

# A Survey of Large Language Models for Autonomous Driving

Zhenjie Yang<sup>1\*</sup>, Xiaosong Jia<sup>1\*</sup>, Hongyang Li<sup>1,2</sup>, Junchi Yan<sup>1†</sup>

<sup>1</sup> Department of Computer Science and Engineering, Shanghai Jiao Tong University

<sup>2</sup> OpenDriveLab

{yangzhenjie, jiaxiaosong, hongyangli, yanjunchi}@sjtu.edu.cn

\* Equal contributions, order decided by a coin toss †Correspondence author

## Abstract

Autonomous driving technology, a catalyst for revolutionizing transportation and urban mobility, has the tend to transition from rule-based systems to data-driven strategies. Traditional module-based systems are constrained by cumulative errors among cascaded modules and inflexible pre-set rules. In contrast, end-to-end autonomous driving systems have the potential to avoid error accumulation due to their fully data-driven training process, although they often lack transparency due to their “black box” nature, complicating the validation and traceability of decisions. Recently, large language models (LLMs) have demonstrated abilities including understanding context, logical reasoning, and generating answers. A natural thought is to utilize these abilities to empower autonomous driving. By combining LLM with foundation vision models, it could open the door to open-world understanding, reasoning, and few-shot learning, which current autonomous driving systems are lacking. In this paper, we systematically review a research line about *Large Language Models for Autonomous Driving (LLM4AD)*. This study evaluates the current state of technological advancements, distinctly outlining the principal challenges and prospective directions for the field. For the convenience of researchers in academia and industry, we provide real-time updates on the latest advances in the field as well as relevant open-source resources via the designated link: <https://github.com/Thinklab-SJTU/Awesome-LLM4AD>.

## 1 Introduction

Autonomous driving is rapidly reshaping our understanding of transportation, heralding a new era of technological revolution. This transformation means not only the future of transportation but also a fundamental shift across various industries. In conventional autonomous driving systems, algorithms typically adopt the modular design [Liang *et al.*, 2020; Luo *et al.*, 2018; Sadat *et al.*, 2020], with separate components responsible for critical tasks such as perception [Li *et al.*, 2022c; Liu *et al.*, 2023d], prediction [Shi *et al.*, 2022;

Jia *et al.*, 2022a; Jia *et al.*, 2023b], and planning [Treiber *et al.*, 2000; Dauner *et al.*, 2023]. Specifically, the perception component handles object detection [Li *et al.*, 2022c; Liu *et al.*, 2023d], tracking [Zeng *et al.*, 2022], and sophisticated semantic segmentation tasks [Cheng *et al.*, 2022]. The prediction component analyzes the external environment [Jia *et al.*, 2021] and estimates the future states of the surrounding agents [Jia *et al.*, 2022b]. The planning component, often reliant on rule-based decision algorithms [Treiber *et al.*, 2000], determines the optimal and safest route to a predetermined destination. While the module-based approach provides reliability and enhanced security in a variety of scenarios, it also presents challenges. The decoupled design between system components may lead to key information loss during transitions and potentially redundant computation as well. Additionally, errors may accumulate within the system due to inconsistencies in optimization objectives among the modules, affecting the vehicle’s overall decision-making performance [Chen *et al.*, 2023a].

Rule-based decision systems, with their inherent limitations and scalability issues, are gradually giving way to data-driven methods. End-to-end autonomous driving solutions are increasingly becoming a consensus in the field [Wu *et al.*, 2022b; Chitta *et al.*, 2023; Chen and Krähenbühl, 2022; Jia *et al.*, 2023c; Jia *et al.*, 2023a; Hu *et al.*, 2023b]. By eliminating integration errors between multiple modules and reducing redundant computations, the end-to-end system enhances the expression of visual [Wu *et al.*, 2022a] and sensory information while ensuring greater efficiency. However, this approach also introduces the “black box” problem, meaning a lack of transparency in the decision-making process, complicating interpretation and validation.

Simultaneously, the explainability of autonomous driving has become an important research focus [Jin *et al.*, 2023a]. Although smaller language models (like early versions of BERT [Devlin *et al.*, 2018] and GPT [Brown *et al.*, 2020]) employed in massive data collection from driving scenarios help address this issue, they often lack sufficient generalization capabilities to perform optimally. Recently, large language models [OpenAI, 2023; Touvron *et al.*, 2023] have demonstrated remarkable abilities in understanding context, generating answers, and handling complex tasks. They are also now integrated with multimodal models [Brohan *et al.*, 2023a; Liu *et al.*, 2023a; Driess *et al.*, 2023; Xu *et al.*, 2023;

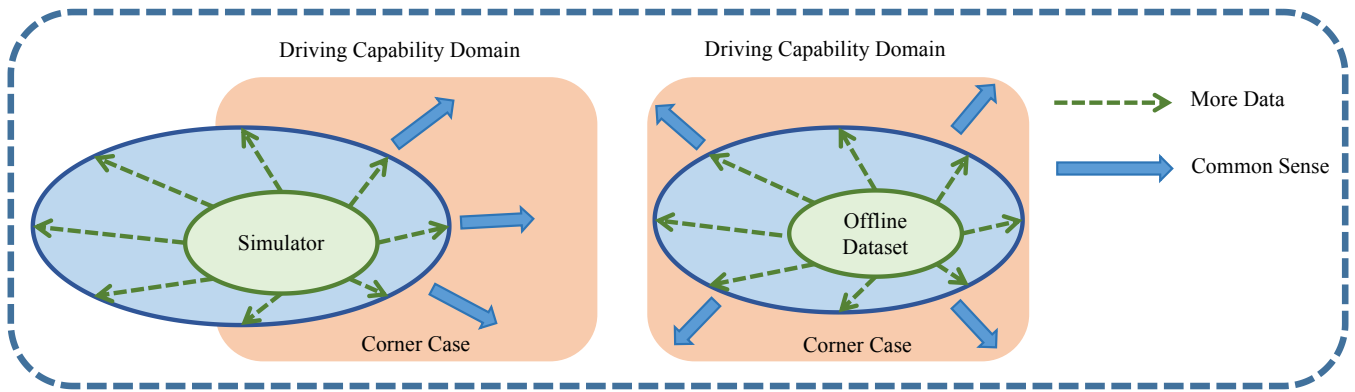


Figure 1: The limitation of current autonomous driving paradigm (green arrow) and where LLMs can potentially enhance autonomous driving ability (blue arrow).

Chen *et al.*, 2023b]. This integration achieves a unified feature space mapping for images, text, videos, point clouds, etc. Such consolidation significantly enhances the system’s generalization capabilities and equips them with the capacity to quickly adapt to new scenarios in a zero-shot or few-shot manner.

In this context, developing an interpretable and efficient end-to-end autonomous driving system has become a research hotspot [Chen *et al.*, 2023a]. Large language models, with their extensive knowledge base and exceptional generalization, could facilitate easier learning of complex driving behaviors. By leveraging the visual-language model (VLM)’s robust and comprehensive capabilities of open-world understanding and in-context learning [Bommasani *et al.*, 2021; Brohan *et al.*, 2023b; Liu *et al.*, 2023a; Driess *et al.*, 2023], it becomes possible to address the long-tail problem for perception networks, assist in decision-making, and provide intuitive explanations for these decisions.

This paper aims to provide a comprehensive overview of this rapidly emerging research field, analyze its basic principles, methods, and implementation processes, and introduce in detail regarding the application of LLMs for autonomous driving. Finally, we discuss related challenges and future research directions.

## 2 Motivation of LLM4AD

In today’s technological landscape, large language models such as GPT-4 and GPT-4V [OpenAI, 2023; Yang *et al.*, 2023b] are drawing attention with their superior contextual understanding and in-context learning capabilities. Their enriched common sense knowledge has facilitated significant advancements in many downstream tasks. We ask the question: *how do these large models assist in the domain of autonomous driving, especially in playing a critical role in the decision-making process?*

In Fig. 1, we give an intuitive demonstration of the limitation of current autonomous driving paradigm and where LLMs can potentially enhance autonomous driving ability. We summarize two primary aspects of driving skills. The orange circle represents the ideal level of driving competence, akin to that possessed by an experienced human driver.

There are two main methods to acquire such proficiency: one, through learning-based techniques within simulated environments; and two, by learning from offline data through similar methodologies. It’s important to note that due to discrepancies between simulations and the real world, these two domains are not fully the same, i.e. sim2real gap [Höfer *et al.*, 2021]. Concurrently, offline data serves as a subset of real-world data since it’s collected directly from actual surroundings. However, it is difficult to fully cover the distribution as well due to the notorious long-tailed nature [Jain *et al.*, 2021] of autonomous driving tasks.

The final goal of autonomous driving is to elevate driving abilities from a basic green stage to a more advanced blue level through extensive data collection and deep learning. However, the high cost associated with data gathering and annotation, along with the inherent differences between simulated and real-world environments, mean there’s still a gap before reaching the expert level of driving skills. In this scenario, if we can effectively utilize the innate common sense embedded within large language models, we might gradually narrow this gap. Intuitively, by adopting this approach, we could progressively enhance the capabilities of autonomous driving systems, bringing them closer to, or potentially reaching, the ideal expert level of driving proficiency. Through such technological integration and innovation, we anticipate significant improvements in the overall performance and safety of autonomous driving.

The application of large language models in the field of autonomous driving indeed covers a wide range of task types, combining depth and breadth with revolutionary potential. LLMs in autonomous driving pipelines is shown in the Fig. 2.

## 3 Application of LLM4AD

In the following sections, we divide existing works based on the perspective of applying LLMs: planning, perception, question answering, and generation. The corresponding taxonomy tree is shown in Fig. 3.

### 3.1 Planning

Large language models (LLMs) have achieved great success with their open-world cognitive and reasoning capa-

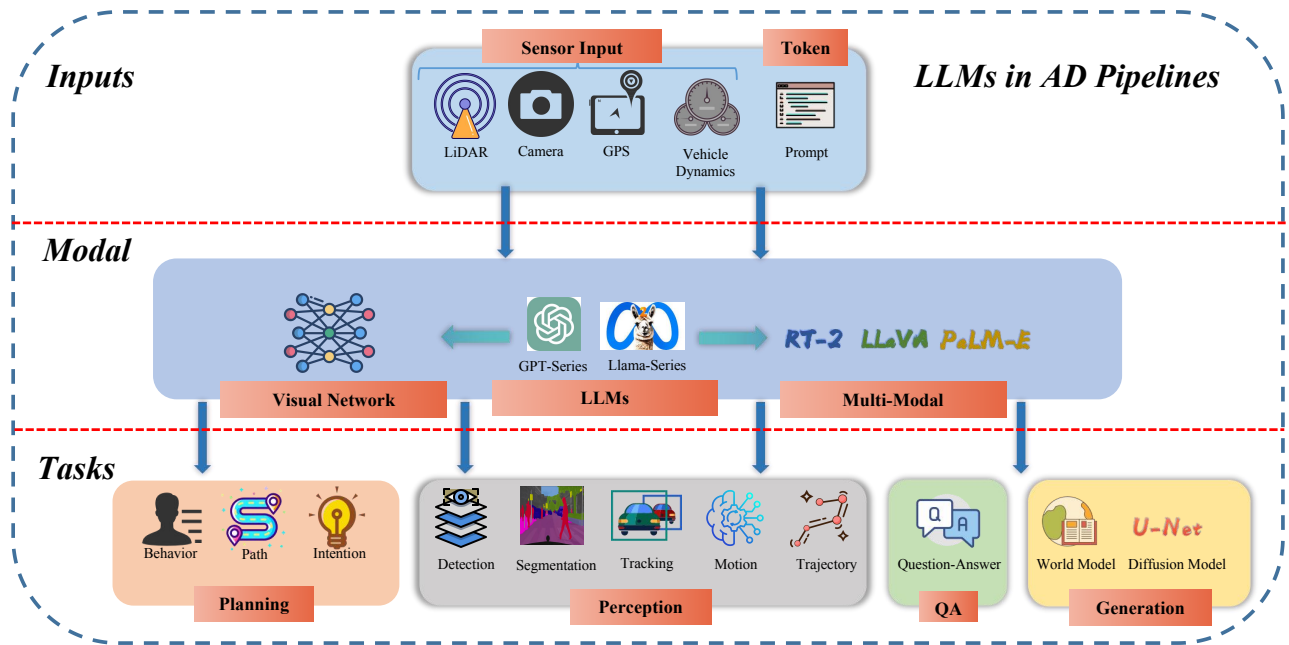


Figure 2: LLMs in Autonomous Driving Pipelines.

bilities [Radford *et al.*, 2018; Radford *et al.*, 2019; Brown *et al.*, 2020; Ouyang *et al.*, 2022; OpenAI, 2023]. These capabilities could provide a transparent explanation of the autonomous driving decision-making process, significantly enhancing system reliability and user trust in the technology [Deruyttere *et al.*, 2019; Kim *et al.*, 2019a; Atakishiyev *et al.*, 2023; Jin *et al.*, 2023a; Malla *et al.*, 2023]. Within this domain, based on whether tuning the LLM, related research can be categorized into two main types: fine-tuning pre-trained models and prompt engineering.

In the application of fine-tuning pre-trained models, MTD-GPT [Liu *et al.*, 2023b] translates multi-task decision-making problems into sequence modeling problems. Through training on a mixed multi-task dataset, it addresses various decision-making tasks at unsignaled intersections. Although this approach outperforms the performance of single-task decision-making RL models, the used scenes are limited to unsignaled intersections, which might be enough to demonstrate the complexity of the real world application. Driving with LLMs [Chen *et al.*, 2023b] designs an architecture that fuses vectorized inputs into LLMs with a two-stage pretraining and fine-tuning method. Due to the limitation of vectorized representations, their method are only tested in the simulation. DriveGPT4 [Xu *et al.*, 2023] presents a multimodal LLM based on Valley [Luo *et al.*, 2023a] and develops a visual instruction tuning dataset for interpretable autonomous driving. Beside predicting a vehicle’s basic control signals, it also responds in real-time, explaining why the action was taken. It outperforms baseline models in a variety of QA tasks while the experiments about planning is simple.

In the prompt engineering perspective, some methods tried to tap into the deep reasoning potential of the LLMs through clever prompt design. DiLu [Wen *et al.*, 2023] designs a

framework of LLMs as agents to solve closed-loop driving tasks. This method introduces a memory module to record experience, to leverage LLMs to facilitate reasoning and reflection processes. DiLu exhibits strong generalization capabilities compared with SOTA RL-based methods. However, the reasoning and reflection processes require multiple rounds of question-answering, and its inference time cannot be ignored. Similarly, Receive Reason and React [Cui *et al.*, 2023b] and Drive as You Speak [Cui *et al.*, 2023a] integrate the language and reasoning capabilities of LLMs into autonomous vehicles. In addition to memory and reflection processes, these methods introduce additional raw sensor information such as camera, GNSS, lidar, and radar. However, the inference speed is unsolved as well. Furthermore, SurréalDriver [Jin *et al.*, 2023b] divides the memory module into short-term memory, long-term guidelines, and safety criteria. Meanwhile, it interviews 24 drivers and uses their detailed descriptions of driving behaviors as chain-of-thought prompts to develop a ‘coach agent’ module. However, there is a lack of comparison with traditional algorithms to prove that large language models indeed bring performance improvements. LanguageMPC [Sha *et al.*, 2023] also designs a chain-of-thought framework for LLMs in driving scenarios and it integrates with low-level controllers by guided parameter matrix adaptation. Although its performance exceeds MPC and RL-based methods in the simplified simulator environments, it lacks validation in complex environments. TrafficGPT [Zhang *et al.*, 2023b] is a fusion of ChatGPT and traffic foundation models which can tackle complex traffic-related problems and provide insightful suggestions. It leverages multimodal data as a data source, offering comprehensive support for various traffic-related tasks. Talk2BEV [Dewangan *et al.*, 2023] introduces a large vision-language model (LVLM) interface

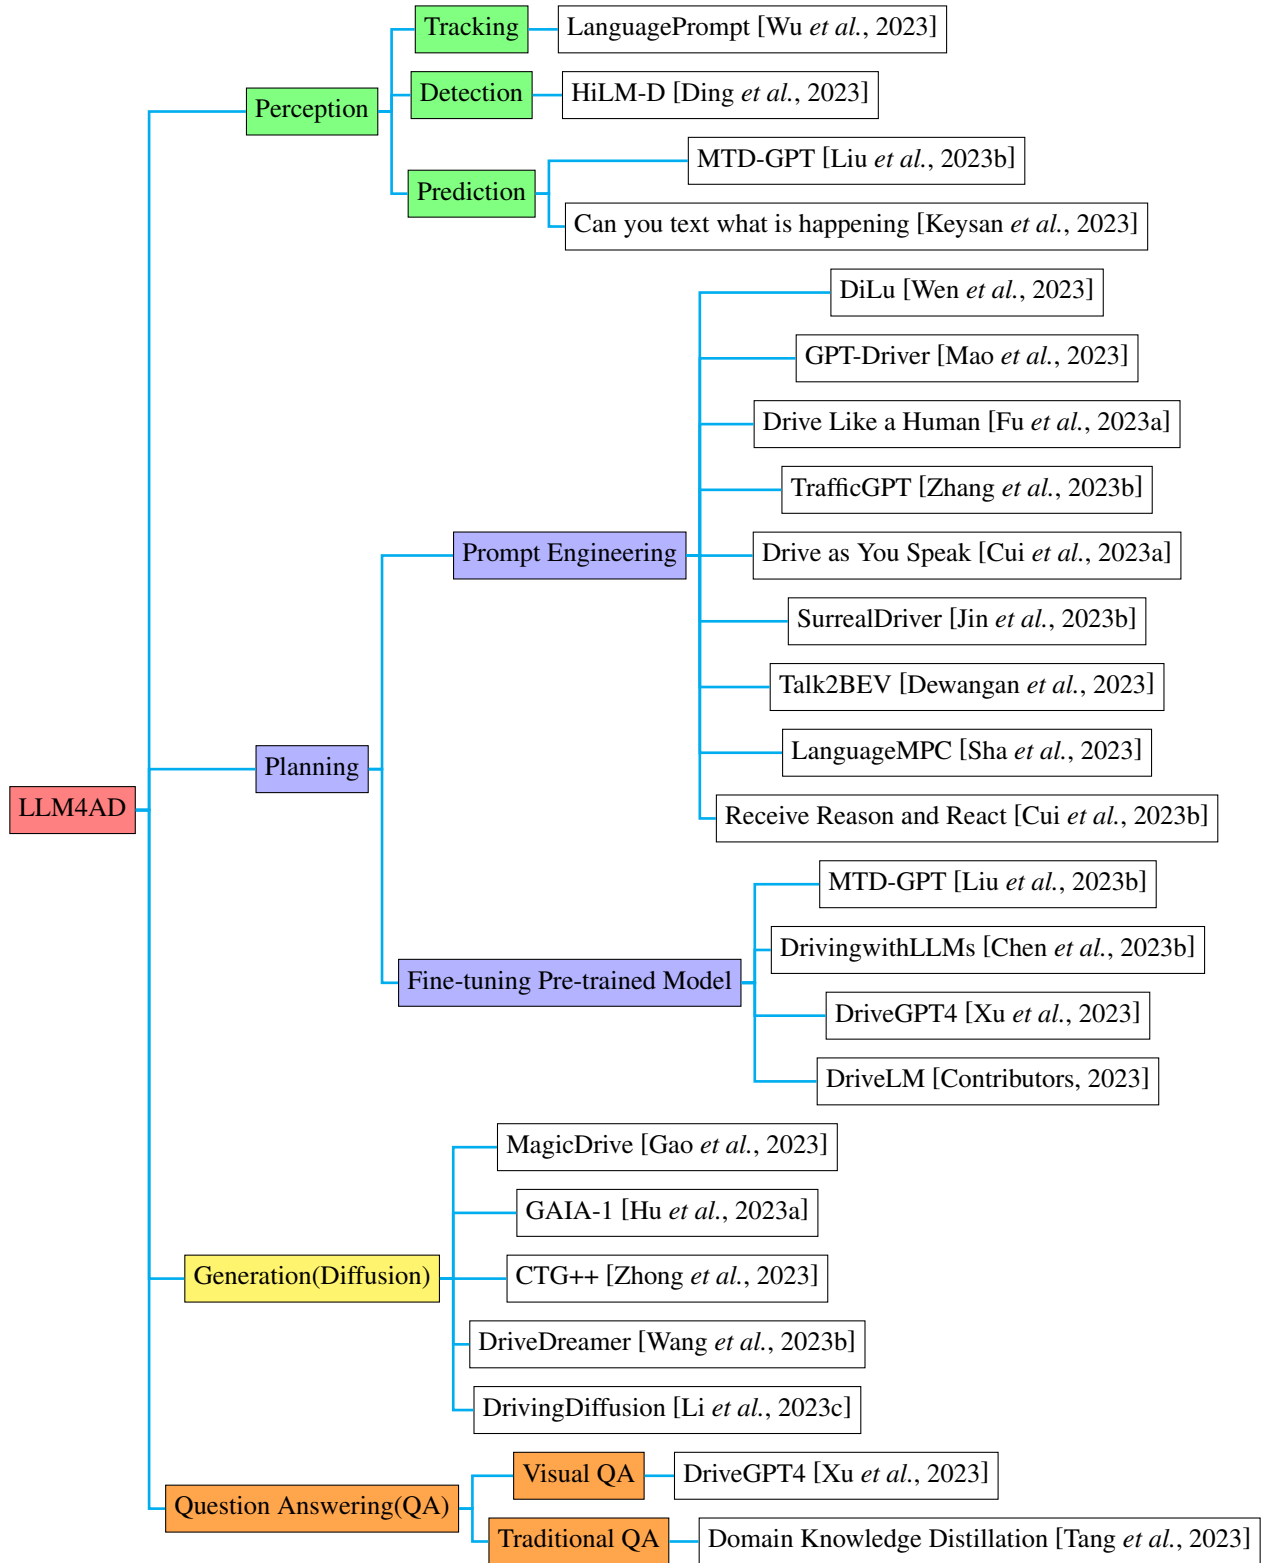


Figure 3: Large Language Models for Autonomous Driving Research Tree

for bird’s-eye view (BEV) maps in autonomous driving contexts. It does not require any training or fine-tuning, only relying on pre-trained image-language models. In addition, it presents a benchmark for evaluating subsequent work in LVLMs for AD applications. GPT-Driver [Mao *et al.*, 2023] transforms the motion planning task into a language modeling problem. It exceeds the UniAD [Hu *et al.*, 2023b] in the L2 metric. Nevertheless, since it uses past speed and acceleration information, there is concern about unfair comparison with UniAD. Additionally, L2 only reflects the fitting degree of the driving route and might not reflect the driving performance [Dauner *et al.*, 2023]. DriveLM [Contributors, 2023] uses a trajectory tokenizer to process ego-trajectory signals to texts, making them belong to the same domain space. Such a tokenizer can be applied to any general vision language models. Moreover, they utilize a graph-structure inference with multiple QA pairs in logical order, thus improving the final planning performance.

#### **Metric:**

MTD-GPT [Liu *et al.*, 2023b] uses single-subtask success rates as the metric in simulation and it exceeds RL expert. DriveGPT4 [Xu *et al.*, 2023] uses root mean squared error (RMSE) and threshold accuracies for evaluation. In vehicle action description, justification, and full sentences, it uses BLEU-4 [Papineni *et al.*, 2002], METEOR [Banerjee and Lavie, 2005], CIDER [Vedantam *et al.*, 2015] and chatgpt score [Fu *et al.*, 2023b]. Driving with LLMs [Chen *et al.*, 2023b] uses the Mean Absolute Error (MAE) for the predictions of the number of cars and pedestrians, normalized acceleration, and brake pressure. Additionally, it measures the accuracy of traffic light detection as well as the mean absolute distance error in meters for traffic light distance prediction. Besides perception-related metrics, it also uses GPT-3.5 to grade their model’s answers which is a recently emerging technique - grading natural language responses [Fu *et al.*, 2023b; Wang *et al.*, 2023a; Liu *et al.*, 2023c]. DiLu [Wen *et al.*, 2023] uses Success Steps in simulation as a metric to evaluate generalization and transformation abilities. SurrealDriver [Jin *et al.*, 2023b] evaluates agents based on two main dimensions: safe driving ability and humanness. Safe driving capabilities are assessed through collision rates, while human likeness is assessed through user experiments with 24 adult participants (age  $29.3 \pm 4.9$  years, male = 17 years) who are legal to drive. LanguageMPC [Sha *et al.*, 2023] customizes some metrics: failure/collision cases, the efficiency of traffic flow, time cost by ego vehicle, and the safety of the ego vehicle’s driving behavior. Similarly, Talk2BEV [Dewangan *et al.*, 2023] measures their methods from the perspective of spatial reasoning, instance attribute, instance counting, and visual reasoning. GPT-Driver [Mao *et al.*, 2023] and DriveLM [Contributors, 2023] contains two metrics: L2 error (in meters) and collision rate (in percentage). The average L2 error is calculated by measuring the distance of each waypoint in the planned trajectory and the offline recorded human driver trajectory. It reflects the fitting of the planned trajectory to the human driving trajectory. The collision rate is calculated by placing an ego vehicle box at each waypoint of the planned trajectory and then checking for collisions with the ground truth bounding boxes of other objects. It reflects

the safety of the planned trajectory.

At present, LLM4AD regarding the planning task lacks a unified metric and cannot uniformly evaluate the pros and cons between each method and traditional counterparts.

### **3.2 Perception**

Large language models have demonstrated their unique value and strong capabilities in “perception” tasks [Radford *et al.*, 2021; Li *et al.*, 2022b; Li *et al.*, 2023a; Li *et al.*, 2023b; Li *et al.*, 2022a]. Especially in environments where data is relatively scarce, these models can rely on their few-shot learning characteristics to achieve fast and accurate learning and reasoning [P *et al.*, 2023; Lin *et al.*, 2023]. This learning ability is of significance in the perception stage of the autonomous driving system, which greatly improves the system’s adaptability and generalization capabilities in changing and complex driving environments. PromptTrack [Wu *et al.*, 2023] fuses cross-modal features in a prompt reasoning branch to predict 3D objects. It uses language prompts as semantic cues and combines LLMs with 3D detection tasks and tracking tasks. Although it achieves better performance compared to other methods, the advantages of LLMs do not directly affect the tracking task. Rather, the tracking task serves as a query to assist LLMs in performing 3D detection tasks. HiLM-D [Ding *et al.*, 2023] incorporates high-resolution information into multimodal large language models for the Risk Object Localization and Intention and Suggestion Prediction (ROLISP) task. It combines LLMs with 2D detection tasks and obtains better performance in detection tasks and QA tasks compared to other multi-modal large models such as Video-LLaMa [Zhang *et al.*, 2023a], ePALM [Shukor *et al.*, 2023]. It is worth noting to point out one potential limitation of the dataset: each video contains only one risk object, which might not capture the complexity of real-world scenarios. [Keysan *et al.*, 2023] integrates pre-trained language models as text-based input encoders for the autonomous driving trajectory prediction task. Joint encoders (image and text) over both modalities perform better than using a single encoder in isolation. While the joint model significantly improves the baseline, its performance has not reached the state-of-the-art level yet [Deo *et al.*, 2021; Gilles *et al.*, 2021].

#### **Metric:**

PromptTrack [Wu *et al.*, 2023] uses the Average Multiple Object Tracking Precision (AMOTA) metric [Bernardin and Stiefelhagen, 2008], the Average MultiObject Tracking Precision (AMOTP) [Bashar *et al.*, 2022] and Identity Switches (IDS) [Huang *et al.*, 2023] metrics. HiLM-D [Ding *et al.*, 2023] uses the BLEU-4 [Papineni *et al.*, 2002], METEOR [Banerjee and Lavie, 2005], CIDER [Vedantam *et al.*, 2015] and SPICE [Anderson *et al.*, 2016], IoU [Rezatofighi *et al.*, 2019] as metrics to compare with the state-of-the-art. [Keysan *et al.*, 2023] uses the standard evaluation metrics that are provided in the nuScenes-devkit [Caesar *et al.*, 2019; Fong *et al.*, 2021]: minimum Average Displacement Error (minADEk), Final Displacement Error (minFDEk), and the miss rate over 2 meters.

### 3.3 Question Answering

Question-Answering is an important task that has a wide range of applications in intelligent transportation, assisted driving, and autonomous vehicles [Xu *et al.*, 2021a; Xu *et al.*, 2021b]. It mainly reflects through different question and answer paradigms, including traditional QA mechanism [Tang *et al.*, 2023] and more detailed visual QA methods [Xu *et al.*, 2023]. [Tang *et al.*, 2023] constructs the domain knowledge ontology by “chatting” with ChatGPT. It develops a web-based assistant to enable manual supervision and early intervention at runtime and it guarantees the quality of fully automated distillation results. This question-and-answer system enhances the interactivity of the vehicle, transforms the traditional one-way human-machine interface into an interactive communication experience, and might be able to cultivate the user’s sense of participation and control. These sophisticated models [Tang *et al.*, 2023; Xu *et al.*, 2023], equipped with the ability to parse, understand, and generate human-like responses, are pivotal in real-time information processing and provision. They design comprehensive questions related to the scene, including but not limited to vehicle states, navigation assistance, and understanding of traffic situations.

#### Metric:

In terms of QA tasks, NLP’s metric is often used. In DriveFPT4 [Xu *et al.*, 2023], it uses BLEU-4 [Papineni *et al.*, 2002], METEOR [Banerjee and Lavie, 2005], CIDER [Vedantam *et al.*, 2015] and chatgpt score [Fu *et al.*, 2023b].

### 3.4 Generation

In the realm of “generation” task, large language models leverage their advanced knowledge-base and generative capabilities to create realistic driving videos or intricate driving scenarios under specific environmental factors [Khachatryan *et al.*, 2023; Luo *et al.*, 2023b]. This approach offers revolutionary solutions to the challenges of data collection and labeling for autonomous driving, also constructing a safe and easily controllable setting for testing and validating the decision boundaries of autonomous driving systems. Moreover, by simulating a variety of driving situations and emergency conditions, the generated content becomes a crucial resource for refining and enriching the emergency response strategies of autonomous driving systems.

The common generative models include the Variational Auto-Encoder(VAE) [Kingma and Welling, 2022], Generative Adversarial Network(GAN) [Goodfellow *et al.*, 2014], Normalizing Flow(Flow) [Rezende and Mohamed, 2016], and Denoising Diffusion Probabilistic Model(Diffusion) [Ho *et al.*, 2020]. With diffusion models have recently achieved great success in text-to-image [Ronneberger *et al.*, 2015; Rombach *et al.*, 2021; Ramesh *et al.*, 2022], some research has begun to study using diffusion models to generate autonomous driving images or videos. DriveDreamer [Wang *et al.*, 2023b] is a world model derived from real-world driving scenarios. It uses text, initial image, HDmap, and 3Dbox as input, then generates high-quality driving videos and reasonable driving policies. Similarly, Driving Diffusion [Li *et al.*, 2023c] adopts a 3D layout as a control signal to generate realistic multi-view videos. GAIA-1 [Hu *et al.*, 2023a] leverages video, text, and action inputs to generate traffic scenarios, environmental elements, and potential risks. In these methods, text encoder both adopt CLIP [Radford *et al.*, 2021] which has a better alignment between image and text. In addition to generating autonomous driving videos, traffic scenes can also be generated. CTG++ [Zhong *et al.*, 2023] is a scene-level diffusion model that can generate realistic and controllable traffic. It leverages LLMs for translating a user query into a differentiable loss function and use a diffusion model to transform the loss function into realistic, query compliant trajectories. MagicDrive [Gao *et al.*, 2023] generates highly realistic images, exploiting geometric information from 3D annotations by independently encoding road maps, object boxes, and camera parameters for precise, geometry-guided synthesis. This approach effectively solves the challenge of multi-camera view consistency. Although it achieves better performance in terms of generation fidelity compared to BEVGen [Swerdlow *et al.*, 2023] and BEVControl [Yang *et al.*, 2023a], it also faces huge challenges in some complex scenes, such as night views and unseen weather conditions.

These methods explore the customized authentic generations of autonomous driving data. Although these diffusion-based models achieved good results on video and image-generated metrics, it is still unclear whether they could really be used in closed-loop to really boost the performance of the autonomous driving system.

#### Metric:

DriveDreamer [Wang *et al.*, 2023b] and DrivingDiffusion [Li *et al.*, 2023c] use the frame-wise Frechet Inception Distance (FID) [Parmar *et al.*, 2022] to evaluate the quality of generated images and the Frechet Video Distance (FVD) [Unterthiner *et al.*, 2019] for video quality evaluation. DrivingDiffusion also uses average intersection crossing (mIoU) [Rezatofighi *et al.*, 2019] scores for drivable areas and NDS [Yin *et al.*, 2021] for all the object classes by comparing the predicted layout with the ground-truth BEV layout. CTG++ [Zhong *et al.*, 2023] following [Xu *et al.*, 2022; Zhong *et al.*, 2022], uses the failure rate, Wasserstein distance between normalized histograms of driving profiles, realism deviation (real), and scene-level realism metric (rel real) as metrics. MagicDrive [Gao *et al.*, 2023] utilizes segmentation metrics such as Road mIoU and Vehicle mIoU [Taran *et al.*, 2018], as well as 3D object detection metrics like mAP [Henderson and Ferrari, 2017] and NDS [Yin *et al.*, 2021].

#### Metric:

### 4 Datasets in LLM4AD

Traditional datasets such as nuScenes dataset [Caesar *et al.*, 2019; Fong *et al.*, 2021] lack action description, detailed caption, and question-answering pairs which are used to interact with LLMs. The BDD-x [Kim *et al.*, 2018], Rank2Tell [Sachdeva *et al.*, 2023], DriveLM [Contributors, 2023], DRAMA [Malla *et al.*, 2023], NuPrompt [Wu *et al.*, 2023] and NuScenes-QA [Qian *et al.*, 2023] datasets represent key developments in LLM4AD research, each bringing unique contributions to understanding agent behaviors and urban traffic dynamics through extensive, diverse, and

Dataset	Task	Size	Description
BDD-X [Kim <i>et al.</i> , 2018]	Planning VQA	77 hours, 6970 videos, 8.4M frames, 26228 captions	Ego-vehicle actions description and explanation.
HAD [Kim <i>et al.</i> , 2019b]	Planning Perception	30 hours, 5744 videos 22366 captions	Joint action description for goal-oriented advice and attention description for stimulus-driven advice.
Talk2Car [Deruyttere <i>et al.</i> , 2019]	Planning Perception	15 hours, 850 videos of 20s each 30k frames, 11959 captions	Object referral dataset that contains commands written in natural language for self-driving cars.
DriveLM [Contributors, 2023]	Perception Prediction Planning	15 hours, 30k frames, 360k QA pairs	P3 with reasoning logic; Connect the QA pairs in a graph-style structure; Use “What if”-style questions.
DRAMA [Malla <i>et al.</i> , 2023]	VQA	91 hours, 17785 videos, 77639 question, 102830 answering, 17066 captions	Joint risk localization with visual reasoning of driving risks in a free-form language description.
Rank2Tell [Sachdeva <i>et al.</i> , 2023]	Perception VQA	several hours, 118 videos of 20s each	Joint important object identification, important object localization ranking, and reasoning.
NuPrompt [Wu <i>et al.</i> , 2023]	Perception	15 hours, 35367 prompts for 3D objects	Object-centric language prompt set for perception tasks.
NuScenes-QA [Qian <i>et al.</i> , 2023]	VQA	15 hours, Train(24149 scenes, 459941 QA pairs), Test(6019 scenes, 83337 QA pairs)	Leverage 3D annotations(object category, position, orientation, relationships information) and designed question templates to construct QA pairs.

Table 1: Description of different datasets regarding LLM4AD.

situation-rich annotations. We give a summary of each dataset in Table 1. We give detailed descriptions below.

**BDD-X Dataset** [Kim *et al.*, 2018]: With over 77 hours of diverse driving conditions captured in 6,970 videos, this dataset is a collection of real-world driving behaviors, each annotated with descriptions and explanations. It includes 26K activities across 8.4M frames and thus provides a resource for understanding and predicting driver behaviors across different conditions.

**Honda Research Institute-Advice Dataset (HAD)** [Kim *et al.*, 2019b]: HAD offers 30 hours of driving video data paired with natural language advice and videos integrated with can-bus signal data. The advice includes Goal-oriented advice(top-down signal) which is designed to guide the vehicle in a navigation task and Stimulus-driven advice(bottom-up signal) which highlights specific visual cues that the user expects the vehicle controller to actively focus on.

**Talk2Car** [Deruyttere *et al.*, 2019]: The Talk2Car dataset contains 11959 commands for the 850 videos of the nuScenes [Caesar *et al.*, 2019; Fong *et al.*, 2021] training set as 3D bounding box annotations. Of the commands, 55.94% were from videos recorded in Boston, while 44.06% were from Singapore. On average, each command contains 11.01 words, which includes 2.32 nouns, 2.29 verbs, and 0.62 adjectives. Typically, there are about 14.07 commands in every video. It is a object referral dataset that contains commands written in natural language for self-driving cars.

**DriveLM Dataset** [Contributors, 2023]: This dataset integrates human-like reasoning into autonomous driving systems, enhancing Perception, Prediction, and Planning (P3). It employs a “Graph-of-Thought” structure, encouraging a futuristic approach through “What if” scenarios, thereby promoting advanced, logic-based reasoning and decision-

making mechanisms in driving systems.

**DRAMA Dataset** [Malla *et al.*, 2023]: Collected from Tokyo’s streets, it includes 17,785 scenario clips captured using the video camera, each clipped to 2 seconds in duration. It contains different annotations: Video-level Q/A, Object-level Q/A, Risk object bounding box, Free-form caption, separate labels for ego-car intention, scene classifier, and suggestions to the driver.

**Rank2Tell Dataset** [Sachdeva *et al.*, 2023]: It is captured from a moving vehicle on highly interactive traffic scenes in the San Francisco Bay Area. It includes 116 clips ( 20s each) of 10FPS captured using an instrumented vehicle equipped with three Point Grey Grasshopper video cameras with a resolution of 1920 × 1200 pixels, a Velodyne HDL-64E S2 LiDAR sensor, and high precision GPS. The dataset includes Video-level Q/A, Object-level Q/A, LiDAR and 3D bounding boxes (with tracking), Field of view from 3 cameras (stitched), important object bounding boxes (multiple important objects per frame with multiple levels of importance-High, Medium, Low), free-form captions (multiple captions per object for multiple objects), ego-car intention.

**NuPrompt Dataset** [Wu *et al.*, 2023]: It represents an expansion of the nuScenes dataset, enriched with annotated language prompts specifically designed for driving scenes. This dataset includes 35,367 language prompts for 3D objects, averaging 5.3 instances per object. This annotation enhances the dataset’s practicality in autonomous driving testing and training, particularly in complex scenarios requiring linguistic processing and comprehension.

**NuScenes-QA dataset** [Qian *et al.*, 2023]: It is a dataset in autonomous driving, containing 459,941 question-answer pairs from 34,149 distinct visual scenes. They are partitioned into 376,604 questions from 28,130 scenes for training, and



83,337 from 6,019 scenes for testing. NuScenes-QA showcases a wide array of question lengths, reflecting different complexity levels, making it challenging for AI models. Beyond sheer numbers, the dataset ensures a balanced range of question types and categories, from identifying objects to assessing their behavior, such as whether they are moving or parked. This design inhibits the model's tendency to be biased or rely on linguistic shortcuts.

## 5 Conclusion

In this paper, we have provided a comprehensive survey on LLM4AD. We classify and introduce different applications employing LLMs for autonomous driving and summarize the representative approaches in each category. At the same time, we summarize the latest datasets related to LLM4AD. We will continue to monitor developments in the field and highlight future research directions.

## Ethical Statement

When applying LLMs to the field of autonomous driving, we must deeply consider their potential ethical implications. First, the illusion of the model may cause the vehicle to misunderstand the external environment or traffic conditions, thus causing safety hazards. Second, model discrimination and bias may lead to vehicles making unfair or biased decisions in different environments or when facing different groups. Additionally, false information and errors in reasoning can cause a vehicle to adopt inappropriate or dangerous driving behaviors. Inductive advice may leave the vehicle vulnerable to external interference or malicious behavior. Finally, privacy leakage is also a serious issue, as vehicles may inadvertently reveal sensitive information about the user or the surrounding environment. To sum up, we strongly recommend that before deploying a large language model to an autonomous driving system, an in-depth and detailed ethical review should be conducted to ensure that its decision-making logic is not only technically accurate but also ethically appropriate. At the same time, we call for following the principles of transparency, responsibility, and fairness to ensure the ethics and safety of technology applications. We call on the entire community to work together to ensure reliable and responsible deployment of autonomous driving technology based on large language models.

## Acknowledgments

This work was partly supported by NSFC (92370201, 62222607, 61972250) and Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102).

## References

[Anderson *et al.*, 2016] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation, 2016.

[Atakishiyev *et al.*, 2023] Shahin Atakishiyev, Mohammad Saleh, Hengshuai Yao, and Randy Goebel. Explainable

artificial intelligence for autonomous driving: A comprehensive overview and field guide for future research directions, 2023.

- [Banerjee and Lavie, 2005] Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In Jade Goldstein, Alon Lavie, Chin-Yew Lin, and Clare Voss, editors, *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics.
- [Bashar *et al.*, 2022] Mk Bashar, Samia Islam, Kashifa Kawaakib Hussain, Md. Bakhtiar Hasan, A. B. M. Ashikur Rahman, and Md. Hasanul Kabir. Multiple object tracking in recent times: A literature review, 2022.
- [Bernardin and Stiefelhagen, 2008] Keni Bernardin and Rainer Stiefelhagen. Evaluating multiple object tracking performance: The clear mot metrics. *EURASIP J. Image Video Process.*, 2008, 2008.
- [Bommasani *et al.*, 2021] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [Brohan *et al.*, 2023a] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- [Brohan *et al.*, 2023b] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-Wei Edward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspiar Singh, Anikait Singh, Radu Soricut, Huong Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-2: Vision-language-action models transfer web knowledge to robotic control, 2023.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.



- [Caesar *et al.*, 2019] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019.
- [Chen and Krähenbühl, 2022] Dian Chen and Philipp Krähenbühl. Learning from all vehicles. In *CVPR*, 2022.
- [Chen *et al.*, 2023a] 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.
- [Chen *et al.*, 2023b] Long Chen, Oleg Sinavski, Jan Hünermann, Alice Karnsund, Andrew James Willmott, Danny Birch, Daniel Maund, and Jamie Shotton. Driving with llms: Fusing object-level vector modality for explainable autonomous driving, 2023.
- [Cheng *et al.*, 2022] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. 2022.
- [Chitta *et al.*, 2023] Kashyap Chitta, Aditya Prakash, Bernhard Jaeger, Zehao Yu, Katrin Renz, and Andreas Geiger. Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE Pattern Analysis and Machine Intelligence (PAMI)*, 2023.
- [Contributors, 2023] DriveLM Contributors. Drivelm: Drive on language. <https://github.com/OpenDriveLab/DriveLM>, 2023.
- [Cui *et al.*, 2023a] Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, and Ziran Wang. Drive as you speak: Enabling human-like interaction with large language models in autonomous vehicles. *arXiv preprint arXiv:2309.10228*, 2023.
- [Cui *et al.*, 2023b] Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, and Ziran Wang. Receive, reason, and react: Drive as you say with large language models in autonomous vehicles. *arXiv preprint arXiv:2310.08034*, 2023.
- [Dauner *et al.*, 2023] Daniel Dauner, Marcel Hallgarten, Andreas Geiger, and Kashyap Chitta. Parting with misconceptions about learning-based vehicle motion planning. In *CoRL*, 2023.
- [Deo *et al.*, 2021] Nachiket Deo, Eric M. Wolff, and Oscar Beijbom. Multimodal trajectory prediction conditioned on lane-graph traversals, 2021.
- [Deruyttere *et al.*, 2019] Thierry Deruyttere, Simon Vandenhende, Dusan Grujicic, Luc Van Gool, and Marie-Francine Moens. Talk2car: Taking control of your self-driving car. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, 2019.
- [Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [Dewangan *et al.*, 2023] Vikrant Dewangan, Tushar Choudhary, Shivam Chandhok, Shubham Priyadarshan, Anushka Jain, Arun K. Singh, Siddharth Srivastava, Krishna Murthy Jatavallabhula, and K. Madhava Krishna. Talk2bev: Language-enhanced bird’s-eye view maps for autonomous driving, 2023.
- [Ding *et al.*, 2023] Xinpeng Ding, Jianhua Han, Hang Xu, Wei Zhang, and Xiaomeng Li. Hilm-d: Towards high-resolution understanding in multimodal large language models for autonomous driving, 2023.
- [Driess *et al.*, 2023] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. Palm-e: An embodied multimodal language model, 2023.
- [Fong *et al.*, 2021] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *arXiv preprint arXiv:2109.03805*, 2021.
- [Fu *et al.*, 2023a] Daocheng Fu, Xin Li, Licheng Wen, Min Dou, Pinlong Cai, Botian Shi, and Yu Qiao. Drive like a human: Rethinking autonomous driving with large language models, 2023.
- [Fu *et al.*, 2023b] Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. Gptscore: Evaluate as you desire, 2023.
- [Gao *et al.*, 2023] Ruiyuan Gao, Kai Chen, Enze Xie, Lanqing Hong, Zhenguo Li, Dit-Yan Yeung, and Qiang Xu. Magicdrive: Street view generation with diverse 3d geometry control, 2023.
- [Gilles *et al.*, 2021] Thomas Gilles, Stefano Sabatini, Dzmityr Tsishkou, Bogdan Stanculescu, and Fabien Moutarde. Gohome: Graph-oriented heatmap output for future motion estimation, 2021.
- [Goodfellow *et al.*, 2014] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [Henderson and Ferrari, 2017] Paul Henderson and Vittorio Ferrari. End-to-end training of object class detectors for mean average precision, 2017.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020.
- [Höfer *et al.*, 2021] Sebastian Höfer, Kostas Bekris, Ankur Handa, Juan Camilo Gamboa, Melissa Mozifian, Florian Golemo, Chris Atkeson, Dieter Fox, Ken Goldberg, John Leonard, et al. Sim2real in robotics and automation: Applications and challenges. *IEEE transactions on automation science and engineering*, 18(2):398–400, 2021.

- [Hu *et al.*, 2023a] Anthony Hu, Lloyd Russell, Hudson Yeo, Zak Murez, George Fedoseev, Alex Kendall, Jamie Shotton, and Gianluca Corrado. Gaia-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023.
- [Hu *et al.*, 2023b] 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.
- [Huang *et al.*, 2023] Junchao Huang, Xiaoqi He, and Sheng Zhao. The detection and rectification for identity-switch based on unfalsified control, 2023.
- [Jain *et al.*, 2021] Ashesh Jain, Luca Del Pero, Hugo Grimmett, and Peter Ondruska. Autonomy 2.0: Why is self-driving always 5 years away? *arXiv preprint arXiv:2107.08142*, 2021.
- [Jia *et al.*, 2021] Xiaosong Jia, Liting Sun, Masayoshi Tomizuka, and Wei Zhan. Ide-net: Interactive driving event and pattern extraction from human data. *IEEE Robotics and Automation Letters*, 6(2):3065–3072, 2021.
- [Jia *et al.*, 2022a] Xiaosong Jia, Li Chen, Penghao Wu, Jia Zeng, Junchi Yan, Hongyang Li, and Yu Qiao. Towards capturing the temporal dynamics for trajectory prediction: a coarse-to-fine approach. In *CoRL*, 2022.
- [Jia *et al.*, 2022b] Xiaosong Jia, Liting Sun, Hang Zhao, Masayoshi Tomizuka, and Wei Zhan. Multi-agent trajectory prediction by combining egocentric and allocentric views. In *Conference on Robot Learning*, pages 1434–1443. PMLR, 2022.
- [Jia *et al.*, 2023a] 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, 2023.
- [Jia *et al.*, 2023b] Xiaosong Jia, Penghao Wu, Li Chen, Yu Liu, Hongyang Li, and Junchi Yan. Hdgt: Heterogeneous driving graph transformer for multi-agent trajectory prediction via scene encoding. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2023.
- [Jia *et al.*, 2023c] 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, 2023.
- [Jin *et al.*, 2023a] Bu Jin, Xinyu Liu, Yupeng Zheng, Pengfei Li, Hao Zhao, Tong Zhang, Yuhang Zheng, Guyue Zhou, and Jingjing Liu. Adapt: Action-aware driving caption transformer, 2023.
- [Jin *et al.*, 2023b] Ye Jin, Xiaoxi Shen, Huiling Peng, Xiaohan Liu, Jingli Qin, Jiayang Li, Jintao Xie, Peizhong Gao, Guyue Zhou, and Jiangtao Gong. Surrealdriver: Designing generative driver agent simulation framework in urban contexts based on large language model, 2023.
- [Keysan *et al.*, 2023] Ali Keysan, Andreas Look, Eitan Kosman, Gonca Gürsun, Jörg Wagner, Yu Yao, and Barbara Rakitsch. Can you text what is happening? integrating pre-trained language encoders into trajectory prediction models for autonomous driving, 2023.
- [Khachatryan *et al.*, 2023] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators, 2023.
- [Kim *et al.*, 2018] Jinkyu Kim, Anna Rohrbach, Trevor Darrell, John Canny, and Zeynep Akata. Textual explanations for self-driving vehicles. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [Kim *et al.*, 2019a] Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, and John Canny. Grounding human-to-vehicle advice for self-driving vehicles, 2019.
- [Kim *et al.*, 2019b] Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, and John Canny. Grounding human-to-vehicle advice for self-driving vehicles. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [Kingma and Welling, 2022] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022.
- [Li *et al.*, 2022a] Hongyang Li, Chonghao Sima, Jifeng Dai, Wenhai Wang, Lewei Lu, Huijie Wang, Jia Zeng, Zhiqi Li, Jiazhi Yang, Hanming Deng, Hao Tian, Enze Xie, Jiangwei Xie, Li Chen, Tianyu Li, Yang Li, Yulu Gao, Xiaosong Jia, Si Liu, Jianping Shi, Dahua Lin, and Yu Qiao. Delving into the devils of bird’s-eye-view perception: A review, evaluation and recipe. *arXiv preprint arXiv:2209.05324*, 2022.
- [Li *et al.*, 2022b] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation, 2022.
- [Li *et al.*, 2022c] 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 *ECCV*, pages 1–18. Springer, 2022.
- [Li *et al.*, 2023a] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, 2023.
- [Li *et al.*, 2023b] Tianyu Li, Li Chen, Huijie Wang, Yang Li, Jiazhi Yang, Xiangwei Geng, Shengyin Jiang, Yuting Wang, Hang Xu, Chunjing Xu, Junchi Yan, Ping Luo, and Hongyang Li. Graph-based topology reasoning for driving scenes. *arXiv preprint arXiv:2304.05277*, 2023.
- [Li *et al.*, 2023c] Xiaofan Li, Yifu Zhang, and Xiaoqing Ye. Drivingdiffusion: Layout-guided multi-view driving scene video generation with latent diffusion model. *arXiv preprint arXiv:2310.07771*, 2023.
- [Liang *et al.*, 2020] Ming Liang, Bin Yang, Wenyuan Zeng, Yun Chen, Rui Hu, Sergio Casas, and Raquel Urtasun.

- Pnpnet: End-to-end perception and prediction with tracking in the loop. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11553–11562, 2020.
- [Lin *et al.*, 2023] Zhiqiu Lin, Samuel Yu, Zhiyi Kuang, Deepak Pathak, and Deva Ramanan. Multimodality helps unimodality: Cross-modal few-shot learning with multi-modal models, 2023.
- [Liu *et al.*, 2023a] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023.
- [Liu *et al.*, 2023b] Jiaqi Liu, Peng Hang, Xiao qi, Jianqiang Wang, and Jian Sun. Mtd-gpt: A multi-task decision-making gpt model for autonomous driving at unsignalized intersections, 2023.
- [Liu *et al.*, 2023c] Yang Liu, Dan Iter, Yichong Xu, Shuo-hang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: Nlg evaluation using gpt-4 with better human alignment, 2023.
- [Liu *et al.*, 2023d] Zhijian Liu, Haotian Tang, Alexander Amini, Xingyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *ICRA*, 2023.
- [Luo *et al.*, 2018] Wenjie Luo, Bin Yang, and Raquel 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 conference on Computer Vision and Pattern Recognition*, pages 3569–3577, 2018.
- [Luo *et al.*, 2023a] Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Da Li, Pengcheng Lu, Tao Wang, Linmei Hu, Minghui Qiu, and Zhongyu Wei. Valley: Video assistant with large language model enhanced ability, 2023.
- [Luo *et al.*, 2023b] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation, 2023.
- [Malla *et al.*, 2023] Srikanth Malla, Chiho Choi, Isht Dwivedi, Joon Hee Choi, and Jiachen Li. Drama: Joint risk localization and captioning in driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1043–1052, 2023.
- [Mao *et al.*, 2023] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023.
- [OpenAI, 2023] OpenAI. Gpt-4 technical report, 2023.
- [Ouyang *et al.*, 2022] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.
- [P *et al.*, 2023] Jishnu Jaykumar P, Kamalesh Palanisamy, Yu-Wei Chao, Xinya Du, and Yu Xiang. Proto-clip: Vision-language prototypical network for few-shot learning, 2023.
- [Papineni *et al.*, 2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics, 2002.
- [Parmar *et al.*, 2022] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation, 2022.
- [Qian *et al.*, 2023] Tianwen Qian, Jingjing Chen, Linhai Zhuo, Yang Jiao, and Yu-Gang Jiang. Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. *arXiv preprint arXiv:2305.14836*, 2023.
- [Radford *et al.*, 2018] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [Radford *et al.*, 2019] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.
- [Ramesh *et al.*, 2022] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022.
- [Rezatofighi *et al.*, 2019] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression, 2019.
- [Rezende and Mohamed, 2016] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows, 2016.
- [Rombach *et al.*, 2021] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- [Ronneberger *et al.*, 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.
- [Sachdeva *et al.*, 2023] Enna Sachdeva, Nakul Agarwal, Suhas Chundi, Sean Roelofs, Jiachen Li, Behzad Dariush, Chiho Choi, and Mykel Kochenderfer. Rank2tell: A multimodal driving dataset for joint importance ranking and reasoning. *arXiv preprint arXiv:2309.06597*, 2023.
- [Sadat *et al.*, 2020] Abbas Sadat, Sergio Casas, Mengye Ren, Xinyu Wu, Pranaab Dhawan, and Raquel Urtasun. Perceive, predict, and plan: Safe motion planning

- through interpretable semantic representations. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16*, pages 414–430. Springer, 2020.
- [Sha *et al.*, 2023] Hao Sha, Yao Mu, Yuxuan Jiang, Li Chen, Chenfeng Xu, Ping Luo, Shengbo Eben Li, Masayoshi Tomizuka, Wei Zhan, and Mingyu Ding. Language-mpc: Large language models as decision makers for autonomous driving. *arXiv preprint arXiv:2310.03026*, 2023.
- [Shi *et al.*, 2022] 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.
- [Shukor *et al.*, 2023] Mustafa Shukor, Corentin Dancette, and Matthieu Cord. ep-alm: Efficient perceptual augmentation of language models, 2023.
- [Swerdlow *et al.*, 2023] Alexander Swerdlow, Runsheng Xu, and Bolei Zhou. Street-view image generation from a bird’s-eye view layout. *arXiv preprint arXiv:2301.04634*, 2023.
- [Tang *et al.*, 2023] Yun Tang, Antonio A. Bruto da Costa, Jason Zhang, Irvine Patrick, Siddhartha Khastgir, and Paul Jennings. Domain knowledge distillation from large language model: An empirical study in the autonomous driving domain, 2023.
- [Taran *et al.*, 2018] Vlad Taran, Nikita Gordienko, Yuriy Kochura, Yuri Gordienko, Alexandr Rokovyi, Oleg Alienin, and Sergii Stirenko. Performance evaluation of deep learning networks for semantic segmentation of traffic stereo-pair images. In *Proceedings of the 19th International Conference on Computer Systems and Technologies*. ACM, sep 2018.
- [Touvron *et al.*, 2023] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023.
- [Treiber *et al.*, 2000] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical review E*, 62(2):1805, 2000.
- [Unterthiner *et al.*, 2019] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric and challenges, 2019.
- [Vedantam *et al.*, 2015] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation, 2015.
- [Wang *et al.*, 2023a] Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. Is chatgpt a good nlg evaluator? a preliminary study, 2023.
- [Wang *et al.*, 2023b] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xinze Chen, and Jiwen Lu. Drivedreamer: Towards real-world-driven world models for autonomous driving. *arXiv preprint arXiv:2309.09777*, 2023.
- [Wen *et al.*, 2023] Licheng Wen, Daocheng Fu, Xin Li, Xinyu Cai, Tao Ma, Pinlong Cai, Min Dou, Botian Shi, Liang He, and Yu Qiao. Dilu: A knowledge-driven approach to autonomous driving with large language models. *arXiv preprint arXiv:2309.16292*, 2023.
- [Wu *et al.*, 2022a] Penghao Wu, Li Chen, Hongyang Li, Xiaosong Jia, Junchi Yan, and Yu Qiao. Policy pre-training for autonomous driving via self-supervised geometric modeling. In *The Eleventh International Conference on Learning Representations*, 2022.
- [Wu *et al.*, 2022b] 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, 2022.
- [Wu *et al.*, 2023] Dongming Wu, Wencheng Han, Tiancai Wang, Yingfei Liu, Xiangyu Zhang, and Jianbing Shen. Language prompt for autonomous driving, 2023.
- [Xu *et al.*, 2021a] Li Xu, He Huang, and Jun Liu. Sutd-trafficqa: A question answering benchmark and an efficient network for video reasoning over traffic events. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9878–9888, June 2021.
- [Xu *et al.*, 2021b] Li Xu, He Huang, and Jun Liu. Sutd-trafficqa: A question answering benchmark and an efficient network for video reasoning over traffic events, 2021.
- [Xu *et al.*, 2022] Danfei Xu, Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Bits: Bi-level imitation for traffic simulation, 2022.
- [Xu *et al.*, 2023] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K. Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model, 2023.
- [Yang *et al.*, 2023a] Kairui Yang, Enhui Ma, Jibin Peng, Qing Guo, Di Lin, and Kaicheng Yu. Bevcontrol: Accurately controlling street-view elements with multi-perspective consistency via bev sketch layout. *arXiv preprint arXiv:2308.01661*, 2023.
- [Yang *et al.*, 2023b] Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. The dawn of Imms: Preliminary explorations with gpt-4v(ision), 2023.
- [Yin *et al.*, 2021] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking, 2021.
- [Zeng *et al.*, 2022] Fangao Zeng, Bin Dong, Yuang Zhang, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. Motr: End-to-end multiple-object tracking with transformer. In *European Conference on Computer Vision (ECCV)*, 2022.
- [Zhang *et al.*, 2023a] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual

language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023.

[Zhang *et al.*, 2023b] Siyao Zhang, Daocheng Fu, Zhao Zhang, Bin Yu, and Pinlong Cai. Trafficgpt: Viewing, processing and interacting with traffic foundation models. *arXiv preprint arXiv:2309.06719*, 2023.

[Zhong *et al.*, 2022] Ziyuan Zhong, Davis Rempe, Danfei Xu, Yuxiao Chen, Sushant Veer, Tong Che, Baishakhi Ray, and Marco Pavone. Guided conditional diffusion for controllable traffic simulation, 2022.

[Zhong *et al.*, 2023] Ziyuan Zhong, Davis Rempe, Yuxiao Chen, Boris Ivanovic, Yulong Cao, Danfei Xu, Marco Pavone, and Baishakhi Ray. Language-guided traffic simulation via scene-level diffusion, 2023.