

LMDrive: Closed-Loop End-to-End Driving with Large Language Models

Hao Shao^{1,2} Yuxuan Hu⁵ Letian Wang⁴
Guanglu Song² Steven L. Waslander⁴ Yu Liu^{2,3} ✉ Hongsheng Li^{1,3,5} ✉

¹CUHK MMLab ²SenseTime Research

³Shanghai Artificial Intelligence Laboratory ⁴University of Toronto ⁵CPII under InnoHK

Abstract

Despite significant recent progress in the field of autonomous driving, modern methods still struggle and can incur serious accidents when encountering long-tail unforeseen events and challenging urban scenarios. On the one hand, large language models (LLM) have shown impressive reasoning capabilities that approach “Artificial General Intelligence”. On the other hand, previous autonomous driving methods tend to rely on limited-format inputs (e.g., sensor data and navigation waypoints), restricting the vehicle’s ability to understand language information and interact with humans. To this end, this paper introduces LMDrive, a novel language-guided, end-to-end, closed-loop autonomous driving framework. LMDrive uniquely processes and integrates multi-modal sensor data with natural language instructions, enabling interaction with humans and navigation software in realistic instructional settings. To facilitate research in language-based closed-loop autonomous driving, we also publicly release the corresponding dataset which includes approximately 64K instruction-following data clips, and the LangAuto benchmark that tests the system’s ability to handle complex instructions and challenging driving scenarios. Extensive closed-loop experiments are conducted to demonstrate LMDrive’s effectiveness. To the best of our knowledge, we’re the very first work to leverage LLMs for closed-loop end-to-end autonomous driving. Code is available on our [webpage](#).

1. Introduction

Remarkable progress in autonomous driving has been witnessed in recent years with an increasing number of commercial autonomous vehicles (AVs) deployed on public roads. Generally, state-of-the-art autonomous driving systems can be categorized into two primary approaches: 1) a modular approach where the system is decomposed into several sub-modules such as perception, prediction, and planning, and fixed interfaces are designed to integrate them

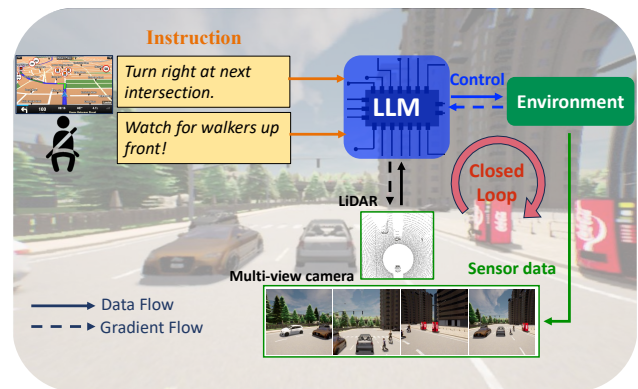


Figure 1. We present LMDrive, the first language-guided closed-loop end-to-end driving framework. LMDrive takes as input the language instruction and multi-modal multi-view sensor data, and outputs control signals in real time to drive in complex scenarios.

together [2, 26, 41, 57]; and 2) an end-to-end approach that directly converts sensor data to control signals via a neural network [18, 39]. While both of these approaches are widely adopted and constantly making breakthroughs on challenging benchmarks, both of them share a limitation in that they solely rely on fixed-format inputs such as the sensor data, target waypoints, and action commands, which restricts the agent’s ability to comprehend multi-modal information and to interact with humans and the environment. On the other hand, large language models (LLMs) have shown an impressive range of capabilities that approach “Artificial General Intelligence.” This encompasses language comprehension, knowledge retrieval, and reasoning. Such capabilities could greatly enhance the safety, controllability, and explainability of autonomous agents. In this work, we seek to answer the question for the first time: *Can we build cognitive autonomous driving systems on top of LLMs, that can interact with human passengers or navigation software simply by natural language?*

Making autonomous systems understand natural language opens profound opportunities for advanced reasoning in complex scenarios and efficient interaction with humans, addressing many previously non-trivial problems. To name

✉ Corresponding author.

a few: 1) in long-tail unforeseen events and challenging urban situations (*e.g.*, complex and dense intersections) where modern AV systems typically struggle [47] or even incur serious accidents [43], the language-aware AVs can easily survive by following navigation instructions from passengers or navigation software. 2) adapting to passengers’ sudden notices (*e.g.*, small objects missed by perception systems) using natural language, previously difficult and rule-dependent.

Toward these appealing properties, many pioneering works have explored the potential of using large language models to enhance the AV system’s reasoning abilities, interpretability, and overall performance in open-loop settings. One of the most common strategies [6, 11, 32, 36] is to 1) first use LLMs to transform the scene perception results and navigation commands into textual descriptions; 2) feed these textural descriptions into LLMs to generate textual driving decisions; and then 3) transfer textual driving decisions into executable control commands. While good preliminary results are shown, this type of approach, where different LLMs tackle sub-tasks individually, is hard to be trained in an end-to-end manner, loses the capability to scale with a large amount of data, and is not robust to perception errors and uncertainties. For example, since the LLMs in the latter two stages do not have access to the sensor data, inaccurate or missed detections in the first stage can lead to large accumulative errors in the latter stages. Towards addressing these issues, end-to-end language-based driving methods [52] have been proposed. However, all of these methods undergo training and evaluation in the open-loop setting, where actions are generated and evaluated against the expert actions, but not executed in the actual environments. Notably, when executing navigation instructions such as “turn right”, the AV agent should not only generate a sequence of actions, but also consider the changes that the actions bring to the environment. The absence of closed-loop evaluation leads to insufficient consideration of critical issues such as cumulative errors, human-robot interaction, and temporal consistency of actions, which makes the resulting methods difficult to scale beyond a short time horizon making them ineffective in actual systems. To the best of our knowledge, there is no existing paper that leverages LLMs for closed-loop end-to-end autonomous driving.

In this work, we introduce LMDrive, an instruction-following multi-modal LLM model for end-to-end closed-loop autonomous driving. The proposed model can process camera-LiDAR sensor data, comprehend driving instructions in natural language, and directly generate vehicle control signals. A pre-trained LLM model is adopted and kept frozen to maintain its reasoning capability. To adapt the LLM for autonomous driving, multiple camera-LiDAR data encoders and learnable input/output adapters are integrated. A pre-training strategy specifically designed for the driving task is also introduced for the multi-modal

vision encoder. To facilitate the training of LMDrive in a closed-loop setting, we develop a language-guided driving dataset based on the CARLA simulator [12], which simulates the dynamic world and realistic challenging scenarios. To better test the driving model’s capability in realistic conversation scenarios, where the instructions might come from humans and navigation software, we 1) consider both navigation instruction and notice instructions; 2) diversify instructions into various phrases via prompt engineering on LLMs; 3) incorporate misleading and unreasonable instructions which are infeasible due to safety concerns or traffic regulations; 4) extend the navigation instructions to include multiple consecutive segments, such as “Take a left here, then another left at the next one”. Besides, we provide the corresponding LangAuto evaluation benchmark and the pre-trained LMDrive model for reproducibility, and we hope these can facilitate further research in end-to-end closed-loop language-based autonomous driving.

To summarize, this paper makes the following contributions:

- We propose a novel end-to-end, closed-loop, language-based autonomous driving framework, LMDrive, which interacts with the dynamic environment via multi-modal multi-view sensor data and natural language instructions.
- We provide a dataset with about 64K data clips, where each clip includes one navigation instruction, several notice instructions, a sequence of multi-modal multi-view sensor data, and control signals. The duration of the clip spans from 2 to 20 seconds.
- We present the benchmark LangAuto for evaluating the autonomous agents that take language instructions as navigation inputs, which include misleading/long instructions and challenging adversarial driving scenarios.
- We conduct extensive closed-loop experiments to demonstrate the effectiveness of the proposed framework, and analyze different components of LMDrive to shed light on continuing research along this direction.

2. Related Works

2.1. End-to-End Autonomous Driving

Much progress [7, 8] has been achieved recently in the field of end-to-end autonomous driving. UniAD [18] devised a framework that incorporates full-stack driving tasks and utilizes query-unified interfaces to communicate between different tasks. ThinkTwice [22] designed a Look Module to retrieve information from critical regions and utilize the features to refine the coarse prediction. ReasonNet [39] exploited temporal and global information of the driving scene to improve perception performance and benefit occlusion detection. InterFuser [38] proposed a transformer-based framework to fully fuse and process information from multi-modal multi-view sensors for comprehensive scene understanding. TCP [51] proposed an ap-

proach that integrates the two branches for trajectory planning and direct control by involving a novel multi-step prediction. LAV [3] introduced supervisory tasks for learning viewpoint-invariant representation, offering richer supervision at training and aiding complex reasoning during inference. Besides imitation training, various approaches incorporate reinforcement learning strategies. Latent DRL [44] trained in a supervised way to get a latent representation of the observation and conducts reinforcement learning using the representation as the input. Roach [54] employed a reinforcement learning agent with privileged access to environmental information and distilled a model as the final agent. ASAPRL [50] and TaEcRL [56] exploited abstracted skills to effectively improve reinforcement learning efficiency and the final performance by facilitating effective exploration and reward signaling. ScenerioNet [25] and TrafficGen [14] enhanced driving performance via generating and testing on diverse driving scenarios. However, these end-to-end methods lack the ability to verbally or textually interact with humans (passengers), and usually have low explainability in the decision-making process.

2.2. LLMs in Driving Tasks

Emerging advancements in large language models (LLMs) [10, 19, 21, 45, 46] have been witnessed over the last few months. Furthermore, vision large language models (VLLM) [15, 28, 40] further introduce vision encoders and open the doors for LLMs to interpret not only textual data but also images and data in other modalities [5, 29, 58]. In the field of autonomous driving (AD), recent research has integrated LLMs into the AD system for better explainability and natural interaction with humans. Some studies adopt the visual language model approach, which can handle multi-modal input data and provide the textual description as well as the control signal for the driving scenarios. For example, DRIVEGPT4 [52] proposed a multimodal LLM framework, which takes a sequence of frames as input, then generates responses to human inquiries and predicts control signals for the next step. However, since the framework lacks an input command, the predicted control can not follow the specific navigation command, which denotes that the framework is hard to deploy in real scenarios. Meanwhile, more researchers focus on transforming the driving situations into textual descriptions as the input for the LLM, for directly interpreting and reasoning about comprehensive driving situations. In this thread of works, GPT-Driver [32] reformulated motion planning as the task of natural language modeling by converting heterogeneous scene input to language tokens. LanguageMPC [36] leveraged a LLM to reason the complex scenarios and output high-level driving decisions. Then the method tunes a parameter matrix to convert the decision into the low-level control signals. LLM-Driver [6] utilized the numeric vector as the



Figure 2. Two examples of the collected data with corresponding labeled navigation instructions and optional notice instructions.




Figure 3. Distribution of parsed clips in terms of clip length and the corresponding navigation instruction type.

input modality and fused vectorized object-level 2D scene representation to enable the LLM to answer the questions based on the current environment.

However, this line of work only considered the driving problem in the open-loop settings, and ignored questions such as cumulative error, temporal action consistency, and end-to-end trainability, which are critical for bringing the models into actual closed-loop driving tasks. To the best of our knowledge, we are the first language-based end-to-end autonomous driving method in the closed-loop setting. Relevant datasets, benchmarks, and trained models are also open-sourced to facilitate further research.

3. Dataset generation

We aim to develop an intelligent driving agent that can generate driving actions based on three sources of input: 1) sensor data (multi-view camera and LiDAR), so that the agent can generate actions that are aware of and compliant with the current scene; 2) navigation instructions (*e.g.*, lane changing, turning), so that the agent can drive to meet the requirement in natural language (instruction from humans or navigation software); and 3) human notice instruction, so that the agent can interact with humans and adapt to human’s suggestions and preferences (*e.g.*, pay attention to adversarial events, deal with long-tail events, *etc.*). In this section, we describe how to generate the multi-modal dataset needed to train the agent, and the prompt design for the navigation instruction and human notice instruction. Specifically, we choose the CARLA [12] as the simulator, because it can simulate a realistic dynamic closed-loop world and it is widely adopted in the field of end-to-end autonomous driving. The data collection consists of two stages: 1) col-

Type	Three randomly chosen instructions of each instruction type
Follow	Maintain your current course until the upcoming intersection. In [x] meters, switch to left lane. Ease on to the left and get set to join the highway.
Turn	After [x] meters, take a left. At the next intersection, just keep heading straight, no turn. You'll be turning left at the next T-junction, alright?
Others	Feel free to start driving. Slow down now. Head to the point, next one's [x] meters ahead, [y] meters left/right.
Notice	Watch for walkers up front. Just a heads up, there's a bike ahead. Please be aware of the red traffic signal directly in front of you.

Table 1. Examples of considered navigation instructions (follow, turn, others) and notice instructions. $[x]$ and $[y]$ represent the float number for a specific distance.

lecting sensor data and control signals with an expert agent; and 2) parsing and labeling collected data with instructions.

Sensor and control data collection. We utilize a rule-based expert agent [38] to create a dataset including about 3M driving frames. Since the expert agent can access the privileged information in the CARLA, this dataset will include camera data, LiDAR data, and control actions for each frame. To enhance the diversity of the collected dataset, the agent runs on 2.5k routes, 8 towns, and 21 kinds of environmental conditions (e.g., weather, time of the day). We use four RGB cameras (left, front, right, rear) and one LiDAR. The side cameras are angled at 60° . Besides, we center-crop the front image as an additional focus-view image to capture the status of the distant traffic light. The LiDAR has 64 channels and generates 600K points per second.

Parsing and language annotation. In the second stage, we parse the collected data into clips, and label each clip with proper navigation instructions and optional notice instructions. The parsing process takes a sequence of frames as input, and segments these frames into clips, where each clip corresponds to one navigation instruction. For instance, if the agent started to turn left at frame T_0 and ended at frame T_n , we will label (T_0, T_n) as a new clip with the instruction “*Hang a left at the next crossroads*”. Besides, if an adversarial event¹ occurs at time T_a , we will add one notice instruction into this clip, simulating a real-life scenario where a passenger or a side assistance system would communicate with the driver when an emergency happens. As shown in Figure 2, each clip includes sensor data, control signals, the corresponding navigation instruction, and optional notice instructions. The distribution of the parsed clips in terms of clip length and corresponding instruction is shown in Figure 3. In our dataset, we collect 64K parsed clips and 464K notice instructions.

Instruction design. We consider three types of navigation

¹The adversarial events include bad road conditions, the front vehicle’s sudden brake, unexpected entities rushing into the road from occluded regions, vehicles running a red traffic light, etc.

instructions (follow, turn, and others) along with one type of notice instruction, consisting of a total of 56 different instructions. Table 1 shows some examples and the full list can be found in supplementary material. To enable the agent to drive in realistic instructional settings where the instructions come from navigation software or humans, we

- *Diversifying the instructions:* Considering the inherent richness of natural language, for each type of instruction, we utilize ChatGPT to generate 8 different variants, each carrying the same semantic meaning but varying in phrasing. This enables more comprehensive coverage and flexibility in language interpretation, accommodating the diverse ways the same instruction can be conveyed.
- *Incorporating misleading instructions:* in real-world cases, the navigation software or passengers may give misleading instructions to the AV that violate traffic rules or raise safety concerns. For example, on a single-lane road, following an instruction “*Change to left lane*” is dangerous. To improve the robustness of our model against misleading instructions, we simulate these scenarios and add them to our dataset.
- *Connecting multiple instructions:* In many cases, the instructions may consist of two to three consecutive instructions, such as “*Turn right at this intersection, then go straight to the next intersection and turn right again.*” We also construct some consecutive complex instruction data to simulate real navigation-based driving scenarios.

4. LMDrive methodology

In this work, we propose LMDrive, a framework that can understand and follow high-level driving instructions through natural language. As illustrated in Figure 4, LMDrive is composed of two major components: 1) a vision encoder that processes multi-view multi-modal sensor data (camera and LiDAR) for scene understanding and generating visual tokens; 2) a LLM and its associated component (tokenizer, Q-Former, and adapters) that takes in the visual tokens and language instruction, to predict the control signal and whether the given instruction is completed. We will introduce the vision encoder in Section 4.1, and the language model with its associated components in Section 4.2. Finally, we describe the training details in Section 4.3.

4.1. Vision encoder

In the visual language community [5, 29, 58], the most common way to align vision and language could be using pre-trained CLIP models [35] to encode image features. However, the large flops and parameter size of the CLIP models increase the difficulty of its deployment in AV systems. Also, AV perception systems are usually in 3D to include LiDAR input. Hence, inspired by InterFuser [38] and TF++ [20], we design a multi-view multi-modality vision encoder to encode/fuse the sensor data. As shown in Figure 5, the vision encoder consists of the sensor encoding

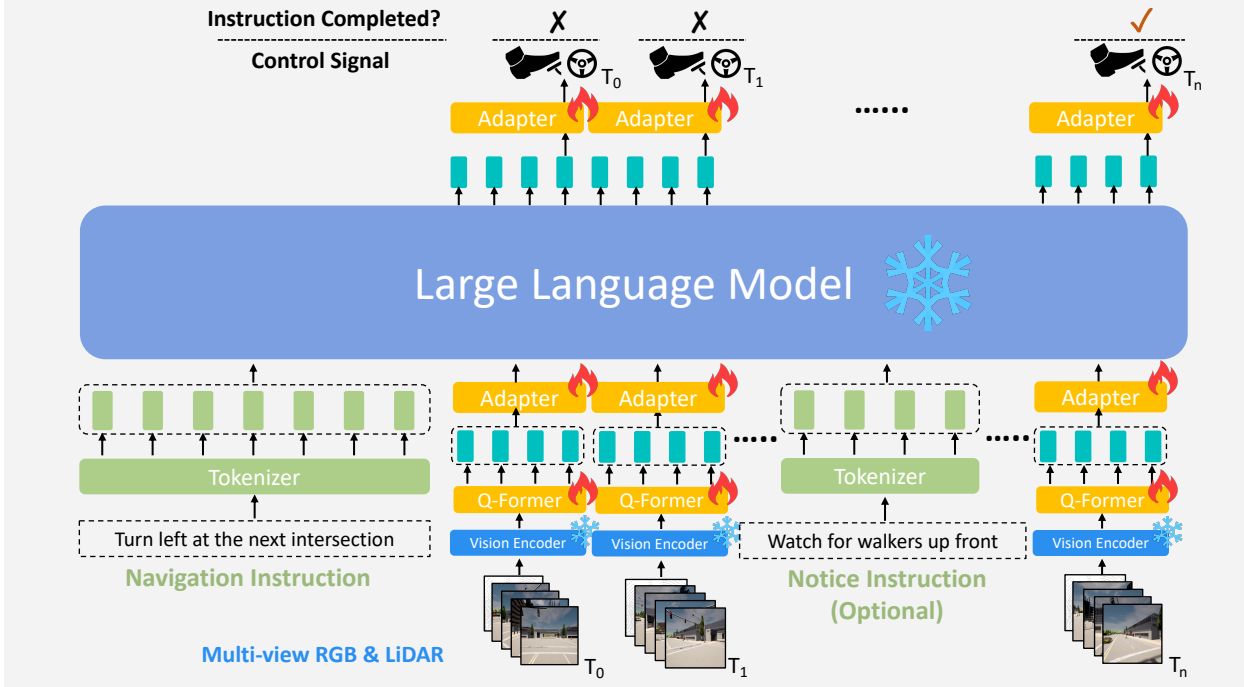


Figure 4. The structure of the proposed LMDrive model, which consists of two major components: 1) a vision encoder that processes multi-view multi-modal sensor data (camera and LiDAR) for scene understanding and generating visual tokens; 2) a large language model and its associated component (tokenizer, Q-Former, and adapters) that processes all the historic visual tokens and the language instructions (navigation instruction and optional notice instruction), to predict the control signal and whether the given instruction is completed.

part which encodes image and LiDAR input respectively, and a BEV decoder that fuses image and point cloud features to generate visual tokens which are then passed to the language model. Notably, the vision encoder is pre-trained on perception tasks by adding additional prediction heads, then the encoder is frozen for later use by the LLM.

Sensor encoding. For each image input, we apply a 2D backbone ResNet [17] to extract the image feature map. The feature map is then flattened into one-dimensional tokens. For a comprehensive understanding of the global context from multiple viewpoints, tokens from different views will be fused by a standard K_{enc} -layer transformer encoder, each layer containing Multi-Headed Self-Attention [48], MLP blocks and layer normalization [1]. For the LiDAR input, we adopt a 3D backbone PointPillars [23] to process the raw point cloud data into ego-centered LiDAR features, where each pillar encompasses points within a $0.25\text{m} \times 0.25\text{m}$ area. PointNet [34] is then used to aggregate features and downsample the feature map to $C \times H \times W$, which subsequently serve as BEV queries.

BEV decoder. The encoded sensor features above are then passed into the BEV decoder to generate visual tokens. Specifically, the BEV decoder is designed as a standard transformer with K_{dec} layers. The BEV point cloud features are fed into the BEV decoder as $H \times W$ queries to attend to the multi-view image features and generate BEV

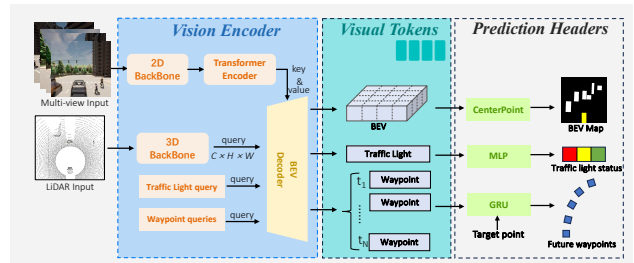


Figure 5. The detailed structure of the vision encoder, which takes as input the multi-view multi-modality sensor data. In the pre-training stage, the vision encoder is appended with prediction headers to perform pre-training tasks (object detection, traffic light status classification, and future waypoint prediction). In the instruction-finetuning stage and inference stage, the prediction headers are discarded, and the vision encoder is frozen to generate visual tokens to feed into the LLM.

tokens. We also feed N learnable queries and 1 learnable query into the BEV decoder to generate N waypoint point tokens and 1 traffic light token respectively. Thus the three types of visual tokens (BEV, waypoint, and traffic light) will contain rich scene information and will be then presented to the large language model.

Pre-training with prediction headers. we consider three vision encoder pre-training tasks: object detection, future waypoint prediction, and traffic light status classification. For object detection, the BEV tokens will pass through a

one-stage CenterPoint [53] to predict the bounding boxes and velocity of the objects in an $H_m \times W_m$ area. For waypoint prediction, we pass the N waypoint tokens along with the navigation waypoint into the GRU network [9] sequentially to predict N future waypoint. For the traffic light status classification, a 2-layer MLP is applied to the traffic light token. Three corresponding loss terms are considered: 1) the detection loss as in InterFuser [39]; 2) the l_1 waypoint loss; and 3) the cross-entropy traffic light state loss. Note that these prediction headers are only used in the pre-training of the vision encoder, and will be discarded in the training of the LLMs and inference of the whole model.

4.2. LLM for instruction-following auto driving

As illustrated in Figure 4, in our framework, the LLM functions as the “brain” throughout the entire driving procedure, processing sensor tokens generated by the frozen vision encoder for each frame, comprehending natural language instructions, generating the necessary control signal, and predicting whether the given instruction is completed. Specifically, we choose LLaMA [45] as the language backbone, which has been widely adopted in many language [16, 55] and vision [29, 58] instruction-tuning models. We also have three associated components to bridge LLM with instruction, visual information input, and action prediction: 1) a tokenizer, 2) a Q-Former, 3) two adapters.

Instruction and visual tokenization. Given the navigation instruction and optional notice instruction, we apply the LLaMA tokenizer [45] to convert the instruction into textual tokens. Note that the duration of executing one instruction would span from a few seconds to a few minutes, and our model is deployed in the closed-loop setting. Thus at each frame, we utilize all historic sensor information (with maximum limit T_{max}) to depress cumulative error and improve the temporal consistency of the model. Specifically, for each frame’s multi-view multi-modality sensor input, we utilize the vision encoder pre-trained in the previous section to generate visual tokens ($H \times W$ BEV tokens, N waypoint tokens, and one traffic light token). However, the number of visual tokens (e.g., 406 tokens for each frame) quickly grows too large for the LLM because usually hundreds of frames are needed to complete one instruction. To overcome this, we follow BLIP-2 [24] to use the Q-Former to reduce the number of visual tokens. Specifically, for each frame, we employ M learnable queries to attend to the visual tokens via cross-attention layers, which can reduce each frame’s visual token number to M . Subsequently, we use a 2-layer MLP adapter to convert the tokens extracted by the Q-Former to share the same dimension as the language token, which can then be fed into the LLM.

Action prediction. After receiving a sequence of instructional and visual tokens, the LLM predicts the action tokens. One another 2-layer MLP adapter is then applied to predict

future waypoints, as well as a flag to indicate whether the given instruction has been completed. Note that to enhance the supervision signal, we will also conduct prediction for every historic frame during training, and only the prediction of the latest frame will be executed at inference time. To get the final control signal, which includes braking, throttling, and steering, following LBC [4], we use two PID controllers for latitudinal and longitudinal control to track the heading and velocity of predicted waypoints respectively.

Training objectives. When finetuning the LLM and its associated components, we consider two loss terms: 1) the l_1 waypoint loss; 2) the classification loss (cross-entropy), indicating if the current frame completes the instruction.

4.3. Training details

LMDrive’s training consists of two stages: 1) the vision encoder pre-training stage; and 2) the instruction-finetuning stage, to align the instruction/vision and control signal.

Vision encoder pre-training stage. The vision encoder takes a single frame’s sensor data as input, and we use the dataset collected in Section 3 for training. Specifically, since the instruction annotation process will drop some frames, we use the raw dataset before the instruction annotation for the vision encoder pre-training, which includes data of around 3M frames. Only the vision encoder is pre-trained with perception tasks for scene understanding.

Instruction-finetuning stage. The entire system is trained for end-to-end autonomous driving under the guidance of the instruction, where the Q-Former and Adapters are trainable and the other components are frozen. While our LMDrive takes a sequence of frames as input, during training we set a fixed sequence length T_{max} for building up batch data. The training utilizes the instruction-following data generated in Section 3. To enable the model to reject misleading instructions, we label the corresponding data as ‘completed’ after the misleading instruction is given for about 1 second. Since the dataset is collected at a high frequency (~ 10 Hz), the data in adjacent frames are highly similar. To encourage efficient training, following video prediction methods [13, 37, 49], we sample the training frames in a fixed interval, and apply temporal augmentation which randomly shifts the training frames either forward or backward, with the random shift less than the fixed interval.

5. LangAuto Benchmark

We propose the LangAuto (Language-guided Autonomous Driving) CARLA benchmark, the first benchmark that evaluates closed-loop driving performance under language instructions. Compared with the previous CARLA benchmark, Town05 [33] and Longest6 [8] that navigate the agent with discrete driving commands or target waypoints, our benchmark only provides the AV with a navigation instruction and optional notice instructions in natural language.

LLM Backbone	LangAuto			LangAuto-Short			LangAuto-Tiny		
	DS \uparrow	RC \uparrow	IS \uparrow	DS \uparrow	RC \uparrow	IS \uparrow	DS \uparrow	RC \uparrow	IS \uparrow
Random Init.	10.7 \pm 3.8	16.2 \pm 4.9	0.63 \pm 0.04	14.2 \pm 4.4	20.1 \pm 4.4	0.72 \pm 0.04	20.1 \pm 4.1	24.7 \pm 5.1	0.75 \pm 0.03
LLaMA [45]	31.3 \pm 1.5	37.1 \pm 1.6	0.82 \pm 0.01	42.8 \pm 7.2	49.1 \pm 8.5	0.87\pm0.03	52.2 \pm 5.3	57.8 \pm 8.0	0.91\pm0.05
LLaMA2 [46]	32.8 \pm 2.1	40.1 \pm 2.2	0.81 \pm 0.02	44.8 \pm 6.2	53.5 \pm 5.5	0.84 \pm 0.02	56.1 \pm 4.1	64.2 \pm 4.7	0.87 \pm 0.04
Vicuna [55]	33.5 \pm 1.9	39.3 \pm 1.9	0.83 \pm 0.02	45.3 \pm 4.9	54.3 \pm 3.9	0.83 \pm 0.03	55.5 \pm 3.9	63.1 \pm 4.2	0.88 \pm 0.04
Vicuna-v1.5 [55]	34.0 \pm 3.8	39.0 \pm 3.3	0.85\pm0.06	47.0 \pm 4.3	56.5 \pm 2.4	0.83 \pm 0.04	59.0 \pm 2.6	69.9 \pm 2.3	0.84 \pm 0.02
LLaVA-v1.5 [27]	36.2\pm2.3	46.5\pm4.3	0.81 \pm 0.03	50.6\pm1.7	60.0\pm3.4	0.84 \pm 0.04	66.5\pm3.6	77.9\pm2.3	0.85 \pm 0.02

Table 2. Performance comparison of 6 LLM backbones on the LangAuto benchmark. We report the metrics for 3 evaluation runs.

Module design	DS \uparrow	RC \uparrow	IS \uparrow
Baseline (LLaVA-v1.5)	36.2\pm2.3	46.5\pm4.3	0.81\pm0.03
w/o Q-Former	31.7 \pm 3.5	41.2 \pm 4.4	0.79 \pm 0.02
w/o using BEV tokens	33.9 \pm 3.9	45.9 \pm 5.1	0.72 \pm 0.03
w/o visual pre-training	16.9 \pm 5.1	24.1 \pm 4.7	0.70 \pm 0.04

Table 3. Ablation study on the module design.

Specifically, the LangAuto benchmark covers all 8 publicly available towns in CARLA to include various scenarios (e.g., highways, intersections, roundabouts). We also consider 16 kinds of environmental conditions, encompassing the combinations of 7 weather conditions (Clear, Cloudy, Wet, MidRain, WetCloudy, HardRain, SoftRain) and 3 daylight conditions (Night, Noon, Sunset). Besides, LangAuto consists of three tracks to fully test the agent’s instruction-following abilities:

- LangAuto track: For each route, navigation instructions are given and updated to the agent based on the agent’s current position. We also divide this track into three sub-tracks with different route lengths, to better distinguish the performance. LangAuto where the routes are longer than 500 meters, LangAuto-Short where the route length is between 150 and 500 meters, and LangAuto-Tiny where the route length is shorter than 150 meters.
- LangAuto-Notice track: Based on the LangAuto track, we additional add notice instructions to the agent. This setting simulates real cases where passengers or other assistance systems can give real-time notice in long-trail complex or adversarial scenarios, which is usually hard for the AV system to handle by itself. Ideally, the agent that can comprehend and leverage instruction can achieve better performance.
- LangAuto-Sequential track: Based on the LangAuto track, we merge 10% of consecutive 2 to 3 instructions into a single long instruction. This setting mimics the realistic scenarios where the multi-sentence instructions come from the passengers or navigation software.

Note that misleading instructions will be randomly ($\sim 5\%$) and intermittently given to the driving agent, which lasts for a certain duration (1-2 seconds). The driving agent is expected to reject these misleading instructions and execute safe actions that are compliant to the current scene, until the next correct instruction is spawned.

Metrics. We consider three major metrics introduced by the

CARLA LeaderBoard [42]: route completion (RC), infraction score (IS), and driving score (DS). The route completion refers to the percentage of the total route length that has been completed. It only takes into account the distance traveled along the predetermined route, where each segment of the predetermined route corresponds to a navigation instruction. If the agent deviates too far from the route, the agent is regarded as violating the instruction, and this episode is marked as a failure and terminated. The infraction score measures infractions triggered by the agent. When collisions or traffic rule violations occur, the infraction score is decayed by a corresponding discount factor. The driving score is the product of the route completion ratio and the infraction score, describing both driving progress and safety. It is generally recognized as the primary ranking metric.

6. Experiments

6.1. Experiment Setup

We implement and evaluate our approach on the open-source CARLA simulator of version 0.9.10.1 [12]. For the 2D backbone and 3D backbone of the vision encoder, ResNet-50 [17] backbone is pre-trained on ImageNet, and PointPillars [23] is trained from scratch. C, H, W, N are set as 256, 50, 50, 5 respectively. For the Q-former token number M , we empirically found that $M = 4$ tokens achieve decent performance. During the instruction-finetuning stage, we sample the training frames in a fixed interval of 2. Because each parsed clip typically contains multiple notice instructions, to avoid overfitting to notice following and generating over-conservative behaviors, during training we randomly removed notices for 75% of the clips. For the remaining clips, we ensure that a maximum of one notice is included. We refer readers to the appendix for more details.

6.2. Quantitative Results

LLM backbones. In Table 2, we evaluate our method with different pretrained or randomly initialized LLM models of 7B parameters: LLaMA [45] and LLaMA2 [46] were pretrained with large public language datasets; Vicuna [55] and Vicuna-v1.5 [55] are additionally finetuned with human conversation data; LLaVA-v1.5 [27] incorporate multi-modal data (e.g., language, images) for training. We observe that LLaVA-v1.5 [27] surpasses the other

LLM Backbone	Benchmark Type	Infraction Score \uparrow	Vehicle Collisions \downarrow	Pedestrian Collisions \downarrow	Layout Collisions \downarrow	Red light Violations \downarrow	Offroad Infractions \downarrow	Blocked Infractions \downarrow
LLaVA-v1.5	LangAuto	0.81	0.33	0.03	0.50	0.92	0.36	0.22
	LangAuto-Notice	0.87	0.17	0.02	0.31	0.50	0.17	0.27
Vicuna-v1.5	LangAuto	0.85	0.30	0.03	0.43	1.18	0.24	0.19
	LangAuto-Notice	0.91	0.15	0.01	0.28	0.56	0.26	0.19

Table 4. Performance comparison on LangAuto and LangAuto-Notice benchmarks. The metrics (except infraction score) are normalized by the driven distance (km).

LLM Backbone	Benchmark Type	DS \uparrow	RC \uparrow	IS \uparrow
LLaVA-v1.5	LangAuto	36.2	46.5	0.81
	LangAuto-Sequential	34.0	43.7	0.81
Vicuna-v1.5	LangAuto	34.0	39.0	0.85
	LangAuto-Sequential	31.9	37.1	0.84

Table 5. Performance comparison on LangAuto and LangAuto-Sequential benchmarks.

LLM models, which demonstrates the importance of adopting pretrained multi-modal LLMs for instruction-following autonomous driving. Besides, the models finetuned with instruction data perform better than others (Vicuna-v1.5 > LLaMA2 \approx Vicuna > LLaMA). We also test a randomly initialized 7B LLM model, which struggles to drive properly with the same amount of training data, which demonstrates the necessity of finetuning pretrained LLM for instruction-following driving.

Module Design. In Table 3, we conduct ablation studies on different components. First, instead of reducing the number of BEV tokens with the Q-Former, we directly down-sample the BEV features to 4×4 , then feed them into the LLM (denoted as “w/o Q-Former”). The average driving score dropped from 36.2 to 31.7. Second, excluding the BEV tokens input into LLM decoder (denoted as “w/o using BEV tokens”) results in a decreased infraction score (0.81 to 0.72). One explanation is that the BEV tokens are important for detecting and reasoning the surrounding obstacles and road structures. Third, we remove the pre-training stage of the vision encoder and directly train it from scratch in the instruction-finetuning stage. The driving score drops to 16.9 (denoted as “w/o visual pre-training”), demonstrating the importance of our proposed vision encoder pre-training.

LangAuto-Notice Benchmark. The LangAuto-Notice benchmark provides some notice instructions to the agent when adversarial events happen. As shown in Table 4, our agents can effectively leverage the information of the notice in real-time, resulting in a significant decrease in both collisions and traffic rule violations.

LangAuto-Sequential Benchmark. In Table 5, we demonstrate the effect of the LangAuto-Sequential benchmark where some consecutive 2 to 3 navigation instructions are merged together to a long and complex instruction. This setting additionally requires the agent to be able to be temporally aware of which instruction has been completed and

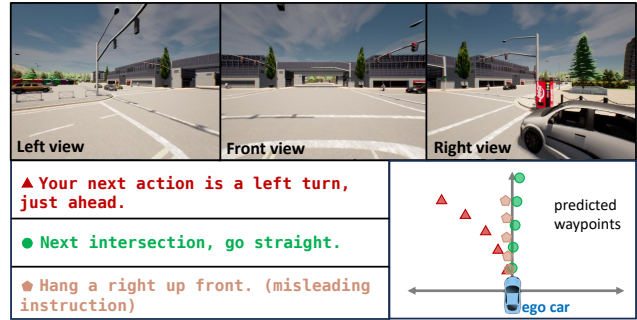


Figure 6. An example of how our LMDrive predicts future waypoints, given sensor inputs and varied navigational instructions. Under the first two instructions, our model predicts different waypoints accordingly. The third instruction is a misleading one (turn right on a left-only lane). The model appropriately rejects the incorrect instruction, generating a slower speed and a safe path that is compliant with the scenario.

which has not. Our agents based on LLaVA and Vicuna both have a performance drop.

7. Conclusion

In this paper, we introduced LMDrive, a language-guided, end-to-end, closed-loop autonomous driving framework. LMDrive incorporates natural language instructions along with sensor data, enabling human-like interaction and navigation in complex driving scenarios. We also propose the language-guided driving dataset, comprising around 64K multi-modal data clips along with corresponding navigation instructions. We established the LangAuto benchmark for evaluating autonomous driving systems considering natural language instructions. The effectiveness of LMDrive was demonstrated through extensive closed-loop experiments, underlining the potential of improving the interaction of autonomous vehicles with humans and the environment. Our work paves the way for future advancements in language-based autonomous driving.

Acknowledgement. This project is funded in part by National Key R&D Program of China Project 2022ZD0161100, by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK, by General Research Fund of Hong Kong RGC Project 14204021. Hongsheng Li is a PI of CPII under the InnoHK.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. **5**
- [2] Sergio Casas, Cole Gulino, Simon Suo, Katie Luo, Renjie Liao, and Raquel Urtasun. Implicit latent variable model for scene-consistent motion forecasting. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16*, pages 624–641. Springer, 2020. **1**
- [3] Dian Chen and Philipp Krähenbühl. Learning from all vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17222–17231, 2022. **3, 1**
- [4] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *Conference on Robot Learning*, pages 66–75. PMLR, 2020. **6**
- [5] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multi-modal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023. **3, 4**
- [6] Long Chen, Oleg Sinavski, Jan Hünemann, Alice Karnsund, Andrew James Willmott, Danny Birch, Daniel Maund, and Jamie Shotton. Driving with llms: Fusing object-level vector modality for explainable autonomous driving. *arXiv preprint arXiv:2310.01957*, 2023. **2, 3**
- [7] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: Challenges and frontiers. *arXiv preprint arXiv:2306.16927*, 2023. **2**
- [8] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. **2, 6**
- [9] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. **6**
- [10] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022. **3**
- [11] DriveLM Contributors. Drivelm: Drive on language. <https://github.com/OpenDriveLab/DriveLM>, 2023. **2**
- [12] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. **2, 3, 7**
- [13] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. **6**
- [14] Lan Feng, Quanyi Li, Zhenghao Peng, Shuhan Tan, and Bolei Zhou. Trafficgen: Learning to generate diverse and realistic traffic scenarios. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3567–3575. IEEE, 2023. **3**
- [15] Peng Gao, Renrui Zhang, Chris Liu, Longtian Qiu, Siyuan Huang, Weifeng Lin, Shitian Zhao, Shijie Geng, Ziyi Lin, Peng Jin, Kaipeng Zhang, Wenqi Shao, Chao Xu, Conghui He, Junjun He, Hao Shao, Pan Lu, Hongsheng Li, and Yu Qiao. Sphinx-x: Scaling data and parameters for a family of multi-modal large language models, 2024. **3**
- [16] Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, Pieter Abbeel, Sergey Levine, and Dawn Song. Koala: A dialogue model for academic research. *Blog post, April*, 1, 2023. **6**
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **5, 7, 1**
- [18] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, et al. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17853–17862, 2023. **1, 2**
- [19] Bernhard Jaeger, Kashyap Chitta, and Andreas Geiger. Hidden biases of end-to-end driving models. 2023. **3**
- [20] Bernhard Jaeger, Kashyap Chitta, and Andreas Geiger. Hidden biases of end-to-end driving models. *arXiv preprint arXiv:2306.07957*, 2023. **4**
- [21] Xiaosong Jia, Yulu Gao, Li Chen, Junchi Yan, Patrick Langechuan Liu, and Hongyang Li. Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7953–7963, 2023. **3**
- [22] Xiaosong Jia, Penghao Wu, Li Chen, Jiangwei Xie, Conghui He, Junchi Yan, and Hongyang Li. Think twice before driving: Towards scalable decoders for end-to-end autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21983–21994, 2023. **2**
- [23] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. **5, 7**
- [24] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. **6, 1**
- [25] Quanyi Li, Zhenghao Mark Peng, Lan Feng, Zhizheng Liu, Chenda Duan, Wenjie Mo, and Bolei Zhou. Scenarionet: Open-source platform for large-scale traffic scenario simulation and modeling. *Advances in neural information processing systems*, 36, 2024. **3**
- [26] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, pages 1–18. Springer, 2022. **1**

- [27] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. 7
- [28] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. 3
- [29] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 3, 4, 6
- [30] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 1
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 1
- [32] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023. 2, 3
- [33] Aditya Prakash, Kashyap Chitta, and Andreas Geiger. Multi-modal fusion transformer for end-to-end autonomous driving. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
- [34] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018. 5, 1
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 4
- [36] Hao Sha, Yao Mu, Yuxuan Jiang, Li Chen, Chenfeng Xu, Ping Luo, Shengbo Eben Li, Masayoshi Tomizuka, Wei Zhan, and Mingyu Ding. Languagempc: Large language models as decision makers for autonomous driving. *arXiv preprint arXiv:2310.03026*, 2023. 2, 3
- [37] Hao Shao, Shengju Qian, and Yu Liu. Temporal interlacing network. *AAAI*, 2020. 6
- [38] Hao Shao, Letian Wang, Ruobing Chen, Hongsheng Li, and Yu Liu. Safety-enhanced autonomous driving using interpretable sensor fusion transformer. In *Conference on Robot Learning*, pages 726–737. PMLR, 2023. 2, 4
- [39] Hao Shao, Letian Wang, Ruobing Chen, Steven L Waslander, Hongsheng Li, and Yu Liu. Reasonnet: End-to-end driving with temporal and global reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13723–13733, 2023. 1, 2, 6
- [40] Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. Visual cot: Unleashing chain-of-thought reasoning in multi-modal language models, 2024. 3
- [41] Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. Motion transformer with global intention localization and local movement refinement. *Advances in Neural Information Processing Systems*, 35:6531–6543, 2022. 1
- [42] CARLA team. Carla autonomous driving leaderboard. <https://leaderboard.carla.org/>, 2020. Accessed: 2021-02-11. 7
- [43] Trisha Thadani. Cruise recalls all its driverless cars after pedestrian hit and dragged. *The Washintong Post*. 2
- [44] Marin Toromanoff, Emilie Wirbel, and Fabien 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*, pages 7153–7162, 2020. 3
- [45] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 3, 6, 7
- [46] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 3, 7
- [47] Jordan Valinsky. ‘complete meltdown’: Driverless cars in san francisco stall causing a traffic jam. *CNN Business*. 2
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 5
- [49] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, pages 20–36. Springer, 2016. 6
- [50] Letian Wang, Jie Liu, Hao Shao, Wenshuo Wang, Ruobing Chen, Yu Liu, and Steven L Waslander. Efficient reinforcement learning for autonomous driving with parameterized skills and priors. *arXiv preprint arXiv:2305.04412*, 2023. 3
- [51] Penghao Wu, Xiaosong Jia, Li Chen, Junchi Yan, Hongyang Li, and Yu Qiao. Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline. *Advances in Neural Information Processing Systems*, 35:6119–6132, 2022. 2
- [52] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kenneth KY Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv preprint arXiv:2310.01412*, 2023. 2, 3
- [53] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11784–11793, 2021. 6
- [54] Zhejun Zhang, Alexander Liniger, Dengxin Dai, Fisher Yu, and Luc Van Gool. End-to-end urban driving by imitating a reinforcement learning coach. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15222–15232, 2021. 3
- [55] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhaghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena, 2023. 6, 7

- [56] Tong Zhou, Letian Wang, Ruobing Chen, Wenshuo Wang, and Yu Liu. Accelerating reinforcement learning for autonomous driving using task-agnostic and ego-centric motion skills. *arXiv preprint arXiv:2209.12072*, 2022. [3](#)
- [57] Yang Zhou, Hao Shao, Letian Wang, Steven L. Waslander, Hongsheng Li, and Yu Liu. Smartrefine: A scenario-adaptive refinement framework for efficient motion prediction, 2024. [1](#)
- [58] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. [3](#), [4](#), [6](#)