CSG-Driver: Common Sense Guided Autonomous Driving under Legal Compliance and Practical Flexibility in Dilemma Situations

Yerin Lee[1], Seonggyu Lee[1], Shiho Kim[2]

[1] BK21 Graduate Program in Intelligent Semiconductor Technology, Yonsei University,
[2] Yonsei University

Abstract

Despite significant advancements in autonomous driving research, addressing long-tail cases remains a critical challenge. In this context, LLMs have gained attention for their interpretability and explainability, leading to increasing efforts to integrate them into autonomous driving tasks.

In this paper, we propose **CSG-Driver**, a human-like driving agent that combines compliance with road traffic laws and adaptability through the application of human common sense. We developed a closed-loop driving system within the CARLA simulator, which converts sensory data into natural language descriptions, incorporates road traffic laws, and utilizes prompts based on human driving behavior and past experiences. To address challenges in decision-making, the system employs common sense prompts and Chain-of-Thought reasoning to handle complex scenarios such as intersections with yellow lights, illegal parking avoidance, and highway driving.

Our experimental results demonstrate that CSG-Driver effectively resolves long-tail cases by leveraging LLMs to balance safety, traffic law compliance, and practical adaptability.

Introduction

Imagine you are passing an intersection at high-speed while the signal changes from green to yellow. You must decide whether to stop or pass based on the speed of your car, how many vehicles are around you, and how far from the stop line etc. This concept is called "dilemma zone" by (Gazis, Herman, and Maradudin 1960). In order to cope with this kind of dilemma, some research encodes traffic laws to be interpretable for a machine and this is used to train the reinforcement learning framework for autonomous driving strategies (Zhang et al. 2022), (Lin et al. 2022). With the advancements in LLMs, recent studies have proposed enhancing their reasoning capabilities through prompt engineering techniques such as Chain of Thought (CoT) and Tree of Thought (ToT), enabling human-like driving in AD systems (Touvron et al. 2023),(Achiam et al. 2023),(Wang et al. 2022).

In this paper, we argue that bridging the gap in handling long-tail cases and ambiguities requires equipping autonomous driving models with "common sense," allowing them to act more like humans when facing dilemma situations. To address this, we propose a "CSG-Driver" system, where the vehicle strikes a balance between human-like behavior—occasionally bending rules when the situation warrants—and strictly adhering to traffic laws, utilizing the common sense of human drivers.

Related Works

For drive planning, some research makes progress in the end-to-end autonomous driving field. In (Hu et al. 2023) paper, UniAD, a comprehensive end-to-end system, was presented to avoid accumulative errors by using a unified-query mechanism between the perception, prediction, and planning components. (Chen and Krähenbühl 2022) presented, a mapless, end-to-end driving system which trains from all nearby vehicles and predicts trajectories called LAV. This paper introduced supervisory tasks like 3D detection and segmentation for viewpoint-invariant feature extraction. However, it is almost impossible to train the model with all driving scenarios for solving long-tail cases which causes vulnerable handling.

Recent research focuses on interaction with humans and explainability for decision-making. (Huang, Wu, and Lv 2022) integrates human prior knowledge and deep reinforcement learning by taking state-action pairs by collecting the expert performs and their actions. Some studies utilize LLMs and build some frameworks to improve human-like reasoning for some ambiguous situations. (Wang et al. 2023) proposed DriveMLM that leverages Multi-modal LLM (MLLM) taking input including user command and gives driving decisions and explanations for these decisions. (Jin et al. 2023) used human driver interview data to learn a framework. This paper learns directly from driver conversations and LLMs help to give feedback and guidance for a coach agent. (Wen et al. 2023) proposed DiLu utilizing three key modules, including memory, reasoning, and reflection. This framework gives driving decisions based on reinforcement learning. This model tries to act like a human by reflecting the past experiences and taking real-world datasets.

Method

Motivation

(Reed et al. 2021) emphasizes the ethical necessity of "Rule-Breaking for Safety" in predefined conditions and advocates for integrating human-like flexibility into decision-making

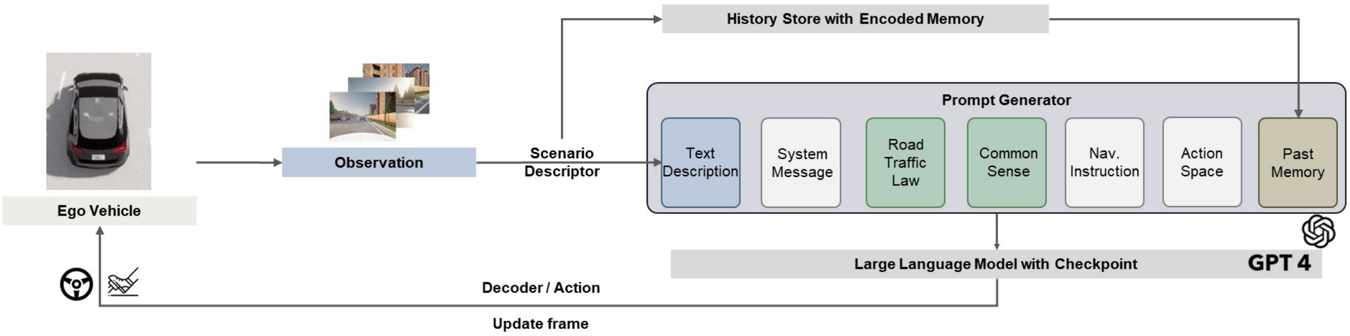

Figure 1: Architecture of **'CSG-Driver'**, integrating road traffic laws, common sense, and past experiences via an LLM in the CARLA simulator for adaptive, rule-flexible decision-making.

frameworks for autonomous systems. Inspired by this concept, we propose 'CSG-Driver', a human-like driving agent that adheres to road traffic laws while exercising the common sense and adaptability of human drivers through the use of a Large Language Model (LLM).

As illustrated in Fig. 1, we developed a closed-loop driving system within the CARLA simulator, which integrates both road traffic laws and the common sense of human drivers. The system begins by observing the surrounding environment from multiple viewpoints and converting into textual descriptions that capture the current scenario. Additionally, we incorporated the road traffic laws, which often introduce ambiguity in decision-making due to variations in interpretation.

To handle such dilemmas, we designed a mechanism inspired by human behavior in "dilemma zones" and converted these behaviors into textual guides representing common sense. Simultaneously, past experiences are integrated into the input prompt, enabling the LLM to contextualize its understanding of closed-loop driving tasks and make precise, adaptive decisions in continuous scenarios. Detailed implementations are discussed in the following sections.

Describing Scenarios Through Observation

In this section, we focus on converting real-time road situations into natural language descriptions that are interpretable by the LLM. By following predefined sentence structures, we generate text descriptions tailored to the current scenario. These descriptions include details such as the positions and speeds of the ego vehicle and surrounding vehicles, the map data of the ego vehicle's current path, traffic light statuses, and navigation instructions.

Information about the ego vehicle and surrounding vehicles is described relative to the ego vehicle's position. For instance, objects and vehicles in the left lane, ahead, behind, and in the right lane are contextualized to reflect their spatial relationships to the ego vehicle. These detailed descriptions are subsequently used as inputs for prompt generation and as part of the current frame's input for memory storage in past experiences.

Common Sense in Human Driving

Common sense, derived from human experience, enables drivers to make context-aware decisions beyond strict rule-following. For example, when approaching a yellow light, drivers evaluate factors like speed, distance to the stop line, and traffic conditions. While LLMs inherently possess general common sense, their effectiveness in autonomous driving requires domain-specific adaptation. To address this, we designed common sense prompts using natural language to emulate human-like reasoning.

Inspired by (Papaioannou 2007), we extracted key behavioral elements, such as speed and distance in dilemma zones, and translated them into prompts for LLMs. These prompts enhance the system's ability to interpret complex scenarios and balance rule compliance with practical safety, as shown in Fig. 2. This approach bridges the gap between rigid rules and human-like adaptability, improving the safety and flexibility of autonomous driving systems.

Incorporating Past Experiences

To enable closed-loop simulation and handle continuous driving scenarios, we designed a mechanism to store the ego vehicle's past experiences and integrate them into the current scenario. This memory system ensures that prior actions and observations influence the decision-making process, creating a more context-aware autonomous driving system.

Our approach involves encoding and storing observations from previous time steps in a structured format that can be efficiently retrieved when needed. This structured memory enables the system to organize past experiences into a searchable format, allowing relevant information to be seamlessly integrated into the decision-making process. These stored experiences are formatted in a few-shot structure and used as inputs for prompt generation.

Generating Prompts

As shown in Fig. 1, the Prompt Generator combines multiple inputs to create a comprehensive prompt for the LLM, with each component playing a distinct role in shaping the LLM's understanding and decision-making process. The system message outlines key contextual elements, such as the traffic rules the ego vehicle must follow, its driving style, and its operational role. We also predefined the action space to ensure clarity in decision-making. The action space includes options such as [IDLE, Turn-left, Turn-right, Acceleration, Deceleration, Emergency-stop, and Overtaking]. To maintain consistency in outputs, we established a predefined structure for the LLM's responses. This structure guides the LLM to generate both the appropriate action and its reasoning for every frame.

Applying Prompts to LLMs

For this part, we utilized 'GPT-4' as the Large Language Model (LLM). To enhance the reasoning process, we predefined several checkpoints to guide the LLM's thought flow and ensure it validates critical aspects during decision-making. These checkpoints included tasks such as 'Check Road Traffic Law,' 'Check Traffic Light Status,' 'Check Common Sense,' 'Check Front Status,' and three additional checks, making a total of seven. Each checkpoint represented a key verification step that the LLM had to address systematically during its reasoning process.

Additionally, we implemented the Chain-of-Thought (CoT) prompting technique introduced in (Wei et al. 2022). This method enables step-by-step reasoning, allowing the LLM to emulate human-like problem-solving processes. By combining these checkpoints with CoT prompting, the LLM could methodically analyze scenarios and make well-informed decisions, even in complex situations such as dilemma zones. The outputs generated by the LLM demonstrated its ability to carefully assess traffic conditions and select appropriate actions aligned with road safety and situational demands. This approach ensures a robust and adaptive decision-making system for autonomous driving.

Common Sense when yellow traffic light

1. Drivers closer to the stop line (e.g., within 40 m) are more likely to proceed, while those farther away (e.g., over 50 m) tend to stop.
2. Drivers traveling at higher speeds (e.g., over 50 km/h) are more likely to proceed, whereas those at lower speeds (e.g., under 30 km/h) are more inclined to stop.
3. Drivers with faster reaction times are more likely to make a quick decision to stop, while those with slower reaction times may proceed, even when it's risky.
4. Aggressive drivers are more likely to accelerate and proceed through the yellow light, while cautious drivers tend to decelerate and stop.
5. Drivers who perceive they can clear the intersection safely before the light turns red are more likely to proceed.
6. Drivers who feel pressured by closely following vehicles are more likely to proceed to avoid abrupt braking and potential rear-end collisions.
7. Drivers familiar with the intersection's timing are more likely to make accurate stop-or-go decisions compared to those unfamiliar with it.
8. Drivers on wet or slippery roads are more likely to stop due to the increased risk of skidding, while those on dry roads may proceed.
9. Larger vehicles (e.g., trucks or buses) are more likely to proceed, as stopping requires longer braking distances, whereas smaller cars may stop more easily.
10. In heavy traffic conditions, drivers are more likely to stop to avoid blocking the intersection, whereas in light traffic, they are more likely to proceed.
11. Drivers are more likely to stop when pedestrians are visible at the intersection or crosswalk.
12. Drivers who perceive the yellow light duration to be longer are more likely to proceed, while those who think it's shorter are more likely to stop.
13. Drivers who perceive higher risks (e.g., presence of red-light cameras or police monitoring) are more likely to stop, even if they are close to the stop line.
14. Distracted drivers (e.g., using a phone) may react slower and are more likely to make poor decisions at yellow lights.
15. At night, drivers are more cautious and tend to stop more often due to reduced visibility, whereas during the day, they may feel more confident to proceed.

Figure 2: **Example prompt of common sense at yellow traffic light.** A prompt created by analyzing the behavioral patterns of human drivers in natural language.

Experimental Results

To demonstrate how an autonomous driving system can effectively apply common sense and learn to drive like a human, while also exposing the limitations of strict compliance with road traffic laws, we conducted experiments using the CARLA simulator (Dosovitskiy et al. 2017) — an open-source platform for developing, training, and validating autonomous driving systems in realistic urban environments.

Intersection with Yellow Traffic Light

The experiment was conducted in Town 05 of the CARLA simulator, which features a "Squared-grid town with cross junctions and a bridge." The ego vehicle was initialized at random points and driven at varying speeds ranging from 30 km/h to 80 km/h. As the vehicle approached an intersection within a distance of 50 meters from the stop line, the traffic light was switched to yellow to evaluate its decision-making process in response to the scenario.

Fig. 3 illustrates the ratio of the ego vehicle's decisions to either stop at or cross the stop line, based on its speed. At lower speeds, such as 30 km/h, the vehicle frequently opted to stop due to sufficient time and space to decelerate safely. However, as speeds increased, the frequency of stopping decreased significantly, as abrupt braking at higher speeds could lead to hazardous situations. Instead, the ego vehicle demonstrated a preference for crossing the stop line, prioritizing safety and maintaining traffic flow.

Fig. 4(a) captures a snapshot of the experiment and the corresponding decision-making process. In this scenario, the vehicle identified a close following distance from the car behind, where an abrupt stop could increase the risk of a rear-end collision. Consequently, it accelerated to safely navigate through the intersection, demonstrating an adaptive approach to handling the dilemma zone by balancing safety and situational demands.

Driving Beyond the Rules: Handling Dilemmas

We also implemented scenarios where human drivers rely on common sense to evaluate the system's decision-making capabilities.

avoiding illegally parked vehicles Although road traffic laws specify vehicle behavior based on lane markings, certain situations, such as the one shown in **Fig. 4(b) - A**, require temporary deviation from these rules. For instance, bypassing a parked vehicle by momentarily crossing a yellow line represents a common-sense action that avoids excessive disruption to traffic flow. Without this flexibility, autonomous vehicles would remain stationary, causing severe congestion. By implementing this scenario, we evaluated the ego vehicle's ability to utilize the 7 predefined actions in the action space. The results demonstrated that the model could adaptively perform decisions resembling human drivers, effectively avoiding obstacles while maintaining smooth traffic flow.

highway driving Although traffic laws mandate strict adherence to speed limits, real-world highway environments often see vehicles adjusting their speed beyond legal limits to align with traffic flow. In **Fig. 4(b) - B**, the speed limit is 90 km/h, yet the ego vehicle, emulating human-like decision-making, considered the relative speed of surrounding vehicles and maintained sufficient following distance while accelerating slightly beyond the limit. This approach allowed the system to achieve a natural and seamless integration into the traffic flow, ensuring safety and efficiency while minimizing disruptions.

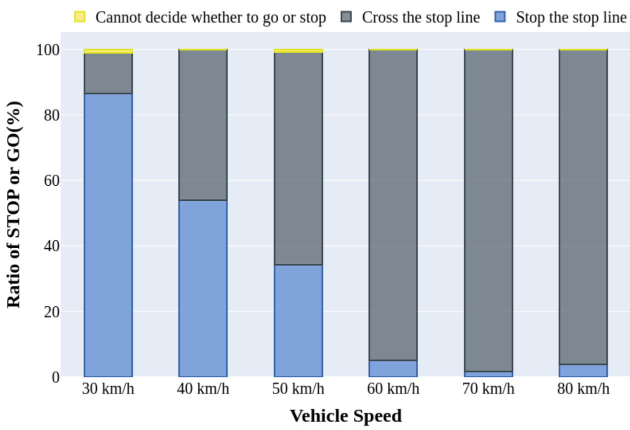

Figure 3: Decision ratio (STOP/GO) at yellow light across vehicle speeds.

Discussions and Conclusions

In this paper, we introduced **CSG-Driver**, a framework designed to balance flexibility in breaking traffic laws with maintaining safety and minimizing road congestion. By incorporating common sense derived from human research, we demonstrated how the system effectively deals with dilemma situations, ensuring human-like decision-making in challenging scenarios. Using the CARLA simulator, we evaluated the system in three key scenarios (1) Intersection with Yellow Light (2) Avoiding Illegally Parked Vehicles (3) Highway Driving. However, while applying common sense enhances flexibility, strict adherence to traffic laws remains a fundamental requirement. Striking the right balance between traffic law compliance and common sense application is critical to developing a realistic and effective driving system.

As future work, we aim to develop a learning mechanism that dynamically adjusts the proportion of traffic law and common sense application based on driving contexts. This will create a more realistic autonomous driving system capable of addressing complex scenarios. Additionally, we plan to conduct co-simulations with macro traffic simulations to analyze how the system influences overall traffic flow and congestion patterns. By pursuing these directions, we expect to validate the robustness and applicability of the proposed decision-making system further.

Ultimately, we hope this research contributes to resolving edge cases in scenarios where human and autonomous drivers coexist, aiding the commercialization of autonomous driving systems by ensuring smoother integration into real-world traffic environments.

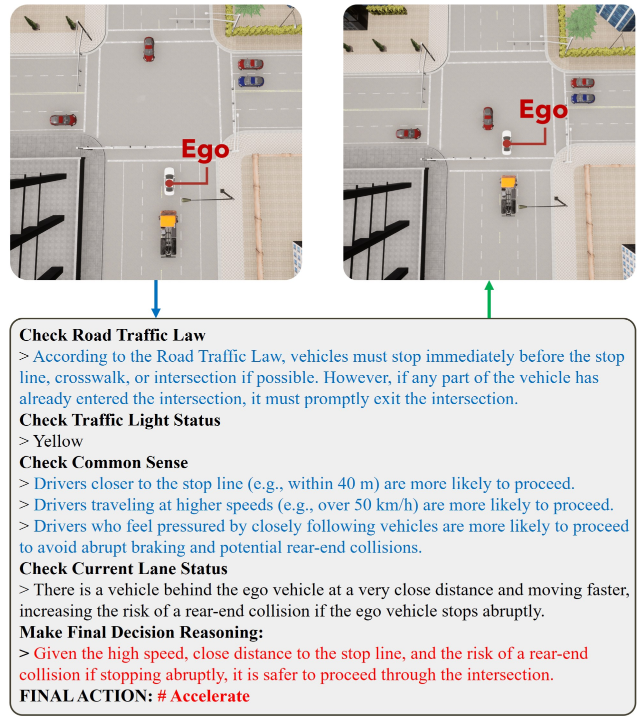

(a) **Intersection with Yellow Traffic Light**

(A) Avoiding illegally parked vehicles

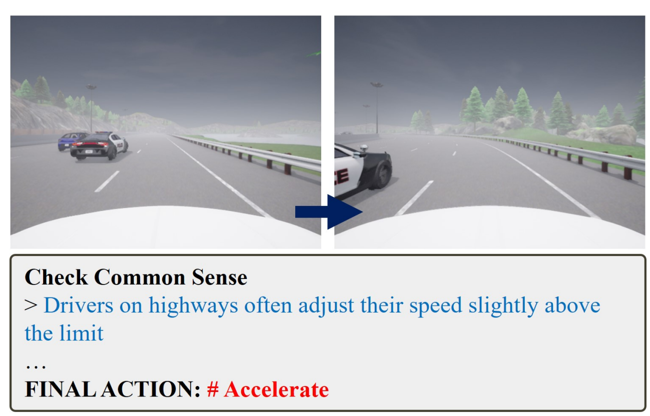

(B) Highway Driving

(b) **Driving Beyond the Rules: Handling Dilemmas**

Figure 4: An example of solving a dilemma zone using CSG-Driver's common sense.

Acknowledgments

This research was supported by Korea Institute for Advancement of Technology(KIAT) grant funded by the Korea Government(MOTIE) (P0020535, The Competency Development Program for Industry Specialist)

References

Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Chen, D.; and Krähenbühl, P. 2022. Learning from all vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17222–17231.

Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An open urban driving simulator. In *Conference on robot learning*, 1–16. PMLR.

Gazis, D.; Herman, R.; and Maradudin, A. 1960. The problem of the amber signal light in traffic flow. *Operations research*, 8(1): 112–132.

Hu, Y.; Yang, J.; Chen, L.; Li, K.; Sima, C.; Zhu, X.; Chai, S.; Du, S.; Lin, T.; Wang, W.; et al. 2023. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17853–17862.

Huang, Z.; Wu, J.; and Lv, C. 2022. Efficient deep reinforcement learning with imitative expert priors for autonomous driving. *IEEE Transactions on Neural Networks and Learning Systems*, 34(10): 7391–7403.

Jin, Y.; Shen, X.; Peng, H.; Liu, X.; Qin, J.; Li, J.; Xie, J.; Gao, P.; Zhou, G.; and Gong, J. 2023. Surrealdriver: Designing generative driver agent simulation framework in urban contexts based on large language model. *arXiv preprint arXiv:2309.13193*.

Lin, J.; Zhou, W.; Wang, H.; Cao, Z.; Yu, W.; Zhao, C.; Zhao, D.; Yang, D.; and Li, J. 2022. Road traffic law adaptive decision-making for self-driving vehicles. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 2034–2041. IEEE.

Papaioannou, P. 2007. Driver behaviour, dilemma zone and safety effects at urban signalised intersections in Greece. *Accident Analysis & Prevention*, 39(1): 147–158.

Reed, N.; Leiman, T.; Palade, P.; Martens, M.; and Kester, L. 2021. Ethics of automated vehicles: breaking traffic rules for road safety. *Ethics and Information Technology*, 23(4): 777–789.

Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Wang, W.; Xie, J.; Hu, C.; Zou, H.; Fan, J.; Tong, W.; Wen, Y.; Wu, S.; Deng, H.; Li, Z.; et al. 2023. Drivemlm: Aligning multi-modal large language models with behavioral planning states for autonomous driving. *arXiv preprint arXiv:2312.09245*.

Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.

Wen, L.; Fu, D.; Li, X.; Cai, X.; Ma, T.; Cai, P.; Dou, M.; Shi, B.; He, L.; and Qiao, Y. 2023. Dilu: A knowledge-driven approach to autonomous driving with large language models. *arXiv preprint arXiv:2309.16292*.

Zhang, C.; Kacem, K.; Hinz, G.; and Knoll, A. 2022. Safe and rule-aware deep reinforcement learning for autonomous driving at intersections. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 2708–2715. IEEE.