

SANDRA: Safe Large-Language-Model-Based Decision Making for Automated Vehicles Using Reachability Analysis

Yuanfei Lin*, Sebastian Illing*, and Matthias Althoff

Abstract—Large language models have been widely applied to knowledge-driven decision-making for automated vehicles due to their strong generalization and reasoning capabilities. However, the safety of the resulting decisions cannot be ensured due to possible hallucinations and the lack of integrated vehicle dynamics. To address this issue, we propose SANDRA, the first safe large-language-model-based decision making framework for automated vehicles using reachability analysis. Our approach starts with a comprehensive description of the driving scenario to prompt large language models to generate and rank feasible driving actions. These actions are translated into temporal logic formulas that incorporate formalized traffic rules, and are subsequently integrated into reachability analysis to eliminate unsafe actions. We validate our approach in both open-loop and closed-loop driving environments using off-the-shelf and finetuned large language models, showing that it can provide provably safe and, where possible, legally compliant driving actions, even under high-density traffic conditions. To ensure transparency and facilitate future research, all code and experimental setups are publicly available at github.com/CommonRoad/SanDRA.

Index Terms—Formal methods, autonomous driving, methods for safety, large language models, decision making, formalization of traffic rules.

I. INTRODUCTION

DECISION-making is essential for automated vehicles to navigate safely and interact appropriately with other traffic participants. Traditional methods rely on finite state machines for simple scenarios [1, Sec. II-B] and neural networks to handle more complex behaviors like lane changes [2]. Recent efforts have shifted toward a knowledge-driven paradigm with the emergence of large language models (LLMs) that enable human-like and context-sensitive decision-making [3], [4]. However, their limited understanding of underlying physical models, tendency to hallucinate, and implicit consideration of traffic rules pose critical safety concerns that are unacceptable for real-world deployment. To address this, we propose using reachability analysis to verify the safety of decisions generated by LLMs before their execution (cf. Fig. 1).

A. Related Work

Below, we concisely review related works on LLM-based decision-making and safety verification approaches for auto-

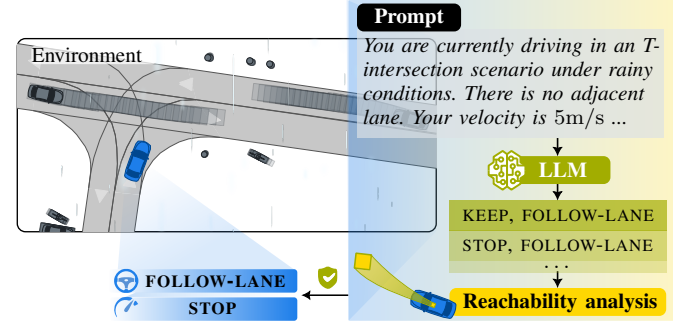


Fig. 1: An example usage of SANDRA, where the LLM is prompted to generate a ranked list of longitudinal and lateral action pairs ordered from best to worst. The action pair corresponding to stopping while staying on the current lane is verified as safe using reachability analysis and then executed.

ated vehicles.

a) *LLMs as Decision Makers*: LLMs are trained on vast and diverse internet-scale data, enabling them to understand common driving logs and reason about traffic scenarios [3], [4]. For this reason, they have been applied in automated driving to generate software programs, high-level actions, waypoints, and control signals; an overview is shown in Tab. I. The majority of these approaches employ LLMs to choose suitable driving actions given sensor inputs and/or narrative descriptions of the driving context. However, the safety of the selected actions is typically unverified or delegated to hardcoded rules [5], [6], another LLM [7]–[9], or software modules evaluating criticality measures [10]–[12]. Despite being instructed to generate safe actions, LLMs often fail to fully capture critical scenario information and may hallucinate, leading to frequent rule violations and collisions [7], [12], [13]. A similar fragility arises from their limited understanding of vehicle dynamics when directly generating waypoints to be followed, resulting in kinematically infeasible trajectories [14]–[16]. In contrast, another category of work employs LLMs to adjust key components of the software stack without altering its overall structure [17]–[19], allowing existing safety mechanisms to remain effective. However, these approaches still utilize only a limited portion of the full potential of LLMs and lack formal guarantees of correctness.

b) *Safety Verification*: Formal verification is the process of verifying whether system executions comply with a formal specification. For a review on verifying motion plans for automated vehicles, we refer to [26]. Among the various techniques, reachability analysis is the most widely

* The first two authors have contributed equally to this work.

Manuscript received Month Date, Year; revised Month Date, Year.

The authors are with the Department of Computer Engineering, Technical University of Munich, 85748 Garching, Germany. Matthias Althoff is also with the Munich Center for Machine Learning (MCML), 80538 Munich, Germany. {yuanfei.lin, sebastian.illing, althoff}@tum.de (Corresponding author: Yuanfei Lin.)

TABLE I: Related work on LLMs as decision makers and their support for safety-verified decision outputs.

Work	Output	Safety check	Formally correct
[5]–[10]	High-level actions	✓	✗
[11]–[13]	High-level actions	✗	✗
[14], [15]	High-level actions & waypoints	✗	✗
[18], [19]	High-level actions & controller parameters	✓	✗
[16]	Waypoints	✗	✗
[17], [20]	Programs	✓	✗
[21], [22]	Control signals	✗	✗
[23], [24]	Traffic descriptions	✗	✗
[25]	Traffic descriptions & waypoints	✗	✗
Ours	High-level actions & trajectories	✓	✓

used approach for online safety verification of automated vehicles [27]–[30]. The specification is typically expressed as a temporal logic formula, such as linear temporal logic (LTL) or metric temporal logic (MTL), encompassing safety specifications [31], [32], formalized traffic rules [33]–[35], and driving actions [36]–[38]. These formulas can be integrated into the software stack of automated vehicles to generate specification-compliant trajectories, monitor traffic rule compliance, and repair trajectories that violate these specifications [30], [39], [40]. Specifically, by coupling reachability analysis with model checking, one can obtain driving corridors for automated vehicles that satisfy the given formulas [41], [42]. If the reachable set defining these corridors becomes empty within the planning horizon, a specification-compliant action no longer exists. In addition, verification techniques can be integrated into deep learning-based frameworks as safety layers, such as reinforcement learning [2], [43], [44]. However, to the best of our knowledge, no such integration exists yet for an LLM-based decision-making framework.

B. Contributions

In this work, we present SANDRA, the first framework for integrating LLMs safely into decision making for automated vehicles through reachability analysis, combining the strengths of machine learning and formal methods. In particular, our contributions are:

- 1) designing a lightweight prompt that guides LLMs to generate and prioritize driving action candidates;
- 2) mapping driving actions to LTL formulas and using them to label human trajectories for LLM finetuning;
- 3) integrating reachability analysis with the LTL formulas of action candidates, along with formalized traffic rules, to verify their safety. The computed reachable sets can be used as constraints for the downstream trajectory planning module; and
- 4) evaluating safety performance in both open- and closed-loop simulations, with the closed-loop setup incorporating both most-likely and set-based predictions of surrounding traffic participants.

The remainder of this article is structured as follows: Sec. II describes the general setup and necessary preliminaries.




Fig. 2: Exemplary road network with lanes defined by lanelets. The ego vehicle is currently following lane L_{cur} , as determined by the high-level routing module to reach the goal. Alternatively, it may switch to the left-adjacent lane L_{adj_l} instead of remaining on L_{cur} .

Sec. III details the proposed framework for safe LLM-based decision making. In Sec. IV, we demonstrate the advantages of our approach. Finally, Sec. V concludes the article.

II. PRELIMINARIES

A. General Setup

The vehicle for which the decision-making is performed is referred to as the *ego vehicle*. We use the index $k \in \mathbb{N}_0$ to denote the discrete time step corresponding to the continuous time $t_k = k\Delta t$, where $\Delta t \in \mathbb{R}_+$ is a fixed time increment. The dynamics of the ego vehicle is modeled as:

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k), \quad (1)$$

where $\mathbf{x}_k \in \mathcal{X}_k$ is the state within the permissible state space $\mathcal{X}_k \in \mathbb{R}^{n_x}$, and $\mathbf{u}_k \in \mathcal{U}_k$ is the control input within the admissible input space $\mathcal{U}_k \in \mathbb{R}^{n_u}$. Without loss of generality, we set the current time step to 0 and denote the decision-making horizon by h . The solution of (1) at time step k is denoted as $\chi_k(\mathbf{x}_0, \mathbf{u}_{[0,k]})$, where $\mathbf{u}_{[0,k]}$ is the input trajectory. The solution over the time interval $[0, k]$ is denoted by $\chi_{[0,k]}(\mathbf{x}_0, \mathbf{u}_{[0,k]})$. The operator $\text{occ}(\cdot)$ returns the spatial occupancy of the entity.

As shown in Fig. 2, we describe the road network as a set of lanelets l_i , $i \in \mathbb{N}_+$, each defined by left and right boundaries represented as polylines [45]. We then define a lane L_j as a set of connected lanelet sections. A lane L_j is adjacent to another lane $L_{j'}$ if there exists a lanelet in L_j whose border fully overlaps with the border of a lanelet in $L_{j'}$. We denote the lane currently occupied by the ego vehicle as L_{cur} , with L_{adj_l} and L_{adj_r} referring to its left and right adjacent lanes, respectively. Using the centerline of L_{cur} , we can construct a curvilinear coordinate system in which vehicles are described by their longitudinal position s and lateral deviation d [46] (cf. Fig. 2). In addition, we denote the relative longitudinal distance, orientation, velocity, steering angle, and acceleration as Δs , θ , v , δ , and a , respectively.

B. Specification-Compliant Reachability Analysis

Since decision making typically focuses on short-term driving behaviors, we employ reachability analysis over a finite time horizon and formalize driving actions as formulas in LTL over finite traces (LTL_f) [47]. In this setting, temporal operators are interpreted over finite traces bounded by the




Fig. 3: SANDRA overview diagram. Our tool SANDRA takes planning and environmental information as inputs and processes them to generate a structured description of the current scenario. The description is then used to prompt the LLM to produce a ranked sequence of longitudinal and lateral action pairs, ordered from best to worst. After converting the actions and traffic rules into LTL_f formulas, we take their conjunction and apply reachability analysis to verify the safety of the resulting behavior. Once verified, the action pair is either executed directly or its corresponding reachable sets are passed to the motion planner to generate safe trajectories. If no verified actions are available, a fail-safe plan is executed.

decision-making horizon h . To ensure safety over an infinite time horizon, fail-safe planning is required [30, Sec. X].

1) *Linear Temporal Logic over Finite Traces*: Given a set of atomic propositions $\mathcal{AP}$ and formulas φ , φ_1 , and φ_2 , the syntax of LTL_f is defined according to the grammar [47, Sec. 2]:

$$\varphi ::= \sigma \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \mathbf{G}\varphi \mid \mathbf{F}\varphi, \quad (2)$$

where $\sigma \in \mathcal{AP}$ is an atomic proposition, and the symbols $\neg$ and $\wedge$ denote Boolean *negation* and *conjunction* connectives, respectively. The temporal operator $\mathbf{G}\varphi$ requires that φ will *always* hold over the finite trace, whereas $\mathbf{F}\varphi$ requires that φ *eventually* holds within the finite trace. Further logic operators can be derived from (2) (see [48, Sec. 5.1.1]), such as $\varphi_1 \vee \varphi_2 := \neg(\neg\varphi_1 \wedge \neg\varphi_2)$ (*disjunction*) and $\varphi_1 \rightarrow \varphi_2 := \neg\varphi_1 \vee \varphi_2$ (*implication*). If a finite trace $\chi_{[0,k]}(\mathbf{x}_0, \mathbf{u}_{[0,k]})$ satisfies the specification φ , we write $\chi_{[0,k]}(\mathbf{x}_0, \mathbf{u}_{[0,k]}) \models \varphi$. For formalized driving actions and traffic rules, each atomic proposition is typically associated with a predicate, defined as a function of vehicle states.

2) *Specification-Compliant Reachable Sets*: We assume that the predicted time-varying occupancies of all obstacles are given as $\mathcal{O}_k \subset \mathbb{R}^2$ at each time step k . Based on this, we define the set of forbidden states of the ego vehicle as $\mathcal{X}_k^F = \{\mathbf{x}_k \in \mathcal{X}_k \mid \text{occ}(\mathbf{x}_k) \cap \mathcal{O}_k \neq \emptyset\}$. The specification-compliant reachable set of the ego vehicle $\mathcal{R}_k^e$ at time step k is defined as the set of states reachable from an initial state in $\mathcal{X}_0$, while avoiding the forbidden states $\mathcal{X}_\tau^F$ for every step $\tau \in \{0, \dots, k\}$, via a drivable trajectory that satisfies φ . Formally, we have [42, Sec. II-C]:

$$\begin{aligned} \mathcal{R}_k^e(\varphi) := \left\{ \chi_k(\mathbf{x}_0, \mathbf{u}_{[0,k]}) \mid \left(\exists \mathbf{x}_0 \in \mathcal{X}_0, \forall \tau \in \{0, \dots, k\}, \right. \right. \\ \left. \exists \mathbf{u}_\tau \in \mathcal{U}_\tau : \chi_\tau(\mathbf{x}_0, \mathbf{u}_{[0,\tau]}) \notin \mathcal{X}_\tau^F \right) \wedge \\ \left. \chi_{[0,k]}(\mathbf{x}_0, \mathbf{u}_{[0,k]}) \models \varphi \right\}. \end{aligned} \quad (3)$$

Obtaining $\mathcal{R}_k^e(\varphi)$ is generally infeasible [49]. Therefore, we use a tight overapproximation $\mathcal{R}_k(\varphi) \supseteq \mathcal{R}_k^e(\varphi)$ that encloses all safe and specification-compliant states [42].

C. Problem Formulation

This article addresses two key challenges: designing lightweight scenario representations for prompting LLMs in the decision making of automated vehicles, and bridging their natural language outputs with safety verification techniques. For the latter, we focus on translating these outputs into LTL_f formulas φ , and verifying the legal safety of each candidate action by ensuring that its corresponding reachable set for the ego vehicle remains non-empty throughout the decision-making horizon h , i.e.,

$$\forall k \in \{0, \dots, h\} : \mathcal{R}_k(\varphi) \neq \emptyset.$$

We emphasize that the degree of safety guarantee depends on the accuracy of predictions for other traffic participants. If an inaccurate prediction leads to a safety-critical situation, we apply the safety concept described in [30, Sec. X].

III. SANDRA

This section begins with an introduction to the overall algorithm of SANDRA in Sec. III-A, followed by detailed explanations of its key components in Secs. III-B-III-E.

A. Overall Algorithm

As shown in Fig. 3, at each decision-making cycle, SANDRA takes as input both environmental and planning data, including high-level routes from navigation systems, drivable lanes, detected obstacles, and other contextual information such as traffic rules [50] and weather conditions.¹ To enable personalization, different driving styles can be incorporated through predefined driving preferences [12] or verbal commands from humans [20]. The above input is then automatically transformed into a narrative description, which is used to prompt the LLM to perform reasoning and generate a sequential ordering of $\kappa \in \mathbb{N}_+$ feasible longitudinal and lateral action pairs (cf. Tab. II and Sec. III-B). Although the actions are presented separately for clarity, they are in fact considered jointly to generalize across diverse road structures, unlike the other approaches for generating driving decisions listed in Tab. I. Note that to obtain a comprehensive representation of

¹<https://openweathermap.org/api>

the scenario, other sensor modalities can also be tokenized as vision tokens [51] and be combined with textual tokens from the narrative in a shared embedding space, enabling the use of multi-modal LLMs [12]. Afterwards, the returned action pairs are converted into LTL_f formulas (cf. Sec. III-C), which are then conjuncted with the formalized traffic rules for legal safety (cf. Sec. III-D). With this, the safety verification is carried out through the derivation of reachable sets that satisfy the formulas (cf. Sec. III-E). The first verified action pair can then be executed by the controller, or its corresponding reachable sets can be used as a constraint by various trajectory planners to generate safe trajectories [30, Sec. VIII]. When the connection to the cloud is lost or other failures occur that cause the LLM to fail to respond in time, produce hallucinations, or when no feasible verified action is available, we revert to a fail-safe trajectory to guarantee safety [30, Sec. X] (cf. Sec. III-E).

B. Prompt Design and LLM Querying

We construct a describer module that converts driving context information into descriptive text, structured into four key components, with full details provided in Fig. 4; the latter three form the user prompt. In addition to the automatically converted content, minimal manual inputs specify driving styles, hyperparameters, and regulation-related aspects. The modular design of the prompt enables easy adaptation and extension.

1) *System Prompt*: In the system prompt, we introduce the decision-making task for LLMs, which may also include personalized commands, e.g., “drive faster”. The LLM is guided to generate longitudinal and lateral action pairs ranked from best to worst by prompting it to decide in a chain-of-thought manner [52]. Specifically, we prefilter and provide the feasible actions based on current lane information, generalizing the approaches from [5], [7], [12]. The longitudinal actions are determined by whether the maneuvers are allowed in the current driving context, whereas the lateral actions are obtained by querying the lane structure at the current position of the ego vehicle (cf. Fig. 2). For instance, in Fig. 1, all longitudinal actions are contextually allowed, since the scenario is set in a rural environment; however, as the occupied lane of the ego vehicle has no adjacent lanes, the only feasible lateral action is FOLLOW-LANE.

2) *Ego Vehicle Prompt*: This section provides dynamic information about the ego vehicle, including environmental context and road network configurations. It also covers the current vehicle state, including velocity, orientation, steering angle, and acceleration. A snippet of the ego vehicle prompt is shown in Fig. 1.

3) *Traffic Rule Prompt*: To enhance rule compliance in decision-making, regulation-related information is incorporated into the prompt. The traffic rules cover both general regulations (indexed by R_G) [34, Sec. III-A] and rules for specific use cases or road types, such as interstate [34, Sec. III-B] and intersection [35]. For instance, at intersections, the LLM is informed about the right-of-way situation, relevant traffic signs, traffic lights, or the presence of traffic controllers [35]. Similarly, local regulations expressed in natural language

<System>: You are driving a car and need to make a high-level driving decision. [User command]. First, carefully observe the environment; then, reason through your decision step by step and present it in natural language, and finally, return the top [κ] advisable longitudinal-lateral action pairs ranked from best to worst. {Feasible longitudinal actions}. {Feasible lateral actions}. The past action pairs are {previous actions}.

<Ego vehicle>: You are currently driving in a {road type} scenario with {weather condition}. (There is a {left / right}-adjacent lane with the {same / opposite} direction. There are incoming lanes on the {left / right}.) Your velocity is { v }, your orientation is { θ }, your steering angle is { δ }, and your acceleration is { a }.

<Traffic rules>: Please adhere to the traffic regulations in {country} :
 [(R_G1: Safe distance to preceding vehicle rule.)]
 [(R_G2: Unnecessary braking rule.)]
 [(R_G3: Maximum speed limit rule.)] ...
 [(Traffic rules for specific use cases or types of roads)] .
 [(Local regulations)] .

<Obstacles>: Here is an overview of all relevant obstacles surrounding you: {Obstacle 1} . {Obstacle 2}

Fig. 4: LLM prompt design. Automatically generated content in the prompt is marked with {curly brackets}, while manual inputs are in [square brackets], and (round brackets) denote optional elements depending on the road structure and its regulatory features.

can be retrieved and integrated via an online retrieval agent, as demonstrated in [6]. Note that we use concretized traffic rules in the prompt, which specify, particularize, or exemplify the national traffic law texts and serve as an intermediate representation before their formalization in a formal language [30, Sec. IV-A].

Running Example: Consider the German Road Traffic Regulation (StVO), where we select the safe distance rule R_G1 as an example of the general traffic rules [34, Sec. III-A] (translated from German):

The distance to a vehicle ahead must generally be large enough that one can stop safely even if that vehicle brakes suddenly. [...] (§ 4(1) StVO).

Given the reaction time t_d of the ego vehicle, the *large enough* safe distance is defined as [34, Sec. IV-C]:

$$d_{\text{safe}}(v, v_{\text{obs}}) = \frac{v_{\text{obs}}^2}{-2|a_{\text{obs},\min}|} - \frac{v^2}{-2|a_{\min}|} + vt_d, \quad (4)$$

where v_{obs} denotes the velocity of the obstacle obs, and $a_{\min}$ and $a_{\text{obs},\min}$ represent the minimum accelerations of the ego and leading vehicles, respectively, with $a_{\text{obs},\min} < a_{\min} < 0$.

4) *Obstacle Prompt*: This prompt block describes the surrounding traffic participants in the scenario, including their IDs, types, and states. These obstacles are identified either through lane-based adjacency or using Lidar-like beams relative to the ego vehicle [53, Fig. 2]. In addition, the description is enriched with criticality measures [54], e.g., time-to-collision, to quantify the threat level posed by the obstacles

TABLE II: High-level action space for the ego vehicle and the corresponding LTL_f formula.

Action	Direction	LTL _f formula
KEEP	Longitudinal action	$\mathbf{G}(a \leq a_{\text{lim}})$
ACCELERATE		$\mathbf{G}(a > a_{\text{lim}})$
DECELERATE		$\mathbf{G}(a < -a_{\text{lim}})$
STOP		$\mathbf{FG}(\text{in_standstill}(\mathbf{x}))$
FOLLOW-LANE	Lateral action	$\mathbf{G}(\text{in_lane}(\mathbf{x}, L_{\text{cur}}))$
LEFT-LANE		$\mathbf{FG}(\text{in_lane}(\mathbf{x}, L_{\text{adj}_l}))$
RIGHT-LANE		$\mathbf{FG}(\text{in_lane}(\mathbf{x}, L_{\text{adj}_r}))$

to the ego vehicle. For instance, for dynamic vehicles at interstates, we can have the description of

Car {id}: it is driving on the {same / left adjacent / right adjacent} lane in the {same / opposite} direction and is {Δs in front of / behind} you. Its velocity is {v}, its orientation is {θ}, its steering angle is {δ}, and its acceleration is {a}. (The time-to-collision is {time-to-collision}).

The prompt is then fed into the LLM, which can be either an off-the-shelf model or a model finetuned on recorded driving datasets. To enhance reliability and facilitate parsing, we define the desired output as a schema, i.e., a sequence of longitudinal-lateral action pairs, and guide the LLM to produce actions in a format that conforms to this schema.² Moreover, for finetuning, we use recorded driving datasets and automatically label each scenario, using the prompt design in Fig. 4 as input and the action pair of a randomly selected vehicle – labeled according to the LTL_f formula shown in Tab. II – as output. When $\kappa > 1$ action pairs are required, the labeled action pair is used as the top-1 option, and the remaining positions are filled either with randomly selected feasible action pairs or according to other selection metrics. The LLM is then finetuned based on the collected input-output pairs to adapt it for our decision-making tasks in automated driving [55].

C. Action to LTL_f

Incorporating string representations of actions into reachability analysis requires their systematic conversion into LTL_f formulas. To the best of our knowledge, such a conversion of driving actions has not yet been presented in the literature. Analogous to the traffic rule formalization in [34], [35], we propose a two-step process: first, concretizing the string descriptions, and second, formalizing them into logical formulas. For instance, the action STOP can be concretized as “the ego vehicle eventually maintaining a velocity close to zero within the decision-making horizon”, where the closeness to zero is quantified by the measurement uncertainty $v_{\text{err}} \in \mathbb{R}_0$. Afterward, we extract and specify predicates in higher-order logic based on the concretization, reusing existing ones from the formalized traffic rules [34], [35] where applicable. With this, the action STOP is formalized as $\mathbf{FG}(\text{in_standstill}(\mathbf{x}))$, where the predicate $\text{in_standstill}(\mathbf{x}) \leftrightarrow -v_{\text{err}} \leq v \leq v_{\text{err}}$ indicates whether the ego vehicle is in a standstill state [35, Sec. IV-B].

As summarized in Tab. II, we formalize the most commonly used longitudinal and lateral actions based on their temporal and spatial concretizations, which can be easily extended to additional action definitions. The longitudinal actions KEEP, ACCELERATE, and DECELERATE are interpreted as confining the acceleration of the ego vehicle to a certain range and are distinguished based on a predefined acceleration threshold $a_{\text{lim}} \in \mathbb{R}_0$. In contrast, lateral actions are based on the predicate $\text{in_lane}(\mathbf{x}, L_j) \leftrightarrow \bigvee_{l_i \in L_j} \text{in_lanelet}(\mathbf{x}, l_i)$ that specifies whether the ego vehicle is occupying a particular lane L_j , determined by checking its associated lanelets (cf. Fig. 2), with $\text{in_lanelet}(\mathbf{x}, l_i) \leftrightarrow \text{occ}(\mathbf{x}) \cap \text{occ}(l_i) \neq \emptyset$ [41, Sec. V]. The action FOLLOW-LANE is interpreted as the ego vehicle always stays in the currently occupied lane L_{cur} . Instead, for the lane change actions LEFT-LANE and RIGHT-LANE, we concretize them such that the ego vehicle eventually remains in the corresponding adjacent lane L_{adj_l} or L_{adj_r} within the decision-making horizon; thus, $\mathbf{FG}$ is used as the outermost temporal operator. By translating the LLM action pairs into LTL_f formulas, we obtain a sequence of specifications $\varphi_1, \varphi_2, \dots, \varphi_\kappa$ through the conjunction of the longitudinal and lateral actions.

Running Example: The LTL_f formula corresponding to the action pair ACCELERATE and FOLLOW-LANE is formalized as

$$\varphi_{\text{acc_foll}} = \mathbf{G}(a > a_{\text{lim}}) \wedge \mathbf{G}(\text{in_lane}(\mathbf{x}, L_{\text{cur}})). \quad (5)$$

D. Formalized Traffic Rules

As the traffic rules in the prompt (cf. Sec. III-B3) may be ambiguous and are only implicitly considered by the LLM, we explicitly incorporate them into the decision-making process to ensure legal safety. To achieve this, each action specification $\varphi_1, \varphi_2, \dots, \varphi_\kappa$ is individually augmented by conjoining it with these formalized rules expressed in LTL_f [33], which are themselves conjoined. When these rules are formalized in other temporal logics, such as MTL [34], [35], they can be rewritten in LTL_f using conversion tools, e.g., Spot [56].

Running Example: The concretized safe distance rule can be formalized in LTL_f as [30, (2)]:³

$$\varphi_{\text{R-G1}} = \mathbf{G}(\text{precedes}(\mathbf{x}_{\text{obs}}, \mathbf{x}) \rightarrow \text{keeps_safe_distance}(\mathbf{x}, \mathbf{x}_{\text{obs}})), \quad (6)$$

where the meaning of the predicate $\text{precedes}(\mathbf{x}_{\text{obs}}, \mathbf{x})$ is self-explanatory, and $\text{keeps_safe_distance}(\mathbf{x}, \mathbf{x}_{\text{obs}})$ holds if and only if the relative longitudinal distance Δs between the ego vehicle and the preceding vehicle obs exceeds the threshold d_{safe} (cf. (4)):

$$\text{keeps_safe_distance}(\mathbf{x}, \mathbf{x}_{\text{obs}}) \leftrightarrow \Delta s > d_{\text{safe}}(v, v_{\text{obs}}).$$

E. Safety Verification

The online safety verification relies on dynamic models that describe the future behavior of both the ego vehicle and other traffic participants [27]. For the behavior of other

²<https://openai.com/index/introducing-structured-outputs-in-the-api/>

³For simplicity, the case without a prior cut-in is omitted. Details of the rule formalization process can be found in [30, Sec. IV].




Fig. 5: Prediction of other traffic participants and reachable set of the ego vehicle. The specifications $\varphi_{\text{acc_foll}}$ and $\varphi_{\text{R_G1}}$ are given in (5) and (6), respectively.

vehicles, we use their most likely trajectories [57] to prune traffic-rule-violating solution spaces [41], [42], while set-based predictions define the forbidden states $\mathcal{X}_k^F$ to ensure safety during decision-making (cf. (3)) [58]. Examples of both types of predictions are illustrated in Fig. 5a. These predictions serve complementary purposes: the most likely trajectories primarily capture interactions with other traffic participants, while the set-based predictions are required to ensure safety. However, other prediction strategies can be integrated into our approach as well. With this setup, we perform reachability analysis combined with model checking to compute the reachable sets $\mathcal{R}_k(\varphi)$ for $k \in [0, h]$ that satisfy the joint action and traffic-rule specification φ [41], [42]. We iterate through the specification sequence until we identify an action pair that is verified as safe, which is then returned.

To ensure safety for an infinite time horizon, we compute fail-safe solutions in parallel that can always safely transfer the ego vehicle to an invariably safe state [59] (i.e., a state that remains safe for an infinite time horizon) and thereby guarantee safety (i.e., collision-free) for all legal behaviors of other obstacles [30, Sec. X]. An example of such an invariably safe state is pulling over to the highway shoulder and coming to a stop.

Running Example: As shown in Fig. 5b, the reachable set at time step k with respect to $\varphi_{\text{acc_foll}}$ (cf. (5)) considers the set-based prediction of the preceding vehicle as the forbidden states, with colliding regions removed. By further considering $\varphi_{\text{R_G1}}$ along with the most likely prediction of the preceding vehicle (cf. (6)), reachable states that violate the safe distance rule are additionally pruned.

IV. EVALUATION

In this section, we evaluate the effectiveness of our proposed framework under the settings described in Sec. IV-A. We begin with a case study in Sec. IV-B, followed by several ablation studies in Sec. IV-C to validate the design choices of the SANDRA framework. Finally, we assess its performance in a closed-loop decision-making setting in Sec. IV-D.

A. General Settings

We evaluate SANDRA using the open-source CommonRoad framework [60]. In our implementation, we use the

TABLE III: General parameters for the evaluations.

Description	Notation and value
Threshold	$a_{\text{lim}} = 0.2\text{m/s}^2$; $v_{\text{err}} = 0.1\text{m/s}$
Reachability analysis	longitudinal acceleration: $[-6, 6]\text{m/s}^2$; lateral acceleration: $[-4, 4]\text{m/s}^2$; maximum velocity: 30m/s
Set-based prediction	maximum acceleration: 12m/s^2
Safe distance d_{safe}	$t_d = 0.4\text{s}$; $a_{\text{min}} = -6\text{m/s}^2$; $a_{\text{obs,min}} = -12\text{m/s}^2$
Dataset evaluation	
Temporal setting	$h = 25$; $\Delta t = 0.04\text{s}$
Finetuning scenarios	2000 training scenarios (highD dataset); 800 test scenarios (MONA dataset)
Highway-env simulation	
Action mapping (highway-env vs. our approach)	
LANE_LEFT	LEFT-LANE, KEEP/DECELERATE/ACCELERATE
IDLE	FOLLOW-LANE, KEEP
LANE_RIGHT	RIGHT-LANE, KEEP/DECELERATE/ACCELERATE
FASTER	FOLLOW-LANE, ACCELERATE
SLOWER	FOLLOW-LANE, DECELERATE
Hyperparameter	$\kappa = 3$; number of historical actions: 5
Frequency setting	simulation: 15Hz; policy: 5Hz
Temporal setting	scenario duration: 30 steps; $\Delta t = 0.2\text{s}$; $h = 15$ (most likely), $h = 8$ (set-based)

CommonRoad-Reach toolbox [61] for reachability analysis and the Spot tool [56] for LTL_f model checking. Additionally, we use time-to-collision as the criticality measure, computed with the CommonRoad-CriMe toolbox [54]. In closed-loop evaluations, the widely-used highway-env [62] simulation environment is employed to assess the safety of driving actions, where the most likely trajectories are predicted under a constant-velocity assumption. The set-based predictions of the other vehicles are computed based on [58]. We use GPT-4o [63] as the cloud-based LLM and Qwen3 [64] as the local LLM. The experiments are conducted on a desktop equipped with an AMD Ryzen 7 9800X 8-core processor and an NVIDIA RTX 5090 GPU. The parameters used in the numerical experiments are shown in Tab. III; all remaining parameters are as defined in the original paper.

1) *Scenarios:* For open-loop batch evaluation, we use scenarios from the highD dataset [65] to finetune LLMs and scenarios from the MONA dataset [66] to evaluate performance. Both datasets contain recordings of naturalistic vehicle trajectories, where the trajectories of other obstacles are treated as their most likely predictions, while the trajectory of the ego vehicle is labeled with high-level actions based on the LTL_f formulas in Tab. II.

2) *Traffic Rules:* We adopt the German traffic rules from [34, Sec. III-A], which are formalized in MTL and rewritten in LTL_f. In addition to the safe distance rule R_G1 introduced in the running example (cf. (6)), the following general traffic rules are considered (see Fig. 4):

- R_G2: “The ego vehicle is not allowed to brake abruptly without reason.”, which is formalized as [67, (11)]:

$$\varphi_{\text{R_G2}} = \mathbf{G}(\text{brakes_abruptly}(\mathbf{x}) \rightarrow \text{braking_justification}(\mathbf{x})), \quad (7)$$

where we define $\text{braking_justification}(\mathbf{x})$ as the case in




Fig. 6: CommonRoad scenarios⁴ in which the ego vehicle approaches a signalized intersection. (6b) and (6c) use the most likely predictions of other obstacles, while (6d) employs set-based predictions. When formalized traffic rules are incorporated in (6e), no safe action is found and a fail-safe mode is activated. The trajectory of the ego vehicle is referenced from the rear axle, and, in the 3D view, the upward axis denotes time steps.

which the ego vehicle executes a fail-safe action.

- R_G3: “The ego vehicle must not exceed the speed limit.⁵”, which is formalized as [34, Tab. I]:

$$\varphi_{R_G3} = \mathbf{G}(\text{keeps_lane_speed_limit}(\mathbf{x}) \wedge \text{keeps_fov_speed_limit}(\mathbf{x}) \wedge \text{keeps_type_speed_limit}(\mathbf{x}) \wedge \text{keeps_braking_speed_limit}(\mathbf{x})). \quad (8)$$

⁴CommonRoad-ID: DEU_Goeppingen-37_1_T-4

⁵The speed limit value is automatically obtained from the environment, e.g., from traffic signs, and integrated into the rule description as “the maximum allowed speed is `speed limit`”.

B. Case Study

Consider a German urban scenario (cf. Fig. 6a) where the ego vehicle is heading toward a signalized intersection with a velocity of 7.7m/s and acceleration of -2.1m/s^2 . Given the lane structure – with a right adjacent lane to the ego vehicle – all longitudinal actions are feasible, while the available lateral actions are limited to FOLLOW-LANE and RIGHT-LANE. The user command is set to “drive faster”, with the number of ranked action pairs limited to $\kappa = 3$, a decision-making horizon of $h = 25$, and a time increment of $\Delta t = 0.1\text{s}$. After automatically describing the environment and providing the prompt to GPT-4o, we obtain a ranked sequence of three longitudinal-lateral action pairs:

- 1) ACCELERATE, RIGHT-LANE,
- 2) KEEP, RIGHT-LANE, and
- 3) DECELERATE, FOLLOW-LANE.

We verify the safety of the first action pair by computing the reachable sets that satisfy the corresponding action formula in LTL_f , as shown in Fig. 6b. The reachable set at time step 23 is notably small and becomes empty at time step 24 due to complete overlap with the predicted occupancy of the preceding vehicle. As a result, the first action pair is verified to be unsafe. The second-best result, KEEP and RIGHT-LANE, likewise results in empty reachable sets due to collisions. Afterwards, we compute the reachable sets satisfying the last action pair DECELERATE and FOLLOW-LANE, as shown in Fig. 6c, and use a sampling-based planner [68] to find the optimal trajectory within these sets. By sampling only within the obtained reachable sets, the trajectory can be found faster than sampling the entire solution space.

Fig. 6d shows that the last action pair is provably safe within the decision-making horizon, provided the reachable sets remain non-empty under the set-based prediction of other obstacles. However, the action pair is not legally safe, as the reachable sets become empty when the action formula is conjoined with $\varphi_{R_G1} \wedge \varphi_{R_G2} \wedge \varphi_{R_G3}$. This occurs because the initial velocity of the ego vehicle is too high for it to brake smoothly and avoid a collision if the preceding vehicle brakes suddenly; that is, the safe distance rule φ_{R_G1} cannot be satisfied without violating the unnecessary braking rule φ_{R_G2} . In such a situation, a fail-safe solution is required to guarantee safety over an infinite time horizon, bringing the ego vehicle to a standstill, as shown in Fig. 6e.

C. Ablation Study

To facilitate the evaluation process, we transform both the highD and MONA datasets into CommonRoad scenarios⁶, using the setup shown in Tab. III. In each scenario, we designate a randomly selected vehicle as the ego vehicle and formulate a decision-making problem based on its initial state. It is important to note that the final state recorded in the dataset is withheld from the LLM to avoid revealing the ground-truth decision outcome. Moreover, we do not include the formalized traffic rules in the reachability analysis, as it has been shown that many human drivers violate rules in the dataset [34], [35],

⁶<https://commonroad.in.tum.de/tools/dataset-converters>




Fig. 7: Ablation studies on the effectiveness of the main components of SANDRA.

making it unsuitable as an upper-limit performance indicator. With this, we conduct ablation studies on the main components of the SANDRA framework, evaluating their effectiveness using two metrics in an open-loop setting:

- **Safe@ κ** : the proportion of scenarios where at least one of the returned κ action pairs is verified as safe by SANDRA, using the recorded trajectories of other obstacles; and
- **Top- κ accuracy**: the proportion of scenarios where the dataset label appears among the top κ action pairs.

a) User Commands: We include three different setups for driving styles: no command, the command “driving aggressively”, and the command “driving cautiously”. All styles use GPT-4o as the underlying LLM but do not include traffic rules in the prompt. The results in Fig. 7a show that, on average, a safe action can be found with $\kappa = 3$ in 99.8% of cases, where 100% is achieved when using fail-safe planning. Since our focus is on the safety of driving decisions, $\kappa = 3$

is recommended as it balances the maximum number of iterations with the effectiveness of finding safe decisions. Among all command options, driving cautiously is the most effective for identifying safe actions at both $\kappa = 1$ and $\kappa = 3$. In addition, the top- κ accuracy results indicate that the behaviors of human drivers in the MONA dataset most closely align with a cautious driving style, which is consistent with typical driving patterns in urban environments. As the off-the-shelf LLM is not pretrained on structured driving data, receives inputs in a textual format unlike how humans process decisions, and must generate explicit driving actions – which humans typically do not – the top-1 accuracy remains relatively low across all commands.

b) Traffic Rule Prompt: Here, we compare our prompt design with and without the narrative concretization of the three general rules in the traffic rule prompt (see Fig. 4), using no user command and employing GPT-4o as the underlying LLM. As shown in Fig. 7b, including the traffic rules in the prompt makes it slightly easier to identify safe actions among the first returned action pairs. With respect to accuracy against the dataset labels, the presence or absence of the traffic rule prompt does not make any noticeable difference. Given this minor effect – and the fact that not all safe actions are traffic-rule-compliant – it is necessary to incorporate formalized traffic rules into safety verification to enhance legal safety in decision-making.

c) LLMs: We conduct supervised finetuning of both GPT-4o and Qwen3-8B on the highD dataset. The input for each training example consists of a system prompt and a user prompt describing a driving scenario (cf. Fig. 4), deliberately excluding any traffic rules or explicit user commands. The output is a JSON-formatted action ranking, with the labeled action pair ranked highest and the remaining positions filled with randomly ordered feasible action pairs. GPT-4o is finetuned via the OpenAI API, while low-rank adaptation [55] is applied to the locally hosted Qwen3 model. As shown in Fig. 7c, both finetuned models achieve high safe@1 and safe@3 scores exceeding 99%, along with a significant increase in top-1 accuracy compared to their pretrained counterparts. Additionally, Qwen3 achieves comparable top- κ accuracy to GPT-4o despite having a much smaller model size. The runtime performance of the GPT-4o and Qwen3-0.8B models with $\kappa = 3$, including both their original and finetuned versions, is shown in Fig. 7d. Qwen3-0.8B achieves an average computation time of 0.67s, which is considerably shorter than that of GPT-4o at 5.56s, and is suitable for local deployment with low latency. For both models, safety verification consumes only a negligible portion of the overall decision-making process, with an average runtime of 0.033s and a standard deviation of 0.018s.

D. Closed-Loop Evaluation

In the highway-env simulation, our longitudinal and lateral action pairs are mapped to the embedded discrete meta-actions, as listed in Tab. III. Note that the longitudinal action STOP is excluded, as it is not applicable in the given highway driving context. Due to the absence of explicit emergency maneuvers in the simulation environment and the modeling

TABLE IV: Closed-loop evaluation across three settings, where values in bold denote the best performance. Note that the traveled distance is computed only for runs that were successfully completed.

Method	Setting	Success rate	Collision-free steps	Rule-compliant steps			Success steps	Traveled distance	Fail-safe rate
				φ_{R_G1}	φ_{R_G2}	φ_{R_G3}			
DiLu [7]	①	93.3%	27.7	15.7	24.4	27.7	12.3	116.5m	-
	②	53.3%	18.8	12.1	14.7	18.8	8.0	94.8m	-
	③	76.7%	23.8	13.8	19.6	23.8	9.6	102.4m	-
SANDRA + most likely pred. + w/o traffic rules in verif.	①	83.3%	28.5	23.3	21.7	28.5	15.3	114.4m	6.5%
	②	73.3%	25.6	17.5	19.5	25.6	10.9	93.7m	6.8%
	③	83.3%	25.8	16.9	20.6	25.8	11.1	98.7m	7.3%
SANDRA + most likely pred. + w/ traffic rules in verif.	①	100.0%	30.0	28.5	26.8	30.0	20.0	84.4m	25.2%
	②	96.7%	29.4	27.2	26.3	29.4	19.9	68.4m	25.0%
	③	100.0%	30.0	27.5	27.2	30.0	20.3	72.5m	26.1%
SANDRA + set-based pred. + w/o traffic rules in verif.	①	100.0%	30.0	25.0	16.2	30.0	16.5	113.6m	3.8%
	②	100.0%	30.0	21.4	14.3	30.0	14.1	88.6m	7.4%
	③	100.0%	30.0	19.6	13.7	30.0	13.1	86.4m	7.9%
SANDRA + set-based pred. + w/ traffic rules in verif.	①	100.0%	30.0	28.4	25.3	30.0	20.8	85.6m	25.9%
	②	100.0%	30.0	27.7	23.9	30.0	20.9	70.1m	26.9%
	③	100.0%	30.0	27.6	23.3	30.0	20.2	72.2m	25.8%

of other vehicles using the intelligent driver model [62], the meta-action SLOWER is adapted as the fail-safe behavior by increasing its braking intensity.⁷ Furthermore, we utilize GPT-4o as the LLM, incorporating the user command “do not change lanes too often” alongside the traffic rule prompt encoding the general rules listed in Sec. IV-A2.

1) *Evaluation Settings*: We compare our approach to the state-of-the-art framework DiLu [7], which features a decision-making agent equipped with memory, reasoning, and reflection capabilities, as well as a memory module for retaining experience. Note that we disable memory collection during evaluation to maintain a fair comparison. Following [7], we adapt six metrics to evaluate the performance of safe decision-making in a closed-loop setting:

- **Success rate**: the percentage of scenarios where the ego vehicle completes the full duration without collisions;
- **Collision-free steps**: the average collision-free time steps across all scenarios;
- **Rule-compliant steps**: the average number of time steps across all scenarios during which the ego vehicle remains compliant with the given traffic rule;
- **Success steps**: the average number of time steps across all scenarios that are both collision-free and compliant with all traffic rules;
- **Traveled distance**: the average longitudinal distance traveled from the initial state to the end of the scenario simulation; and
- **Fail-safe rate**: percentage of time steps during which a fail-safe action is executed.

The evaluation is conducted under three different environmental settings:

- ① highway with four lanes and a vehicle density⁸ of 2;
- ② highway with four lanes and a vehicle density of 3; and

- ③ highway with five lanes and a vehicle density of 3.

Each setting is evaluated using 10 different random seeds⁹ and repeated three times.

2) *Safety Evaluation*: The results in Tab. IV show that all variants of SANDRA achieve performance comparable to or better than DiLu in terms of success rate. This holds even though it does not incorporate few-shot examples from past driving experiences into the prompt. By relying solely on the zero-shot prompt, our approach not only reduces token usage but also eliminates dependency on the quality and availability of prior examples. Regarding formal safety, relying solely on most likely predictions is insufficient, particularly in critical scenarios where the actual behavior of the other obstacles may deviate significantly, e.g., under setting ②. Although formalized traffic rules are considered and fail-safe actions are frequently triggered, these interventions can still come too late. In contrast, set-based prediction allows the ego vehicle to account for all possible behaviors of other vehicles over the decision-making horizon, thereby ensuring safety throughout.

3) *Rule Compliance*: For traffic rule compliance, we observe that SANDRA substantially improves adherence to the safe distance rule φ_{R_G1} compared to DiLu, irrespective of the specific variant. The greatest improvement is achieved when formalized rules are incorporated into the reachability analysis. However, this also results in more fail-safe actions, i.e., emergency braking, which in turn reduces the traveled distance, particularly in our dense traffic settings. When combining most likely predictions with traffic rules in the verification, compliance with φ_{R_G2} is generally higher, whereas set-based predictions maintain larger longitudinal gaps between vehicles, often resulting in unnecessary braking. Furthermore, all runs comply with the speed limit rule φ_{R_G3} . Overall, integrating set-based predictions with formalized traffic rules into the safety verification yields the highest number of success steps,

⁷The control parameter for the speed response time constant is set to 0.6s, and the speed increment step is set to 15m/s.

⁸The larger the density, the smaller the vehicle longitudinal spacing.

⁹Highway-env seeds: 5838, 2421, 7294, 9650, 4176, 6382, 8765, 1348, 4213, and 2572.




Fig. 8: Action distribution of the ego vehicle in the highway-env simulation, where fail-safe actions are not taken into account.

i.e., the best collision-free and legally compliant performance.

4) *Conservatism*: The traveled distance reported in Tab. IV indicates that, without integrating rules into the verification process, the set-based approach does not induce overly conservative behavior for the ego vehicle. This is particularly true in less dense scenarios (e.g., setting ①). This is further supported by the low fail-safe rate and by the fact that SANDRA allows more lane-change options than DiLu, regardless of the prediction model (cf. Fig. 8). In contrast, incorporating traffic rules in the reachability analysis increases the rate of fail-safe actions, since most highway-env scenarios are not initialized with the ego vehicle maintaining a safe distance to preceding vehicles; consequently, the first few time steps are often required for emergency braking. This also results in additional acceleration actions (cf. Fig. 8), as the ego vehicle needs to recover the distance lost during the fail-safe period. Reducing the number of fail-safe actions would require either a relaxation of the traffic rules or the adoption of trajectory repair for rule compliance [40], both of which are, however, beyond the scope of this work.

5) *Example Scenario*: An example scenario from highway-env is shown in Fig. 9a, where DiLu fails to account for the dynamic coupling between ego vehicle braking and the distance to the preceding vehicle, resulting in a rear-end crash at simulation step 8 (cf. Fig. 9b). This issue is mitigated by SANDRA through set-based prediction of surrounding obstacles, which considers possible emergency braking maneuvers of the preceding vehicle during decision-making. Since the ego vehicle is initially too close to the preceding vehicle (cf. Fig. 9a), it first executes a series of fail-safe actions to create sufficient space before following the LLM-planned actions (cf. Figs. 9c and 9d). When formalized traffic rules are incorporated in the verification process, this phase is prolonged (cf. Fig. 9d) to satisfy additional safe-distance requirements in the rule R_G1 (cf. (6)).

To quantify the extent to which the ego vehicle complies with or violates the three traffic rules in MTL, we additionally present the robustness [30, Sec. VI] of trajectories obtained from the example scenarios in Fig. 9e. A positive robustness value indicates satisfaction of the rules, whereas a negative value indicates violation. The larger the absolute value of the robustness, the stronger the satisfaction or violation of the rules. For a systematic measure of robustness for the general traffic rules in Sec. IV-A2, we refer the readers to [69]. Without explicitly considering rules in decision-making, DiLu exhibits




Fig. 9: Illustrative comparison in highway-env¹⁰ under setting ②, in which the past occupancies of vehicles are shown for the previous five time steps.

a high degree of violations from the very beginning, ultimately leading to a collision. In contrast, safety verification using set-based predictions in SANDRA improves rule compliance of the ego vehicle from the very beginning. Moreover, the variant that incorporates formalized traffic rules achieves the highest robustness in rule compliance on average across scenarios.

V. CONCLUSION

We present the first framework for safe decision-making in automated vehicles using LLMs. By translating natural language driving actions into LTL_f formulas, we enable the coupling of set-based reachability analysis with model checking to verify the safety of each action. Unlike existing methods, it supports lightweight LLM prompting, addresses inherent safety gaps in LLMs, and provides provable safety guarantees through set-based predictions of surrounding obstacles and the verification of formalized traffic rules. Additionally, our approach bridges decision-making and downstream trajectory planning by leveraging the computed reachable sets as planning constraints. Future work includes exploring the selection of optimal decision-making horizons, employing vision-language models with surround-view images or point clouds

¹⁰Highway-env seed: 5838

as inputs to achieve a broader multimodal understanding of the environment, and deploying our approach on a research vehicle in real-world scenarios.

ACKNOWLEDGMENT

The authors gratefully acknowledge the financial support provided by the German Research Foundation (DFG) under grant AL 1185/9-1, as well as the provision of API credits through the OpenAI Research Access Program.

REFERENCES

- [1] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," vol. 1, no. 1, pp. 33–55, 2016.
- [2] H. Krasowski, X. Wang, and M. Althoff, "Safe reinforcement learning for autonomous lane changing using set-based prediction," in *Proc. of the IEEE Int. Conf. on Intell. Transp. Syst.*, 2020, pp. 1–7.
- [3] C. Cui, Y. Ma, X. Cao, W. Ye, Y. Zhou, K. Liang, J. Chen, J. Lu, Z. Yang, K.-D. Liao, *et al.*, "A survey on multimodal large language models for autonomous driving," in *Proc. of the IEEE/CVF Winter Conf. on Applications of Computer Vision*, 2024, pp. 958–979.
- [4] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, *et al.*, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.
- [5] Y. Jin, R. Yang, Z. Yi, X. Shen, H. Peng, X. Liu, J. Qin, J. Li, J. Xie, P. Gao, *et al.*, "SurrealDriver: Designing LLM-powered generative driver agent framework based on human drivers' driving-thinking data," in *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2024, pp. 966–971.
- [6] T. Cai, Y. Liu, Z. Zhou, H. Ma, S. Z. Zhao, Z. Wu, and J. Ma, "Driving with regulation: Interpretable decision-making for autonomous vehicles with retrieval-augmented reasoning via LLM," *arXiv preprint arXiv:2410.04759*, 2024.
- [7] L. Wen, D. Fu, X. Li, X. Cai, T. Ma, P. Cai, M. Dou, B. Shi, L. He, and Y. Qiao, "DiLu: A knowledge-driven approach to autonomous driving with large language models," in *Proc. of the Int. Conf. on Learning Representations*, 2024.
- [8] N. Shinn, F. Cassano, A. Gopinath, K. Narasimhan, and S. Yao, "Reflexion: Language agents with verbal reinforcement learning," *Proc. of the Advances in Neural Information Processing Syst.*, vol. 36, pp. 8634–8652, 2023.
- [9] K. Jiang, X. Cai, Z. Cui, A. Li, Y. Ren, H. Yu, H. Yang, D. Fu, L. Wen, and P. Cai, "KoMa: Knowledge-driven multi-agent framework for autonomous driving with large language models," *IEEE Trans. on Intell. Veh.*, pp. 1–15, 2024, early Access.
- [10] Z. Zhou, H. Huang, B. Li, S. Zhao, Y. Mu, and J. Wang, "SafeDrive: Knowledge-and data-driven risk-sensitive decision-making for autonomous vehicles with large language models," *arXiv preprint arXiv:2412.13238*, 2024.
- [11] H. Pang, Z. Wang, and G. Li, "Large language model guided deep reinforcement learning for decision making in autonomous driving," *arXiv preprint arXiv:2412.18511*, 2024.
- [12] G. Kou, F. Jia, W. Mao, Y. Liu, Y. Zhao, Z. Zhang, O. Yoshie, T. Wang, Y. Li, and X. Zhang, "PADriver: Towards personalized autonomous driving," *arXiv preprint arXiv:2505.05240*, 2025.
- [13] W. Wang, J. Xie, C. Hu, H. Zou, J. Fan, W. Tong, Y. Wen, S. Wu, H. Deng, Z. Li, *et al.*, "DriveMLM: Aligning multi-modal large language models with behavioral planning states for autonomous driving," *arXiv preprint arXiv:2312.09245*, 2023.
- [14] J. Mao, Y. Qian, H. Zhao, and Y. Wang, "GPT-driver: Learning to drive with GPT," *arXiv preprint arXiv:2310.01415*, 2023.
- [15] 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," in *Proc. of the European Conf. on Computer Vision*, 2024, pp. 256–274.
- [16] H. Shao, Y. Hu, L. Wang, S. L. Waslander, Y. Liu, and H. Li, "LMDrive: Closed-loop end-to-end driving with large language models," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2024, pp. 15 120–15 130.
- [17] Y. Lin, C. Li, M. Ding, M. Tomizuka, W. Zhan, and M. Althoff, "DrPlanner: Diagnosis and repair of motion planners for automated vehicles using large language models," *IEEE Robot. and Automation Letters*, vol. 9, no. 10, pp. 8218–8225, 2024.
- [18] N. Baumann, C. Hu, P. Sivasothilingam, H. Qin, L. Xie, M. Magno, and L. Benini, "Enhancing autonomous driving systems with on-board deployed large language models," in *Proc. of Robot.: Science and Syst.*, 2025.
- [19] 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*, 2025.
- [20] C. Cui, Z. Yang, Y. Zhou, Y. Ma, J. Lu, L. Li, Y. Chen, J. Panchal, and Z. Wang, "Personalized autonomous driving with large language models: Field experiments," in *Proc. of the IEEE Int. Conf. on Intell. Transp. Syst.*, 2024, pp. 20–27.
- [21] L. Chen, O. Sinavski, J. Hünemann, A. Karnsund, A. J. Willmott, D. Birch, D. Maund, and J. Shotton, "Driving with LLMs: Fusing object-level vector modality for explainable autonomous driving," in *Proc. of the IEEE Int. Conf. on Robot. and Autom.*, 2024, pp. 14 093–14 100.
- [22] Z. Xu, Y. Zhang, E. Xie, Z. Zhao, Y. Guo, K.-Y. K. Wong, Z. Li, and H. Zhao, "DriveGPT4: Interpretable end-to-end autonomous driving via large language model," *IEEE Robot. and Automation Letters*, vol. 9, no. 10, pp. 8186–8193, 2024.
- [23] C. Cui, Y. Ma, X. Cao, W. Ye, and Z. Wang, "Drive as you speak: Enabling human-like interaction with large language models in autonomous vehicles," in *Proc. of the IEEE/CVF Winter Conf. on Applications of Computer Vision*, 2024, pp. 902–909.
- [24] X. Han, Z. Wu, X. Xia, and J. Ma, "Traffic regulation-aware path planning with regulation databases and vision-language models," in *Proc. of the IEEE Int. Conf. on Robot. and Autom.*, 2025.
- [25] J.-J. Hwang, R. Xu, H. Lin, W.-C. Hung, J. Ji, K. Choi, D. Huang, T. He, P. Covington, B. Sapp, *et al.*, "EMMA: End-to-end multimodal model for autonomous driving," *arXiv preprint arXiv:2410.23262*, 2024.
- [26] N. Mehdipour, M. Althoff, R. D. Tebbens, and C. Belta, "Formal methods to comply with rules of the road in autonomous driving: State of the art and grand challenges," *Automatica*, vol. 152, no. 110692, 2023.
- [27] M. Althoff and J. M. Dolan, "Online verification of automated road vehicles using reachability analysis," *IEEE Trans. on Robot.*, vol. 30, no. 4, pp. 903–918, 2014.
- [28] C. Pek, S. Manzingier, M. Koschi, and M. Althoff, "Using online verification to prevent autonomous vehicles from causing accidents," *Nature Machine Intell.*, vol. 2, no. 9, pp. 518–528, 2020.
- [29] H. Ahn, K. Bernertorp, P. Inani, A. J. Ram, and S. Di Cairano, "Reachability-based decision-making for autonomous driving: Theory and experiments," *IEEE Trans. on Control Syst. Technology*, vol. 29, no. 5, pp. 1907–1921, 2020.
- [30] M. Althoff, S. Maierhofer, G. Würsching, Y. Lin, F. Lercher, and R. Stolz, "No more traffic tickets: A tutorial to ensure traffic-rule compliance of automated vehicles," *Proc. of the IEEE*, pp. 1–30, 2025, early Access.
- [31] N. Aréchiga, "Specifying safety of autonomous vehicles in signal temporal logic," in *Proc. of the IEEE Intell. Veh. Symp.*, 2019, pp. 58–63.
- [32] M. Hekmatnejad, S. Yaghoubi, A. Dokhanchi, H. B. Amor, A. Shrivastava, L. Karam, and G. Fainekos, "Encoding and monitoring responsibility sensitive safety rules for automated vehicles in signal temporal logic," in *Proc. of the ACM/IEEE Int. Conf. on Formal Methods and Models for Syst. Design*, 2019, pp. 1–11.
- [33] K. Esterle, L. Gressenbuch, and A. Knoll, "Formalizing traffic rules for machine interpretability," in *Proc. of the IEEE Connected and Automated Veh. Symp.*, 2020, pp. 1–7.
- [34] S. Maierhofer, A.-K. Rettinger, E. C. Mayer, and M. Althoff, "Formalization of interstate traffic rules in temporal logic," in *Proc. of the IEEE Intell. Veh. Symp.*, 2020, pp. 752–759.
- [35] S. Maierhofer, P. Moosbrugger, and M. Althoff, "Formalization of intersection traffic rules in temporal logic," in *Proc. of the IEEE Intell. Veh. Symp.*, 2022, pp. 1135–1144.
- [36] A. Corso and M. J. Kochenderfer, "Interpretable safety validation for autonomous vehicles," in *Proc. of the IEEE Int. Conf. on Intell. Transp. Syst.*, 2020, pp. 1–6.
- [37] Y. E. Sahin, R. Quirynen, and S. Di Cairano, "Autonomous vehicle decision-making and monitoring based on signal temporal logic and mixed-integer programming," in *Prof. of the American Control Conf.*, 2020, pp. 454–459.
- [38] M. Klischat and M. Althoff, "Synthesizing traffic scenarios from formal specifications for testing automated vehicles," in *Prof. of the IEEE Intell. Veh. Symp.*, 2020, pp. 2065–2072.

- [39] E. Plaku and S. Karaman, "Motion planning with temporal-logic specifications: Progress and challenges," *AI communications*, vol. 29, no. 1, pp. 151–162, 2016.
- [40] Y. Lin, Z. Xing, X. Han, and M. Althoff, "Traffic-rule-compliant trajectory repair via satisfiability modulo theories and reachability analysis," *IEEE Trans. on Robot.*, pp. 1–18, 2025, early Access.
- [41] E. Irani Liu and M. Althoff, "Specification-compliant driving corridors for motion planning of automated vehicles," *IEEE Trans. on Intell. Veh.*, vol. 8, no. 9, pp. 4180–4197, 2023.
- [42] F. Lercher and M. Althoff, "Specification-compliant reachability analysis for autonomous vehicles using on-the-fly model checking," in *Proc. of the IEEE Intell. Veh. Symp.*, 2024, pp. 1484–1491.
- [43] X. Li, Z. Serlin, G. Yang, and C. Belta, "A formal methods approach to interpretable reinforcement learning for robotic planning," *Science Robot.*, vol. 4, no. 37, p. eaay6276, 2019.
- [44] H. Krasowski and M. Althoff, "Provable traffic rule compliance in safe reinforcement learning on the open sea," *IEEE Trans. on Intell. Veh.*, vol. 9, no. 12, pp. 7617–7634, 2024.
- [45] P. Bender, J. Ziegler, and C. Stiller, "Lanelets: Efficient map representation for autonomous driving," in *Proc. of the IEEE Intell. Veh. Symp.*, 2014, pp. 420–425.
- [46] G. Würsching and M. Althoff, "Robust and efficient curvilinear coordinate transformation with guaranteed map coverage for motion planning," in *Proc. of the IEEE Intell. Veh. Symp.*, 2024, pp. 2694–2701.
- [47] G. De Giacomo, M. Y. Vardi, et al., "Linear temporal logic and linear dynamic logic on finite traces," in *Proc. of the Int. Joint Conf. on Artificial Intell.*, vol. 13, 2013, pp. 854–860.
- [48] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press, 2008.
- [49] A. Platzer and E. M. Clarke, "The image computation problem in hybrid systems model checking," in *Proc. of the Int. Workshop on Hybrid Systems: Computation and Control*. Springer, 2007, pp. 473–486.
- [50] United Nations Economic Commission for Europe, "Convention on road traffic," United Nations Conference on Road Traffic, 1968, (consolidated version of 2006). [Online]. Available: https://www.unece.org/fileadmin/DAM/trans/conventn/Conv_road_traffic_EN.pdf
- [51] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al., "Learning transferable visual models from natural language supervision," in *Proc. of the Int. Conf. on Machine Learning*, 2021, pp. 8748–8763.
- [52] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al., "Chain-of-thought prompting elicits reasoning in large language models," *Proc. of the Advances in Neural Info. Processing Syst.*, vol. 35, pp. 24 824–24 837, 2022.
- [53] X. Wang, H. Krasowski, and M. Althoff, "Commonroad-RL: A configurable reinforcement learning environment for motion planning of autonomous vehicles," in *Proc. of the IEEE Int. Conf. on Intell. Transp. Syst.*, 2021, pp. 466–472.
- [54] Y. Lin and M. Althoff, "CommonRoad-CriMe: A toolbox for criticality measures of autonomous vehicles," in *Proc. of the IEEE Intell. Veh. Symp.*, 2023, pp. 1–8.
- [55] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al., "LoRA: Low-rank adaptation of large language models," *Proc. of the Int. Conf. on Learning Representations*, vol. 1, no. 2, p. 3, 2022.
- [56] A. Duret-Lutz, A. Lewkowicz, A. Fauchille, T. Michaud, E. Renault, and L. Xu, "Spot 2.0 – a framework for LTL and ω -automata manipulation," in *Proc. of the Int. Symposium on Automated Technology for Verification and Analysis*, 2016, pp. 122–129.
- [57] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *ROBOMECH J.*, vol. 1, no. 1, pp. 1–14, 2014.
- [58] M. Koschi and M. Althoff, "Set-based prediction of traffic participants considering occlusions and traffic rules," *IEEE Trans. on Intell. Veh.*, vol. 6, no. 2, pp. 249–265, 2020.
- [59] C. Pek and M. Althoff, "Efficient computation of invariably safe states for motion planning of self-driving vehicles," in *Proc. of the IEEE/RSJ Int. Conf. on Intell. Robot. and Syst.*, 2018, pp. 3523–3530.
- [60] M. Althoff, M. Koschi, and S. Manzing, "CommonRoad: Composible benchmarks for motion planning on roads," in *Proc. of the IEEE Intell. Veh. Symp.*, 2017, pp. 719–726.
- [61] E. Irani Liu, G. Würsching, M. Klischat, and M. Althoff, "CommonRoad-Reach: A toolbox for reachability analysis of automated vehicles," in *Proc. of the IEEE Int. Intell. Transp. Systems Conf.*, 2022, pp. 2313–2320.
- [62] E. Leurent, "An environment for autonomous driving decision-making," <https://github.com/eleurent/highway-env>, 2018.
- [63] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford, et al., "GPT-4o system card," *arXiv preprint arXiv:2410.21276*, 2024.
- [64] A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, et al., "Qwen3 technical report," *arXiv preprint arXiv:2505.09388*, 2025.
- [65] R. Krajewski, J. Bock, L. Kloecker, and L. Eckstein, "The highD dataset: A drone dataset of naturalistic vehicle trajectories on German highways for validation of highly automated driving systems," in *Proc. of the IEEE Int. Conf. on Intell. Transp. Syst.*, 2018, pp. 2118–2125.
- [66] L. Gressenbuch, K. Esterle, T. Kessler, and M. Althoff, "MONA: The Munich motion dataset of natural driving," in *Proc. of the IEEE Int. Conf. on Intell. Transp. Syst.*, 2022, pp. 2093–2100.
- [67] Y. Lin and M. Althoff, "Rule-compliant trajectory repairing using satisfiability modulo theories," in *Proc. of the IEEE Intell. Veh. Symp.*, 2022, pp. 449–456.
- [68] G. Würsching and M. Althoff, "Sampling-based optimal trajectory generation for autonomous vehicles using reachable sets," in *Proc. of the IEEE Int. Intell. Transp. Syst. Conf.*, 2021, pp. 828–835.
- [69] Y. Lin, H. Li, and M. Althoff, "Model predictive robustness of signal temporal logic predicates," *IEEE Robot. and Automation Letters*, vol. 08, no. 12, pp. 8050–8057, 2023.



Yuanfei Lin received his B.Eng. degree in Automotive Engineering from Tongji University, China, in 2018, and dual M.Sc. degrees in Mechanical Engineering and Mechatronics and Robotics at the Technical University of Munich, Germany, in 2020. In 2023, he was a visiting scholar at the University of California, Berkeley, USA. He completed his Ph.D. in Computer Science at the Technical University of Munich, Germany, in 2025. His research interests include motion planning, formal methods, and large language models for automated vehicles.



Sebastian Illing is currently pursuing a M.Sc. degree in Computer Science at the Technical University of Munich, Germany. He received his B.Sc. in Informatics: Games Engineering in 2024, also from the Technical University of Munich, Germany. His research interests include large language models for autonomous driving and 3D Computer Vision.



Matthias Althoff is an associate professor in computer science at the Technical University of Munich, Germany. He received his diploma engineering degree in Mechanical Engineering in 2005, and his Ph.D. degree in Electrical Engineering in 2010, both from the Technical University of Munich, Germany. From 2010 to 2012 he was a postdoctoral researcher at Carnegie Mellon University, Pittsburgh, USA, and from 2012 to 2013 an assistant professor at Technische Universität Ilmenau, Germany. His research interests include formal verification of continuous and hybrid systems, reachability analysis, planning algorithms, nonlinear control, automated vehicles, and power systems.