



Bridging Transportation
Researchers (BTR)



USC University of
Southern California



MAPLE: Macro–Micro Aligned Pigouvian Lane Control for Mixed-Autonomy Traffic Efficiency

Kexin Wang^{a,*}, Gavin Huang^a, Zehao Wang^b, Jiachen Li^b, Ruolin Li^a


^a University of Southern California, CA, USA, ^b Georgia Institute of Technology, GA, USA

* kwang255@usc.edu, <https://kexinwang.me/>

August 12th, 2026

Highway weaving
creates severe lane
imbalance and localized
congestion.

Lane changes can
redistribute traffic, but
uncoordinated
maneuvers may further
disrupt traffic flow.



Related Work: Lane-Change Coordination for Bottleneck Congestion Management

Microscopic models

Describe immediate interactions. Their costs focus on local incentives.

- MOBIL^[1]: weighs acceleration gains against effects on followers.
- Ntousakis et al.[2] jointly optimize vehicle trajectories to coordinate freeway merging.
- Zhou et al.^[3] learn cooperative lane-changing policies through multi-agent actor–critic training.


Limited guarantee that learned policies will align with system efficiency goals

Macroscopic models

Distribute demand across lanes.

- Roncoli et al.^[4] develop multilane optimal control and feedback lane assignment.
- Pan et al.^[5] jointly optimize lane-change guidance, speed regulation, and ramp metering using model-based reinforcement learning.

Operate on aggregated traffic, do not directly influence individual actions.



[1] A. Kesting, M. Treiber, and D. Helbing, "General lane-changing model MOBIL for car-following models," *Transportation Research Record*, 2007.

[2] I. A. Ntousakis, I. K. Nikolas, and M. Papageorgiou, "Optimal vehicle trajectory planning in the context of cooperative merging on highways," *Transportation Research Part C: Emerging Technologies*, 2016.

[3] W. Zhou, D. Chen, J. Yan, Z. Li, H. Yin, and W. Ge, "Multi-agent reinforcement learning for cooperative lane changing of connected and autonomous vehicles in mixed traffic," *Autonomous Intelligent Systems*, 2022.

[4] C. Roncoli, M. Papageorgiou, and I. Papamichail, "Traffic flow optimisation in presence of vehicle automation and communication systems, Part I: A first-order multi-lane model for motorway traffic," *Transportation Research Part C: Emerging Technologies*, 2015.

[5] T. Pan, R. Guo, W. H. K. Lam, R. Zhong, W. Wang, and B. He, "Integrated optimal control strategies for freeway traffic mixed with connected automated vehicles: A model-based reinforcement learning approach," *Transportation Research Part C: Emerging Technologies*, 2021.

Related Work: Lane-Change Coordination for Bottleneck Congestion Management

Micro-Macro models

Yang et al.^[6] integrate tactical lane control with game-theoretic trajectory planning, using predicted maneuver costs to refine lane-flow redistribution.



This motivates a complementary **interface** that assigns each candidate maneuver an explicit marginal congestion cost while **retaining** the vehicle's existing motion controller.

guarantee that learned policies
will align with system efficiency goals




Operate on aggregated traffic,
directly influence individual actions.

Main Contributions

- **A modular macro–micro coordination framework.** We coordinate connected automated vehicles' (CAV) lane changes through V2I signals while retaining existing rule-based, learning-based, or model predictive controllers.
- **Predictive marginal congestion pricing.** We translate maneuver-induced density and speed changes in both source and target lanes into individual decision costs using gradients of an arrival-horizon congestion objective accounting for delay, speed deterioration, and excess occupancy.
- **One-pass equilibrium with bounded cost alignment.** We establish a pure Nash equilibrium of the assigned maneuver costs after one forward pass within a fixed decision epoch. A computable bound quantifies the discrepancy between individual cost comparisons and nonlinear predicted system-level congestion-cost comparisons.
- **Evaluation across traffic regimes and CAV penetration levels.** In 420 simulations across six demand regimes, MAPLE achieves the highest mean throughput and the lowest mean delay, queue, jerk, and harsh-braking rate among seven methods. At 50% CAV penetration, it increases throughput by **7.3–17.9%** and reduces delay by **9.5–21.3%** relative to four conventional baselines. US-101 replay further shows decreasing delay with higher CAV penetration.

Model pipeline




The lane-choice game

A *decision epoch* consists of one pass through the finite player set $\mathcal{I} = \{1, \dots, N^{CAV}\}$. Each player i selects a maneuver from a nonempty candidate set $\mathcal{M}_i \subseteq \{keep, left, right\}$ satisfying lane-adjacency and route constraints.

Fix the predicted arrival order $1 \prec \dots \prec N^{CAV}$, with deterministic tie breaking, and call vehicles $j \prec i$ the predecessors of i . Write $\mathbf{m} = (m_i)_{i \in \mathcal{I}}$ for the maneuver profile and $m_{\prec i} = (m_j)_{j \prec i}$ for the predecessor maneuvers. The lane-choice interaction is a finite cost-minimization game:

$$\mathcal{G} = (\mathcal{I}, \{\mathcal{M}_i\}_{i \in \mathcal{I}}, \{J_i\}_{i \in \mathcal{I}}),$$

$$J_i(\mathbf{m}) = J_i(m_i; m_{\prec i}).$$



The cost function design

Let $C_n^i(\rho, v)$ be a differentiable lane congestion cost for CAV i 's decision context. Its gradient at the keep-lane reference defines the marginal price:

$$\lambda_n^i = \nabla C_n^i(\bar{z}_n^i) = (\lambda_n^{i,\rho}, \lambda_n^{i,v})^\top$$

The two components measure the cost sensitivity to density and speed. The maneuver cost weights the corresponding state changes and sums over lanes:

$$J_i(m_i; m_{\prec i}) = \sum_{n \in \mathcal{L}} (\lambda_n^i)^\top \Delta z_n^i(m_i)$$

Lower values indicate a more favorable predicted congestion effect.

- 1 Predicted the arrival state $\rho_n^i, \bar{v}_n^i$
- 2 Design lane congestion cost $C_n^i(\rho, v)$
- 3 Compute marginal price $\lambda_n^{i,\rho}, \lambda_n^{i,v}$
- 4 Model state changes Δz_n^i caused by a maneuver
- 5 Evaluate the maneuver cost J_i for each admissible candidate $m_i \in \mathcal{M}_i$, where $\mathcal{M}_i \subseteq \{keep, right, left\}$, and select the minimum-cost maneuver m_i^*

Arrival-state prediction

1 Estimate arrival time T_i


The RSU estimates when each CAV will reach the reference boundary based on its remaining distance and the observed traffic speed.

2 Account for traffic changes $N_n^{in,x}$, $N_n^{in,y}$, $N_n^{out,x}$, $N_n^{out,y}$ during travel


Before the vehicle arrives, the predictor accounts for expected lane entries and exits, committed lane changes, and mandatory ramp movements.

3 Predict lane conditions ρ_n^i , $\bar{v}_n^i$ at arrival


These movements update the predicted density and speed of each lane. Lane changes also introduce a temporary speed disturbance that gradually dissipates.



Density (left) and speed (right) are paired with observations at the predicted arrival times



Design long-horizon congestion cost



State changes caused by a maneuver

A lane change removes one vehicle from its source lane and adds it to its target lane. Let $tgt(m)$ denote the target of $m \in \mathcal{M}_i$, with $tgt(keep) = n_i$. Using 1_E for the indicator of condition E , the density change is

$$\Delta \rho_n^i(m) = \frac{\mathbf{1}_{n=tgt(m)} - \mathbf{1}_{n=n_i}}{L_n}$$

Applying the same speed model $\bar{v}_n^i$ to the candidate and reference gives

$$\Delta \bar{v}_n^i(m) = \frac{\tau_i}{\tau_{rel}} [V_n(\rho_n^i + d_n) - V_n(\rho_n^i)] - \frac{\delta \bar{v}_n(t_i)}{L_n} \left[\frac{A_n^i + \Delta N_n^i}{\rho_n^i + d_n + \kappa} - \frac{A_n^i}{\rho_n^i + \kappa} \right]$$

Together, these effects define $\Delta z_n^i(m) = (\Delta \rho_n^i(m), \Delta \bar{v}_n^i(m))^T$ for the pricing rule.

Sequential Decisions

A vehicle's assigned cost depends on its predecessors:

$$J_i(\mathbf{m}) = J_i(m_i; m_{\prec i})$$


In the fixed order, each vehicle chooses

$$\boxed{m_i^*} \in \arg \min_{m_i \in \mathcal{M}_i} J_i(m_i; m_{\prec i}^*), \quad i = 1, \dots, N^{\text{CAV}}$$

Numerical Setting

Synthetic demand regimes; inflows are in veh/h. Table covers steady demand (D1-D3), congestion formation (D4), a demand reduction (D5), and a ramp pulse (D6).


Case	Mainline	On-ramp
D1	1200	450
D2	2400	900
D3	2880	1080
D4	1600 → 2640	600 → 990
D5	2640 → 1400	990 → 450
D6	2000	450 → 1200 → 450




D1-D3



D4



D5



D6

Performance Evaluation

AGGREGATE PERFORMANCE ACROSS SIX DEMAND REGIMES

Method	Throughput (veh/h)	Delay (veh·h)	Queue (veh)	Abs. jerk (m/s ³)	Harsh braking (events/1000 veh-km)	Actual LC (per run)
No LC	1718.5	6.836	61.00	0.208	265.2	19.4
Rule-based LC	1856.0	6.048	52.56	0.192	182.0	57.0
SUMO LC2013	1688.0	6.953	62.37	0.216	272.9	39.6
MOBIL-style	1717.3	6.846	61.28	0.212	258.7	44.3
Yang-style adapter	1962.9	5.625	47.76	0.179	119.7	46.0
Payment-based adapter	1750.0	6.696	59.62	0.209	233.8	59.6
MAPLE	1990.7	5.474	46.17	0.178	112.0	60.0

Values are means over six demand regimes and ten paired seeds (60 runs per method; 420 simulations in total). Higher throughput is better; lower delay, queue, jerk, and harsh-braking frequency are better. Actual LC is the mean number of completed lane changes per run. Bold indicates the best mean for each performance metric.

- No LC: suppresses discretionary CAV lane changes but retains route-mandatory movements.
- Rule-based LC: requires an ego acceleration gain above $0.2m/s^2$ and a target-lane follower acceleration no lower than $-4m/m^2$.
- SUMO LC2013 uses strategic and speed-gain motives, with cooperative and keep-right motives disabled.
- MOBIL-style rule: adds the target follower's acceleration change with politeness weight 0.3, using the same incentive threshold and follower-braking bound.

* *All methods with the same MAPPO lower-level controller.*

Performance Evaluation

AGGREGATE PERFORMANCE ACROSS SIX DEMAND REGIMES

Method	Throughput (veh/h)	Delay (veh·h)	Queue (veh)	Abs. jerk (m/s ³)	Harsh braking (events/1000 veh-km)	Actual LC (per run)
No LC	1718.5	6.836	61.00	0.208	265.2	19.4
Rule-based LC	1856.0	6.048	52.56	0.192	182.0	57.0
SUMO LC2013	1688.0	6.953	62.37	0.216	272.9	39.6
MOBIL-style	1717.3	6.846	61.28	0.212	258.7	44.3
Yang-style adapter	1962.9	5.625	47.76	0.179	119.7	46.0
Payment-based adapter	1750.0	6.696	59.62	0.209	233.8	59.6
MAPLE	1990.7	5.474	46.17	0.178	112.0	60.0

Values are means over six demand regimes and ten paired seeds (60 runs per method; 420 simulations in total). Higher throughput is better; lower delay, queue, jerk, and harsh-braking frequency are better. Actual LC is the mean number of completed lane changes per run. Bold indicates the best mean for each performance metric.

- Yang-style adapter weighs lane-density imbalance, maneuver frequency, and costs from predicted trajectory interactions.
- Payment-based adapter ranks maneuvers by the joint utility of the ego vehicle and target-lane follower.


* *All methods with the same MAPPO lower-level controller.*

Real-Scenario Evaluation

We replay the recorded US-101 periods. Because the data contain no CAVs, we randomly convert vehicles at penetration rates of 0%, 40%, and 100%.

- Mean delay decreases monotonically from 95.64s/veh at 0% penetration to 83.01s/veh at 40% and 66.59s/veh at 100%, reductions of 12.63% and 30.4%.
- CAV delay reaches approximately 66s/veh by 60% penetration, whereas HDV delay continues to decline through 80%.

CAV penetration (%)	All vehicles	CAV	HDV
0	95.64	–	95.64
20	89.64	74.02	93.54
40	83.01	70.41	91.41
60	75.95	66.10	90.72
80	70.36	66.63	85.26
100	66.59	66.59	–




Compatibility and Robustness of Lower-Level Execution

LANE-CHOICE REPLACEMENT WITH FIXED MPC PROXY; MEANS OVER D1–D6.



Metric	Yang-style	MAPLE	Change
Throughput (veh/h)	1807.27	1810.30	+0.17%
Delay (veh-h)	6.67	6.59	−1.20%
Harsh braking ^a	1657.40	1504.08	−9.25%
Mean jerk (m/s ³)	0.31	0.30	−3.22%

^aEvents per 1000 veh-km.



- The supplied point estimates show similar throughput (0.17%)
- Modestly lower delay (-1.20%)
- A larger reduction in harsh-braking frequency (-9.25%)
- Mean absolute jerk also decreases
- Compatibility with different lower-level controllers in scenarios D2 and D3.
- The evaluated lower-level controllers include a rule-based controller, MAPPO, and MPC.


Robustness to an Unannounced Emergency Stop




Without **prior warning**, MAPLE reacts through ordinary traffic feedback and substantially limits the propagation of the emergency-stop disturbance.



Extensions...



Network-Level MAPLE




Uncertainty-Aware Robust Analysis

Full paper forthcoming. Visit <https://kexinwang.me/> for updates.

SIAS Societally-aware Interactive Autonomous Systems Lab

- Unlocking the Potential of AVs in Intelligent Transportation Systems
- Re-engineering Transportation Systems for AV Integration
- Game Theory/Controls/Robotics/AI



Thanks!

