# Contextual-Personalized Adaptive Cruise Control via Fine-Tuned Large Language Models

Anonymous ICCV submission

Paper ID *****

## Abstract

*Adaptive cruise control (ACC) is a widely adopted technique within advanced driver assistance systems (ADAS) to alleviate driver workload and fatigue in long-distance driving or stop-and-go traffic scenarios. However, conventional ACC systems typically fail to account for drivers' preferences or changing environmental conditions, limiting their adaptability in adjusting headway. To bridge this gap, this study introduces a novel contextual-personalized ACC (CP-ACC) framework that leverages the contextual reasoning and adaptive customization potential of large language models (LLMs). Specifically, LLMs including LLaMA-3-8B and Mistral-7B are fine-tuned with a synthetically generated dataset encompassing diverse drivers' preferences (e.g., energy efficiency, comfort) and real-time contextual information (e.g., weather, traffic conditions). CP-ACC demonstrates the ability to identify, quantify, and balance competing objectives (e.g., safety, mobility) compared to linear feedback ACC and the intelligent driver model (IDM). Supervised fine-tuning (SFT) further enhances the LLMs' ability to recognize driving objectives and generate safe, context-aware longitudinal control commands, outperforming zero-shot and few-shot prompting. Overall, the proposed CP-ACC framework presents a promising direction for delivering smart, adaptive, and personalized driving assistance tailored to varying drivers' preferences and dynamic traffic environments.*

## 1. Introduction

The rapid advancement of autonomous driving technologies has driven an increased reliance on advanced driver assistance systems (ADAS), with adaptive cruise control (ACC) playing a key role in enabling automatic longitudinal control to maintain safe inter-vehicle headways [35, 36]. Note that real-world driving is characterized by diverse individual preferences, such as prioritizing "safety", "comfort", "mobility", or "energy efficiency", necessitating personalized and context-sensitive system responses [39]. However, existing ACC systems exhibit limitations in integrating two critical dimensions essential for informed decision-making: (i) the diversity of drivers' preferences, and (ii) the dynamically evolving surrounding environment. This motivates a research question: *How can ACC systems be designed to interpret and quantify these dimensions, thus providing smart and customized driving assistance?*

Large language models (LLMs) have emerged as a transformative tool to deliver personalized recommendations, driven by their sophisticated semantic reasoning, contextual synthesis, and knowledge transfer abilities [33]. Their superior performance on tasks such as summarization, inference, and question answering significantly streamlines the generation of high-quality advice or content [38]. To date, LLM-related techniques have been applied across various intelligent transportation systems (ITS) domains. For instance, chain-of-thought (CoT) prompting has been employed to understand real-world traffic scenarios, which mimics human reasoning by decomposing complex tasks into logical stepwise sequences [5]. Fine-tuning refers to the continued training of a pre-trained LLM on task-specific data. It has demonstrated effectiveness in traffic detection and generation tasks [4]. The prompting-reasoning-finetuning framework has been proposed for motion planning in autonomous vehicles [20]. Moreover, retrieval-augmented generation (RAG) extends LLM input with externally retrieved up-to-date knowledge and has been leveraged to develop personalized warning systems [32].

Synthesizing the potential of LLMs for customized mobility services through their task-adaptive capabilities and addressing the unmet need for ACC systems to adapt to diverse travel scenarios, this study introduces a novel contextual-personalized ACC (CP-ACC) framework. The proposed CP-ACC leverages supervised fine-tuning (SFT) to enable scenario-specific automatic longitudinal control to improve user satisfaction. The main contributions of this study are bi-folded: (i) A question and answer (Q&A) is designed to fine-tune LLMs for generating CP-ACC outputs, and (ii) Implemented with fine-tuned LLMs, the CP-ACC framework quantifies driver preferences, contextual information, and vehicle states to provide ACC service that manages multiple driving objectives, such as "safety", "comfort", "mobility", and "energy efficiency" in a balanced and adaptive manner.

# 2. Related work

## 2.1. LLMs for Autonomous Driving

In the domain of autonomous driving, LLMs serve as sophisticated reasoning engines capable of interpreting driving objectives, user preferences, and nuanced contextual information that are often difficult to encode by traditional rule-based approaches [2]. Building on recent advances, studies have investigated the use of LLMs to derive motion planning strategies or control commands directly from sensor data and/or natural language directives. For example, *DriveGPT4* processes multimodal inputs to generate driving decisions along with interpretable reasoning, thereby enhancing system transparency and user trust [31]. *Talk2Drive* was developed to translate high-level driving instructions into executable vehicle actions, effectively linking human language with vehicle behavior [3]. Fu et al. employed an LLM (GPT-3.5) to interpret the driving environment in a human-like manner, particularly in complex situations such as long-tail corner cases [9]. Within cooperative driving automation (CDA), LLMs have shown considerable promise for addressing key challenges such as inferring driver intent, predicting trajectories, and coordinating complex multi-agent interactions [8, 30]. Yang et al. developed the language-to-trajectory dataset to support the training of *Traj-LLM*, which predicts multi-vehicle trajectories based on textual descriptions of vehicle interactions [34]. Jiang et al. proposed *KoMA* for cooperative decision-making among autonomous vehicles by understanding the intentions of surrounding vehicles, using LLM agents to mimic human cognition [14].

## 2.2. Personalized ADAS

Personalization in ADAS has attracted substantial attention recently, as customizing driving assistance for drivers' preferences markedly improves user comfort, trust, and system acceptance [19]. Studies have shown that adaptive systems can capture and respond to heterogeneity in driving behavior and style, establishing the foundation for truly intelligent and human-centered autonomous driving technologies [11, 26]. Liao et al. introduced a digital twin framework to replicate individual driving patterns for real-time lane-change prediction [18]. Li et al. designed a graph neural network-based ramp-merging trajectory predictor that incorporates driver-specific nodes, achieving an 11.4% improvement over non-personalized baselines [17]. Li et al. [16] encoded driver-specific preferences into motion planning, thereby aligning vehicle responses with individual expectations. Addressing personalized ACC, Wang et al. employed Gaussian Processes to mimic personalized car-following behavior [29]. Zhu et al. developed a Kullback-Leibler (KL) divergence-based clustering algorithm to categorize driving styles and provide personalized ACC services accordingly [41]. Furthermore, inverse reinforcement learning has emerged as a powerful method for identifying individual driving preferences to enable personalized ACC systems [23, 39].

Leveraging recent advancements in foundation models, such as LLMs, vision-language models (VLMs), and vision-language-action (VLA) frameworks, personalized driving systems now possess enhanced abilities to adapt to drivers' preferences via natural language interaction and multimodal reasoning. *TravelPlanner+* tailors travel plans using detailed user profiles[25], and *LLM-PDA* delivers real-time driving suggestions based on behavioral risk profiles [32]. Cui et al. developed an on-board VLM system that learns user preferences via a retrieval-augmented memory, reducing takeover rates by 65.2% [3]. *PADriver* utilizes a multi-modal LLM to enable personalized autonomous driving, incorporating a risk-aware assessment of potential actions [15].

Despite these advances, the integration of LLMs into autonomous driving and personalized ADAS faces two key limitations. First, most personalized control methods rely on classical modeling approaches that are limited in interpreting complex contextual and preference-based nuances. Second, comprehensive research on the specific application of LLMs in ACC, particularly in scenarios requiring customization based on drivers' preferences and dynamic traffic environments, remains scarce. This study aims to bridge these research gaps by introducing the CP-ACC framework, which uniquely synergizes the contextual reasoning capabilities of LLMs with the real-time operational requirements of ACC, offering a scalable and adaptable paradigm for smart and human-centered driving assistance.

# 3. Methodology

The proposed CP-ACC framework is outlined in Fig. 1. The process begins with the construction of a supervised fine-tuning (SFT) dataset composed of question-and-answer (Q&A) pairs that integrate user profiles and real-time vehicle states across specific driving scenarios. Details of this dataset are provided in Sec. 3.1. Pre-trained LLMs are then fine-tuned using this dataset to develop a task-specific model tailored for CP-ACC applications, as discussed in Sec. 3.2. When prompted with inputs resembling those in the fine-tuning dataset, the CP-ACC model interprets natural language queries and generates longitudinal speed commands by effectively balancing multiple driving objectives, namely, "safety", "comfort", "mobility", and "energy effi-

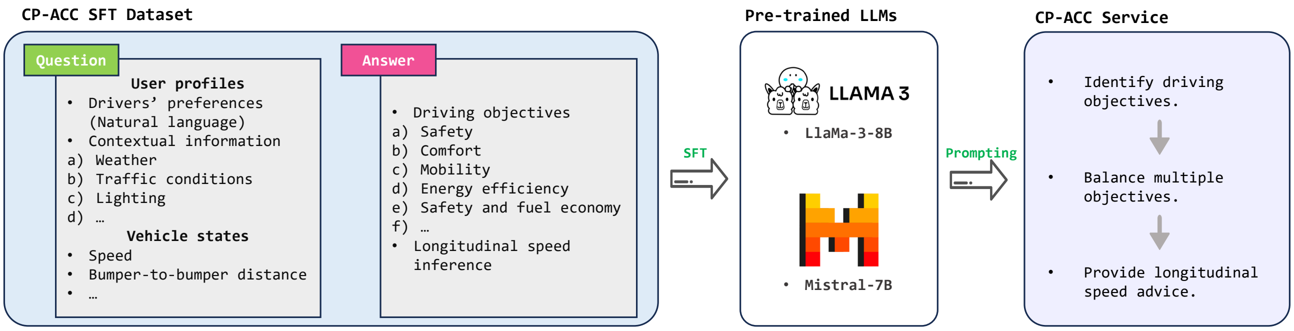


Figure 1. An Overview of the CP-ACC Framework.

ciency".

### 3.1. CP-ACC SFT Dataset Generation

A Q&A dataset is developed under the CP-ACC settings to fine-tune the LLMs, with an example illustrated in Fig. 2. Each question in the dataset includes both user profiles and vehicle states. The user profiles capture drivers' preferences and contextual information such as traffic conditions, lighting, and weather. The vehicle states consist of the ego vehicle's position $x_j(t)$, speed $v_j(t)$, and previous acceleration $a_j(t-1)$, as well as the preceding vehicle's speed $v_{j-1}(t)$ and the bumper-to-bumper distance $s(t)$ between the two vehicles, as defined in Eq. (1).

$$s(t) = x_{j-1}(t) - x_j(t) - L_{j-1}, \tag{1}$$

where $x_{j-1}(t)$ and $L_{j-1}$ are the position and length of the preceding vehicle, respectively.

With regard to the answer, CP-ACC first identifies one or more driving objectives, including "safety", "comfort", "mobility", and "energy efficiency". For example,

- **Please improve overall fuel efficiency and sustainability."** is categorized under **"energy efficiency"**.
- **It is snowing, keep a large following distance."** is categorized under **"safety"**.
- **I care about energy usage but need to maintain a good pace."** addresses both **"energy efficiency"** and **"mobility"**.
- **Maximize fuel savings and time efficiency with a smooth ride."** concerns **"energy efficiency"**, **"mobility"**, and **"comfort"**.

To accommodate diverse driving objectives, the CP-ACC framework updates the ego vehicle's speed based on the standards illustrated in Fig. 3. It combines default settings with objective-specific settings derived from expert knowledge. Specifically, time-to-collision (TTC) serves as a primary metric for safety and efficiency evaluation, defined in Eq. (2) [24]. The thresholds include $\text{TTC}(t) = 4.0$ s for "safety" [12], $|\text{jerk}| \leq 1.5$ m/s$^3$ for "comfort" [27], $\text{TTC}(t) = 2.0$ s for "mobility" [22], and a target speed of $v_{\text{target}} = 20$ m/s for "energy efficiency" [28]. For multi-objective cases, these settings are integrated accordingly to balance competing priorities.

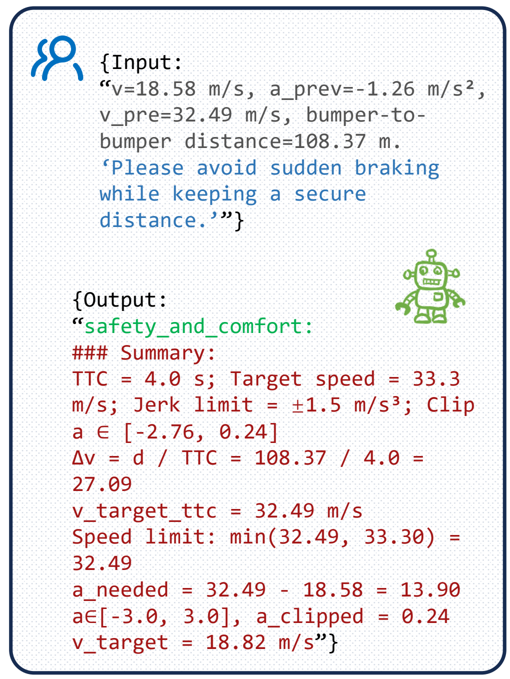


Figure 2. A Q&A Pair Example.

$$\text{TTC}(t) = \begin{cases} \frac{s(t)}{\Delta v(t)}, & \text{if } v_{j-1}(t) < v_j(t); \\ \text{Inf}, & \text{otherwise}, \end{cases} \tag{2}$$

where $\Delta v(t) = v_j(t) - v_{j-1}(t)$ denotes the relative speed between the ego vehicle and the preceding vehicle.

### 3.2. Supervised Fine-tuning (SFT)

SFT employs task-specific Q&A pairs to refine pre-trained LLMs, enhancing both their capabilities and controllability. By aligning the training objectives with human instructions rather than generic next-token prediction, SFT constrains

```
Default settings:

•   v_target: 33.3m/s
•   Acceleration range: [-3.0, 3.0]m/s²
•   TTC: 2.5s

Objective-specific settings:

•   "safety":                                              {"TTC": 4.0s, "v_target": 33.3m/s,  "jerk_limit": None},
•   "comfort":                                             {"TTC": 2.5s, "v_target": 33.3m/s,  "jerk_limit": ±1.5m/s³},
•   "mobility":                                            {"TTC": 2.0s, "v_target": 33.3m/s,  "jerk_limit": None},
•   "energy_efficiency":                                   {"TTC": 2.5s, "v_target": 20.0m/s,  "jerk_limit": None},
•   "safety_and_comfort":                                  {"TTC": 4.0s, "v_target": 33.3m/s,  "jerk_limit": ±1.5m/s³},
•   "safety_and_mobility":                                 {"TTC": 3.0s, "v_target": 33.3m/s,  "jerk_limit": None},
•   "comfort_and_mobility":                                {"TTC": 2.0s, "v_target": 33.3m/s,  "jerk_limit": ±1.5m/s³},
•   "safety_and_energy_efficiency":                        {"TTC": 4.0s, "v_target": 20.0m/s,  "jerk_limit": None},
•   "energy_efficiency_and_comfort":                       {"TTC": 2.5s, "v_target": 20.0m/s,  "jerk_limit": ±1.5m/s³},
•   "energy_efficiency_and_mobility":                      {"TTC": 2.0s, "v_target": 20.0m/s,  "jerk_limit": None},
•   "safety_and_mobility_and_comfort":                     {"TTC": 3.0s, "v_target": 33.3m/s,  "jerk_limit": ±1.5m/s³},
•   "safety_and_energy_efficiency_and_comfort":            {"TTC": 4.0s, "v_target": 20.0m/s,  "jerk_limit": ±1.5m/s³},
•   "safety_and_energy_efficiency_and_mobility":           {"TTC": 3.0s, "v_target": 20.0m/s,  "jerk_limit": None},
•   "energy_efficiency_and_mobility_and_comfort":          {"TTC": 2.0s, "v_target": 20.0m/s,  "jerk_limit": ±1.5m/s³},
•   "safety_and_energy_efficiency_and_comfort_and_mobility": {"TTC": 3.0s, "v_target": 20.0m/s,  "jerk_limit": ±1.5m/s³}.
```

Figure 3. CP-ACC Settings.

model outputs to desired response characteristics and domain knowledge. This approach offers an efficient pathway for human intervention in model behavior and enables rapid adaptation to specialized domains without extensive retraining or architectural modification [37].

To provide CP-ACC services, LLMs including LLaMA-3-8B [10] and Mistral-7B [13] were fine-tuned on the CP-ACC SFT dataset. Quantized low-rank adaptation (QLoRA) was employed to minimize the memory footprint without degrading performance. First, NormalFloat4 (NF4) quantization compressed the base model's weights to 4-bit precision [6]. A secondary quantization step was then applied to the low-rank adaptation matrices and optimizer constants to further reduce the memory usage. Additionally, NVIDIA's unified memory feature was utilized to offload optimizer states to CPU RAM when GPU memory capacity is exceeded, thereby preventing out-of-memory errors during training [37].

## 4. Experiment Setup

The CP-ACC models were fine-tuned using the LLaMA Factory framework [40] on $4\times$ NVIDIA Tesla T4 GPUs (16 GB each). The key hyperparameters used during fine-tuning are summarized in Tab. 1. QLoRA parameters were configured with a rank of 8 and an alpha value of 16 to balance model capacity and memory efficiency. A 4-bit quantization was employed to reduce the model size while maintaining acceptable fine-tuning performance.

### 4.1. Simulation Scenario

To capture acceleration, cruising, and deceleration phases, the preceding vehicle's trajectory is generated using the fifth-order polynomial specified in Eq. (3).

| Parameters | Values |
|---|---|
| Learning rate | 5e-5 |
| Number of epochs | 2 |
| Numerical precision | fp16 |
| Batch size per device | 1 |
| QLoRA rank | 8 |
| QLoRA alpha | 16 |
| Quantization bits | 4 |

Table 1. The setting of hyperparameters.

$$x(t) = b_0 + b_1 t + b_2 t^2 + b_3 t^3 + b_4 t^4 + b_5 t^5, \qquad (3)$$

where $x(t)$ denotes the vehicle's position at time $t$. $b_0$, and the coefficients $b_1$, $b_2$, $b_3$, $b_4$, and $b_5$ are determined by enforcing the boundary conditions listed below.

$$\begin{aligned} x(0) = x_0;\ \dot{x}(0) = v_0;\ \ddot{x}(0) = a_0; \\ x(T) = x_T;\ \dot{x}(T) = v_T;\ \ddot{x}(T) = a_T, \end{aligned} \qquad (4)$$

where $v(0)$ and $a(0)$ are the initial speed and acceleration, respectively. $T$ is the time duration. By specifying the boundary conditions as $x(0) = 0$, $x(T) = 150.0$ m, $v(0) = v(T) = 0$ m/s, $a(0) = a(T) = 0$ m/s$^2$, and $T = 15.0$ s, a smooth trajectory can be generated for the preceding vehicle from the initial to the final state. The ego vehicle is generated at position $-20$ m, and with an initial speed of $5.0$ m/s and zero acceleration (0 m/s$^2$).

### 4.2. Benchmarks

To comprehensively assess the performance of the proposed CP-ACC framework, benchmark comparisons are conducted against established car-following models commonly used in both real-world applications and simulation studies, including a linear feedback-based ACC model and intelligent driver models (IDM) car-following models. In addition, the evaluation incorporates zero-shot and few-shot LLMs to examine the impact and effectiveness of SFT within the proposed framework.

#### 4.2.1. Linear feedback-based ACC Model

A linear feedback-based ACC model, which regulates acceleration based on gap and speed errors, is formulated in Eq. (5).

$$a_j^{\mathrm{ACC}}(t) = k_1(s(t) - s_0 - t_{hw}v_j(t)) + k_2\Delta v(t), \quad (5)$$

where $s_0$ represents the minimum bumper-to-bumper distance. $t_{hw} = 2$ s is the current time-gap setting, and $k_1 = 0.23\ \mathrm{s}^{-2}$ and $k_2 = 0.07\ \mathrm{s}^{-1}$ are the feedback gains for the gap and speed errors, respectively [21].

#### 4.2.2. IDM Car-following Model

The IDM is a collision-free car-following model capable of generating realistic acceleration profiles, as expressed by Eq. (6).

$$a_j^{\mathrm{IDM}}(t) = a\left[1 - \left(\frac{v_j(t)}{v_{\mathrm{target}}}\right)^{\delta} - \left(\frac{s^*(v_j(t), \Delta v(t))}{s_j(t)}\right)^2\right],$$
$$s^*(v_j(t), \Delta v(t)) = s_0 + \max(0, v_j(t)T_g + \frac{v_j(t)\Delta v(t)}{2\sqrt{ab}}), \quad (6)$$

where $v_{\mathrm{target}}$ represents the target speed, $T_g$ is the bumper-to-bumper time gap to the preceding vehicle. $\delta$ is the acceleration exponent, and $b$ denotes the comfortable deceleration. The term $s^*$ indicates the desired distance [27].

Since the IDM car-following model inherently lacks the capability to interpret driving objectives expressed in natural language, we introduce multiple IDM configurations with different parameter settings as benchmarks to emulate specific driving objectives. For instance, parameter settings such as $T_g = 2.0$ s represents "mobility", $T_g = 4.0$ s corresponds to "safety", $v^e = 20.0$ m/s targets "energy efficiency", $\delta = 1$ aligns with "comfort" [27].

#### 4.2.3. Zero-shot LLMs

Distinct from fine-tuning, zero-shot prompting requires only a natural language instruction to guide the LLMs, without necessitating gradient-based parameter updates for task adaptation [1]. Despite the absence of task-specific weight adjustments, LLMs can still demonstrate notable zero-shot performance across various domains [7]. To generate ACC results using the zero-shot approach, the CP-ACC settings and a question formatted as illustrated in Fig. 2 are directly provided to the LLMs as prompts.

#### 4.2.4. Few-shot LLMs

In the few-shot prompting paradigm, LLMs receive a natural language task description and a limited number of demonstration examples, without any model parameter updates [1]. Specifically, several Q&A pairs from the CP-ACC SFT dataset are then presented to the LLMs to guide the subsequent generation of ACC output.

## 5. Results

The performance of zero-shot LLMs, few-shot LLMs, and fine-tuned LLMs was evaluated on a driving objective identification task using 20 randomly generated user profiles. As illustrated in Fig. 5, both LLaMA-3-8B and Mistral-7B achieved over 50.0% accuracy in the zero-shot setting, demonstrating a baseline understanding of user goals without any task-specific examples. Adding a few-shot context, i.e., including several examples in the prompt, led to improved results for both models, with each correctly identifying 14 out of 20 objectives. This suggests that few-shot prompting helps the models better interpret user inputs by providing relevant cues. The best performance was achieved through SFT, particularly with LLaMA-3-8B, which achieved an accuracy of 80.0%. Mistral-7B maintained its performance at 14 correct predictions after SFT. The models performed more accurately when the user profile contained a single objective (e.g., "safety", "comfort", or "energy efficiency") compared to when multiple objectives were combined (e.g., "safety and energy efficiency" or "energy efficiency, mobility, and comfort"). Some errors were attributed to ambiguous language in the prompts. For example, the term "efficient" could refer to either "energy efficiency" or "mobility", leading to misclassification. These findings highlight the importance of precise and unambiguous language when designing prompts or datasets for objective identification tasks.

Fig. 4 compares the longitudinal control performance of a linear feedback-based ACC, IDM car-following models, and LLMs across three sequential stages representing different driving objectives: (i) Stage 1 (0–5 s): "Energy efficiency and comfort"; (ii) Stage 2 (6–10 s): "Safety"; and (iii) Stage 3 (11–15 s): "Safety and mobility". This setup simulates realistic shifts in driving priorities, such as a sudden change in road or weather conditions requiring a transition from comfort-focused to safety-critical behavior. Conventional models, including linear feedback-based ACC and objective-specific IDM variants, apply fixed control strategies without adapting to changing objectives. Although they maintained smooth trajec-

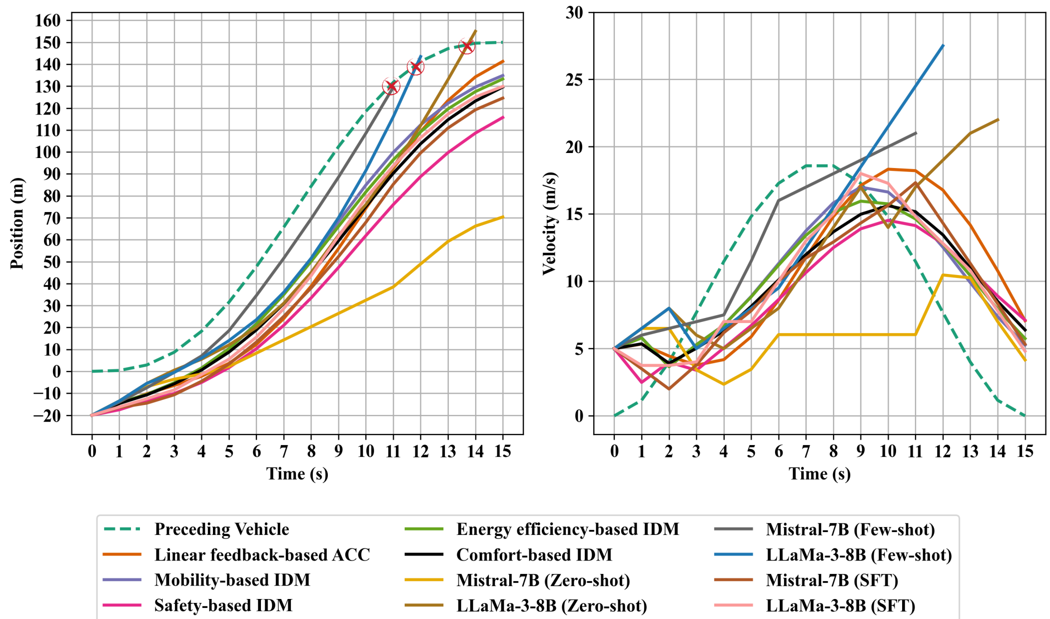


Figure 4. Comparison of Vehicle Trajectories under Different ACC Strategies.

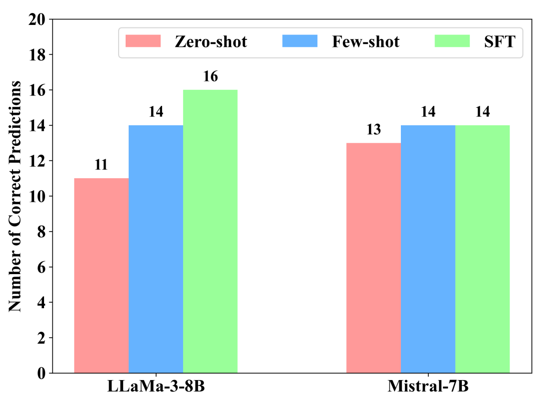


Figure 5. Results of Driving Objectives Identification.

tories and safe spacing comparable to the preceding vehicle, they lacked flexibility in adjusting behavior dynamically. LLM-based controllers exhibited varied performance. Specifically, LLaMA-3-8B (Zero-shot), Mistral-7B (Few-shot), and LLaMA-3-8B (Few-shot) failed to recognize the change in safety requirements, resulting in unsafe following behavior and eventual collisions. Mistral-7B (Zero-shot) adopted a conservative approach, maintaining a nearly constant speed even as the preceding vehicle accelerated, which is an overcautious behavior driven by the "safety" objective. Notably, the CP-ACC models (i.e., Mistral-7B (SFT) and LLaMA-3-8B (SFT)) successfully adapted to the changing driving objectives and consistently followed the expected output format without unnecessary reasoning steps. This demonstrates the strong potential of SFT in enabling LLMs to handle domain-specific control tasks and dynamically adjust to evolving driving goals.

Tab. 2 presents the TTC statistics across three driving stages. Two key metrics are reported per stage: $N_{\text{Inf}}$ is the number of cases where the ego vehicle's speed is less than or equal to the preceding vehicle's speed, representing an "absolute safety" condition; and Avg. TTC denotes the average TTC value when the ego vehicle is faster, indicating potential collision risk. Notably, both fine-tuned models, i.e., Mistral-7B (SFT) and LLaMA-3-8B (SFT), demonstrate strong safety performance, comparable to the Safety-based IDM. In Stage 2, where "safety" is prioritized, Mistral-7B (SFT) achieves the highest average TTC (54.9 s) among all models except the Safety-based IDM, confirming its ability to maintain safe following distances. In Stage 3, requiring a balance between "safety" and "mobility", the LLaMA-3-8B (SFT) maintains an average TTC of 5.1 s, which is comparable to the Mobility-based IDM. Mistral-7B (SFT) shows a more aggressive behavior aligned with "mobility", achieving a lower TTC of 4.8 s, but still above the critical safety threshold of 2.0 s.

Tab. 3 presents the average of absolute values of jerk, which reflects "comfort". In Stage 1, where comfort and energy efficiency are the primary objectives, Mistral-7B (SFT) maintains low jerk values (1.0 $m/s^3$), similar to the Comfort-based IDM and linear feedback-based ACC, indicating smooth acceleration and high passenger comfort. In contrast, LLaMA-3-8B (SFT) records a higher jerk (2.1 $m/s^3$), suggesting more aggressive acceleration in pursuit of

Table 2. TTC Statistics.

| Models | Stage 1 | | Stage 2 | | Stage 3 | |
|---|---|---|---|---|---|---|
| | $N_{Inf}$ | Avg. TTC | $N_{Inf}$ | Avg. TTC | $N_{Inf}$ | Avg. TTC |
| Linear feedback-based ACC | 3 | 10.4 | 4 | 11.2 | 0 | 2.3 |
| Mobility-based IDM | 4 | 3.1 | 4 | 18.5 | 0 | 5.1 |
| Safety-based IDM | 4 | 13.5 | 5 | N/A | 0 | 9.6 |
| Energy efficiency-based IDM | 4 | 3.2 | 4 | 39.1 | 0 | 5.6 |
| Comfort-based IDM | 4 | 3.6 | 4 | 53.9 | 0 | 5.8 |
| Mistral-7B (SFT) | 4 | 5.1 | 4 | 54.9 | 0 | 4.8 |
| LLaMa-3-8B (SFT) | 4 | 4.5 | 3 | 31.5 | 0 | 5.1 |

$N_{Inf}$: Number of infinity values; Avg. TTC: Average value of TTC; N/A: Not applicable.

Table 3. Jerk Statistics.

| Models | Stage 1 | Stage 2 | Stage 3 |
|---|---|---|---|
| Linear feedback-based ACC | 1.0 | 0.8 | 0.8 |
| Mobility-based IDM | 1.6 | 0.8 | 0.4 |
| Safety-based IDM | 1.7 | 0.5 | 0.5 |
| Energy efficiency-based IDM | 1.6 | 0.7 | 0.5 |
| Comfort-based IDM | 1.0 | 0.5 | 0.6 |
| Mistral-7B (SFT) | 1.0 | 0.3 | 1.4 |
| LLaMa-3-8B (SFT) | 2.1 | 1.5 | 0.7 |

Avg. Abs. Jerk: Average of absolute values of jerk.

energy-efficient speeds.

In summary, the findings underscore CP-ACC's capacity to understand and modulate driving objectives. SFT LLMs effectively balance "safety", "comfort", "mobility", and "energy efficiency" by adapting longitudinal control strategies in accordance with evolving driver preferences, contextual information, and vehicle states.

## 6. Conclusions

The increasing advancements in intelligent transportation systems (ITS) and autonomous driving technologies have significantly raised public expectations for mobility service quality. Existing advanced driver assistance systems (ADAS), particularly adaptive cruise control (ACC), struggle to deliver automatic longitudinal control services that holistically account for drivers' preferences and dynamic traffic environments. Motivated by the emergent capabilities of large language models (LLMs) in natural language processing, contextual reasoning, and high customizability, this paper proposes a contextual-personalized ACC (CP-ACC) framework. The framework was realized via supervised fine-tuning (SFT) of Llama-3-8B and Mistral-7B using a CP-ACC task-specific dataset. The proposed CP-ACC enables natural language interaction to interpret and balance multiple, often competing driving objectives such as "safety", "comfort", "mobility", and "energy efficiency".

This work marks a crucial step toward the realization of smart and human-centered ACC systems. Future development of this framework will integrate multimodal data, including naturalistic driving data, camera, and LiDAR inputs, to enhance contextual awareness and support more sophisticated decision-making. Crucially, real-world testing will be deployed to validate the practical efficacy of this framework and facilitate its integration into future autonomous driving systems.

## References

[1] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 5

[2] 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*, pages 902–909, 2024. 2

[3] Can Cui, Zichong Yang, Yupeng Zhou, Yunsheng Ma, Juanwu Lu, Lingxi Li, Yaobin Chen, Jitesh Panchal, and Ziran Wang. Personalized autonomous driving with large language models: Field experiments. In *2024 IEEE 27th Inter-*

*national Conference on Intelligent Transportation Systems (ITSC)*, pages 20–27. IEEE, 2024. 2

[4] Tianyu Cui, Xinjie Lin, Sijia Li, Miao Chen, Qilei Yin, Qi Li, and Ke Xu. Trafficllm: Enhancing large language models for network traffic analysis with generic traffic representation. *arXiv preprint arXiv:2504.04222*, 2025. 1

[5] Longchao Da, Minquan Gao, Hao Mei, and Hua Wei. Prompt to transfer: Sim-to-real transfer for traffic signal control with prompt learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 82–90, 2024. 1

[6] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36: 10088–10115, 2023. 4

[7] Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. How abilities in large language models are affected by supervised fine-tuning data composition. *arXiv preprint arXiv:2310.05492*, 2023. 5

[8] Shiyu Fang, Jiaqi Liu, Mingyu Ding, Yiming Cui, Chen Lv, Peng Hang, and Jian Sun. Towards interactive and learnable cooperative driving automation: a large language model-driven decision-making framework. *IEEE Transactions on Vehicular Technology*, 2025. 2

[9] Daocheng Fu, Xin Li, Licheng Wen, Min Dou, Pinlong Cai, Botian Shi, and Yu Qiao. Drive like a human: Rethinking autonomous driving with large language models. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW)*, pages 910–919. IEEE, 2024. 2

[10] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 4

[11] Martina Hasenjäger and Heiko Wersing. Personalization in advanced driver assistance systems and autonomous vehicles: A review. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7. IEEE, 2017. 2

[12] Stephen Hirst and Robert Graham. The format and presentation of collision warnings. In *Ergonomics and safety of intelligent driver interfaces*, pages 203–219. CRC Press, 2020. 3

[13] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023. 4

[14] Kemou Jiang, Xuan Cai, Zhiyong Cui, Aoyong Li, Yilong Ren, Haiyang Yu, Hao Yang, Daocheng Fu, Licheng Wen, and Pinlong Cai. Koma: Knowledge-driven multi-agent framework for autonomous driving with large language models. *IEEE Transactions on Intelligent Vehicles*, 2024. 2

[15] Genghua Kou, Fan Jia, Weixin Mao, Yingfei Liu, Yucheng Zhao, Ziheng Zhang, Osamu Yoshie, Tiancai Wang, Ying Li, and Xiangyu Zhang. Padriver: Towards personalized autonomous driving. *arXiv preprint arXiv:2505.05240*, 2025. 2

[16] Haoran Li, Wangling Wei, Sifa Zheng, Chuan Sun, Yunpeng Lu, and Tuqiang Zhou. Personalized driving behavior oriented autonomous vehicle control for typical traffic situations. *Journal of the Franklin Institute*, 361(10):106924, 2024. 2

[17] Siyan Li, Chuheng Wei, Guoyuan Wu, Matthew J Barth, Amr Abdelraouf, Rohit Gupta, and Kyungtae Han. Personalized trajectory prediction for driving behavior modeling in ramp-merging scenarios. In *2023 Seventh IEEE International Conference on Robotic Computing (IRC)*, pages 1–4. IEEE, 2023. 2

[18] Xishun Liao, Xuanpeng Zhao, Ziran Wang, Zhouqiao Zhao, Kyungtae Han, Rohit Gupta, Matthew J Barth, and Guoyuan Wu. Driver digital twin for online prediction of personalized lane-change behavior. *IEEE Internet of Things Journal*, 10 (15):13235–13246, 2023. 2

[19] Xishun Liao, Zhouqiao Zhao, Matthew J Barth, Amr Abdelraouf, Rohit Gupta, Kyungtae Han, Jiaqi Ma, and Guoyuan Wu. A review of personalization in driving behavior: Dataset, modeling, and validation. *IEEE Transactions on Intelligent Vehicles*, 2024. 2

[20] Jiageng Mao, Yuxi Qian, Junjie Ye, Hang Zhao, and Yue Wang. Gpt-driver: Learning to drive with gpt. *arXiv preprint arXiv:2310.01415*, 2023. 1

[21] Vicente Milanés and Steven E Shladover. Modeling cooperative and autonomous adaptive cruise control dynamic responses using experimental data. *Transportation Research Part C: Emerging Technologies*, 48:285–300, 2014. 5

[22] Michiel M Minderhoud and Piet HL Bovy. Extended time-to-collision measures for road traffic safety assessment. *Accident Analysis & Prevention*, 33(1):89–97, 2001. 3

[23] Mehmet Fatih Ozkan and Yao Ma. Personalized adaptive cruise control and impacts on mixed traffic. In *2021 American Control Conference (ACC)*, pages 412–417. IEEE, 2021. 2

[24] Ziye Qin, Ang Ji, Zhanbo Sun, Guoyuan Wu, Peng Hao, and Xishun Liao. Game theoretic application to intersection management: A literature review. *IEEE Transactions on Intelligent Vehicles*, 2024. 3

[25] Harmanpreet Singh, Nikhil Verma, Yixiao Wang, Manasa Bharadwaj, Homa Fashandi, Kevin Ferreira, and Chul Lee. Personal large language model agents: A case study on tailored travel planning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 486–514, 2024. 2

[26] Zhanbo Sun, Xue Yao, Ziye Qin, Peitong Zhang, and Ze Yang. Modeling car-following heterogeneities by considering leader–follower compositions and driving style differences. *Transportation research record*, 2675(11):851–864, 2021. 2

[27] Martin Treiber and Arne Kesting. *Traffic flow dynamics*. Springer, 2013. 3, 5

[28] Meng Wang, Winnie Daamen, Serge Hoogendoorn, and Bart Van Arem. Potential impacts of ecological adaptive cruise

control systems on traffic and environment. *IET Intelligent Transport Systems*, 8(2):77–86, 2014. 3

[29] Yanbing Wang, Ziran Wang, Kyungtae Han, Prashant Tiwari, and Daniel B Work. Personalized adaptive cruise control via gaussian process regression. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 1496–1502. IEEE, 2021. 2

[30] Mengyao Wu, F Richard Yu, Peter Xiaoping Liu, and Ying He. Facilitating autonomous driving tasks with large language models. *IEEE Intelligent Systems*, 2024. 2

[31] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 2024. 2

[32] Zixuan Xu, Tiantian Chen, Zilin Huang, Yang Xing, and Sikai Chen. Personalizing driver agent using large language models for driving safety and smarter human–machine interactions. *IEEE Intelligent Transportation Systems Magazine*, 2025. 1, 2

[33] Fan Yang, Zheng Chen, Ziyan Jiang, Eunah Cho, Xiaojiang Huang, and Yanbin Lu. Palr: Personalization aware llms for recommendation. *arXiv preprint arXiv:2305.07622*, 2023. 1

[34] Kairui Yang, Zihao Guo, Gengjie Lin, Haotian Dong, Zhao Huang, Yipeng Wu, Die Zuo, Jibin Peng, Ziyuan Zhong, Xin Wang, et al. Trajectory-llm: A language-based data generator for trajectory prediction in autonomous driving. In *The Thirteenth International Conference on Learning Representations*, 2025. 2

[35] Lei Yang, Yafei Liu, Zhanbo Sun, Qiruo Yan, Ziye Qin, and Xuting Wang. Stability analysis of connected automated vehicle platoon controller under unreliable communication conditions. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pages 3647–3652. IEEE, 2023. 1

[36] Xue Yao, Simeon C. Calvert, and Serge P. Hoogendoorn. A novel framework for identifying driving heterogeneity through action patterns. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–13, 2025. 1

[37] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*, 2023. 4

[38] Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, et al. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*, 2024. 1

[39] Zhouqiao Zhao, Xishun Liao, Amr Abdelraouf, Kyungtae Han, Rohit Gupta, Matthew J Barth, and Guoyuan Wu. Real-time learning of driving gap preference for personalized adaptive cruise control. In *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 4675–4682. IEEE, 2023. 1, 2

[40] Yaowei Zheng, Richong Zhang, Junhao Zhang, YeYanhan YeYanhan, and Zheyan Luo. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 400–410, 2024. 4

[41] Bing Zhu, Yuande Jiang, Jian Zhao, Rui He, Ning Bian, and Weiwen Deng. Typical-driving-style-oriented personalized adaptive cruise control design based on human driving data. *Transportation research part C: emerging technologies*, 100: 274–288, 2019. 2