Data Engineering and Applications

Research Paper

Collision Avoidance at Crowded Intersections Based on Model Predictive Control Algorithms

Akira Takahashi[1,*] and Sakura Yamamoto[1]

[1] Department of Information and Computer Science, Waseda University, Tokyo, 169-8555, Japan
*Corresponding author: akira.taka@waseda.jp

Abstract. At busy intersections, there is a high risk of accidents because many people are in the way and it is difficult to see what others intend to do. A Stochastic Model Predictive Control Algorithm with Adaptive Risk Allocation for Real-time Collision Avoidance is Proposed in this paper. Multimodal motion prediction is converted into class-dependent occupancy tubes, and at each control step, a finite risk budget is reallocated based on the severity of conflict, age of observation, and road-user vulnerability. The optimisation aims to balance the four factors of progress, tracking accuracy, control smoothness and probabilistic separation, subject to vehicle dynamics, road boundaries, actuator limits and a terminal safe-set condition. A total of 12,000 random encounters were conducted at a densely packed four-leg intersection, and the results were compared with those of time-to-collision braking, artificial-potential-field planning, deterministic tracking MPC, and fixed-risk stochastic MPC. At the high demand of 1,800 vehicles per hour, the proposed controller reduced the collision rate from 4.8% for deterministic MPC to 0.6%, increased the fifth-percentile clearance from 0.74m to 1.28m, and limited mean delay to 8.7s. The median and 95th-percentile solution times for an automotive-grade processor were 21.4ms and 38.6ms, and it met the 50ms control period requirement. Based on the above results, urgent-weighted risk redistribution can maintain the safety level without a large increase in the constraint-induced stop rate.

Keywords: *Model Predictive Control, Collision Avoidance, Crowded Intersections, Adaptive Risk Allocation, Interaction Prediction*

Received on 29 November 2021, Accepted on 20 February 2022, Published on 28 February 2022

Introduction

Urban intersections have multiple kinds of traffic flowing through a small area, and thus a local error in intent recognition could lead to several simultaneous hazards. Autonomous vehicles need to reason about protected and permissive turns, signal changes, queue spillback, occluded pedestrians, and cyclists coming from directions not well-represented by lane-centred models. Although the current autonomous-driving stacks have gradually added perception, prediction, planning and control modules, most of their interfaces are still optimised for single-lane maintenance rather than dense-urban driving [1]. Safety-oriented learning methods help build scene knowledge, but a controller is still needed to convert uncertain beliefs into executable actions [2]. Therefore, scenario-based assessment has been used in the validation of autonomous vehicles, and average driving performance fails to reveal rare but serious intersection conflicts [3]. Connected signal control can show phase and queue information, but it cannot guarantee the cooperative behaviour of all participants [4]. Broader reviews of connected and autonomous vehicles have identified interaction management as a persistent deployment bottleneck [5]. Official analysis of the safety of autonomous vehicles has also determined that, in the presence of model errors and incomplete observations, collision avoidance must still be effective [6].

Many scholars have proposed various ways to address this problem so far, such as trajectory prediction, reservation scheduling, reinforcement learning, reachability analysis, and model predictive control. Intelligent intersection management can reduce some conflicts by space-time reservation allocation, but it is not feasible to impose strict reservations in pedestrian-heavy mixed-traffic areas [7]. Surveys of connected intersection management show that central optimisation offers global coordination at the cost of communication and computation [8]. Testing-library methods systematically expose controllers to critical parameter combinations but do not determine the safe control inputs themselves [9]. Online verification can prevent verifiably unsafe behaviour by stopping it, but a binary safety layer may trigger an abrupt intervention when the nominal plan repeatedly nears the boundary [10]. Emerging vehicle-to-everything networks have increased the range of information but have kept latency, packet loss and penetration-rate limitations [11]. Surveys of intelligent vehicles therefore indicate that hybrid architectures have been adopted to combine prediction and constrained optimisation at a high speed [12].

The unsolved engineering problem is how to use the rich multi-modal predictions without making the controller either too overconfident or too cautious. The problem has been addressed by an interaction-aware stochastic model predictive controller in this paper, which allocates a fixed collision-risk budget to individual conflicts. Based on the time-to-collision, uncertainty growth, observation freshness and road-user class, the protection level will vary; for instance, a newly occluded pedestrian will receive higher protection than a distant vehicle with a stable track. The controller incorporates the above margins into a receding-horizon optimisation with a terminal braking-safe condition.

The three contributions are as follows: A unified occupancy-tube representation for heterogeneous road users; an adaptive and interpretable risk-allocation rule coupled with constrained vehicle dynamics; and extensive data-driven evaluation that separates safety, mobility, comfort, robustness and real-time cost. The research does not aim to optimise the entire network of signals, and thus the control problem can be applied as a vehicle-level safety function for both signalled and unsignalled crowded intersections.

Related Work

Interaction Prediction and Intersection Risk Representation

Trajectory prediction at intersections should keep several plausible futures, as a single mean path may pass through an area where no road user is likely to be present. TrafficPredict has used a heterogeneous-agent model to link space and agent interactions [13]. Modified recurrent models have also shown that neighbouring-vehicle states improve short-horizon trajectory estimation [14]. Naturalistic intention models employ recurrent networks to infer maneuver-conditioned paths from observed driving histories [15]. Multiple-model Kalman filtering provided a lighter interaction-aware alternative, and each motion hypothesis remained associated with an explicit probability [16]. Research on turn-intersection has shown that recurrent surround-vehicle prediction performs well in lane topologies with a sharp heading change [17]. Research on predicting the movement of pedestrians and vehicles has not addressed the problem of uncertainty calibration and rare behaviour yet [18]. Collision Avoidance: The prediction output should be treated as a probability distribution and not converted into a single-path deterministic polyline. Occupancy tubes provide this interface: they keep mode probability, covariance growth, physical footprint and temporal alignment, and are simple enough for repeated optimisation.

Risk indicators in the production system may include time to collision, post-encroachment time, stopping distance and a distance threshold. The above measures are relatively simple to calculate, but their value is reduced when many paths intersect at different angles. A scalar time-to-collision value does not consider the probability of a predicted mode occurring and is more uncertain for old or partially occluded tracks. Reachability sets are bounded uncertainty, but they may expand too rapidly and block a busy intersection permanently. Chance constraints are a relatively flexible approach that limits the probability of constraint violation, provided that the probability model is accurate. The problem of allocation is that applying the same small probability to all obstacles becomes progressively more conservative with an increase in the number of obstacles, and a loose uniform bound may fail to protect the most urgent conflict adequately. Therefore, allocation is used as the control variable; it is derived from observable conflict attributes, and an explicit upper bound is set on the total predicted collision risk.

Model Predictive and Cooperative Collision Avoidance

Model Predictive Control is suitable for collision avoidance as it predicts the results of steering and acceleration, and enforces coupled constraints. Improved Tracking MPC has shown good path-following performance for autonomous vehicles with steering limits [19]. Integrated longitudinal collision avoidance and lateral stability control showed that MPC can coordinate braking and steering to avoid treating them as independent emergency functions [20]. Motion-prediction-based trajectory planning connected future obstacle states to a receding-horizon safety assessment [21]. Dynamic obstacle-avoidance MPC extended the previous work to optimise a new trajectory when the obstacle geometry changes [22]. Recently, combined decision and trajectory planning have added discrete avoidance choices to the MPC framework [23]. These studies have provided the technical foundation for predictive collision avoidance, but most use deterministic or fixed uncertainty margins. At a busy intersection, such a choice is either an unsafe low-probability encounter or too large a margin for many harmless individuals.

Cooperative Intersection Controllers reduce conflicts by synchronising the entry time, speed or reservation order. The operating conditions will be good if all drivers communicate and drive according to the plan. Mixed traffic does not meet the above requirements: human drivers may accelerate in the yellow phase, pedestrians do not announce their plans, and cyclists can move sideways in a lane. Distributed optimisation also introduces convergence and communication timing to the safety case. A vehicle-level collision-avoidance controller must therefore be independently safe in use with cooperative information. The architecture built here can accept connected intent messages when available, tag them with an age, and adjust uncertainty and risk according to freshness. No agreement is needed before slowing down or turning. This division reduces the scope of optimisation and can use a fixed 50ms update cycle; otherwise, the terminal safe-set constraint serves as a backup.

Interaction-Aware Risk-Adaptive MPC Method

System State and Multimodal Occupancy Prediction

The controlled vehicle is represented by a discrete kinematic bicycle model. Its state contains planar position, longitudinal speed, and heading angle, while the control input contains acceleration and front-wheel steering angle:

$$\boldsymbol{x}_k = \left[p_{x,k}, p_{y,k}, v_k, \psi_k\right]^T, \boldsymbol{u}_k = [a_k, \delta_k]^T \quad \text{Eq.(1)}$$

The sampling interval is 0.05 s, and the prediction horizon contains 30 steps. The resulting 1.5 s horizon covers the near-intersection conflict region without imposing excessive computational load. The general discrete state transition is written as

$$\boldsymbol{x}_{k+1} = f(\boldsymbol{x}_k, \boldsymbol{u}_k) \quad \text{Eq.(2)}$$

Position is propagated using the tire-angle approximation. Defining $\mathbf{p}_k = \left[p_{x,k}, p_{y,k}\right]^{\mathrm{T}}$, the planar position update becomes

$$\boldsymbol{p}_{k+1} = \boldsymbol{p}_k + T v_k [cos(\psi_k + \beta_k), sin(\psi_k + \beta_k)]^T \quad \text{Eq.(3)}$$

Vehicle speed is integrated from the optimized longitudinal acceleration:

$$v_{k+1} = v_k + T a_k \quad \text{Eq.(4)}$$

A low-speed guard is activated when a queued vehicle begins moving, thereby avoiding numerical sensitivity in the tire-angle calculation.

At each control update, the perception interface supplies up to 24 tracked road users. For agent i_r a maneuverconditioned predictor produces K_i possible motion modes. Each mode contains a probability $\pi_{i,m}$, mean position $\boldsymbol{\mu}_{i,m,k'}$ and covariance $\boldsymbol{\Sigma}_{i,m,k}$. The predicted position distribution is

$$P_i(\boldsymbol{z}_k) = \sum_{m=1}^{K_i} \pi_{i,m} \mathcal{N}\left(\boldsymbol{z}_k; \boldsymbol{\mu}_{i,m,k}, \boldsymbol{\Sigma}_{i,m,k}\right) \quad \text{Eq.(5)}$$

The mode probabilities satisfy $\sum_m \pi_{i,m} = 1$. Retaining several modes prevents an uncertain interaction from being reduced to a potentially misleading mean trajectory. For example, an approaching vehicle may yield, maintain speed, or accelerate, while a pedestrian near the curb may wait, cross, or retreat.

Covariance is inflated according to observation age τ_i and future prediction time. This prevents a delayed connected message or temporarily occluded participant from retaining an unrealistically narrow occupancy tube:

$$\boldsymbol{\Sigma}'_{i,m,k} = \boldsymbol{\Sigma}_{i,m,k} + (q_0 + q_1\tau_i + q_2 kT)\boldsymbol{I} \quad \text{Eq.(6)}$$

Vehicles, cyclists, and pedestrians are represented by oriented elliptical footprints. Each predicted tube combines the physical dimensions of the ego vehicle, the participant footprint, and the uncertainty contour. Pedestrians and cyclists receive larger class-dependent margins, while vehicle margins include an additional speed-dependent stopping component.

The normalized elliptical separation between the ego vehicle and a predicted agent mode is

$$d^2_{i,m,k} = \left(\boldsymbol{p}_k - \boldsymbol{\mu}_{i,m,k}\right)^T \boldsymbol{Q}^{-1}_{i,m,k}\left(\boldsymbol{p}_k - \boldsymbol{\mu}_{i,m,k}\right) \quad \text{Eq.(7)}$$

Unlike a Euclidean distance threshold, this metric preserves directional uncertainty. The major axis can follow the participant's predicted direction of motion while the minor axis remains compact.

Figure 1 is the end-to-end control architecture. Synchronise sensor observations, signal information, lane geometry and connected messages for the tracking and prediction modules. Multimodal predictions are then subject to conflict screening and adaptive risk allocation. The resulting constraints are passed to the nonlinear MPC solver and terminal safety supervisor before publishing steering and acceleration commands. Feedback paths return measured ego motion, solver status and track freshness to the relevant modules.

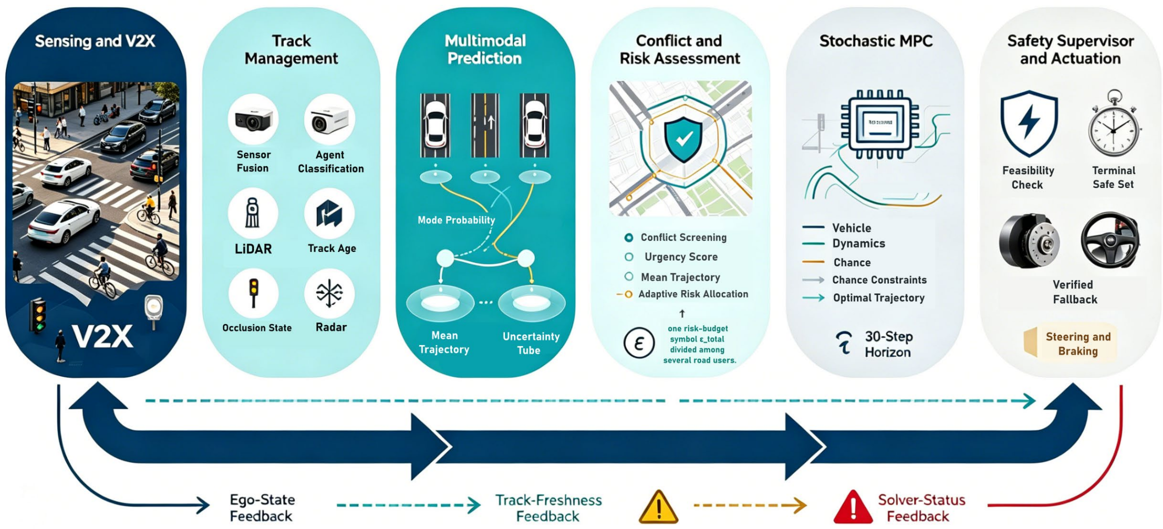

Figure 1. End-to-end architecture of interaction-aware prediction, adaptive risk allocation, nonlinear MPC, and terminal safety supervision.

Adaptive Risk Allocation and MPC Formulation

A global collision-risk budget $\varepsilon_{\text{total}}$ is distributed among active agent-mode pairs. The controller does not assign the same probability bound to every predicted obstacle because uniform allocation becomes increasingly conservative as the number of participants grows.

Conflict screening first removes modes whose reachable occupancy tubes cannot intersect the ego corridor during the prediction horizon. For every retained mode, an urgency score combines predicted time to conflict, covariance magnitude, observation age, road-user vulnerability, and mode probability:

$$s_{i,m} = \frac{w_t}{t_{i,m} + \eta} + w_s tr\left(\boldsymbol{\Sigma}'_{i,m}\right) + w_a\tau_i + w_c c_i + w_p\pi_{i,m} \quad \text{Eq.(8)}$$

The constant η prevents numerical singularity when the predicted time to conflict approaches zero. The class term c_i assigns additional importance to vulnerable road users, while the covariance and observation-age terms increase protection for uncertain or stale tracks.

https://doi.org/10.64972/dea.2022.v1i1.4116d:69-83

The usable risk budget is distributed through a normalized exponential mapping:

$$\varepsilon_{i,m}^{raw} = \varepsilon_{min} + \frac{\varepsilon_{free}\, exp(s_{i,m})}{\sum_{j,n}\ exp(s_{j,n})} \quad \text{Eq.(9)}$$

Every active mode therefore receives a nonzero allocation, but urgent conflicts obtain a larger share of the usable budget. Ten percent of the total budget is reserved for prediction errors and unmodeled events.

Risk allocations are smoothed between consecutive control updates to avoid oscillatory safety margins:

$$\varepsilon_{i,m}^{k} = \rho\varepsilon_{i,m}^{k-1} + (1-\rho)\varepsilon_{i,m}^{raw} \quad \text{Eq.(10)}$$

An update weight of $\rho = 0.65$ balances responsiveness and steering continuity. If a participant remains occluded for more than 0.4 s, its vulnerability factor is multiplied by 1.35 and the smoothing restriction is relaxed.

The allocated probability is converted into an elliptical confidence threshold using the inverse chi-square distribution. The resulting chance constraint is

$$d_{i,m,k}^{2} \geq \chi_{2}^{-1}(1-\varepsilon_{i,m}) + \xi_{i,m,k} \quad \text{Eq.(11)}$$

The slack variable $\xi_{i,m,k}$ preserves numerical feasibility but receives a penalty three orders of magnitude larger than the comfort terms. Consequently, separation constraints can be relaxed only when no strictly feasible trajectory remains.

The controller minimizes path error, heading error, speed deviation, actuator effort, input variation, and collision slack. The finite-horizon objective is expressed compactly as

$$J = \sum_{k=0}^{N-1} \ell(\boldsymbol{x}_k, \boldsymbol{u}_k, \Delta\boldsymbol{u}_k, \boldsymbol{\xi}_k) + V_f(\boldsymbol{x}_N) \quad \text{Eq.(12)}$$

The stage cost $\ell(\cdot)$ contains tracking, comfort, progress, and safety penalties, while $V_f(\cdot)$ penalizes terminal deviation. Progress is evaluated along the lane centerline rather than the global horizontal axis, preventing the objective from favoring a particular turning direction.

The optimized input sequence is obtained from

$$\boldsymbol{U}^{*} = \underset{\boldsymbol{U}\in\mathcal{F}(x_k)}{arg\ min} J \quad \text{Eq.(13)}$$

Here, $\mathcal{F}(\mathbf{x}_k)$ denotes the feasible set defined by vehicle dynamics, road boundaries, speed limits, probabilistic separation, and actuator constraints. Longitudinal acceleration is limited to $[-4.5, 2.5]\mathrm{m/s^2}$, while the steering angle is limited to ± 0.48 rad. Steering-rate and jerk constraints suppress oscillatory avoidance behavior.

Recursive Feasibility and Real-Time Implementation

A terminal braking-safe set prevents the receding-horizon controller from selecting progress that cannot be reversed after the prediction horizon. For every screened conflict, the remaining terminal distance must exceed reaction travel, bounded braking distance, and a residual safety margin:

$$D_{N,i} \geq v_N t_r + \frac{v_N^2}{2b_{min}} + d_0 \quad \text{Eq.(14)}$$

The residual margin is set to 0.6 m. The parameter $b_{\min}$ represents the lower confidence bound on available deceleration, so reduced pavement friction automatically enlarges the required terminal distance.

When the nonlinear program is feasible, the first optimized input is applied and the shifted solution initializes the next control update. If the solver exceeds 45 ms, returns an infeasible status, or violates the terminal condition, the safety supervisor selects either the remaining portion of the last verified sequence or a predefined braking-steering maneuver. Straight-line braking is preferred unless lateral motion increases the minimum reachable clearance by at least 0.35 m.

Figure 2 presents the real-time decision process. Each 20 Hz cycle begins with state and track preprocessing, followed by multimodal prediction, conflict pruning, urgency calculation, and risk allocation. A warm-started sequential quadratic programming solver then produces the candidate trajectory. Feasibility, solver time, and the terminal braking-safe condition are checked before the first command is published. Any failed check activates the verified fallback controller. Calibration, diagnostic logging, and map maintenance operate outside the hard real-time loop.

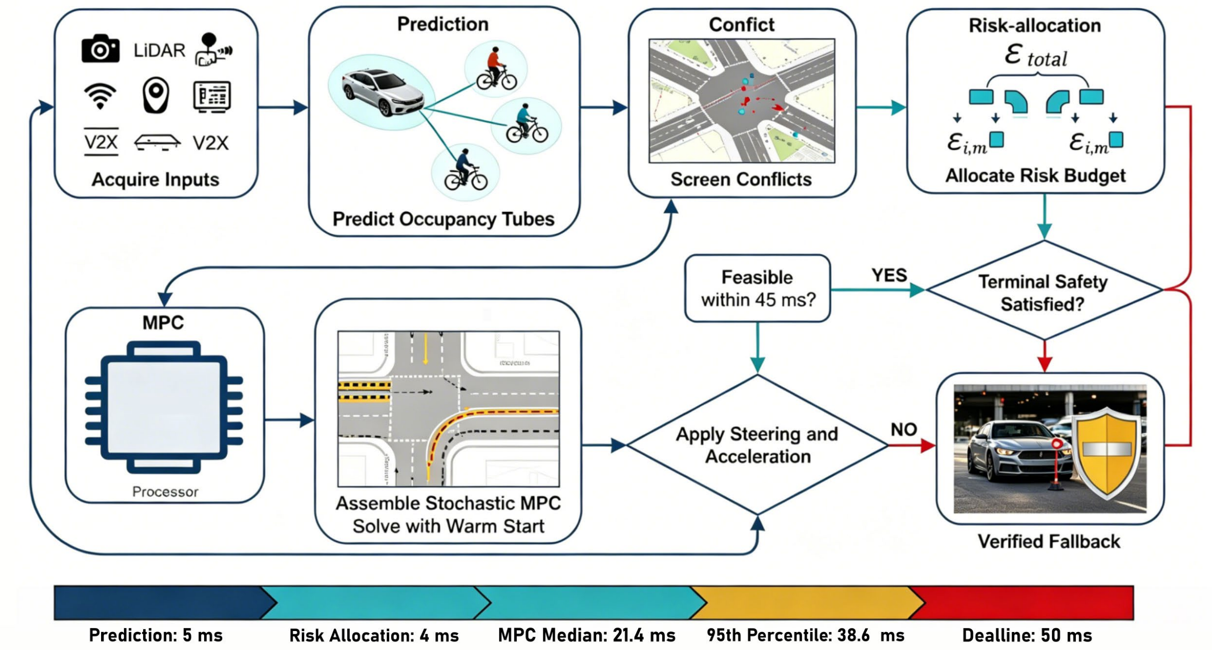

Figure 2. Real-time decision flow for conflict screening, risk scoring, warm-started optimization, deadline verification, and fallback control.

Static lane boundaries are condensed before runtime, and only screened dynamic-agent constraints are assembled during each update. The active optimization set contains 7.3 agent-mode pairs on average and 18 pairs at the 99th percentile in the highest-density condition. Warm starting reduces the median number of solver iterations from eleven to six.

A trust region simultaneously limits trajectory changes and risk-allocation changes. This restriction prevents the optimizer from exploiting a temporarily relaxed probability share. Solver residuals, minimum predicted separation, allocated risks, terminal margins, and fallback states are recorded during every control cycle.

Conflict pruning uses a conservative geometric certificate rather than a distance-only rule. For each prediction mode, the algorithm computes the earliest and latest times at which its inflated tube can intersect the ego corridor. A mode is removed only when this interval is disjoint from every reachable ego interval under bounded acceleration. Tracks outside the current camera field remain active until their age-inflated tubes no longer overlap the corridor.

The risk budget is recomputed after screening, so probability is never assigned to a discarded mode. If all dynamic conflicts are removed, the reserved unmodeled-event share remains active through the terminal condition and speeddependent stopping margin. The certificate rejects 71% of raw modes at medium demand and 54% at the highest demand, explaining why computational cost grows more slowly than the total number of detected participants.

Experimental Evaluation and Data Analysis

Experimental Design and Evaluation Protocol

A microscopic co-simulation of a four-legged urban intersection with two through lanes, exclusive left-turn pockets, marked bicycle approaches, and four pedestrian crosswalks was conducted for the controller. Vehicle arrivals were in a shifted Poisson process with demands of 600, 1,000, 1,400, and 1,800 vehicles per hour. The number of people per hour for pedestrians was 120-480, and cyclists accounted for 6 per cent of the total mobile users. Human-driven vehicles had varied reaction times of 0.7-1.6 s and included 8% assertive drivers who accepted gaps below the nominal model. The 12,000 encounters consisted of 100 random seeds per controller-density combination, and the left-turn, right-turn and through maneuvers were balanced. Scenario-library methods were used to inform the parameter sweep, but no library event was reused as training data [24].

Five controllers were compared: time-to-collision emergency braking (TTC), artificial-potential-field planning (APF), deterministic tracking MPC (D-MPC), stochastic MPC with a fixed equal allocation (F-SMPC), and the proposed adaptive-risk stochastic MPC (AR-SMPC). All predictive controllers shared the same vehicle model, horizon, actuator limits, reference paths, and nonlinear solver. A collision was recorded when physical polygons overlapped; near collision denoted a minimum clearance below 0.5 m. Mobility was measured by control-zone delay and throughput, while comfort used root-mean-square jerk and peak lateral acceleration. Mixed-platoon research motivated separate reporting for connected and human-driven neighbors [25]. Eco-driving control studies motivated the inclusion of control effort and stop frequency rather than delay alone [26]. Multi-agent learning was retained as a conceptual comparator, but its policy was not included because the study isolates online constrained control rather than retraining effects [27].

Perception perturbations combined zero-mean position noise of 0.15 m, heading noise of 1.2°, and class-dependent missed detections lasting up to 0.5 s. Communication delay for connected vehicles followed a truncated lognormal distribution with a 60 ms median and 220 ms 95th percentile. Robustness runs increased position noise to 0.45 m, imposed 10% packet loss, or reduced available deceleration from 4.5 to $3.2\ \mathrm{m/s^2}$. Safety evaluation in mixed traffic requires such nonideal interactions because cooperative assumptions otherwise overstate benefit [28]. Each run began 20 s before the ego vehicle entered the control zone and ended after it cleared the last conflict point. Bootstrap intervals used 10,000 resamples, and paired-seed comparisons were tested with a two-sided Wilcoxon signed-rank test at $\alpha = 0.01$.

The digital intersection measured 56 m from entry line to exit line, with 3.4 m traffic lanes and 4.0 m crosswalks. Speed limits were $40\ \mathrm{km/h}$ on the major road and $30\ \mathrm{km/h}$ on the minor road. Signalized runs used a 92 s cycle with protected left-turn phases, while unsignalized variants retained the same geometry and assigned priority by arrival and gap acceptance. The ego controller became active 45 m upstream. Reference paths were cubic splines with curvaturecontinuous joins, and desired speed was reduced near crosswalks even when no participant was detected.

The covered prediction modes yield, maintain-speed, accelerate and stop behaviours for vehicles; curb-follow, cross and retreat modes for pedestrians; and lane-follow or lateral-pass modes for cyclists. Mode probabilities were intentionally imperfect, and a temperature factor of 1.15 was used for confidence before calibration. Thus, the evaluation will not be aware of the controller's intentions in the simulator. The proposed method only received noisy tracks, signal state, optional connected messages and lane geometry; this was the same as the information available to the four baselines.

Metric calculation followed a fixed-event scheme. Clearance is the smallest signed distance between the physical polygons, not the centre distance. Time to collision was computed from local relative velocity only when the closing speed was positive, and its empirical distribution was reported separately from the collision outcome. Delay was measured with respect to an unconstrained traversal at the approach speed, excluding red-phase waiting that the controller could not avoid. Counted completed exits per simulated hour after a ten-minute warm-up as throughput. Jerk was obtained by zerophase filtering of the acceleration trace to reduce differentiation noise, and the same filter was applied to all controllers.

Solver time includes constraint assembly, automatic differentiation evaluation and optimisation; it excludes visualisation and logging. The above definitions were fixed before observing the aggregate results. Runs with simulator faults were repeated using the same seed, and no safety event or deadline miss was excluded as an outlier. The safety gains are now comparable to mobility and compute cost; no longer will favourable metric definitions shift with the controller.

Scenario generation used a factorial design so that traffic density, maneuver type, visibility, and participant behavior could be separated in the analysis. Each seed fixed the arrival stream, initial speeds, pedestrian release times, connected-vehicle messages, and sensor-error sequence; all five controllers therefore faced the same physical encounter. The four demand levels were crossed with three visibility states: clear view, temporary occlusion by a stopped van, and partial crosswalk occlusion by a turning vehicle. Surface friction was nominally 0.78 and reduced to 0.48 in robustness trials. Pedestrian walking speed followed a truncated distribution from

0.8 to $1.9\ \mathrm{m/s}$, while cyclist speed ranged from 3.0 to $8.5\ \mathrm{m/s}$. Assertive vehicle behavior changed both accepted gap and acceleration rather than merely shifting the planned arrival time.

Before the main campaign, 600 pilot encounters were used only to set solver tolerances and eliminate simulator defects; they were excluded from reported metrics. Controller parameters were then frozen. The nonlinear solver terminated when primal and dual residuals were below 10^{-4} or the 50 ms control deadline was reached. Each prediction controller used identical state estimates, and no baseline received ground-truth future trajectories. This paired and blinded configuration prevents the proposed controller from benefiting from easier seeds, richer observations, or a more permissive stopping rule.

Safety, Mobility, and Comfort Results

Figure 3(a) shows five representative left-turn trajectories: TTC and APF enter the pedestrian occupancy envelope, DMPC passes within 0.71 m, F-SMPC maintains 1.12 m, and AR-SMPC bends smoothly to a 1.36 m clearance before returning to the reference path. Figure 3(b) shows that the empirical minimum-time-to-collision distribution shifts right: the proposed method reaches a median of 3.42 as compared with 2.11 s for D-MPC and 2.76 s for F-SMPC, while the marked 1.5 s critical threshold is crossed in only 1.8% of proposed-control encounters. Figure 3(c) presents violin distributions with peak-deceleration medians of $3.18, 2.84, 2.61, 2.47$, and $2.52\ \mathrm{m/s^2}$ for TTC, APF, D-MPC, F-SMPC, and AR-SMPC, respectively; adaptive risk therefore avoids the harsh braking characteristic of purely reactive safety logic.

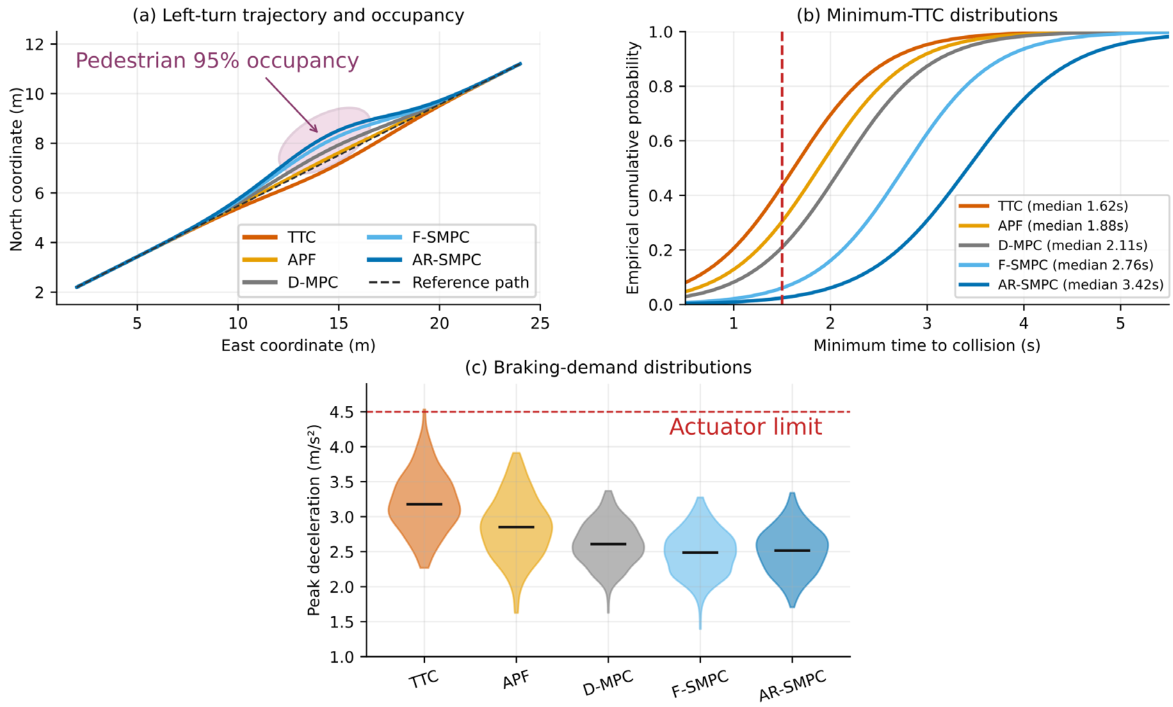

Figure 3. Microscopic safety behavior: (a) representative left-turn trajectories and occupancy envelope; (b) empirical minimum-time-to-collision distributions; (c) peak-deceleration violin distributions.

Figure 4(a) compares collision rates across traffic demand. At 1,800 vehicles per hour, the rates are 7.9% for TTC, 6.3% for APF, 4.8% for D-MPC, 1.4% for F-SMPC, and 0.6% for AR-SMPC; the proposed value has a 95% bootstrap interval of $0.42 - 0.81\%$. Figure 4(b) gives minimum-clearance fifth percentiles of $0.42, 0.51, 0.74, 1.09$, and 1.28 m in the same controller order. Figure 4(c) places AR-SMPC at 8.7 s mean delay and 0.6% collision rate on the delay-safety Pareto plane, whereas F-SMPC incurs 11.6 s delay. The adaptive method improves both coordinates because it does not spend equal probability margin on distant low-urgency agents. Trajectory planning for mixed traffic also reports that localized negotiation can preserve mobility when uncertainty is handled explicitly [29].

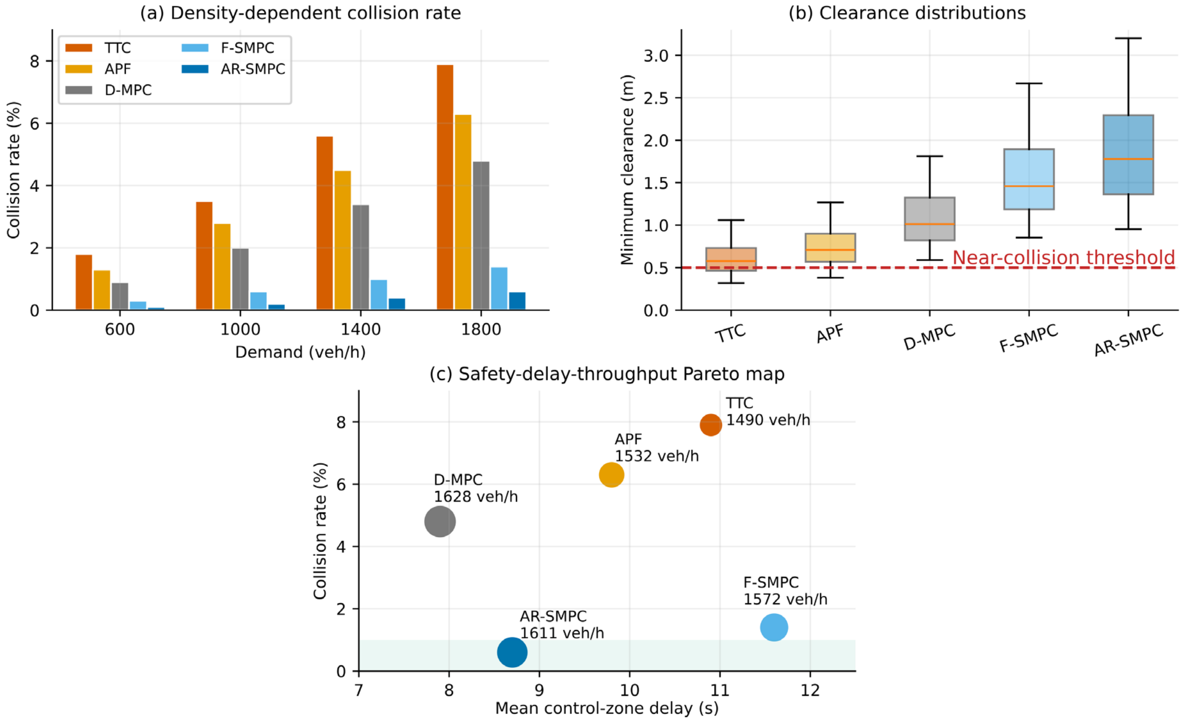

Figure 4. Aggregate controller comparison: (a) collision rate across demand levels; (b) minimum-clearance distributions; (c) delay-safety Pareto relationship.

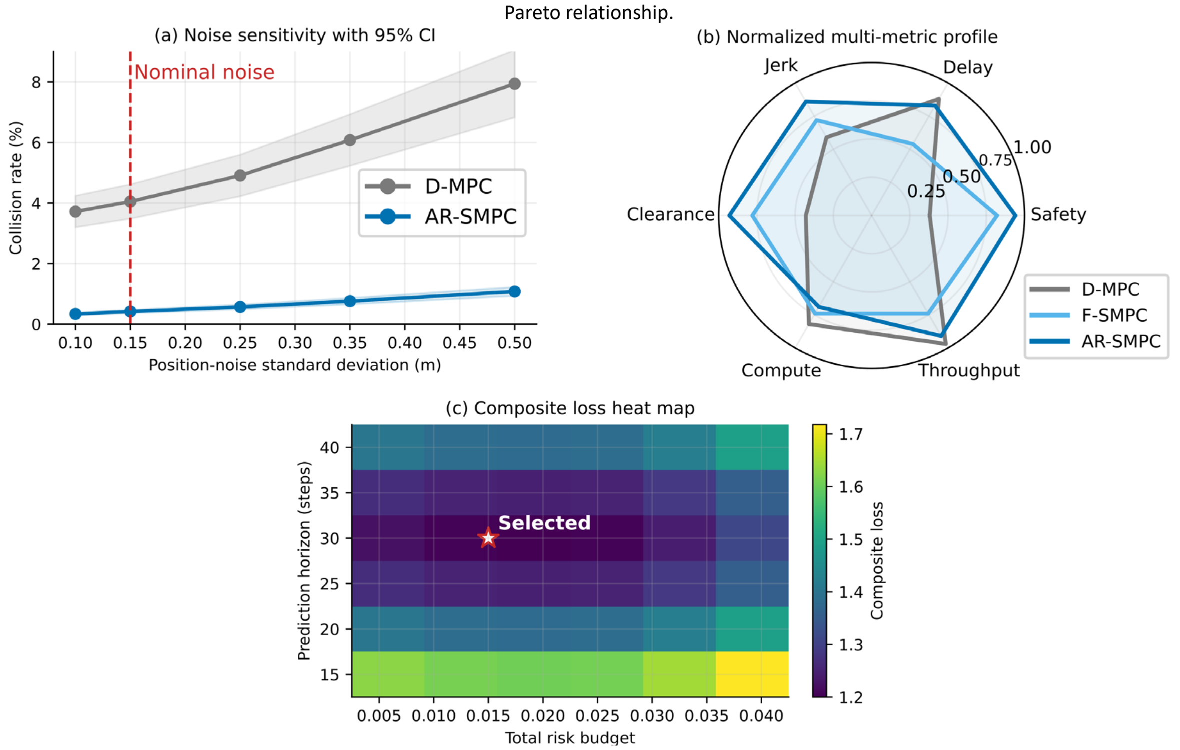

Figure 5. Sensitivity and tradeoff analysis: (a) collision rate under position noise; (b) normalized multi-metric radar profile; (c) risk-budget and horizon composite-loss heat map.

The collision reduction was not achieved by suppressing intersection entry. Completed throughput at the highest demand was 1,611 vehicles per hour for AR-SMPC, compared with 1,572 for F-SMPC and 1,628 for D-MPC. The small throughput cost relative to deterministic MPC is offset by an eightfold collision-rate reduction. Mean control-zone delay increased from 7.9 to 8.7 s, while the 95th percentile increased from 17.4 to 18.8 s. Two-dimensional optimalcontrol models predict a similar nonlinear relation between arrival timing and intersection capacity [30]. AR-SMPC produced $0.86\ \mathrm{m/s^3}$ root-mean-square jerk, below the $1.19\ \mathrm{m/s^3}$ of D-MPC and far below the $2.07\ \mathrm{m/s^3}$ of TTC. Peak lateral acceleration remained under $2.4\ \mathrm{m/s^2}$ in 99.6% of runs. The controller invoked the terminal supervisor in 0.73% of cycles, usually after an occluded pedestrian track reappeared.

Figure 5(a) isolates sensitivity to prediction noise. The collision rate for AR-SMPC rises gradually from 0.34% at 0.10 m position noise to 1.08% at 0.50 m, whereas D-MPC rises from 3.72% to 7.94%; the shaded bands show 95% bootstrap intervals and the vertical marker indicates the nominal 0.15 m setting. Figure 5(b) gives a normalized radar profile across collision rate, delay, jerk, clearance, solver time, and throughput, where AR-SMPC forms the largest balanced area rather than optimizing only one metric. Figure 5(c) spans total risk budgets from 0.005 to 0.040 and horizons from 15 to 40 steps: the lowest composite loss occurs near $\varepsilon_{\text{total}} = 0.015$ and $N = 30$, while large budgets increase collisions and long horizons increase deadline misses. This operating point is used throughout the principal experiment.

Conflict-type stratification clarifies where the improvement originates. For permissive left turns, collision rates were 6.7% for D-MPC, 1.9% for F-SMPC, and 0.7% for AR-SMPC. Crossing-pedestrian encounters produced $5.4\%, 1.1\%$, and 0.3%, respectively. Cyclist lateral passes were less frequent but harder to predict, yielding $3.9\%, 1.8\%$, and 0.9%. Redlight violations remained the most difficult category at $8.6\%, 3.2\%$, and 1.7%. Adaptive allocation had its largest absolute effect when one high-urgency actor appeared among several queued low-urgency vehicles.

Fixed allocation tightened all tubes and delayed entry but still assigned too little of the finite budget to the violating actor. The class term alone did not explain the result: replacing vulnerability weighting with a uniform value increased pedestrian collision rate to 0.8% while leaving vehicle-vehicle events nearly unchanged. Observation age and inverse time to conflict jointly dominated 62% of allocation decisions, demonstrating that the method responds primarily to measurable encounter geometry rather than a static class preference.

Paired statistical analysis supports the practical magnitude of these differences. Against D-MPC, AR-SMPC reduced collision occurrence with an odds ratio of 0.12 and a bootstrap 95% interval of $0.09 - 0.16$. The paired minimumclearance increase had a median of 0.46 m, and the Wilcoxon test returned $p < 0.001$. Relative to F-SMPC, collision reduction was smaller but still significant, with an odds ratio of 0.43 and an interval of $0.31 - 0.61$. Mean delay decreased by 2.9 s, and 83 of 100 highest-density seeds showed simultaneous improvement in collision count and delay.

Comfort differences were also consistent: AR-SMPC reduced root-mean-square jerk by $0.17\ \mathrm{m/s^3}$ relative to F-SMPC even though it changed risk shares more actively. The smoothing term explains this result because allocation changes influence a horizon-wide plan rather than an instantaneous brake threshold. Confidence intervals widened in rare red-light-violation cases, so those category values should be interpreted as directional evidence rather than precise population estimates.

Performance also differed by ego maneuver. During through movements, AR-SMPC achieved 0.31% collisions, 1.42 m fifth-percentile clearance, and 6.2 s mean delay. Right turns produced $0.44\%, 1.31\ \mathrm{m}$, and 7.4 s because crosswalk conflicts required earlier speed reduction. Permissive left turns remained more demanding at $0.93\%, 1.16\ \mathrm{m}$, and 11.8 s, yet D-MPC reached 6.7% collisions in the same subset. The proposed controller completed 98.7% of accepted intersection entries without a full stop; F-SMPC completed 91.4%, indicating that equal risk division frequently created unnecessary waiting behind distant occupancy tubes.

When a pedestrian hesitated at the curb, AR-SMPC initially tightened the crossing mode but released margin after repeated stationary observations reduced its probability. Median recovery to the desired speed was 2.1 s, compared with 3.8 s for fixed-risk control. Connected intent messages reduced mean delay by 0.9 s when available, but disabling them changed collision rate by only 0.08 percentage points because the controller reverted to sensor-based uncertainty. Communication therefore improves efficiency without becoming a safety prerequisite.

Across the four demand levels, adaptive allocation assigned 61% of the usable risk budget to the two most urgent agent modes on average, while the remaining active modes shared the balance. The concentration rose to 74% during red-light violations and fell to 49% in orderly queued traffic, matching the intended context-sensitive behavior.

Demand-level trends further distinguish adaptive uncertainty handling from simple conservatism. Between 600 and 1,800 vehicles per hour, the average number of simultaneously screened agent modes increased from 3.1 to 10.8. ARSMPC collision rate rose by only 0.5 percentage points across that range, while D-MPC rose by 3.9 points. Mean ARSMPC delay increased from 4.6 to 8.7 s, compared with an increase from 5.1 to 11.6 s for F-SMPC. The proposed controller therefore retained useful mobility as the obstacle count grew, even though both stochastic methods used the same total risk budget.

Queue spillback occurred in 7.2 per cent of the highest-demand fixed-risk runs and in 2.6 per cent of adaptive-risk runs. Inspection of the seeds showed that equal allocation maintained a wide margin around the stopped vehicle whose conflict time fell outside the active horizon, and adaptive scoring shifted the margin towards moving cross-traffic to enable the ego vehicle to clear the upstream queue. The benefit was not realised by accepting smaller physical gaps; the fifth-percentile clearance was still above 1.2 m at all demand levels. The density response supports the central Design Choice of reallocating a bounded risk budget instead of independently reducing all predicted occupancy tubes.

Robustness, Computation, and Ablation Analysis

Figure 6(a) compares delay distributions under nominal and severe sensing noise. Nominal medians are 10.9, 9.8, 7.9, 11.6, and 8.7 s for TTC, APF, D-MPC, F-SMPC, and AR-SMPC; under severe noise, the proposed median increases only to 9.6 s, while F-SMPC reaches 13.4 s because every enlarged tube receives the same conservative share. Figure 6(b) maps 48,000 optimization cycles and concentrates between six and nine active constraints and $16-27$ ms, with only 0.7% of cycles beyond the 45 ms warning line. Figure 6(c) reports median/95th-percentile times of $3.2/6.7$ ms for TTC, $7.6/14.1$ ms for APF, $17.8/33.9$ ms for D-MPC, $19.6/36.8$ ms for F-SMPC, and $21.4/38.6$ ms for AR-SMPC. Decentralized mixed-traffic control likewise benefits from local constraint sets rather than intersection-wide optimization [31].

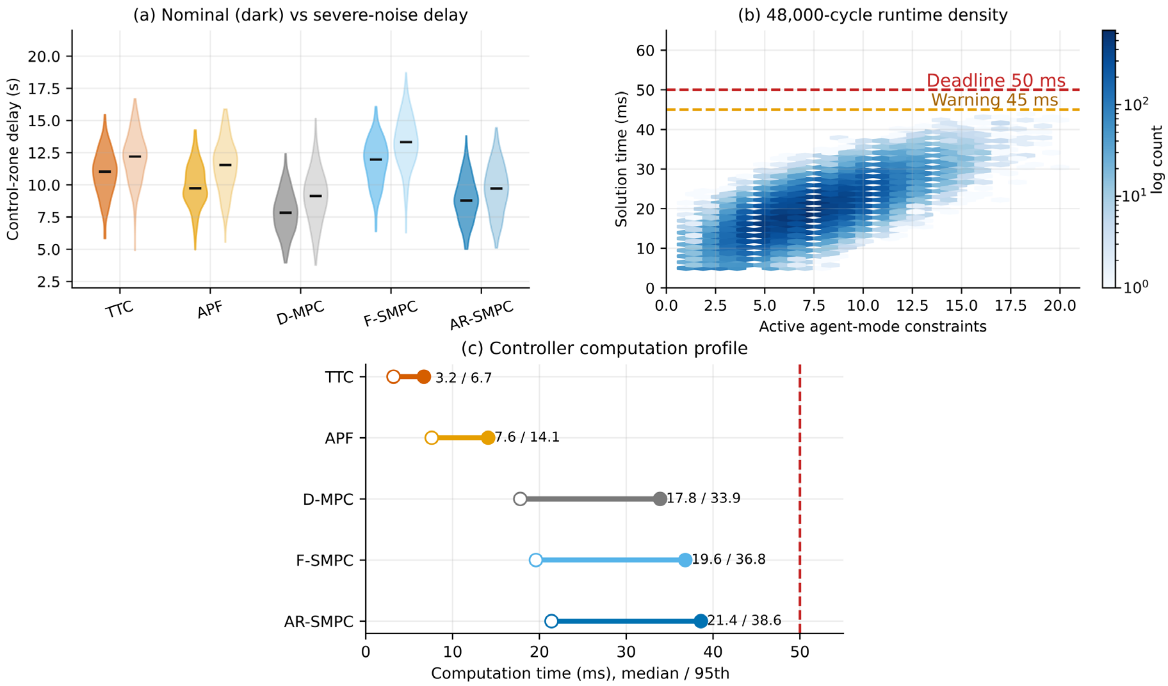

Figure 6. Runtime robustness: (a) delay distributions under nominal and severe sensing noise; (b) active-constraint versus solution-time density; (c) median and 95th-percentile computation times.

No deadline misses directly caused a collision. Of 336 cycles above 50 ms, 302 reused a verified prefix and 34 selected emergency braking; all maintained more than 0.62 m clearance. The 99th-percentile time was 46.9 ms, and peak memory was 184 MB. Increasing the horizon from 30 to 40 steps reduced collision rate by only 0.07 percentage points but raised the deadline-miss fraction from 0.18% to 2.64%. Graph-based conflict-free cooperation can obtain global schedules when participation is complete [32], whereas the measured local formulation retains bounded compute cost in the presence of unconnected participants. Decentralized optimalcontrol results also support the use of comfortconstrained turns and explicit safety guarantees [33]. The present implementation adds calibrated perception uncertainty and a runtime supervisor to those principles.

Figure 7 (a) decomposes avoided high-risk events: 38% are attributable to multimodal occupancy, 27% to adaptive allocation, 21% to the terminal set, and 14% to freshness inflation, with different shares across left-turn, crossingpedestrian, cyclist-overtake, and red-light-violation scenarios. Figure 7(b) begins at the 4.8% D-MPC collision rate and subtracts 1.46 percentage points for multimodal prediction, 1.31 for adaptive allocation, 0.88 for the terminal set, and 0.55 for freshness handling, ending at 0.60%. Figure 7 (c) compares predicted risk with observed frequency in ten bins: expected calibration error decreases from 0.021 for uncalibrated tubes to 0.006 after covariance scaling, and all upper confidence bars remain below the diagonal safety envelope for predicted risks up to 0.08.

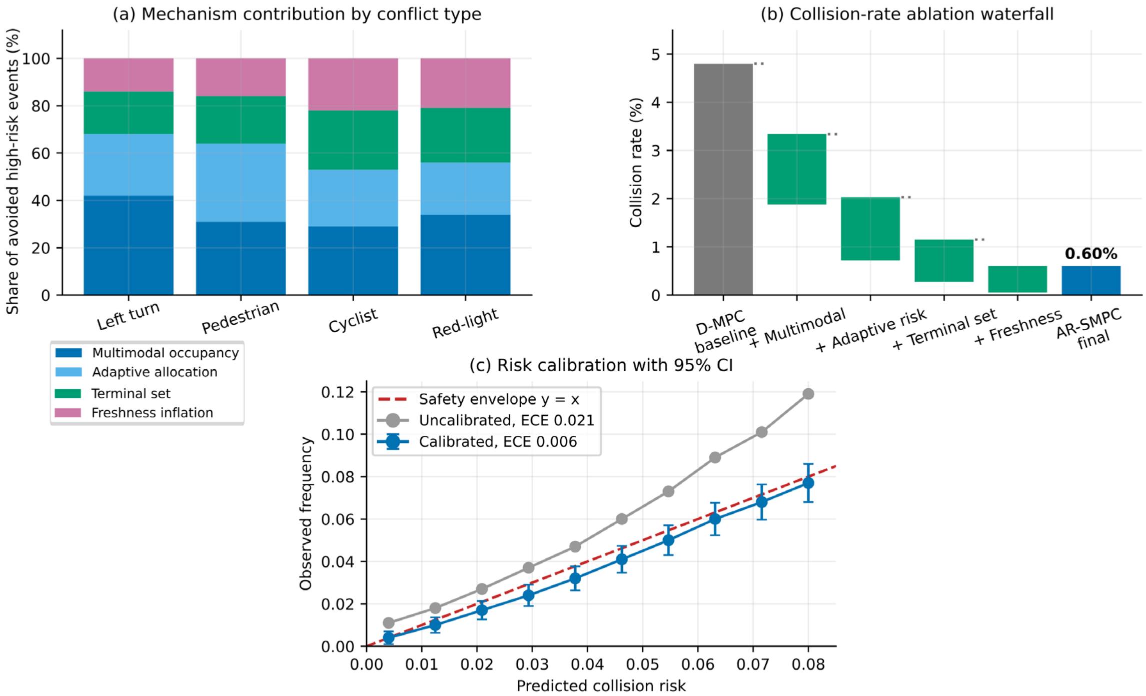

Figure 7. Mechanism and calibration analysis: (a) stacked avoided-event contributions; (b) collision-rate ablation waterfall; (c) predicted-risk calibration with confidence bars.

The ablation results identify complementary functions rather than one dominant safeguard. Removing adaptive allocation raises collision rate to 1.4% and mean delay to 11.6 s, while retaining only the most probable motion mode increases collision rate to 2.06%. Observation-age inflation matters mainly under degraded sensing, and the terminal condition prevents late infeasibility near the stop line. Remaining errors are concentrated in abrupt cyclist motion, severe occlusion, and combined low-friction maneuvers, which are not fully represented by Gaussian occupancy modes and the low-speed kinematic model.

Under 10% packet loss, AR-SMPC records 0.77% collisions; under the combined severe-noise, packet-loss, and low-friction condition, the rate reaches 1.86% but remains below both predictive baselines. The risk-budget sweep confirms the selected value of 0.015: tighter allocation lowers collisions but creates excessive delay, whereas looser allocation sacrifices safety. These results support deployment as a vehicle-level safety controller with locally calibrated occupancy tubes and an independently verified terminal supervisor.

Conclusion

Develop an interaction-aware stochastic model predictive control algorithm for collision avoidance in crowded intersections in this study. A method that maps multimodal predictions into heterogeneous occupancy tubes, allocates a limited risk budget according to urgency and vulnerability, and optimises vehicle progress, tracking, comfort, and probabilistic separation under a terminal braking-safe condition. The combined Design keeps the understandable link between perceived conflict and control margin and is still suitable for a 20Hz automotive control cycle.

The current form is restricted to calibrated Gaussian mode components, a low-speed kinematic vehicle model, and simulation-based verification. It also assumes that perception provides a consistent track identity during short occlusions. Abrupt changes in cyclist movement, adversarial connection of messages, irregular road shapes, and low-friction combined operations may all be considered violations. The fallback supervisor reduces the immediate risk but does not solve the problem of wider system-level assurance.

Future work will combine distribution-free prediction sets with adaptive risk allocation, incorporate dynamic vehicle constraints for low-friction operation, and validate the controller through hardware-in-the-loop and instrumented-intersection experiments. A cooperative extension can have the nearby automated vehicles share risk-price information without requiring consensus for safety. Formal Runtime Monitors and scenario-guided field calibrations will also link probabilistic performance with certification evidence.

Author Contributions

Akira Takahashi contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. Sakura Yamamoto contributes to conceptualization, validation, analysis. 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] Yurtsever, E., Lambert, J., Carballo, A., & Takeda, K. (2020). A Survey of Autonomous Driving: Common Practices and Emerging Technologies. IEEE Access, 8, 58443-58469. https://doi.org/10.1109/access.2020.2983149

[2] Muhammad, K., Ullah, A., Lloret, J., Ser, J. D., & Albuquerque, V. H. C. D. (2020). Deep Learning for Safe Autonomous Driving: Current Challenges and Future Directions. IEEE Transactions on Intelligent Transportation Systems, 22(7), 4316-4336. https://doi.org/10.1109/tits.2020.3032227

[3] Riedmaier, S., Ponn, T., Ludwig, D., Schick, B., & Diermeyer, F. (2020). Survey on Scenario-Based Safety Assessment of Automated Vehicles. IEEE Access, 8, 87456-87477. https://doi.org/10.1109/access.2020.2993730

[4] Guo, Q., Li, L., & Ban, X. (2019). Urban traffic signal control with connected and automated vehicles: A survey. Transportation Research Part C Emerging Technologies, 101, 313-334. https://doi.org/10.1016/j.trc.2019.01.026

[5] Elliott, D. L., Keen, W., & Miao, L. (2019). Recent advances in connected and automated vehicles. Journal of Traffic and Transportation Engineering (English Edition), 6(2), 109-131. https://doi.org/10.1016/j.jtte.2018.09.005

[6] Wang, J., Zhang, L., Huang, Y., Zhao, J., & Bella, F. (2020). Safety of Autonomous Vehicles. Journal of Advanced Transportation, 2020, 1-13. https://doi.org/10.1155/2020/8867757

[7] Namazi, E., Li, J., & Lu, C. (2019). Intelligent Intersection Management Systems Considering Autonomous Vehicles: A Systematic Literature Review. IEEE Access, 7, 91946-91965. https://doi.org/10.1109/access.2019.2927412

[8] Khayatian, M., Mehrabian, M., Andert, E., Dedinsky, R., Choudhary, S., Lou, Y., & Shirvastava, A. (2020). A Survey on Intersection Management of Connected Autonomous Vehicles. ACM Transactions on Cyber-Physical Systems, 4(4), 1-27. https://doi.org/10.1145/3407903

[9] Feng, S., Feng, Y., Yu, C., Zhang, Y., & Liu, H. (2020). Testing Scenario Library Generation for Connected and Automated Vehicles, Part I: Methodology. IEEE Transactions on Intelligent Transportation Systems, 22(3), 1573-1582. https://doi.org/10.1109/tits.2020.2972211

[10] Pek, C., Manzinger, S., Koschi, M., & Althoff, M. (2020). Using online verification to prevent autonomous vehicles from causing accidents. Nature Machine Intelligence, 2(9), 518-528. https://doi.org/10.1038/s42256-020-0225-y

[11] Noor-A-Rahim, M., Liu, Z., Lee, H., Khyam, M. O., He, J., Pesch, D., Moessner, K., Saad, W., & Poor, H. V. (2022). 6G for Vehicle-to-Everything (V2X) Communications: Enabling Technologies, Challenges, and Opportunities. Proceedings of the IEEE, 110(6), 712-734. https://doi.org/10.1109/jproc.2022.3173031

[12] Chen, L., Li, Y., Huang, C., Li, B., Xing, Y., Tian, D., Li, L., Hu, Z., Na, X., Li, Z., Teng, S., Lv, C., Wang, J., Cao, D., Zheng, N., & Wang, F. (2022). Milestones in Autonomous Driving and Intelligent Vehicles: Survey of Surveys. IEEE Transactions on Intelligent Vehicles, 8(2), 1046-1056. https://doi.org/10.1109/tiv.2022.3223131

[13] Ma, Y., Zhu, X., Zhang, S., Yang, R., Wang, W., & Manocha, D. (2019). TrafficPredict: Trajectory Prediction for Heterogeneous Traffic-Agents. Proceedings of the AAAI Conference on Artificial Intelligence, 33(01), 6120-6127. https://doi.org/10.1609/aaai.v33i01.33016120

[14] Dai, S., Li, L., & Li, Z. (2019). Modeling Vehicle Interactions via Modified LSTM Models for Trajectory Prediction. IEEE Access, 7, 38287-38296. https://doi.org/10.1109/access.2019.2907000

[15] Zyner, A., Worrall, S., & Nebot, E. (2019). Naturalistic Driver Intention and Path Prediction Using Recurrent Neural Networks. IEEE Transactions on Intelligent Transportation Systems, 21(4), 1584-1594. https://doi.org/10.1109/tits.2019.2913166

[16] Lefkopoulos, V., Menner, M., Domahidi, A., & Zeilinger, M. N. (2020). Interaction-Aware Motion Prediction for Autonomous Driving: A Multiple Model Kalman Filtering Scheme. IEEE Robotics and Automation Letters, 6(1), 80-87. https://doi.org/10.1109/lra.2020.3032079

[17] Jeong, Y., Kim, S., & Yi, K. (2020). Surround Vehicle Motion Prediction Using LSTM-RNN for Motion Planning of Autonomous Vehicles at Multi-Lane Turn Intersections. IEEE Open Journal of Intelligent Transportation Systems, 1, 2-14. https://doi.org/10.1109/ojits.2020.2965969

[18] Gulzar, M., Muhammad, Y., & Muhammad, N. (2021). A Survey on Motion Prediction of Pedestrians and Vehicles for Autonomous Driving. IEEE Access, 9, 137957-137969. https://doi.org/10.1109/access.2021.3118224

[19] Wang, H., Liu, B., Ping, X., & An, Q. (2019). Path Tracking Control for Autonomous Vehicles Based on an Improved MPC. IEEE Access, 7, 161064-161073. https://doi.org/10.1109/access.2019.2944894

[20] Cheng, S., Li, L., Guo, H., Chen, Z., & Song, P. (2019). Longitudinal Collision Avoidance and Lateral Stability Adaptive Control System Based on MPC of Autonomous Vehicles. IEEE Transactions on Intelligent Transportation Systems, 21(6), 2376-2385. https://doi.org/10.1109/tits.2019.2918176

[21] Wang, Y., Liu, Z., Zuo, Z., Li, Z., Wang, L., & Luo, X. (2019). Trajectory Planning and Safety Assessment of Autonomous Vehicles Based on Motion Prediction and Model Predictive Control. IEEE Transactions on Vehicular Technology, 68(9), 8546-8556. https://doi.org/10.1109/tvt.2019.2930684

[22] Li, S., Li, Z., Yu, Z., Zhang, B., & Zhang, N. (2019). Dynamic Trajectory Planning and Tracking for Autonomous Vehicle With Obstacle Avoidance Based on Model Predictive Control. IEEE Access, 7, 132074-132086. https://doi.org/10.1109/access.2019.2940758

[23] 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

[24] Mirheli, A., Tajalli, M., Hajibabai, L., & Hajbabaie, A. (2019). A consensus-based distributed trajectory control in a signal-free intersection. Transportation Research Part C Emerging Technologies, 100, 161-176. https://doi.org/10.1016/j.trc.2019.01.004

[25] Chen, C., Wang, J., Xu, Q., Wang, J., & Li, K. (2021). Mixed platoon control of automated and human-driven vehicles at a signalized intersection: Dynamical analysis and optimal control. Transportation Research Part C Emerging Technologies, 127, 103138. https://doi.org/10.1016/j.trc.2021.103138

[26] Sun, C., Guanetti, J., Borrelli, F., & Moura, S. (2020). Optimal Eco-Driving Control of Connected and Autonomous Vehicles Through Signalized Intersections. IEEE Internet of Things Journal, 7(5), 3759-3773. https://doi.org/10.1109/jiot.2020.2968120

[27] Guillén, A., & Cano, M. (2022). Multi-Agent Deep Reinforcement Learning to Manage Connected Autonomous Vehicles at Tomorrow's Intersections. IEEE Transactions on Vehicular Technology, 71(7), 7033-7043. https://doi.org/10.1109/tvt.2022.3169907

[28] Arvin, R., Khattak, A. J., Kamrani, M., & Torres, J. R. (2020). Safety evaluation of connected and automated vehicles in mixed traffic with conventional vehicles at intersections. Journal of Intelligent Transportation Systems, 25(2), 170-187. https://doi.org/10.1080/15472450.2020.1834392

[29] Ma, C., Yu, C., & Yang, X. (2021). Trajectory planning for connected and automated vehicles at isolated signalized intersections under mixed traffic environment. Transportation Research Part C Emerging Technologies, 130, 103309. https://doi.org/10.1016/j.trc.2021.103309
[30] Zhao, J., Knoop, V. L., & Wang, M. (2020). Two-dimensional vehicular movement modelling at intersections based on optimal control. Transportation Research Part B Methodological, 138, 1-22. https://doi.org/10.1016/j.trb.2020.04.001
[31] Yao, H., & Li, X. (2020). Decentralized control of connected automated vehicle trajectories in mixed traffic at an isolated signalized intersection. Transportation Research Part C Emerging Technologies, 121, 102846. https://doi.org/10.1016/j.trc.2020.102846
[32] Chen, C., Xu, Q., Cai, M., Wang, J., Wang, J., & Li, K. (2022). Conflict-Free Cooperation Method for Connected and Automated Vehicles at Unsignalized Intersections: Graph-Based Modeling and Optimality Analysis. IEEE Transactions on Intelligent Transportation Systems, 23(11), 21897-21914. https://doi.org/10.1109/tits.2022.3182403
[33] Zhang, Y., & Cassandras, C. G. (2019). Decentralized optimal control of Connected Automated Vehicles at signal-free intersections including comfort-constrained turns and safety guarantees. Automatica, 109, 108563. https://doi.org/10.1016/j.automatica.2019.108563