LLM-POWERED CONSENSUS FOR INTELLIGENT TRANSPORTATION SYSTEM

Anonymous authors
Paper under double-blind review

ABSTRACT

In the pursuit of Artificial General Intelligence (AGI), the development of Intelligent Transportation Systems (ITS) represents a critical milestone. While current research is largely focused on ego-centric autonomous systems, this approach is limited by the constraints of on-board sensors and local decision-making, leading to potential risks and accidents. Addressing these challenges, this paper introduces an innovative LLM-powered consensus framework for ITS. By leveraging an improved Raft algorithm and incorporating Large Language Models, we facilitate dynamic grouping, inter-vehicle information sharing, and global planning. This decentralized system enhances flexibility, resilience, and the ability to achieve global optimality, resulting in improved traffic flow and reduced collisions. The experimental results validate the effectiveness of our approach in creating a safer and more efficient intelligent transportation ecosystem.

1 INTRODUCTION

On our way to Artificial General Intelligence (AGI), intelligent transportation systems (ITS) is the intermediate barrier to be solved. Autonomous driving vehicles are the integration of environmental perception, reasonable planning, social interaction, etc. It is the epitome of what a general AI should be able to master. Currently, most research focus on the ego-only autonomous driving system, with hundreds of sensors collecting data and and helping make decisions. Despite the multitude of proposed solutions from Uber, Tesla, Waymo and Baidu Somerville et al. (2018); Tesla (2021); Gibbs (2017); Sun et al. (2020); Baidu (2021), the perceptual scope remains constrained to line-of-sight limitations, primarily attributed to the inherent drawbacks of on-board sensors. Furthermore, the functionality of devices within a vehicle predominantly contributes to localized decision-making processes. In simpler terms, neighboring vehicles may lack awareness of specific decisions made by a particular vehicle. Moreover, the susceptibility of sensors to failures introduces a significant risk factor, potentially resulting in casualties. This becomes particularly dangerous in ITS, where human lives are at stake. The potential conflicts arising from disparate local decisions or erroneous sensor readings among different vehicles can escalate into accidents. A painful illustration of this is evident in a tragic incident involving a Tesla car, wherein a sensor failure failed to detect a sizable truck and trailer crossing the highway, causing the vehicle to accelerate uncontrollably under the truck Yadron & Tynan (2016).

To address the aforementioned issues, there has been an introduction of Vehicle-to-Vehicle (V2V) networks, or a broader scope referred to as Vehicle-to-Everything (V2X) networks. These networks leverage communication infrastructures, such as cellular networks, to facilitate the exchange of information among vehicles and various infrastructural elements Chen et al. (2017); Torrent-Moreno et al. (2009); Cheng et al. (2017). Presently, numerous Internet of Things (IoT)-related systems utilize centralized architectures, exemplified by systems like autonomous vehicle communication and household appliance control. These systems necessitate connected nodes to communicate data with central controllers. However, the centralized architecture encounters challenges related to reliability and performance Sun et al. (2019). A prevalent issue with centralized systems is the susceptibility to a single point of failure (SPOF), wherein the overall system's performance is significantly dependent on the stability of the central node. In the context of addressing the challenges posed by centralized systems and facilitating Vehicle-to-Vehicle (V2V) communication for environmental sensing in autonomous driving, a promising alternative is the utilization of distributed consensus networks. This approach offers a decentralized, resilient, and trustworthy framework, effectively mitigating

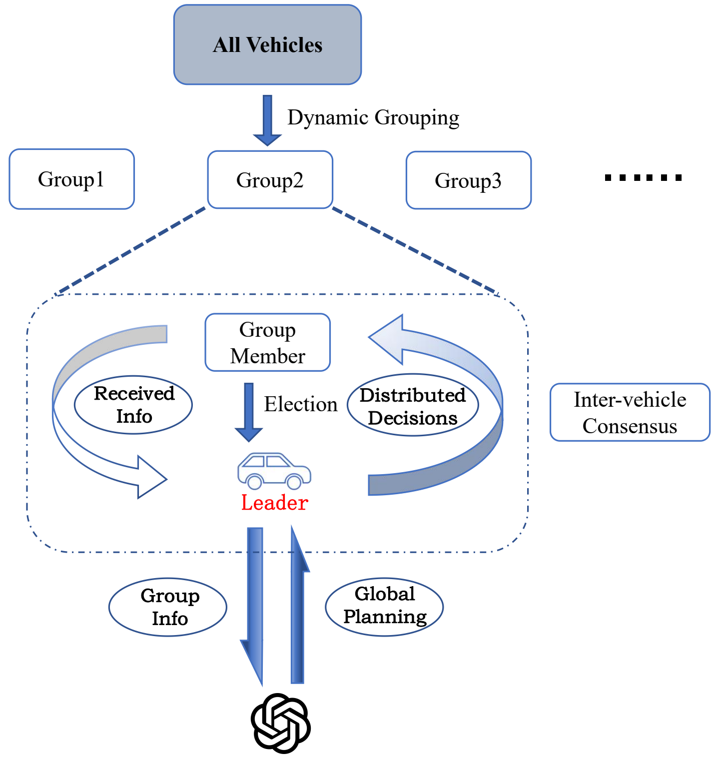

Figure 1: Overview of our proposed LLM-powered consensus system. The system consists of three parts, dynamic grouping, inter-vehicle consensus and LLM planning. All vehicles are divided into several groups with a dynamic grouping algorithm according to the vehicles' physical location. In the stage of group consensus, we realize information sharing based on Raft algorithm. With all the useful data from every vehicle in the group, LLM can plan globally for each group member.

the limitations associated with centralized systems and catering to the requirements of large-scale, less reliable communication scenarios.

Currently there are already some work focusing on distributed consensus network among ITS. Li et al. (2022) and Feng et al. (2023) leverage traditional consensus algorithm in a network, achieving data sharing among vehicles. They also implement their work on hardware. However, they only care about data security and transfer efficiency. In this paper, we go one step further. With the development of Large Language Models (LLM), more and more LLM-based planning algorithms for autonomous driving are emerging. These algorithms are more flexible and have stronger generalization ability because of the broad knowledge from LLM.

Therefore, we propose an LLM-powered consensus for intelligent transportation system, leveraging an improved Raft algorithm Ongaro & Ousterhout (2014). As illustrated in Figure1, the system consists of three parts, dynamic grouping, group consensus and LLM planning. Considering that vehicles are flowing on the road, there is no need to let two cars far apart plan routes together. Thus we divide all vehicles into several groups with a dynamic grouping algorithm according to the vehicles' physical location (see section 3.3), which makes the system more flexible and have the ability to handle various situations. In the stage of group consensus, we realize information sharing based on Raft algorithm Ongaro & Ousterhout (2014), and we also draw on Li et al. (2022) and Feng et al. (2023)'s strengths to enhance the stability of our system. With all the useful data from every vehicle in the group, LLM can make decisions for each group member.

As the experimental results show in section 4, our system can achieve global optimality. Currently most autonomous vehicles are only planning the optimal route for themselves, but we know that greedy can only reach the local optimum. That is why traffic congestion will occur. While in our system, all vehicles in the same group will share the information, and make decisions from the global view. Therefore, our system can achieve better traffic flow in the intelligent transportation system.

The main contributions of this paper are as follows:

- We propose an LLM-powered consensus for intelligent transportation system. Based on the hardware realization, we go one step further, utilizing Large Language Models to flexibly make decisions.
- We achieve global optimality in our system, resulting in better traffic flow in the intelligent transportation system. The experiments show that there are fewer collisions and less traffic congestion with our system.

The remainder of this paper is organized as follows: Section 2 investigates a literature review of consensus in autonomous systems and the use of LLM in autonomous driving planning. In section 3, we elaborate on the architecture of our system, with particular emphasis on its consensus design and the incorporation of dynamic group algorithm. Section 4 presents a detailed experimental study, meticulously designed to benchmark and validate the performance of our system. Finally, Section 5 synthesizes the key contributions of this study and casts a spotlight on potential directions for future research, as well as its broader implications for society at large.

2 Related Work

2.1 Consensus Problem

Consensus refers to the fundamental problem of achieving agreement among distributed processes or nodes in a network, even in the presence of failures or malicious actors Lamport (2019b). It plays a pivotal role in various domains, including distributed systems, blockchain technology, and fault-tolerant computing. Achieving consensus is crucial for ensuring data consistency and reliability in distributed systems. The Paxos consensus algorithm Lamport (2019a) and the Practical Byzantine Fault Tolerance (PBFT) algorithm Castro et al. (1999) are significant contributions to the field, addressing the consensus problem in practical, fault-tolerant settings. These concepts and algorithms have had a profound impact on the design and implementation of distributed systems, ensuring that they function reliably and consistently in complex, decentralized environments.

Consensus algorithms have garnered significant attention within autonomous systems and a variety of domains such as intelligent transportation systems (ITS), Cooperative Adaptive Cruise Control (CACC), Connected Autonomous Vehicles (CAV), Vehicular Ad-hoc Networks (VANET), and self-driving cars Wang et al. (2018). Extensive research has been conducted on foundational consensus approaches that primarily deal with maintaining optimal spacing between vehicles in platoons Jia et al. (2016). These approaches, which operate under various traffic dynamic models Helbing (2002), necessitate the exchange of parameters between neighboring vehicles to depict their interaction on the road and to ensure both a stable platoon formation and enhanced traffic flow. Based on known vehicle speed, acceleration, and inter-vehicle distances, necessary adjustments are made accordingly Ren (2006); Dao et al. (2008); Fernandes & Nunes (2012); Jia & Ngoduy (2016).

To enhance road efficiency, researchers have proposed advanced consensus strategies such as multi-platoon coordination]Li et al. (2019) and other generalized formations Porfiri et al. (2007). While these methods lead to convergence and stability in inter-vehicle spacing when incorporating traffic dynamics, their limited capacity for handling complex vehicle maneuvers restricts their broader usage.

Concurrently, the application of blockchain-based consensus mechanisms has emerged for fostering more comprehensive information exchange within vehicle-to-vehicle (V2V) networks Kang et al. (2019a). Numerous studies have focused on proof-based consensus methods. For instance, Kang et al. (2019b), as detailed in their work, employed distributed smart contracts based on proof-of-work (PoW) and proof-of-storage consensus. These enable data auditing and verification by incentivizing PoW nodes to collect and authenticate local metadata in exchange for vehicle coins, a specific cryptocurrency, which can be utilized for enhancing computational resources. Similarly, proof-of-storage nodes contribute storage capacity for the same rewards and subsequent upgrades. In Yang et al. (2019), the paper introduced a hybrid consensus mechanism combining PoW and proof-of-stake (PoS), followed by a trust management system for assessing data credibility through a Bayesian inference model. However, these methods are computationally intensive due to the competitive mining nature of the participating nodes.

2.2 LLM-BASED AUTONOMOUS DRIVING PLANING

Significant advancements have been made by large language models (LLMs) in demonstrating open-world cognitive and reasoning skills Radford et al. (2018; 2019); Brohan et al.; Luo et al. (2018); Achiam et al. (2023). These skills have the potential to clarify the decision-making mechanisms in autonomous driving, thereby substantially bolstering the reliability of the systems and the trust users place in them Deo et al. (2021); Kim et al. (2019); Atakishiyev et al. (2021); Jia et al. (2023); Luo et al. (2018). Regarding the application of LLMs, the field can be divided into two primary approaches: the fine-tuning of pre-trained models and the practice of prompt engineering.

In the realm of fine-tuning pretrained models for application, Liu et al. (2023) introduced MTD-GPT, which effectively converts multi-task decision-making challenges into sequence modeling issues. By training on a composite dataset involving multiple tasks, it effectively handles diverse decision-making scenarios at unregulated intersections. While this method surpasses the performance of traditional single-task reinforcement learning models, its application scope is largely confined to scenarios at unsignaled crossings, which may not fully capture the intricacies of real-world situations. Chen et al. (2023) proposed the concept of Driving with LLMs, where an architecture was developed to integrate vector-based inputs into LLMs through a two-phase pretraining and adjustment process. However, due to the constraints of vectorized input representations, the approach was exclusively evaluated in a simulated environment. Xu et al. (2023) presented DriveGPT4, a multimodal LLM based on the Valley framework Luo et al. (2023), and they curated a dataset for visual instruction tuning to enhance interpretability in autonomous driving. In addition to predicting fundamental control commands for vehicles, DriveGPT4 provides real-time explanations for its actions. While it shows superior performance in various question-answering tasks, the planning experiments conducted were relatively straightforward.

From the viewpoint of prompt engineering, various approaches have attempted to harness the deep reasoning capabilities of LLMs through the creation of intelligent prompts. Wen et al. (2023a) developed a framework named DiLu, which treats LLMs as agents for closed-loop driving missions. This approach incorporates a memory module to capture experiences, thus enhancing the reasoning and reflective abilities of LLMs. DiLu demonstrates remarkable generalization when contrasted with state-of-the-art reinforcement learning-based methods. Nonetheless, its reasoning and reflection involve numerous cycles of question-and-answer, leading to a notable increase in inference time. In a similar vein, Cui et al. (2023) and Cui et al. (2024) integrated LLMs' language and reasoning skills into self-driving cars with approaches such as Receive Reason and React and Drive as You Speak. In addition to memory and reflection, these methods incorporate raw sensor data from cameras, GNSS, lidar, and radar. Nonetheless, they still grapple with the challenge of slow inference speed. Moreover, Keysan et al. (2023) introduced SurrealDriver, which structures the memory module into short-term memory, long-term guidelines, and safety criteria. The study also conducted interviews with 24 drivers, utilizing their detailed driving behavior descriptions as chain-of-thought prompts to create a 'coach agent' module.

3 METHOD

3.1 AUTONOMOUS AGENTS BY LLM

The advent of Large Language Models (LLMs) has opened new frontiers in the field of autonomous driving, bringing closer the vision of vehicles that not only navigate and respond to their environment with high precision but also understand and interpret complex scenarios in a human-like manner. This section delves into the role of LLMs as autonomous agents in driving systems, highlighting their potential to revolutionize the way vehicles make decisions, plan routes, and interact with both their passengers and surroundings.

Our code is based on the LimSim environment Wen et al. (2023b), with the addition of LLM as the ultimate decision-making terminal. We assign specific NPC roles to each vehicle and through fine-tuning, conduct multiple rounds of agent verification for each vehicle. LLM serves not only as the core of decision-making but also enhances the adaptability and granularity of the simulation. By integrating LLM, the simulator is capable of considering a range of dynamic variables from psychological attributes and adherence to social norms to sudden incidents. We leverage the powerful capabilities of LLM to describe the driving state in natural language, thereby improving the transparency and interpretability of the decision-making process. Employing Large Language Models (LLM) as drivers simulates the complexity and diversity of human driving behavior. Through high-fidelity simulation environments and precise vehicle dynamics, it enhances the evaluation and refinement of autonomous driving algorithms. By adopting large language models as autonomous agents, the framework offers significant improvements and high adaptability in simulating the complexity and diversity of human driving behavior. This method not only enhances the fidelity and analytical rigor of autonomous driving simulations but also pushes the boundaries of our ability to replicate real-world driving scenarios. The innovativeness and effectiveness of LLM as autonomous agents in enhancing the realism and interactivity of simulators.

3.2 GROUP CONSENSUS

This section details the methodology for achieving information consensus within crowded urban driving scenarios. Drawing on the Raft protocol Ongaro & Ousterhout (2014), we demonstrate how driving situations are integrated with computer systems. Initially, our primary objective mirrors that of enabling a collection of machines to function as a cohesive unit, capable of withstanding failures among its constituents. We classify these machines into a collective that includes diverse vehicle types. The method of categorizing these groups will not be discussed at this stage. Within this group, each vehicle is considered a node, with three possible states corresponding to the three states in the Raft algorithm: follower, leader, and candidate. Their roles align with those defined in the Raft protocol. In the absence of a leader within the group, vehicles engage in a dialogue to elect a new leader, facilitating the selection of vehicle roles. The election aims to achieve a quasi-centralized control within the group. The election process follows the principles of the Raft algorithm, whereby each vehicle is assigned a random number. Vehicles then compare these numbers: the one with a smaller number becomes a follower, and this comparison process continues until all vehicles are followers except for the leader.

Upon selecting a leader, it is crucial to maintain communication among vehicles. This involves keeping a heartbeat signal active. If a vehicle fails to send a heartbeat to the leader within a specified timeframe, it is considered offline and is removed from the current group to join another. Likewise, the leader must send heartbeat signals. If it fails to do so in time, it is considered offline, prompting a re-election. In this case, every follower automatically becomes a candidate, and the selection process is repeated.

After completing these election tasks, we can ensure that the information within the group are always synchronized, and a leader is present at any given moment. Next, we need to assign tasks to the leader. The premise is that continuous communication within the group is mandatory; otherwise, vehicles will be removed from the group. Since communication within the group is ensured, followers can send their location information to the leader, including essential data such as position, speed, acceleration, vehicle length, and width. Thus, the leader obtains information about all vehicles in the group. The leader then compiles this information, including their own, into a package and sends it to the LLM, or Large Language Model. The model analyzes the received information and makes comprehensive judgments, subsequently providing action decisions for each vehicle, including acceleration and turning information. These decisions are also sent to the leader. It's worth noting that we assume the decision-making process, including reading and sending information, is instantaneous, meaning the environmental time is considered stationary at that moment. Similarly to the Raft algorithm, if the leader goes offline at any point, a log system is employed to ensure that every node, or vehicle, can access the latest decision information.

Once the leader receives the decision information, they distribute it according to the ID of each vehicle within the group, ensuring every vehicle receives the instructions for their next actions. This process completes a full cycle of consensus communication.

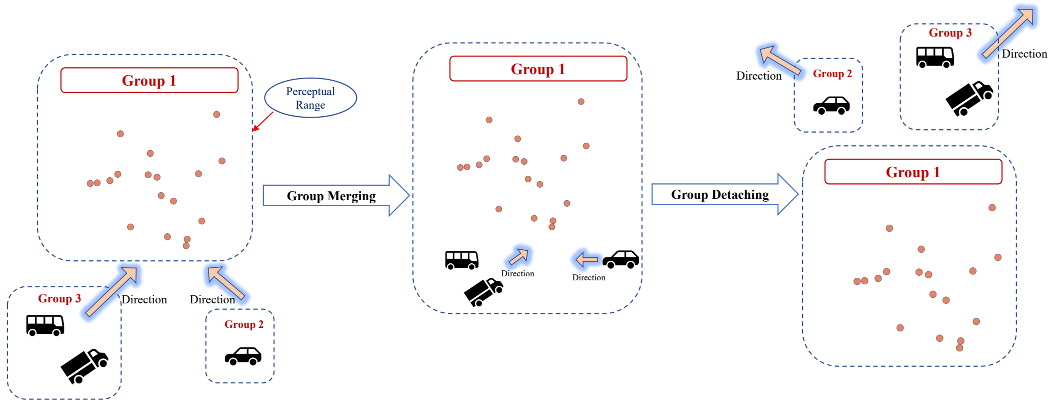

Figure 2: Overview of Dynamic Group algorithm. If a vehicle is alone, it forms its own group; upon encountering other vehicles, it integrates with them. When groups meet, vehicles can merge groups, with the smaller group integrating into the larger one, eventually forming a larger group. If vehicles separate, they can also detach from the group.

3.3 Dynamic Group

We've explained how to perform Raft consensus operations within a group to synchronize vehicle information. However, grouping all vehicles in an environment into a single group would lead to centralization, imposing an unsustainable operational burden. Therefore, it's necessary to divide vehicles into groups to achieve a decentralized approach. But how do we categorize these groups? We considered various methods, including static grouping based on the environment, random grouping by assigning vehicles to random groups, and dynamic grouping where each vehicle communicates with others within a certain distance. A vehicle joins a group if it's not already in one or if the new group is larger; if there are no nearby vehicles, it forms a group on its own. Ultimately, we opted for dynamic grouping due to several reasons.

Our experiments showed that static grouping faced significant inter-group communication issues, with delays and inconvenience when vehicles moved across groups, increasing the communication load within groups. Random grouping yielded average communication outcomes with certain random errors and exceptional cases. Therefore, dynamic grouping was chosen.

In dynamic grouping, vehicles must constantly communicate externally. If a vehicle is alone, it forms its own group; upon encountering other vehicles, it integrates with them. When groups meet, vehicles can merge groups, with the smaller group integrating into the larger one, eventually forming a larger group. If vehicles separate, they can also detach from the group.

This approach allows for dynamic grouping, followed by Raft consensus operations within each group to synchronize vehicle information. The LLM then returns decision-making for each group separately, achieving synchronization of information for all vehicles within the environment.

In summary, the operation of the entire system, as illustrated in Figure 2, employs dynamic grouping to categorize vehicles into several groups. Within this system, vehicles automatically merge into or detach from groups, ensuring that adjacent vehicles are within the same group. Within a group, vehicles automatically elect a leader. The leader collects information sent by each vehicle and then forwards this data to the LLM. The LLM makes decisions based on this information, and the leader disseminates the returned decisions back to each vehicle, maintaining information synchronization within each group member. If a leader goes offline, a new election commences to ensure continuous synchronization of information within the group. This process allows the LLM to make comprehensive and efficient decision simulations by always maintaining up-to-date and synchronized data across the groups.

4 EXPERIMENT

4.1 EXPERIMENTAL SETUP

To rigorously validate our consensus verification problem, we established a variety of driving environments, encapsulating the diverse scenarios vehicles encounter during normal driving conditions. These environments included:

- **Normal Road:** Sparse traffic conditions with few vehicles on the road
- **Crowded Road:** Busy roads with heavy vehicle traffic
- **Highway:** High-speed highways where vehicles travel at high speeds
- **Intersection:** Complex intersections with intricate vehicular movements

These four settings comprehensively represent the spectrum of challenges vehicles face in real-world driving situations. Within these environments, we conducted experiments under three distinct configurations: with consensus, with partial consensus (part of consensus), and without consensus. For the configurations involving with and without consensus, vehicles were driven five times in each road setting. We recorded the completion rates and the number of collisions, averaging the driving speeds to assess performance.For the partial consensus configuration, we undertook three ablation experiments, each omitting one of the three components of the consensus mechanism. Each scenario was tested by driving five times across the five road settings. The results from these three conditions were aggregated and averaged to calculate the overall impact of each component on the consensus mechanism's effectiveness.

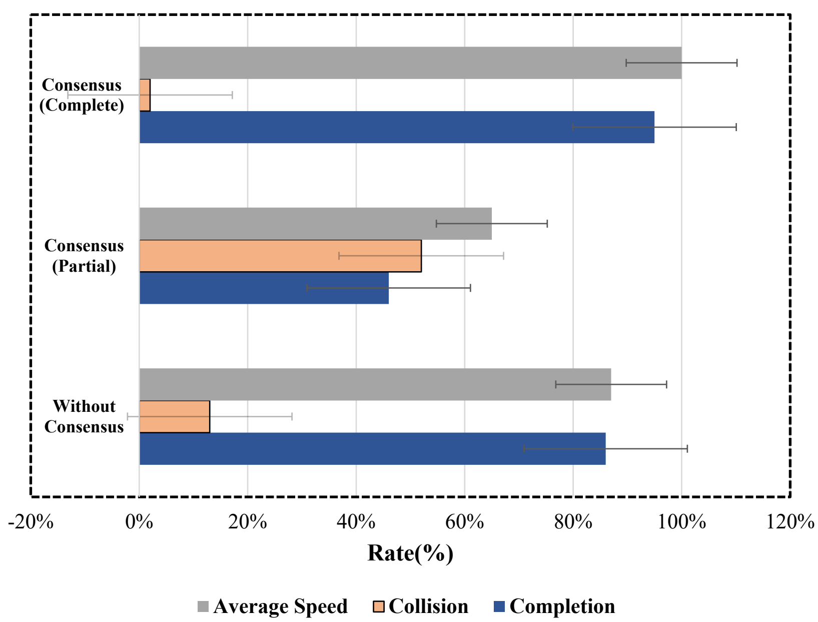

Figure 3: The results of the experiments under three distinct configurations: with consensus, with partial consensus (part of consensus), and without consensus.

4.2 GLOBAL OPTIMALITY

Initially, we needed to validate the global optimality of our method, which means we had to experimentally verify that our method could effectively optimize vehicle operation under different circumstances compared to scenarios where the method is not used. We utilized a multi-vehicle intersection environment with numerous vehicles, where our egocar, the vehicle focused on by the camera (which is not the leader), needed to move from one multi-vehicle node to another. The results in Table 1 showed that using the Raft consensus method could effectively avoid vehicle collisions and congestion, allowing vehicles to smoothly reach their final destinations while ensuring high-speed and stable operation. In contrast, vehicles not utilizing the consensus method experienced sudden changes in speed or emergency braking during operation due to not receiving relevant information in a timely manner. Conversely, the consensus allows vehicles to promptly obtain relevant information and make superior decisions. Additionally, the consensus enables vehicles to demonstrate more exceptional driving performance, resulting in smoother and more stable driving.

Table 1: Comparison of Different Consensus Usage Scenarios

	Consensus (Full)	Consensus (Partial)	Without Consensus
Completion Status	High	Medium	Low
Collision Probability	Low	Medium	High
Running Speed	High	Low	Medium
Lane Change	Low	High	High

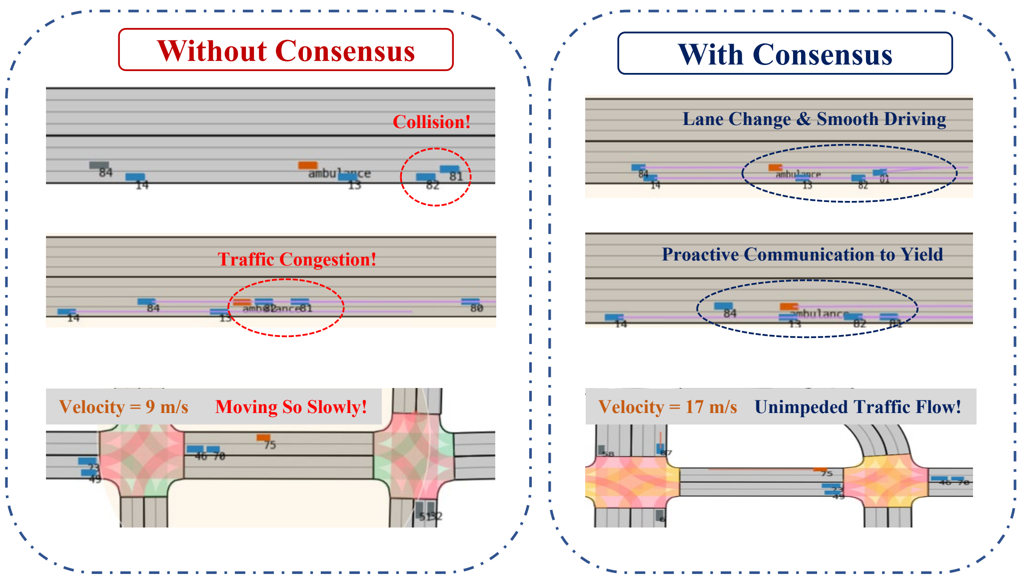

Figure 4: The comparison between situations with and without consensus. With consensus, the traffic is smoother and faster, and there are fewer collisions.

4.3 Ablation Study

Next, we verified the completeness of our method, meaning we needed to confirm that no steps in our method were redundant. We tested the code with certain steps omitted, including the absence of intra-group selection (random leader within the group, free information transmission), lack of intra-group synchronization (missing timely information synchronization within the group), and absence of group merging (maintaining consistent group identities). The results in Table 2 showed that in environments characterized by these three scenarios, when facing multiple groups and multiple vehicles, collisions, interference, or various problems would occur, such as collisions at intersections or during lane changes. However, the complete code did not encounter similar issues under the same conditions. Therefore, we ultimately demonstrated that each step of the Raft consensus operations is indispensable, with no room for omissions, to achieve the final synchronization of decision information.

Table 2: Comparison of consensus experimental results with different missing components

	Intersection	**Crowded Road**	**Highway**
Lack of Intra-Group Selection	Collision	Collision	Congestion
Lack of Intra-Group Merging	Collision	Congestion	Collision
Lack of Intra-Group Synchronization	Congestion	Collision	Collision

For individual vehicles or in scenarios where vehicles were sparse and scattered, the consensus method struggled to make an impact, performing similarly to vehicles without consensus, which also makes sense. Therefore, it appears that consensus offers unique advantages for multi-vehicle or complex environmental scenarios, demonstrating superior functionality, primarily in terms of vehicle driving stability and communication coordination. Compared to non-consensus vehicles, the Raft consensus method performs better in these two aspects.

5 CONCLUSION

This study has made significant strides in addressing the critical challenges faced by intelligent transportation systems. By introducing an LLM-powered consensus framework, we have enhanced the decision-making capabilities of autonomous vehicles, moving beyond the constraints of ego-only systems and localized planning. The integration of Large Language Models into the consensus process not only enriches the decision-making knowledge base but also increases flexibility and generalization in route planning. Our system's dynamic grouping and consensus mechanisms ensure that vehicles can effectively share information, leading to globally optimal decisions and improved traffic flow. The experimental results validate the efficacy of our approach in reducing collisions and traffic congestion, thereby offering a safer and more efficient transportation experience.

Looking ahead, this research opens up several potential directions for future studies. Firstly, the scalability and robustness of the LLM-powered consensus in larger and more diverse traffic scenarios need to be investigated. Secondly, the integration of additional context-awareness and real-time data into the decision-making process can further refine the performance of autonomous vehicles. Moreover, research into the ethical dimensions of decision-making algorithms in ITS is necessary to ensure that these systems align with societal values.

The broader implications of this work extend beyond the realm of transportation. The consensus mechanisms and LLM-based planning algorithms developed here can be applied to a wide range of Internet of Things (IoT) systems, enhancing the performance and reliability of decentralized networks. The success of our system in achieving global optimality also underscores the transformative potential of AI in solving complex societal challenges, paving the way to Artificial General Intelligence (AGI).

In summary, this study not only contributes to the advancement of autonomous driving technology but also provides a model for how AI can be leveraged to create more intelligent, efficient, and safe systems. As we continue to navigate the path towards Artificial General Intelligence, the lessons learned from developing ITS will undoubtedly play a crucial role in shaping the future of technology and society.

Author Contributions

If you'd like to, you may include a section for author contributions as is done in many journals. This is optional and at the discretion of the authors.

Acknowledgments

Use unnumbered third level headings for the acknowledgments. All acknowledgments, including those to funding agencies, go at the end of the paper.

References

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.

Shahin Atakishiyev, Mohammad Salameh, Hengshuai Yao, and Randy Goebel. Explainable artificial intelligence for autonomous driving: A comprehensive overview and field guide for future research directions. *arXiv preprint arXiv:2112.11561*, 2021.

Baidu. Apolloauto. `https://github.com/ApolloAuto/apollo`, 2021. [Online].

Anthony Brohan, Noah Brown, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, MontseGonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-WeiEdward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspiar Singh, Anikait Singh, Radu Soricut, Huong Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-2: Vision-language-action models transfer web knowledge to robotic control.

Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OsDI*, volume 99, pp. 173–186, 1999.

Long Chen, Oleg Sinavski, Jan Hünermann, Alice Karnsund, Andrew James Willmott, Danny Birch, Daniel Maund, and Jamie Shotton. Driving with llms: Fusing object-level vector modality for explainable autonomous driving. *arXiv preprint arXiv:2310.01957*, 2023.

Shanzhi Chen, Jinling Hu, Yan Shi, Ying Peng, Jiayi Fang, Rui Zhao, and Li Zhao. Vehicle-to-everything (v2x) services supported by lte-based systems and 5g. *IEEE Communications Standards Magazine*, pp. 70–76, Jan 2017. doi: 10.1109/mcomstd.2017.1700015. URL `http://dx.doi.org/10.1109/mcomstd.2017.1700015`.

Xiang Cheng, Chen Chen, Wuxiong Zhang, and Yang Yang. 5g-enabled cooperative intelligent vehicular (5genciv) framework: When benz meets marconi. *IEEE Intelligent Systems*, 32(3): 53–59, May 2017. doi: 10.1109/mis.2017.53. URL `http://dx.doi.org/10.1109/mis.2017.53`.

Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, and Ziran Wang. Receive, reason, and react: Drive as you say with large language models in autonomous vehicles. *arXiv preprint arXiv:2310.08034*, 2023.

Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, and Ziran Wang. Drive as you speak: Enabling human-like interaction with large language models in autonomous vehicles. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 902–909, 2024.

Thanh-Son Dao, Christopher Michael Clark, and Jan Paul Huissoon. Distributed platoon assignment and lane selection for traffic flow optimization. In *2008 IEEE Intelligent Vehicles Symposium*, Jun 2008. doi: 10.1109/ivs.2008.4621202. URL `http://dx.doi.org/10.1109/ivs.2008.4621202`.

Nachiket Deo, EricM. Wolff, and Oscar Beijbom. Multimodal trajectory prediction conditioned on lane-graph traversals. *arXiv: Computer Vision and Pattern Recognition,arXiv: Computer Vision and Pattern Recognition*, Jun 2021.

Chenglin Feng, Zhangchen Xu, Xincheng Zhu, Paulo Valente Klaine, and Lei Zhang. Wireless distributed consensus in vehicle to vehicle networks for autonomous driving. *IEEE Transactions on Vehicular Technology*, pp. 8061–8073, Jun 2023. doi: 10.1109/tvt.2023.3243995. URL `http://dx.doi.org/10.1109/tvt.2023.3243995`.

Pedro Fernandes and Urbano Nunes. Platooning with ivc-enabled autonomous vehicles: Strategies to mitigate communication delays, improve safety and traffic flow. *IEEE Transactions on Intelligent Transportation Systems*, 13(1):91–106, Mar 2012. doi: 10.1109/tits.2011.2179936. URL `http://dx.doi.org/10.1109/tits.2011.2179936`.

Samuel Gibbs. Google sibling waymo launches fully autonomous ride-hailing service. *The Guardian*, 7, 2017.

Dirk Helbing. Traffic and related self-driven many-particle systems. *Reviews of Modern Physics*, pp. 1067–1141, Jul 2002. doi: 10.1103/revmodphys.73.1067. URL `http://dx.doi.org/10.1103/revmodphys.73.1067`.

Dongyao Jia and Dong Ngoduy. Enhanced cooperative car-following traffic model with the combination of v2v and v2i communication. *Transportation Research Part B: Methodological*, pp. 172–191, Aug 2016. doi: 10.1016/j.trb.2016.03.008. URL `http://dx.doi.org/10.1016/j.trb.2016.03.008`.

Dongyao Jia, Kejie Lu, Jianping Wang, Xiang Zhang, and Xuemin Shen. A survey on platoon-based vehicular cyber-physical systems. *IEEE Communications Surveys amp; Tutorials*, pp. 263–284, Jan 2016. doi: 10.1109/comst.2015.2410831. URL `http://dx.doi.org/10.1109/comst.2015.2410831`.

Xiaosong Jia, Yulu Gao, Li Chen, Junchi Yan, PatrickLangechuan Liu, and Hongyang Li. Driveadapter: Breaking the coupling barrier of perception and planning in end-to-end autonomous driving. Aug 2023.

Jiawen Kang, Zehui Xiong, Dusit Niyato, Dongdong Ye, Dong In Kim, and Jun Zhao. Toward secure blockchain-enabled internet of vehicles: Optimizing consensus management using reputation and contract theory. *IEEE Transactions on Vehicular Technology*, pp. 2906–2920, Mar 2019a. doi: 10.1109/tvt.2019.2894944. URL `http://dx.doi.org/10.1109/tvt.2019.2894944`.

Jiawen Kang, Rong Yu, Xumin Huang, Maoqiang Wu, Sabita Maharjan, Shengli Xie, and Yan Zhang. Blockchain for secure and efficient data sharing in vehicular edge computing and networks. *IEEE Internet of Things Journal*, pp. 4660–4670, Jun 2019b. doi: 10.1109/jiot.2018.2875542. URL `http://dx.doi.org/10.1109/jiot.2018.2875542`.

Ali Keysan, Andreas Look, Eitan Kosman, Gonca Gürsun, Jörg Wagner, Yao Yu, and Barbara Rakitsch. Can you text what is happening? integrating pre-trained language encoders into trajectory prediction models for autonomous driving. Sep 2023.

Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, and John Canny. Grounding human-to-vehicle advice for self-driving vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10591–10599, 2019.

Leslie Lamport. The part-time parliament. In *Concurrency: the Works of Leslie Lamport*, pp. 277–317. 2019a.

Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. In *Concurrency: the Works of Leslie Lamport*, pp. 179–196. 2019b.

Yongfu Li, Chuancong Tang, Kezhi Li, Xiaozheng He, Srinivas Peeta, and Yibing Wang. Consensus-based cooperative control for multi-platoon under the connected vehicles environment. *IEEE Transactions on Intelligent Transportation Systems*, 20(6):2220–2229, Jun 2019. doi: 10.1109/tits.2018.2865575. URL `http://dx.doi.org/10.1109/tits.2018.2865575`.

Zongyao Li, Lei Zhang, Xiaoshuai Zhang, and Muhammad Imran. Design and implementation of a raft based wireless consensus system for autonomous driving. In *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*, Dec 2022. doi: 10.1109/globecom48099.2022.10000789. URL http://dx.doi.org/10.1109/globecom48099.2022.10000789.

Jiaqi Liu, Peng Hang, Jianqiang Wang, Jian Sun, et al. Mtd-gpt: A multi-task decision-making gpt model for autonomous driving at unsignalized intersections. *arXiv preprint arXiv:2307.16118*, 2023.

Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Minghui Qiu, Pengcheng Lu, Tao Wang, and Zhongyu Wei. Valley: Video assistant with large language model enhanced ability. *arXiv preprint arXiv:2306.07207*, 2023.

Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, Jun 2018. doi: 10.1109/cvpr.2018.00376. URL http://dx.doi.org/10.1109/cvpr.2018.00376.

Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX annual technical conference (USENIX ATC 14)*, pp. 305–319, 2014.

Maurizio Porfiri, D. Gray Roberson, and Daniel J. Stilwell. Tracking and formation control of multiple autonomous agents: A two-level consensus approach. *Automatica*, 43(8):1318–1328, Aug 2007. doi: 10.1016/j.automatica.2007.01.004. URL http://dx.doi.org/10.1016/j.automatica.2007.01.004.

Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

Wei Ren. Consensus based formation control strategies for multi-vehicle systems. In *2006 American Control Conference*, Jan 2006. doi: 10.1109/acc.2006.1657384. URL http://dx.doi.org/10.1109/acc.2006.1657384.

Heather Somerville, Paul Lienert, and Alexandria Sage. Uber's use of fewer safety sensors prompts questions after arizona crash. *Business news, Reuters*, 2018.

Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2446–2454, 2020.

Yao Sun, Lei Zhang, Gang Feng, Bowen Yang, Bin Cao, and Muhammad Ali Imran. Blockchain-enabled wireless internet of things: Performance analysis and optimal communication node deployment. *IEEE Internet of Things Journal*, pp. 5791–5802, Jun 2019. doi: 10.1109/jiot.2019.2905743. URL http://dx.doi.org/10.1109/jiot.2019.2905743.

Tesla. Full self-driving capability, 2021. URL https://www.tesla.com/autopilot.

M. Torrent-Moreno, J. Mittag, P. Santi, and H. Hartenstein. Vehicle-to-vehicle communication: Fair transmit power control for safety-critical information. *IEEE Transactions on Vehicular Technology*, 58(7):3684–3703, Sep 2009. doi: 10.1109/tvt.2009.2017545. URL http://dx.doi.org/10.1109/tvt.2009.2017545.

Ziran Wang, Guoyuan Wu, and Matthew Barth. A review on cooperative adaptive cruise control (cacc) systems: Architectures, controls, and applications. *Cornell University - arXiv,Cornell University - arXiv*, Sep 2018.

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*, 2023a.

Licheng Wen, Daocheng Fu, Song Mao, Pinlong Cai, Min Dou, Yikang Li, and Yu Qiao. Limsim: A long-term interactive multi-scenario traffic simulator, 2023b.

Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kenneth KY Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv preprint arXiv:2310.01412*, 2023.

Danny Yadron and Dan Tynan. Tesla driver dies in first fatal crash while using autopilot mode. *the Guardian*, 1, 2016.

Zhe Yang, Kan Yang, Lei Lei, Kan Zheng, and Victor C. M. Leung. Blockchain-based decentralized trust management in vehicular networks. *IEEE Internet of Things Journal*, pp. 1495–1505, Apr 2019. doi: 10.1109/jiot.2018.2836144. URL `http://dx.doi.org/10.1109/jiot.2018.2836144`.

A Appendix

You may include other additional sections here.