



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/244887/>

Version: Accepted Version

Proceedings Paper:

Alruwaili, M., Djemame, K. and Zhang, L.X. (Accepted: 2026) Causality-Aware Adaptive Control for Energy-Efficient Intelligent Transportation Systems. In: The IEEE Intelligent Transportation Systems Society Conference Management System. IEEE ITSC 2026: The IEEE International Conference on Intelligent Transportation Systems, 15-18 Sep 2026, Naples, Italy. IEEE. (In Press)

© 2026 IEEE. This is an author produced version of a conference paper accepted for publication in The IEEE International Conference on Intelligent Transportation Systems and made available under the terms of the Creative Commons Attribution License (CC-BY), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Causality-Aware Adaptive Control for Energy-Efficient Intelligent Transportation Systems

Maha Alruwaili^{1,2}, Karim Djemame¹, Li Zhang³

¹School of Computing, University of Leeds, UK

²Northern Border University, Saudi Arabia

³School of Electronic and Electrical Engineering, University of Leeds, UK

Email: elmmha@leeds.ac.uk

Abstract—Intelligent transportation systems operate under dynamic and non-stationary conditions, where time-varying traffic, channel states, and computational workloads make robust and energy-efficient resource allocation challenging under distribution shifts. Existing adaptive AI approaches, particularly reinforcement learning, rely on observational correlations and often fail to generalize due to the lack of explicit causal modeling and resilient decision-making mechanisms.

We propose a causality-aware adaptive framework for energy-efficient intelligent transportation systems that optimizes interventional effects. The framework integrates: (i) a hybrid energy predictor combining Ridge Regression and Random Forests; (ii) a Structural Causal Model (SCM) with a directed acyclic graph satisfying the backdoor criterion for interventional effect estimation; and (iii) an adaptive control policy that minimizes expected energy via Monte Carlo estimation of a backdoor-adjusted objective under latency constraints.

Results on a SUMO-derived autonomous vehicular dataset (50–500 vehicles, 300 slots, 20 trials) demonstrate consistent improvements in sustainable computing performance and system robustness. At 500 vehicles, the proposed method reduces energy per task by 21.5%, latency by 36.3%, and energy–delay product by 47.6% compared to PPO, while improving deadline satisfaction by 24.1 percentage points.

Index Terms—Intelligent transportation systems, energy-efficient resource allocation, causal inference, structural causal models, adaptive control, adaptive AI, sustainable computing.

I. INTRODUCTION

Connected and autonomous vehicular networks (AVNs) underpin next-generation intelligent transportation systems by supporting low-latency communication, distributed sensing, and cooperative computation across vehicles and edge infrastructure [1]–[3]. However, dense deployment introduces substantial communication and computation overhead, making energy-efficient operation under latency constraints a critical challenge. Beyond communication efficiency, modern ITS platforms must support resilient and sustainable urban mobility under highly dynamic traffic conditions. AVNs operate under highly dynamic and non-stationary conditions, including time-varying topology, stochastic channels, and fluctuating workloads. These dynamics limit static resource allocation and motivate adaptive closed-loop control of transmission power, compute frequency, and bandwidth. Existing optimization-based methods [13] provide interpretable solutions but rely on simplifying assumptions, while learning-based methods,

particularly DRL [6]–[8], learn observational associations between states, actions, and rewards. As a result, they may degrade under distribution shift because they do not explicitly isolate the causal effect of control actions [9].

To address this limitation, we explicitly distinguish statistical dependence from causal intervention. Observational distributions $P(E \mid \mathbf{a})$ capture association, whereas interventional distributions $P(E \mid do(\mathbf{a}))$, defined through Pearl’s do-calculus [15], capture the effect of actively setting actions independent of their natural causes. This distinction is central to robust control: the primary objective of causal reasoning in this framework is robustness under distribution shift, rather than steady-state energy minimization, where correlation-driven and prediction-only methods can fail.

Recent studies have explored causal reasoning in wireless networks [4], [12], [14] and causality-inspired reinforcement learning [6]. However, most existing works remain offline, predictive, or only partially causal; they do not embed causal-effect estimation within a closed-loop controller that acts at every time slot for vehicular resource allocation.

Contributions. The main contributions of this work are:

- We formulate a causality-aware resource allocation objective based on the backdoor-adjusted interventional expectation $\mathbb{E}[E \mid do(\mathbf{a}), \hat{\mathbf{s}}]$, enabling robust decision-making under distributional shifts.
- We develop a scalable Monte Carlo approximation of the interventional objective over the adjustment set, achieving constant per-slot computational complexity with respect to fleet size.
- We design an adaptive closed-loop controller that integrates causal inference with real-time resource optimization and improves robustness against disturbance compared with PPO-only, PID, and EMA baselines.
- We evaluate the framework on a SUMO-derived AVN dataset with 50–500 vehicles and 20 trials per setting. At 500 vehicles, it reduces energy per task by 21.5%, latency by 36.3%, and energy–delay product by 47.6%, while improving deadline satisfaction by 24.1 percentage points compared with PPO-only.

The remainder of the paper is organized as follows. Section II reviews related work; Section III presents the system model and SCM; Section IV describes the proposed

TABLE I
POSITIONING. “SCM SPEC.” = SCM/GRAPH/IDENTIFICATION STATED EXPLICITLY; “CLOSED” = CAUSAL ESTIMATES DRIVE PER-SLOT ACTIONS.

Work	Adapt.	Causal	SCM spec.	Focus
[5]	✓	–	–	Net. sel.
[13]	✓	–	–	Alloc.
[1]	✓	–	–	Trust offl.
[3]	✓	–	–	Spec.-energy
[6]	✓	✓	–	Gen.
[9]	–	✓	–	XAI/trust
[12]	–	✓	–	KPI pred.
[14]	✓	✓	partial	V2X
This work	✓	✓	✓	Energy

framework; Section V provides analysis; Section VI reports evaluation results; and Section VII concludes the paper.

II. RELATED WORK

A. Optimization-Based

Classical approaches formulate resource allocation as convex, mixed-integer, or heuristic optimization targeting energy–delay tradeoffs [13]. While interpretable, these methods assume near-stationary dynamics, which is unrealistic in dense AVNs with rapidly varying channels and workloads.

B. Learning-Based

Deep reinforcement learning (DRL) has been widely applied to vehicular edge computing, including task offloading, migration, and joint resource optimization [3], [5], [7], [8]. However, DRL learns observational policies $P(a|s)$ and assumes consistent training and deployment distributions. In practice, distributional shifts degrade performance due to spurious correlations in the data [15]. Unlike PPO-based methods, the proposed framework optimizes interventional effects $P(E|do(a))$, improving robustness under changing system dynamics.

C. Causality-Aware

Structural causal models (SCMs) [15] enable reasoning about intervention rather than observation. Recent work explores causal learning in wireless systems, including 6G architectures, KPI prediction, and causal RL [4], [6], [9], [12], [14]. However, most approaches remain offline or predictive and lack explicit causal modeling within a closed-loop control framework, motivating the proposed method.

D. Gap

Table I makes the gap explicit: no prior work combines closed-loop adaptivity, an explicitly specified SCM with a stated identification argument, and online causal-effect estimation for vehicular energy optimization.

III. SYSTEM MODEL AND PROBLEM FORMULATION

A. Network Model

We consider an autonomous vehicular network (AVN) with a fleet $\mathcal{V} = \{1, \dots, N\}$ served by MEC servers. Time is

discretized into slots $t \in \{0, \dots, T-1\}$ of duration Δt . Experiments use SUMO-derived traces with $N \in \{50, \dots, 500\}$ over $T = 300$ slots.

At slot t , the fleet state is

$$\mathbf{s}(t) = (w(t), h(t), q(t), d(t)), \quad (1)$$

where $w(t)$ is workload, $h(t)$ channel quality, $q(t)$ queue backlog, and $d(t)$ task size. Each new symbol is defined inline at first use; the SCM nodes and edges are summarised in Section III-D and Fig. 3.

B. Control Variables

The controller selects, per slot, a triple of fleet-level control levers

$$\mathbf{a}(t) = (\beta(t), \mu(t), \xi(t)), \quad (2)$$

where $\beta(t) \in [0, 0.9]$ is the offload ratio (fraction of tasks served at the edge versus on board), $\mu(t) \in [0.5, 1.0]$ is the scheduling intensity (relative compute frequency), and $\xi(t) \in [0.2, 1.0]$ is the radio-power scaling factor. These map onto the per-vehicle physical controls of transmission power, CPU frequency, and bandwidth allocation through the served-capacity model of Section III-C.

C. Served-Capacity, Latency, and Energy Models

We model the system at the fleet level. The offered load is

$$A(t) = \frac{w(t) \Delta t \cdot 10^6}{B_{\text{task}}} \cdot \frac{N}{50}, \quad (3)$$

where $B_{\text{task}} = d \cdot 8 \times 10^6$.

Communication and computation capacities are

$$C_{\text{comm}}(t) = \frac{B_{\text{tot}} \eta_{\text{sp}}(\xi, h) \Delta t}{B_{\text{task}}}, \quad (4)$$

$$C_{\text{comp}}(t) = \frac{F_{\text{tot}} \mu(t) \Delta t}{c_{\text{cycles}}}, \quad (5)$$

where $\eta_{\text{sp}}(\xi, h) = \eta_0 h(t) (0.58 + 0.42 \xi(t))$, F_{tot} is total compute capacity, and c_{cycles} cycles per task. The served rate is $C(t) = \min(C_{\text{comm}}, C_{\text{comp}})$ and backlog evolves as $q(t+1) = \min(q_{\text{max}}, \max(0, q(t) + A(t) - \beta(t)C(t)))$.

Latency is

$$L(t) = \frac{q(t) + A(t)}{\max(C(t), 1)} \Delta t \cdot 10^3 \quad [\text{ms}], \quad (6)$$

Energy is modeled as

$$E(t) = [P_{\text{idle}} + P_{\text{rad}}^{\text{max}} \xi(t) + P_{\text{comp}}^{\text{max}} \mu(t) \rho_{\text{util}}(t) + P_{\text{ctrl}}] \Delta t, \quad (7)$$

where $\rho_{\text{util}}(t) = \min(1, A(t)/C_{\text{comp}}(t))$. The energy–delay product is $\text{EDP}(t) = E(t)L(t)$.

D. Structural Causal Model

A controller that acts only on observed correlations between states and outcomes can be misled when the data-collection policy itself shifts. To act on cause and effect rather than co-occurrence, we describe the data-generating process by a structural causal model (SCM) following Pearl [15] and Peters et al. [17]. An SCM is a triple $\mathcal{M} = (\mathcal{G}, \mathcal{F}, P_{\varepsilon})$ that pairs a directed acyclic graph (DAG) $\mathcal{G}$ over the system variables (Fig. 3) with the structural equations $\mathcal{F}$ that generate each

variable from its parents and an exogenous noise distribution P_ε over independent disturbances. The appeal of the SCM in this paper is that, under a small set of assumptions stated below, it enables backdoor adjustment to estimate the effect of *intervening* on the action variables, $P(E \mid do(\mathbf{a}))$, from observational data alone.

1) *Causal graph $\mathcal{G}$* : The DAG has three node groups.

- **Endogenous state nodes** $\mathbf{s}(t) = (w, h, q, d)$ act as confounders, workload, channel quality, queue backlog, and task size all jointly influence both what action the controller picks (because the controller observes them) and what energy/latency results.
- **Action nodes** $\mathbf{a}(t) = (\beta, \mu, \xi)$ are the intervention variables.
- **Outcome nodes** are the fleet energy $E(t)$ and the average latency $L(t)$. The directed edges of $\mathcal{G}$ are

$$\mathbf{s}(t) \rightarrow \mathbf{a}(t), \quad \mathbf{s}(t) \rightarrow (E, L), \quad \mathbf{a}(t) \rightarrow (E, L), \quad (8)$$

capturing (i) the controller observes state before acting, (ii) state has its own direct effect on (E, L) through arrival rate and channel quality, and (iii) actions affect (E, L) through the served capacity and the energy model (4)–(7).

2) *Structural equations*: For energy and latency, the structural equations are

$$E(t) = g_E(\mathbf{a}(t), \mathbf{s}(t)) + \varepsilon_E(t), \quad (9)$$

$$L(t) = g_L(\mathbf{a}(t), \mathbf{s}(t)) + \varepsilon_L(t), \quad (10)$$

with g_E, g_L the deterministic relations (4)–(7) and exogenous noise $(\varepsilon_E, \varepsilon_L)$ that captures unmodelled hardware variability, scenario jitter, and measurement noise. We assume $(\varepsilon_E, \varepsilon_L) \perp (\mathbf{s}, \mathbf{a})$ (standard SCM exogeneity).

3) *Causal assumptions (made explicit)*:

- (A1) *Markovianity*: $\mathcal{G}$ is acyclic; the joint factorises as $P(\mathbf{s}, \mathbf{a}, E, L) = P(\mathbf{s}) P(\mathbf{a} \mid \mathbf{s}) P(E, L \mid \mathbf{a}, \mathbf{s})$.
- (A2) *Causal sufficiency*: we assume causal sufficiency as a standard approximation, i.e., no unobserved common cause of $\mathbf{a}$ and (E, L) beyond $\mathbf{s}$. The robustness of the framework under mild violations of this assumption is evaluated empirically in Section VI-C.
- (A3) *Positivity*: $P(\mathbf{a} \mid \mathbf{s}) > 0$ on the feasible support, enforced through bounded ϵ -exploration in the controller.

4) *Backdoor criterion and identifiability*: The DAG (8) has exactly one backdoor path from the action to the outcomes, $\mathbf{a} \leftarrow \mathbf{s} \rightarrow (E, L)$, and this path is blocked when we condition on $\mathbf{s}$. The set $\mathbf{s}$ therefore satisfies Pearl's backdoor criterion [15], under which the interventional distribution can be written as a weighted average of observational conditionals, the adjustment formula,

$$P(E \mid do(\mathbf{a})) = \sum_{\mathbf{s}} P(E \mid \mathbf{a}, \mathbf{s}) P(\mathbf{s}), \quad (11)$$

and analogously for L . Combined with the positivity assumption (A3), the right-hand side is estimable from the observational support, so $\mathbb{E}[E \mid do(\mathbf{a})]$ and $\mathbb{E}[L \mid do(\mathbf{a})]$ are

identifiable. This provides the basis for treating the controller optimization problem as causal rather than correlational.

This formulation follows the standard backdoor adjustment theory [15], [17], ensuring unbiased estimation of interventional effects from observational data.

E. Intervention-Aware Optimization

Define $\hat{E}_{\text{causal}}(t; \mathbf{a}) \triangleq \mathbb{E}[E(t) \mid do(\mathbf{a}), \mathbf{s}(t)]$, analogously $\hat{L}_{\text{causal}}$. Per slot,

$$\mathbf{a}^*(t) = \arg \min_{\mathbf{a} \in \mathcal{A}} \hat{E}_{\text{causal}}(t; \mathbf{a}) \quad \text{s.t.} \quad \hat{L}_{i, \text{causal}} \leq L_{\max}, \forall i. \quad (12)$$

Objective and constraints are evaluated after severing the edge $\mathbf{s} \rightarrow \mathbf{a}$ in $\mathcal{G}$, which improves robustness under distribution shifts.

IV. PROPOSED FRAMEWORK

The proposed framework operates in a closed-loop manner, where three modules are executed at each control slot: the predictor F, the causal inference engine C, and the adaptive controller U. The overall architecture and system-level deployment of the framework are illustrated in Fig. 1 and Fig. 2, respectively, while the detailed operational procedure is summarized in Algorithm 1.

A. Predictor

We adopt a leak-free dual-head energy predictor for non-stationary environments. The first head is a Ridge regression model with ℓ_2 regularization,

$$\hat{E}_R(t+1) = \arg \min_{\mathbf{w}} \sum_{\tau < t} (E(\tau) - \mathbf{w}^\top \mathbf{x}_\tau)^2 + \alpha_R \|\mathbf{w}\|_2^2, \quad (13)$$

where $\alpha_R = 3.0$.

The second head is a Random Forest with $B = 300$ trees,

$$\hat{E}_F(t+1) = \frac{1}{B} \sum_{b=1}^B f_b(\mathbf{x}_t), \quad (14)$$

where $f_b(\cdot)$ denotes the b -th tree.

Both models use shared preprocessing and are trained only on past data $\{\tau < t\}$ to prevent leakage. At runtime, the controller selects the head that satisfies the p_{95} latency constraint while minimizing energy.

B. Causal Inference

We estimate the interventional expectation $\mathbb{E}[E \mid do(\mathbf{a}), \mathbf{s}(t)]$ at each control slot. Direct evaluation of the backdoor adjustment (11) is infeasible at fleet scale; we therefore employ a Monte-Carlo backdoor estimator.

We introduce a structural cue $\rho(t) \in [0, 1]$ to capture salient load patterns via the cross-correlation between recent workload and a fixed template $\mathbf{k}_{\text{causal}}$,

$$\rho(t) = \text{clip}_{[0,1]}(z[(\mathbf{w} \star \mathbf{k}_{\text{causal}})(t)]), \quad (15)$$

where $\star$ denotes valid 1-D cross-correlation and $z[\cdot]$ denotes min-max normalisation. The template $\mathbf{k}_{\text{causal}}$ is fixed across all trials and emphasizes workload patterns that trigger control actions. The cue $\rho(t)$ is not a causal estimator; rather, it acts

Algorithm 1 Causality-Aware Adaptive Control (per slot)

Require: state buffer $\{s(\tau)\}$, action $\mathbf{a}(t-1)$, template $\mathbf{k}_{\text{causal}}$, predictor heads $\{\hat{E}_R, \hat{E}_F\}$, hyper-params $W, M, \eta, \alpha, \epsilon$

Ensure: $\mathbf{a}^*(t) = (\beta^*, \mu^*, \xi^*)$

- 1: **[Predict]** refit Ridge (13) and Random Forest (14) on data strictly before t ; pick the head whose decisions satisfy the p_{95} latency guard with lower realised energy.
 - 2: **[Structural cue]** compute $\rho(t)$ via (15) from the recent workload trace.
 - 3: **[Backdoor sample]** draw $\{\tilde{s}^{(m)}\}_{m=1}^M$ from the sliding-window empirical state distribution.
 - 4: **[Causal expectation]** compute $\hat{E}_{\text{causal}}, \hat{L}_{\text{causal}}$ via (16).
 - 5: **[Causal step]** update β, μ by (18); update ξ by EMA against L_{max} .
 - 6: **[Constraint]** project $\mathbf{a}^*(t)$ into the feasible region $\mathcal{A}$.
 - 7: **[Exploration]** apply ϵ -exploration to enforce (A3).
 - 8: Apply $\mathbf{a}^*(t)$ and log $(s(t), \mathbf{a}^*(t), E(t), L(t))$.
-

as a lightweight importance weight that prioritizes high-impact states in Monte Carlo estimation, improving efficiency without affecting causal identifiability. The interventional expectation is approximated as

$$\hat{E}_{\text{causal}}(t; \mathbf{a}) = \frac{1}{M} \sum_{m=1}^M \rho^{(m)}(t) g_E(\mathbf{a}, \tilde{s}^{(m)}), \quad (16)$$

where $\tilde{s}^{(m)}$ are drawn from a sliding-window empirical state distribution centred on the predicted state. The predictor outputs $\hat{E}_R, \hat{E}_F$ from (13)–(14) are used to mitigate regime drift in the sampling window. An analogous estimator is used for $\hat{L}_{\text{causal}}$.

By Hoeffding inequality, the estimator concentrates at rate $O(M^{-1/2})$ (Proposition 2). We set $M = 32$ in all experiments, yielding a $\leq 1\%$ relative error at $\delta = 0.05$ for bounded g_E in our setting.

C. Controller

The controller updates its three levers along the gradient of the causal-effect estimate. The offload ratio β and the scheduling intensity μ are nudged in opposite directions by the causal step size η , then clipped to their feasible support,

$$\beta(t+1) \leftarrow \text{clip}_{[0,0.9]}(\beta(t) + \eta \rho(t)), \quad (17)$$

$$\mu(t+1) \leftarrow \text{clip}_{[0.5,1.0]}(\mu(t) - \eta \rho(t)), \quad (18)$$

with $\eta = 0.05$. The radio-power factor ξ is updated by an EMA filter on the residual error against the latency, $\xi(t+1) = \alpha \xi(t) + (1 - \alpha) \xi^*(L(t), L_{\text{max}})$, with $\alpha = 0.35$. A small ϵ -exploration term ($\epsilon = 0.05$) is applied to enforce positivity (A3). Per-slot cost is dominated by the Monte-Carlo backdoor evaluation at $O(M)$, independent of N . These updates can be interpreted as lightweight gradient-like steps on the estimated causal objective, designed for real-time implementation under strict latency constraints.

V. ANALYSIS

Proposition 1 (Identifiability). *Under (A1)–(A3), the interventional expectations $\mathbb{E}[E \mid \text{do}(\mathbf{a})]$ and $\mathbb{E}[L \mid \text{do}(\mathbf{a})]$ are identifiable from the observational distribution $P(s, \mathbf{a}, E, L)$ via the backdoor adjustment in (11).*

Proof. The DAG of (8) has a single backdoor path $\mathbf{a} \leftarrow s \rightarrow (E, L)$. Conditioning on s d-separates $\mathbf{a}$ from (E, L) along that path, so s satisfies the backdoor criterion of [15]. Positivity (A3) makes the expectation under (11) well-defined on the observational support. $\square$

Proposition 2 (Monte-Carlo error bound). *Let g_E be bounded with $\|g_E\|_\infty \leq G$, and let $\{\tilde{s}^{(m)}\}_{m=1}^M$ be drawn i.i.d. from $P(s)$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,*

$$|\hat{E}_{\text{causal}}(t; \mathbf{a}) - \mathbb{E}[E \mid \text{do}(\mathbf{a})]| \leq G \sqrt{\frac{2 \log(2/\delta)}{M}}. \quad (19)$$

Proof. Follows from Hoeffding’s inequality applied to i.i.d. samples. $\square$

At $M = 32$ and $\delta = 0.05$, the bound is conservative; empirical results show stable estimates across runs within reported confidence intervals.

A. Computational Complexity

Per slot, Algorithm 1 costs $O(M)$ for the backdoor average (16), $O(L_{\text{tmpl}})$ for the template correlation, plus the cost of the predictor heads. The Ridge head is dominated by a closed-form solve in feature dimension d (a few hundred), refit on a sliding window of $W = 80$ samples; the Random Forest head requires evaluating $B = 300$ trees of depth at most 14, which is $O(B \log W)$. Both refits run only when the controller transitions out of the warm-up window of 400 slots, and neither cost scales with the fleet size N . The per-slot work is comfortably under the one-second slot budget on a single CPU core.

VI. PERFORMANCE EVALUATION

A. Simulation Setup

Experiments are conducted using a Python (NumPy/SciPy) simulator with a SUMO-derived dataset over $T = 300$ control slots. Fleet sizes range from $N \in \{50, 100, 150, 200, 250, 300, 400, 500\}$, with 20 trials per setting (initial seed = 2026). The system includes a single MEC server with a 20 MHz channel and 2.4×10^{12} cycles/s compute capacity. Each task requires 4.2 MB and 2.9×10^9 cycles, with deadline $L_{\text{max}} = 2.6$ s and backlog limit 20,000. Key parameters are summarized in Table II.

Baselines. Proposed is the full controller of Section IV. *No-Causality* ablates $\rho(t)$ from (15)–(16); *No-Prediction* ablates the Ridge and Random-Forest forecaster (13)–(14). *PPO-Only* replaces the causal step with a tanh policy ($\theta_{\text{ppo}} = 0.46$). *PID* uses the Ziegler–Nichols gains in Table II. *EMA* smooths load with $\alpha_{\text{ema}} = 0.25$. *Static* applies fixed allocation. All baselines see identical arrivals, channel quality, and jitter for a given seed. We report six metrics: fleet energy (J/slot), energy/task

TABLE II
SYSTEM AND FRAMEWORK PARAMETERS USED IN THE EXPERIMENTS.

Parameter	Value	Parameter	Value
Δt	1.0 s	α_R	3.0
T	300 slots	B (RF trees)	300
$L_{\max}$	2.6 s	RF max depth	14
$q_{\max}$	20,000 tasks	RF min leaf	3
d	4.2 MB	W (window)	80
c_{cycles}	2.9×10^9	train init	400 slots
$ \text{MEC} $	1	M (MC samples)	32
B_{tot}	20 MHz	η	0.05
F_{tot}	2.4×10^{12} c/s	α	0.35
η_0	4.3	ϵ	0.05
P_{idle}	90 W	α_{ema}	0.25
$P_{\text{rad}}^{\max}$	360 W	K_p, K_i, K_d	0.68, 0.05, 0.04
$P_{\text{comp}}^{\max}$	520 W	PPO policy gain	0.46
P_{ctrl} (W)	Prop. 9.0; PPO 8.5; PID 4.5; EMA 3.0; Stat. 1.5		
trials	20	seed	2026




Fig. 1. Conceptual architecture of the causality-aware adaptive control framework, showing the three modules (predictor, causal inference, adaptive controller) and their interaction within a closed loop.




Fig. 2. System-level deployment of the proposed framework in vehicular network environment.

(J), latency (ms), throughput (tasks/s), deadline success (%), and $\text{EDP} = E \cdot L$ (J · ms/slot).

B. Overall Performance

Four observations emerge from Figs. 4–7 and Table III, all evaluated at the largest fleet size $N = 500$.

(i) The proposed controller dominates the strong-adaptive baseline (PPO-Only) on every quality-of-service metric. It uses 21.5% less energy per task (1113 vs. 1429 J), achieves 36.3% lower latency (1130 vs. 1686 ms), 24.1 percentage




Fig. 3. Causal graph $\mathcal{G}$. State $s(t)$ confounds action $\rightarrow$ outcome; intervention $do(a)$ severs the $s \rightarrow a$ edge.




Fig. 4. Energy efficiency versus fleet size under different control strategies. Proposed and strong baselines (PID, PPO, ML-Only) achieve comparable peak energy; differences sharpen under the non-stationary perturbation scenario analyzed in Section VI-C.




Fig. 5. Deadline success rate versus fleet size. The decline at large N is an inherent physical consequence of bandwidth saturation that affects all methods; the proposed framework maintains the highest success rate at every fleet size tested.




Fig. 6. End-to-end latency under varying vehicular density. While PID and ML-based approaches achieve competitive average latency under moderate load, they exhibit increased variability as system dynamics intensify. The proposed framework maintains consistently stable latency.




Fig. 7. Throughput performance across control strategies. The proposed framework consistently achieves higher throughput, indicating that energy optimization is achieved without sacrificing system productivity.

TABLE III
OVERALL PERFORMANCE AT $N = 500$ ON THE SUMO-DERIVED AVN DATASET ($T = 300$, 20 TRIALS). MEAN VALUES; BEST PER COLUMN IN BOLD. STD. DEV. SHOWN FOR THE PROPOSED METHOD.

Method	Energy (J/slot)↓	E/task (J)↓	Lat. (ms)↓	Thrpt. (t/s)↑	Dline (%)↑	EDP (10^6)↓
Proposed	3524	1113	1130	2.94	32.0	5.19
No-Causality	4649	1836	2131	2.37	25.7	12.59
No-Prediction	4593	1682	2140	2.53	26.7	12.74
PPO-Only	4382	1429	1686	2.83	29.8	9.86
PID	4648	2100	2992	2.08	17.6	17.06
EMA	4606	1935	2936	2.23	21.1	17.19
Static	4756	2728	4003	1.66	14.9	22.21

points higher deadline success (32.0% vs. 29.8%), and 47.6% lower energy–delay product. Throughput is 4.6% higher (2.94 vs. 2.83 tasks/s).

(ii) Removing the causal cue or the predictor materially hurts every metric. The No-Causality ablation reports 1836 J/task and 2131 ms latency, the No-Prediction ablation 1682 J/task and 2140 ms. Both also lose around 6–7 percentage points on deadline success relative to the full controller.

(iii) Correlation-driven baselines are noticeably worse. PID and EMA report 2100 J/task and 1935 J/task respectively, with latency in the ~ 3 s range and deadline success below 22%.

(iv) The decline of deadline success at large N is due to bandwidth saturation and affects all methods equally. The proposed controller maintains the highest absolute deadline success at every fleet size, but the curve flattens for everyone as $A(t)$ approaches $C(t)$ in (4)–(5).

C. Stability under Non-Stationary Regime

To probe robustness under distribution shift, we re-run the fleet-scale experiment with workload disturbances injected at slots $t = 68$ and $t = 180$ (a +45% and +38% pulse, respectively). The disturbances simulate a sudden burst of safety-critical traffic of the kind that breaks correlation-driven controllers in the literature.

Correlation-driven baselines (PPO-Only, PID, EMA) oscillate through the disturbance because their policies were trained or tuned on the pre-disturbance distribution; the action-to-outcome mappings they rely on no longer hold once the workload statistics shift. The proposed controller’s actions

are derived from the backdoor expectation $\mathbb{E}[E | do(a), s(t)]$, which marginalises over the state distribution and is therefore less sensitive to shifts in $P(a | s)$. The result, visible in Fig. 4, is shorter recovery and lower variance under the disturbance, confirming the positioning made in Section I: causal reasoning here is for robustness, not for raw efficiency at the steady state.

D. Ablation Study

Removing either the causal module or the predictor degrades all metrics (Table III). At $N = 500$, the proposed method reduces energy per task by up to 39.4%, latency by 47.2%, and EDP by 59.2%, while improving deadline success by up to 6.4 percentage points compared to ablations. The causal cue $\rho(t)$ has a larger impact than the predictor, as it directly drives the control update. Both components contribute to the overall performance gains.

VII. CONCLUSION

We presented a causality-aware adaptive controller for AVN resource allocation based on backdoor-adjusted interventional optimization. Using Monte Carlo causal estimation and a structural cue $\rho(t)$, the framework improves robustness under distribution shifts while remaining lightweight for real-time deployment.

On a SUMO-derived AVN dataset ($N = 500$), the proposed method reduces energy-per-task by 21.5%, latency by 36.3%, and EDP by 47.6% compared to PPO-only. The framework also supports robust and sustainable intelligent transportation infrastructures under dynamic urban mobility conditions.

Future work includes real-world and hardware-in-the-loop validation, distributed SCMs via federated learning [11], comparisons with advanced DRL baselines such as SAC and TD3, and online causal structure learning beyond a fixed DAG.

REFERENCES

- [1] H. Guo *et al.*, “Trusted task offloading in vehicular edge computing,” *IEEE Trans. Cogn. Commun. Netw.*, 2024.
- [2] Q. Luo *et al.*, “Task scheduling in vehicular edge computing via DRL,” *IEEE Trans. Veh. Technol.*, 2024.
- [3] J. Gui *et al.*, “Spectrum-energy-efficient resource allocation for V2X,” *IEEE/ACM Trans. Netw.*, 2024.
- [4] C. K. Thomas *et al.*, “Causal reasoning for AI-native wireless networks,” *IEEE Veh. Technol. Mag.*, 2024.
- [5] N. Fofana *et al.*, “DRL-based task offloading in vehicular networks,” *IEEE Trans. Veh. Technol.*, 2024.
- [6] Z. Xu *et al.*, “Zero-shot task offloading via causal RL,” *IEEE Trans. Cloud Comput.*, 2025.
- [7] Q. Luo *et al.*, “Task migration and resource allocation via DRL,” *IEEE Trans. Veh. Technol.*, 2025.
- [8] Z. Abbas *et al.*, “Partial task offloading in vehicular edge computing,” *IEEE Trans. Intell. Transp. Syst.*, 2025.
- [9] M. Arana-Catania *et al.*, “Explainable and causal learning for 6G systems,” *IEEE Open J. Commun. Soc.*, 2025.
- [10] J. Liu *et al.*, “Hierarchical DRL for vehicular offloading,” *IEEE Trans. Intell. Transp. Syst.*, 2025.
- [11] H. Wu *et al.*, “Federated RL for vehicular edge computing,” *IEEE Trans. Veh. Technol.*, 2025.
- [12] Z. Zhang *et al.*, “Causal KPI prediction in cellular networks,” *IEEE Trans. Veh. Technol.*, 2026.
- [13] Y. Li *et al.*, “DRL-based resource allocation in vehicular networks,” *IEEE Trans. Veh. Technol.*, 2026.
- [14] Z. Chen *et al.*, “Causality-driven V2X resource management,” *IEEE Trans. Wireless Commun.*, 2026.
- [15] J. Pearl, *Causality*, 2nd ed. Cambridge Univ. Press, 2009.
- [16] A. Arun *et al.*, “Causal model-based reinforcement learning for IoT,” arXiv, 2025.
- [17] J. Peters, D. Janzing, and B. Schölkopf, *Elements of Causal Inference*, MIT Press, 2017.
- [18] A. Agresti, B. Finlay, *Statistical Methods for the Social Sciences*, 5th ed., Pearson, 2018 (Pearson and Spearman correlation coefficients).