

# Enhancing Autonomous Driving Systems with On-Board Deployed Large Language Models

Nicolas Baumann<sup>\*,†</sup>, Cheng Hu<sup>‡</sup>, Paviththiren Sivasothilingam<sup>\*</sup>, Haotong Qin<sup>\*</sup>, Lei Xie<sup>‡</sup>, Michele Magno<sup>\*</sup>, Luca Benini<sup>†</sup>

<sup>\*</sup>Center for Project-Based Learning, ETH Zurich, Switzerland

<sup>†</sup>Integrated Systems Laboratory, ETH Zurich, Switzerland

<sup>‡</sup>Department of Control Science and Engineering, Zhejiang University, China

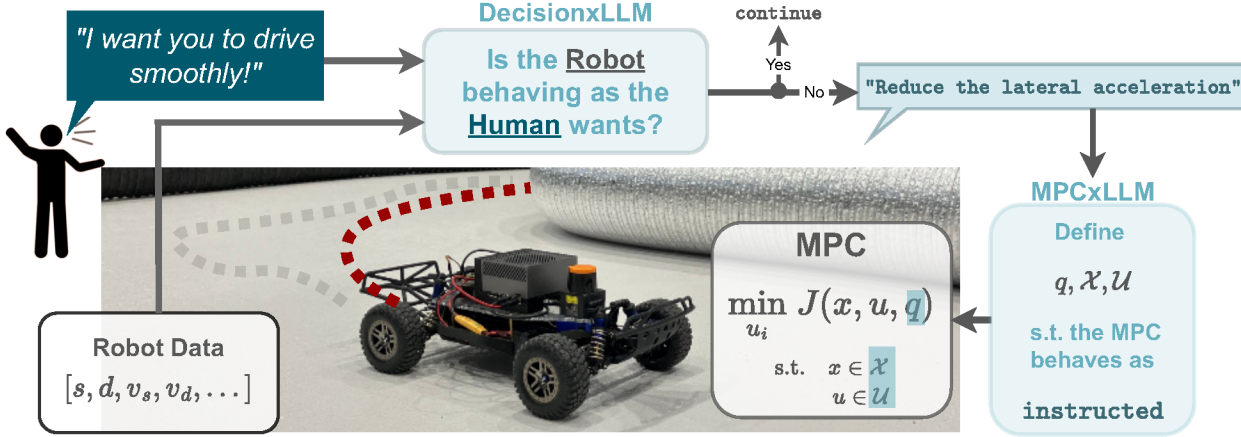


Fig. 1: Schematic overview of the proposed Large Language Model (LLM)-enhanced Autonomous Driving Systems (ADS). The LLM enables natural language-based Human Machine Interaction (HMI), with a *DecisionxLLM* stage analyzing robotic state information to ensure alignment with human preferences. If misalignment is detected, *DecisionxLLM* instructs the *MPCxLLM* stage to adjust the cost  $J(x, u, q)$  and constraint  $(\mathcal{X}, \mathcal{U})$  parameters  $x, u, q$  of a low-level Model Predictive Controller (MPC) controller, where safety and constraint satisfaction is ensured through the MPC, while task adaption and decision-making are managed by a LLM.

**Abstract**—Neural Networks (NNs) trained through supervised learning, struggle with managing edge-case scenarios common in real-world driving due to the intractability of exhaustive datasets covering all edge-cases, making knowledge-driven approaches, akin to how humans intuitively detect unexpected driving behavior, a suitable complement to data-driven methods. This work proposes a hybrid architecture combining low-level Model Predictive Controller (MPC) with locally deployed Large Language Models (LLMs) to enhance decision-making and Human Machine Interaction (HMI). The *DecisionxLLM* module evaluates robotic state information against natural language instructions to ensure adherence to desired driving behavior. The *MPCxLLM* module then adjusts MPC parameters based on LLM-generated insights, achieving control adaptability while preserving the safety and constraint guarantees of traditional MPC systems. Further, to enable efficient on-board deployment and to eliminate dependency on cloud connectivity, we shift processing to the on-board computing platform: We propose an approach that exploits Retrieval Augmented Generation (RAG), Low Rank Adaptation (LoRA) fine-tuning, and quantization. Experimental results demonstrate that these enhancements yield significant improvements in reasoning accuracy by up to 10.45%, control adaptability by as much as 52.2%, and up to 10.5× increase in computational efficiency (tokens/s), validating the proposed framework’s practicality for real-time deployment even on down-scaled robotic platforms. This work bridges high-level decision-making with low-level control adaptability, offering a synergistic framework for knowledge-driven and adaptive Autonomous Driving Systems (ADS).

## I. INTRODUCTION

In the early 2010s [5, 22], it was widely anticipated by experts that research in ADS would soon lead to the widespread adoption of fully autonomous vehicles, fundamentally transforming the automotive sector. However, progress toward full autonomy proved more challenging than initially predicted.

Before the adoption of Machine Learning (ML), many autonomous driving systems were primarily addressed using classical robotic algorithms for perception, planning, and control, adhering to the *See-Think-Act* cycle [26]. While strongly principled, these approaches exhibited significant sensitivity to heuristics and parameter tuning. With the advent of ML and especially NNs, the ability to implicitly learn heuristics and improve robustness against parameter sensitivity was demonstrated [21, 32]. Consequently, efforts were directed towards substituting individual components of the *See-Think-Act* cycle with NNs or bypassing the cycle entirely through end-to-end learning paradigms, such as Reinforcement Learning (RL) [12, 30].

Nowadays data-driven ML approaches remain the predominant methodology in ADS [7, 13]. However, despite the considerable progress achieved through these approaches, full autonomy remains elusive. ML systems inherently rely on extensive amounts of training data to generalize effectively, but

edge-case scenarios are typically underrepresented in datasets. Consequently, data-driven approaches struggle in these contexts, requiring human intervention to address situations where little or no relevant data exists [6, 16, 25].

These limitations suggest that driving is not solely a data-driven problem but partially relies on knowledge-driven reasoning [29]. For example, when encountering anomalous scenarios, data-driven solutions must have been explicitly trained on such examples [16], whereas human drivers rely on common sense and situational reasoning to handle these situations effectively [29]. Furthermore, the simulation or synthetic generation of every possible peculiar road scenario is intractable, highlighting the necessity for knowledge-driven methodologies in ADS.

In recent years, significant advancements have been made in LLMs, which represent the closest approximation to artificial knowledge systems to date [1, 2, 3]. While LLMs have demonstrated their capabilities in robotic tasks such as manipulation and scene understanding [11, 28], their adoption in the ADS domain remains relatively limited especially because existing robotic embodied Artificial Intelligence (AI) systems predominantly depend on cloud-based models, such as *GPT4* [2]. However, reliance on cloud infrastructure introduces concerns regarding latency, connection stability, security, and privacy [9, 23, 27]. As a result, local deployment of LLMs on robotic platforms emerges as a more robust and secure alternative for ADS.

Concerns persist regarding the deployment of LLMs for critical tasks, given their susceptibility to hallucinations [18]. Within this work, their integration into every facet of autonomous driving is neither suggested nor advisable. Instead, emphasis is placed on leveraging their knowledge-driven reasoning capabilities in specific scenarios where their strengths are most applicable.

Hence, a hybrid system architecture is proposed that adheres to the classical *See-Think-Act* cycle while incorporating a locally deployed, knowledge-based LLM. The architecture in this work is designed to enable LLMs to support HMI, decision-making, and dynamic control adjustments, while the evaluation has been performed on a 1:10 scaled autonomous car platform [4]. The underlying controller operates based on a low-level MPC, ensuring safety through constraint satisfaction. A *DecisionxLLM* module monitors robotic state data sampled over recent time intervals, analyzing adherence to user instructions. If discrepancies are detected, the *MPCxLLM* stage interacts with the MPC controller, adjusting cost function weights and constraints as needed. This approach facilitates seamless HMI while maintaining the safety and reliability inherent in MPC-based systems. The proposed framework allows switching between LLMs; for instance, *GPT4o* could be utilized for tasks requiring extensive cloud-based resources, while local LLMs such as *Qwen2.5-7b* can be deployed depending on connectivity constraints and other operational requirements.

To summarize, the contributions of this work are as follows:

**I Knowledge-based Decision Making:** A framework is

proposed for leveraging LLMs to interpret robotic data conditioned on human desired driving behavior, enabling decision-making based on behavioral adherence. By implementing the proposed RAG and LoRA fine-tuning techniques, decision-making accuracy is improved by up to 10.45% on local LLMs. Open-source code: [github.com/ForzaETH/LLMxRobot](https://github.com/ForzaETH/LLMxRobot).

**II Human-Machine Interaction:** Adherence is identified in relation to human prompts, enabling natural language-based HMI through dynamic adjustments of cost and constraint parameters in the low-level MPC controller. This approach enables an increase in control adaptability by up to 52.2%.

**III Embodied AI on the Edge:** The proposed framework avoids reliance on cloud services by deploying LLMs locally on embedded platforms, such as the *Jetson Orin AGX*, ensuring reliability, enhanced privacy, and improved security on computationally constrained devices. By employing Q5\_k\_m quantization and the `llama.cpp` inference engine, up to a 10.5-fold increase in computational efficiency (tokens/s) can be achieved on embedded On-Board Computers (OBCs).

## II. RELATED WORK

This section reviews relevant work on the use of LLMs for robotic control (Section II-A) and decision-making (Section II-B), concluding with a contextual summary (Section II-C).

### A. LLMs and Robot Control

Recent studies highlight that direct control of robotic actuators by LLMs is unsuitable due to their lack of training data on actuator-level commands, incompatibility with real-time control frequencies, and limited suitability for classical control paradigms [31]. Instead, approaches such as [20, 24, 31] emphasize the role of reward functions as intermediaries, enabling interaction between LLMs and low-level controllers like MPC. Here, the LLM interacts with the cost function and system constraints, allowing for interpretable and flexible control adaptation.

Building on this, Ismail et al. [20] propose using an LLM to generate objective functions and constraints for manipulation tasks based on human prompts. This architecture combines the adaptability of LLMs with the safety and constraint guarantees of classical MPC controllers.

Similarly, Ma et al. [24] demonstrates the significance of reward functions within LLM-RL interactions. Instead of focusing on cost functions and constraints as in [20, 31], the LLM iteratively designs reward functions and domain randomization strategies during the training of an RL locomotion policy, this has shown to yield great flexibility in designing effective *Sim-to-Real* policies [24].

Collectively, these works demonstrate that LLMs are more effective at interpreting and adjusting reward functions and constraints rather than acting as low-level controllers. This approach combines the interpretability and flexibility of LLMs

with the safety guarantees of traditional controllers. However, existing solutions depend on cloud-based *GPT4* models, introducing concerns regarding privacy, latency, and internet reliability [9, 23, 27]. Additionally, these approaches limit the LLMs role to adapting the behavior of the low-level controller rather than actively participating in any decision-making processes.

### B. Robotic Decision Making with LLMs

Wen et al. [29] propose the *DiLU* framework, which utilizes an LLM for decision-making in autonomous driving. *DiLU* incorporates three core components: a reasoning module that interprets the current driving scenario and generates high-level decisions, a reflection module that evaluates and refines these decisions based on previous outcomes, and a memory module that accumulates experiences from previous interactions. This architecture enables *DiLU* to integrate reasoning with iterative refinement, allowing it to handle driving scenarios effectively.

While *DiLU* demonstrates reasoning and adaptability capabilities, it operates within a discrete action space, which limits its applicability to continuous control tasks common in real-world robotics. Furthermore, it relies on cloud-based *GPT4* models, with the aforementioned cloud-reliance downsides [9, 23, 27]. Lastly, the framework has only been validated in simulation within a simple highway lane-switching RL environment. It remains to be seen whether this approach can generalize effectively to physical robotic systems operating in real-world conditions.

### C. Summary of LLMs and Robotics

As summarized in Table I, existing works like Yu et al. [31] and Ismail et al. [20] focus on manipulation and locomotion tasks, relying on classical low-level controllers and lacking robotic reasoning. *DiLU* explores LLM-based decision-making for autonomous driving but is limited to a discrete action space in a simulated highway lane-switching environment, restricting its applicability from real-world continuous control tasks. Additionally, these approaches depend on cloud-based *GPT4*, posing challenges with latency, privacy, and connectivity [9, 23, 27].

In contrast, our proposed approach emphasizes local edge deployment of the LLM, and decision-making in a continuous action space on a physical robotic car, directly grounded on robotic sensor data.

| Related Work       | Task         | Decision | Controls | LLM     |
|--------------------|--------------|----------|----------|---------|
| Ismail et al. [20] | Manipulation | ✗        | ✓        | GPT4†   |
| Yu et al. [31]     | Locomotion   | ✗        | ✓        | GPT4†   |
| Ma et al. [24]     | Locomotion   | ✗        | ✓        | GPT4†   |
| Wen et al. [29]    | Driving      | ✓        | ✗        | GPT4†   |
| <b>ours</b>        | Driving      | ✓        | ✓        | Qwen-7b |

TABLE I: Comparison of related work integrating LLMs with robotics, highlighting differences in tasks, decision-making, control influence, and LLM models used. † denotes cloud dependency.

## III. METHODOLOGY

This section describes the hardware and robotic platform used for onboard LLM processing in Section III-A. Section III-B describes how the decision-making module *DecisionxLLM* utilizes robotic data for reasoning and integrates with the overall stack. Finally, Section III-C explains the interaction between the LLM and the controller within *MPCxLLM*, including the MPC formulation.

### A. Robotic System and Computational Hardware

In this work, the robotic platform of Figure 2, along with its autonomy algorithms as detailed in the open-source *FITENTH* autonomy stack [4], is employed. The autonomy stack has been extended with the integration of a kinematic MPC controller, further described in Section III-C1. Consequently, positional state information is represented in the *Frenet* co-ordinate frame, where the *s* coordinate indicates longitudinal progress along a global trajectory, and the *d* coordinate denotes lateral deviation from that trajectory, following the conventions in [4].

A key hardware component is the *Jetson Orin AGX* serving as the OBC. This OBC incorporates a 2048-core NVIDIA Ampere architecture Graphics Processing Unit (GPU) with 64 Tensor Cores, delivering 275 TOPS, and is utilized for LLM inference. Additionally, the Central Processing Unit (CPU), a 12-core Arm Cortex-A78AE, is responsible for running the autonomy stack, including the MPC. The OBC is equipped with 64GB of shared Random Access Memory (RAM), providing ample memory for computational tasks.

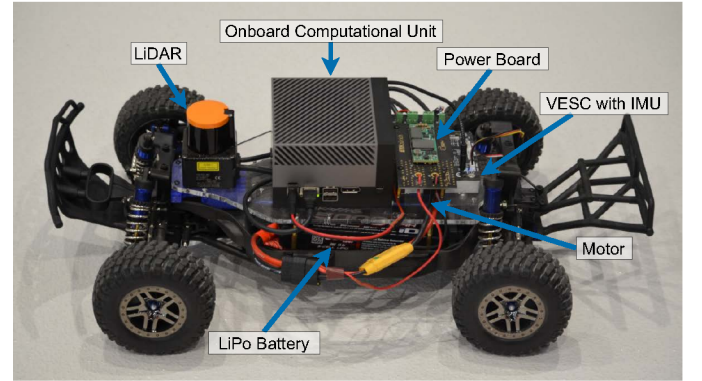


Fig. 2: The 1:10 scaled robotic platform, utilizing the *Jetson Orin AGX* as the OBC for executing computations related to the locally deployed LLM and the ADS autonomy stack.

In this work, *GPT4o* is utilized via the *OpenAI* API [19] as a cloud-based upper baseline, while two locally deployable models, *Phi-3-mini-4k-instruct*<sup>1</sup> (referred to as *Phi3-mini*), a 3B parameter LLM, and *Qwen2.5-7b-Instruct*<sup>2</sup>, a 7B parameter LLM, are sourced directly from *HuggingFace*. These models were intentionally selected to validate the robustness and

<sup>1</sup><https://huggingface.co/unsloth/Phi-3-mini-4k-instruct>

<sup>2</sup><https://huggingface.co/unsloth/Qwen2.5-7B-Instruct>

adaptability of the proposed framework, demonstrating its effectiveness across diverse architectures and parameter scales.

### B. DecisionxLLM — Decision Making with Robotic Data

The decision-making mechanism of *DecisionxLLM*, shown in Figure 3, enables dynamic evaluation of robotic data against desired driving behavior expressed through natural language prompts. The system processes a brief temporal snapshot (e.g., 2 seconds) of the robot’s state, including position, velocity, and proximity to environmental boundaries.

Given this robotic data and a human-defined driving behavior prompt, the LLM assesses whether the robot adheres to the specified behavior. For example, if a passenger in an autonomous taxi requests a smoother ride due to discomfort, the LLM could infer a reduction in lateral acceleration and prioritize gentle maneuvers, enhancing overall user experience through natural language HMI.

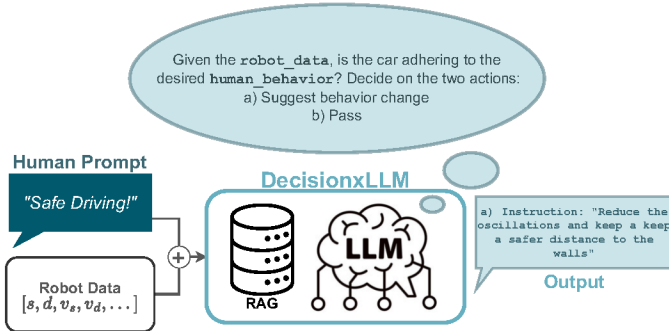


Fig. 3: Diagram illustrating the decision-making process of the LLM, where it evaluates robot data conditioned on a desired driving behavior based on a human prompt. The LLM determines whether the behavior aligns with expectations or suggests necessary adjustments.

The RAG module within the *DecisionxLLM* architecture, inspired by [29], optionally enhances the system by enabling memory modules to enrich the prompt with relevant context. This includes safety-critical and robot-specific information, such as nominal operating ranges (e.g., speed limits, distance thresholds). This capability allows human users to define custom safety and preference profiles while significantly improving the LLM’s decision-making abilities by augmenting robot-specific constraints into the prompt. This augmentation is particularly valuable on computationally constrained embedded OBCs, as the performance improvement comes without having to employ larger compute heavy LLMs. An example of the decision RAG is provided in Appendix C, Listing 3.

If the *DecisionxLLM* determines that the robot behavior aligns with the desired behavior, no further action is taken. However, if deviations are detected, the module generates a concise adjustment instruction in natural language, specifying how the behavior should be corrected. This instruction seamlessly integrates with the *MPCxLLM* module, where it serves as input to dynamically adjust relevant parameters, ensuring alignment with the desired behavior.

### C. MPCxLLM — Controller Interaction

The interaction between the *MPCxLLM* controller and the LLM, depicted in Figure 4, follows principles from [20, 31]. This integration enables an LLM, aware of the MPC formulation and its adjustable parameters, to interface with the low-level controller. As a result, task flexibility is achieved through natural language-based HMI, while the MPC ensures safety and constraint satisfaction at the low level.

Importantly, the inference latency of the LLM is decoupled from the control frequency of the MPC. Operating at a higher abstraction level, the LLM intermittently adjusts the MPC parameters without interfering with the MPC fixed-frequency control loop. This ensures control stability and safety, while the LLM focuses on task-level adaptations.

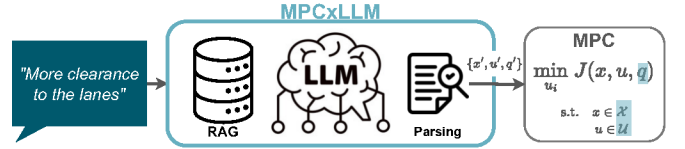


Fig. 4: Illustration of the *MPCxLLM* architecture: a natural language instruction serves as input, optionally enhanced by a RAG, processed by the LLM, and finally parsed to extract relevant parameters, which are then transmitted to the MPC via Robot Operating System (ROS) dynamic reconfigure.

Similar to the RAG module described in Section III-B, an optional RAG can be integrated here. Given the computational constraints of the OBC and the limited size of the locally deployed LLM, the RAG can significantly improve performance by enriching prompts with context-specific hints tailored for the MPC. An example of such a MPC RAG memory is provided in Appendix B, Listing 2. It is worth noting that the RAG module is optional and can be disabled, as evaluated in Table III.

The MPC is based on a kinematic model, with its cost function primarily designed to ensure accurate tracking of the given trajectory and velocity.

#### 1) Kinematic Model:

$$\begin{aligned} \dot{s} &= \frac{v \cos(\Delta\phi)}{1 - \kappa_r(s)n}, \\ \dot{n} &= v \sin(\Delta\phi), \\ \Delta\dot{\phi} &= \frac{v \tan(\delta)}{L} - \kappa_r(s) \cdot \frac{v \cos(\Delta\phi)}{1 - \kappa_r(s)n}. \end{aligned} \quad (1)$$

where  $\kappa_r$  indicates the curvature of the reference trajectory.  $L$  is the distance from the front axle to the rear axle.  $s$  denotes the distance traveled along the reference trajectory, and  $n$  is the lateral deviation from this trajectory.  $\Delta\phi$  represents the heading angle error.  $\delta$  denotes the steering angle and  $v$  is the velocity. The incremental model is utilized to smooth the control inputs. Consequently, the state of the model is comprised of the following five variables:

$$x = [s \quad n \quad \Delta\phi \quad \delta \quad v]^T \quad (2)$$

The input variables include the steering angle difference  $\Delta\delta$  and the longitudinal acceleration  $a$ :

$$u = [\Delta\delta \quad a]^T \quad (3)$$

2) *MPC Formulation*:

$$\begin{aligned} \min_{u_{k+i|k}} J(x, u, q) = & \sum_{i=0}^{N-1} q_n \cdot n_{k+i|k}^2 + q_v \cdot (v_{k+i|k} - v_{ref})^2 \\ & + q_\alpha \cdot \Delta\phi_{k+i|k}^2 + \|\Delta u_{k+i|k}\|_{q_R} + q_n \cdot n_{k+N|k}^2 \\ & + q_v \cdot (v_{k+N|k} - v_{ref})^2 + q_\alpha \cdot \Delta\phi_{k+N|k}^2 \end{aligned} \quad (4)$$

s.t.

$$\begin{aligned} x_{k+1+i|k} &= f(x_{k+i|k}, u_{k+i|k}) \\ x_{k+i|k} &\in \mathcal{X}, u_{k+i|k} \in \mathcal{U} \\ \forall i &= 0, 1, \dots, N-1 \end{aligned} \quad (5)$$

where  $N$  is the prediction horizon.  $q_n, q_\alpha, q_v$  are state weight parameters.  $q_R$  represents the weight matrix for the control inputs. Its diagonal elements respectively indicate the penalties on the difference of the steering angle and the longitudinal acceleration, denoted by  $q_{\Delta\delta}$  and  $q_{ac}$ .  $\mathcal{X}$  and  $\mathcal{U}$  represent the sets of constraints for the states and inputs, respectively. Among them, the velocity  $v$  and the steering angle  $\delta$  are each limited in magnitude. The vehicle's lateral error is constrained within the road boundaries, and an online tuning parameter  $\epsilon$  is introduced as a boundary inflation factor to ensure the vehicle's driving safety. The steering angle difference and the longitudinal acceleration are also limited in their respective ranges. To simplify the expression of the optimization problem, the lateral acceleration constraint  $a_c(x, u)$  based on the vehicle's kinematic model is not explicitly included in the formulation.

Overall, the weight parameters  $q$ , the constraints on states  $\mathcal{X}$  and inputs  $\mathcal{U}$ , and the boundary inflation  $\epsilon$ , can all be treated as adjustable parameters for the LLM to tune so that the vehicle can exhibit the desired behavior.

#### D. LoRA Finetuning of the LLMs

While *GPT4o* together with robot or task-specific RAG have shown to be zero-shot capable, small LLMs which would be locally deployable perform significantly worse than the large cloud-bound *GPT4o* and thus have to be LoRA finetuned with synthetic data derived from *GPT4o*. Although the correctness of the synthetic data is not formally guaranteed, observed response quality was deemed high enough for LoRA fine-tuning on smaller LLMs. Hence, finetuning via Parameter Efficient Fine-Tuning (PEFT) methods, particularly LoRA [17], is employed. LoRA can reduce the number of trainable parameters by thousands of times and the GPU memory requirement by up to threefold, significantly simplifying the finetuning process [17]. For LoRA-based PEFT, the `unsloth` framework is utilized [10]. Training utilizes synthetic datasets generated by zero-shot prompting *GPT-4o* [19] across each domain:

**I Synthetic Data for DecisionxLLM:** *GPT4o* generates state summaries using randomized parameters derived

from the decision-RAG, detailed in Listing 3. By leveraging randomized nominal operation ranges and thresholds (e.g., speed ranges, critical distances to boundaries, etc.), *GPT4o* synthesizes robot state representations and identifies deviations from expected behavior according to the RAG. These randomized parameters encourage the local LLM to focus more emphasis on the decision-RAG hints. This enhances the robot's ability to ground its decision-making process in the provided hints, facilitating customization of safety and preference profiles by the user. More information in Appendix A.

**II Synthetic Data for MPCxLLM:** *GPT4o* generates data using randomized parameters in a base MPC formulation (Listing 1). The model adapts these parameters to enforce specific driving behaviors while providing brief justifications. This dataset enables LoRA to learn interactions between MPC elements, with training data designed to be out-of-distribution due to the parameter randomization and an altered MPC formulation being employed during inference. More information in Appendix A.

Post-training quantization is applied to enhance inference speed and efficiency by converting the finetuned LLMs into the Q5\_k\_m GGUF format. This compression reduces both memory usage and computational demands, enabling notably faster inference when utilizing the `llama.cpp` inference engine [8, 14] on the resource-constrained OBC, as demonstrated in Table IV.

This combined approach ensures efficient training, optimized inference, and enhanced onboard decision-making capabilities for both *MPCxLLM* and *DecisionxLLM* modules.

## IV. RESULTS

This section presents the experimental setup, along with qualitative and quantitative results of the *DecisionxLLM* modules's ability to assess whether the vehicle adheres to instructed driving behaviors based on robotic state information in Section IV-A. Next, the controller adaptability through *MPCxLLM* is evaluated in Section IV-B. Finally, the computational performance of the edge-deployed LLM is evaluated on the resource-constrained OBC in Section IV-D.

### A. Reasoning on Robotic Observations

Bubble 1 illustrates a qualitative example of the *DecisionxLLM* output. The `human_prompt` instructs the system to "*Drive normally*". The LLM evaluates the robot state against this directive, leveraging optional `rag_hints` that provide contextual information, such as for example nominal vehicle speed and safe distances from track boundaries. Detailed `rag_hints` are provided in Listing 3.

In this scenario, the `rag_hints` define "normal" driving as minimal lateral oscillations, safe clearance from boundaries, and sufficient speed. The LLM identifies anomalies within the robot state: significant lateral oscillations, proximity to the right wall at 0.1m, and a low speed of  $1\text{ms}^{-1}$ . Based on these observations, the module concludes that the state does not conform to the human directive of "*driving normally*".



Subsequently, the LLM generates a natural language instruction for the *MPCxLLM* module to correct these issues, enhancing interpretability and ensuring alignment with the desired driving behavior.

#### Bubble 1a: DecisionxLLM – Human Instruction

Drive normally!

#### Bubble 1b: DecisionxLLM – Prompt

The human wants to: human\_prompt.

The data has been sampled for 2.0 seconds in 4 samples.

- s-coordinate: 19, 20, 21, 22
- d-coordinate: -0.6, 0.6, -0.65, 0.61
- s-speed: 1.0, 1.0, 0.9, 1.1
- d-speed: 1.2, -1.2, 1.21, -1.2
- distance to left wall: 1, 1, 1, 1
- distance to right wall: 0.1, 0.1, 0.1, 0.1
- crashed: False

Here are some guides to help you reason: rag\_hints.

Check if the car is doing what the human wants. Choose one of the following actions to command the car:

- a) Continue behavior
- b) Change behavior: instruction

#### Bubble 1c: DecisionxLLM – Response

The car is driving with s-speeds close to 1.0 m/s, which is **below the normal driving speed** range of 5-7 m/s. The d-coordinate oscillates between positive and negative values, indicating **oscillation around the racing line**. The d-speed is **high**, exceeding 0.5 m/s, which is not normal. The car is not crashed, facing the wall, or reversing. The distances to the **walls are very close**, which is unsafe. The car is **not adhering to the human's** desired driving style of "drive normally."

Action: **b) Change behavior**

Instruction: The car should **increase its s-speed** to a normal range of 5-7 m/s, **reduce the oscillation** in d-coordinate, and **move closer to the centerline to increase safety**.

On a quantitative note, the *DecisionxLLM* module is evaluated based on its ability to determine whether a robotic state adheres to the human's desired driving behavior. A dataset containing 200 sets of robot states was curated and labeled according to 8 example human prompts, allowing for quantitative evaluation. Example prompts include:

- I "Drive faster than 3 m/s!"
- II "Normal driving on the racing line."
- III "Reverse the car!"

These prompts were selected for their suitability to programmatically classify the robot states based on the predefined driving characteristics listed above. The LLM performs a binary classification over 1600 robot state samples, to determine adherence conditioned on a desired driving behavior. More information on the definition of decision-making accuracy in Appendix C.

| LLM       | Params | Quant | LoRA | RAG | Accuracy [%]↑ |
|-----------|--------|-------|------|-----|---------------|
| GPT4o     | ?      | ?     | ✗    | ✗   | <b>81.68</b>  |
| Phi3-mini | 3.8B   | FP16  | ✗    | ✗   | 72.15         |
| Qwen2.5   | 7B     | FP16  | ✗    | ✗   | 77.75         |
| GPT4o     | ?      | ?     | ✗    | ✓   | <b>92.48</b>  |
| Phi3-mini | 3.8B   | FP16  | ✗    | ✓   | 78.69         |
| Qwen2.5   | 7B     | FP16  | ✗    | ✓   | 82.47         |
| Phi3-mini | 3.8B   | FP16  | ✓    | ✓   | 82.60         |
| Qwen2.5   | 7B     | FP16  | ✓    | ✓   | <b>87.32</b>  |
| Phi3-mini | 3.8B   | Q5    | ✓    | ✓   | 84.95         |
| Qwen2.5   | 7B     | Q5    | ✓    | ✓   | <b>87.02</b>  |

TABLE II: Decision-making accuracy across various LLMs, illustrating the impact of quantization, LoRA fine-tuning, and RAG on the performance of the *DecisionxLLM* module in evaluating whether a robotic state aligns with desired driving behavior. Higher accuracy values indicate improved adherence to the specified driving characteristics. ? indicates that this is proprietary knowledge, not known to the public.

Table II demonstrates the performance of the *DecisionxLLM* module in determining whether a robotic state adheres to the desired human driving behavior. All performance improvements are stated in absolute percentage points. The results indicate that the inclusion of RAG consistently enhances model performance across all tested LLMs, with an average improvement of 7.35% (*GPT4o/Phi3/Qwen2.5* – +10.8/+6.54/+4.72%). Fine-tuning locally deployable LLMs via *GPT4o* distillation, as detailed in Section III-D, further improves model accuracy by an average of 4.38% (*Phi3/Qwen2.5* – +3.91/+4.85%). Lastly, quantization, which is essential for deployment on computationally constrained OBCs, does not substantially degrade decision-making performance, maintaining accuracy within a margin of 1.02% relative to full-precision models (*Phi3/Qwen2.5* – +2.35/-0.3%).

In summary, the proposed system effectively reasons over robotic state information conditioned on human driving instructions, benefiting from RAG, LoRA fine-tuning (*Phi3/Qwen2.5* – +10.45/+9.57%), and demonstrating resilience to the effects of model quantization for local deployment. Further, comparing the two local LLMs, *Qwen2.5-7b* consistently outperforms *Phi3-mini-3.8b* across all settings, making it the preferred choice. That said, *GPT4o* achieves the highest performance overall, and the framework's flexibility enables its use in scenarios where dependency on cloud connectivity is warranted.

While Table II presents the average decision-making accuracy on the custom dataset, Table V provides a detailed breakdown across sub-categories and additional highly con-

temporary LLMs, including a *DeepSeek R1* distilled *Qwen2.5-7b* model [15].

### B. MPCxLLM — Control Adaptability

The *MPCxLLM* module is evaluated, in a closed-loop simulation environment of the open-source *FITENTH* autonomy stack [4], where the MPC handles vehicle control. In this context, simulation is preferred over physical testing due to time efficiency and access to ground-truth data.

The LLM governs interaction scenarios to enable the computation of quantifiable Root Mean Square Error (RMSE) metrics. We assess changes in closed-loop behavior relative to the default MPC formulation outlined in Section III-C1, using the following criteria:

- I **Centerline:** RMSE of the distance from the centerline, denoted as  $E_C[m]$ . Instruction prompt: "Drive as far away from the walls as possible!"
- II **Reference Velocity:** RMSE deviation from a human-provided reference velocity, denoted as  $E_V[ms^{-1}]$ . Instruction prompt: "Follow the reference velocity of 1.25 m/s as closely as possible!"
- III **Reversing:** RMSE deviation from a negative reference velocity specified by the human, denoted as  $E_R[ms^{-1}]$ . Instruction prompt: "Drive the track in reverse at -1 m/s!"
- IV **Smooth Driving:** RMSE of acceleration deviating from  $0 ms^{-2}$ , denoted as  $E_S[ms^{-2}]$ . Instruction prompt: "Reduce jerkiness!"

#### Bubble 2a: MPCxLLM – Instruction

Reverse the car!

#### Bubble 2b: MPCxLLM – Prompt

Adapt the tuneable parameters of the MPC so that the car achieves the following: instruction.  
This is the MPC formulation: base\_memory.  
Make use of these memories: rag\_memories.  
Return format:  
new\_mpc\_params = {param1: new\_value1,  
param2: new\_value2, ...}

#### Bubble 2c: MPCxLLM – Response

```
new_mpc_params = {qv: 0.1, 'qn': 40, qalpha: 50, ddelta_min: -5, ddelta_max: 0, dv_min: -50, dv_max: -1, v_min: -1, v_max: -1, boundary_inflation: 0.1}
```

**Explanation:** The new parameters are tuned to achieve the goal of driving the track in reverse at -1.0 m/s. The minimum and maximum velocities ( $v_{min}$  and  $v_{max}$ ) are set to -1 and -1, respectively, to ensure the velocity is at -1.0 m/s.

It is important to note that while the evaluation scenarios presented here are specifically designed to ensure the computation of quantifiable metrics, such as RMSE, the *MPCxLLM* module is capable of processing and responding to a wide range of other natural language instructions. The chosen scenarios are simply representative examples where measurable outcomes — like centerline deviation, velocity tracking, and smooth driving — allow for clear and reproducible analysis. However, the methodology and interaction flow described are not limited to these examples and can extend to other instructions.

Table III demonstrates the controller adaptability through the evaluated LLMs. The MPC baseline, represents the different RMSEs adherences with default parameters, tuned for nominal operation of the vehicle. All percentage improvements are reported in absolute percentage points. The results show that the inclusion of RAG generally improves instruction adherence for locally deployable LLMs (*Phi3-mini*, *Qwen2.5*) by approximately 40%, although performance without RAG can be negligible or even detrimental. *GPT4o* cannot be deployed via LoRA fine-tuning and is therefore excluded. For locally deployable models, LoRA fine-tuning offers an additional 15–20% improvement over RAG-only setups. Hence, in terms of controller adaptability, both RAG and LoRA demonstrate a 45.1% and 52.2% (for *Phi3-mini* and *Qwen2.5* respectively) improvement over the nominal MPC baseline. Lastly, quantization introduces a minor 1-6% performance drop but is essential for achieving acceptable computational performance on the OBC, as detailed in Section IV-D. Note that the numerical values should be interpreted with caution, as the closed-loop experiment is not deterministic, since it depends on the MPC's adherence quality and the specific map used in the open-source simulation environment [4], hence the relative improvement should be regarded. From the local LLMs, Table III shows that *Qwen2.5-7b* consistently outperforms *Phi3-mini-3.8b* in terms of controller adaptability.

Bubble 2 illustrates a qualitative result of the *MPCxLLM* module. The input instruction is passed to the LLM, where the prompt demonstrates how the LLM is guided by the instruction and optionally enriched by the *MPCxLLM* RAG. Finally, the module's output displays a new set of MPC cost and constraint parameters that are subsequently parsed and transmitted to the MPC as a ROS dynamic reconfigurable parameter.

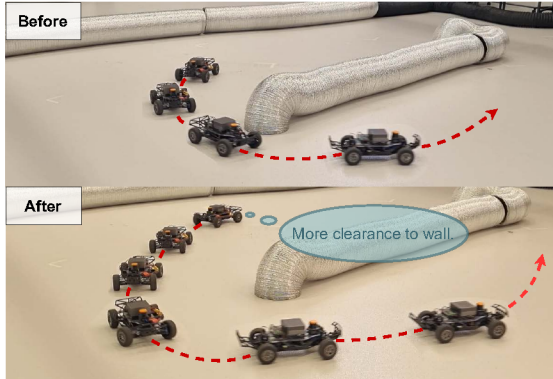
### C. Physical Robot

Multiple qualitative examples of the proposed framework operating on the physical robot are shown in Figure 5. In Figure 5a, the human instructs the robot to increase its distance from the wall. The *before* image illustrates how close the robot was initially driving to the wall, while the *after* image demonstrates a much safer clearance achieved through adjustments made by the *MPCxLLM* module. The *MPCxLLM* was prompted with: "Drive further away from the wall."

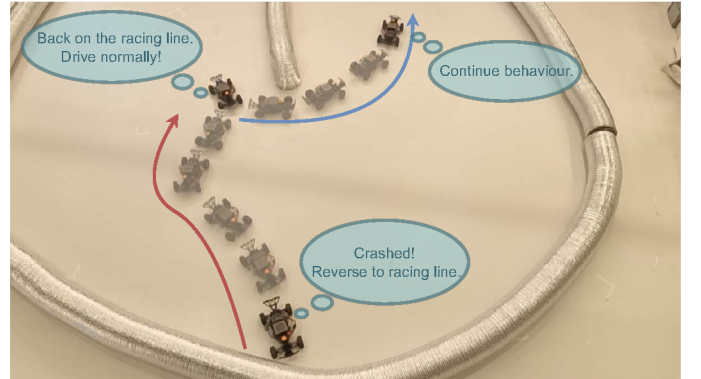
Figure 5b showcases the *DecisionxLLM* and *MPCxLLM* modules combined, by detecting a crash and subsequently instructing the robot to reverse and then safely resume its

| LLM            | Params | Quant | LoRA | RAG | $E_C[m] \downarrow$ | $E_V[ms^{-1}] \downarrow$ | $E_R[ms^{-1}] \downarrow$ | $E_S[ms^{-2}] \downarrow$ | Improve [%] $\uparrow$ |
|----------------|--------|-------|------|-----|---------------------|---------------------------|---------------------------|---------------------------|------------------------|
| MPC (baseline) | -      | -     | -    | -   | 0.7                 | 1.6                       | 4.6                       | 1.5                       | -                      |
| GPT4o          | ?      | ?     | ✗    | ✗   | <b>0.5 (19.3%)</b>  | <b>1.8 (-11.8%)</b>       | 1.5 (68.1%)               | <b>1.4 (5.1%)</b>         | <b>20.2%</b>           |
| Phi3-mini      | 3.8B   | FP16  | ✗    | ✗   | N.C.                | 1.9 (-17.1%)              | 2.8 (39.5%)               | 11.8 (-711.4%)            | -229.6% <sup>†</sup>   |
| Qwen2.5        | 7B     | FP16  | ✗    | ✗   | 0.7 (-14.6%)        | 1.8 (-14.4%)              | <b>1.3 (70.9%)</b>        | 2.0 (-40.4%)              | 0.5%                   |
| GPT4o          | ?      | ?     | ✗    | ✓   | <b>0.5 (21.3%)</b>  | 0.3 (81.5%)               | <b>1.2 (74.7%)</b>        | 1.3 (7.5%)                | <b>46.3%</b>           |
| Phi3-mini      | 3.8B   | FP16  | ✗    | ✓   | 0.6 (6.6%)          | 0.6 (60.8%)               | 4.2 (9.8%)                | <b>0.3 (82.5%)</b>        | 39.9%                  |
| Qwen2.5        | 7B     | FP16  | ✗    | ✓   | 0.6 (8.6%)          | <b>0.1 (91.7%)</b>        | 1.5 (67.7%)               | 1.4 (4.8%)                | 43.2%                  |
| Phi3-mini      | 3.8B   | FP16  | ✓    | ✓   | 0.7 (-0.3%)         | 0.6 (59.3%)               | <b>1.0 (78.8%)</b>        | <b>0.8 (42.8%)</b>        | 45.1%                  |
| Qwen2.5        | 7B     | FP16  | ✓    | ✓   | <b>0.4 (37.2%)</b>  | <b>0.4 (71.9%)</b>        | 1.1 (76.4%)               | 1.1 (23.2%)               | <b>52.2%</b>           |
| Phi3-mini      | 3.8B   | Q5    | ✓    | ✓   | 0.57 (12.6%)        | 0.5 (70.7%)               | 1.2 (73.6%)               | 1.1 (-22.3%)              | 44.%                   |
| Qwen2.5        | 7B     | Q5    | ✓    | ✓   | <b>0.43 (33.4%)</b> | <b>0.2 (85.8%)</b>        | <b>1.3 (71.3%)</b>        | <b>1.5 (-3.3%)</b>        | <b>46.8%</b>           |

TABLE III: Quantitative Comparison of LLM configurations with MPC. Performance metrics include deviation from the centerline ( $E_C[m]$ ), reference velocity ( $E_V[ms^{-1}]$ ), reversing accuracy ( $E_R[ms^{-1}]$ ), and driving smoothness ( $E_S[ms^{-2}]$ ), with percentage improvements shown relative to the MPC baseline. The average improvement column summarizes overall performance across all metrics. <sup>†</sup> indicates an average over completed runs, excluding N.C. (*Not Completed*). ? indicates that this is proprietary knowledge, not known to the public.



(a) Prompt: "Drive further away from the wall."



(b) Prompt: "Drive normally!"

Fig. 5: Visualization of physical robot behaviors: (a) increasing clearance from walls, (b) recovering from a crash scenario. All experiments were performed utilizing the edge-deployed *Qwen2.5-7b* LLM.

path. The *DecisionxLLM* module is instructed by the human to: "Drive normally!". The full model outputs of the LLM is shown in Bubble 3. More qualitative experiments on the physical robot are depicted in Appendix D.

#### D. Computation

Given the need for efficient hardware utilization in autonomous driving scenarios and the necessity for real-time interaction, the efficiency based on two models is evaluated and discussed in Table IV for the locally deployable *Phi3-mini-3.8b* and *Qwen2.5-7b*. The same input prompt was used for all compute evaluations and performed 60 times sequentially.

For the framework based on *Phi3*, when deployed on RTX 3090 hardware using FP16, the model with 3.8 billion parameters achieves a token output rate of 25.23 tokens per second and utilizes 4.3 GB of memory. In contrast, when quantized to Q5\_k\_m, the memory usage decreases to 3.9 GB, and the throughput speed significantly increases to 148.36 tokens per second. On the computationally constrained *Jetson Orin* hardware, the FP16 *Phi3* achieves an inference time of 15.29

seconds, while the Q5\_k\_m quantized model shows improved performance with an inference time of 4.80 seconds for 154 tokens. For the framework based on Qwen, the outcome is similar. On the RTX 3090, when the model is quantized to Q5\_k\_m, the throughput speed rises to nearly  $9.7\times$  that of FP16, reaching 107.52 tokens per second. Thus, the inference time decreases significantly to 0.91 seconds, approximately 20% of the time required for FP16. On the *Jetson Orin*, the Q5\_k\_m configuration increases the throughput to about  $10.5\times$  compared to the FP16, at 22.12 tokens per second. The inference time is lowered to 5.52 seconds, less than 20% of that of the FP16. These results highlight the substantial efficiency gains achievable through post-training quantization, enabling the deployment of our LLMs-based framework in computationally constrained hardware while maintaining real-time capabilities on robotic platforms.

#### V. CONCLUSION

This work introduces a hybrid architecture that enables the integration of low-level MPC and edge-deployed LLMs to



| HW              | LLM  | Quant | Param<br>[#B] | Mem<br>[GB] | Tokens<br>[#] | Tokens/s<br>[s <sup>-1</sup> ]↑ | Latency [s]↓ |             |
|-----------------|------|-------|---------------|-------------|---------------|---------------------------------|--------------|-------------|
|                 |      |       |               |             |               |                                 | $\mu_t$      | $\sigma_t$  |
| RTX 3090        | Phi3 | FP16  | 3.8           | 4.3         | 72            | <b>25.23</b>                    | <b>2.85</b>  | <b>0.05</b> |
|                 | Qwen | FP16  | 7             | 7.8         | 50            | 11.14                           | 4.51         | 0.08        |
|                 | Phi3 | Q5    | 3.8           | 3.9         | 110           | <b>148.36</b>                   | <b>0.75</b>  | <b>0.06</b> |
| Jetson Orin AGX | Qwen | Q5    | 7             | 5.6         | 97            | 107.52                          | 0.91         | 0.08        |
|                 | Phi3 | FP16  | 3.8           | 4.2         | 72            | <b>4.71</b>                     | <b>15.29</b> | <b>0.48</b> |
|                 | Qwen | FP16  | 7             | 6.9         | 64            | 2.11                            | 30.48        | 0.59        |
|                 | Phi3 | Q5    | 3.8           | 3.6         | 154           | <b>32.47</b>                    | <b>4.80</b>  | <b>0.58</b> |
|                 | Qwen | Q5    | 7             | 5.3         | 121           | 22.12                           | 5.52         | <b>0.58</b> |

TABLE IV: Comparison of computational performance for locally deployable models, *Phi3-mini-3.8b* and *Qwen2.5-7b*. The LLMs were deployed on both an RTX 3090 GPU and the *Jetson Orin AGX* robotic OBC. FP16 denotes full-precision models, while Q5 represents the Q5\_k\_m quantized models implemented via the `llama.cpp` inference engine. The number of tokens denotes the output tokens generated for the given inference. The average inference latency with the standard deviation is denoted with  $\mu_t, \sigma_t$  respectively.

enhance robotic decision-making and HMI through natural language in ADS. It offers the flexibility to choose between different LLMs based on operational requirements such as cloud connectivity, privacy considerations, and latency constraints. This approach bridges the gap between high-level reasoning and low-level control adaptability.

On locally deployed LLMs, the *DecisionxLLM* module demonstrates up to 10.45% improvement in reasoning accuracy when augmented with RAG and LoRA fine-tuning. The *MPCxLLM* module showcases controller adaptability, achieving up to a 52% improvement in controller adaptability, highlighting how natural language can adjust low-level MPC parameters to achieve flexible robotic behaviors while maintaining safety and constraint satisfaction through MPC systems.

Furthermore, this work demonstrates the deployment of embodied AI locally on embedded platforms, highlighting the importance of quantization in enabling real-time performance. Through post-training quantization, a 10.5 $\times$  improvement in throughput are achieved for the *Qwen2.5-7b* model on the *Jetson Orin AGX* OBC, allowing for efficient deployment of LLMs on resource-constrained hardware.

## VI. LIMITATIONS

One limitation of this approach is the relatively slow decision-making and controller adaptation process, which may fail to capture subtle (or high-frequency) behavioral nuances in the robotic state. Additionally, the reliance on text-based LLMs constrains the reasoning capabilities to state-based information alone. In contrast, multimodal LLMs could significantly enhance performance by incorporating visual data, enabling richer and more context-aware reasoning. However, the computational constraints of the embedded OBC necessitated the use of standard LLMs as an initial step. Future

work may address these constraints by exploring efficient deployment strategies for multimodal models on resource-limited hardware and investigating their reasoning capabilities. Lastly, locally deployed LLMs are not without flaws. They occasionally introduce reasoning errors and inconsistencies in controller adaptability, highlighting areas for improvement by retraining with a larger amount of distillation data. The proposed framework should therefore be viewed as a potential approach to integrating knowledge into ADS.

## ACKNOWLEDGMENTS

This work is funded in part by the dAIEDGE project supported by the EU Horizon Europe research and innovation programme under Grant Agreement Number: 101120726.

## REFERENCES

- [1] Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*, 2024.
- [2] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [3] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- [4] Nicolas Baumann, Edoardo Ghignone, Jonas Kühne, Niklas Bastuck, Jonathan Becker, Nadine Imholz, Tobias Kränzlin, Tian Yi Lim, Michael Lötscher, Luca Schwarzenbach, et al. Forzaeth race stack—scaled autonomous head-to-head racing on fully commercial off-the-shelf hardware. *Journal of Field Robotics*, 2024.
- [5] Christian Berger and Bernhard Rumpel. Engineering autonomous driving software. *Experience from the DARPA Urban Challenge*, pages 243–271, 2012.
- [6] Jan-Aike Bolte, Andreas Bar, Daniel Lipinski, and Tim Fingscheidt. Towards corner case detection for autonomous driving. In *2019 IEEE Intelligent vehicles symposium (IV)*, pages 438–445. IEEE, 2019.
- [7] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [8] Arnav Chavan, Raghav Magazine, Shubham Kushwaha, Mérouane Debbah, and Deepak Gupta. Faster and lighter llms: A survey on current challenges and way forward. *arXiv preprint arXiv:2402.01799*, 2024.

- [9] Jiasi Chen and Xukan Ran. Deep learning with edge computing: A review. *Proceedings of the IEEE*, 107(8): 1655–1674, 2019. doi: 10.1109/JPROC.2019.2921977.
- [10] Michael Han Daniel Han and Unsloth team. Unsloth, 2023. URL <http://github.com/unslothai/unsloth>.
- [11] Jiafei Duan, Wentao Yuan, Wilbert Pumacay, Yi Ru Wang, Kiana Ehsani, Dieter Fox, and Ranjay Krishna. Manipulate-anything: Automating real-world robots using vision-language models. *arXiv preprint arXiv:2406.18915*, 2024.
- [12] Florian Fuchs, Yunlong Song, Elia Kaufmann, Davide Scaramuzza, and Peter Dürri. Super-human performance in gran turismo sport using deep reinforcement learning. *IEEE Robotics and Automation Letters*, 6(3):4257–4264, 2021. doi: 10.1109/LRA.2021.3064284.
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11): 1231–1237, 2013.
- [14] Georgi Gerganov and Ope-Source Contributors. Llama.cpp, 2023. URL <https://github.com/ggerganov/llama.cpp>.
- [15] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- [16] Florian Heidecker, Jasmin Breitenstein, Kevin Rösch, Jonas Löhdefink, Maarten Bieshaar, Christoph Stiller, Tim Fingscheidt, and Bernhard Sick. An application-driven conceptualization of corner cases for perception in highly automated driving. In *2021 IEEE Intelligent Vehicles Symposium (IV)*, pages 644–651. IEEE, 2021.
- [17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [18] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 2023.
- [19] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- [20] Seif Ismail, Antonio Arbues, Ryan Cotterell, René Zurbügg, and Carmen Amo Alonso. Narrate: Versatile language architecture for optimal control in robotics. *arXiv preprint arXiv:2403.10762*, 2024.
- [21] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science Robotics*, 5(47):eabc5986, 2020. doi: 10.1126/scirobotics.abc5986. URL <https://www.science.org/doi/abs/10.1126/scirobotics.abc5986>.
- [22] Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, et al. Towards fully autonomous driving: Systems and algorithms. In *2011 IEEE intelligent vehicles symposium (IV)*, pages 163–168. IEEE, 2011.
- [23] Zhenyan Lu, Xiang Li, Dongqi Cai, Rongjie Yi, Fangming Liu, Xiwen Zhang, Nicholas D Lane, and Mengwei Xu. Small language models: Survey, measurements, and insights. *arXiv preprint arXiv:2409.15790*, 2024.
- [24] Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. *arXiv preprint arXiv:2310.12931*, 2023.
- [25] Marco Pavone. Decision making and control with llms. Lecture presented at NVIDIA GTC 2024, 2024. URL <https://www.nvidia.com/en-us/on-demand/session/gtc24-s62855/>. Accessed: 2024-07-02.
- [26] Roland Siegwart, Illah R. Nourbakhsh, and Davide Scaramuzza. *Introduction to Autonomous Mobile Robots*. The MIT Press, 2nd edition, 2011. ISBN 0262015358.
- [27] Tushar Prasanna Swaminathan, Christopher Silver, and Thangarajah Akilan. Benchmarking deep learning models on nvidia jetson nano for real-time systems: An empirical investigation. *arXiv preprint arXiv:2406.17749*, 2024.
- [28] Tai Wang, Xiaohan Mao, Chenming Zhu, Runsen Xu, Ruiyuan Lyu, Peisen Li, Xiao Chen, Wenwei Zhang, Kai Chen, Tianfan Xue, et al. Embodiedscan: A holistic multi-modal 3d perception suite towards embodied ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19757–19767, 2024.
- [29] Licheng Wen, Daocheng Fu, Xin Li, Xinyu Cai, Tao Ma, Pinlong Cai, Min Dou, Botian Shi, Liang He, and Yu Qiao. Dilu: A knowledge-driven approach to autonomous driving with large language models. *arXiv preprint arXiv:2309.16292*, 2023.
- [30] Peter R Wurman, Samuel Barrett, Kenta Kawamoto, James MacGlashan, Kaushik Subramanian, Thomas J Walsh, Roberto Capobianco, Alisa Devlic, Franziska Eckert, Florian Fuchs, et al. Outracing champion gran turismo drivers with deep reinforcement learning. *Nature*, 602(7896):223–228, 2022.
- [31] Wenhao Yu, Nimrod Gileadi, Chuyuan Fu, Sean Kirmani, Kuang-Huei Lee, Montse Gonzalez Arenas, Hao-Tien Lewis Chiang, Tom Erez, Leonard Hasenclever, Jan Humplik, et al. Language to rewards for robotic skill synthesis. *arXiv preprint arXiv:2306.08647*, 2023.
- [32] Baha Zarrouki, Verena Klös, Nikolas Heppner, Simon Schwan, Robert Ritschel, and Rick Voßwinkel. Weights-varying mpc for autonomous vehicle guidance: a deep reinforcement learning approach. In *2021 European Control Conference (ECC)*, pages 119–125. IEEE, 2021.

## APPENDIX A

### LORA FINE-TUNING EXAMPLES

For *DecisionxLLM*, we generate synthetic data via a two-stage process. First, in simulation, the vehicle collects diverse states (e.g., centerline tracking, reversing, unsafe maneuvers). These states are input to *GPT4o*, with prompts enriched by a RAG context (Listing 3) that injects robot-specific details (e.g., speed ranges, safe distances). Parameter values are randomized programmatically to avoid overfitting, ensuring diversity in the synthetic dataset. *GPT4o* then produces behavior descriptions based on queries such as: “The human wants to: Drive Safely. Check if this state: robot\_state adheres to the command. Additional context: RAG\_info. Decide: (a) Yes, (b) No, adjust behavior.” This yields 626 state-prompt response pairs for *DecisionxLLM* fine-tuning.

For *MPCxLLM*, *GPT4o* is prompted with the human instruction (e.g., “Do not exceed speeds of 10 km/h.”), *base\_memory* (Listing 1), and *RAG\_memories* (Listing 2). Programmatic randomization of the numeric inputs prevents overfitting, producing 150 prompt response pairs for LoRA fine-tuning of MPC adaptation.

## APPENDIX B

### ADDITIONAL MPC INFORMATION

The baseline MPC parameters within Table III have been set to ensure nominal tracking of the racing line in the simulation environment, which are the default parameters visible in Listing 1. Hence, in the controller adaptability experiment Section IV-B, the *MPCxLLM* module was tasked to adapt the nominal behavior (e.g., from tracking the racing line to “Reverse the car!”) as described in Section IV-B — thus yielding a quantitative and measurable alteration of the nominal (baseline) MPC behavior.

The base memory of Listing 1 within the *MPCxLLM* prompt Bubble 2 serves as a predefined knowledge foundation for mapping high-level natural language instructions into precise MPC parameters. It encodes the cost structure, tuneable parameters, and safety constraints essential for dynamic control adjustments. Each parameter is strictly defined with names, valid ranges, and default values that are defined as ROS dynamic reconfigure parameters.

```
# Cost expression with adjustable weights:
model.cost_expr_ext_cost = (
    weight_qn * n**2 +
    weight_qalpha * alpha**2 +
    weight_qv * (v - V_target)**2 +
    weight_qac * der_v**2 +
    weight_ddelta * derDelta**2 +
    u.T @ R @ u
)

# Tuneable cost weights and constraints (USE EXACT NAMES, DO NOT CREATE NEW ONES!):
# param: min, max, default # description
qv 0, 2, 10 # Velocity weight: minimizes speed tracking error
qn 0, 100, 20 # Lateral weight: minimizes deviation from the track
qalpha 0, 100, 7 # Heading weight: minimizes orientation error
qac 0, 1, 0.01 # Acceleration weight: penalizes high acceleration
qdelta 0, 100, 0.1 # Steering weight: penalizes fast steering changes
alat_max 0, 20, 10 # Max lateral acceleration: limits side force
a_min -20, 0, -5 # Min acceleration: lower acceleration bound
a_max 0, 20, 5 # Max acceleration: upper acceleration bound
v_min -2, 5, 1 # Min velocity: lower speed bound
v_max -1, 10, 5 # Max velocity: upper speed bound
track_safety_margin 0, 1.0, 0.45 # Safety margin: increases track boundary margin
```

Listing 1: MPCxLLM Base Memory

Listing 2 shows the *MPCxLLM* RAG memories establish context-specific mappings between natural language instructions and corresponding MPC parameter adjustments. Each memory entry provides guidance on how parameters from the *base\_memory* influence the MPC, resembling simplified *driving school* instructions, akin to [29]. These memory entries are modular and can be combined to address more complex scenarios. The RAG mechanism ensures effective retrieval by performing similarity matching between the user instruction and scenario descriptions.

```
# Memory Entry 0:
Scenario:
To force going forwards v_min should be positive. If you want it to be able to reverse, then set v_min to
negative.
MPC Action:
mpc_params = {
    'v_min': positive, if you want to go forwards, else negative to reverse
```

```

}

# Memory Entry 1:
Scenario:
Always have v_max be higher than v_min.
MPC Action:
mpc_params = {
    'v_max': higher than v_min
}

...

# Memory Entry 10:
Scenario:
To minimize the lateral acceleration and jerk, set alat_max to a low value and a_min and a_max close to
zero. If you want to drive more aggressive, then set alat_max to a higher value.
MPC Action:
mpc_params = {
    'alat_max': low for minimizing lateral acceleration, else high for aggressive driving
}

```

Listing 2: MPCxLLM RAG Memories

## APPENDIX C ADDITIONAL DECISION-MAKING INFORMATION

Decision accuracy was evaluated using a dataset of 200 state samples (using a different state-dataset as the LoRA fine-tuning data), yielding 1600 state-command pairs across 8 driving commands Section IV-A. Each state corresponds to a command (e.g., "Drive faster than 3 m/s."). The *DecisionxLLM* module performs binary classification to determine if the robot's state meets the command (Figure 1 or Figure 3). Adherence is verified computationally; for instance, for "Reverse the car!", the system checks if the velocity is negative. Accuracy is measured as the percentage of correctly classified state-command pairs.

Listing 3 is an excerpt is from the *DecisionxLLM* RAG entries. The RAG enhances the `human_prompt` by augmenting it with context retrieved through similarity searches across various hints. The top-k most relevant hints are selected to enrich the prompt, providing additional contextual grounding. This approach allows users to ground the robot using natural language. For instance, a user might specify a different desired nominal speed range, such as operating the car between 2-4 m/s<sup>-1</sup>, thereby supplying *DecisionxLLM* with updated operational constraints.

```

# Hint 1:
If the d-speed is above than 0.5m/s is high.

# Hint 2:
Unless specified differently by the human, the car is usually driving at speeds between 5 and 7m/s.

# Hint 3:
If the distance to a wall is smaller than 0.4m, the car is close to that wall.

...

# Hint 8:
A d-coordinate above 0.3m is considered not to be on the racing line.

# Hint 9:
The car is oscillating if the d-coordinate oscillates between positive and negative values exceeding a
magnitude of 0.3 metres.

# Hint 10:
Oscillations in d-coordinate under 0.3m or d-speed under 0.5m/s are negligible.

```

Listing 3: DecisionxLLM RAG Memories

## APPENDIX D CRASH RECOVERY OUTPUT EXAMPLES ON THE PHYSICAL ROBOT

Bubble 3 captures the chat log from the experiment detailed in Figure 5b. The robot begins in a crashed state, facing a wall. The *DecisionxLLM* identifies this condition and adjusts the MPC to initiate a reversing maneuver. Once the robot successfully reverses out of the crash, it resumes normal driving along the track. This log presents the LLM output generated during the experiment, where the locally deployed *Qwen2.5-7b* was running on the OBC.

### Bubble 3a: DecisionxLLM – Instruction

Drive normally!

### Bubble 3b: DecisionxLLM – Prompt

The human wants to: *"Drive normally!"*

The data has been sampled for 2.0 seconds in 5 samples.

- **s-speed:** -0.009, -0.026, -0.03, -0.03, -0.013
- **d-speed:** 0.096, 0.091, 0.088, 0.09, 0.096
- **distance to left wall:** 2.953, 2.954, 2.954, 2.955, 2.955
- **distance to right wall:** 0.149, 0.148, 0.148, 0.147, 0.147
- **crashed:** True

Here are some guides to help you reason: `rag_hints`.

Check if the car is doing what the human wants. Choose one of the following actions to command the car:

- a) Continue behavior
- b) Change behavior: `instruction`

### Bubble 3c: DecisionxLLM – Response

The car is **facing a crash** and is reversing, which is **not the normal driving behavior**. The car should be facing forward and driving at a normal speed. The s-speed is close to 0m/s, indicating the car is stopped, and the d-speed is high, suggesting the car is oscillating around the racing line. The car is also close to the right wall, which is a sign of a crash. **Reversing is necessary** to get the car back on the racing line and back to normal driving.

Action:

**Change behavior:** Reverse the car to get back on the racing line.

### Bubble 3d: MPCxLLM – Response

```
new_mpc_params      =      {qv: 0.5, qn: 50.0, qalpha: 5.0, qac: 0.01, qddelta:
0.1, alat_max: 10.0, a_min: -5.0, a_max: 5.0, v_min: -2.0, v_max: -1.0,
track_safety_margin: 0.45}
```

**Explanation:** Setting parameters to reverse the car.

### Human Narration

The car was stuck in the wall. The *DecisionxLLM* stage made the correct decision, albeit a reasoning error on stating that the d-speed is high, but this did not change the overall decision-making. Then the *MPCxLLM* stage correctly adapts the MPC to induce reversing behavior.

### Bubble 3e: DecisionxLLM – Prompt

The human wants to: *"Drive normally!"*

The data has been sampled for 2.0 seconds in 5 samples.

- **s-speed:** -0.022, -0.014, -0.037, -0.005, -0.038
- **d-speed:** 0.091, 0.095, 0.088, 0.097, 0.088
- **distance to left wall:** 2.956, 2.955, 2.955, 2.955, 2.956
- **distance to right wall:** 0.146, 0.147, 0.147, 0.147, 0.146
- **crashed:** True

Here are some guides to help you reason: `rag_hints`.

Check if the car is doing what it needs to ensure recovery. Choose one of the following actions to command the car:

- a) Continue behavior
- b) Change behavior: `instruction`



### Bubble 3f: DecisionxLLM – Response

The **car has crashed and is reversing, so it should continue to reverse** until it is safe to proceed. The car's speed is very low, indicating it is stopped, and it is close to the right wall. Therefore, the car should reverse slowly to get back on the racing line.

Action:

**Continue behavior**

### Human Narration

Even though reversing behavior was already induced, the state was sampled while the car was still crashed into the wall. Hence the decision of continuing the reversing behavior is correct.

### Bubble 3g: DecisionxLLM – Prompt

The human wants to: *"Drive normally!"*

The data has been sampled for 2.0 seconds in 5 samples.

- **s-speed:** -0.254, -0.149, -0.086, -0.088, -0.029
- **d-speed:** 0.748, 0.281, 0.163, 0.155, 0.112
- **distance to left wall:** 0.372, 0.379, 0.386, 0.393, 0.391
- **distance to right wall:** 2.82, 2.846, 2.839, 2.832, 2.834
- **crashed:** False

Here are some guides to help you reason: *rag\_hints*.

Check if the car is doing what the human wants. Choose one of the following actions to command the car:

- a) Continue behavior
- b) Change behavior: *instruction*

### Bubble 3h: DecisionxLLM – Response

The car is reversing, which is **not the normal driving behavior**. The car should be facing forward to maintain normal driving. Therefore, the action should be to **stop the car from reversing and resume normal driving**.

Action:

**Change behavior:** Stop reversing and resume normal driving.

### Bubble 3j: MPCxLLM – Response

```
new_mpc_params = {'v_min': 1.5, 'v_max': 1.5, 'qn': 100, 'qalpha': 100, 'qac': 0.01, 'qddelta': 0.1, 'alat_max': 10, 'a_min': -5, 'a_max': 5}
```

**Explanation:** The parameters are tuned to smoothly resume normal driving.

### Human Narration

It is correctly detected that the car is no longer crashed and should now drive forwards again. The MPC parameters are then again adjusted to do so.

## A. Velocity Adherence on the Physical Robot

Fig. 6 illustrates the robot's ability to respond to human commands effectively by adjusting its velocity to remain within a specific range. This experiment is performed on the physical robot using the edge-deployed *Qwen2.5* model. Upon receiving the instruction to maintain a target velocity window, the robot promptly modifies its speed, demonstrating precise control and adaptability. This behavior highlights the system's capability to monitor and regulate its performance, ensuring that the velocity consistently stays within the designated limits.

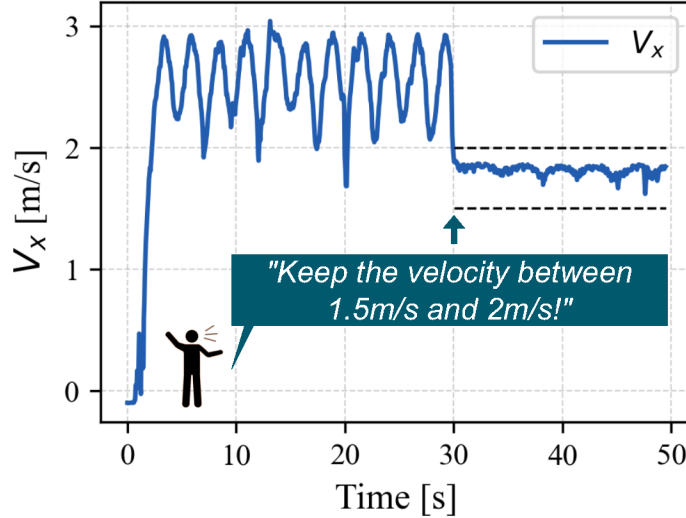


Fig. 6: Illustration of the robot’s velocity behavior before and after (at 30s) the human prompts the LLM: “Drive at speeds between 1.5 and 2.0 m/s”.

#### APPENDIX E EXTENDED DECISION-MAKING RESULTS

Table V extends Table II by presenting the performance of various LLMs across individual decision-making test cases, with the average accuracy reported in the final column. Additionally, this table includes evaluations of contemporary models such as *Phi4-14b*<sup>3</sup>, *Gemma2-9b*<sup>4</sup>, and *DeepSeek R1-distilled Qwen2.5-7b*<sup>5</sup>. These models represent state-of-the-art architectures at the time of this submission. However, due to time constraints, the quantization and Q5\_k\_m GGUF inference pipeline could not be implemented.

Furthermore, the large parameter count (or slower inference architecture) of *Phi4-14b* and *Gemma2-9b* — and in the case of *Qwen2.5-7b-R1distill*, the extended Chain of Thought (CoT) reasoning process — renders them infeasible for closed-loop MPC and simulation experiments. As a result, only decision-making results are presented here.

A key observation from the results is that, with and without RAG, *Qwen2.5-7b-R1distill* demonstrates a significant performance boost over its non-R1-distilled counterpart of the same parameter count. The † symbol denotes instances where the LLM was occasionally cut off during evaluation — despite doubling the maximum token limit to 1024 — directly leading to incorrect decision scores. This issue stems from the computational overhead associated with R1 CoT reasoning, resulting in a significantly higher number of output tokens and consequently longer inference and evaluation times. This means, that given more time, the performance values would probably be even higher.

Moreover, *Qwen2.5-7b-R1distill* exhibits a notable drop in performance when fine-tuned with LoRA. This finding is particularly intriguing, as the LoRA fine-tuning was based on Supervised Fine-Tuning (SFT) using *GPT-4o* data, suggesting a potential degradation in reasoning capability rather than improvement and that the LoRA tuning step could perhaps be improved by utilizing R1 generated data for SFT. Further, it is to note, that once LoRA tuned on *GPT4o* data, the LLM now again learns to be concise and is no longer cut-off, however this comes at the cost of reasoning accuracy performance.

<sup>3</sup><https://huggingface.co/unsloth/phi-4>

<sup>4</sup><https://huggingface.co/unsloth/gemma-2-9b>

<sup>5</sup><https://huggingface.co/unsloth/DeepSeek-R1-Distill-Qwen-7B>

| Model                | Params | Quant | LoRA | RAG | Centerline   | Close Wall  | Forward   | Oscillating | Racingline  | Reversed     | Speed       | Stop        | Avg. Accuracy |
|----------------------|--------|-------|------|-----|--------------|-------------|-----------|-------------|-------------|--------------|-------------|-------------|---------------|
| GPT4o                | ?      | ✗     | ✗    | ✗   | <b>94.57</b> | 80.5        | <b>81</b> | 80.5        | <b>81</b>   | 70.83        | 77.5        | <b>87.5</b> | <b>81.68</b>  |
| Phi3-mini            | 3.8B   | ✗     | ✗    | ✗   | 83.7         | 56.5        | 76        | 76.5        | 74          | 48.96        | 74.5        | 87          | 72.15         |
| Phi4-14b             | 14B    | ✗     | ✗    | ✗   | 91.3         | <b>83.5</b> | 74.5      | 75.5        | 71          | 62.5         | 71          | 77          | 75.79         |
| Gemma2-9b            | 9B     | ✗     | ✗    | ✗   | 84.78        | 86          | 80        | <b>81.5</b> | 78          | <b>84.9</b>  | <b>77.5</b> | 83.5        | 82.02         |
| Qwen2.5-7b           | 7B     | ✗     | ✗    | ✗   | 82.07        | 71          | 76        | 84          | 75.5        | 72.4         | 71.5        | 89.5        | 77.75         |
| Qwen2.5-7b-R1 distil | 7B     | ✗     | ✗    | ✗   | 92.93        | 84          | 75        | 79          | 80          | 72.4         | 74.5        | 82.5        | 80.04†        |
| GPT4o                | ?      | ✗     | ✗    | ✓   | <b>100</b>   | <b>98</b>   | <b>97</b> | <b>85.5</b> | <b>95.5</b> | <b>83.85</b> | <b>96</b>   | 84          | <b>92.48</b>  |
| Phi3-mini            | 3.8B   | ✗     | ✗    | ✓   | 90.76        | 66          | 81        | 84          | 78          | 68.23        | 77          | 84.5        | 78.69         |
| Phi4-14b             | 14B    | ✗     | ✗    | ✓   | 97.83        | 84          | 86.5      | 75          | 85          | 65.62        | 77          | 85          | 81.99         |
| Gemma2-9b            | 9B     | ✗     | ✗    | ✓   | 96.2         | 82          | 79.5      | 83          | 79          | 71.35        | 74          | <b>85</b>   | 81.26         |
| Qwen2.5-7b           | 7B     | ✗     | ✗    | ✓   | 90.79        | 78.5        | 83        | 84          | 81          | 73.44        | 78.5        | 90.5        | 82.47         |
| Qwen2.5-7b-R1 distil | 7B     | ✗     | ✗    | ✓   | 94.02        | 89          | 88.5      | 82.5        | 86.5        | 78.12        | 79          | 86.5        | 85.52†        |
| Phi3-mini            | 3.8B   | ✗     | ✓    | ✓   | 94.57        | 68.5        | 82.5      | 80.5        | 82          | 79.69        | 75.5        | <b>97.5</b> | 82.60         |
| Phi4-14b             | 14B    | ✗     | ✓    | ✓   | <b>100</b>   | <b>90.5</b> | <b>94</b> | 80.5        | <b>93.5</b> | 79.17        | <b>91.5</b> | 89.5        | <b>89.83</b>  |
| Gemma2-9b            | 9B     | ✗     | ✓    | ✓   | 82.61        | 72          | 76        | 75          | 73          | 72.4         | 66.5        | 72          | 73.69         |
| Qwen2.5-7b           | 7B     | ✗     | ✓    | ✓   | 97.28        | 78.5        | 90        | <b>86</b>   | 88.5        | <b>81.77</b> | 81          | 95.5        | 87.32         |
| Qwen2.5-7b-R1 distil | 7B     | ✗     | ✓    | ✓   | 85.33        | 73          | 81.5      | 79.5        | 85.5        | 85.42        | 83          | 93          | 83.28         |
| Phi3-mini            | 3.8B   | ✓     | ✓    | ✓   | 88.59        | 84.5        | 84        | 74.5        | 90          | 86.98        | 85          | 86          | 84.95         |
| Qwen2.5-7b           | 7B     | ✓     | ✓    | ✓   | <b>98.91</b> | <b>81.5</b> | <b>85</b> | <b>87.5</b> | <b>85.5</b> | 80.21        | <b>82</b>   | <b>95.5</b> | <b>87.02</b>  |

TABLE V: Extended performance evaluation of Table II showcasing the single decision making test cases and additional LLMs. Accuracy units are reported in %. ? denotes proprietary information. † denotes that the evaluation was occasionally clipped to the max tokens set to 1024 per inference, which might indicate that a higher score could potentially have been achieved with more time.