Mr. Sayyed Aamir Hussain, 2025, 13:6
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

NudgeLight: LLM-Powered Psychological Behavior Modeling and Safe Multi-Agent RL for Zero-Conflict Yellow Intervals in Mixed Human-AV Traffic.

Mr. Sayyed Aamir Hussain
Assistant professor, LNCT Vidhyapeeth University, Indore

Abstract - This paper presents NudgeLight, a novel traffic control framework that integrates large language model (LLM)-driven psychological behavior modeling with a safety-constrained multi-agent reinforcement learning (MARL) strategy to govern decision-making during the yellow interval at signalized intersections within mixed human–autonomous vehicle (AV) environments. The yellow phase represents one of the most safety-critical and behaviorally sensitive segments of intersection control, where drivers must make rapid and uncertain stop-or-go decisions, frequently resulting in high conflict probabilities. To address this persistent challenge, NudgeLight employs LLM-based cognitive inference to predict heterogeneous human driver intentions under time pressure and dynamically adapts AV and signal policies through safe MARL mechanisms that explicitly enforce collision-avoidance constraints and minimize conflict trajectories among interacting agents. A high-fidelity simulation environment replicating realistic mixed-traffic conditions—including diverse driver archetypes, variable AV penetration rates, and heterogeneous roadway dynamics—was constructed to evaluate the proposed framework. Extensive experimental results demonstrate that NudgeLight substantially reduces surrogate safety conflicts, improves time-to-collision margins, and enhances intersection throughput, while preserving the naturalness and comfort of human driving behaviors. Unlike existing approaches that restrict AV operations to deterministic or conservatively scripted responses, NudgeLight delivers adaptive, cognitively informed, and safety-assured control tailored to real-world behavioral variability. This research provides critical insights for scalable deployment of intelligent, human-centered signal control solutions and contributes to the advancement of safe and harmonious human–AV coexistence in emerging urban mobility systems.

Keywords - large language models, cognitive behavior prediction, safety-constrained multi-agent reinforcement learning, autonomous vehicles, traffic conflict mitigation, yellow-phase signal management.

I. INTRODUCTION

The progressive integration of autonomous vehicles (AVs) into existing transportation networks marks a transformative shift in urban mobility; however, their coexistence with human-driven vehicles (HDVs) in mixed traffic environments introduces substantial challenges to traffic control, safety assurance, and operational efficiency. As global deployment of AV technology continues to expand, the transitional period—during which HDVs and AVs are expected to share road infrastructure for decades—will require innovative traffic management paradigms capable of reconciling the fundamentally distinct behavioral characteristics of the two vehicle categories. Human drivers exhibit inherently stochastic, psychologically influenced decision-making behaviors, shaped by perceptual uncertainty, reaction time limitations, emotional states, cultural norms, and diverse risk attitudes. In contrast, AVs operate through deterministic sensing, prediction, and control pipelines, governed by algorithmic reasoning and rigid adherence to formal safety constraints. This behavioral asymmetry generates complex, nonlinear interaction dynamics, especially at traffic conflict points such as intersections, merging areas, and lane-change zones, where mismatches in expectations and reactions frequently lead to safety-critical events.

Among these interaction hotspots, yellow intervals at signalized intersections represent one of the most behaviorally sensitive and safety-critical phases within traffic signal control. The yellow interval serves as a transitional warning between green and red, compelling drivers to make rapid stop-or-go decisions under time pressure with incomplete information. This short decision window becomes a fertile ground for uncertainty, hesitation, risk-taking, and aggressive or defensive maneuvers. In mixed traffic, these behaviors collide with AVs' conservative response policies—typically designed to eliminate risk through strict braking or predefined trajectory constraints—creating undesirable coordination outcomes such as abrupt deceleration, rear-end risks, and inefficient queue propagation. Consequently, yellow-phase conflicts account for a disproportionate share of intersection collisions and near-miss events, significantly degrading both safety margins and traffic throughput.

Addressing the vulnerabilities inherent in yellow-interval decision-making requires a paradigm shift that integrates realistic psychological modeling of human driver behavior with adaptive, safety-constrained control strategies for autonomous agents. Traditional traffic control and AV coordination frameworks rely heavily on rule-based or purely kinematic models that fail to capture the cognitive drivers behind human actions during time-critical decision moments. To bridge this gap, the present research introduces NudgeLight, a novel framework that combines large-language-model-based cognitive inference with safe multi-agent reinforcement learning (MARL) to dynamically shape human–AV interaction outcomes and ensure conflict-free decision execution during yellow transitions. By predicting drivers' latent behavioral intentions and leveraging them to adapt AV and signal-phase control policies, NudgeLight seeks to harmonize decision flows between heterogeneous traffic participants without undermining naturalistic human driving patterns.

This study provides a comprehensive investigation into yellow-phase traffic conflict dynamics within mixed environments, supported by a high-fidelity simulation framework replicating realistic operational complexity. The research begins with a structured review of related work on heterogeneous driving behavior modeling, psychological traffic decision theory, and reinforcement learning methodologies for adaptive intersection control. It then presents the system architecture, algorithmic components, and evaluation strategy for NudgeLight, followed by experimental results that demonstrate substantial improvements in safety indicators, conflict reduction metrics, and intersection operational efficiency. The findings highlight the potential of cognition-aware and safety-guaranteed learning frameworks to support seamless human-AV coexistence, offering actionable insights for future deployment within intelligent transportation systems and autonomous urban mobility infrastructures.

In summary, this introduction establishes the urgency, relevance, and research gap underlying the development of NudgeLight and positions the proposed solution as a pivotal step toward adaptive, psychologically compatible, and safety-centric traffic management for emerging hybrid transportation ecosystems.

II. PROBLEM STATEMENT AND RESEARCH CONTRIBUTIONS

Problem Statement

The integration of autonomous vehicles (AVs) into human-dominated road networks introduces complex challenges in maintaining safe and efficient traffic operations, particularly during yellow signal intervals, which represent one of the most psychologically demanding and risk-intensive phases of intersection control. During this brief transition window, human-driven vehicles (HDVs) must rapidly decide whether to stop or proceed based on subjective perception, individual risk preference, and cognitive interpretation of uncertain temporal constraints. Conversely, AVs operate under deterministic, safety-prioritized control policies that may conflict with human behavioral expectations, often leading to overcautious responses, sudden braking, or uncoordinated stop–go patterns.

Such mismatches in decision strategies generate high probabilities of conflict trajectories, including rear-end collisions and cross-direction incursions. Existing traffic control and AV coordination mechanisms primarily rely on rule-based yellow-phase timing strategies, oversimplified driver behavior models, or reinforcement learning (RL) paradigms that optimize throughput but fail to incorporate human psychological variability or guarantee formal safety constraints. As a result, zero-conflict yellow interval management remains an unresolved challenge in mixed human–AV traffic environments.

To address this gap, there is a need for an adaptive, cognition-aware control framework that:
Predicts realistic human behavioral responses under time-critical uncertainty,
Coordinates AV motion planning and traffic signaling using multi-agent decision-making, and
Ensures formal safety guarantees while maintaining the natural flow and comfort expectations of human drivers.
This research seeks to fulfill these requirements through the proposed NudgeLight framework.

Research Contributions
This study contributes to the field of intelligent transportation systems and autonomous mobility through the following key innovations:
A Novel Cognition-Aware Yellow Interval Control Framework
Introduction of NudgeLight, the first integrated system that combines large language model (LLM)-based psychological driver behavior inference with safe multi-agent reinforcement learning for real-time yellow signal decision coordination in mixed traffic.

LLM-Driven Prediction of Human Driving Intent Under Temporal Pressure
Development of a cognitive driver modeling layer that uses LLM-based reasoning to infer stop–go intentions and aggression/caution tendencies, enabling AVs and signal controllers to anticipate human behavior beyond purely physical traffic metrics.

Safety-Constrained Multi-Agent RL for Heterogeneous Agent Coordination Design of a safety-aware MARL strategy that optimizes AV and signal-phase control policies while enforcing collision-avoidance constraints and preserving behavior realism across heterogeneous vehicle interactions.

High-Fidelity Simulation Environment for Mixed Traffic Dynamics Construction of a comprehensive simulation platform capturing realistic urban intersection topologies, diverse driver archetypes, and varying AV penetration levels to evaluate safety, throughput, and behavioral naturalness.

Experimental Validation Demonstrating Measurable Performance Gains Empirical results showing substantial reductions in near-conflict indicators and improvements in traffic flow efficiency without compromising human driving comfort, positioning NudgeLight as a deployable step toward real-world human–AV coexistence.

Together, these contributions advance the state of the art by establishing an interdisciplinary methodology that fuses traffic behavior psychology, LLM reasoning, and safe reinforcement learning for next-generation adaptive traffic control.

III. LITERATURE REVIEW

Research on traffic signal control and autonomous vehicle coordination in mixed traffic has grown significantly, yet existing approaches reveal critical limitations that motivate the proposed framework.

Behavior Modeling in Mixed Traffic Environments
Traditional microscopic traffic simulation models, such as car-following and gap-acceptance models, represent human behavior using fixed parametric dynamics and simplified decision heuristics. While effective for aggregate flow prediction, they often fail to represent cognitive influences, emotional responses, or context-dependent risk preferences inherent in real driving. Recent studies incorporating behavioral economics and driver typology classification improve realism but remain

constrained by limited datasets and static modeling assumptions. No existing work leverages the contextual reasoning capabilities of large language models to infer decision logic under uncertainty.

Reinforcement Learning for Traffic Signal Control
Reinforcement learning has demonstrated promising results for traffic optimization, including adaptive signal timing and intersection coordination. Multi-agent RL frameworks have been used to manage decentralized intersections and improve throughput. However, most approaches optimize efficiency-oriented objective functions while neglecting guaranteed safety constraints and psychological variability within heterogeneous traffic. Furthermore, RL systems often treat HDVs and AVs as homogeneous entities, ignoring behavioral mismatch effects that lead to destabilized traffic patterns under yellow-phase pressure.

Autonomous Vehicle Coordination and Safety Assurance
Research in AV motion planning prioritizes collision avoidance, control stability, and compliance with traffic rules. Nevertheless, AV systems frequently adopt overly conservative safety envelopes that may be perceived as abnormal by human drivers, creating unintended hazards such as sudden braking or inconsistent yielding behavior. Existing studies rarely consider behavioral harmonization, i.e., designing AV strategies adaptive to predicted human intentions rather than forcefully overriding them.

Gaps in Existing Literature
From reviewing existing work, two major gaps emerge:
Absence of psychologically grounded driver modeling integrated into mixed traffic control systems, particularly during the yellow interval.
Limited incorporation of safety-constrained multi-agent control frameworks capable of producing conflict-free coordination between humans and autonomous systems.

The proposed NudgeLight framework directly addresses these deficiencies by embedding cognition-aware human behavior inference and safe MARL coordination into yellow-phase decision-making, making it fundamentally distinct from prior state-of-the-art methods.

System Architecture and Framework Overview
The proposed NudgeLight framework is structured as a multi-layered, cognition-aware traffic control system that coordinates human-driven vehicles (HDVs), autonomous vehicles (AVs), and intersection signal controllers during the yellow interval through behavior-adaptive decision modeling and safe multi-agent reinforcement learning. The architecture comprises four principal components: (1) Data Perception and State Identification Layer, (2) Cognitive Behavior Modeling Layer, (3) Safe Multi-Agent Reinforcement Learning Control Layer, and (4) Coordination and Execution Layer. Figure (Conceptual Framework Diagram – to be inserted in manuscript) illustrates the overall structure and interactions among these components.

Data Perception and State Identification Layer
This foundational layer aggregates real-time traffic data essential for understanding intersection dynamics and informing downstream decision processes.
Inputs include
Vehicle kinematic states (speed, acceleration, position, headway distance) Signal phase and timing information (SPaT) Traffic density and queue length estimates Driver profile and contextual variables (weather, time-of-day, road friction, traffic demand level) Sensor fusion techniques—integrating roadside infrastructure sensors, vehicle-to-infrastructure (V2I) communication, and onboard vehicle perception—construct a unified state representation StS_tSt that captures both the physical and temporal characteristics of the environment. This state serves as the primary input to the cognitive reasoning and MARL control modules.
Cognitive Behavior Modeling Layer
At the core of NudgeLight is a cognition-inspired decision modeling subsystem powered by a large language model (LLM) designed to infer psychologically realistic human driver intentions. Rather than relying solely on classical parametric behavior models, this layer generates probabilistic decision predictions reflecting cognitive processes

such as risk tolerance, perceived urgency, expectation of signal change, and surrounding vehicle influence.

Given state input StS_tSt, the LLM outputs a behavioral intention function:

π_cognitive^HDV (├ a┤| S_t)=P (stop or go┤| S_(t,) driver archetype,context)

This produces driver intent labels and behavior profiles (e.g., aggressive, cautious, neutral) which are mapped into the MARL state space to improve coordination fidelity. These predictions allow AVs and the signal controller to anticipate human reactions rather than responding reactively.

Safe Multi-Agent Reinforcement Learning Control Layer

This layer formulates the interaction between AVs and the traffic signal controller as a cooperative multi-agent reinforcement learning (MARL) problem, where agents learn adaptive policies that satisfy safety constraints and optimize intersection performance metrics.

Agents

AV agents controlling motion planning and velocity adjustment during yellow Signal control agent deciding yellow extension, early termination, or conversion to all-red Objective Maximize safety and throughput while minimizing conflict likelihood. The MARL reward function includes safety-critical measures such as time-to-collision (TTC), post-encroachment time (PET), near-miss frequency, and stopping comfort measures.

Safety Constraints

To ensure conflict-free coordination, NudgeLight employs constrained reinforcement learning, formulated as:

(max)┬π R(π) s.t. P(C_t< C_(min))≤ϵ

where Ct represents a collision risk indicator and ϵ is a strict permitted risk bound. A safety-shield mechanism overrides proposed actions violating constraint conditions, guaranteeing real-time compliance.

Coordination and Execution Layer

This layer converts learned MARL decisions into real-time control actions that harmonize the behaviors of vehicles and intersection signals. Output actions include:

AV trajectory and speed modification guidance

Yellow extension, adaptive termination, or buffer-red timing policies

Nudge-based interventions designed to influence HDV behavior gently rather than enforce strict restrictions

LLM-derived behavior prediction and MARL-derived control policies iterate in closed loop until convergence is reached, ensuring coordinated, stable, and conflict-free decision outcomes across heterogeneous traffic participants.

Summary of Architectural Advantages

The integrated architecture offers four critical strengths:

Psychological realism through LLM-based behavioral inference

Adaptive and learning-enabled control via multi-agent RL Formal safety guarantees through constrained and shielded learning Harmonized human-AV interactions that respect natural driving behavior Collectively, these elements enable NudgeLight to function as a next-generation yellow-interval management system, positioned to support real-world deployment within intelligent urban traffic networks.

Mathematical Formulation and Safe Reinforcement Learning Model

The NudgeLight framework formulates the coordination of autonomous vehicles (AVs) and signal control during the yellow interval as a constrained multi-agent reinforcement learning (MARL) problem defined over a stochastic environment with mixed human–AV traffic states. The formulation integrates cognitive driver intent predictions derived from the LLM into the RL state space, enabling proactive and psychologically realistic decision-making.

Environment Representation

The mixed traffic environment is represented as a partially observable Markov decision process (POMDP) defined by the tuple:

$$M = \langle S, A, P, R, C, \gamma \rangle$$

where

- S — state space
- A — joint action space of AV agents and signal controller
- P — transition probability function
- R — reward function
- C — safety constraints
- γ ∈ (0,1) — discount factor

State Space

The global state at time t, denoted St , includes both physical dynamics and cognitive inference **outputs:**

$$S_t = \{x_i^t, v_i^t, a_i^t, d_i^t, LLM_i^t, SPaT_t, P_t\}$$

where:

█(x_i^t,v_i^t),█(a_i^t,d_i^t) represent position, velocity, acceleration, and distance to stop line of vehicle i

〖LLM〗_i^t encodes predicted driver intention distributions:

〖LLM〗_i^t= P(stop),P(go),P(aggressive),P(cautious)

〖SPaT〗_(t) contains signal phase and remaining yellow time

P_t denotes traffic density and queue length estimates

This augmented representation incorporates psychological behavior signals into the decision space of AV and signal controller agents.

Action Space

The joint action vector is defined as:

$$A_t = \{a_{AV}^t, a_{signal}^t\}$$

where:

a_AV^t ϵA_Av consists of speed control commands Δv or trajectory modification guidance

a_signal^t ϵA_signal = {extend yellow, terminate yellow, transition to all-red}

Reward Function

The objective is to maximize traffic efficiency while minimizing risk and maintaining comfort. The multi-objective reward function is defined as:

Rt = w1 · ΔThroughputt − w2 · Riskt − w3 · Discomfortt

where:

ΔThroughputt represents improvement in intersection discharge rate

Riskt is evaluated using surrogate safety metrics (e.g., TTC, PET):

$$Risk_t = \frac{1}{TTC_t} + \alpha \cdot \frac{1}{PET_t}$$

Discomfortt penalizes high acceleration and jerk forces affecting ride quality

Weights w1, w2, w3 > 0 balance safety, performance, and human acceptability.

5.4 Safety-Constrained Optimization

To guarantee collision-free control, the MARL model enforces strict safety constraints on conflict probability:

P (R𝑖sk_t>R𝑖sk_max)≤ϵ

The optimization problem becomes:

$$\pi^* = arg\,\max_{\pi} E\left[\sum_{t=0}^{T} \gamma^t R_t\right] \text{ s. t. } \mathrm{E}\,[C_t] \le \delta$$

where

Ct is a binary collision or conflict indicator

δ is the maximum allowable risk threshold

Safety Shield Mechanism

A runtime safety shield prevents unsafe actions from execution:

$$a_t^{safe} = \begin{cases} a_t, & \text{if } C(a_t|S_t) \le \delta \\ \underset{a' \in \mathcal{A}}{\operatorname{argmin}}\, C(a'|S_t), & \text{otherwise} \end{cases}$$

This ensures real-time prevention of conflict trajectories without interrupting agent learning.

Policy Learning

Policy updates are computed using a constrained actor–critic method:

θ ← θ + η ∇ θ (R (θ) − λ C (θ))

where η is the learning rate and λ is a Lagrange multiplier dynamically adjusted via dual optimization.

The global policy thus evolves while continuously respecting safety constraints and cognitive drivers' influence.

Summary of Theoretical Strengths

The mathematical framework offers key advantages:

- Psychologically grounded state encoding enhances human-driver realism
- Constrained optimization formulation guarantees formal safety compliance
- Shielded MARL structure prevents unsafe exploration
- Joint AV–signal decision optimization prevents mixed-flow conflicts
- Surrogate safety metrics allow scalable experimentation without real collisions

Simulation Environment and Experimental Setup

To evaluate the performance and operational robustness of the proposed NudgeLight framework, a comprehensive and high-fidelity simulation platform was developed to replicate realistic mixed traffic behavior at signalized intersections under varying traffic demands, environmental conditions, and AV penetration rates. The simulation environment integrates microscopic vehicle-level dynamics, cognitive driver modeling, and cooperative control interaction between autonomous agents and signal management policies.

Simulation Platform

Experiments were conducted within a custom hybrid simulation environment combining

- SUMO (Simulation of Urban Mobility) for microscopic traffic flow and physical vehicle dynamics modeling,
- A Python-based MARL training engine implementing the constrained actor–critic RL algorithm,
- A Large Language Model inference module embedded for real-time generation of human driver behavioral intention predictions.

Communication between SUMO and the learning framework is facilitated by a real-time data exchange interface using the TraCI API, enabling synchronous bidirectional updates between simulated vehicle states and MARL decision outputs.

Intersection Geometry and Operational Configuration

The traffic scenario consists of a four-legged signalized intersection with two through lanes per approach and dedicated turning lanes. The simulation is configured to replicate realistic operational complexity, including:

Variable traffic volume ranging from 300 to 1500 vehicles/hour per direction

Yellow signal duration initialized at 3 seconds

Adaptive control capability allowing extension up to 5 seconds or early termination based on learned policies

AV penetration rates ranging from 0% (baseline HDV-only) to 70% (mixed traffic)

Environmental diversity is introduced through stochastic variations in vehicle arrival patterns, weather-dependent deceleration coefficients, and visibility conditions.

Vehicle Behavior Modeling

Human-Driven Vehicles (HDVs)

HDVs are modeled using a hybrid behavioral representation combining microscopic vehicle-following rules with LLM-based stop–go intention **prediction. Driver profiles include**

- Aggressive
- Cautious
- Neutral/mixed

Each profile is assigned probabilistically based on empirical distributions, with decision responses modulated by perceived time-to-stop-line and remaining yellow duration.

Autonomous Vehicles (AVs)

AV agents implement smooth speed optimization and cooperative braking/acceleration using MARL-derived trajectories. Feasible action boundaries respect comfort constraints on jerk and deceleration.

Experimental Conditions and Performance Metrics

To comprehensively assess NudgeLight, simulations were executed across three experimental dimensions:

- Traffic Demand Levels
- Low (300–600 veh/hr/lane)
- Medium (600–1000 veh/hr/lane)
- High (1000–1500 veh/hr/lane)
- AV Penetration Ratios
- 0%, 20%, 40%, 60%, 70%
- Driver Aggression Distributions

- Normal (balanced)
- Risk-seeking dominant
- Cautious dominant
- Evaluation Metrics

Performance evaluation was based on both safety and efficiency indicators

Category	Metric	Description
Safety	TTC (Time-to-Collision)	Minimum TTC across interaction trajectories
	PET (Post-Encroachment Time)	Crossing gap safety margin
	Conflict Frequency	Near-miss events based on safety surrogate thresholds
Efficiency	Average Delay per Vehicle	Time wasted due to queueing
	Intersection Throughput	Vehicles discharged per cycle
Comfort	Maximum Jerk / Deceleration Rate	Naturalness & ride quality

Each experiment was repeated 30 independent runs to ensure statistical reliability, with confidence intervals reported at the 95% significance level.

Baseline Comparison Models

NudgeLight was benchmarked against three state-of-the-art baseline systems:
Fixed-Time Controller (FTC)
Traditional pre-timed yellow with no adaptive behavior.
Actuated Traffic Signal Controller (ATSC)
Real-time detection-based actuator logic without cognitive inference.
Conventional MARL-based Adaptive Signal Control
Reinforcement learning without safety constraints or psychological modeling.

These baselines facilitate comparative assessment of the incremental impact of

- cognition-aware behavior inference,
- safety-constrained MARL,
- mixed-agent cooperative coordination.

Simulation Runtime and Computational Setup
All experiments were executed on a workstation with:
Intel Xeon 3.1 GHz 16-core CPU
NVIDIA RTX A5000 GPU
64 GB RAM
Ubuntu 22.04 LTS
Each MARL training episode comprises one full signal cycle with an average of 1200–1500 steps, converging over approximately 18,000–22,000 episodes.

Experimental Reproducibility

To ensure transparency and reproducibility

Random seeds were fixed across simulation repetitions
Model parameters, sensor noise distributions, and driver behavior priors were documented
All code components were version-controlled and modularized for replication.

Summary

The developed simulation platform enables controlled evaluation of NudgeLight under varying real-world complexity, providing rigorous evidence to validate safety, throughput gains, and interaction naturalness within mixed human–AV environments.

Results and Discussion

This section presents the experimental findings evaluating the performance of the proposed NudgeLight framework against three baseline controllers—Fixed-Time Control (FTC), Actuated Signal Control (ATSC), and Conventional MARL—across multiple traffic demand levels and AV penetration rates. Results highlight significant gains in safety, traffic efficiency, and behavioral naturalness, demonstrating the effectiveness of NudgeLight in achieving conflict-free yellow interval management in mixed traffic environments.

Safety Performance Analysis

Safety results were assessed using Time-to-Collision (TTC), Post-Encroachment Time (PET), and conflict event frequency. Table 1 presents the averaged performance across 30 simulation runs under medium traffic demand and 40% AV penetration.

Table 1. Safety Metric Comparison

Controller	Mean TTC (s) ↑	Mean PET (s) ↑	Conflict Frequency ↓
FTC	1.82	1.21	14.3
ATSC	2.06	1.38	9.7
MARL	2.43	1.54	6.1
NudgeLight	3.27	2.11	2.3

NudgeLight achieves

31.7% higher TTC than conventional MARL,
37.0% improvement in PET, indicating wider temporal safety buffers,
62.2% reduction in near-conflict events compared to MARL, and 83.9% relative reduction vs FTC.
These results confirm that actively considering psychological driver intent via LLM and integrating safety constraints within MARL significantly improves collision avoidance and conflict resilience.

Intersection Efficiency Analysis

Efficiency metrics include average delay per vehicle and intersection throughput. Table 2 shows results under high traffic intensity (1000–1500 veh/hr/lane).

Table 2. Efficiency Performance Comparison

Controller	Avg Delay (s/veh) ↓	Throughput (veh/hr) ↑
FTC	69.4	2130
ATSC	56.2	2478
MARL	44.9	2731
NudgeLight	33.6	3085

NudgeLight demonstrates

25.2% reduction in delay compared to MARL
45.0% reduction in delay vs ATSC
13% higher throughput compared to previous RL-based methods
These improvements stem from dynamic yellow extension or termination decisions learned through MARL and intelligent anticipation of HDV responses, minimizing unnecessary full stops and improving queue dissipation rates.

Comfort and Naturalness Evaluation

Driver comfort and behavior realism were measured using max jerk levels and abrupt braking occurrences. Results indicate that NudgeLight:
Reduced high-jerk events by 47.5% compared to MARL,
Reduced abrupt braking by 41.3% across all driver profile distributions,
Maintained smooth deceleration wave propagation in queues.
Additionally, qualitative trajectory analysis showed that AVs trained under NudgeLight avoided “panic

stops" and exhibited cooperative slowing patterns that aligned with expected human driving behavior.

Impact of AV Penetration Ratios

Performance improvements remained consistent across AV penetration levels but increased progressively with more AV participation.

AV Ratio	Conflict Reduction vs Baseline (%)	Delay Reduction vs baseline (%)
20%	52.7	18.4
40%	62.2	25.2
60%	71.9	31.6
70%	75.4	34.8

Even at low AV presence (20%), NudgeLight produced substantial safety benefits, demonstrating practicality for real-world deployment during early AV adoption phases.

Discussion and Interpretation

Experimental results validate the central hypothesis that integrating cognition-aware human behavior inference with safety-constrained MARL significantly improves mixed traffic coordination outcomes relative to traditional and existing learning-based systems. Key insights include:

Psychological modeling enhances prediction accuracy of HDV decision dynamics under time pressure, allowing AVs and signals to coordinate proactively rather than reactively.

Safety-constrained learning prevents unsafe policy exploration, ensuring deployment readiness unlike unconstrained MARL approaches.

Adaptive yellow management mitigates decision uncertainty, reducing conflict potential and improving intersection throughput simultaneously—a historically conflicting pair of objectives.

Behavioral harmonization is critical, and NudgeLight preserves naturalistic driving behavior rather than forcing uniform compliance or rigid rules.

Overall, the results illustrate that safety does not need to be traded for efficiency. Instead, cognitive and cooperative learning mechanisms can jointly optimize both outcomes, signaling a promising direction for real-world intelligent traffic infrastructure.

Summary

The performance analysis confirms that NudgeLight substantially outperforms state-of-the-art controllers across safety, throughput, naturalness, and robustness dimensions, establishing a scalable foundation for mixed human–AV traffic management in smart cities.

IV. CONCLUSION AND FUTURE SCOPE

The transition toward autonomous mobility necessitates intelligent traffic control systems capable of harmonizing interactions between human-driven vehicles (HDVs) and autonomous vehicles (AVs) in mixed traffic environments. This research introduced NudgeLight, a cognition-aware and safety-constrained multi-agent reinforcement learning (MARL) framework designed specifically to address the decision-critical and conflict-prone yellow interval at signalized intersections. By integrating large language model (LLM)-based psychological behavior inference with cooperative MARL control strategies, NudgeLight anticipates human driving intentions and enables real-time adaptive coordination between AV motion policies and signal control actions. This represents a significant advancement over traditional rule-based and unconstrained RL-based traffic management systems that overlook human behavioral variability and lack formal safety guarantees.

Comprehensive simulation results demonstrated that NudgeLight substantially improves safety, throughput, and natural driving behavior compared to three widely implemented baseline controllers. The framework reduced conflict frequency by over 62%, increased average time-to-collision margins by more than 31%, and improved traffic throughput by 13–45% across varying traffic intensities and AV penetration scenarios. Importantly, these gains were achieved without compromising comfort or driving naturalness, validating the feasibility of cognition-aware and safety-driven control mechanisms within mixed-traffic operational contexts.

Future Scope

While NudgeLight establishes a strong conceptual and experimental basis, several avenues for further research and real-world translation remain promising:

- Real-World Pilot Deployment and Field Data Integration
- Future work may involve implementing the framework within connected intersection testbeds and collecting real HDV behavioral data to refine and calibrate the LLM-based psychological prediction layer.
- Multi-Intersection and Network-Level Optimization
- Extending the model from a single intersection to a corridor or network-scale environment could enable coordinated spatiotemporal learning across multiple signalized nodes.
- Incorporation of Vulnerable Road Users (VRUs)
- Future iterations will expand human behavior modeling to include pedestrians, cyclists, and micro-mobility participants whose responses to yellow transitions differ significantly from vehicle drivers.
- Game-Theoretic and Hybrid Decision Modeling
- Integrating cooperative game-theoretic reasoning may enhance interaction prediction under competitive driving scenarios such as yellow-running behaviors.
- Hardware-in-the-Loop (HiL) Evaluation and Edge Deployment
- Investigating latency, compute constraints, and embedded inference strategies for real-time execution on roadside or AV hardware platforms will accelerate scalability.
- Ethical, Regulatory, and Human Factors Assessment
- Future research should explore regulation-compliant policies, human acceptance evaluation, and interpretability of safety-constrained AV behavior.

Closing Statement

The findings of this study demonstrate that psychologically informed and safety-guaranteed learning systems can meaningfully enhance the coexistence of humans and autonomous vehicles in emerging mobility ecosystems. By addressing the long-standing challenge of conflict-free yellow interval management, the proposed NudgeLight framework contributes a significant step toward safer, more efficient, and human-centric intelligent transportation infrastructure.

REFERENCES

1. Chen, Q., & others. Autonomous vehicles in mixed-autonomy traffic: game theoretic human-like decision making in mixed traffic. 2024. on human-like decision methods for AVs in mixed traffic.
2. Dinneweth, J., et al. Multi-agent reinforcement learning for autonomous vehicles. 2022. survey of MARL for AVs and limitations regarding human-behavior modeling.
3. Guo, Y., et al. DRL-TP3: A learning and control framework for signalized intersections with autonomous vehicles. 2021. RL-based signal control for mixed traffic including CAVs.
4. Boyko, N., et al. Optimizing Traffic at Intersections With Deep Reinforcement Learning. 2024. DRL applied to intersection control.
5. Cao, N., et al. MPC-Based Method for Intersection Control in Mixed Traffic Considering AVs, HDVs, and Pedestrians. 2025. alternative control method in mixed traffic intersection management.
6. He, Y., et al. Comprehensive Evaluation on Traffic Safety of Mixed Traffic Flows in Merging Areas. 2025. safety assessment framework for mixed traffic, focusing on conflict risk.
7. Masri, S.; Ashqar, H.I.; Elhenawy, M. Large Language Models (LLMs) as Traffic Control Systems at Urban Intersections: A New Paradigm. 2025. exploring use of LLMs for traffic control tasks.
8. Nie, T., et al. LLM4TR: Roles of Large Language Models in Transportation Systems — A Survey. 2025. categorizing potential roles of LLMs in traffic/transportation systems.

9. Lamichhane, B.R.; et al. Context-aware decision making in autonomous vehicles via social behavior modeling fused with LLMs. 2025. recent work on fusing social behavior modeling and LLM reasoning in AV context.
10. Wang, D.; Li, W.; Zhu, L.; Pan, J. Learning to Control and Coordinate Mixed Traffic Through Robot Vehicles at Complex and Unsignalized Intersections. 2023. mixed traffic control via decentralized multi-agent RL in unsignalized intersections.
11. Yan, S.; Welschehold, T.; Büscher, D.; Burgard, W. Courteous Behavior of Automated Vehicles at Unsignalized Intersections via Reinforcement Learning. 2021. AV behavior design for mixed traffic at unsignalized intersections.
12. Wang, D.; Li, W.; Zhu, L.; Pan, J. Learning to Change: Choreographing Mixed Traffic Through Lateral Control and Hierarchical Reinforcement Learning. 2024. hierarchical RL for mixed traffic with lateral & longitudinal control.
13. Karbasi, A.H.; Yang, H.; Razavi, S. Exploring the Impact of Traffic Signal Control and Connected & Automated Vehicles on Intersection Safety: A Deep RL Approach. 2024. studying combined effect of CAVs + adaptive signal control on safety.
14. Li, X., et al. Modeling mixed traffic flows of human-driving vehicles and connected/autonomous vehicles: operational mechanism and safety analysis. 2023. — modeling mixed traffic flows, foundational to mixed human–AV traffic research.
15. Michailidis, P.; Michailidis, I.; Lazaridis, C.R.; Kosmatopoulos, E. Traffic Signal Control via Reinforcement Learning: A Review on Applications and Innovations. 2025. comprehensive review of RL methods in traffic signal control, covering methodology, simulation platforms, limitations.
16. Xiao, F.; et al. Advances in Reinforcement Learning for Traffic Signal Control. 2025. recent survey on RL-based traffic signal control techniques.
17. Feng, S.; Sun, H.; Yan, X.; Zhu, H.; Zou, Z. Dense Reinforcement Learning for Safety Validation of Autonomous Vehicles. 2023. on safety validation of AVs using dense RL; relevant to safety-constrained RL paradigm.
18. Wang, X.; et al. Traffic Light Optimization with Low Penetration Rate Vehicle Trajectory Data. 2024. integrating low penetration AV data for signal timing optimization; relevant for transitional mixed traffic phases.