

# iMotion-LLM: Instruction-Conditioned Trajectory Generation

Abdulwahab Felemban<sup>1</sup> Nussair Hroub<sup>1</sup> Jian Ding<sup>1</sup> Eslam Abdelrahman<sup>1</sup> Xiaoqian Shen<sup>1</sup>  
 Abdullah Mohamed<sup>2\*</sup> Mohamed Elhoseiny<sup>1</sup>

<sup>1</sup>King Abdullah University of Science and Technology (KAUST) <sup>2</sup>Meta Reality Labs

<sup>1</sup>{firstname.lastname}@kaust.edu.sa <sup>2</sup>abdullahadel@meta.com

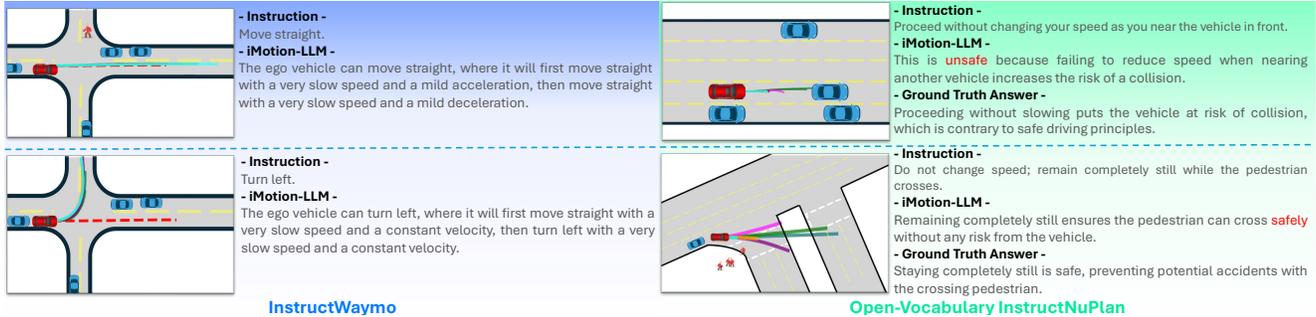


Figure 1. iMotion-LLM generates feasible, safety-aligned trajectories from human instructions. It uses two datasets, InstructWaymo (direction-based) and Open-Vocabulary InstructNuPlan (safety-focused), to support instruction-conditioned generation and justification.

## Abstract

We introduce *iMotion-LLM*, a large language model (LLM) integrated with trajectory prediction modules for interactive motion generation. Unlike conventional approaches, it generates feasible, safety-aligned trajectories based on textual instructions, enabling adaptable and context-aware driving behavior. It combines an encoder-decoder multimodal trajectory prediction model with a pre-trained LLM fine-tuned using LoRA, projecting scene features into the LLM input space and mapping special tokens to a trajectory decoder for text-based interaction and interpretable driving. To support this framework, we introduce two datasets: 1) *InstructWaymo*, an extension of the Waymo Open Motion Dataset with direction-based motion instructions, and 2) *Open-Vocabulary InstructNuPlan*, which features safety-aligned instruction-caption pairs and corresponding safe trajectory scenarios. Our experiments validate that instruction conditioning enables trajectory generation that follows the intended condition. *iMotion-LLM* demonstrates strong contextual comprehension, achieving 84% average accuracy in direction feasibility detection and 96% average accuracy in safety evaluation of open-vocabulary instructions. This work lays the foundation for text-guided motion generation in autonomous driving, supporting simulated data generation, model interpretability, and robust safety alignment testing for trajectory generation models. Our code, pre-trained model, and datasets are available at: [vision-cair.github.io/iMotion-LLM/](https://vision-cair.github.io/iMotion-LLM/).

\*This work was done outside of Meta in a personal capacity.

## 1. Introduction

Trajectory prediction is a core component of autonomous driving pipelines, enabling the ego vehicle to anticipate surrounding agents and plan safe maneuvers [27]. Large-scale benchmarks such as the Waymo Open Motion Dataset (WOMD) [19] and the Waymo Sim Agents Challenge [45] have advanced prediction accuracy and realism, but they provide limited controllability: models forecast futures from past motion without explicit mechanisms to direct agent behavior. As autonomous systems mature, the ability to *control* generated trajectories, particularly through natural language, becomes essential for scenario generation, safety testing, and interpretability.

Recent research has explored language-conditioned traffic and trajectory generation, including LANG-TRAJ [7], InteractTraj [61], ProSim-Instruct [56], and others [20, 55, 64, 65]. While these methods demonstrate progress in scene-level realism or simulation-based evaluation, none provide a rigorous framework for measuring controllability under diverse language inputs. In contrast, our work establishes *Instruction-Conditioned Trajectory Generation* as a task: we pair large-scale real-world datasets with textual instructions, introduce a direct metric for instruction adherence (IFR), and compare against strong trajectory prediction backbones to ensure preserving trajectory quality.

Our framework lays the foundation for next-generation AV systems that can reason and act safely based on language. While NuPlan [6] enables sampling of challenging

cases through scenario mining, our setting emphasizes semantically controlled trajectories. This is similar in spirit to concurrent efforts [7, 56], but distinguishes itself by providing a rigorous evaluation of controllability. As a result, developers can not only represent challenging scenarios, but also manipulate agent motion to stress-test downstream planning under diverse conditions, where prior work such as LCTGen [55] has shown that such manipulation can significantly degrade the performance of policy-based planners.

**Our contributions are:**

- We introduce two new datasets: **InstructWaymo** and **Open-Vocabulary InstructNuPlan**, which augment open-loop driving scenarios with direction-based and free-form, safety-aligned instructions.
- We develop instruction-conditioned models, including an LLM-based approach that integrates feasibility and safety reasoning for trajectory generation.
- We propose **Instruction Following Recall (IFR)**, an effective metric to evaluate adherence to instructions.

## 2. Related Work

**Multimodal Large Language Models.** Large Language Models (LLMs) have significantly advanced in recent years [1, 5, 16, 48, 57, 58], with models like GPT-4 [1] demonstrating remarkable abilities in generating coherent, contextually relevant text across numerous domains. With the strong performance of LLMs, there is an emergence of multi-modal LLMs (MLLMs) [3], which extend the LLMs with reasoning abilities across diverse modalities. Notable works include VisualGPT [8], Flamingo [3], MiniGPT-4 [9, 66], LLaVA [37, 38], and InstructBLIP [15]. These works used visual instruction tuning to align with human intentions. There are some extensions that focus on detection and segmentation [4, 35, 59, 66], videos [36, 40, 63], and 3D [23, 25, 62]. Our work focuses on MLLMs for motion prediction tasks.

**Trajectory Prediction Models for Driving Scenarios.** Trajectory prediction analyzes agents’ historical tracks on maps to forecast joint future positions. Early approaches used LSTMs [2, 24] for encoding agent states and CNNs [13, 21, 50] for rasterized scene images. GNNs [11, 29, 43] later improved agent interaction modeling. Transformer-based models like SceneTransformer [47] and WayFormer [46] enhanced prediction efficiency but focused mainly on encoding vectorized representations. Motion Transformer [53, 54] and GameFormer [31] achieved better accuracy through improved decoding, while MotionLM [51] adopted LLM-like structures without incorporating language reasoning capabilities. Earlier work also examined social-awareness in [49].

**Multimodal Large Language Models for Autonomous Driving.** The emergence of LLMs has sparked adaptations for autonomous driving [10, 17, 26, 28]. GPT-Driver [41]

Table 1. Direction categories with their corresponding presence proportion in the train set.

Stationary	Straight	Right	Left	Left U-turn
16,748	863,910	184,286	221,762	4,575

and SurrealDriver [33] demonstrate LLMs’ transformative impact on motion planning and maneuver generation. However, most methods focus on text or image inputs, overlooking vector representation’s benefits for motion prediction. Vector representation abstractly captures essential driving scenario information. Like Driving with LLMs [10], we integrate LLMs with vector-based data, but while [10] introduced QA-focused benchmarks representing motion as single quantized actions (acceleration, braking, steering), we focus on multi-modal multi-agent trajectories, better aligning with existing trajectory prediction modules for safer, more reliable motion prediction.

**Conditional and Promptable Trajectory Generation.**

Several works explore controllable or instruction-driven trajectory generation. **InteractTraj** [61] and **TrafficGen** [20] generate synthetic data with controllability over traffic density or interactions but evaluate mainly on realism. Diffusion-based methods like CTG and CTG++ [64, 65] add goal- or LLM-guided constraints, while LCTGen [55] edits scenarios from free-form text, though without reproducible benchmarks or controllability metrics. Closer to our setting, LANGTRAJ [7] and ProSim-Instruct [56] pair trajectories with instructions, but support limited instruction types and report only realism-based simulation metrics. In contrast, we provide large-scale datasets, a dedicated controllability metric (IFR), and rigorous benchmarking against strong trajectory prediction backbones, while also enabling textual justification of feasibility and safety alignment with driving scenarios, establishing Instruction-Conditioned Trajectory Generation as a standalone task.

## 3. Motion Instruction Datasets

### 3.1. InstructWaymo Dataset

InstructWaymo offers a new perspective on WOMD [18] by making trajectory generation instructable and language descriptive. Inspired by the mAP calculation used in WOMD, which evaluates model performance across various driving behaviors, we designed a module that categorizes future motion based on direction, speed, and acceleration. InstructWaymo uses future direction information as instructions alongside detailed motion descriptions such as two-step direction, speed, and acceleration, which are presented as captions to be used as output text for driving behavior interpretation by models. Additionally, the feasibility of each instruction (direction) is calculated, adding an extra layer of comprehension by identifying feasible and infeasible directions for each driving scenario.

InstructWaymo is publicly available as a script to aug-

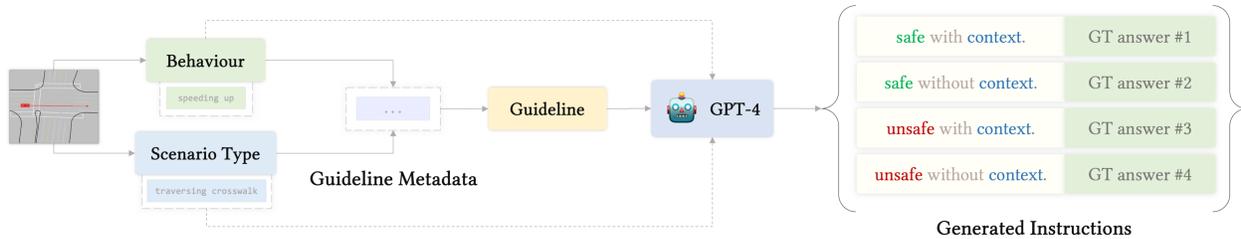


Figure 2. Open-Vocabulary InstructNuPlan pipeline for generating safe/unsafe instruction-caption pairs using GPT-4o mini, guided by scenario type, motion behavior, and safety metadata.

ment WOMD. This data augmentation is applied to approximately 400K driving scenarios, which are preprocessed similarly to the GameFormer preprocessing method. Each scenario includes up to 32 neighboring agents, with a total of 33 agents, including the ego agent. Each agent in the scene is considered a focal agent (the ego-view agent), resulting in 4.2M samples across different focal agents. Among these, 1.3M samples involve a unique vehicle as the focal agent with a valid trajectory and detected instructions.

**Direction.** Direction is fundamental for instructing navigation; we adopted the WOMD direction bucketing script to obtain five conceivable direction conditions encompassing five classes listed in Table 1 with their statistics. The statistics shows a bias toward some behaviors like moving straight. We use driving directions as instructions in the experiments on InstructWaymo.

**Speed and Acceleration.** Following the intuition used in [44], we categorize trajectories of moving vehicles based on speeds and relative change in speeds. For that, we defined 5-speed categories and 9-acceleration categories; the suggested upper threshold and the categories are listed in the supplementary.

**Feasibility of directions.** We define the feasibility of directions into three categories: 1) Ground-Truth direction (GT), which is based on the ground-truth future trajectory and hence is always assumed to be a feasible direction; 2) Feasible directions (F), which are derivable directions but not the GT direction; 3) Infeasible directions (IF), which is the complement set of feasible directions. To assess feasibility, we consider a set of candidate destination lanes relative to the ego vehicle’s current location and heading. These candidate destinations are possible locations on associated lanes within a range determined by the vehicle’s speed (maximum range  $r$ ). This range is calculated based on a maximum speed increase of 15 km/h within 8 seconds, not exceeding the road speed limit, and within a maximum range of 60 meters, indicating a safe range of acceleration or deceleration. For the feasibility of staying stationary, we detect if the vehicle current speed is within 65 km/h, allowing it to slow down to stop in 8 seconds. These heuristics are grounded in general safe driving practices and are configurable in the codebase that is publicly available.

**LLM Instruction and caption.** Using the extracted at-

tributes, we construct a templated pair comprising an input instruction and an output caption for the LLM. The input instruction specifies the target (final) direction the ego vehicle should reach. The output caption—generated autoregressively—covers the target direction and an interpretable execution plan: two intermediate wayfinding steps, along with indicative speeds and accelerations that realize the final direction. This pairing makes the model’s adherence to the instruction explicit and verifiable.

### 3.2. Open-Vocabulary InstructNuPlan Dataset

Open-Vocabulary InstructNuPlan extends NuPlan [6] to explore instruction-conditioned trajectory generation beyond predefined instruction sets. Unlike InstructWaymo, which focuses on structured direction-based commands, this dataset introduces diverse open-vocabulary instructions, enabling more flexible and context-aware motion generation. It also emphasizes safety alignment, distinguishing between safe and unsafe behaviors to evaluate an autonomous vehicle’s ability to follow safe instructions while rejecting unsafe ones. This dataset serves as a benchmark for assessing instruction-following capabilities in real-world driving scenarios. The full process of generating open-vocabulary instruction-caption pairs for a driving scenario is illustrated in Figure 2. The scenario type is provided in the original NuPlan dataset [6].

**Selected Scenario Types.** To capture diverse driving contexts, we select 14 scenario types from the NuPlan dataset, spanning a wide range of real-world decision-making challenges such as interactions with pedestrians, vehicles, and road infrastructure. The chosen scenarios include: Accelerating at crosswalk, Traversing crosswalk, Waiting for pedestrian to cross, Following lane (with/without/with slow lead), Stopping with lead, Starting protected/non-protected turns (cross/non-cross), Behind bike, Behind long vehicle, and Traversing intersection. We sample an equal number of examples per type, yielding 7,119 training and 124 testing scenarios. From each, we generate multiple safe and unsafe instruction–caption pairs using GPT-4o mini, resulting in 115,097 training and 2,078 testing pairs. Following WOMD conventions, we use 1.1 seconds of past motion to predict 8 seconds into the future.

**Safety Grounding and Instruction Generation.** To gen-

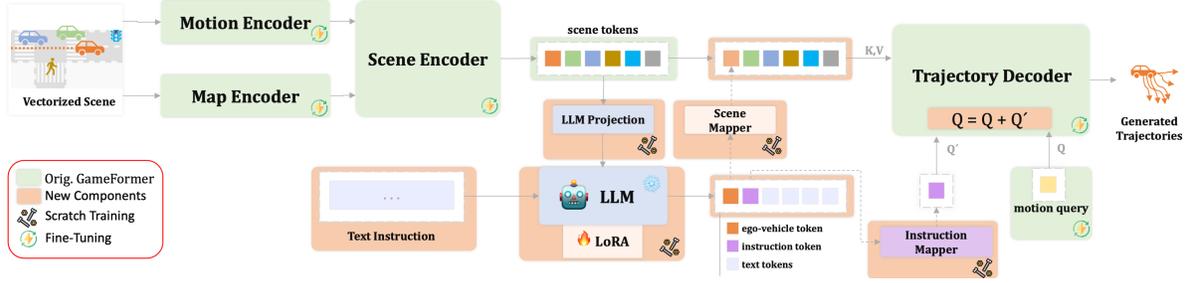


Figure 3. Given a textual instruction and scene context embeddings, iMotion-LLM employs an LLM Projection module to map the encoded scene token embeddings from the Scene Encoder into the LLM input space. The LLM then generates an updated ego-vehicle token, an instruction token, and caption text tokens. The instruction token is projected into a query  $Q'$ , while the ego-vehicle token is projected to form the key and value representing the ego-vehicle embedding, which is subsequently used by the Multimodal Trajectory Decoder. Newly added components are highlighted in orange.

erate safety-aligned instruction–caption pairs, we construct metadata that explicitly defines safe and unsafe behaviors for each scenario type, following established traffic rules and autonomous driving best practices. Each scenario is annotated with up to 10 safe and 10 unsafe behaviors. Given a scenario, we load its metadata, determine the ego vehicle’s future motion, and match it against predefined behavior categories (*not moving*, *stopping*, *waiting then moving*, *slowing down*, *speeding up*, *slowing down then speeding up*, *speeding up then slowing down*, *maintaining speed*). From this, we retrieve guidelines specifying safe and unsafe outcomes. For instance, in a *waiting for a pedestrian to cross* scenario, if the ego vehicle’s behavior is *not moving*, the safe instruction is “Do not move; the vehicle should remain stationary while a pedestrian is crossing”, whereas any movement is considered unsafe. This process ensures that instructions generated by GPT-4o mini are both context-aware and safety-aligned.

**Context-Enriched Instruction Generation.** Beyond labeling safe vs. unsafe behaviors, we generate instructions in two formats to test the model’s ability to handle explicit and implicit cues: 1. **With context:** explicitly references scene elements (e.g., “Slow down before the intersection due to oncoming traffic.”). 2. **Without context:** abstract, context-agnostic phrasing (e.g., “Reduce speed to 30 km/h.”). This dual formulation enables a more comprehensive assessment of contextual understanding. For captions, we always include contextual details to remain aligned with the instruction.

## 4. iMotion-LLM

### 4.1. Baselines

**Non-Conditional Baselines.** Transformer-based trajectory prediction models like GameFormer [30] and MTR [52] typically consist of two components: a Scene Encoder and a Multimodal Trajectory Decoder. The Scene Encoder fuses

observed map and agent information into scene context embeddings ( $scene\ tokens$ )  $\in \mathbb{R}^{R \times d_{scene}}$ , where  $d_{scene}$  is the embedding dimension. Each embedding maintains a fixed correspondence to its source (e.g., ego vehicle motion is always the first token), which preserves spatial grounding. The decoder then uses cross-attention with these embeddings as keys and values ( $K, V$ ), and  $M$  learnable queries  $Q \in \mathbb{R}^{M \times d_{scene}}$  to predict a Gaussian Mixture Model (GMM) of future multimodal trajectories. These two modules—Scene Encoder and Multimodal Trajectory Decoder—form the backbone of such models and are illustrated in Figure 3 in green.

**Conditional Baselines.** cGAN [42] introduces a foundational approach for incorporating conditional labels in image generation. Inspired by this, we enable decoding in originally non-conditional transformer-based baselines to be conditioned on a given instruction label. In the decoder, we introduce an additional learnable query,  $Q' \in \mathbb{R}^{1 \times d_{scene}}$ , which is fused with the motion generation queries,  $Q$ . The fusion is performed by adding  $Q'$  element-wise to each of the  $M$  motion queries. When training the conditional baselines without the use of LLM components (i.e., the orange components in Figure 3),  $Q'$  is learned through a simple embedding layer that takes a categorical class as input. When integrating an LLM with the base model—leveraging its capability to process diverse natural language instructions— $Q'$  is instead derived from the LLM’s output embeddings. Further details on training the conditional baseline are provided in the pseudo-code in the supplementary material.

### 4.2. Text-Guided Trajectory Generation

To enable natural-form text-based guidance and interaction, we incorporate an LLM into pre-trained conditional baselines consisting of a Scene Encoder and a Multimodal Trajectory Decoder. To enable this integration design, illustrated in Figure 3, four main blocks are required: 1.LLM

Projection module. 2.LLM itself. 3.Scene Mapper. 4.Instruction Mapper.

**LLM Projection.** Inspired by Vision Language Models (VLMs) [14, 67], we employ a simple linear projection layer to map input scene embeddings to  $\mathbb{R}^{R \times d_{LLM}}$ , aligning with the LLM embedding dimension  $d_{LLM}$ .

**LLM.** All projected scene embeddings, along with the input text instruction, are fed into the LLM, which generates output tokens representing the ego-vehicle token, the instruction token, and additional textual tokens that form an output caption after grounding the input instruction.

**Scene Mapper.** We map the instruction-grounded ego token from  $\mathbb{R}^{1 \times d_{LLM}}$  back to  $\mathbb{R}^{1 \times d_{scene}}$  using a multilayer perceptron (MLP), referred to as the Scene Mapper. This mapped token replaces the corresponding ego token and is combined with the remaining keys and values from the Scene Encoder, which bypass the LLM. Together, they serve as the keys and values for the Multimodal Trajectory Decoder.

**Instruction Mapper.** Following the Scene Mapper, we project the instruction token back to the motion generation model’s embedding space ( $\mathbb{R}^{1 \times d_{scene}}$ ) using an MLP. The resulting vector is then fused with the motion query  $Q$  via element-wise addition

**Output Caption.** Along with generating scene and instruction tokens, the LLM has the capability of outputting text that describes how the instruction is executed, a textual decision (‘[Accept]’ or ‘[Reject]’) to indicate whether an instruction is feasible or not, and a safety justification for accepting or rejecting the instruction.

## 5. Experiments

### 5.1. Datasets Setup

**Waymo Open Motion Dataset (WOMD).** For training For training unconditional baseline models, we use the WOMD dataset [18], which consists of more than 400K driving scenarios. Considering each agent out of 32 agents as the focal agent, we construct 4.2M training examples to reproduce the results of GameFormer [30] and MTR [52]. We consider a trajectory window of 9 seconds, with 1.1 seconds of observed history.

**InstructWaymo.** Using InstructWaymo, we train conditional models on 1.3M samples, each pairing the ego vehicle with one of five direction instructions. Each scenario is annotated with a ground-truth (GT) instruction, additional feasible (F) alternatives, and infeasible (IF) instructions. For non-LLM conditional baselines, we train from scratch using GT only. In contrast, iMotion-LLM additionally incorporates IF instructions during fine-tuning to strengthen feasibility awareness. For evaluation, we use 1,500 examples balanced across GT/F/IF. Trajectory training follows WOMD settings, except we restrict the dataset to cases

where the ego agent is a vehicle.

**Open-Vocabulary InstructNuPlan.** For open-vocabulary fine-tuning and testing of iMotion-LLM, we use this dataset with 115,097 training examples and 2,078 test examples. We split the test samples into four categories: safe without context, safe with context, unsafe without context, and unsafe with context. We use the same trajectory training setup as WOMD.

**Metrics.** Traditional trajectory metrics such as minADE/minFDE capture geometric accuracy but do not assess whether a model follows the instructed direction. We therefore introduce *Instruction Following Recall (IFR)*, which measures the fraction of generated trajectories whose final direction matches the instructed direction. We extract directions from predictions using the same direction-extraction module as for ground truth. Formally, for  $N$  scenarios and  $M$  generated trajectories per scenario,

$$IFR = \frac{1}{N} \sum_{i=1}^N \frac{1}{M} \sum_{j=1}^M \mathbb{1}[D_{i,j}^{\text{pred}} = D_i^{\text{instr}}], \quad (1)$$

where  $D_i^{\text{instr}}$  is the instructed direction bucket for scenario  $i$ , and  $D_{i,j}^{\text{pred}}$  is the direction bucket extracted from the  $j$ -th generated trajectory. Higher IFR indicates stronger instruction adherence, while minADE/minFDE quantify trajectory quality. In addition, we evaluate feasibility and safety detection with classification accuracy, and report justification-quality metrics in the supplementary material.

### 5.2. Qualitative Results of iMotion-LLM

In Figure 4, iMotion-LLM generates context-aware trajectories for **unsafe** (left) and **safe** (right) instructions, effectively assessing safety and providing grounded, diverse predictions with safety-aligned justifications.

**Unsafe Instructions Results.** The first unsafe case (top-left) instructs an abrupt stop in a crosswalk, risking pedestrians. iMotion-LLM flags this, with five of six trajectories avoiding the stop. The second case (middle-left) instructs maintaining speed while turning, risking loss of control. iMotion-LLM detects this, generating varied-speed trajectories and reinforcing the need to decelerate. The third case (bottom-left) instructs acceleration near a crosswalk, ignoring pedestrians. iMotion-LLM flags the risk, producing both cautious and aggressive trajectories, with textual reasoning aiding safety filtering.

**Safe Instructions Results.** The first safe case (top-right) instructs slowing for a slow-moving lead vehicle and stopping. iMotion-LLM correctly follows this, generating adaptive trajectories and reinforcing its understanding of dynamic agent interactions. The second safe case (middle-right) instructs a left-favoring trajectory at a controlled speed. iMotion-LLM correctly follows this, recognizing the unprotected turn even though it is not explicitly mentioned

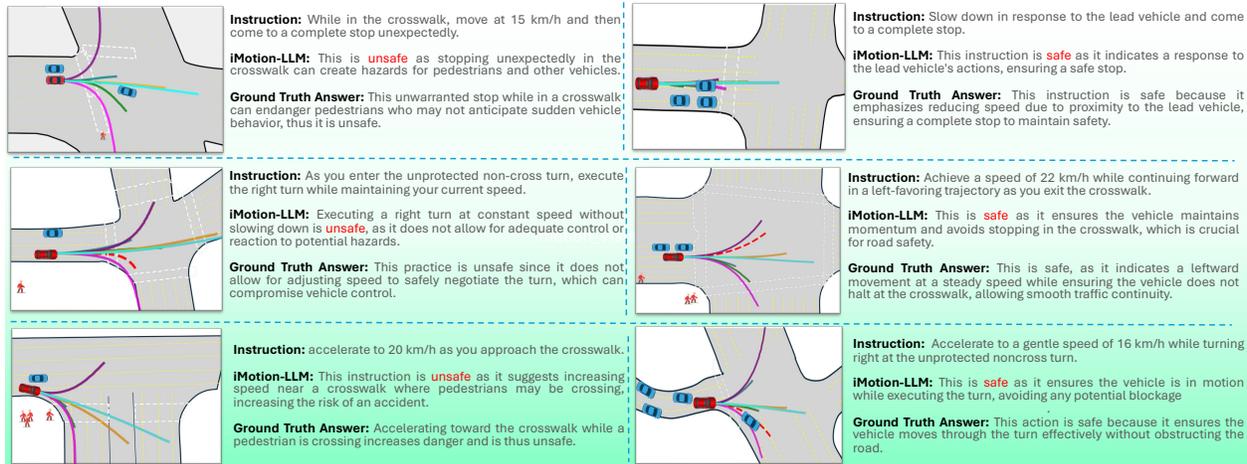


Figure 4. **iMotion-LLM Qualitative Results.** Qualitative results showcasing three unsafe instructions (left) and three safe instructions (right). The results demonstrate iMotion-LLM’s capability to generate relevant trajectories, assess the safety of given instructions, and provide reasoning for its decisions.

Table 2. **InstructWaymo Results.** Instruction Following Recall (IFR) on feasible instruction types: Ground-Truth (GT), Feasible (F), minADE and minFDE, and the feasibility detection capability of different models.

Model	GT-IFR $\uparrow$	F-IFR $\uparrow$	minADE/minFDE $\downarrow$	Feasibility Detection
GameFormer [30]	66.89	15.21	0.78 / 1.64	$\times$
MTR [52]	53.46	16.64	0.74 / 1.62	$\times$
C-GameFormer	83.10	48.70	<b>0.65 / 1.20</b>	$\times$
C-MTR	64.29	52.22	0.67 / 1.39	$\times$
ProSim-Instruct	86.32	24.91	3.52 / 3.9	$\times$
iMotion-LLM	<b>87.30</b>	<b>52.24</b>	0.67 / 1.25	$\checkmark$

in the ground truth response, showcasing strong contextual reasoning. The final safe case (bottom-right) instructs accelerating to 16 km/h while turning. iMotion-LLM classifies this as safe, generating diverse paths. Notably, one trajectory takes a left turn, highlighting the model’s ability to explore multiple reasonable options.

### 5.3. Results on InstructWaymo

**Task and Metrics.** To evaluate iMotion-LLM, we assess the ability of trajectory generation models to follow a given *categorical direction* using the Instruction-Following Recall (IFR) metric. Additionally, we evaluate the model’s ability to reason about instruction feasibility by measuring the accuracy of accepting feasible instructions (GT or F) and rejecting infeasible ones (IF). To complement the instruction-following evaluation, we also report minADE and minFDE on GT-instruction examples to assess the quality of the generated trajectories. These conventional metrics ensure that the model not only follows instructions but also produces accurate and realistic motion plans.

Table 3. **Open-Vocabulary InstructNuPlan Results.** comparison of two different types of input complex instructions in NuPlan: Safe and Unsafe. We also split the evaluation into instructions with and without context. IFR is not reported for unsafe instructions, as corresponding ground-truth trajectories are unavailable.

I/P Instruction	Context	Accuracy (%)	IFR $\uparrow$
Safe	$\times$	96.61	68.96
Safe	$\checkmark$	<b>98.39</b>	<b>70.22</b>
Unsafe	$\times$	95.00	N/A
Unsafe	$\checkmark$	96.88	N/A

**Evaluated Models.** Although unconditional models such as GameFormer [30] and MTR [52] do not take instructions or direction categories as input, we evaluate them as lower-bound baselines to assess the effectiveness of the conditional baselines and iMotion-LLM. This comparison is not intended to demonstrate superiority over unconditional models, as the tasks differ fundamentally, but rather to analyze whether incorporating instruction-following capabilities enables intractable trajectory generation, as implemented in C-GameFormer, C-MTR, and iMotion-LLM.

For training iMotion-LLM, we use a conditionally adapted and pretrained version of GameFormer, referred to as C-GameFormer. We adopt a mixed training strategy: 70% of the data consists of ground-truth (GT) instructions paired with their corresponding trajectories, while the remaining 30% consists of infeasible (IF) instructions, for which no ground-truth trajectory is provided—only the textual instruction indicating that it should be rejected. During training, the text loss is always backpropagated, while the trajectory loss is applied only when a corresponding ground-truth trajectory is available (i.e., for GT instructions). Table 2 summarizes the evaluation results of the

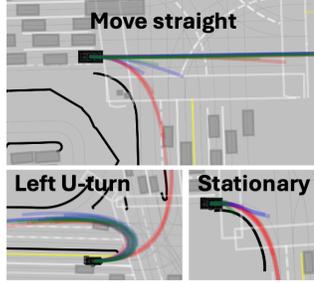


Figure 5. Qualitative comparison of three challenging InstructWaymo scenarios. Non-conditional GameFormer in red, C-GameFormer in blue, and iMotion-LLM in green.

models.

**Results of Non-Conditional Baselines.** While non-conditional MTR and GameFormer can generate realistic driving trajectories, their generated trajectories do not necessarily follow the ground truth direction (GT-IFR), as shown in Table 2. Moreover, as expected, the recall of trajectories generated in other feasible directions (F-IFR) is very low, as these models were neither optimized to cover a diverse range of directions for the same driving scenario nor prompted to do so during inference.

**Results of Conditional Baselines.** Table 2 shows that making models conditional, namely C-GameFormer and C-MTR, improve IFR for both GT and F instructions, while also reducing trajectory error as indicated by lower minADE and minFDE scores. This demonstrates that providing the intended direction as an input signal helps models generate more accurate and instruction-aligned trajectories. For comparison with recent approaches, we evaluate ProSim-Instruct, pretrained on ProSim-Instruct-450k, which is based on WOMD and uses instructions similar to InstructWaymo; in our setting, however, we evaluate single-agent conditioning under the same scenarios and instructions. Although ProSim-Instruct achieves comparable GT-IFR, it suffers from degraded trajectory quality as measured by minADE/minFDE, reflecting its emphasis on standard realism-oriented metrics. In contrast, iMotion-LLM maintains strong trajectory quality while further improving IFR, and introduces a key advantage: the ability to generate textual justifications and detect instruction feasibility. Qualitative results in Figure 5 illustrate these improvements: GameFormer fails to cover diverse directions due to its non-conditional nature, C-GameFormer enables instruction adherence, and iMotion-LLM achieves the best alignment, including in challenging cases such as *stop*. Notably, iMotion-LLM reaches 52.24% F-IFR despite not observing such instructions during training, compared to only 24.91% for ProSim-Instruct, highlighting both the gains of our approach and the persistent challenge of generalizing to alternative controllable driving modes in instruction-conditioned trajectory generation.

## 5.4. Results on Open-Vocabulary InstructNuPlan

**Task and Metrics.** In open-vocabulary instructions, iMotion-LLM is the only model that can address this challenge through its language processing capabilities. We evaluate its performance by assessing its ability to determine instruction safety in a given driving scenario, the primary metric for this task. Additionally, we compute Instruction-Following Recall (IFR) based on the alignment of safe instructions with the ground-truth trajectory, as they are derived from it and the scenario type.

**Safety Detection.** As shown in Table 3, the model demonstrates strong performance in distinguishing between safe and unsafe instructions, with slightly higher accuracy in detecting unsafe instructions. This may be attributed to certain examples containing explicitly unsafe directives, such as “speed up without stopping for the pedestrian crossing.” Interestingly, the model maintains high safety detection accuracy even without additional context. When context is included, accuracy improves slightly, which aligns with the evaluation setup where contextual information serves as an additional hint rather than an inherent advantage.

**Open-Vocabulary Instruction Following.** The IFR performance in Table 3 indicates that iMotion-LLM retains a reasonable level of IFR, surpassing the non-conditional GameFormer baseline reported on InstructWaymo (Table 2), though not reaching the performance of iMotion-LLM on InstructWaymo. Given that the instructions are in natural language and do not necessarily correspond to a predefined direction category, unlike in InstructWaymo where IFR is computed based on structured directional commands, the results remain reasonable. Additionally, qualitative examples suggest that the model exhibits greater diversity in the generated trajectories on Open-Vocabulary InstructNuPlan compared to InstructWaymo.

## 5.5. Ablation Studies

We conducted a series of ablation studies to analyze the influence of different key components on iMotion-LLM performance.

**Impact of Class-Balanced Sampling.** Given the strong class imbalance in motion directions, we compare results with and without class-balanced sampling for GT instructions for fine-tuning iMotion-LLM. As shown in Table 5, balanced sampling leads to consistent improvements in both GT-IFR and F-IFR across most instruction classes.

**LLM Backbone Comparison.** To assess the impact of the LLM’s capabilities, we evaluate iMotion-LLM with various backbones. While LLaMA-7B achieves the best instruction-following recall overall, Mistral-7B shows stronger safety justification accuracy in NuPlan. Table 4 presents a full comparison. All other experiments use LLaMA-7B as the LLM model.

Table 4. Comparison of LLMs on InstructWaymo (left) and Open-Vocabulary InstructNuPlan (right). GT/F/IF-Acc denotes the feasibility detection accuracy. Safe/Safe+Context refer to instructions without and with context, and Avg Acc reports overall safety detection accuracy.

LLM	InstructWaymo				Open-Vocabulary InstructNuPlan				
	minADE/FDE	GT-IFR	GT-Acc	F-IFR	F-Acc	IF-Acc	Safe IFR	Safe+context IFR	Avg Acc
LLaMA-7B [58]	<b>0.67/1.25</b>	<b>87.30</b>	97.33	<b>52.24</b>	62.27	92.73	<b>68.96</b>	<b>70.22</b>	96.25
Mistral-7B [32]	<b>0.67/1.25</b>	86.61	<b>97.93</b>	47.23	<b>64.87</b>	<b>95.00</b>	63.69	63.95	<b>97.25</b>
LLaMA-1B [22]	0.69/1.33	79.31	97.20	35.36	63.93	91.13	62.05	63.20	96.50
Vicuna-7B [12]	0.70/1.35	80.81	94.47	36.56	39.80	06.13	62.26	62.75	96.25

Table 5. Effect of class-balanced sampling. GT-IFR and F-IFR across instruction classes, comparing fine-tuning iMotion-LLM with and without class-balanced sampling.

Class	Without Balance		With Balance	
	GT-IFR ↑	F-IFR ↑	GT-IFR ↑	F-IFR ↑
Stationary	27.61	0.83	51.00	16.89
Straight	92.11	26.61	<b>98.83</b>	<b>70.39</b>
Left	80.78	19.56	95.94	68.11
Right	82.28	19.83	97.61	68.56
Left U-turn	<b>94.00</b>	<b>32.33</b>	93.11	37.28

**Effect of LLM Mapper Depth.** We compare projection modules of varying depth. As shown in Table 6, the combination of linear and 2-layer MLP achieves the best performance across all metrics except GT-Acc and F-Acc.

**Training Strategy for Backbone Modules.** We evaluate different fine-tuning configurations: freezing the backbone, fine-tuning only the trajectory decoder, only the scene encoder, and full fine-tuning. Training only the decoder provides the best balance between performance and stability, as shown in Table 7.

## 6. Limitations

**IFR Metric.** Our IFR metric uses discrete directional thresholds (as in WOMD), which capture basic direction-following but not alignment under complex or open-vocabulary instructions. VLM-based evaluation is a promising alternative, though current VLMs struggle with BEV-rendered inputs.

**Dataset Bias.** The datasets may inherit biases from pre-processing and class imbalance. Vectorized maps introduce lossy simplifications, and Open-Vocabulary InstructNuPlan relies on guided prompts to reduce hallucination, which may limit generalization to unconstrained instructions.

**Multi-Agent Support.** iMotion-LLM is designed for single-agent controllability and shows reduced performance in multi-agent prediction. Richer datasets and models explicitly supporting interactive instructions are needed; initial multi-agent results are included in the supplementary.

**Generalization.** The lower performance on alternative feasible instructions suggests limited robustness. Strong results on ground-truth instructions may partly reflect biases from future log conditioning rather than true controllable

Table 6. Projection Layer Ablation. Comparison of different input/output projection architectures. 2L and 4L denote 2-layer and 4-layer MLPs.

In-proj	Out-proj	minADE/FDE ↓	GT-IFR ↑	GT-Acc ↑	F-IFR ↑	F-Acc ↑	IF-Acc ↑
Linear	Linear	0.78 / 1.58	80.78	<b>97.40</b>	36.83	<b>65.87</b>	92.13
Linear	MLP 2L	<b>0.67 / 1.25</b>	<b>87.30</b>	97.33	<b>52.24</b>	62.27	<b>92.73</b>
MLP 2L	MLP 2L	0.82 / 1.77	66.49	97.13	21.42	63.20	92.13
MLP 4L	MLP 4L	0.76 / 1.57	68.73	97.13	16.49	63.93	91.07

Table 7. Ablation on Backbone Fine-Tuning. Comparison of different fine-tuning strategies in iMotion-LLM.

Training Strategy	minADE/FDE ↓	GT-IFR ↑	GT-Acc ↑	F-IFR ↑	F-Acc ↑	IF-Acc ↑
Frozen	0.78 / 1.64	68.46	97.00	27.19	63.00	92.33
Traj Decoder Only	<b>0.67 / 1.25</b>	<b>87.30</b>	<b>97.33</b>	<b>52.24</b>	62.27	<b>92.73</b>
Scene Encoder Only	0.70 / 1.33	82.32	<b>97.33</b>	39.04	64.53	90.33
Fully Finetuned	0.76 / 1.57	74.53	97.20	41.70	<b>66.07</b>	91.60

generation.

**Computation.** LLM-based models incur higher inference costs than lightweight predictors. While advances in sparse attention [68], quantization [34], activation sparsity [39], and hardware-aware frameworks [60] improve efficiency, LLM-based modeling is currently more suited to offline applications.

## 7. Conclusion

In conclusion, we present iMotion-LLM, a large language model for Instruction-Conditioned Trajectory Generation. By leveraging textual instructions, it produces contextually relevant, safety-aligned trajectories while interpreting diverse commands. Through LoRA fine-tuning of a pre-trained LLM, iMotion-LLM maps scene features into the LLM input space, enabling grounded and interpretable motion reasoning. Despite the added compute cost of LLMs, the improvements in interpretability, flexibility, and safety alignment justify the trade-off for safety-critical offline applications such as simulation. Our results show that iMotion-LLM, guided by InstructWaymo, effectively aligns with feasible instructions and rejects infeasible ones. Open-Vocabulary InstructNuPlan further enhances safety-aware decision-making. This work lays the foundation for text-guided motion generation in autonomous driving, enabling simulation of challenging scenarios and robust safety reasoning. Further dataset and training details are provided in the supplementary, and all code and models are publicly released for reproducibility.

## 8. Acknowledgments

Special thanks to Habib Slim, Lucas Bezerra, and Yahia Battach for their valuable assistance during the preparation of this work. For computer time, this research used Ibex managed by the Supercomputing Core Laboratory at King Abdullah University of Science & Technology (KAUST) in Thuwal, Saudi Arabia.

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2
- [2] Alexandre Alahi, Kratharth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016. 2
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 2
- [4] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 1(2):3, 2023. 2
- [5] 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. 2
- [6] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric Wolff, Alex Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles. *arXiv preprint arXiv:2106.11810*, 2021. 1, 3
- [7] Wei-Jer Chang, Wei Zhan, Masayoshi Tomizuka, Manmohan Chandraker, and Francesco Pittaluga. Langtraj: Diffusion model and dataset for language-conditioned trajectory simulation, 2025. 1, 2
- [8] Jun Chen, Han Guo, Kai Yi, Boyang Li, and Mohamed Elhoseiny. Visualgpt: Data-efficient adaptation of pretrained language models for image captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18030–18040, 2022. 2
- [9] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. Minigpt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*, 2023. 2
- [10] 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. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024. 2
- [11] Yuxiao Chen, Boris Ivanovic, and Marco Pavone. Scept: Scene-consistent, policy-based trajectory predictions for planning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17103–17112, 2022. 2
- [12] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, 2023. 8
- [13] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2090–2096. IEEE, 2019. 2
- [14] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. InstructBLIP: Towards general-purpose vision-language models with instruction tuning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 5
- [15] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven C. H. Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *CoRR*, abs/2305.06500, 2023. 2
- [16] 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. 2
- [17] 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. *arXiv preprint arXiv:2310.02251*, 2023. 2
- [18] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. Qi, Y. Zhou, Z. Yang, A. Chouard, P. Sun, J. Ngiam, V. Vasudevan, A. McCauley, J. Shlens, and D. Anguelov. Large scale interactive motion forecasting for autonomous driving : The waymo open motion dataset. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9690–9699, Los Alamitos, CA, USA, 2021. IEEE Computer Society. 2, 5
- [19] Scott Ettinger, Shuyang Cheng, Benjamin Caine, Chenxi Liu, Hang Zhao, Sabeek Pradhan, Yuning Chai, Ben Sapp, Charles R Qi, Yin Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9710–9719, 2021. 1

- [20] 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. 1, 2
- [21] Thomas Gilles, Stefano Sabatini, Dzmitry Tsishkou, Bogdan Stanculescu, and Fabien Moutarde. Home: Heatmap output for future motion estimation. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 500–507. IEEE, 2021. 2
- [22] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, et al. The LLama 3 herd of models, 2024. 8
- [23] Ziyu Guo, Renrui Zhang, Xiangyang Zhu, Yiwen Tang, Xi-anzheng Ma, Jiaming Han, Kexin Chen, Peng Gao, Xi-anzhi Li, Hongsheng Li, et al. Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. *arXiv preprint arXiv:2309.00615*, 2023. 2
- [24] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [25] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *arXiv preprint arXiv:2307.12981*, 2023. 2
- [26] 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. 2
- [27] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, Lewei Lu, Xiaosong Jia, Qiang Liu, Jifeng Dai, Yu Qiao, and Hongyang Li. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1
- [28] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pages 9118–9147. PMLR, 2022. 2
- [29] Zhiyu Huang, Xiaoyu Mo, and Chen Lv. Multi-modal motion prediction with transformer-based neural network for autonomous driving. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2605–2611. IEEE, 2022. 2
- [30] Zhiyu Huang, Haochen Liu, and Chen Lv. Gameformer: Game-theoretic modeling and learning of transformer-based interactive prediction and planning for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3903–3913, 2023. 4, 5, 6
- [31] Zhiyu Huang, Haochen Liu, and Chen Lv. Gameformer: Game-theoretic modeling and learning of transformer-based interactive prediction and planning for autonomous driving. *arXiv preprint arXiv:2303.05760*, 2023. 2
- [32] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. 8
- [33] 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. *arXiv preprint arXiv:2309.13193*, 2023. 2
- [34] Jihoon Kim, Minseok Kim, Haejun Lee, Sunghyun Park, Jinho Lee, Jinwoo Kim, and Minsoo Kang. Fast and efficient 2-bit llm inference on gpu. *arXiv preprint arXiv:2311.16442*, 2023. 8
- [35] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. *arXiv preprint arXiv:2308.00692*, 2023. 2
- [36] KunChang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 2
- [37] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. 2
- [38] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 2
- [39] Zihan Liu, Fan Yang, Yikang Zhang, Tao Lin, Mingjie Sun, Han Wang, and Zhouhan Lin. Q-sparse: All large language models can be fully sparsely-activated. *arXiv preprint arXiv:2407.10969*, 2024. 8
- [40] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023. 2
- [41] Jiageng Mao, Yuxi Qian, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023. 2
- [42] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets, 2014. 4
- [43] Xiaoyu Mo, Zhiyu Huang, Yang Xing, and Chen Lv. Multi-agent trajectory prediction with heterogeneous edge-enhanced graph attention network. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):9554–9567, 2022. 2
- [44] Abduallah Mohamed, Deyao Zhu, Warren Vu, Mohamed El-hoseiny, and Christian Claudel. Social-implicit: Rethinking trajectory prediction evaluation and the effectiveness of implicit maximum likelihood estimation. In *European Conference on Computer Vision*, pages 463–479. Springer, 2022. 3
- [45] Nico Montali, John Lambert, Paul Mougins, Alex Kuefler, Nicholas Rhinehart, Michelle Li, Cole Gulino, Tristan Emrich, Zoey Yang, Shimon Whiteson, Brandyn White, and Dragomir Anguelov. The waymo open sim agents challenge. In *Advances in Neural Information Processing Systems*, pages 59151–59171. Curran Associates, Inc., 2023. 1
- [46] Nigamaa Nayakanti, Rami Al-Rfou, Aurick Zhou, Kratharth Goel, Khaled S Refaat, and Benjamin Sapp. Wayformer:

- Motion forecasting via simple & efficient attention networks. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2980–2987. IEEE, 2023. 2
- [47] Jiquan Ngiam, Vijay Vasudevan, Benjamin Caine, Zhengdong Zhang, Hao-Tien Lewis Chiang, Jeffrey Ling, Rebecca Roelofs, Alex Bewley, Chenxi Liu, Ashish Venugopal, et al. Scene transformer: A unified architecture for predicting future trajectories of multiple agents. In *International Conference on Learning Representations*, 2021. 2
- [48] 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. 2
- [49] Saeed Saadatnejad, Mohammadhossein Bahari, Pedram Khorsandi, Mohammad Saneian, Seyed-Mohsen Moosavi-Dezfooli, and Alexandre Alahi. Are socially-aware trajectory prediction models really socially-aware? *Transportation Research Part C: Emerging Technologies*, 141:103705, 2022. 2
- [50] Tim Salzman, Boris Ivanovic, Punarjay Chakravarty, and Marco Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 683–700. Springer, 2020. 2
- [51] Ari Seff, Brian Cera, Dian Chen, Mason Ng, Aurick Zhou, Nigamaa Nayakanti, Khaled S Refaat, Rami Al-Rfou, and Benjamin Sapp. Motionlm: Multi-agent motion forecasting as language modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8579–8590, 2023. 2
- [52] 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*, 2022. 4, 5, 6
- [53] 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. 2
- [54] Shaoshuai Shi, Li Jiang, Dengxin Dai, and Bernt Schiele. Mtr++: Multi-agent motion prediction with symmetric scene modeling and guided intention querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 2
- [55] Shuhan Tan, Boris Ivanovic, Xinshuo Weng, Marco Pavone, and Philipp Kraehenbuehl. Language conditioned traffic generation. In *7th Annual Conference on Robot Learning*, 2023. 1, 2
- [56] Shuhan Tan, Boris Ivanovic, Yuxiao Chen, Boyi Li, Xinshuo Weng, Yulong Cao, Philipp Krähenbühl, and Marco Pavone. Promptable closed-loop traffic simulation. In *8th Annual Conference on Robot Learning*, 2024. 1, 2
- [57] 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. 2
- [58] 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. 2, 8
- [59] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [60] Hengrui Wu, Xu Zhang, Hongyi Zhang, Zhiwei Lin, Jiawei Li, Mingyu Liu, Bo Yuan, and Yufei He. Flightllm: Efficient large language model inference with a complete mapping flow on fpgas. *arXiv preprint arXiv:2401.03868*, 2024. 8
- [61] Junkai Xia, Chenxin Xu, Qingyao Xu, Yanfeng Wang, and Siheng Chen. Language-driven interactive traffic trajectory generation. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2025. Curran Associates Inc. 1, 2
- [62] Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. *arXiv preprint arXiv:2308.16911*, 2023. 2
- [63] 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. 2
- [64] 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. In *7th Annual Conference on Robot Learning*, 2023. 1, 2
- [65] Ziyuan Zhong, Davis Rempe, Danfei Xu, Yuxiao Chen, Sushant Veer, Tong Che, Baishakhi Ray, and Marco Pavone. Guided conditional diffusion for controllable traffic simulation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3560–3566, 2023. 1, 2
- [66] 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. 2
- [67] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. MiniGPT-4: Enhancing vision-language understanding with advanced large language models. In *The Twelfth International Conference on Learning Representations*, 2024. 5
- [68] Yufei Zhu, Yuekai Huang, Xiaohang Zheng, Yizhi Zhao, Bowen Peng, Xipeng Wang, Shijie Yang, Zhongzhi Chen, and Yang Liu. Star attention: Efficient llm inference over long sequences. *arXiv preprint arXiv:2411.17116*, 2024. 8