

# Can Vehicle Motion Planning Generalize to Realistic Long-tail Scenarios?

Marcel Hallgarten<sup>1,2</sup> Julian Zapata<sup>2,3</sup> Martin Stoll<sup>2</sup> Katrin Renz<sup>1,4</sup> Andreas Zell<sup>1</sup>

<sup>1</sup>University of Tübingen <sup>2</sup>Robert Bosch GmbH <sup>3</sup>University of Duisburg-Essen <sup>4</sup>Tübingen AI Center

<https://github.com/mh0797/interPlan>

**Abstract**— Real-world autonomous driving systems must make safe decisions in the face of rare and diverse traffic scenarios. Current state-of-the-art planners are mostly evaluated on real-world datasets like nuScenes (open-loop) or nuPlan (closed-loop). In particular nuPlan seems to be an expressive evaluation method since it is based on real-world data and closed-loop, yet it mostly covers basic driving scenarios. This makes it difficult to judge a planner’s capabilities to generalize to rarely-seen situations. Therefore, we propose a novel closed-loop benchmark *interPlan* containing several edge cases and challenging driving scenarios. We assess existing state-of-the-art planners on our benchmark and show that neither rule-based nor learning-based planners can safely navigate the *interPlan* scenarios.

A recently evolving direction is the usage of foundation models like large language models (LLM) to handle generalization. We evaluate an LLM-only planner and introduce a novel hybrid planner that combines an LLM-based behavior planner with a rule-based motion planner that achieves state-of-the-art performance on our benchmark.

## I. INTRODUCTION

Generalization to previously unseen driving scenarios is crucial to achieving autonomy and motivates research on learning-based vehicle motion planning methods. Yet, simple rule-based planners outperform learning-based methods in the recently proposed closed-loop nuPlan benchmark [1]. Even more surprisingly, the rule-based method PDM-Closed [2] achieves a nearly perfect score, suggesting that it is capable of tackling the enormous challenge of real-world driving. In this work, we challenge this conception. The fact that PDM-Closed [2] is a reactive centerline planner, unable to do lane changes or interact in a targeted way with surrounding traffic agents, underlines that the nuPlan benchmark mostly covers basic driving scenarios. While solving these is a fundamental requirement for motion planning methods, it falls short of proving that they generalize to rare scenarios from the long-tailed distribution of potential real-world scenarios.

**The *interPlan* benchmark.** During complex traffic maneuvers there is often mutual influence between traffic participants. Mastering these bi-directional interactions is crucial for smooth progress, ultimately making it a fundamental enabler for autonomy [3], [4]. Therefore, we propose a novel benchmark to evaluate vehicle motion planning methods in highly interactive and rare scenarios. We take nuPlan [1] scenarios as a starting point for our benchmark and augment them with additional agents, obstacles, or alternative navigation goals, creating realistic long-tail scenarios. In total, our benchmark comprises 80 scenarios covering the following situations: Nudging around parked vehicles, overtaking obstacles, passing a construction zone, passing an accident site,

	nuPlan [1]	Carla [5]	interPlan (ours)
Task	Planning	E2E Driving	Planning
Based on real-world scenarios	✓	✗	✓
Diverse traffic agents	✗	✗	✓
Long-tail scenarios	✗	✓	✓
Interactive lane-changes	✗	✓	✓
Scenario-generation interface	✗	✓	✓
LLM interface	✗	✗	✓

TABLE I: Comparative summary of planning benchmarks

facing jaywalkers, as well as lane changes in low, medium, and high traffic density. A comparative overview is shown in Tab. I.

**State-of-the-art methods do not generalize to long-tail scenarios.** We conduct experiments with an exhaustive set of state-of-the-art planning methods and reveal critical shortcomings in their ability to generalize to difficult unseen scenarios. Rule-based vehicle motion planning methods rely on predefined rules and behaviors to navigate the vehicle. They fail to handle complex and uncommon driving scenarios not covered by the limited rules. On the other hand, learning-based methods are impaired by a lack of robustness and thus cannot deliver the promise of better generalization.

**LLM-based planning.** Foundation models, especially Large Language Models (LLMs), have shown outstanding world-understanding and generalization ability. This makes them tailored to analyze complex driving environments and make informed and well-founded decisions. Therefore, we propose LLM-based planning baselines for our novel benchmark and analyze their capabilities in closed-loop simulation.

**Contributions.** We summarize our contributions as follows:

- We propose the *interPlan* benchmark, a publicly available closed-loop driving benchmark focused on highly interactive and difficult scenarios. We evaluate and analyze a comprehensive set of state-of-the-art methods.
- We demonstrate that even though some methods achieve excellent results in common driving scenarios, they fail in complex long-tail situations, e.g., passing accident sites.
- Fueled by the rising interest in motion planning based on LLMs, we implement GPT-Driver [6] as a baseline and challenge its abilities.
- We propose a novel two-stage planner that combines an LLM-based behavior planner with a downstream rule-based motion planner. It led to a new state-of-the-art on the novel benchmark and serves as a strong baseline for future research.

## II. RELATED WORK

**Closed-loop vehicle motion planning.** Training a learning-based planner through imitation learning is a straightforward and widely adopted approach [7]–[12] pioneered by ALVINN [13]. Rule-based methods, on the other hand, employ a hand-designed decision-making framework with interpretable rules [14]–[20]. For instance, the Intelligent Driver Model (IDM) [21] always follows a lane-centerline while maintaining a safe distance to the leading vehicle. Two predominant methods exist in evaluating vehicle motion planners. Open-loop ego forecasting compares the planned trajectory to an expert’s ground-truth trajectory using distance-based metrics. In contrast, closed-loop evaluation involves simulating the planned trajectory with realistic environmental feedback, such as vehicle dynamics and reactions of surrounding traffic agents. To assess the interactive behavior in complex and interactive scenarios, closed-loop evaluation is indispensable, particularly as open-loop ego-forecasting evaluation was shown to be misaligned with closed-loop driving [2], [22]. The most widely adopted driving simulators are CARLA [5] and nuPlan [1]. The simplistic CARLA [5] simulator uses synthetic data, which is not guaranteed to produce realistic driving situations. On the other hand, nuPlan [1] is based on real-world data, which comes at the cost of undersampling rare and critical scenarios, since collecting them is prohibitive due to the cost and danger involved. For instance, obtaining recordings of an emergency break caused by a child running on the street is dangerous and unethical. While the Carla Leaderboard fueled research for end-to-end methods that directly process sensor information [8], [9], [23]–[27], the nuPlan competition focused on modular planners that use outputs of upstream perception modules (e.g., bounding boxes, maps) as inputs [28]–[33]. In this work, we make use of the nuPlan closed-loop simulation to assess planners independently from perception noise. In order to assess generalization capabilities to rare driving situations, our proposed *interPlan* benchmark is made up of augmented real-world scenarios.

**Benchmarking in complex interactive scenarios.** Autonomous driving systems need to meet high safety standards. In particular, they must be extraordinarily robust when confronted with unseen scenarios. Various approaches exist to generate safety-critical scenarios for the synthetic Carla simulator [5], [34]–[36]. These methods focus on robustness in the face of rule-breaking or assertive agents and behaviors that mislead the ego-vehicle (EV) into questionable decisions. Other approaches include gradient-based optimization [37], [38] and prior-knowledge based scenario construction [39]–[41]. Similarly, we augment driving scenarios to assess the interactive behavior of planning methods, facing a variety of traffic conditions ranging from sparse conservative traffic to high-density assertive traffic. Moreover, we also test scenarios that are not safety-critical but require profound reasoning and robust generalization, such as passing an accident site or a construction zone. In contrast to previous work, we base our scenario generation on real-world scenarios from the widely

adopted nuPlan dataset and simulator. We thus benefit from an unprecedented amount of typical driving scenarios for training before testing a planner’s generalization to long-tail scenarios.

**LLM-based vehicle motion planning.** With the success of Large Language Models (LLMs), a vital field of research emerged from applying them to vehicle motion planning. GPT-Driver [6] proposes to model motion planning as a language modeling task. By feeding outputs of a perception module [42] into a prompt and instructing the LLM to perform chain-of-thought reasoning to ultimately output a safe motion trajectory, it achieves state-of-the-art performance on the open-loop nuScenes benchmark [43]. This was extended by AgentDriver [44], which uses another LLM to generate the task prompt. We introduce a hybrid approach that combines the outstanding world understanding and common sense of an LLM-based behavior planner with the excellent robustness of a rule-based motion planner. While prior work demonstrated strong reasoning capabilities of LLMs in traffic scenarios [45]–[48], methods are often only evaluated in open-loop evaluation [6], [42], [44], [49], [50], in the synthetic Carla simulator [5], [51], or the simplistic HighwayEnv environment [52]–[54]. In contrast, our closed-loop *interPlan* benchmark is based on augmented real-world data, covers 80 complex driving scenarios, and includes a comprehensive set of baselines.

## III. REALISTIC SCENARIO GENERATION

### A. Problem Setup

Generalizing to difficult unseen scenarios is a crucial capability to enable real-world autonomy. However, extracting useful scenarios for benchmarking from real-world recordings is practically infeasible due to their low probability. Staging them explicitly in real contexts is costly and often morally questionable. At the same time, generating such scenarios from scratch raises the question of realism. Therefore, we re-use scenarios from a large-scale real-world dataset and augment them to create difficult and rare scenarios that represent the long tails of the distribution of real-world situations, such as encountering construction zones, jaywalking pedestrians, and accident sites. We chose the large-scale nuPlan dataset [1] and the respective closed-loop simulator with reactive agents.

Being able to make plans in anticipation of surrounding agents’ intentions and reactions is a fundamental pillar of real-world driving and essential to solving scenarios like unprotected turns, overtakes, or lane changes [4]. Therefore, we also evaluate state-of-the-art planning methods in highly interactive scenarios, such as highway merges at different traffic densities or overtaking parked vehicles. In the following, we describe how we augment the nuPlan scenarios to generate challenging interactive and rare scenarios to build our *interPlan* benchmark.

### B. Scenario Generation

**Realistic long-tail scenarios.** In our benchmark, we aim to test long-tail scenarios, which are underrepresented in real-world driving and thus might not occur in large-scale

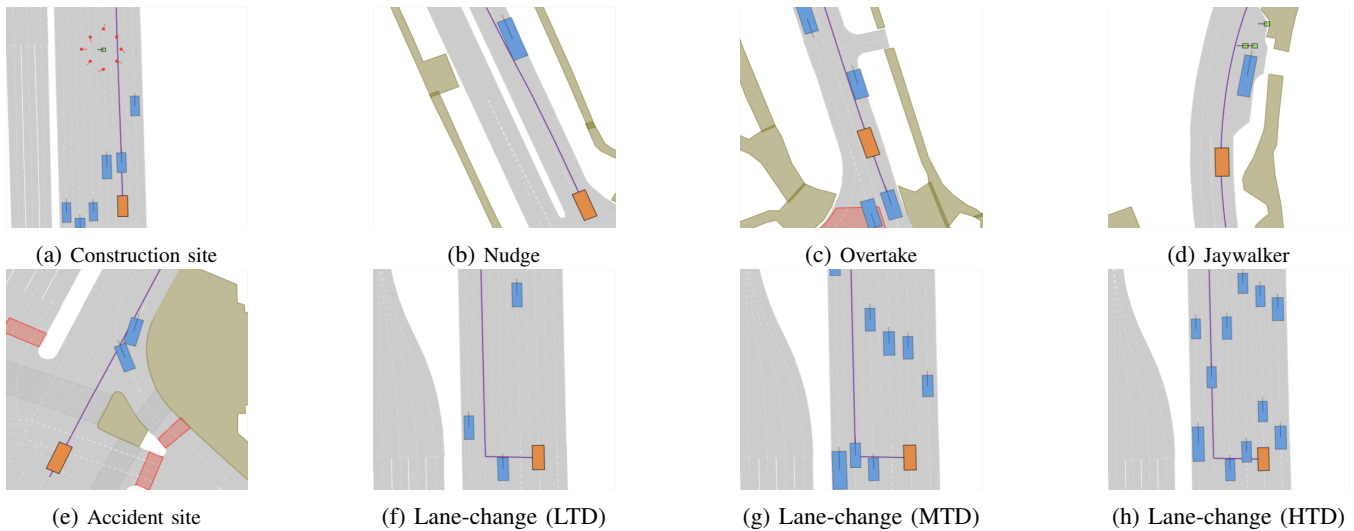


Fig. 1: **interPlan scenario types.** The ego vehicle and its navigation route are shown in orange and purple, surrounding vehicles and pedestrians are blue and green respectively. Traffic cones and stop lines are depicted in red.

training data. Such scenarios include encountering a lane blocked by construction zones (cf. Fig. 1a), nudging around (cf. Fig. 1b) or overtaking (cf. Fig. 1c) a parked vehicle, and facing jaywalkers at bus stops (cf. Fig. 1d). We generate these scenarios for our benchmark by adding objects such as traffic cones, parked vehicles, or stopped busses to the original nuPlan scenarios while leaving the EV state unchanged. Moreover, pedestrians who start crossing the street right when the EV approaches test the planner’s ability to anticipate such non-compliant behaviors. In addition, we include scenarios where the planner has to navigate around an accident site (cf. Fig. 1e). This is particularly challenging because crashed vehicles with intersecting bounding boxes are found rarely, if at all, in real-world datasets. Thus, sophisticated reasoning is required to understand that these vehicles have collided and cannot be expected to start moving, so they must be overtaken carefully. We design the crash sites in such a way that they represent common accidents such as rear-end collisions or collisions with crossing traffic.

**Interactive lane-change scenarios.** A major shortcoming of the original nuPlan benchmark is that it can be solved almost perfectly by simple lane-following, thus with little interaction with surrounding traffic. For instance, the challenge-winning method PDM-Closed [2] is by design unable to perform lane changes but achieves a nearly perfect score. To counteract this, we incentivize lane changes by augmenting the goal of original nuPlan scenarios. For instance, in the scenario shown in Fig. 1f, the EV has to do 3 lane changes. Moreover, we initialize the scenarios with different levels of traffic densities, i.e., low (LTD), medium (MTD), and high traffic density (HTD), by spawning traffic agents around the EV. Agents are randomly placed along each lane with sufficient spacing to avoid collisions, with maximum distances of 33m, 50m, and 100m for HTD, MTD, and LTD.

**Augmented traffic agent behavior.** Real-world drivers exhibit diverse behaviors ranging from conservative and careful driving to assertive and reckless maneuvers. Planning

methods have to be able to handle these varying conditions, especially without knowing what they are about to encounter. Our benchmark tests the ability of planning methods to make safe decisions under these varying conditions by applying different policies to control the surrounding traffic agents. Our primary objective is not realistic traffic simulation, which is an ongoing research field itself, but to create diverse and reasonable traffic agents that can effectively test a planner’s interactive behavior. Hence, in addition to nuPlan’s default reactive IDM policy, which breaks as soon as the EV enters its lane, we implement an assertive policy, which only reacts to the EV when it has fully merged into the lane. Simulation of merging scenarios can be carried out with all agents being controlled by the conservative reactive policy, the assertive policy, or a mixed policy, which randomly assigns a policy to each agent. The lane-change scenarios in our benchmark combine these traffic agent policies with different traffic densities. Thus, lane changes in HTD require complex interaction and potentially even prioritizing safety over progress. At the same time, LTD scenarios test a planner’s general ability to merge into another lane.

### C. The interPlan Benchmark

We test an exhaustive list of state-of-the-art planners, including learning-based (PDM-Open [2], UrbanDriver [28], GC-PGP [30], Gameformer [31], DTPP [32]) and rule-based planners (IDM [21], IDM+MOBIL [55], PDM-Closed [2]). While some strictly follow lane-centerlines, others are not limited regarding their planning space. We intend to evaluate how well these planners can generalize to unseen scenarios, which is an indispensable capability for safe deployment. As a consequence, we intentionally test all planners in a zero-shot setting, i.e., they are trained on the large-scale real-world nuPlan train set without fine-tuning on the augmented benchmark scenarios. For the *interPlan* benchmark, we select 10 scenarios for each of the following types: Avoiding a construction zone (Constr.), encountering an accident (Acc.),

Scenario type	Val14 mining	interPlan (ours)
Total scenarios	51	80
Noisy objects	7	0
Unsafe initialization	13	0
Unsafe IDM behavior	8	0
Basic driving scenarios	13	0
Jaywalkers (non-react./react.)	6 / 0	0 / 10
Non-reactive merge	4	0
Interactive lane change	0	30
Nudge	0	10
Overtake	0	10
Construction site	0	10
Accident site	0	10

TABLE II: Comparison to Val14 mining

jaywalking pedestrian (Jayw.), nudging around a parked vehicle (Nudge), overtaking a parked vehicle through the oncoming lane (Overt.), and lane changes in low (LTD), medium (MTD) and high (HTD) traffic density. Among the 10 LTD, MTD, and HTD scenarios, there are 3, 3, and 4 with conservative agents, assertive agents, and mixed agents, respectively.

#### D. Metrics

The *interPlan* benchmark builds upon the established nuPlan metrics, which comprise compliance with driving direction and drivable area, speed limit, as well as safety (collisions, time-to-collision), and comfort (accelerations, jerk). Metrics for comfortability, speeding, progress, and keeping the TTC within bounds are aggregated into a weighted average. The other metrics (collisions, compliance with driving direction and area, being stationary for too long) do not contribute to the weighted average, but immediately reduce the score to 0 if performance is below a threshold. E.g., if the vehicle causes a collision, the entire scenario will be evaluated with a score of 0. We extend this driving score as follows: We measure which percentage of lane changes required to get to the goal was completed and include it in the weighted average before applying the penalties. Moreover, we add an additional multiplicative penalty based on the minimal progress needed to pass the obstacles (parked cars, construction zones, etc.). Hence, if the vehicle gets stuck in front of an obstacle, the scenario is evaluated with a score of 0%. Finally, we deactivate the driving direction compliance penalty for overtakes (Overt.) and accident sites (Acc.), where the vehicle has to go in an oncoming lane to pass the obstacle.

#### E. Comparison to Val14 Mining

A straightforward alternative to our scenario construction method is mining a large dataset for challenging scenarios [33]. We mine the Val14 test split with the state-of-the-art PDM-Closed planner and identify 51 scenarios where the score is below 60%. See Tab. II for a comparison between the mined scenarios and *interPlan*.

20 of these scenarios turn out to be simulation failures, with initialization of the ego vehicle too close to an obstacle or even off-road (13) or where the sudden appearance of an object leads to an imminent collision (7). Another eight scenarios contain unrealistically unsafe behavior of the IDM agents, such as running a red light.

In 13 of the remaining 23 scenarios, the low performance is merely caused by basic planner failures, such as getting stuck at a stop line or slightly going offroad at a tight turn. Thus, the scenarios are not challenging or of particular interest. Finally, collisions with pedestrians appear in six scenarios, and collisions during merges in four. In nuPlan simulations, pedestrians always non-reactively follow their recorded logs, which can lead to curious collisions if the controlled vehicle does not exactly follow the recording vehicle’s path. Conversely, in our jaywalker scenarios, the movement of pedestrians is triggered by the approaching EV to always allow a reaction, and the locations are selected in such a way that jaywalkers can be anticipated (e.g., at bus stops). Finally, the four collisions at merges are related to a weakness of the IDM policy used to simulate surrounding agents: They do not react to vehicles in the neighboring lane, not even if the lanes are about to merge. Our lane-change scenarios force the EV to enter the lane of a neighboring IDM agent. By entering the neighboring lane, the EV becomes visible to the IDM so that it is able to react to it. Because different agent policies are employed (cf. Sec. III-B), this creates interactive scenarios where the EV must show diverse strategies to succeed, such as nudging into the lane and waiting for the neighboring vehicle to break and create a gap. Not having found a single interactive lane change, we extend the search to all scenarios tagged with lane changes where PDM-Closed achieves a score below 90%. In all cases, the EV either gets stuck in a pickup/dropoff zone, or the bad score stems from uncomfortable driving or slow progress. Although the expert does change lanes, this is not required to follow the route. In contrast, *interPlan*’s LTD, MTD, and HTD scenarios enforce lane changes via the navigation route and explicitly reward them in the score.

## IV. RESULTS

### A. Generalization of state-of-the-art Planning Methods

Our experimental results are shown in Tab. III. We find that PDM-Closed achieves the best results with an overall score of 42%. However, in contrast to previous results on the Val14 test-split and the nuPlan leaderboard, it is far from achieving a perfect score. In particular, it achieves the best score among all state-of-the-art planners in the Construction Zone, Overtake, and Nudge scenarios. This is mainly due to its ability to sample different lateral offsets from the centerline it is following, enabling it to pass some of the obstacles in the lane. Moreover, it outperforms all other methods in the LTD, MTD, and HTD lane change scenarios. However, due to its limitation of not being able to do lane changes, this merely reflects its ability to find an optimal trajectory concerning comfort and safety since it is unable to reach the goal lane in these scenarios. In addition, almost all methods in our benchmark fail in the Acc. and Overt. scenarios, resulting in a score of 0%. Exceptions are DTPP for Acc. and PDM-Open and PDM-Closed for the Overt. scenarios. While DTPP successfully makes a lane change before the stationary vehicles in the Acc. scenario shown in Fig. 2a, PDM-Closed stops before them (see 2e). Similarly,

	Method	Val14 $\uparrow$	interPlan $\uparrow$	Constr. $\uparrow$	Acc. $\uparrow$	Jayw. $\uparrow$	Nudge $\uparrow$	Overt. $\uparrow$	LTD $\uparrow$	MTD $\uparrow$	HTD $\uparrow$	Driv. $\uparrow$	Goal $\uparrow$	No-Col. $\uparrow$
learned	Urban Driver [28]	50	4	0	0	0	0	0	0	29	0	30	30	47
	GC-PGP [30]	55	10	0	0	0	0	0	18	16	44	73	17	67
	GameFormer [31]	75	11	0	0	48	0	0	0	20	21	30	30	87
	PDM-Open [2]	54	25	13	0	56	36	8	29	29	26	60	40	43
	DTPP [32]	73	25	18	18	44	10	0	40	36	34	60	33	93
rule	IDM [21]	77	31	0	0	66	0	0	61	61	61	100	0	100
	IDM+MOBIL [55]	75	31	21	0	66	0	0	71	21	70	93	52	80
	PDM-Closed [2]	92	42	18	0	48	74	9	62	62	62	100	0	100

TABLE III: **The interPlan benchmark.** Val14 denotes the score reported on common scenarios for reference. Besides the aggregated interPlan score, we also report the scores on each scenario type. The sub-scores drivable area compliance (Driv.), reaching the goal lane (Goal), and collision avoidance (No-Col.) refer only to the lane-change scenarios (LTD, MTD, HTD).

PDM-Open and PDM-Closed pass the parked car in the easiest of the 10 overtake scenarios, i.e., the one with no oncoming traffic. For PDM-Open, this is not intentional but instead caused by a failure to follow the curved centerline since this model does not consider surrounding vehicles. Surprisingly, PDM-Closed performs worse than IDM in the scenarios with Jaywalkers even though it employs an IDM policy for centerline following. In contrast to IDM, PDM uses a future horizon of only 2.0s to evaluate potential longitudinal profiles. In two scenarios, a collision is unavoidable when the pedestrian is inside this horizon. In a further scenario, PDM-Closed avoids a collision by evading the pedestrians at high speed (see Fig. 2f instead of stopping and waiting for them to cross (see Fig. 2b). Overall, our results reveal critical shortcomings in the generalization capabilities of state-of-the-art planning methods.

### B. Interactive Lane Changes

In the lane change scenarios (i.e., LTD, MTD, HTD), we often observe trajectories that result in collisions or violations of the drivable area for the learning-based planners UrbanDriver, GameFormer, and PDM-Open (see Fig. 2h). GC-PGP often stops and gets stuck (see Fig. 2d). On the other hand, rule-based centerline planners (PDM-Closed, IDM) do not cause these infractions and thus incur smaller penalties. However, they do not reach the goal lane even once. MOBIL extends IDM with a criterion to achieve safe lane changes, depending on the surroundings. We use a simple implementation that, given that MOBIL approves the lane change, tasks the controller with reaching waypoints on the neighboring lane. It achieves the highest rate of reaching the goal lane at 52%. However, this comes at the cost of some drivable area infractions and collisions. We find that some collisions are explained by the downstream controller not being able to merge to the neighboring lane without oscillating, and others are caused by erroneously expecting vehicles approaching from behind to break (see Fig. 2g). Nonetheless, it has the highest rate of reaching the goal lane and performs well if the surrounding vehicles act conservatively (see Fig. 2c). We conclude that most state-of-the-art planners lack sophisticated lane change trajectory planning, such as synchronizing to a gap in preparation, proactively influencing other vehicles' behavior, and aborting a lane change when necessary.

### C. Rule-Based vs. Learning-Based Planning

Notably, the rule-based planners IDM, IDM+MOBIL, and PDM-Closed outperform the learning-based contestants. This is especially surprising for IDM and PDM-Closed, since they are by design limited to simple lane-following without lane changes or overtakes and thus unable to solve scenarios where a construction zone, a parked vehicle, or an accident blocks the current lane. Moreover, the lowest scores are achieved by planners that regress waypoints (GC-PGP, UrbanDriver) instead of selecting a trajectory among a set of feasible samples. This indicates that they have the lowest generalization capabilities and are susceptible even to small distributional shifts. The best learning-based planner is DTPP, which uses a learned cost function alongside trajectories generated by a sampler. Overall, none of the models demonstrates sufficient world understanding to navigate successfully through these difficult scenarios.

## V. LLM-BASED PLANNING

In the previous section, we assessed state-of-the-art planning methods under difficult conditions and revealed that there is a critical lack of generalization ability to solve rare and difficult scenarios. With the rise of foundation models and the tremendous interest surrounding them recently, exploiting their impressive world understanding and generalization capabilities in the field of autonomous driving became a vital field of research. As the scenarios we proposed in our benchmark demand a thorough understanding of traffic scenarios and complex reasoning (such as anticipating a pedestrian to unlawfully cross the street at a bus stop), approaching the planning task with foundation models appears to be a promising strategy. Prior work achieves impressive results using LLMs for open-loop trajectory planning in real-world driving scenarios. However, the current state of the art on regular nuPlan scenarios suggests that such common scenarios can be solved efficiently with simple rule-based planners, and the promise of better generalization to long-tail scenarios is yet to be put to test. In the following, we explore the capabilities of LLMs in the context of closed-loop behavior planning in such scenarios. First, we establish a simple baseline, which is an adaptation of the open-loop GPT-Driver [6] planner. Then, we describe our LLM-based behavior planning method. We evaluate both methods on the same *interPlan* benchmark we used for the state-of-the-art planning methods. By releasing our code, we hope to

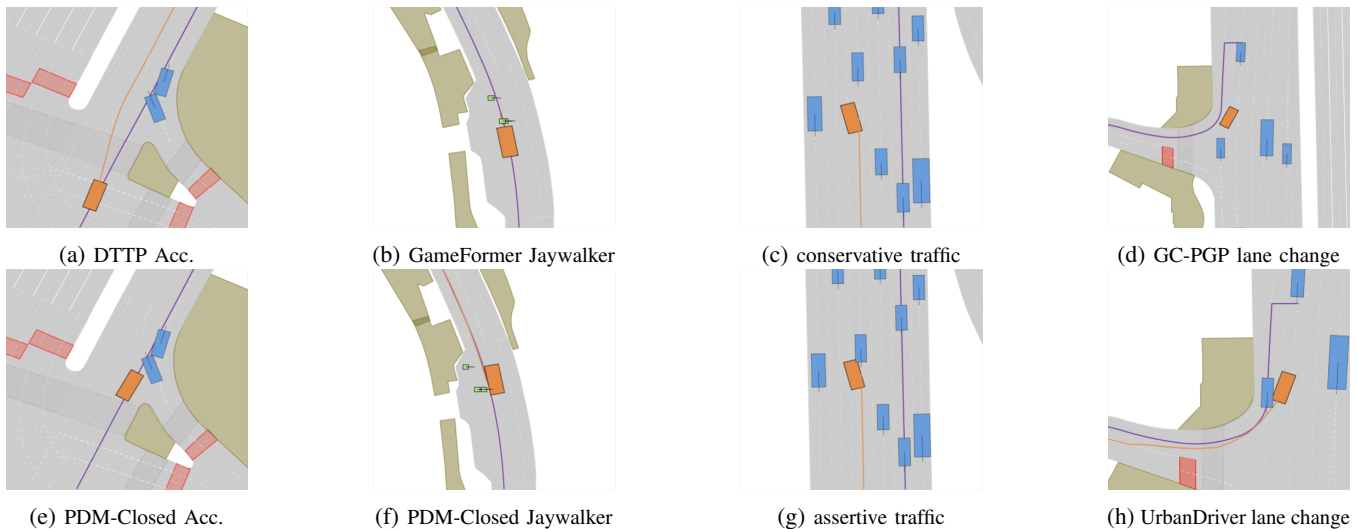


Fig. 2: **Qualitative results.** Fig. (a), (b), (c) show successfully avoiding an accident site, waiting for jaywalkers, and executing a lane change surrounded by conservative agents. Failure cases are stopping before the accident site (e), narrowly avoiding the pedestrians (f), and colliding with an assertive vehicle approaching quickly from behind (g). Fig. (d) and (h) show failures of learning-based planners, such as suddenly stopping (d) or causing a collision (h). The planned trajectory is depicted in orange, the navigation route in purple.

facilitate further research on improved LLM-based planning for automated driving.

#### A. LLM Waypoints Planning

GPTDriver [6] fine-tunes and instructs an LLM to predict a set of future waypoints. The prompt comprises general task instruction, perception context, and ego-states. Chain-of-thought reasoning is used to improve the interpretability and quality of the output. We apply the method with the same prompt after applying minor changes to adapt it to the nuPlan requirements, i.e., we change the output to an 8-second trajectory at 2 Hz and include route information into the prompt. Thus, we provide the planner with a ‘Mission Goal’ describing which lanes are on route. We test various LLMs, including Llama-7B, Llama-13B [56], and GPT-3.5. We call this baseline `LLMWaypointsPlanner` and fine-tune it on a set of 600 diverse and unaugmented nuPlan scenarios for 3 epochs. For Llama we employ QLoRa [57] using a rank of 64 and a learning rate of  $2e-4$ . For GPT, we use openAI’s fine-tuning API.

#### B. LLM Behavior Planning

In addition, we propose a two-stage hybrid planner that combines an LLM behavior planner with a rule-based motion planner. In the first stage, the traffic scenario is described in a natural language prompt. As in the case of our `LLMWaypointsPlanner` baseline, it comprises general task instruction, perception context, and ego states, as well as route information. Finally, the model is instructed to select a suitable behavior from a list of available behaviors. We define the following behaviors: follow the current lane, merge left, merge right, overtake obstacle, stop and wait. We filter the list based on the availability of neighboring lanes and the presence of obstacles in the current lane before presenting the options to the LLM. Each option is associated with a centerline and

an offset from it. For instance, ‘overtake obstacle’ refers to the current centerline and the offset required to pass the obstacle, whereas ‘merge left’ refers to the left neighboring centerline with no offset. The second stage applies a rule-based motion planner to find an optimal trajectory conditioned on the selected behavior parameters. We leverage PDM-Closed [2], the top-performing method in our benchmark, for motion planning. This algorithm samples lateral offsets and longitudinal speed profiles locally around the chosen behavior and selects a trajectory based on a cost function. The motion planner runs at 10 Hz, while the LLM behavior planner is queried at 1 Hz. We employ the LLM in a zero-shot setting, as we want to test the generalization of planning methods to unseen difficult scenarios. Moreover, labeling the regular nuPlan scenarios would almost exclusively result in ‘follow current lane’ samples.

#### C. Rule-Based Motion Planners can be enhanced with LLMs

Tab. IV compares the results of both LLM-based planning methods to the best method on our benchmark, i.e., PDM-Closed. Despite using the same prompt and fine-tuning strategy, the `LLMWaypointsPlanner` baseline cannot reproduce the strong results from other benchmarks (see [6]) and trivially outputs a longitudinal profile at a constant heading angle. Conversely, our `HybridLLMPlanner` outperforms PDM-Closed, setting a new state-of-the-art on our proposed benchmark. Surprisingly, the larger 13B Llama model performs worse than the smaller 7B version often making questionable decisions, resulting in collisions or less progress.

#### D. LLMs lack Traffic Understanding

While the `HybridLLMPlanner` outperforms all existing methods, it is merely a strong baseline and unable to solve all the difficult *interPlan* scenarios. For instance, we observe

Model	LLM	interPlan $\uparrow$	Constr. $\uparrow$	Acc. $\uparrow$	Jayw. $\uparrow$	Nudge $\uparrow$	Overt. $\uparrow$	LTD $\uparrow$	MTD $\uparrow$	HTD $\uparrow$	Driv. $\uparrow$	Goal $\uparrow$	No-Col. $\uparrow$
PDM-Closed [2]	None	42	18	0	48	74	9	62	62	62	100	0	100
HybridLLMPlanner	LLama-7B	53	27	20	48	93	28	81	48	80	100	43	87
HybridLLMPlanner	LLama-13B	48	36	10	48	82	19	71	40	77	97	43	83
HybridLLMPlanner	GPT-3.5	40	52	0	48	25	9	76	34	69	97	47	77
LLMWaypointsPlanner	LLama-7B	16	0	0	0	0	0	29	64	37	90	0	67
LLMWaypointsPlanner	LLama-13B	17	0	0	0	0	0	31	64	38	87	0	70
LLMWaypointsPlanner	GPT-3.5	22	0	0	0	0	0	64	41	69	93	0	100

TABLE IV: LLM baselines on the interPlan benchmark.

that it often toggles between different behaviors, such as overtaking and waiting for oncoming traffic, ultimately getting stuck behind the obstacle. Moreover, the decision is often inconsistent with the situation analysis, indicating a lack of understanding of traffic. Thus, we highlight fine-tuning with auxiliary tasks to enhance traffic understanding as an important direction for future research.

## VI. CONCLUSION

In this work, we presented a novel realistic benchmark that evaluates planning algorithms in highly interactive and rare long-tail scenarios. We found that existing state-of-the-art methods exhibit critical limitations in their ability to generalize to such scenarios. Moreover, we provided an LLM-based planning baseline, proposed a novel LLM-based behavior planner, and explored their capabilities. We hope that our work encourages future research to tackle this generalization problem as we see large potential in the use of multimodal foundation models to better understand traffic scenarios and plan suitable maneuvers, which will be the focus of our future work.

## VII. LIMITATIONS

While interPlan extends the widely adopted nuPlan simulator with challenging scenarios, several aspects remain unaddressed. As all scenarios are hand-designed or generated based on heuristics, the benchmark comprises only a limited number of scenarios. Nonetheless, interPlan is able to reveal critical shortcomings in current state-of-the-art methods. Further, interPlan scenarios don't provide ground truth and thus cannot be used for training. Therefore, planners can only be trained on the regular nuPlan scenarios before testing their generalization capabilities in interPlan. Finally, we believe that data-driven traffic simulation could further increase traffic diversity and make the benchmark more challenging.

## REFERENCES

- [1] H. Caesar, J. Kabzan, K. S. Tan, W. K. Fong, E. Wolff, A. Lang, L. Fletcher, O. Beijbom, and S. Omari, "nuPlan: A closed-loop ml-based planning benchmark for autonomous vehicles," *arXiv preprint arXiv:2106.11810*, 2021.
- [2] D. Dauner, M. Hallgarten, A. Geiger, and K. Chitta, "Parting with misconceptions about learning-based vehicle motion planning," in *Conference on Robot Learning*. PMLR, 2023, pp. 1268–1281.
- [3] S. Hagedorn, M. Hallgarten, M. Stoll, and A. Condurache, "Rethinking integration of prediction and planning in deep learning-based automated driving systems: a review," *arXiv preprint arXiv:2308.05731*, 2023.
- [4] N. Rhinehart, J. He, C. Packer, M. A. Wright, R. McAllister, J. E. Gonzalez, and S. Levine, "Contingencies from observations: Tractable contingency planning with learned behavior models," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 663–13 669.
- [5] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [6] J. Mao, Y. Qian, H. Zhao, and Y. Wang, "Gpt-driver: Learning to drive with gpt," *arXiv preprint arXiv:2310.01415*, 2023.
- [7] K. Chitta, A. Prakash, and A. Geiger, "Neat: Neural attention fields for end-to-end autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 793–15 803.
- [8] K. Chitta, A. Prakash, B. Jaeger, Z. Yu, K. Renz, and A. Geiger, "Transfuser: Imitation with transformer-based sensor fusion for autonomous driving," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [9] W. Zeng, W. Luo, S. Suo, A. Sadat, B. Yang, S. Casas, and R. Urtasun, "End-to-end interpretable neural motion planner," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8660–8669.
- [10] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine, "Precog: Prediction conditioned on goals in visual multi-agent settings," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2821–2830.
- [11] F. Codevilla, E. Santana, A. M. López, and A. Gaidon, "Exploring the limitations of behavior cloning for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9329–9338.
- [12] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4693–4700.
- [13] D. A. Pomerleau, "Alvin: An autonomous land vehicle in a neural network," *Advances in neural information processing systems*, vol. 1, 1988.
- [14] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann *et al.*, "Stanley: The robot that won the darpa grand challenge," *Journal of field Robotics*, vol. 23, no. 9, pp. 661–692, 2006.
- [15] A. Bacha, C. Bauman, R. Faruque, M. Fleming, C. Terwelp, C. Reinholtz, D. Hong, A. Wicks, T. Alberi, D. Anderson *et al.*, "Odin: Team victortango's entry in the darpa urban challenge," *Journal of field Robotics*, vol. 25, no. 8, pp. 467–492, 2008.
- [16] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman *et al.*, "A perception-driven autonomous urban vehicle," *Journal of Field Robotics*, vol. 25, no. 10, pp. 727–774, 2008.
- [17] C. Urmsion, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer *et al.*, "Autonomous driving in urban environments: Boss and the urban challenge," *Journal of field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.
- [18] C. Chen, A. Seff, A. Kornhauser, and J. Xiao, "Deepdriving: Learning affordance for direct perception in autonomous driving," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2722–2730.
- [19] A. Sauer, N. Savinov, and A. Geiger, "Conditional affordance learning for driving in urban environments," in *Conference on Robot Learning*. PMLR, 2018, pp. 237–252.
- [20] H. Fan, F. Zhu, C. Liu, L. Zhang, L. Zhuang, D. Li, W. Zhu, J. Hu, H. Li, and Q. Kong, "Baidu apollo em motion planner," *arXiv preprint arXiv:1807.08048*, 2018.
- [21] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical review E*, vol. 62, no. 2, p. 1805, 2000.
- [22] F. Codevilla, A. M. Lopez, V. Koltun, and A. Dosovitskiy, "On offline evaluation of vision-based driving models," in *Proceedings of the*

- European Conference on Computer Vision (ECCV)*, 2018, pp. 236–251.
- [23] D. Chen and P. Krähenbühl, “Learning from all vehicles,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17 222–17 231.
- [24] B. Jaeger, K. Chitta, and A. Geiger, “Hidden biases of end-to-end driving models,” *arXiv preprint arXiv:2306.07957*, 2023.
- [25] H. Shao, L. Wang, R. Chen, H. Li, and Y. Liu, “Safety-enhanced autonomous driving using interpretable sensor fusion transformer,” in *Conference on Robot Learning*. PMLR, 2023, pp. 726–737.
- [26] P. Wu, X. Jia, L. Chen, J. Yan, H. Li, and Y. Qiao, “Trajectory-guided control prediction for end-to-end autonomous driving: A simple yet strong baseline,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 6119–6132, 2022.
- [27] H. Shao, L. Wang, R. Chen, S. L. Waslander, H. Li, and Y. Liu, “Reasonnet: End-to-end driving with temporal and global reasoning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 13 723–13 733.
- [28] O. Scheel, L. Bergamini, M. Wolczyk, B. Osiński, and P. Ondruska, “Urban driver: Learning to drive from real-world demonstrations using policy gradients,” in *Conference on Robot Learning*. PMLR, 2022, pp. 718–728.
- [29] K. Renz, K. Chitta, O.-B. Mercea, A. Koepke, Z. Akata, and A. Geiger, “Plant: Explainable planning transformers via object-level representations,” *arXiv preprint arXiv:2210.14222*, 2022.
- [30] M. Hallgarten, M. Stoll, and A. Zell, “From prediction to planning with goal conditioned lane graph traversals,” in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2023, pp. 951–958.
- [31] Z. Huang, H. Liu, and C. 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*, 2023, pp. 3903–3913.
- [32] Z. Huang, P. Karkus, B. Ivanovic, Y. Chen, M. Pavone, and C. Lv, “Dtp: Differentiable joint conditional prediction and cost evaluation for tree policy planning in autonomous driving,” *arXiv preprint arXiv:2310.05885*, 2023.
- [33] J. Cheng, Y. Chen, X. Mei, B. Yang, B. Li, and M. Liu, “Rethinking imitation-based planner for autonomous driving,” *arXiv preprint arXiv:2309.10443*, 2023.
- [34] W. Ding, B. Chen, M. Xu, and D. Zhao, “Learning to collide: An adaptive safety-critical scenarios generating method,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2243–2250.
- [35] W. Ding, B. Chen, B. Li, K. J. Eun, and D. Zhao, “Multimodal safety-critical scenarios generation for decision-making algorithms evaluation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1551–1558, 2021.
- [36] J. Wang, A. Pun, J. Tu, S. Manivasagam, A. Sadat, S. Casas, M. Ren, and R. Urtasun, “Advsim: Generating safety-critical scenarios for self-driving vehicles,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9909–9918.
- [37] N. Hanselmann, K. Renz, K. Chitta, A. Bhattacharyya, and A. Geiger, “King: Generating safety-critical driving scenarios for robust imitation via kinematics gradients,” in *European Conference on Computer Vision*. Springer, 2022, pp. 335–352.
- [38] D. Rempe, J. Phillon, L. J. Guibas, S. Fidler, and O. Litany, “Generating useful accident-prone driving scenarios via a learned traffic prior,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17 305–17 315.
- [39] G. Bagschik, T. Menzel, and M. Maurer, “Ontology based scene creation for the development of automated vehicles,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1813–1820.
- [40] D. McDuff, Y. Song, J. Lee, V. Vineet, S. Vemprala, N. A. Gyde, H. Salman, S. Ma, K. Sohn, and A. Kapoor, “Causality: Complex simulations with agency for causal discovery and reasoning,” in *Conference on Causal Learning and Reasoning*. PMLR, 2022, pp. 559–575.
- [41] T. Menzel, G. Bagschik, and M. Maurer, “Scenarios for development, test and validation of automated vehicles,” in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1821–1827.
- [42] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang *et al.*, “Planning-oriented autonomous driving,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 17 853–17 862.
- [43] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nusscenes: A multimodal dataset for autonomous driving,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [44] J. Mao, J. Ye, Y. Qian, M. Pavone, and Y. Wang, “A language agent for autonomous driving,” *arXiv preprint arXiv:2311.10813*, 2023.
- [45] H. Sha, Y. Mu, Y. Jiang, L. Chen, C. Xu, P. Luo, S. E. Li, M. Tomizuka, W. Zhan, and M. Ding, “Languagempc: Large language models as decision makers for autonomous driving,” *arXiv preprint arXiv:2310.03026*, 2023.
- [46] L. Chen, O. Sinavski, J. Hünermann, A. Karnsund, A. J. Willmott, D. Birch, D. Maund, and J. Shotton, “Driving with llms: Fusing object-level vector modality for explainable autonomous driving,” *arXiv preprint arXiv:2310.01957*, 2023.
- [47] J. Kim, S. Moon, A. Rohrbach, T. Darrell, and J. Canny, “Advisable learning for self-driving vehicles by internalizing observation-to-action rules,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9661–9670.
- [48] B. Jin, X. Liu, Y. Zheng, P. Li, H. Zhao, T. Zhang, Y. Zheng, G. Zhou, and J. Liu, “Adapt: Action-aware driving caption transformer,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7554–7561.
- [49] Q. Sun, S. Zhang, D. Ma, J. Shi, D. Li, S. Luo, Y. Wang, N. Xu, G. Cao, and H. Zhao, “Large trajectory models are scalable motion predictors and planners,” *arXiv preprint arXiv:2310.19620*, 2023.
- [50] C. Sima, K. Renz, K. Chitta, L. Chen, H. Zhang, C. Xie, P. Luo, A. Geiger, and H. Li, “Drivelm: Driving with graph visual question answering,” *arXiv preprint arXiv:2312.14150*, 2023.
- [51] W. Wang, J. Xie, C. Hu, H. Zou, J. Fan, W. Tong, Y. Wen, S. Wu, H. Deng, Z. Li *et al.*, “Drivelm: Aligning multi-modal large language models with behavioral planning states for autonomous driving,” *arXiv preprint arXiv:2312.09245*, 2023.
- [52] E. Leurent, “An environment for autonomous driving decision-making,” <https://github.com/eleurent/highway-env>, 2018.
- [53] D. Fu, X. Li, L. Wen, M. Dou, P. Cai, B. Shi, and Y. Qiao, “Drive like a human: Rethinking autonomous driving with large language models,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 910–919.
- [54] Y. Wang, R. Jiao, C. Lang, S. S. Zhan, C. Huang, Z. Wang, Z. Yang, and Q. Zhu, “Empowering autonomous driving with large language models: A safety perspective,” *arXiv preprint arXiv:2312.00812*, 2023.
- [55] A. Kesting, M. Treiber, and D. Helbing, “General lane-changing model mobil for car-following models,” *Transportation Research Record*, vol. 1999, no. 1, pp. 86–94, 2007.
- [56] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, “Llama 2: Open foundation and fine-tuned chat models,” *arXiv preprint arXiv:2307.09288*, 2023.
- [57] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, “Qlora: Efficient finetuning of quantized llms,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.