Journal of Applied Automation Technologies

Research Paper

A New Technical Framework for Safe and Efficient Left Turns in Autonomous Driving

Patrycja Ostrowska[1, *], Sabina Walczak[1] and Lidia Wójcikowa[1]

[1] Faculty of Electrical Engineering, Automatics, Computer Science and Biomedical Engineering, AGH University of Science and Technology, Krakow, 30-059, Poland
*Corresponding author: patrycja.os@agh.edu.pl

Abstract. Unprotected left turns are still a difficult problem for autonomous driving because perception uncertainty, occlusion, and unpredictable road-user behaviour need to be handled in a short decision time. This paper proposes a risk-adaptive technical framework for integrating conflict-zone graph encoding, multimodal arrival-time prediction, dynamic risk allocation, chance-constrained gap selection and jerk-limited trajectory optimisation. The framework estimates the probability distribution of conflict-zone entry and clearance times, adjusts the admissible risk according to visibility, localisation uncertainty, road conditions and prediction calibration. 18,000 closed-loop simulation tests, 2,400 log-replay cases and 320 closed-course experiments were conducted for performance assessment. The proposed method had a 98.1% turn-completion rate and a 0.18% collision rate, and its collision rate was 0.27-0.71% for the four representative baselines. The median stop-line-to-clearance time was 7.6s, and the average intersection throughput reached 428 vehicles per hour per approach. The expected calibration error was 0.012, the median computational latency was 41ms, and 99.2% of the closed-course control cycles maintained a feasible backup trajectory. With combined sensing and behaviour disturbance, the collision-free survival rate after three seconds of continuous occlusion was 99.1%. As shown in the above results, adding probabilistic prediction and risk-constrained planning can improve the safety of left turns without adding much latency or other issues.

Keywords: *Autonomous Driving, Left-Turn Planning, Trajectory Prediction, Risk Allocation, Chance-Constrained Optimization, Motion Planning, Intersection Safety*

Received on 28 June 2022, Accepted on 19 February 2023, Published on 08 March 2023

Introduction

Urban intersections have complex motion, ambiguous priority, and a short decision horizon; thus, they are a primary source of operational risk for autonomous vehicles [1]. An unprotected left turn is particularly difficult because the ego vehicle crosses the opposing lane and has to pay attention to pedestrians, cyclists and vehicles from all sides [2]. Lane keeping does not have such a discrete commitment point; after this point, even conservative braking cannot return to a safe state [3]. Humans negotiate this transition by means of eye contact, speed variation and other non-rule-based priority signals [4]. Occluding vehicles and roadside structures further obscure the relevant actors for a short time before the collision [5]. Prediction errors are thus linked to planning errors and are not isolated perception defects [6]. A planner that anticipates the future may choose an apparently open gap even if a low-probability but feasible path enters the same conflict zone [7]. Excessive caution is not a suitable solution because it prolongs queueing, creates rear-end exposure, and makes it difficult for other road users to recognise automated behaviour [8].

Recently, some research has begun to use graphs, polylines and attention-based scene tokens to represent agents and maps in motion forecasting [9]. Multimodal predictors can describe several plausible futures, but their probability scores are not always calibrated for direct application in safety constraints [10]. Reinforcement-learning policies have shown interactive negotiation, but reward design and rare-event coverage are problems for certification [11]. Model-predictive planners offer physical consistency and receding-horizon feedback, but their computational cost is relatively high when many interacting futures need to be considered [12]. Rule-based gap acceptance is still transparent but cannot handle coupled uncertainty in speed, intent and occlusion gracefully [13]. Reachability and control-barrier methods have strong safety structures, but overly conservative sets can reduce the utility of progress at dense intersections [14]. Therefore, the existing systems generally optimise prediction, decision or control independently and fail to address how to convert uncertain occupancy into interpretable manoeuvre-level risk budgets [15].

This paper builds a unified technical framework for safe and efficient left turns that explicitly represents uncertainty, propagates it, and utilizes it. First, predict the distribution of arrivals and departures in a conflict area; convert the overlap into an adjusted collision probability, allocate an adaptive risk budget, and then optimise a dynamically feasible path. The framework still has the modular interfaces required by the production stack, and prediction and planning can exchange only the relevant parameters for the manoeuvre. It is a conflict-zone graph that preserves interaction structure at a low computational cost, a chance-constrained gap selector with quality-aware risk allocation, and a jerk-limited trajectory optimizer supported by an emergency-safe fallback. All-encompassing experimental procedures are used to evaluate the average success rate and tail risk, calibration, comfort, latency and throughput under controlled perturbations.

Related Work

Interaction-Aware Prediction at Intersections

Trajectory prediction is no longer done by a single recurrent model; rather, scenes that consider map context and agent interaction are now employed [16]. To enhance the geometric accuracy of lanes, boundaries and crosswalks, a structured polyline-based map encoder can be used instead of a pixelated image [17]. Graph neural networks can also extend messages according to the actual adjacency and conflict relationships, and are thus suitable for scenarios where actors in distant areas will soon be in the same intersection region [18]. Transformer-based models can learn long-range dependencies and support multimodal decoding, but their memory and latency are relatively high as the number of scene tokens increases [19]. Probabilistic decoders have multiple intentions, and although they do not guarantee that the predicted probabilities are well-calibrated in the low-risk region required by a safety planner, this is still an issue [20].

Intersection forecasting also needs the ego plan. An opposing driver may yield, accelerate or maintain speed in response to visible ego movement, so an open-loop prediction will be out of date as soon as the autonomous vehicle starts to creep. Conditional forecasting considers how the actors' future behaviour will affect the selected ego trajectory. It has the problem of combinatorial explosion: every new candidate requires an additional prediction round. Therefore, the efficient system should have a small shared representation and a way to update only the interaction variables that have been changed by each candidate. Arrival-time distributions at conflict zones offer such a way to keep the timing relation that drives left-turn risk without requiring repeated decoding of full high-resolution trajectories.

Decision-Making and Motion Planning for Left Turns

Classical gap-acceptance models compare the estimated time-to-arrival with a fixed critical gap and are still popular because their decisions are easy to inspect [21]. However, a fixed threshold does not consider perception confidence, road friction, vehicle response or the asymmetric consequences of entering versus waiting. Sampling-based and lattice planners extend the manoeuvre set and can impose kinematic limits, but safety evaluation for many multimodal predictions is computationally expensive [22]. Model Predictive Control integrates feasibility and feedback, but nonconvex collision constraints and discrete yielding decisions often require approximation or hierarchical decomposition [23]. Learning-based policies can find socially efficient behaviour in the simulation, but they may exploit simulator-specific cues and provide little evidence of rare hazardous interactions [24]. Formal safety filters, such as reachable-set and barrier-function approaches, can supervise a nominal planner; however, conservative invariant sets may cause a deadlock if the uncertainty persists for a long time [25].

https://doi.org/10.64972/jaat.2023v1.284p6e:76-88

The outstanding engineering problem is not that there are no individual tools, but rather that they do not connect stably. A forecasting module generally outputs sampled trajectories, a decision layer uses hand-tuned thresholds, and a controller receives a reference path after uncertainty has been collapsed. The framework put forward will use conflict-zone occupancy time as its common unit. Prediction provides a distribution of arrival and clearance; decision assigns a maneuver risk budget; optimisation spends this budget to track delay and comfort; and the safety supervisor keeps a provably executable stop option. This division maintains interpretability and does not require the whole stack to be a single network.

Proposed Risk-Adaptive Left-Turn Framework

Conflict-Zone Scene Representation and Probabilistic Prediction

The framework runs at 10 Hz and receives tracked actors, lane topology, traffic-control state, localization covariance, and a candidate route. Its first operation constructs a conflict-zone graph rather than a generic fully connected scene graph. Nodes represent the ego vehicle, relevant dynamic agents, lane segments, crosswalks, occlusion boundaries, and geometric conflict zones. Edges encode lane succession, right-of-way, visibility, potential path intersection, and temporal precedence. Actors with no reachable overlap during the planning horizon are pruned, which keeps inference time nearly constant in ordinary urban traffic. Figure 1 presents this information flow from synchronized sensing through graph construction, probabilistic forecasting, risk allocation, gap selection, trajectory optimization, and safety supervision.

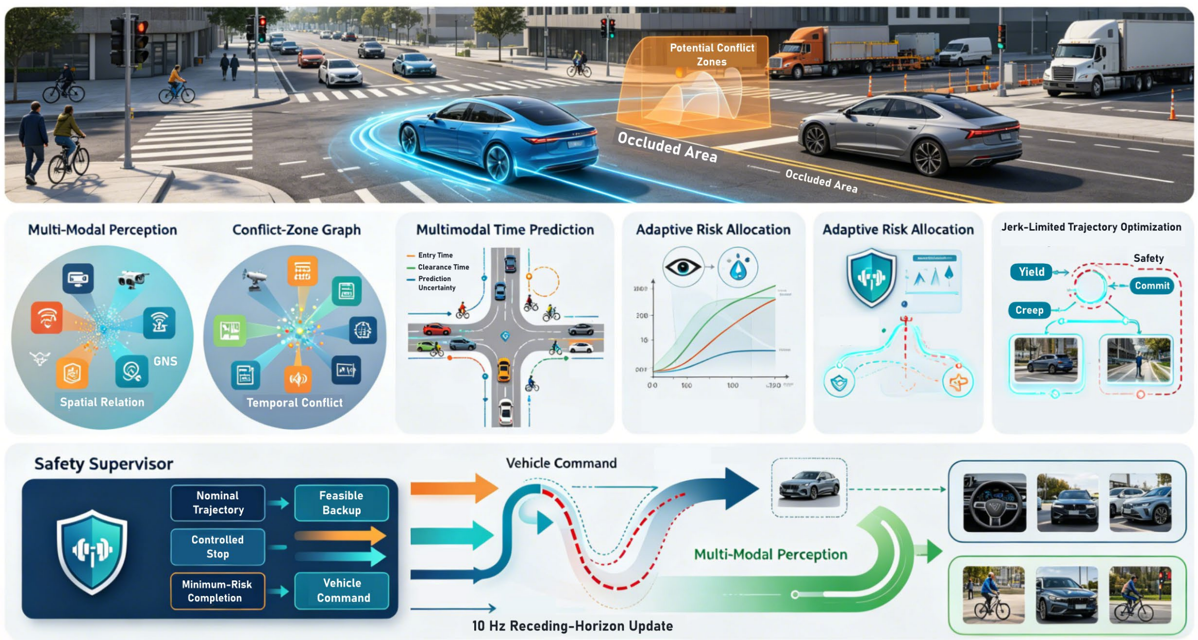

Figure 1. Overall workflow of the proposed risk-adaptive left-turn framework.

For actor i, the input feature combines relative pose, velocity, acceleration, class, observed history, lane association, signal context, visibility, and tracking covariance. A graph encoder produces a state embedding by aggregating typed neighbor messages:

$$h_i^{(l+1)} = \phi_l\left(h_i^{(l)}, \sum_{j\in\mathcal{N}(i)} \alpha_{ij}^{(l)} \psi_l\big(h_i^{(l)}, h_j^{(l)}, e_{ij}\big)\right) \qquad \text{Eq.(1)}$$

The attention weight favors neighbors that share a reachable conflict zone and discounts actors whose projected occupancy is temporally separated:

$$\alpha_{ij} = \frac{exp\left(q_i^\top k_j/\sqrt{d} - \lambda_\tau\left|\hat{t}_i - \hat{t}_j\right|\right)}{\sum_{r\in\mathcal{N}(i)} exp\left(q_i^\top k_r/\sqrt{d} - \lambda_\tau\left|\hat{t}_i - \hat{t}_r\right|\right)} \qquad \text{Eq.(2)}$$

Instead of decoding only Cartesian trajectories, the prediction head estimates a mixture distribution for the time at which each actor enters and leaves each conflict zone:

$$p_i(t^{in}, t^{out}) = \sum_{m=1}^{M} \pi_{im}\mathcal{N}\left(\begin{bmatrix} t^{in} \\ t^{out} \end{bmatrix}; \mu_{im}, \Sigma_{im}\right), \sum_{m=1}^{M} \pi_{im} = 1 \qquad \text{Eq.(3)}$$

The network is trained with a mixture negative log-likelihood augmented by a calibration term evaluated over conflictoccupancy events:

$$\mathcal{L}_{pred} = -\sum_{i} log\, p_i\left(t_i^{in*}, t_i^{out*}\right) + \lambda_c \sum_{b=1}^{B} |acc(b) - conf(b)| \qquad \text{Eq.(4)}$$

To account for an unobserved road user emerging from an occluded area, a virtual actor is added to the graph with a conservative occupancy distribution derived from the visible boundary, speed limit and road-user class prior. The virtual actor is deleted as soon as the free space is visible. This mechanism does not claim to predict an invisible agent; it reserves probability mass for plausible emergence and allows the risk allocator to trade a slow visibility-seeking creep for indefinite waiting.

Adaptive Risk Allocation and Gap Selection

The decision layer enumerates yield, creep, and commit maneuver families. Each candidate defines an ego arrival and clearance distribution at every conflict zone, including uncertainty from state estimation and tracking error. The overlap probability for actor i and candidate k is

$$P_{ik}^{col} = \iint \mathbf{1}\left(\left[t_{e,k}^{in}, t_{e,k}^{out}\right] \cap \left[t_i^{in}, t_i^{out}\right] \neq \emptyset\right) p_{e,k} p_i d\boldsymbol{t} \qquad \text{Eq.(5)}$$

Because collision events are dependent through shared ego timing, the implementation uses a conservative product approximation with a bounded correlation correction:

$$R_k = 1 - \prod_{i=1}^{N} \left(1 - P_{ik}^{col}\right) + \kappa \max_{i\neq j} Cov\left(I_{ik}, I_{jk}\right) \qquad \text{Eq.(6)}$$

The admissible risk decreases with occlusion, perception covariance, prediction miscalibration, wet-road indication, and control-tracking error:

$$\varepsilon_t = clip\left(\varepsilon_0 e^{-\beta_0 o_t - \beta_u u_t - \beta_c c_t - \beta_f f_t} + \delta n_t, \varepsilon_{min}, \varepsilon_{max}\right) \qquad \text{Eq.(7)}$$

A candidate is feasible only when its total risk is below ε_t, its predicted stopping margin remains positive before commitment, and it respects red-light, crosswalk, and lane-boundary constraints. Among feasible candidates, the selector minimizes

$$J_k = w_r R_k + w_d D_k + w_s S_k + w_j J_k^{jerk} - w_p P_k + w_h H_k, R_k \leq \varepsilon_t \qquad \text{Eq.(8)}$$

Here D_k is delay, S_k penalizes low clearance margin, J_k^{jerk} represents comfort, P_k rewards progress, and H_k penalizes rapid maneuver switching. A commit state is entered only when the selected candidate remains feasible for two consecutive cycles and the minimum probabilistic clearance exceeds a speed-dependent threshold.

Figure 2 shows the maneuver-state logic and nested safety gates. The route begins in approach, transitions to yield at the stop line, may creep to improve visibility, and enters commit only after the probability, stopping-margin, and control-feasibility gates are satisfied. Any violation before the no-return boundary sends the vehicle to controlled stop; after that boundary, the supervisor selects the minimum-risk completion trajectory.

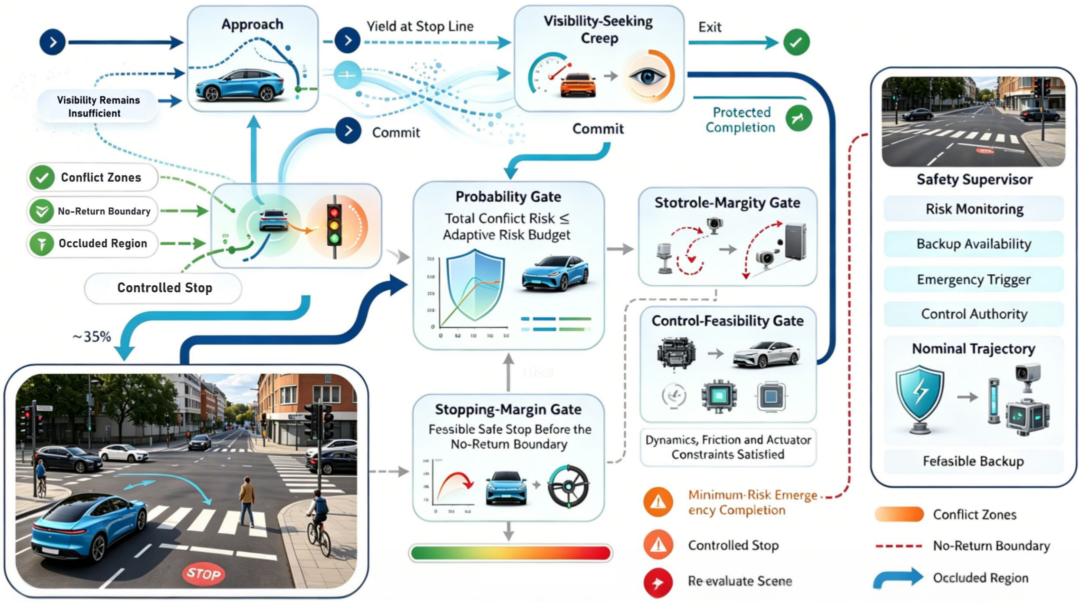

Figure 2. Hierarchical maneuver-state and safety-gate diagram for approach, yield, visibility-seeking creep, commit, protected completion, controlled stop, and minimum-risk emergency completion.

Jerk-Limited Trajectory Optimization and Safety Supervision

Once a maneuver and an admissible gap have been selected, the framework constructs a spatial reference path from the lane centerlines and a quintic lateral connector. Longitudinal motion is optimized in Frenet coordinates over a 5 s receding horizon. The discrete vehicle state contains path progress, velocity, acceleration, and steering angle, while the control vector contains longitudinal jerk and steering rate. The state transition is expressed as

$$x_{n+1} = f_d(x_n, u_n), u_n = [j_n, \dot{\delta}_n]^T \qquad \text{Eq.(9)}$$

The trajectory optimizer tracks the selected conflict-zone occupancy interval while accounting for motion accuracy, passenger comfort, and terminal-state consistency. Its objective function is written in compact form as

$$J_{traj} = \sum_{n=0}^{H} \ell(x_n, u_n) + w_T \left\| x_H - x_H^{ref} \right\|^2 \qquad \text{Eq.(10)}$$

Here, $\ell(x_n, u_n)$ denotes the stage cost associated with velocity tracking, lateral deviation, longitudinal jerk, and steering-rate variation. The terminal term penalizes deviation from the desired state at the end of the prediction horizon. This compact representation allows the relative importance of safety, tracking accuracy, and ride comfort to be adjusted without expanding the main optimization expression.

Dynamic feasibility is maintained by imposing explicit bounds on velocity, acceleration, jerk, steering response, and lateral acceleration:

$$0 \le v_n \le v_{max}, a_{min} \le a_n \le a_{max}, |j_n| \le j_{max}, |v_n^2 \kappa_n| \le a_{y,max} \qquad \text{Eq.(11)}$$

The curvature term κ_n connects the planned path geometry with the admissible lateral acceleration. Under wet pavement or reduced tire-road friction, the allowable acceleration and jerk limits are lowered before the optimization problem is solved. This treatment prevents a geometrically valid trajectory from being accepted when it cannot be tracked reliably by the physical vehicle.

Probabilistic occupancy constraints are converted into temporal separation conditions. For each active conflict zone, the ego vehicle must either clear the zone before the lower arrival-time quantile of the approaching actor or enter after the actor's upper clearance-time quantile:

$$t_{e,k}^{out} + \Delta_t \le Q_i^{in}(\varepsilon_i) \lor t_{e,k}^{in} - \Delta_t \ge Q_i^{out}(1 - \varepsilon_i) \qquad \text{Eq.(12)}$$

In this expression, Δ_t is an additional temporal safety margin, while Q_i^{in} and Q_i^{out} are the quantile functions of the predicted entry and clearance times. The disjunctive constraint distinguishes passing before an actor from yielding until the actor has cleared the shared region. It therefore preserves the operational meaning of gap acceptance while incorporating prediction uncertainty.

The maneuver-level risk budget is distributed among active actors according to their interaction relevance. A minimum allocation is retained for pedestrians and cyclists so that their risk contribution cannot disappear when several motor vehicles are present:

$$\varepsilon_i = \varepsilon_t \frac{\rho_i + \rho_{min}\mathbf{1}_{VRU}(i)}{\sum_j \left(\rho_j + \rho_{min}\mathbf{1}_{VRU}(j)\right)} \quad \text{Eq.(13)}$$

The relevance score ρ_i considers conflict-zone overlap, predicted arrival time, relative speed, visibility, and right-ofway. The indicator $\mathbf{1}_{\mathrm{VRU}}(i)$ equals one for a vulnerable road user and zero otherwise. This allocation concentrates computation on the most immediate conflicts without weakening protection for pedestrians or cyclists.

At every planning cycle, the safety supervisor independently computes a backup trajectory leading to the latest collision-free stopping pose. The braking profile accounts for actuator delay, road friction, current acceleration, and the additional distance required to reduce acceleration smoothly. A safe stop remains feasible only if

$$d_{stop} = v\tau_a + \frac{v^2}{2\mu g} + \frac{|a|^2}{2j_{max}} \leq d_{free} - d_{margin} \quad \text{Eq.(14)}$$

Here, τ_a is the estimated actuator delay, μ is the tire-road friction coefficient, d_{free} is the currently available collisionfree distance, and d_{margin} is a reserved stopping margin. Before the no-return boundary, violation of this condition causes the system to reject the nominal maneuver and execute a controlled stop. After the boundary has been crossed, the supervisor selects a minimum-risk completion trajectory because an abrupt stop inside the opposing lane could produce a more hazardous state.

Real-time performance is maintained by warm-starting the nonlinear program with the preceding solution, limiting each update to three sequential-convexification iterations, and reusing graph embeddings for actors whose interaction neighborhoods remain unchanged. The newly optimized command is blended with the preceding command and projected onto the safe control set:

$$u_t = \Pi_{\mathcal{U}_{safe}}(\eta u_t^* + (1-\eta)u_{t-1}) \quad \text{Eq.(15)}$$

The blending factor η reduces abrupt command changes, whereas the projection operator guarantees that smoothing cannot move the command outside the feasible control region. Prediction therefore supplies calibrated occupancy distributions, trajectory optimization produces a comfortable and dynamically executable maneuver, and the independent supervisor preserves a valid fallback throughout the left-turn process.

Experimental Evaluation and Analysis

Experimental Protocol and Metrics

The three sources of evaluation are 18,000 closed-loop simulated left turns, 2,400 log-replay cases from urban intersection recordings, and 320 closed-course runs with a drive-by-wire research vehicle. The three intersection geometries in the simulated set have opposing traffic flows of 300 to 1,200 vehicles per hour, pedestrian arrival rates of 0 to 240 per hour, dry and wet friction, and occlusion ratios from 0 to 0.65. Ten times, randomly generate a different seed for each scenario. Record cases to store measured actor motion and use a validated interaction model for counterfactual ego actions. Closed-course tests use a robotic target vehicle and soft pedestrian dummies, with speeds of 20, 30, 40 and $50\ \mathrm{km/h}$ and injected localization noise up to 0.45 m.

The proposed framework, RALT, is compared with fixed-threshold gap acceptance, a sampling-based lattice planner, interaction-aware model predictive control, and a reinforcement-learning policy with a safety shield. All the above methods have the same detection results and maps. Safety indices are as follows: collision rate, minimum post-encroachment time, emergency braking frequency, and 95th-percentile peak deceleration. The indices of efficiency include completion rate, stop-line-to-clearance time, gap utilisation and throughput. Comfort uses jerk and lateral acceleration; calibration uses expected calibration error and Brier score.

Parameter tuning is not evaluation. A development split of 2,100 simulated turns is employed to select the loss weight, baseline risk and maximum creep speed. The rest of the cases are unavailable for comparison. The nominal risk budget is 0.006, bounded by 0.002 and 0.012; the maximum creep speed is 1.8"m/s; the stopping reserve is 1.2 m; and the predictor has six modes. Baseline values are adjusted to match the nominal completion rates of the proposed method within a 2% margin of error before safety comparison.

The simulation matrix includes permissive signals, unsignalised T-junctions, four-leg intersections, legal and contraflow cyclists, pedestrians entering before or after commitment, and constant-speed, yielding, assertive, and late-acceleration vehicles. Occlusion is caused by parked vans, queued vehicles, vegetation and building corners. Friction is 0.40-0.95 and the actuator delay is 70-220ms.

Paired bootstrap intervals are generated by 10,000 resamples, and scenario-level differences are assessed rather than treating correlated time steps as independent observations [26]. Rare conflict scenarios are oversampled during generation but are reweighted according to their naturalistic exposure frequencies in the reporting of the main safety indicators [27]. Pre-planning analysis of predictor calibration will be performed independently; otherwise, the probability results from mean trajectory estimates may be inaccurate [28].

Safety, Efficiency, Comfort, and Computational Results

Across the naturalistic simulation mixture, RALT completed 98.1% of turns with a collision rate of 0.18%, compared with 0.71% for fixed-threshold gap acceptance, 0.46% for the lattice planner, 0.31% for interaction-aware MPC, and 0.27% for the shielded reinforcement-learning policy. Median clearance time was 7.6 s, only 0.3 s slower than the learning policy and 1.5 s faster than conservative MPC. Mean throughput reached 428 vehicles per hour per approach, whereas fixed-threshold control achieved 391 and MPC achieved 374 [29].

The paired collision-rate difference between RALT and fixed-threshold control is -0.53 percentage points, with a 95% bootstrap interval from -0.66 to -0.39. Relative to MPC, the difference is -0.13 points with an interval from -0.21 to -0.05. The learning policy is faster by 0.3 s but triggers emergency braking in 1.6% of turns, compared with 0.7% for RALT.

Figure 3(a) shows that RALT kept collision rate below 0.35% across opposing flows from 300 to 1,200 vehicles per hour, while fixed-threshold control rose from 0.22% to 1.34%; Figure 3(b) reports RALT completion times of 6.1, 6.8, 7.7, and 9.4 s ; Figure 3(c) displays a post-encroachment-time distribution centered at 2.41 s , versus 1.72 s for the learning policy; and Figure 3(d) shows that 94.6% of RALT runs stayed below $2.0\ \mathrm{m}/\mathrm{s}^2$ peak deceleration, compared with 81.3% for the lattice planner [30].

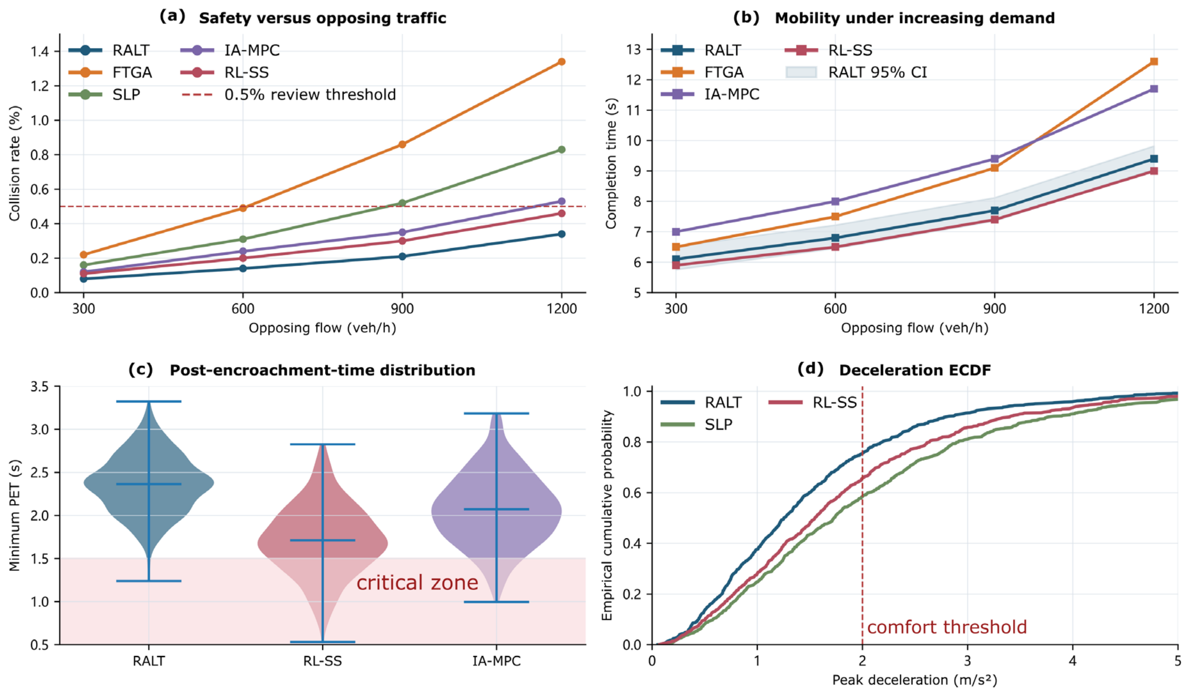

Figure 3. Density-stratified safety and mobility results: (a) collision rate versus opposing flow; (b) completion time versus opposing flow; (c) minimum post-encroachment-time distribution; (d) peak-deceleration empirical cumulative distribution.

https://doi.org/10.64972/jaat.2023v1.284p6e:76-88

Figure 4(a) relates predicted probability to observed frequency over 20 bins, yielding expected calibration errors of 0.012 for RALT, 0.039 for MPC, and 0.067 for the learning policy; Figure 4(b) attributes collision reduction to conflicttime prediction, adaptive allocation, backup supervision, and component interaction; Figure 4(c) shows accepted gaps of 4.8 s on dry pavement, 5.4 s with 35% occlusion, and 6.2 s on wet pavement [31].

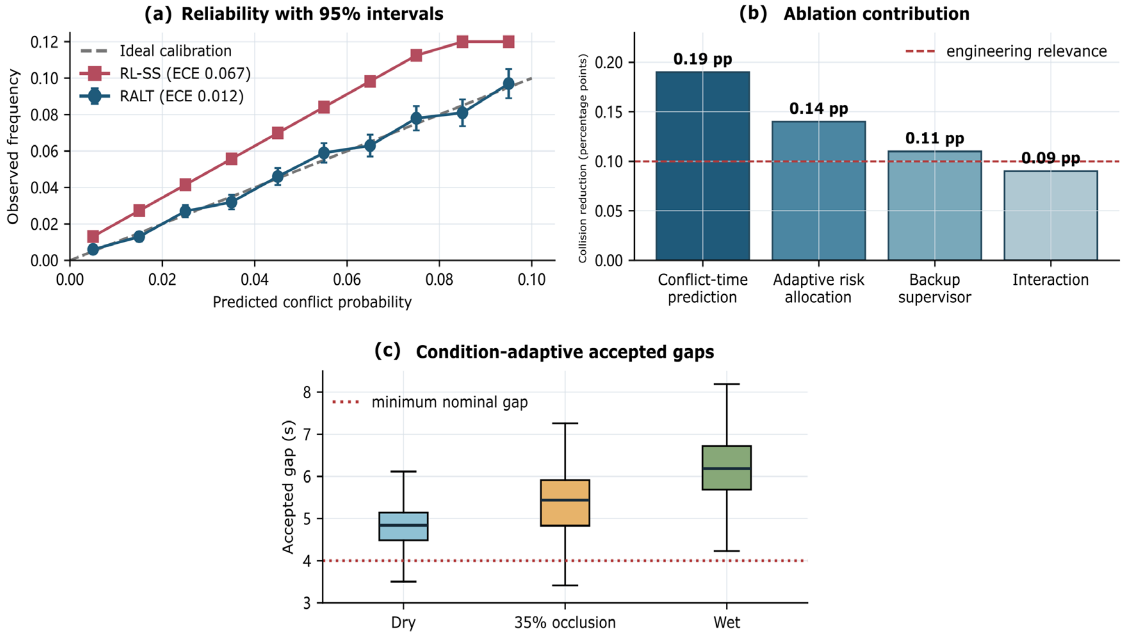

Figure 4. Probability calibration and component analysis: (a) reliability scatter with confidence bands; (b) ablation contributions; (c) accepted-gap distributions.

Figure 5(a) gives RALT normalized scores of 0.93 for safety, 0.88 for efficiency, 0.91 for comfort, 0.86 for calibration, and 0.84 for computation; Figure 5(b) maps 95th-percentile solver latency from 28 to 63 ms as actor count increases from 4 to 24 and mode count from 2 to 8; Figure 5(c) shows that risk consumption is dominated by opposing vehicles before commitment and shifts to crosswalk users around the apex. Median end-to-end latency is 41 ms, its 99th percentile is 83 ms, and deadline misses occur in 0.27% of cycles [32].

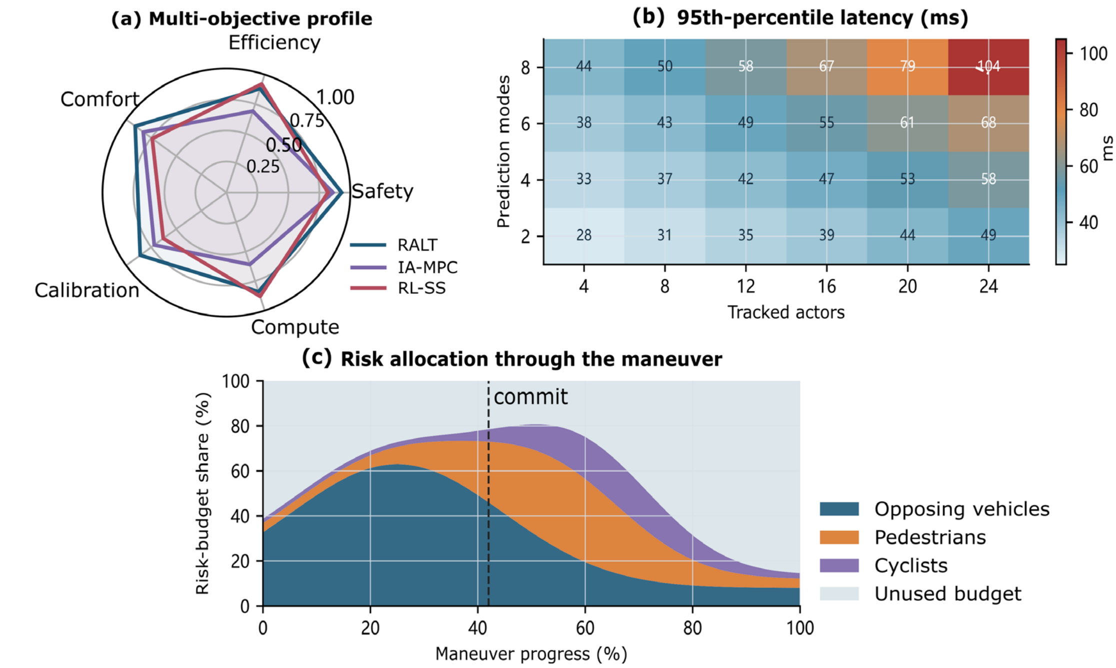

Figure 5. multi-objective and real-time characteristics: (a) normalized radar comparison; (b) solver-latency heatmap; (c) stacked risk-budget composition.

Profiling assigns 13 ms to graph construction and prediction, 4 ms to candidate generation, 18 ms to trajectory optimization, and 6 ms to supervision and message transport. Embedding reuse saves 7.4 ms in 72% of cycles. Warm starting reduces the median iteration count from 4.8 to 2.6 and the 99th-percentile latency by 21 ms.

No contact occurs in 320 closed-course runs. The minimum measured post-encroachment time is 1.38 s during a 50 $\mathrm{km/h}$ target approach with 0.45 m localization noise. Controlled braking is triggered in 11 runs. Peak deceleration remains below $3.4\ \mathrm{m/s^2}$, peak jerk below $3.9\ \mathrm{m/s^3}$, and average lateral tracking error is 0.11 m [33].

Robustness, Failure Analysis, and Generalization

Figure 6(a) shows that temperature scaling reduces expected calibration error from 0.048 to 0.012 in-domain and from 0.091 to 0.029 under joint shift; Figure 6(b) gives collision-free survival of 99.1% for RALT at 3.0 s occlusion, versus 96.8% for MPC and 94.9% for the learning policy; Figure 6(c) plots the safety-delay Pareto frontier, where an adaptive budget of 0.006 produces 0.18% collisions at 7.6 s completion [34].

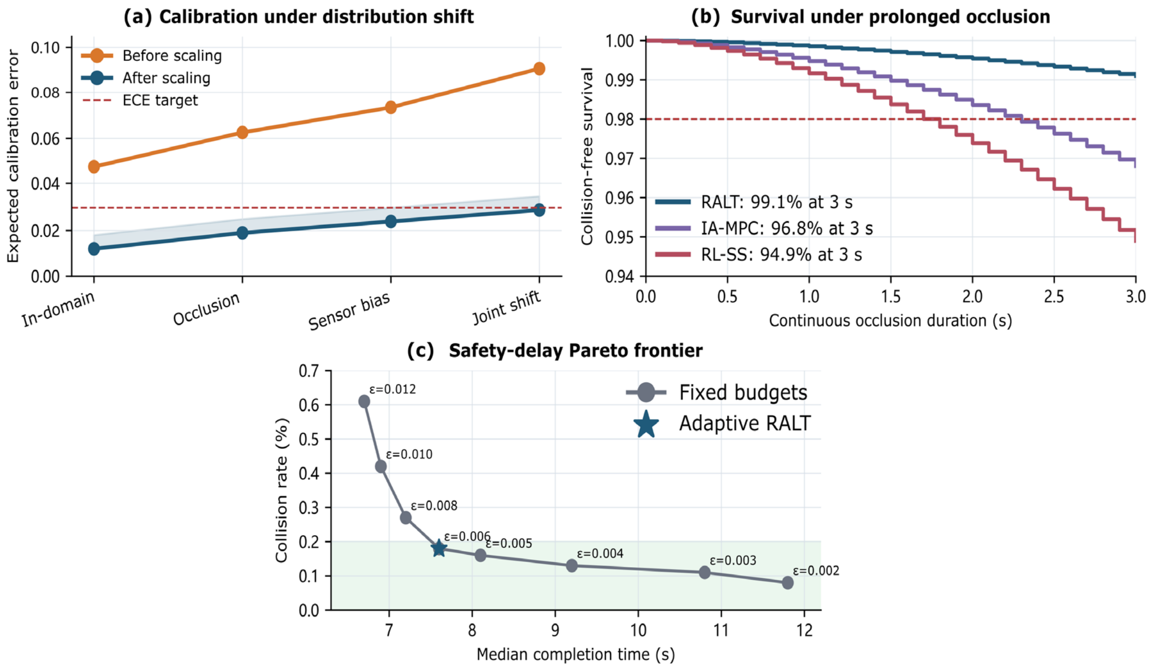

Figure 6. Robustness under distribution shift: (a) calibration error; (b) collision-free survival; (c) safety-delay Pareto frontier.

Failure inspection identifies contraflow cyclists, correlated range bias during rain, and opposing drivers accelerating after apparently yielding. The backup supervisor avoids collision in 37 of 41 injected cases, although four cross the emergency-braking threshold and two exceed $4.0\ \mathrm{m/s^3}$ jerk [35].

Of 186 near-miss cases with post-encroachment time below 1.5 s, 61% involve behavioral change after ego selection, 24% involve correlated sensing error, and 15% involve map or association error. This distribution suggests that increasing predictor capacity alone would address only part of the residual risk.

Figure 7 (a) shows that risk increases sharply when occlusion exceeds 0.50 and target acceleration exceeds $1.5\ \mathrm{m/s^2}$, yet RALT holds 87% of such cases below an 0.5% event rate by selecting creep rather than commit; Figure 7 (b) maps feasible-gap utilization from 0.82 at 0.05 m localization noise and friction 0.9 to 0.54 at 0.45 m noise and friction 0.45 [36].

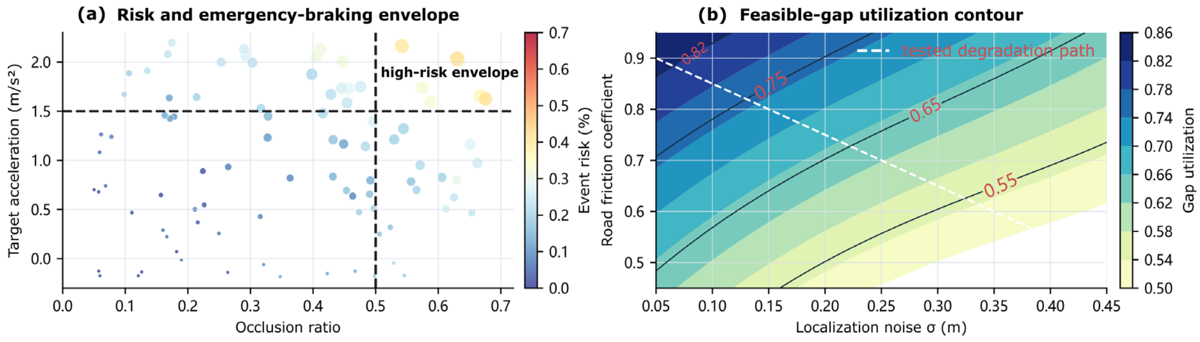

Figure 7. Failure-envelope analysis: (a) event-risk bubble plot; (b) feasible-gap-utilization contour.

Replacing conflict-time prediction with point trajectories increases collision rate from 0.18% to 0.43%. Fixing the risk budget at 0.006 raises the wet-and-occluded collision rate from 0.31% to 0.58%. Removing backup feasibility triples hard-braking events from 0.7% to 2.1%, while removing maneuver hysteresis increases state switching from 1.4 to 3.8 transitions per turn.

Simulation-to-track transfer produces an 6.7% increase in completion time and an $0.09\ \mathrm{m/s^2}$ increase in peak lateral acceleration. Updating only the actuator-delay distribution recovers approximately two thirds of this gap [37]. Sensitivity analysis shows that collision rate is most sensitive to minimum risk and stopping reserve, while completion time is most sensitive to creep speed and progress weight.

The full simulation campaign is rerun with different random seeds, traffic-behavior priors, and solver initialization. Collision rate changes from 0.18% to 0.20%, completion time from 7.6 to 7.7 s, and 99 th-percentile runtimes from 83 to 86 ms, without changing the planner ranking.

When map confidence falls below 0.7, correlated localization error exceeds 0.6 m, or estimated friction falls below 0.35, the tested envelope no longer supports an unprotected commit. The system then requests a protected phase, reroutes, or remains stopped according to traffic rules.

Conclusion

This paper proposed a risk-adaptive framework for autonomous unprotected left turns that connects interaction-aware perception with executable motion via a common conflict-time representation. Several prediction methods are employed to construct a conflict-zone graph; an adaptive chance constraint is used to convert uncertainty into an explicit manoeuvre budget; and finally, a jerk-limited optimizer generates readable trajectories that meet the feasibility requirements for backing up. The evaluation design considers safety, delay, comfort, calibration and computation in simulation, log replay and closed-course operation to determine whether the architecture is an integrated driving function.

The system is not perfect. Its virtual-actor model represents occlusion conservatively but does not infer complex hidden interactions, and its correlation correction is an approximation rather than a full joint occupancy model. The closed-course sample cannot reproduce the behaviour diversity of public roads, especially unconventional cyclist motion and informal negotiation. Performance is also affected by the map-level conflict topology; thus, construction zones or temporary lane guides can reduce the quality of the graph to an extent that is visible.

Future work will learn correlation-aware occupancy distributions, integrate vehicle-to-everything messages with explicit trust estimation, and extend the safety supervisor to mixed traffic containing multiple negotiating automated vehicles. Extend the fleet trial to link scenario-level risk estimation with exposure-weighted safety cases. A further way is to generate compact, human-readable maneuver explanations based on the same logged risk allocation and clearance variables to help engineers and safety operators understand why the vehicle waited, crept, committed, or used its backup.

Author Contributions

Patrycja Ostrowska contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Sabina Walczak and Lidia Wójcikowa contribute 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] Yang, L., Lu, C., Xiong, G., Xing, Y., & Gong, J. (2022). A hybrid motion planning framework for autonomous driving in mixed traffic flow. Green Energy and Intelligent Transportation, 1(3), 100022. https://doi.org/10.1016/j.geits.2022.100022

[2] Chen, T., Chen, Y., Li, H., Gao, T., Tu, H., & Li, S. (2022). Driver Intent-Based Intersection Autonomous Driving Collision Avoidance Reinforcement Learning Algorithm. Sensors, 22(24), 9943. https://doi.org/10.3390/s22249943

[3] Xia, C., Xing, M., & He, S. (2022). Interactive Planning for Autonomous Driving in Intersection Scenarios Without Traffic Signs. IEEE Transactions on Intelligent Transportation Systems, 23(12), 24818-24828. https://doi.org/10.1109/tits.2022.3205250

[4] Zhang, Z., Wang, X., Huang, D., Fang, X., Zhou, M., & Zhang, Y. (2022). MRPT: Millimeter-Wave Radar-Based Pedestrian Trajectory Tracking for Autonomous Urban Driving. IEEE Transactions on Instrumentation and Measurement, 71, 1-17. https://doi.org/10.1109/tim.2021.3139658

[5] Park, S., & Jeong, Y. (2022). Proactive Motion Planning for Uncontrolled Blind Intersections to Improve the Safety and Traffic Efficiency of Autonomous Vehicles. Applied Sciences, 12(22), 11570. https://doi.org/10.3390/app122211570

[6] Yoon, Y., & Yi, K. (2022). Trajectory Prediction Using Graph-Based Deep Learning for Longitudinal Control of Autonomous Vehicles: A Proactive Approach for Autonomous Driving in Urban Dynamic Traffic Environments. IEEE Vehicular Technology Magazine, 17(4), 18-27. https://doi.org/10.1109/mvt.2022.3207305

[7] Yin, B., Menendez, M., & Yang, K. (2021). Joint Optimization of Intersection Control and Trajectory Planning Accounting for Pedestrians in a Connected and Automated Vehicle Environment. Sustainability, 13(3), 1135. https://doi.org/10.3390/su13031135

[8] Zhou, H., Laval, J., Zhou, A., Wang, Y., Wu, W., Qing, Z., & Peeta, S. (2021). Review of Learning-Based Longitudinal Motion Planning for Autonomous Vehicles: Research Gaps Between Self-Driving and Traffic Congestion. Transportation Research Record: Journal of the Transportation Research Board, 2676(1), 324-341. https://doi.org/10.1177/03611981211035764

[9] Jeong, Y., & Yi, K. (2021). Target Vehicle Motion Prediction-Based Motion Planning Framework for Autonomous Driving in Uncontrolled Intersections. IEEE Transactions on Intelligent Transportation Systems, 22(1), 168-177. https://doi.org/10.1109/tits.2019.2955721

[10] 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

[11] Kim, D., Kim, G., Kim, H., & Huh, K. (2022). A Hierarchical Motion Planning Framework for Autonomous Driving in Structured Highway Environments. IEEE Access, 10, 20102-20117. https://doi.org/10.1109/access.2022.3152187

[12] Ammour, M., Orjuela, R., & Basset, M. (2022). A MPC Combined Decision Making and Trajectory Planning for Autonomous Vehicle Collision Avoidance. IEEE Transactions on Intelligent Transportation Systems, 23(12), 24805-24817. https://doi.org/10.1109/tits.2022.3210276

[13] Liu, J., Luo, Y., Zhong, Z., Li, K., Huang, H., & Xiong, H. (2022). A Probabilistic Architecture of Long-Term Vehicle Trajectory Prediction for Autonomous Driving. Engineering, 19, 228-239. https://doi.org/10.1016/j.eng.2021.12.020

[14] Rosati Papini, G. P., Plebe, A., Lio, M. D., & Dona, R. (2022). A Reinforcement Learning Approach for Enacting Cautious Behaviours in Autonomous Driving System: Safe Speed Choice in the Interaction With Distracted Pedestrians. IEEE Transactions on Intelligent Transportation Systems, 23(7), 8805-8822. https://doi.org/10.1109/tits.2021.3086397

[15] Huang, Y., Du, J., Yang, Z., Zhou, Z., Zhang, L., & Chen, H. (2022). A Survey on Trajectory-Prediction Methods for Autonomous Driving. IEEE Transactions on Intelligent Vehicles, 7(3), 652-674. https://doi.org/10.1109/tiv.2022.3167103

[16] Zhang, Y., Hao, R., Zhang, T., Chang, X., Xie, Z., & Zhang, Q. (2022). A Trajectory Optimization-Based Intersection Coordination Framework for Cooperative Autonomous Vehicles. IEEE Transactions on Intelligent Transportation Systems, 23(9), 14674-14688. https://doi.org/10.1109/tits.2021.3131570

[17] Kalatian, A., & Farooq, B. (2022). A context-aware pedestrian trajectory prediction framework for automated vehicles. Transportation Research Part C: Emerging Technologies, 134, 103453. https://doi.org/10.1016/j.trc.2021.103453

[18] Roszyk, K., Nowicki, M. R., & Skrzypczyński, P. (2022). Adopting the YOLOv4 Architecture for Low-Latency Multispectral Pedestrian Detection in Autonomous Driving. Sensors, 22(3), 1082. https://doi.org/10.3390/s22031082

[19] Yu, M., & Long, J. (2022). An Eco-Driving Strategy for Partially Connected Automated Vehicles at a Signalized Intersection. IEEE Transactions on Intelligent Transportation Systems, 23(9), 15780-15793. https://doi.org/10.1109/tits.2022.3145453

[20] Wu, W., Liu, Y., Liu, W., Zhang, F., Dixit, V., & Waller, S. T. (2022). Autonomous Intersection Management for Connected and Automated Vehicles: A Lane-Based Method. IEEE Transactions on Intelligent Transportation Systems, 23(9), 15091-15106. https://doi.org/10.1109/tits.2021.3136910

[21] Zhou, D., Ma, Z., Zhang, X., & Sun, J. (2022). Autonomous vehicles' intended cooperative motion planning for unprotected turning at intersections. IET Intelligent Transport Systems, 16(8), 1058-1073. https://doi.org/10.1049/itr2.12195

[22] Chen, L., Hu, X., Tang, B., & Cheng, Y. (2022). Conditional DQN-Based Motion Planning With Fuzzy Logic for Autonomous Driving. IEEE Transactions on Intelligent Transportation Systems, 23(4), 2966-2977. https://doi.org/10.1109/tits.2020.3025671

[23] Lakshmanan, V. K., Sciarretta, A., & Ganaoui-Mourlan, O. E. (2022). Cooperative Control in Eco-Driving of Electric Connected and Autonomous Vehicles in an Un-Signalized Urban Intersection. IFAC-PapersOnLine, 55(24), 64-71. https://doi.org/10.1016/j.ifacol.2022.10.263

[24] Zhang, M., Abbas-Turki, A., Mualla, Y., Koukam, A., & Tu, X. (2022). Coordination Between Connected Automated Vehicles and Pedestrians to Improve Traffic Safety and Efficiency at Industrial Sites. IEEE Access, 10, 68029-68041. https://doi.org/10.1109/access.2022.3185734

[25] Tang, X., Yang, K., Wang, H., Yu, W., Yang, X., Liu, T., & Li, J. (2022). Driving Environment Uncertainty-Aware Motion Planning for Autonomous Vehicles. Chinese Journal of Mechanical Engineering, 35(1), 120. https://doi.org/10.1186/s10033-022-00790-5

[26] Shu, H., Liu, T., Mu, X., & Cao, D. (2022). Driving Tasks Transfer Using Deep Reinforcement Learning for Decision-Making of Autonomous Vehicles in Unsignalized Intersection. IEEE Transactions on Vehicular Technology, 71(1), 41-52. https://doi.org/10.1109/tvt.2021.3121985

[27] Zhang, S., Jian, Z., Deng, X., Chen, S., Nan, Z., & Zheng, N. (2022). Hierarchical Motion Planning for Autonomous Driving in Large-Scale Complex Scenarios. IEEE Transactions on Intelligent Transportation Systems, 23(8), 13291-13305. https://doi.org/10.1109/tits.2021.3123327

[28] Tang, X., Huang, B., Liu, T., & Lin, X. (2022). Highway Decision-Making and Motion Planning for Autonomous Driving via Soft Actor-Critic. IEEE Transactions on Vehicular Technology, 71(5), 4706-4717. https://doi.org/10.1109/tvt.2022.3151651

[29] Li, P., Pei, X., Chen, Z., Zhou, X., & Xu, J. (2022). Human-like motion planning of autonomous vehicle based on probabilistic trajectory prediction. Applied Soft Computing, 118, 108499. https://doi.org/10.1016/j.asoc.2022.108499

[30] Zhang, Y., Zhang, J., Zhang, J., Wang, J., Lu, K., & Hong, J. (2022). Integrating Algorithmic Sampling-Based Motion Planning with Learning in Autonomous Driving. ACM Transactions on Intelligent Systems and Technology, 13(3), 1-27. https://doi.org/10.1145/3469086

[31] Nan, J., Deng, W., & Zheng, B. (2022). Intention Prediction and Mixed Strategy Nash Equilibrium-Based Decision-Making Framework for Autonomous Driving in Uncontrolled Intersection. IEEE Transactions on Vehicular Technology, 71(10), 10316-10326. https://doi.org/10.1109/tvt.2022.3186976

[32] Karbasi, A., & O'Hern, S. (2022). Investigating the Impact of Connected and Automated Vehicles on Signalized and Unsignalized Intersections Safety in Mixed Traffic. Future Transportation, 2(1), 24-40. https://doi.org/10.3390/futuretransp2010002

[33] Chen, Y., Yu, H., Zhang, J., & Cao, D. (2022). Lane-Exchanging Driving Strategy for Autonomous Vehicle via Trajectory Prediction and Model Predictive Control. Chinese Journal of Mechanical Engineering, 35(1), 71. https://doi.org/10.1186/s10033-022-00748-7

[34] Yin, S., Yang, C., Kawsar, I., Du, H., & Pan, Y. (2022). Longitudinal Predictive Control for Vehicle-Following Collision Avoidance in Autonomous Driving Considering Distance and Acceleration Compensation. Sensors, 22(19), 7395. https://doi.org/10.3390/s22197395

[35] Parseh, M., Nybacka, M., & Asplund, F. (2022). Motion planning for autonomous vehicles with the inclusion of post-impact motions for minimising collision risk. Vehicle System Dynamics, 61(6), 1707-1733. https://doi.org/10.1080/00423114.2022.2088396

[36] Saleh, K. (2022). Pedestrian Trajectory Prediction for Real-Time Autonomous Systems via Context-Augmented Transformer Networks. Sensors, 22(19), 7495. https://doi.org/10.3390/s22197495

[37] Zhang, M. Y., Yang, S. C., Feng, X. J., Chen, Y. Y., Lu, J. Y., & Cao, Y. G. (2022). Route Planning for Autonomous Driving Based on Traffic Information via Multi-Objective Optimization. Applied Sciences, 12(22), 11817. https://doi.org/10.3390/app122211817