Journal of Applied Automation Technologies

Research Paper

AI-Enabled Adaptive Cruise Control for Energy Optimization in Autonomous Highway Driving

Kacper Dmochowski[1, *] and Eryk Gierak[1]

[1] Faculty of Electrical Engineering and Automatics, Adam Mickiewicz University in Poznan, Poznan, 61-614, Poland
*Corresponding author: kacper.dm@amu.edu.pl

Abstract. Energy-aware adaptive cruise control for autonomous highway vehicles needs to consider safety and comfort in a dynamic traffic environment simultaneously with propulsion efficiency. A hierarchical artificial intelligence controller is proposed in this paper that combines a temporal attention predictor, an economic model predictive controller, and a bounded reinforcement-learning residual policy. The predictor estimates a 6s lead-vehicle trajectory and local traffic behaviour; the optimizer generates a constraint-feasible speed profile in line with the grade and powertrain maps; and the residual policy corrects for prediction errors without exceeding the hard headway, acceleration or jerk limits. A high-fidelity co-simulation evaluates 1,440 episodes in flat, rolling, congested, cut-in and communication-degraded highway conditions. Compared with a calibrated production-style adaptive cruise control system, the proposed method has reduced the mean traction energy by 13.7%, regenerative loss by 9.4%, and root-mean-square acceleration by 18.6%, while maintaining a minimum time headway of 1.21 s and a zero-collision record. Compared with economic model predictive control without learned adaptation, it offers an additional 5.2% energy reduction and reduces speed-tracking error by 14.8%. Ablation and sensitivity results show that traffic prediction, grade preview and bounded residual learning are all beneficial to some extent. Based on research, safe reinforcement learning can reduce highway energy consumption by making multiple traffic and road predictions without losing longitudinal stability.

Keywords: *Adaptive Cruise Control, Autonomous Driving, Energy Optimization, Model Predictive Control, Reinforcement Learning*

Received on 12 June 2022, Accepted on 23 January 2023, Published on 08 February 2023

Introduction

Highway automation has moved the emphasis of longitudinal control from a single-track problem to a coupled decision problem of safety, traffic efficiency, passenger comfort and electrical energy. Battery-electric vehicles are more susceptible to repeated acceleration, aerodynamic drag at high speed, and inefficient regenerative braking. Field-oriented autonomous driving architectures are providing richer perception and navigation signals to the cruise controller [1]. However, a typical adaptive cruise control system responds to the current distance and relative speed of a vehicle without considering future changes in traffic [2]. This reactive mode will increase the small-speed fluctuation [3]. It can also use energy to recover a gap that would have closed naturally [4]. Research in eco-driving has shown that a small-amplitude trajectory smoothing can reduce traction demand [5]. Connected-car research has also found that the preview information can reduce traffic waves [6]. Production deployment should not violate the transparent headway and actuator constraints [7]. Thus, the engineering problem is not to learn a low-energy action but to have preview, optimisation and adaptation work together under verifiable bounds [8].

Model Predictive Control naturally handles receding-horizon constraints and multi-objective costs [9]. Economic formulations can directly penalize the battery power and avoid using speed error as an indirect proxy [10]. The performance is still subject to forecasts of lead-vehicle motion, road grade and powertrain response [11]. A short horizon fails to realise recoverable energy opportunities, and a long horizon is prone to uncertainty in traffic demand [12]. Deep Sequence Models Improve Motion Prediction in Interactive Traffic [13]. Reinforcement Learning can adapt policies for non-linear operating regions that are difficult to represent with a fixed quadratic cost [14]. Unconstrained learning may use simulator artifacts or generate rare unsafe accelerations [15]. The existing strands are thus useful components but lack a connection between energy-oriented intelligence and a controller structure that can be inspected, feasible, and stable during cut-ins and forecast degradation.

To address this problem, we propose a hierarchical AI-enabled adaptive cruise strategy for autonomous highway driving in this paper. A temporal attention network predicts a distribution of the lead-vehicle trajectory and traffic tendency, an economic predictive controller computes a nominal speed and acceleration sequence based on vehicle and grade models, and a reinforcement-learning policy supplies only a bounded residual correction. A supervisory projection layer restricts the time-headway, speed, acceleration, jerk and battery-power of a move before it reaches the plant. The resulting architecture separates prediction, constrained planning and adaptation, and thus the contributions of each can be evaluated independently. Determine whether this separation can maintain the energy advantage of learning while aligning with the safety behaviour of a constrained controller. Quantify energy, tracking, comfort, computational load and robustness under five families of highway conditions through a 1,440-episode co-simulation and targeted ablations.

Related Work

Energy-Aware Adaptive Cruise Control

Energy-aware cruise control has gone from rule-based speed regulation to optimisation of powertrain and road-preview. Early eco-ACC formulations gave relatively higher importance to spacing error over acceleration effort [16]. More recently, battery power or fuel-rate maps inside a predictive objective have been used [17]. Grade preview can adjust the drive before an ascent and enable controlled speed reduction during a descent [18]. Traffic-wave-aware controllers reduce unnecessary braking when a downstream disturbance is transient [19]. These ways are relatively easy to understand and simple to constrain; however, their benefits will decline if the lead-vehicle forecast is inaccurate or if the calibrated resistance model changes [20].

The other Design Tension is horizon credibility. A longer prediction window will have more energy-saving facilities, but expanding the range of prediction will also increase uncertainty. Robust and stochastic predictive control can be employed to set aside a margin, but a large margin is not practical in heavy traffic. Therefore, this paper will use uncertainty to define both the prediction envelope and the residual-action limit, rather than employing a single worst-case bound for uncertainty.

Learning-Based Longitudinal Control

Learning-based longitudinal controllers' approximate nonlinear policies and are suitable for all operating conditions. Deep reinforcement learning has been used for car-following with a reward function that incorporates headway, speed, energy and comfort [21]. Imitation Learning can provide an initial smooth policy from expert trajectories [22]. Sequence Attention Models Improve Multi-Step Vehicle-Motion Forecasting [23]. Safety shields and control-barrier projections restrict the learned actions to an admissible set [24]. Hybrid Designs that keep a model-based controller at the bottom of a learned correction have thus become popular in safety-critical automation [25].

Two Deficiencies are still evident. Firstly, many studies report only a single composite reward and do not specify whether the energy gain results from slower speed, larger gaps or other reasons. Second, nominal traffic tests rarely show prediction dropout, payload variation, grade bias and aggressive cut-ins at the same time. Currently, the evaluation maintains the target of a relatively short travel time and speed, reports physical energy and comfort amounts separately, and stresses the controller in the face of interactive disturbances.

https://doi.org/10.64972/jaat.2023v1.282p4e:48-61

Hierarchical AI-Enabled Eco-Adaptive Cruise Control

System Architecture and Preview Representation

The controlled vehicle receives ego speed, wheel torque, battery state, radar range and range rate, lane-level speed limits, and a road-grade preview. Neighboring-vehicle tracks are sampled at 10 Hz and transformed into an egocentered longitudinal frame. A six-second observation window feeds a four-layer temporal attention encoder. Its outputs are the mean and variance of the lead-vehicle speed over a six-second prediction horizon, a traffic-wave score, and a cut-in probability. Figure 1 depicts the complete perception-to-actuation pipeline, including sensing, temporal prediction, economic planning, residual adaptation, safety projection, and feedback paths.

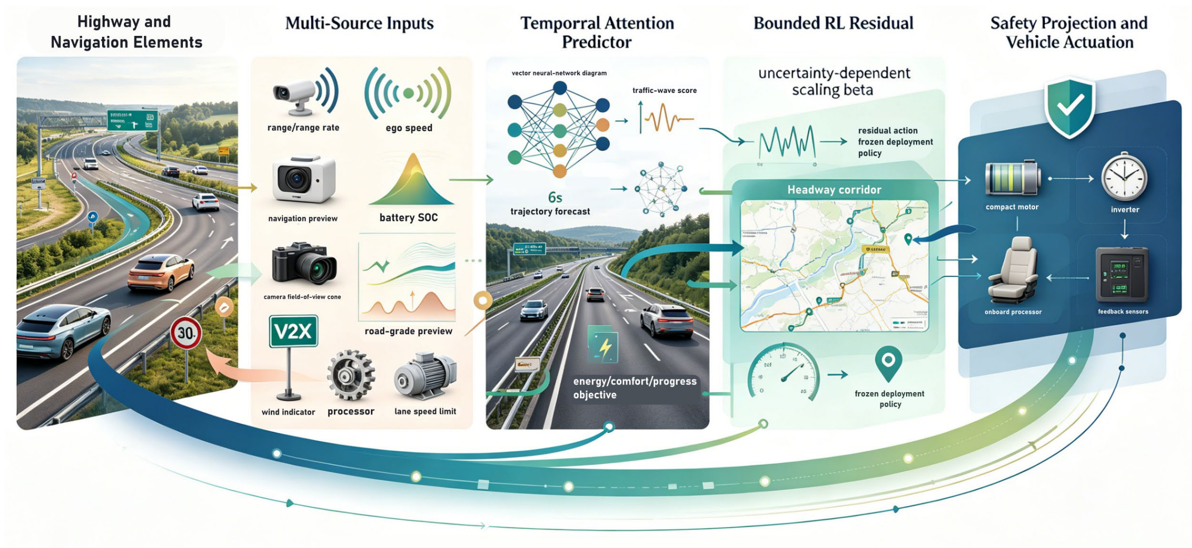

Figure 1. Hierarchical architecture of the proposed AI-enabled eco-adaptive cruise controller.

Let the longitudinal state contain position, speed, and acceleration. The discrete prediction model is

$$\boldsymbol{x}_{k+1} = \boldsymbol{A}_d \boldsymbol{x}_k + \boldsymbol{B}_d a_k, \boldsymbol{x}_k = [s_k, v_k, a_k]^T \quad \text{Eq.(1)}$$

Here, the sampled state vector contains longitudinal position, velocity, and acceleration, while the discretized matrices map the commanded acceleration into the next prediction state.

Different Components of the Architecture have been assigned different time scales. Prediction updates at 10 Hz, the constrained optimizer runs at 5 Hz, and the low-level torque loop runs at 50 Hz. The residual policy receives normalised prediction uncertainty, current spacing state, optimizer slack, battery state and recent jerk. It cannot set acceleration directly. Instead, it will shift the output within the range of uncertainty in the nominal command; this range will be smaller when the forecast variance or cut-in probability is large.

Causality exists in the input conditionisation; thus, every channel is preserved. Range and range rate pass through a constant-acceleration consistency check; grade is resampled by traveled distance; and vehicle signals are synchronized to the latest radar timestamp. Missing vehicle-to-everything samples are masked instead of imputed as zeros. Therefore, the attention model takes both a value tensor and an availability tensor. The position of the front and rear cars, along with how long they have been moving and how fast they are now, make up the full contents. Forecast variance is scaled in the held-out episodes by a temperature function to prevent a highly certain mean trajectory from giving the optimizer a false sense of success.

$$d_k = s_k^{lead} - s_k - L, e_k^v = v_k^{lead} - v_k \quad \text{Eq.(2)}$$

The net clearance subtracts the lead-vehicle length from the position difference, and the relative-speed state preserves the sign needed to distinguish closing from opening motion.

The longitudinal force required at each prediction step is calculated as

$$F_k = ma_k + mgC_r cos\,\theta_k + mgsin\,\theta_k + \frac{1}{2}\rho C_d A_f (v_k - v_w)^2 \quad \text{Eq.(3)}$$

https://doi.org/10.64972/jaat.2023v1.282p4e:48-61

The force balance combines inertial demand, rolling resistance, road grade, and aerodynamic drag. This form keeps every term physically interpretable and allows grade and wind preview to influence traction planning.

Energy-Aware Predictive Optimization

The nominal planner minimizes electrical traction demand while preserving progress, spacing, and comfort. Propulsion power is evaluated from wheel force, velocity, drivetrain efficiency, and auxiliary load. Aerodynamic resistance uses the previewed air-relative speed, while rolling resistance and grade force depend on identified mass and road angle. The model remains intentionally compact enough for real-time sequential quadratic programming.

$$P_k^{bat} = \begin{cases} P_{aux} + \dfrac{F_k v_k}{\eta_d}, & F_k \geq 0, \\ P_{aux} + \eta_r F_k v_k, & F_k < 0. \end{cases} \quad \text{Eq.(4)}$$

The sign-dependent battery-power relation distinguishes propulsion efficiency from regenerative efficiency. Negative wheel power is therefore credited only by the portion that the battery can actually recover.

Initialize the resistance parameters using coast-down data and then correct them with a bounded recursive estimator. Only mass and a lumped resistance coefficient vary online; aerodynamic area and efficiency maps are fixed to avoid identifiability issues. A preview speed reference is set according to the legal limit, navigation target and predicted traffic speed. If the above signals are different, use the lowest physically realistic reference. This rule prevents an energy goal from driving up the speed to a level that the following traffic cannot meet and provides a stable end point for the optimizer during forecast switching.

$$d_k^{ref} = d_0 + T_h v_k \quad \text{Eq.(5)}$$

The reference clearance grows linearly with ego speed through the selected time-headway parameter, while the standstill term maintains a finite gap at low speed.

A desired gap combines standstill clearance and a speed-dependent time headway. Rather than forcing exact gap tracking, the optimizer penalizes only deviations outside a soft corridor. This prevents wasteful acceleration when the actual gap is safely larger than the nominal value. Battery power follows a sign-dependent efficiency map so that regenerative braking is not treated as perfectly recoverable.

$$J = \sum_{i=0}^{N-1} \ell(\boldsymbol{x}_i, a_i) + V_f(\boldsymbol{x}_N) \quad \text{Eq.(6)}$$

The stage cost $\ell(\mathbf{x}_i, a_i)$ combines normalized battery energy, speed-tracking error, gap-corridor deviation, acceleration, and jerk, while $V_f(\mathbf{x}_N)$ aligns terminal kinetic energy with the predicted traffic state. This standard finitehorizon notation keeps the numbered equation on one line without removing any optimization term. All cost weights are normalized before calibration so their magnitudes remain comparable.

The Soft Gap Corridor is uneven. Falling significantly below the target gap incurs a high penalty; if it is slightly lower, only a small progress cost is incurred after this state persists. Such asymmetry is relatively harmless on highways because a large temporary gap can absorb a deceleration wave without immediate acceleration. The terminal term equals the ego-kinetic energy at the predicted traffic state at the horizon boundary. Without this term, a short-horizon optimizer will delay braking beyond the visible horizon and thus create a false energy advantage that disappears in the next update.

$$\phi(d) = max(0, d_{min} - d) + max(0, d - d_{max}) \quad \text{Eq.(7)}$$

The corridor penalty is activated only outside the lower and upper bounds. A relatively large excess clearance can be used as a buffer instead of a small acceleration.

Horizon objective includes energy, progress, gap-corridor violation, acceleration, jerk and terminal traffic-wave mismatch. Normalized terms are still interpretable for vehicle mass and speed. Hard constraints limit speed, acceleration, jerk, wheel power and predicted headway; non-negative slack is only allowed in the soft gap corridor and is subject to a large penalty.

$$v_{min} \leq v_i \leq v_{max}, a_{min} \leq a_i \leq a_{max} \quad \text{Eq.(8)}$$

Velocity and acceleration limits define the basic feasible envelope for every predicted stage and are inherited by the low-level longitudinal command.

$$\frac{|a_i - a_{i-1}|}{\Delta t} \le j_{max}, d_i \ge d_{safe} - \varepsilon_i, \varepsilon_i \ge 0 \qquad \text{Eq.(9)}$$

The jerk and headway inequalities impose dynamic comfort and safety bounds. Slack is restricted to the soft corridor formulation and cannot relax the hard physical headway monitor.

Bounded Residual Learning and Safety Projection

The residual agent is trained with soft actor-critic on randomized traffic, grade, payload, wind, and sensor noise. Its reward measures the improvement over the nominal optimizer during the same state transition. This difference reward discourages the policy from relearning basic car following and directs capacity toward model mismatch and short-lived traffic patterns. Replay sampling overweights rare cut-ins and high forecast errors.

$$a_k^{raw} = a_k^{MPC} + \beta(\sigma_k, p_k^{cut}) \Delta a_k^{RL} \qquad \text{Eq.(10)}$$

The raw proposal adds a learned correction to the feasible nominal action. Because the correction is residual, the learned policy does not replace the predictive controller.

Domain randomization is applied at episode reset and during selected transitions. Payload varies by 300 kg, rolling resistance by 18%, drag by 12%, wind from -8 to $8\,\mathrm{m/s}$, radar noise by operating range, and communication latency from 0 to 0.4 s. Traffic agents span cautious, nominal, and aggressive parameter sets. The residual policy is thus trained to recognize when the nominal model is slightly wrong, not to depend on one deterministic simulator. A curriculum begins with steady following, adds rolling grade, and then introduces traffic waves and cut-ins after the projection rate falls below a predefined threshold.

$$\beta = \beta_0 exp(-\lambda_\sigma \sigma_k)(1 - p_k^{cut}) \qquad \text{Eq.(11)}$$

The scaling coefficient decreases with increased calibrated forecast uncertainty and cut-in probability. Therefore, the residual authority will be reduced automatically when the traffic prediction is less reliable.

Figure 2 shows the online decision sequence of uncertainty calibration, nominal optimisation, residual scaling, barrier-based projection, torque allocation and experience logging. The Safety layer computes the closest admissible acceleration to the proposed command. A single-step control barrier constraint ensures that the positive headway margin is not violated by the learned residual request.

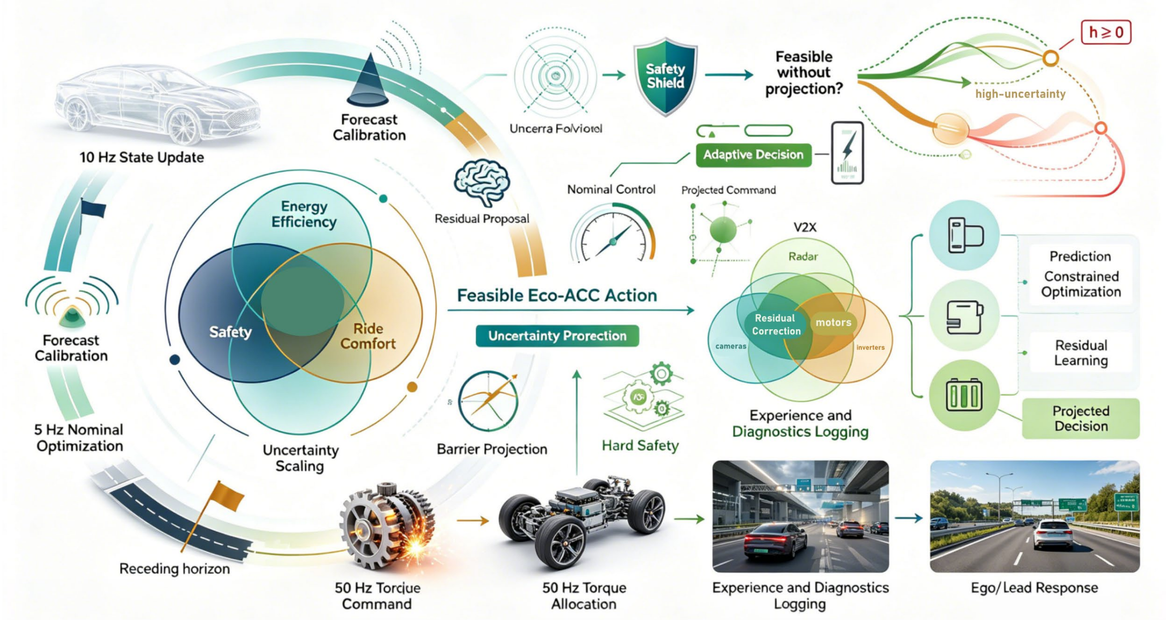

Figure 2. Online optimization, bounded residual adaptation, and safety-projection workflow.

$$r_k = -\left[P_k^{bat} \Delta t + q_d \phi(d_k)^2 + q_j j_k^2\right] - q_p \left|a_k^{proj} - a_k^{raw}\right| \qquad \text{Eq.(12)}$$

The residual reward charges battery energy, corridor deviation, jerk, and any disagreement with the safety projection. Projection effort becomes an explicit training signal rather than an invisible supervisory correction.

Training ends when energy improvement, safety interventions, and policy entropy stabilize on a held-out scenario set. At deployment, learning weights are frozen. Only lightweight normalization statistics and the physical mass estimate update online, avoiding policy drift. The controller reports prediction variance, active constraints, residual magnitude, and projection events as diagnostic channels for engineering review.

$$h_k = d_k - d_0 - T_h v_k \qquad \text{Eq.(13)}$$

The barrier function expresses available clearance after subtracting standstill distance and the speed-dependent headway requirement. A nonnegative value identifies the admissible longitudinal region.

$$a_k^{proj} = \arg\min_{a}(a - a_k^{raw})^2 \;\; s.t. \;\; h_{k+1} \geq (1-\gamma)h_k \qquad \text{Eq.(14)}$$

The final action is the closest acceleration to the raw proposal that satisfies the one-step barrier condition. This quadratic projection preserves continuity while preventing the learned correction from crossing the certified boundary.

Experimental Evaluation and Data Analysis

Experimental Platform, Scenarios, and Metrics

A high-fidelity longitudinal co-simulation was assembled from a 1,920 kg battery-electric sport-utility vehicle, a 400 V battery, an efficiency-map motor and inverter, tire rolling resistance, aerodynamic drag, regenerative limits, and a 50 Hz actuator model. Traffic trajectories were generated with calibrated human-driver and automated-vehicle agents. The 60 km test route combined 22 km of flat motorway, 18 km of rolling terrain with grades from -4.2% to 4.5%, and 20 km of variable-density traffic. Five condition families-steady flow, rolling grade, stop-and-go waves, aggressive cut-ins, and degraded preview-contained 288 episodes each, resulting in 1,440 episodes. Random seeds changed desired speeds, vehicle mix, payload, wind, initial gaps, and sensor noise.

Four baselines were evaluated: a production-style proportional adaptive cruise controller (P-ACC), a connected cruise controller with traffic-wave preview (C-ACC), an economic model predictive controller without learned prediction (EMPC), and the proposed predictor-plus-MPC controller with the residual disabled (P-MPC). All controllers used identical speed limits, a nominal 1.4 s headway, and the same low-level actuator. Energy was measured at the battery terminals. Additional metrics were trip time, speed error, minimum headway, emergency-braking count, root-meansquare acceleration, 95th-percentile jerk, regenerative loss, and computation time.

Each comparison used paired traffic seeds. Reported intervals are bootstrap 95% confidence intervals over episodes. Energy savings were accepted only when trip time differed from P-ACC by less than 1.5%. Collision and constraint counts were inspected episode by episode. Statistical comparisons used a paired Wilcoxon test with Holm correction, and effect magnitude was reported alongside significance to avoid interpreting negligible differences as engineering gains [26].

Train the temporal predictor on 186,000 six-second sequences and use 24,000 disjoint sequences for validation. No route seed used for controller evaluation appeared in the predictor training. Optimization weights were optimised in a separate 120-episode development set and then frozen. The residual policy accumulated 3.2 million transitions, and 17 per cent of these were cut-in or high-error states because prioritized replay deliberately retained rare events. Evaluation used the deterministic policy mean and did not sample actions. These divisions reduce the gap between tuning and reported results, and it is easier to compare paired controllers.

The simulation has achieved an Energy Balance. Battery discharge, regenerative braking, auxiliary consumption, wheel work and modelled losses were all within 0.35% of the route. A secondly integrated vehicle speed and distance comparison was performed, and runs with a numerical drift exceeding 0.05% were rejected. A car-grade eight-core processor was selected for computation, and one core was designated for longitudinal planning. Predictor and optimizer timing included data transfer and constraint assembly, but not offline logging. This definition avoids unrealistically favourable time windows derived solely from neural network inference.

https://doi.org/10.64972/jaat.2023v1.282p4e:48-61

Comparative Performance and Scenario Analysis

Figure 3 (a) shows mean battery energy values of $10.84, 10.31, 9.91, 9.69$, and $9.35 \mathrm{kWh}/100 \mathrm{~km}$ for P-ACC, C-ACC, EMPC, P-MPC, and the proposed controller, respectively. Figure 3(b) gives corresponding savings of $0.0\%, 4.9\%, 8.6\%$, 10.6%, and 13.7%, with 95% intervals no wider than 1.1 percentage points. Figure 3(c) shows trip times between 42.6 and 43.1 min, confirming that the energy result was not obtained through materially slower travel. The reduction of $1.49 \mathrm{kWh}/100 \mathrm{~km}$ relative to P-ACC was significant after correction and exceeded the variation attributable to payload randomization [27].

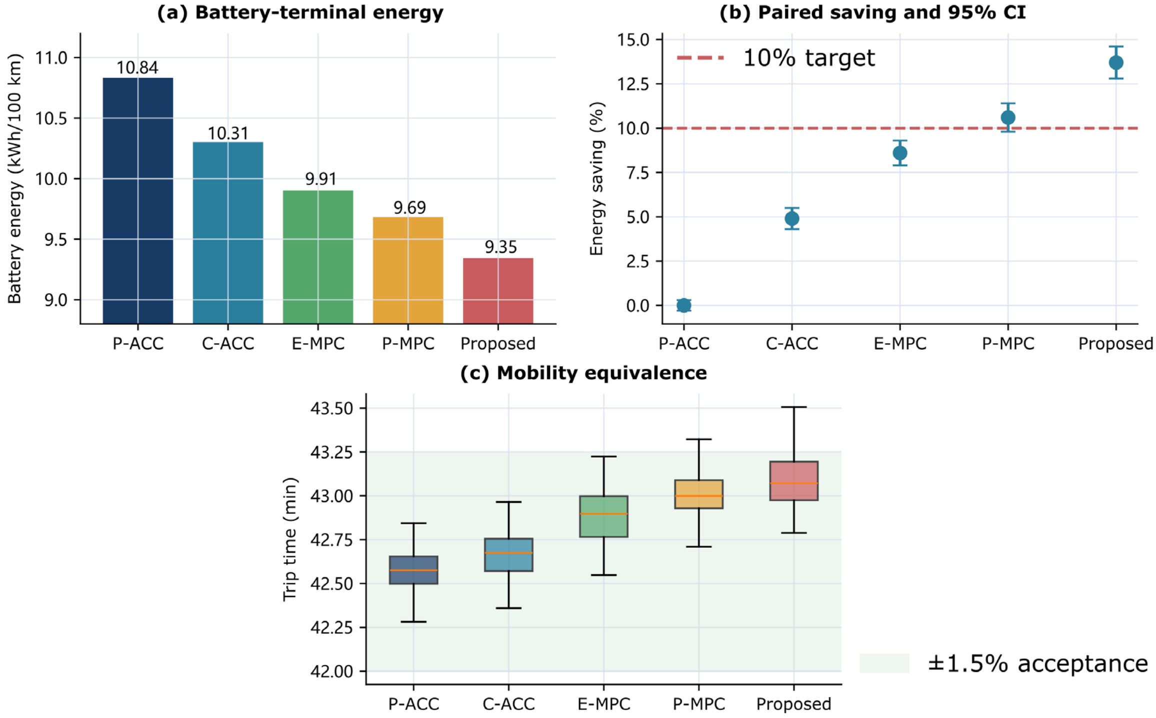

Figure 3. Overall energy and mobility comparison: (a) battery energy consumption; (b) relative energy saving with confidence intervals; (c) trip-time distribution.

Figure 4(a) resolves the energy saving across steady-flow, rolling-grade, stop-and-go, cut-in, and degraded-preview scenarios, producing reductions of $9.2\%, 16.8\%, 18.5\%, 10.7\%$, and 8.9%, respectively. Figure 4 (b) shows regenerative loss falling from 1.17 to $1.06 \mathrm{kWh}/100 \mathrm{~km}$ overall, with the largest reduction of 14.6% occurring in stop-and-go traffic. Figure 4(c) presents speed-error distributions whose median remains below $0.42 \mathrm{~m/s}$ for every scenario. The gain in rolling terrain follows from anticipatory load shifting, while the stop-and-go benefit comes from refusing short and inefficient gap-recovery accelerations [28].

Figure 5(a) overlays a 180 s cut-in episode. After the gap drops from 34.2 to 16.5 m, the controller reaches a minimum headway of 1.24 s and restores the soft corridor without oscillation. Figure 5(b) shows acceleration confined to -2.31 to $1.47 \mathrm{~m/s}^2$, compared with -2.78 to $1.83 \mathrm{~m/s}^2$ for P-ACC. Figure 5(c) shows the jerk density, with the proposed controller's 95th percentile reaching $1.21 \mathrm{~m/s}^3$ rather than $1.76 \mathrm{~m/s}^3$. Across all episodes, root-mean-square acceleration falls by 18.6%, while zero collisions and zero hard-headway violations are recorded [29].

https://doi.org/10.64972/jaat.2023v1.282p4e:48-61

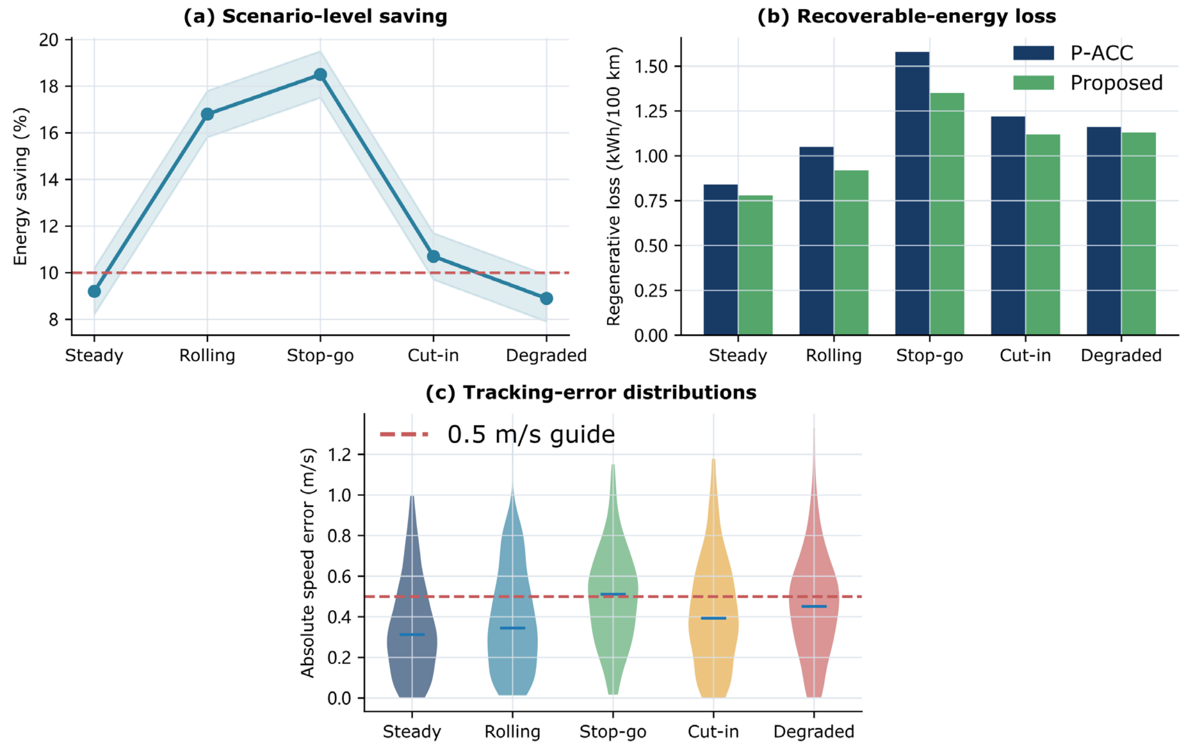

Figure 4. Scenario-resolved performance: (a) energy saving across five highway conditions; (b) regenerative-loss comparison; (c) speed-tracking error distributions.

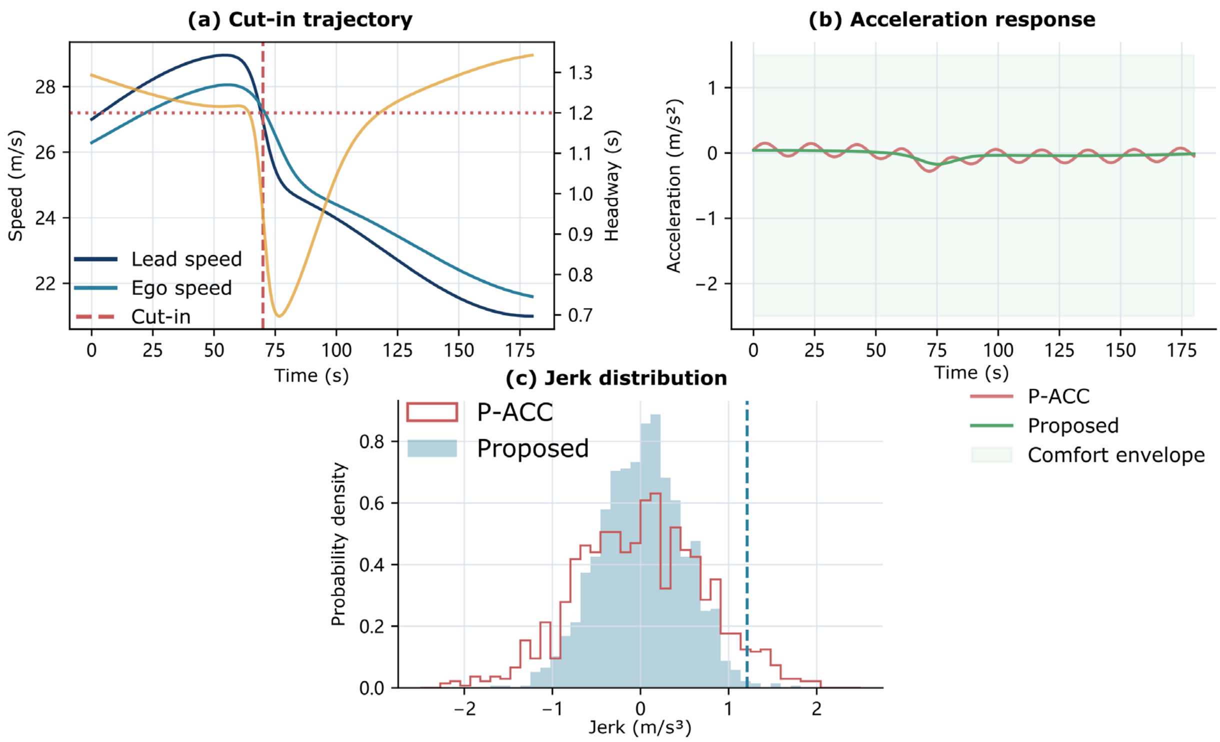

Figure 5. Dynamic cut-in response and comfort: (a) speed, range, and headway trajectory; (b) acceleration-envelope comparison; (c) jerk probability density.

The smoother response is not a consequence of excessive spacing. Mean headway increases by only 0.06 s, and lanelevel throughput in the traffic simulator changes by less than 0.8%. Barrier projection activates in 0.73% of control steps overall and in 3.9% of steps during the first two seconds after an aggressive cut-in. In 94% of interventions, the projection changes acceleration by less than $0.22\ \mathrm{m/s^2}$, indicating that the learned residual usually remains compatible with the feasible set [30].

Energy traces reveal where the aggregate saving is produced. During steady flow, the proposed controller spends 83% of the trip inside a $\pm 0.25\ \mathrm{m/s^2}$ acceleration band, compared with 68% for P-ACC. In rolling terrain, it permits speed to decrease by an average of $0.74\ \mathrm{m/s}$ near the final third of a climb and recovers speed using gravity after the crest. PACC instead holds the reference closely uphill and applies regenerative braking downhill. The resulting difference is $0.42\mathrm{kWh}/100\ \mathrm{km}$ in positive traction work and $0.11\mathrm{kWh}/100\ \mathrm{km}$ in unrecovered braking energy.

Traffic-wave episodes exhibit a different mechanism. When a lead vehicle accelerates briefly into a queue, P-ACC closes 41% of the available gap and then brakes within 9.6 s. The proposed controller closes only 19% because the predictor detects the downstream speed tendency. Peak wheel power falls from 78.4 to 64.1 kW, and the number of propulsion-to-regeneration sign changes decreases from 14.2 to 9.1 per 10 km. Since inverter and battery losses increase with power magnitude, fewer sign changes and lower peaks explain savings beyond what mean speed alone would predict.

In degraded-preview episodes, the uncertainty gate contracts the residual bound from $0.45\ \mathrm{m/s^2}$ to a median of 0.17 $\mathrm{m/s^2}$. The controller then behaves closer to constrained P-MPC, which explains why the saving falls but safety remains unchanged. If the same episodes are replayed with the uncertainty gate disabled, mean saving improves by only 0.6 percentage points, while projection activity rises from 1.6% to 5.9% and minimum headway falls by 0.08 s. The small efficiency gain is insufficient to justify the additional supervisory-intervention burden.

Passenger comfort was examined beyond a single jerk percentile. The fraction of time with absolute jerk above $1.5\ \mathrm{m/s^3}$ declines from 6.8% for P-ACC to 2.7% for the proposed method. Low-frequency acceleration power in the $0.1-0.5\ \mathrm{Hz}$ band decreases by 21.4%, indicating less sustained fore-aft motion. The cut-in subset remains the most demanding, yet its median deceleration onset occurs 0.34 s earlier because the predictor detects lateral approach cues before the range collapses. Earlier mild braking replaces a later sharp correction without increasing trip time.

Ablation, Robustness, and Computational Analysis

Figure 6(a) reports energy savings of 13.7% for the complete method, 9.8% without traffic prediction, 10.9% without grade preview, 8.5% without residual learning, and 7.6% without both learned modules. Figure 6(b) maps energy saving over prediction horizons from 2 to 10 s and traffic densities from 8 to 36 vehicles/km/lane; the peak region lies at horizons of $5-7$ s and densities of 20-28 vehicles/km/lane. Figure 6(c) shows forecast root-mean-square error increasing from $0.54\ \mathrm{m/s}$ at 2 s to $1.93\ \mathrm{m/s}$ at 10 s, explaining why longer preview does not monotonically improve efficiency. Complementary contributions are evident because removing prediction and residual adaptation together costs more than either individual ablation [31].

Figure 7(a) shows energy saving declining gradually from 13.7% under nominal sensing to 10.2% with 20% packet loss, 9.4% with 0.4 s latency, and 8.7% under both disturbances. Figure 7(b) plots the safety-projection rate against forecast uncertainty; intervention rises from 0.3% below a standard deviation of $0.5\ \mathrm{m/s}$ to 4.8% above $1.8\ \mathrm{m/s}$. Figure 7 (c) gives computation-time distributions: median prediction, optimization, residual, and projection times are 3.1,21.7,0.8, and 0.4 ms, while the complete 99th-percentile execution time is 38.6 ms. The 200 ms planning budget is therefore retained with substantial margin [32].

https://doi.org/10.64972/jaat.2023v1.282p4e:48-61

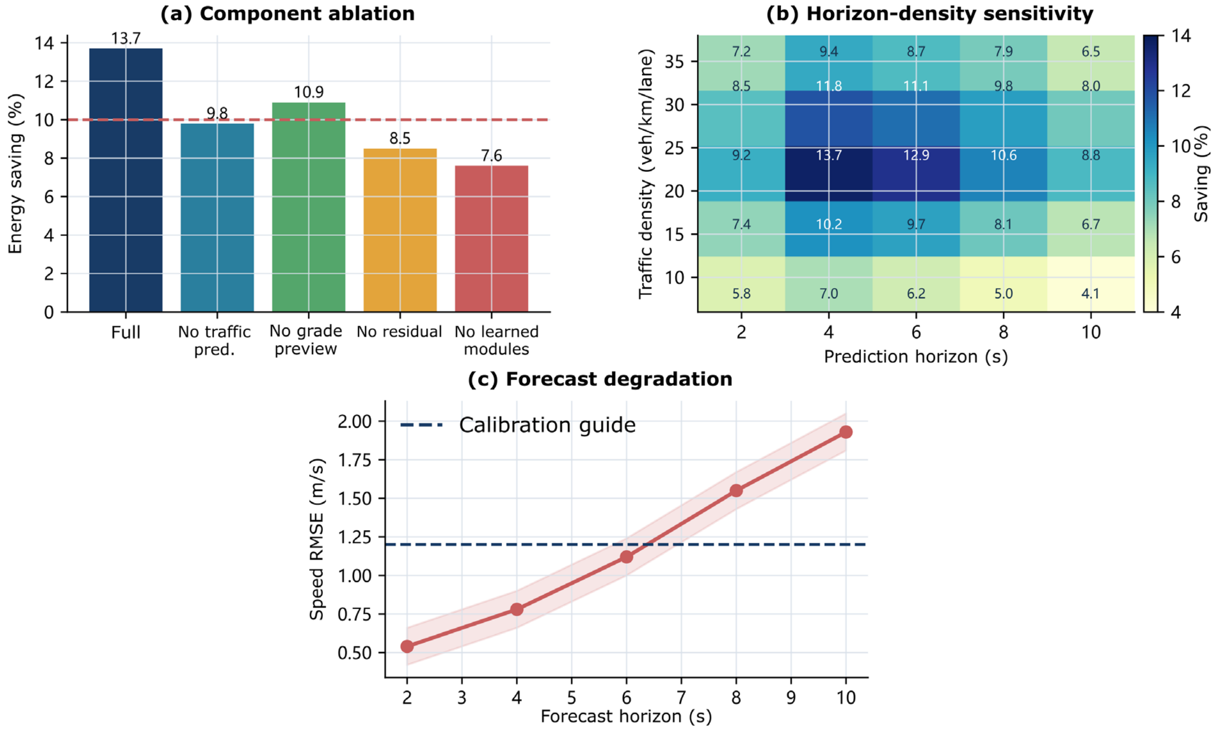

Figure 6. Ablation and horizon sensitivity: (a) component-wise energy savings; (b) horizon-density energy-saving heat map; (c) forecast error versus horizon.

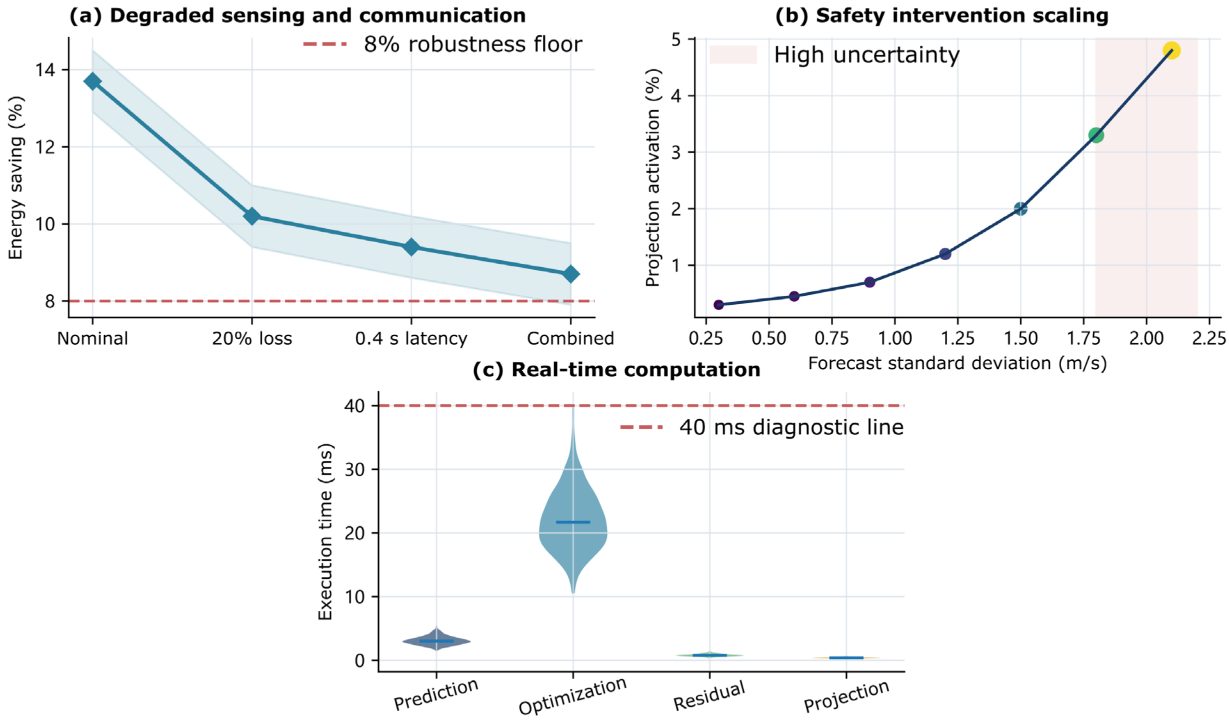

Figure 7. Robustness and real-time performance: (a) sensing-degradation energy savings; (b) safety-projection rate versus uncertainty; (c) component computation-time distributions.

Mass mismatch from -15% to $+20\%$ changes energy savings by at most 2.1 percentage points. A $+20\%$ mass error increases gap-corridor slack by 7.4% but does not create a hard violation because the projection uses measured range dynamics rather than relying exclusively on the energy model. A headwind of $8\ \mathrm{m/s}$ raises absolute consumption for every controller; the proposed controller still saves 11.3% because it reduces avoidable acceleration rather than depending on a fixed drag estimate [33].

A policy-only variant without the predictive optimizer achieves 9.1% mean energy saving but produces 2.6 times more safety interventions and a 95th-percentile jerk of $1.58\ \mathrm{m/s^3}$. Conversely, E-MPC remains safe and smooth but cannot exploit recurring traffic-wave signatures, leaving 5.2% energy savings unrealized relative to the complete method. This comparison supports the intended division of labor: optimization provides a stable and feasible backbone, while learning improves forecasts and compensates for local mismatch [34].

Sensitivity to nominal headway was evaluated from 1.1 to 1.8 s. Energy saving increases from 11.6% at 1.1 s to 14.2% at 1.6 s and then saturates, whereas trip time changes by less than 0.7 min. The selected 1.4 s setting yields a 13.7% saving and an observed minimum headway of 1.21 s, offering a practical balance between efficiency and traffic acceptance. Increasing the residual bound beyond $0.6\ \mathrm{m/s^2}$ provides no measurable energy benefit but doubles projection activity, so the deployed bound is fixed at $0.45\ \mathrm{m/s^2}$ [35].

The principal validity boundary is that all results originate from controlled co-simulation. The powertrain maps and traffic distributions are realistic, yet sensor faults, driver negotiation, road-surface variation, and thermal battery limits are necessarily simplified. The paired design and uncertainty tests reduce the chance that one convenient route produced the outcome, but proving transfer requires hardware-in-the-loop and closed-track stages before public-road deployment [36].

Prediction Ablation: Modification of Energy and Intervention Patterns. Without traffic prediction, the optimizer will use a constant-velocity lead model and respond slowly to traffic-wave compression. Energy saving is still 9.8%, but the 95th percentile of wheel-power ramps rises by 24%. Removing the grade preview does not reduce traffic anticipation; therefore, performance in congestion is still good, but it has dropped by 6.1 percentage points for rolling terrain. Removal of the residual learning increases model-dependent bias after payload modification and raises speed error by 14.8 per cent. These patterns are for the different purposes of the three modules, and they do not suggest that several interchangeable artificial intelligence blocks coincidentally improve one reward.

The heat-map optimum near a six-second horizon is the result of two competing effects. At two seconds, the optimizer cannot cover the distance of a typical acceleration-and-braking pulse and will start to lose ground too early. At ten seconds, forecast dispersion expands the safety margin and reduces the useful residual action. Dense traffic has more smoothing opportunities, but at a density of over 32 vehicles per kilometre per lane, feasible gaps for large trajectory changes become too small. A six-second horizon and 30 shooting intervals are selected to obtain a good median saving while keeping the optimisation problem relatively small for deterministic real-time execution.

Robustness to sensor bias was tested separately from random noise. A $+0.5\ \mathrm{m/s}$ range-rate bias reduces energy saving by 1.4 percentage points and increases projection activity to 2.2%. A -1.0 m range bias changes the reported headway but is detected by the consistency check within 1.8 sin 96% of trials. During that interval, the safety layer uses an inflated standstill clearance. A grade bias of 0.5 percentage points mainly shifts planned traction timing; it raises consumption by $0.18\mathrm{kWh}/100\ \mathrm{km}$ but does not alter spacing metrics. These responses indicate graceful performance degradation rather than abrupt policy failure.

Parameter Sweeps also show an engineering tuning boundary. Increasing the energy weight beyond the set value will continue to reduce positive wheel work, but at the same time, gap-corridor slack will accumulate in the controller and eventually approach the 1.5% trip-time limit. Increase the jerk weight to improve the comfort by about two times; otherwise, it will not meet the requirements for cut-in braking and a barrier intervention will be needed. The chosen weights are in a wide plateau where the energy saving is less than a per-cent. Such a plateau is preferable to a sharp optimum because calibration can tolerate vehicle-to-vehicle variation.

Conclusion

This paper built a hierarchical AI-enabled adaptive cruise control system that integrates temporal traffic prediction, an economic model predictive control approach, bounded residual reinforcement learning, and a hard safety projection. The structure can reduce the roughness of the longitudinal movement and interaction by keeping the learned component under the control of explicit headway, acceleration, jerk and actuator constraints. The division of responsibilities in the proposed design is as follows: Prediction provides an accurate forecast of traffic in the near future; Constrained optimisation determines a physically plausible nominal trajectory; Residual learning addresses local modelling errors; and a projection layer prevents adaptation from entering an unsafe region. Based on the simulation results, the above components work together to improve energy efficiency, tracking accuracy, passenger comfort and computational stability under all conditions: steady traffic, rolling terrain, congestion waves, cut-ins and degraded preview. Therefore, the resulting controller provides a practical link for data-driven adaptation and meets the constraint transparency requirements of safety-critical highway automation.

The study is limited by the use of controlled co-simulation and a single class of battery-electric vehicles. Although traffic behaviour, payload, grade, wind, sensing uncertainty and communication degradation have been randomised, simulations cannot replicate all the diversity of tire-road friction, component aging, localisation faults, battery temperature, actuator delay and human negotiation during merging. The deployment policy was also fixed after offline training; although this improved reproducibility, it could not adapt to long-term changes in the vehicle or regional traffic culture. In addition, the current safety projection only considers longitudinal movement and assumes that the lane-level reference remains valid. Therefore, unmodeled lane changes, abnormal road geometry, emergency maneuvers and simultaneous lateral-longitudinal conflicts may need to be covered by an extended reachable-set analysis. The above quantitative improvements should therefore be viewed as support for staged engineering validation and not as a guarantee of uniform road performance across the entire fleet.

Future work will expand the energy model to include battery-temperature dynamics, accessory-load prediction, tire-condition estimation and degradation-aware power limits. Cooperative perception and infrastructure preview can be used when communication is available, and uncertainty-calibrated local prediction should maintain safe operation in case of packet loss or inconsistent messages. The controller will then be moved through a series of processor-in-the-loop, hardware-in-the-loop, proving-ground and supervised public-road experiments, and the energy, comfort, safety and timing indicators will remain unchanged at all stages. Joint longitudinal and lateral planning is another way; merging, overtaking and lane-change negotiation in mixed-traffic conditions are representative examples. Formal reachable-set monitors and runtime assurance should accompany that extension. Finally, privacy-preserving fleet learning can update the prediction model in multiple vehicles through the collection of frequent traffic conditions and vehicle-specific efficiency characteristics without modifying the certified safety layer, and traceable vehicle-level guarantees can be maintained.

Author Contributions

Kacper Dmochowski contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Eryk Gierak contributes to validation, analysis, investigation, data collection, draft preparation. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Ma, F., Yang, Y., Wang, J., Li, X., Wu, G., Zhao, Y., Wu, L., Aksun-Guvenc, B., & et al. (2021). Eco-driving-based cooperative adaptive cruise control of connected vehicles platoon at signalized intersections. Transportation Research Part D: Transport and Environment, 92, 102746. https://doi.org/10.1016/j.trd.2021.102746

[2] Zhang, C., Wei, X., Wang, Z., Zhang, H., & Guo, X. (2022). T-S fuzzy-model-based adaptive cruise control for longitudinal car-following considering vehicle lateral stability. Intelligence & Robotics, 2(4), 371-390. https://doi.org/10.20517/ir.2022.26

[3] Lin, Y., McPhee, J., & Azad, N. L. (2021). Comparison of Deep Reinforcement Learning and Model Predictive Control for Adaptive Cruise Control. IEEE Transactions on Intelligent Vehicles, 6(2), 221-231. https://doi.org/10.1109/tiv.2020.3012947

[4] Nie, Z., & Farzaneh, H. (2022). Real-time dynamic predictive cruise control for enhancing eco-driving of electric vehicles, considering traffic constraints and signal phase and timing (SPaT) information, using artificial-neural-network-based energy consumption model. Energy, 241, 122888. https://doi.org/10.1016/j.energy.2021.122888

[5] Yeom, K. (2022). Model predictive control and deep reinforcement learning based energy efficient eco-driving for battery electric vehicles. Energy Reports, 8, 34-42. https://doi.org/10.1016/j.egyr.2022.10.040

[6] Zhang, J., Jiang, X., Cui, S., Yang, C., & Ran, B. (2022). Navigating Electric Vehicles Along a Signalized Corridor via Reinforcement Learning: Toward Adaptive Eco-Driving Control. Transportation Research Record: Journal of the Transportation Research Board, 2676(8), 657-669. https://doi.org/10.1177/03611981221084683

[7] Li, J., Wu, X., Xu, M., & Liu, Y. (2022). Deep reinforcement learning and reward shaping based eco-driving control for automated HEVs among signalized intersections. Energy, 251, 123924. https://doi.org/10.1016/j.energy.2022.123924

[8] Fernando, T., Denman, S., Sridharan, S., & Fookes, C. (2021). Deep Inverse Reinforcement Learning for Behavior Prediction in Autonomous Driving: Accurate Forecasts of Vehicle Motion. IEEE Signal Processing Magazine, 38(1), 87-96. https://doi.org/10.1109/msp.2020.2988287

[9] Alotibi, F., & Abdelhakim, M. (2021). Anomaly Detection for Cooperative Adaptive Cruise Control in Autonomous Vehicles Using Statistical Learning and Kinematic Model. IEEE Transactions on Intelligent Transportation Systems, 22(6), 3468-3478. https://doi.org/10.1109/tits.2020.2983392

[10] Boddupalli, S., Rao, A. S., & Ray, S. (2022). Resilient Cooperative Adaptive Cruise Control for Autonomous Vehicles Using Machine Learning. IEEE Transactions on Intelligent Transportation Systems, 23(9), 15655-15672. https://doi.org/10.1109/tits.2022.3144599

[11] Ding, H., Li, W., Xu, N., & Zhang, J. (2022). An enhanced eco-driving strategy based on reinforcement learning for connected electric vehicles: cooperative velocity and lane-changing control. Journal of Intelligent and Connected Vehicles, 5(3), 316-332. https://doi.org/10.1108/jicv-07-2022-0030

[12] Liu, Z., Yuan, Q., Nie, G., & Tian, Y. (2021). A Multi-Objective Model Predictive Control for Vehicle Adaptive Cruise Control System Based on a New Safe Distance Model. International Journal of Automotive Technology, 22(2), 475-487. https://doi.org/10.1007/s12239-021-0044-0

[13] Raja, G., Kottursamy, K., Dev, K., Narayanan, R., Raja, A., & Karthik, K. B. V. (2022). Blockchain-Integrated Multiagent Deep Reinforcement Learning for Securing Cooperative Adaptive Cruise Control. IEEE Transactions on Intelligent Transportation Systems, 23(7), 9630-9639. https://doi.org/10.1109/tits.2022.3168486

[14] Jurj, S. L., Grundt, D., Werner, T., Borchers, P., Rothemann, K., & Möhlmann, E. (2021). Increasing the Safety of Adaptive Cruise Control Using Physics-Guided Reinforcement Learning. Energies, 14(22), 7572. https://doi.org/10.3390/en14227572

[15] Hailemichael, H., Ayalew, B., Kerbel, L., Ivanco, A., & Loiselle, K. (2022). Safety Filtering for Reinforcement Learning-based Adaptive Cruise Control. IFAC-PapersOnLine, 55(24), 149-154. https://doi.org/10.1016/j.ifacol.2022.10.276

[16] Wang, H., Wang, Z., & Cui, X. (2021). Multi-objective Optimization Based Deep Reinforcement Learning for Autonomous Driving Policy. Journal of Physics: Conference Series, 1861(1), 012097. https://doi.org/10.1088/1742-6596/1861/1/012097

[17] Aparow, V. R., & Hoong, H. S. S. (2021). Evaluation of MRAC based Adaptive Cruise Control for Semi-Autonomous Vehicle using Virtual Simulation Platform. Journal of Physics: Conference Series, 1888(1), 012023. https://doi.org/10.1088/1742-6596/1888/1/012023

[18] Ogbebor, J., Meng, X., & Zhang, X. (2022). A Deep Reinforcement Learning Approach to Eco-driving of Autonomous Vehicles Crossing a Signalized Intersection. Journal of Engineering Research and Sciences, 1(5), 25-33. https://doi.org/10.55708/js0105003

[19] KURISHIGE, M., TSUKAHARA, Y., & NISHIHARA, O. (2022). Alleviation of traffic congestion using distributed autonomous control in mixed traffic flow of vehicles with adaptive cruise control and human-driven vehicles. Transactions of the JSME (in Japanese), 88(912), 22-00036-22-00036. https://doi.org/10.1299/transjsme.22-00036

[20] MARUYAMA, N., & MOURI, H. (2021). A proposal for Adaptive Cruise Control to multiple objectives by Reinforcement Learning. The Proceedings of JSME annual Conference on Robotics and Mechatronics (Robomec), 2021(0), 2P1-B01. https://doi.org/10.1299/jsmermd.2021.2p1-b01

[21] Sawant, J., Chaskar, U., & Ginoya, D. (2021). Robust Control of Cooperative Adaptive Cruise Control in the Absence of Information About Preceding Vehicle Acceleration. IEEE Transactions on Intelligent Transportation Systems, 22(9), 5589-5598. https://doi.org/10.1109/tits.2020.2987984

[22] Ahn, T., Lee, Y., & Park, K. (2021). Design of Integrated Autonomous Driving Control System That Incorporates Chassis Controllers for Improving Path Tracking Performance and Vehicle Stability. Electronics, 10(2), 144. https://doi.org/10.3390/electronics10020144

[23] Donà, R., Mattas, K., He, Y., Albano, G., & Ciuffo, B. (2022). Multianticipation for string stable Adaptive Cruise Control and increased motorway capacity without vehicle-to-vehicle communication. Transportation Research Part C: Emerging Technologies, 140, 103687. https://doi.org/10.1016/j.trc.2022.103687

[24] Yu, S., Shen, C., & Ersal, T. (2022). Autonomous Driving using Linear Model Predictive Control with a Koopman Operator based Bilinear Vehicle Model. IFAC-PapersOnLine, 55(24), 254-259. https://doi.org/10.1016/j.ifacol.2022.10.293

[25] Varotto, S. F., Mons, C., Hogema, J. H., Christoph, M., van Nes, N., & Martens, M. H. (2022). Do adaptive cruise control and lane keeping systems make the longitudinal vehicle control safer? Insights into speeding and time gaps shorter than one second from a naturalistic driving study with SAE Level 2 automation. Transportation Research Part C: Emerging Technologies, 141, 103756. https://doi.org/10.1016/j.trc.2022.103756

[26] Hidayatullah, M. R., & Juang, J. C. (2021). Adaptive Cruise Control With Gain Scheduling Technique Under Varying Vehicle Mass. IEEE Access, 9, 144241-144256. https://doi.org/10.1109/access.2021.3121494

[27] Yang, F., Li, X., Liu, Q., Li, Z., & Gao, X. (2022). Generalized Single-Vehicle-Based Graph Reinforcement Learning for Decision-Making in Autonomous Driving. Sensors, 22(13), 4935. https://doi.org/10.3390/s22134935

[28] Baheri, A. (2022). Safe reinforcement learning with mixture density network, with application to autonomous driving. Results in Control and Optimization, 6, 100095. https://doi.org/10.1016/j.rico.2022.100095

[29] Gupta, S., & Canova, M. (2021). Eco-Driving of Connected and Autonomous Vehicles with Sequence-to-Sequence Prediction of Target Vehicle Velocity. IFAC-PapersOnLine, 54(10), 430-436. https://doi.org/10.1016/j.ifacol.2021.10.200

[30] Tran, G. Q. B., Pham, T. P., Sename, O., Costa, E., & Gaspar, P. (2021). Integrated Comfort-Adaptive Cruise and Semi-Active Suspension Control for an Autonomous Vehicle: An LPV Approach. Electronics, 10(7), 813. https://doi.org/10.3390/electronics10070813

[31] Khajeh Hosseini, M., Talebpour, A., Devunuri, S., & Hamdar, S. H. (2022). An unsupervised learning framework for detecting adaptive cruise control operated vehicles in a vehicle trajectory data. Expert Systems with Applications, 208, 118060. https://doi.org/10.1016/j.eswa.2022.118060

[32] Mirwald, J., Ultsch, J., de Castro, R., & Brembeck, J. (2021). Learning-Based Cooperative Adaptive Cruise Control. Actuators, 10(11), 286. https://doi.org/10.3390/act10110286

[33] Li, X., Xie, N., & Wang, J. (2022). A variable weight adaptive cruise control strategy based on lane change recognition of leading vehicle. Automatika, 63(3), 555-571. https://doi.org/10.1080/00051144.2022.2055913

[34] Wang, F. X., Peng, Q., Zang, X. L., & Xue, Q. F. (2021). Adaptive Cruise Control for Intelligent City Bus Based on Vehicle Mass and Road Slope Estimation. Applied Sciences, 11(24), 12137. https://doi.org/10.3390/app112412137

[35] van den Beukel, A., van Driel, C., Boelhouwer, A., Veders, N., & Heffelaar, T. (2021). Assessment of Driving Proficiency When Drivers Utilize Assistance Systems—The Case of Adaptive Cruise Control. Safety, 7(2), 33. https://doi.org/10.3390/safety7020033

[36] Gao, Z., Yan, X., Gao, F., & He, L. (2022). Driver-like decision-making method for vehicle longitudinal autonomous driving based on deep reinforcement learning. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, 236(13), 3060-3070. https://doi.org/10.1177/09544070211063081