The multi-modality hybrid fusion network uses cameras and LiDAR as the sensory inputs, where image input is monotonic with FOV 120 degree and reshaped into shape (160, 704) and the LiDAR point cloud is converted to BEV representation [55] and reshaped into 256, 256. We

apply angular viewpoint augmentation [22] for the training data, by randomly rotating the LiDAR inputs by $\pm 20^\circ$ and adjusting ground truth labels accordingly. Self-attention modules are applied between feature maps after the first two convolutional blocks, where the resolution is respectively $(C_1, 40, 176)$ and $(C_2, 20, 88)$, where C_1, C_2 denotes the channel dimension 72, 216 under our setting. For the attention module, two transformer layers with hidden dimensions 512 and 4 attention heads are respectively applied as encoder and decoder sequentially. For polar ray grid sampling details, we follow the settings of previous work [66]. The joint token encoding is applied on feature maps with resolutions 20, 88 and 32, 32 after two times downsampling. InterFuser [30] and TCP [31]. For controller and safety loss, we directly follow structures from InterFuser [30].

Training of the LLMs: The driving data are gathered by autonomous agents as detailed in [30], ensuring comprehensive logging of all safety metrics. These metrics are then transformed into "driving languages" through a set of predefined prompts, which utilize task token 1 for representation. Subsequently, the driving model is trained employing the loss functions delineated in InterFuser [30]. Error data are compiled by assessing the training set, capturing all outputs from both the driving model and the ground truth (GT). Notably, these driving languages incorporate task tokens 2 and 3 at sampling rates of 20% and 80%, respectively. The uncertainty prediction layer undergoes training by categorizing this duo of data into binary labels and then leveraging the cross-entropy loss for optimization. As for the language model, we employ Vicuna 33B³ as the pre-trained checkpoints and perform finetuning

TABLE I
LOCAL DRIVING EVALUATION ON LONGSET6 ON CARLA 0.9.14

Method	DS \uparrow	RC \uparrow (%)	IS \uparrow
NEAT [20]	20.63 \pm 2.34	45.31 \pm 3.25	0.54 \pm 0.03
WOR [67]	21.48 \pm 2.09	52.28 \pm 4.07	0.51 \pm 0.05
GRIAD [68]	26.98 \pm 4.23	71.43 \pm 3.57	0.56 \pm 0.04
LAV [69]	37.62 \pm 1.77	83.94 \pm 2.69	0.59 \pm 0.05
GF [22]	27.44 \pm 2.11	80.43 \pm 3.98	0.32 \pm 0.02
Late TF [22]	35.42 \pm 4.07	83.13 \pm 1.01	0.39 \pm 0.04
TransFuser [22]	46.95 \pm 5.81	89.64 \pm 2.08	0.52 \pm 0.08
TCP [31]	49.95 \pm 3.78	92.12 \pm 2.08	0.59 \pm 0.08
InterFuser [30]	52.68\pm6.23	92.08 \pm 2.78	0.62\pm0.09
PromptGPT	52.34 \pm 5.11	92.37\pm1.09	0.60 \pm 0.09
Expert	75.83 \pm 2.45	89.82 \pm 0.59	0.85 \pm 0.03

C. Driving Performance

Comparison Baseline: We compare this method with state-of-the-art (SOTA) E2E driving baselines below: 1) NEAT [20] the attention field fusion on image inputs. 2) WOR [67], multi-stage Q-function based learning framework; 3) GRIAD [68] general reinforced imitation baseline; 4) LAV [69], the extension of WOR learning from multiple observed agents; 5) Transfuser [22], the solid SOTA baseline

²<https://carla.org/>

³<https://huggingface.co/lmsys/vicuna-33b-v1.3>

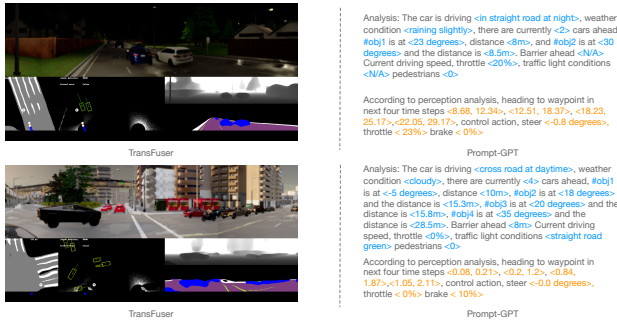


Fig. 5. Visualization of driving states between traditional methods and language model.

which introduce transformer fusion into E2E driving. Since MMFN [23] are Transfuser baseline if only with image and Lidar, we do not list the evaluation results of MMFN. Since the target of this paper is to enhance the modality fusion, we also do not list InterFuser [30] as it introduces an additional controller and decision module. We also include Late-Transfuser (Late TF) and geometrical fusion (GF) in [22] to report the early-late fusion comparison. TCP [31] and InterFuser [30] are selected as the representative of adding additional safety guidance.

Metrics Analysis: The performance is reported in the offline benchmark LongSet6 [22] in Table I. The mean and variance metric values are calculated by evaluating each agent three times. PromptGPT reached a mean driving score (DS) of 49.08 and route completion (RC) of 91.87. Although the general performance is lower than current SOTA baselines InterFuser [30] and TCP [31], the PromptGPT-based driving could still reach competitive performance compared to other previous methods. Meanwhile, the safety score (IS) is with the same level of TCP while route completion (RC) is lower, which is rational since comprehensive language-based driving could better consider the corner cases.

Driving Visualization To give more intuitive explanations, we visualize the driving mediate state for both traditional waypoints prediction networks (TransFuser [22] and LLM-based E2E driving methods in Fig. 5. Given two typical scenarios, we could observe that language prompts could describe most of the information in the driving scenario. Yet, it still suggests that LLMs are still not very good at recognizing small and partially visible vehicles.

D. Ablation Study

To assess the impact of LLM-facilitated driving correction, we conduct driving experiments across three distinct tasks: 1) Driving solely with LLMs, 2) Driving exclusively with the driving model, and 3) Employing both LLMs and the driving model for driving. Furthermore, we introduce various scales of LLMs to evaluate the influence of model size on performance. The results are presented in Table II

Our observations indicate that when relying solely on language model-based driving, the performance attains a driving score of 43.08. This is a reduction of 9.26 compared to driving with a pure language model. When using the

TABLE II
ABLATION STUDY ON DIFFERENT SETTINGS FOR DRIVELM.

Component	Method	DS \uparrow	RC \uparrow (%)
LLMs	Vicuna 7B	16.54 \pm 2.34	39.28 \pm 3.42
	Vicuna 13B	34.55 \pm 2.21	79.43 \pm 4.05
	Vicuna 33B	43.08 \pm 5.28	91.87 \pm 1.08
Task-Type	Drive (Vicuna 33B)	43.08 \pm 5.28	87.43 \pm 3.98
	Drive + w/o re-query	39.62 \pm 4.30	83.44 \pm 2.03
	Drive + Correction	52.34 \pm 5.11	92.37 \pm 1.09

initially generated output directly, the performance further declines to 39.62. However, when employing LLMs as correction advisers, the performance elevates to state-of-the-art levels. These results underscore the effectiveness of our approach. Additionally, concerning the influence of model scales, there’s a distinct correlation between model size and performance. As we scale the LLM from 7B to 33B, there’s a notable performance boost. This trend aligns with our expectations, as intricate driving logic necessitates substantial model capacity.

V. LIMITATION

While language models exhibit competitive performance in driving simulators, they are not without limitations: The inference speed, combined with the re-query mechanism, results in sluggish response times, rendering it unsuitable for real-time driving. Additionally, the complexity of the driving framework poses challenges for evaluation on on-line benchmarks. The unpredictability inherent in language prompts is also a concern. On occasion, the LLM may not generate the necessary driving actions, prematurely halting and solely predicting perception parameters. Also, in order to train LLMs for perform correction, the current framework need the driving model pre-trained and evaluated in a ‘off-shelf’ formation to collect language training corpus, which is not yet End-to-End as well as heavy to train.

VI. CONCLUSION

In this work, we delve into the innovative approach of utilizing LLMs with multi-modality tokens for end-to-end (E2E) autonomous driving. Rather than letting LLMs take the wheel directly, we suggest harnessing them to aid driving models in rectifying erroneous behaviors, especially in complex scenarios. Our experiments indicate that this corrective approach can prevent LLMs from generating non-existent actions. While the current language model hasn’t yet surpassed state-of-the-art performance, it demonstrates a commendable level of competitiveness. This implies the significant potential for integrating LLMs within the realm of autonomous driving.

ACKNOWLEDGMENT

We would like to extend our heartfelt thanks to the Beijing Innovation Center of Humanoid Robotics Co., Ltd. for their substantial support in our research endeavors.

REFERENCES

- [1] I. Gog, S. Kalra, P. Schafhalter, M. A. Wright, J. E. Gonzalez, and I. Stoica, "Pylot: A modular platform for exploring latency-accuracy tradeoffs in autonomous vehicles," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8806–8813, IEEE, 2021.
- [2] T. Liu, Q. Liao, L. Gan, F. Ma, J. Cheng, X. Xie, Z. Wang, Y. Chen, Y. Zhu, S. Zhang, *et al.*, "Hercules: An autonomous logistic vehicle for contact-less goods transportation during the covid-19 outbreak," *arXiv preprint arXiv:2004.07480*, 2020.
- [3] J. Jiao, Y. Zhu, H. Ye, H. Huang, P. Yun, L. Jiang, L. Wang, and M. Liu, "Greedy-based feature selection for efficient lidar slam," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5222–5228, IEEE, 2021.
- [4] Y. Liu, Y. Yixuan, and M. Liu, "Ground-aware monocular 3d object detection for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 919–926, 2021.
- [5] Y. Li and J. Ibanez-Guzman, "Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 50–61, 2020.
- [6] H. Song, W. Ding, Y. Chen, S. Shen, M. Y. Wang, and Q. Chen, "Pip: Planning-informed trajectory prediction for autonomous driving," in *European Conference on Computer Vision*, pp. 598–614, Springer, 2020.
- [7] M. Wen, J. Park, and K. Cho, "A scenario generation pipeline for autonomous vehicle simulators," *Human-centric Computing and Information Sciences*, vol. 10, no. 1, pp. 1–15, 2020.
- [8] L. Claussmann, M. Revilloud, D. Gruyer, and S. Glaser, "A review of motion planning for highway autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 5, pp. 1826–1848, 2019.
- [9] T. Qin, Y. Zheng, T. Chen, Y. Chen, and Q. Su, "A light-weight semantic map for visual localization towards autonomous driving," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11248–11254, IEEE, 2021.
- [10] A. Woo, B. Fidan, and W. W. Melek, "Localization for autonomous driving," *Handbook of Position Location: Theory, Practice, and Advances, Second Edition*, pp. 1051–1087, 2018.
- [11] C. R. Dyer, "Volumetric scene reconstruction from multiple views," in *Foundations of Image Understanding*, pp. 469–489, Springer, 2001.
- [12] A. Bozic, P. Palafox, J. Thies, A. Dai, and M. Nießner, "Transformer-fusion: Monocular rgb scene reconstruction using transformers," *Advances in Neural Information Processing Systems*, vol. 34, pp. 1403–1414, 2021.
- [13] J. Letchner, J. Krumm, and E. Horvitz, "Trip router with individualized preferences (trip): Incorporating personalization into route planning," in *AAAI*, pp. 1795–1800, 2006.
- [14] N. Abu, W. Bukhari, M. Adli, S. Omar, and S. Sohaimeh, "A comprehensive overview of classical and modern route planning algorithms for self-driving mobile robots," *Journal of Robotics and Control (JRC)*, vol. 3, no. 5, pp. 666–678, 2022.
- [15] S. Tsugawa, "Vision-based vehicles in japan: Machine vision systems and driving control systems," *IEEE Transactions on industrial electronics*, vol. 41, no. 4, pp. 398–405, 1994.
- [16] J. Kong, M. Pfeiffer, G. Schilb, and F. Borrelli, "Kinematic and dynamic vehicle models for autonomous driving control design," in *2015 IEEE intelligent vehicles symposium (IV)*, pp. 1094–1099, IEEE, 2015.
- [17] 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*, pp. 533–549, Springer, 2022.
- [18] A. Tampuu, T. Matiisen, M. Semkin, D. Fishman, and N. Muhammad, "A survey of end-to-end driving: Architectures and training methods," *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [19] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [20] 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*, pp. 15793–15803, 2021.
- [21] M. Toromanoff, E. Wirbel, and F. Moutarde, "End-to-end model-free reinforcement learning for urban driving using implicit affordances," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7153–7162, 2020.
- [22] A. Prakash, K. Chitta, and A. Geiger, "Multi-modal fusion transformer for end-to-end autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7077–7087, 2021.
- [23] Q. Zhang, M. Tang, R. Geng, F. Chen, R. Xin, and L. Wang, "Mmfnet: Multi-modal-fusion-net for end-to-end driving," *arXiv preprint arXiv:2207.00186*, 2022.
- [24] J. Huang, G. Huang, Z. Zhu, and D. Du, "Bevdet: High-performance multi-camera 3d object detection in bird-eye-view," *arXiv preprint arXiv:2112.11790*, 2021.
- [25] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. Rus, and S. Han, "Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," *arXiv preprint arXiv:2205.13542*, 2022.
- [26] S. Fadadu, S. Pandey, D. Hegde, Y. Shi, F.-C. Chou, N. Djuric, and C. Vallespi-Gonzalez, "Multi-view fusion of sensor data for improved perception and prediction in autonomous driving," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 2349–2357, 2022.
- [27] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3d object detection network for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1907–1915, 2017.
- [28] Y. Zhou, P. Sun, Y. Zhang, D. Anguelov, J. Gao, T. Ouyang, J. Guo, J. Ngiam, and V. Vasudevan, "End-to-end multi-view fusion for 3d object detection in lidar point clouds," in *Conference on Robot Learning*, pp. 923–932, PMLR, 2020.
- [29] Y. Li, Z. Ge, G. Yu, J. Yang, Z. Wang, Y. Shi, J. Sun, and Z. Li, "Bevdepth: Acquisition of reliable depth for multi-view 3d object detection," *arXiv preprint arXiv:2206.10092*, 2022.
- [30] H. Shao, L. Wang, R. Chen, H. Li, and Y. Liu, "Safety-enhanced autonomous driving using interpretable sensor fusion transformer," *arXiv preprint arXiv:2207.14024*, 2022.
- [31] P. Wu, X. Jia, L. Chen, J. Yan, H. Li, and Y. Qiao, "Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline," in *NeurIPS*, 2022.
- [32] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang, L. Lu, X. Jia, Q. Liu, J. Dai, Y. Qiao, and H. Li, "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [33] X. Jia, P. Wu, L. Chen, J. Xie, C. He, J. Yan, and H. Li, "Think twice before driving: Towards scalable decoders for end-to-end autonomous driving," in *CVPR*, 2023.
- [34] B. Jin, X. Liu, Y. Zheng, P. Li, H. Zhao, T. Zhang, Y. Zheng, G. Zhou, and J. Liu, "Adapt: Action-aware driving caption transformer," *arXiv preprint arXiv:2302.00673*, 2023.
- [35] D. Driess, F. Xia, M. S. M. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, W. Huang, Y. Chebotar, P. Sermanet, D. Duckworth, S. Levine, V. Vanhoucke, K. Hausman, M. Toussaint, K. Greff, A. Zeng, I. Mordatch, and P. Florence, "Palm-e: An embodied multimodal language model," in *arXiv preprint arXiv:2303.03378*, 2023.
- [36] D. Zhu, J. Chen, X. Shen, X. Li, and M. Elhoseiny, "Minigpt-4: Enhancing vision-language understanding with advanced large language models," *arXiv preprint arXiv:2304.10592*, 2023.
- [37] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," 2023.
- [38] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*, pp. 1–16, PMLR, 2017.
- [39] D. A. Pomerleau, "Alvin: An autonomous land vehicle in a neural network," *Advances in Neural Information Processing Systems*, vol. 1, 1988.
- [40] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [41] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *IEEE International Conference on Robotics and Automation*, pp. 4693–4700, IEEE, 2018.

- [42] M. Bansal, A. Krizhevsky, and A. Ogale, “Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst,” *arXiv preprint arXiv:1812.03079*, 2018.
- [43] A. Kendall, J. Hawke, D. Janz, P. Mazur, D. Reda, J.-M. Allen, V.-D. Lam, A. Bewley, and A. Shah, “Learning to drive in a day,” in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8248–8254, IEEE, 2019.
- [44] S. Hecker, D. Dai, and L. Van Gool, “End-to-end learning of driving models with surround-view cameras and route planners,” in *Proceedings of the European Conference on Computer Vision*, pp. 435–453, 2018.
- [45] D. Chen, V. Koltun, and P. Krähenbühl, “Learning to drive from a world on rails,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15590–15599, 2021.
- [46] M. Toromanoff, E. Wirbel, and F. Moutarde, “End-to-end model-free reinforcement learning for urban driving using implicit affordances,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7153–7162, 2020.
- [47] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- [48] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, “Backpropagation applied to handwritten zip code recognition,” *Neural computation*, vol. 1, no. 4, pp. 541–551, 1989.
- [49] J. Redmon and A. Farhadi, “Yolov3: An incremental improvement,” *arXiv*, 2018.
- [50] 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,” *arXiv preprint arXiv:2303.12077*, 2023.
- [51] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, “Deep reinforcement learning from human preferences,” *Advances in neural information processing systems*, vol. 30, 2017.
- [52] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, *et al.*, “Training language models to follow instructions with human feedback,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 27730–27744, 2022.
- [53] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, *et al.*, “Palm: Scaling language modeling with pathways,” *arXiv preprint arXiv:2204.02311*, 2022.
- [54] M. Dehghani, J. Djolonga, B. Mustafa, P. Padlewski, J. Heek, J. Gilmer, A. Steiner, M. Caron, R. Geirhos, I. Alabdulmohsin, *et al.*, “Scaling vision transformers to 22 billion parameters,” *arXiv preprint arXiv:2302.05442*, 2023.
- [55] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, “Pointpillars: Fast encoders for object detection from point clouds,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12697–12705, 2019.
- [56] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine, “Precog: Predictions conditioned on goals in visual multi-agent scenarios,” in *Proceedings of International Conference on Computer Vision*, vol. 2, p. 4, 2019.
- [57] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, “Vectormet: Encoding hd maps and agent dynamics from vectorized representation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11525–11533, 2020.
- [58] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl, “Learning by cheating,” in *Conference on Robot Learning*, pp. 66–75, 2020.
- [59] C. Zhao, Q. Sun, C. Zhang, Y. Tang, and F. Qian, “Monocular depth estimation based on deep learning: An overview,” *Science China Technological Sciences*, vol. 63, no. 9, pp. 1612–1627, 2020.
- [60] A. Hu, Z. Murez, N. Mohan, S. Dudas, J. Hawke, V. Badrinarayanan, R. Cipolla, and A. Kendall, “Fiery: Future instance prediction in bird’s-eye view from surround monocular cameras,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15273–15282, 2021.
- [61] W. Luo, B. Yang, and R. Urtasun, “Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3569–3577, 2018.
- [62] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.
- [63] W. Nie, Q. Liang, Y. Wang, X. Wei, and Y. Su, “Mmf: Multimodal information fusion networks for 3d model classification and retrieval,” *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 16, no. 4, pp. 1–22, 2020.
- [64] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [65] H. Shao, L. Wang, R. Chen, H. Li, and Y. Liu, “Safety-enhanced autonomous driving using interpretable sensor fusion transformer,” in *Conference on Robot Learning*, pp. 726–737, PMLR, 2023.
- [66] A. Saha, O. Mendez, C. Russell, and R. Bowden, “Translating images into maps,” in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 9200–9206, IEEE, 2022.
- [67] D. Chen, V. Koltun, and P. Krähenbühl, “Learning to drive from a world on rails,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15590–15599, 2021.
- [68] R. Chekroun, M. Toromanoff, S. Hornauer, and F. Moutarde, “Gri: General reinforced imitation and its application to vision-based autonomous driving,” *arXiv preprint arXiv:2111.08575*, 2021.
- [69] D. Chen and P. Krähenbühl, “Learning from all vehicles,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17222–17231, 2022.