Probabilistic Game-theoretic Control for Autonomous Vehicles on Lane-free Roads

Zhuopeng Xie*[1], Mohsen Ramezani[2], and David Levinson[3]

[1]PhD student, School of Civil Engineering, University of Sydney, Australia
[2]Associate professor, School of Civil Engineering, University of Sydney, Australia
[3]Professor, School of Civil Engineering, University of Sydney, Australia

Short summary

We propose a probabilistic game-theoretic controller (PGC) for autonomous vehicles on lane-free roads in decentralized and communication-free environments. At each time step, each vehicle solves a nonlinear Model Predictive Control problem and predicts the behavior of nearby vehicles through multi-level game-theoretic reasoning. We represent interaction uncertainty using structured trajectory hypotheses, and we incorporate collision risk through a Conditional Value-at-Risk-based cost term and iterative risk allocation constraints. We then compare the proposed PGC with the deterministic counterpart. The results show that the PGC handles uncertainty more effectively. In particular, it maintains safety while reducing unnecessary hesitation, avoiding overly conservative maneuvers, improving overtaking behavior, and increasing speeds along track. This study establishes a novel framework for probabilistic lane-free autonomous vehicle control under uncertainty.
Keywords: automated driving, game theory, lane-free traffic, model predictive control, probabilistic control, traffic simulation

1 Introduction

Lane-free roads allow autonomous vehicles (AVs) to exploit the full road width rather than remain constrained to fixed lane centers, thereby creating greater path-planning flexibility on existing road infrastructure and offering the potential to reduce travel time, reduce energy consumption, and increase road capacity Papageorgiou et al. (2021); Levy & Haddad (2022); Beza et al. (2025). However, this additional flexibility also makes microscopic control for AVs substantially more challenging, because each vehicle must coordinate longitudinal and lateral motions in a two-dimensional space while satisfying road-geometry, vehicle-dynamics, and inter-vehicle safety constraints in real time.

Existing microscopic lane-free control studies have explored several methodological directions. Among them, optimization-based approaches have been studied most extensively Yanumula et al. (2021, 2023); Levy & Haddad (2021, 2022); Naderi et al. (2025). Other studies have investigated flocking-based coordination, artificial-force-based driving strategies, layered behavior models, and deep reinforcement learning for lane-free autonomous driving Rostami-Shahrbabaki et al. (2023); Berahman et al. (2022); Karalakou et al. (2023). These methods have been applied to a range of lane-free traffic scenarios, including road segments, intersections, and roundabouts Yanumula et al. (2021); Levy & Haddad (2022); Naderi et al. (2025, 2023).

Despite this progress, two gaps remain in the current lane-free AV control literature. First, most studies have been developed for connected or centralized settings Levy & Haddad (2022). From a deployment perspective, however, decentralized and communication-free control is also important, because real-world lane-free traffic is unlikely to consist entirely of connected vehicles in the foreseen future, and communication-free control reduces reliance on external communication and supporting infrastructure. Second, most existing lane-free control studies remain deterministic in their representations of surrounding-vehicle motion and safety constraints Yanumula et al. (2021); Levy & Haddad (2022); Karalakou et al. (2023). By contrast, explicit probabilistic treatment of interaction uncertainty has been studied more extensively in the broader autonomous-driving literature than in lane-free AV control Tang et al. (2022); Jeong (2021); Bonzanini et al. (2024); Yin et al. (2023).

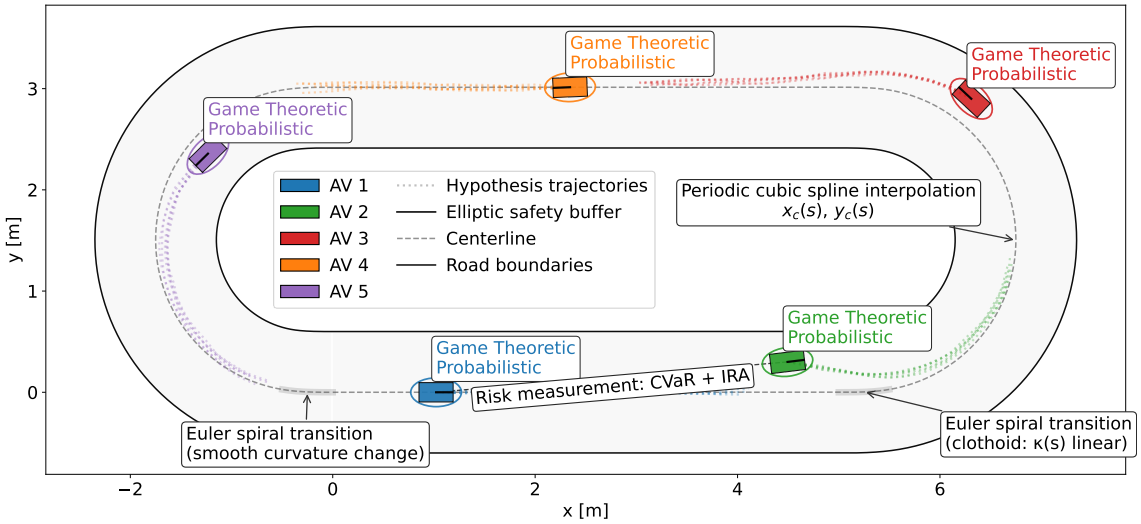
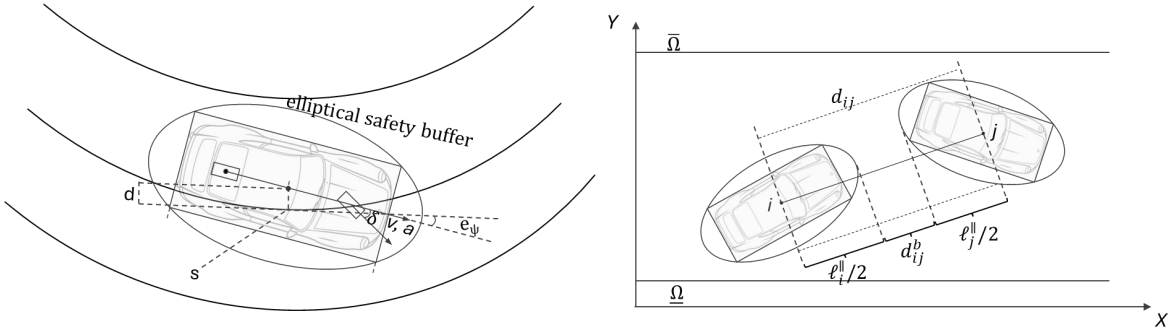


Figure 1: Method overview of the proposed PGC algorithm.

To address these gaps, this paper proposes a probabilistic game-theoretic controller (PGC) for AVs on lane-free roads. In our implementation of the proposed framework, each vehicle solves a nonlinear Model Predictive Control (NMPC) problem and predicts the behavior of nearby vehicles through multi-level game-theoretic reasoning. To account for interaction uncertainty, we construct structured trajectory hypotheses for surrounding vehicles and incorporate safety both in the cost function and in the constraints through probabilistic formulations.

2 Methodology

Figure 1 shows the overall framework of the proposed PGC algorithm. The considered road geometry consists of two straight segments and two semicircular curves forming a closed-loop track, and can potentially be extended to roads with more general curved shapes.



(a) (b)

Figure 2: Illustrations of the Frenet frame and static clearance: (a) bicycle model in Frenet coordinate system, (b) static clearance.

Vehicle kinematic model

Each vehicle is modeled using a bicycle model in the Frenet frame, and its state and control input are given by (1) and (2), respectively. To avoid curvature discontinuities and reduce the complexity, Euler spiral transitions and cubic spline interpolation are applied. The resulting geometry is precomputed to enable efficient conversion from the Frenet coordinates (s, d) to the global Cartesian coordinates (x, y) when necessary.

$$\boldsymbol{\xi}(t) = [s(t),\, d(t),\, e_\psi(t),\, v(t)]^\mathsf{T} \tag{1}$$

$$\mathbf{u}(t) = [a(t),\ \delta(t)]^{\mathsf{T}} \tag{2}$$

where $s(t)$ denotes the longitudinal progress of the robot along the track centerline, $d(t)$ is the lateral offset from the centerline, $e_\psi(t)$ is the heading error with respect to the tangent direction of the centerline, $v(t)$ is the linear speed, $a(t)$ is the longitudinal acceleration, and $\delta(t)$ is the front steering angle.
The Frenet frame is adopted because we extend the straight roads to curved roads. In this case, using the distance along the centerline and the lateral deviation from it is more natural than describing the motion directly in the global Cartesian frame.
The heading error is defined as the local tangent direction of the reference path shown in (3).

$$e_\psi(t) = \theta(t) - \psi_{\text{ref}}\big(s(t)\big) \tag{3}$$

where $\theta(t)$ is the heading angle of the vehicle body, $\psi_{\text{ref}}(s)$ is the tangent angle of the centerline at arc-length position s, and $\kappa(s)$ denotes the curvature of the centerline.
The dynamic update $\boldsymbol{\xi}(t + \Delta t_p) = f(\boldsymbol{\xi}(t), \mathbf{u}(t))$ unfolds component-wise as in (4)–(7).

$$s(t + \Delta t_p) = s(t) + \frac{v(t)\cos e_\psi(t)}{1 - \kappa\big(s(t)\big)d(t)}\,\Delta t_p \tag{4}$$

$$d(t + \Delta t_p) = d(t) + v(t)\sin e_\psi(t)\,\Delta t_p \tag{5}$$

$$e_\psi(t + \Delta t_p) = e_\psi(t) + \left[\frac{v(t)}{l_B}\tan\big(\delta(t)\big) - \kappa\big(s(t)\big)\frac{v(t)\cos e_\psi(t)}{1 - \kappa\big(s(t)\big)d(t)}\right]\Delta t_p \tag{6}$$

$$v(t + \Delta t_p) = v(t) + a(t)\,\Delta t_p \tag{7}$$

where $l_B = l_f + l_r$ is the wheelbase, l_f and l_r are the distance from the front axle and rear axle to the vehicle center, respectively, Δt_p is the prediction interval, and $f(\cdot)$ represents the discrete-time vehicle dynamics.

Geometric modeling of spatial safety

We approximate each vehicle's footprint using an ellipse that encloses its rectangular body. The ellipse passes through the four corners $(\pm l/2, \pm w/2)$ in the vehicle's local body frame, forming a tight safety boundary, as shown in Figure 2 (a) and (8).

$$\frac{2x^2}{l^2} + \frac{2y^2}{w^2} = 1 \tag{8}$$

where l and w denote the length and width of the vehicle.

Road boundary constraint. The road boundary constraint is imposed directly on the lateral offset d. The feasible range of the vehicle lateral position is given by (9). This constraint ensures that the vehicle keeps a minimum clearance from both road edges while moving along the curved track.

$$-\frac{r_w}{2} + r_s \le d \le \frac{r_w}{2} - r_s \tag{9}$$

where r_w is the road width, and r_s is the prescribed safety margin to the boundary.

Centroid-based spatial clearance. To estimate whether two vehicles are too close to each other, we define the static clearance d_{ij}^b as the centroid distance d_{ij} minus the projected safety buffers of both vehicles along the connecting direction, as shown in Figure 2 (b). The bearing from vehicle i to j is denoted by ψ, and the centroid distance is given in (11). The projected buffers are computed by (12), and the resulting static clearance is defined in (13).

$$\psi = \tan^{-1}\left(\frac{y_j - y_i}{x_j - x_i}\right) \tag{10}$$

$$d_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \tag{11}$$

$$\ell_k^{\parallel}(\psi) = 2\sqrt{\frac{l_k^2}{2}\sin^2(\psi - \theta_k) + \frac{w_k^2}{2}\cos^2(\psi - \theta_k)} \tag{12}$$

$$d_{ij}^b = d_{ij} - \frac{\ell_i^{\parallel}(\psi)}{2} - \frac{\ell_j^{\parallel}(\psi)}{2} \tag{13}$$

where x_i, y_i, x_j, y_j are centroid coordinates, l_k and w_k are the vehicle's length and width, and θ_k is the heading of vehicle $k \in \{i, j\}$.

Structured sampling. At each control step, the ego vehicle anticipates the motion of the closest K neighbors, and each neighbor is represented by M hypothesis trajectories over the prediction horizon. For neighbor j, the hypothesis set is denoted by (14), where $\boldsymbol{\xi}_j^{(m)}(h)$ is the predicted state of the mth hypothesis at prediction step h.

$$\mathcal{X}_j = \left\{\boldsymbol{\xi}_j^{(m)}(0), \boldsymbol{\xi}_j^{(m)}(1), \ldots, \boldsymbol{\xi}_j^{(m)}(H)\right\}_{m=1}^{M}, \quad j = 1, \ldots, K \tag{14}$$

To generate the hypothesis trajectories, random perturbations are added to the nominal acceleration and steering input of each neighbor, as in (15)–(18).

$$a_j^{(m)}(h) = \bar{a}_j(h) + \eta_a^{(m)}(h) \tag{15}$$

$$\delta_j^{(m)}(h) = \bar{\delta}_j(h) + \eta_\delta^{(m)}(h) \tag{16}$$

$$\eta_a^{(m)}(h) \sim \mathcal{N}(0, \sigma_a^2) \tag{17}$$

$$\eta_\delta^{(m)}(h) \sim \mathcal{N}(0, \sigma_\delta^2) \tag{18}$$

Each hypothesis trajectory is then propagated using the Frenet-frame vehicle dynamics in (19).

$$\boldsymbol{\xi}_j^{(m)}(h+1) = f\left(\boldsymbol{\xi}_j^{(m)}(h),\, \mathbf{u}_j^{(m)}(h)\right) \tag{19}$$

where $\bar{a}_j(h)$ and $\bar{\delta}_j(h)$ are the nominal acceleration and steering input of neighbor j at prediction step h, respectively, $\eta_a^{(m)}(h)$ and $\eta_\delta^{(m)}(h)$ are zero-mean Gaussian perturbations, and $\mathbf{u}_j^{(m)}(h)$ denotes the perturbed control input used to propagate the mth hypothesis trajectory.

Conditional Value-at-Risk (CVaR). To account for collision risk under motion uncertainty, we evaluate the static clearance between the ego vehicle and each anticipated neighbor over all hypothesised trajectories. The minimum clearance over the horizon is approximated using the soft minimum operator as in (20). This formulation captures the most critical interaction within each hypothesis trajectory while remaining smooth for optimization. The risk associated with neighbor j is measured using the conditional value-at-risk (CVaR), which is defined as the average of the worst $(1-\alpha)$ fraction of the hypothesis outcomes. The CVaR term is given by (22).

$$d_{j,\min}^b(m) = \operatorname{softmin}_{0 \le h \le H} d_j^b(m, h) \tag{20}$$

$$q_j(m) = \exp\left(-d_{j,\min}^b(m)\right) \tag{21}$$

$$\mathrm{CVaR}_{\alpha,j}(q_j) = \frac{1}{|\mathcal{T}_{\alpha,j}|} \sum_{m \in \mathcal{T}_{\alpha,j}} q_j(m) \tag{22}$$

where $d_j^b(m, h)$ is the predicted static clearance between the ego vehicle and neighbor j at prediction step h under hypothesis m, $d_{j,\min}^b(m)$ is the minimum clearance over the horizon, $q_j(m)$ is the corresponding exponential risk value of the mth hypothesis, $\mathcal{T}_{\alpha,j}$ is the index set corresponding to the largest values of $q_j(m)$, and α is the CVaR confidence level.

Iterative risk allocation (IRA). In addition to the CVaR-based risk term, we impose a joint risk constraint using iterative risk allocation (IRA). For each neighbor j, the violation probability is estimated from the M hypothesised trajectories by checking whether the minimum predicted static clearance over the horizon is smaller than a safety threshold R, as in (23).

$$\hat{p}_j = \frac{1}{M} \sum_{m=1}^{M} \mathbf{1} \left\{ \min_{0 \le h \le H} d_j^b(m,h) < R \right\} \tag{23}$$

The risk allocation assigns an individual risk budget ϵ_j to each neighbor while preserving the total joint risk bound, as in (24) and (25).

$$\hat{p}_j \le \epsilon_j, \quad j = 1, \dots, K \tag{24}$$

$$\sum_{j=1}^{K} \epsilon_j = \epsilon, \qquad \epsilon_j \ge 0 \tag{25}$$

Initially, all neighbors are assigned the same risk budget, as in (26). After each iteration, the budgets are redistributed so that more risky neighbors receive larger budgets, as in (27).

$$\epsilon_1^{(0)} = \epsilon_2^{(0)} = \cdots = \epsilon_K^{(0)} = \frac{\epsilon}{K} \tag{26}$$

$$\epsilon_j^{(t+1)} = \frac{\exp\big(\rho \hat{p}_j^{(t)}\big)}{\sum_{k=1}^{K} \exp\big(\rho \hat{p}_k^{(t)}\big)} \, \epsilon, \qquad j = 1, \dots, K \tag{27}$$

where $d_j^b(m,h)$ is the predicted static clearance between the ego vehicle and neighbor j at prediction step h under hypothesis m, $\hat{p}_j$ is the empirical violation probability of neighbor j, ϵ is the total risk budget, ϵ_j is the allocated budget for neighbor j, ρ is the allocation parameter, and R is the safety threshold for the static clearance.
The above allocation approximates the joint chance constraint by limiting the total probability of unsafe events across all selected neighbors. This allows the controller to distribute the available risk budget adaptively according to the predicted interaction risk. As a result, a larger portion of the collision probability budget can be allocated to more critical interactions, which helps reduce overly conservative driving maneuvers.

Finite-horizon NMPC

At each control step, the ego vehicle solves a finite-horizon NMPC problem over a prediction horizon of N steps. The NMPC seeks an optimal control sequence $\mathbf{U} = \{\mathbf{u}(t_n)\}_{n=0}^{N-1}$ that minimizes a cumulative cost while respecting the vehicle dynamics and safety constraints. The stage cost consists of a progress-tracking term, control-effort penalties, a heading-error penalty, and a CVaR-based interaction risk term. In addition, IRA-based joint risk constraints are imposed to limit the total empirical violation probability across all anticipated neighbors. The formulation is given in (28)–(30).

$$\min_{\mathbf{U}} \; J(\boldsymbol{\Xi}, \mathbf{U}) = \sum_{n=0}^{N-1} L\big(\boldsymbol{\xi}(t_n),\, \mathbf{u}(t_n)\big) \tag{28}$$

$$L(\boldsymbol{\xi}, \mathbf{u}) = \underbrace{w_s \left(\frac{\dot{s} - \dot{s}_{\text{des}}}{\dot{s}_{\text{des}}} \right)^2}_{\text{progress tracking}} + \underbrace{w_a \left(\frac{a}{a_{\max}} \right)^2}_{\text{acceleration effort}} + \underbrace{w_\delta \left(\frac{\delta}{\delta_{\max}} \right)^2}_{\text{steering effort}} + \underbrace{w_\psi \left(\frac{e_\psi}{e_{\psi,\max}} \right)^2}_{\text{heading error penalty}} + \underbrace{\sum_{j \in \mathcal{J}^-} w_r \, \text{CVaR}_{\alpha,j}(q_j)}_{\text{interaction risk penalty}} \tag{29}$$

$$
\text{s.t.}\quad
\begin{cases}
\boldsymbol{\xi}(t_{n+1}) = f\big(\boldsymbol{\xi}(t_n),\, \mathbf{u}(t_n),\, \Delta t_p\big), & \text{vehicle dynamics},\\
a_{\min} \le a(t_n) \le a_{\max}, & \text{acceleration bound},\\
\delta_{\min} \le \delta(t_n) \le \delta_{\max}, & \text{steering bound},\\
e_{\psi,\min} \le e_\psi(t_n) \le e_{\psi,\max}, & \text{heading-error bound},\\
-\dfrac{r_w}{2} + r_s \le d(t_n) \le \dfrac{r_w}{2} - r_s, & \text{road-boundary constraint},\\
\hat{p}_j \le \epsilon_j, \quad \forall j \in \mathcal{J}^-, & \text{individual risk constraint},\\
\sum_{j\in\mathcal{J}^-} \epsilon_j = \epsilon, \quad \epsilon_j \ge 0 & \text{IRA budget constraint}
\end{cases}
\tag{30}
$$

where $\boldsymbol{\Xi} = \{\boldsymbol{\xi}(t_n)\}_{n=1}^{N}$ denotes the predicted state sequence, and $n = 0, \ldots, N-1$ is the prediction-step index. The term $\dot{s}$ is the longitudinal progress speed along the centerline, and $\dot{s}_{\text{des}}$ is its desired value. The set $\mathcal{J}^-$ denotes the set of selected neighboring vehicles. The parameters w_s, w_a, w_δ, w_ψ, and w_r are weighting coefficients, while ϵ is the total collision-probability budget.
At each control step, only the first control input of the optimized sequence is applied to the ego vehicle.

Game-theoretic interaction

To model local interactions in an unconnected and decentralized environment, a level-k game-theoretic reasoning scheme is adopted. At each control step, the ego vehicle anticipates the motion of the selected neighboring vehicles $j \in \mathcal{J}^-$ and constructs their future trajectories recursively. For neighbor j, the level-k hypothesis trajectory set is denoted by (31), where $\boldsymbol{\xi}_j^{(k,m)}(h)$ is the predicted state of the mth hypothesis at prediction step h under reasoning level k.

$$
\mathcal{X}_j^{(k)} = \Big\{\boldsymbol{\xi}_j^{(k,m)}(0),\, \boldsymbol{\xi}_j^{(k,m)}(1),\, \ldots,\, \boldsymbol{\xi}_j^{(k,m)}(H)\Big\}_{m=1}^{M}, \quad j \in \mathcal{J}^- \tag{31}
$$

For level-1 reasoning, each neighbor is assumed to move at constant speed and gradually return to the centerline. Therefore, the nominal longitudinal acceleration is set to zero, while the nominal steering input consists of a feedforward term that follows the road curvature and a feedback term that drives the vehicle back toward the centerline. Random perturbations are then added to both inputs, as in (32)–(35).

$$
a_j^{(1,m)}(h) = 0 + \eta_a^{(m)}(h) \tag{32}
$$

$$
\delta_j^{(1,m)}(h) = \delta_{\text{ff},j}(h) + \delta_{\text{fb},j}(h) + \eta_\delta^{(m)}(h) \tag{33}
$$

$$
\delta_{\text{ff},j}(h) = \arctan\big(l_B \kappa(s_j(h))\big) \tag{34}
$$

$$
\delta_{\text{fb},j}(h) = -k_d d_j(h) - k_\psi e_{\psi,j}(h) \tag{35}
$$

For higher-level reasoning, the motion of each neighbor is generated by perturbing its level-$(k-1)$ motion, as in (36) and (37). In this way, the level-k prediction is constructed recursively from the previous reasoning layer.

$$
a_j^{(k,m)}(h) = a_j^{(k-1)}(h) + \eta_a^{(m)}(h), \qquad k \ge 2 \tag{36}
$$

$$
\delta_j^{(k,m)}(h) = \delta_j^{(k-1)}(h) + \eta_\delta^{(m)}(h), \qquad k \ge 2 \tag{37}
$$

Each hypothesis trajectory is then propagated using the Frenet-frame vehicle dynamics, as in (38). Based on the resulting level-$(k-1)$ neighbor predictions, the ego vehicle solves its own NMPC problem to obtain a level-k best response.

$$
\boldsymbol{\xi}_j^{(k,m)}(h+1) = f\Big(\boldsymbol{\xi}_j^{(k,m)}(h),\, \mathbf{u}_j^{(k,m)}(h)\Big) \tag{38}
$$

$$
\mathbf{U}^{(k)} = \text{BR}\Big(\big\{\mathcal{X}_j^{(k-1)}\big\}_{j\in\mathcal{J}^-}\Big) \tag{39}
$$

where $\mathbf{u}_j^{(k,m)}(h) = [a_j^{(k,m)}(h),\, \delta_j^{(k,m)}(h)]^\mathsf{T}$ is the control input of the mth hypothesis of neighbor j at reasoning level k, $\kappa(s)$ is the road curvature, k_d and k_ψ are feedback gains, and $\text{BR}(\cdot)$ denotes the best-response operator implemented by the NMPC solver.

Table 1: Main parameters used in the proposed PGC algorithm.

Symbol	Meaning	Unit	Value	Symbol	Meaning	Unit	Value
Track and vehicle geometry							
L_s	Length of each straight segment	m	5.0	R_{trk}	Radius of each semicircular curve	m	1.5
l	Vehicle length	m	0.3375	l_B	Wheelbase	m	0.235
l_f	Distance from front axle to vehicle center	m	0.1175	l_r	Distance from rear axle to vehicle center	m	0.1175
r_s	Safety margin to road boundary	m	0.20	r_w	Road width	m	1.2
w	Vehicle width	m	0.1911				
Control and prediction settings							
H	Game-theoretic prediction horizon	steps	12	K	Number of selected neighbors		4
M	Number of hypotheses per neighbor		8	N	NMPC prediction horizon	steps	12
k	Game-theoretic reasoning level		2	Δt_c	Control update interval	s	0.10
Δt_p	Prediction interval	s	0.20				
State and input constraints							
$a_{\max}$	Maximum longitudinal acceleration	m/s^2	3.0	$a_{\min}$	Minimum longitudinal acceleration	m/s^2	-4.0
$\bar{a}_j$	Level-1 nominal neighbor acceleration	m/s^2	0	$e_{\psi,\max}$	Maximum heading error	rad	0.70
$e_{\psi,\min}$	Minimum heading error	rad	-0.70	$\dot{s}_{\text{des}}$	Desired progress speed along centerline	m/s	0.2-1.0
$\delta_{\max}$	Maximum steering angle	rad	1.0	$\delta_{\min}$	Minimum steering angle	rad	-1.0
Structured sampling and level-1 prediction							
k_d	Feedback gain for lateral offset in level-1		0.60	k_ψ	Feedback gain for heading error in level-1		1.0
σ_a	Standard deviation of acceleration noise	m/s^2	0.14	σ_δ	Standard deviation of steering noise	rad	0.035
CVaR and IRA							
R	Safety threshold for static clearance in IRA	m	0.05	α	CVaR confidence level		0.90
ϵ	Total collision-probability budget		0.05	ρ	IRA allocation parameter		1.5
NMPC cost weights							
w_a	Weight on acceleration effort term		0.015	w_r	Weight on CVaR-based risk term		7.5
w_s	Weight on progress-speed tracking term		8.0	w_δ	Weight on steering effort term		0.025
w_ψ	Weight on heading-error term		1.5				

3 Results and discussion

We formulate the NMPC problem in Python using CasADi (a Python framework for nonlinear optimization and automatic differentiation) and solve the resulting nonlinear optimization problem with IPOPT (Interior Point OPTimizer). We set the game-theoretic reasoning level to 2 and evaluate the resulting probabilistic game-theoretic lane-free AV controller (PGC-Level 2) against its level-1 counterpart (PGC-Level 1) as well as a deterministic game-theoretic controller (DGC). The deterministic counterpart removes the structured sampling process and all probability-related components, and retains only a single predicted trajectory for each surrounding vehicle to evaluate the interaction risk and impose the corresponding risk-based constraints.

Table 1 summarizes the main parameters of the proposed algorithm. Note that the vehicle dimensions considered in this study are relatively small, since our ultimate goal is to transfer the proposed algorithm to mobile robots and further evaluate it on our robotic testbed. The desired speeds along track $\dot{s}_{\text{des}}$ of Vehicle 1-5 are 1.0, 0.8, 0.6, 0.4, 0.2 m/s, respectively, to observe the overtaking behavior as well as more interactions.

The complete simulation animations of the three algorithms are available at Zenodo[1]. Table 2 reports the average speed along track of each vehicle and the minimum static clearance under the three algorithms. In these simulations, the experiment ends when the slowest vehicle, namely Vehicle 5, completes five laps. The average speed of each vehicle is calculated as its cumulative travelled distance divided by the total simulation time. The minimum static clearance is defined as the smallest static clearance observed among all pairs of vehicles during the entire simulation.

Deterministic game-theoretic controller

As shown in the animation, the DGC algorithm exhibits a relatively conservative driving style in the considered scenario. At the beginning of the simulation, the green vehicle is cut off by the red vehicle and therefore performs a relatively large detour to avoid a potential conflict, as shown in Figure 3 (a).

Another representative case occurs between 02:40 and 02:55, as shown in Figure 3 (b), where the green vehicle follows the orange vehicle for a relatively long period without attempting to overtake. Only after the orange vehicle starts entering the curve does the green vehicle identify a feasible

[1] https://zenodo.org/records/19326316

Table 2: Speeds along track and static clearance under different algorithms.

	DGC	PGC-Level 1	PGC-Level 2
Vehicle 1 speed (m/s)	0.908	0.927	0.915
Vehicle 2 speed (m/s)	0.731	0.742	0.752
Vehicle 3 speed (m/s)	0.567	0.594	0.578
Vehicle 4 speed (m/s)	0.364	0.391	0.389
Vehicle 5 speed (m/s)	0.207	0.216	0.212
Minimum static clearance (m)	0.313	0.081	0.159

overtaking opportunity and eventually complete the maneuver. These observations indicate that, although the deterministic controller can maintain safe interactions, it tends to behave cautiously and may become hesitant in complex multi-vehicle situations, especially when the overtaking opportunity is not sufficiently clear.
The quantitative results in Table 2 further support this observation. Specifically, the global minimum static clearance remains at 0.313 m, which suggests the overall safety level is acceptable, but the vehicle speed along the track is relatively lower.

Probabilistic game-theoretic controller with level-1 reasoning

The PGC-Level 1 is able to complete most overtaking interactions efficiently, although some oscillations and relatively close encounters can still be observed. For example, between 01:00 and 01:10, as shown in Figure 4 (a), the blue, red, and purple vehicles engage in a multi-vehicle interaction, in which the red and purple vehicles respectively give way to the blue vehicle from the inner and outer sides of the road, allowing the blue vehicle to pass through. However, both the red and purple vehicles also exhibit some oscillatory behavior during this interaction.
Another representative case occurs between 03:15 and 03:25, as shown in Figure 4 (b), where the blue vehicle attempts to overtake the orange vehicle. During this process, the following distance becomes slightly small, but the blue vehicle still successfully completes the overtaking maneuver while maintaining a safe separation from the red vehicle on the outer side.
Table 2 shows that the global minimum static clearance is 0.081 m, indicating that the PGC-Level 1 is riskier than the deterministic algorithm. However, the vehicle speeds under this algorithm are generally higher than those under the deterministic algorithm and are also closer to the desired speed.

Probabilistic game-theoretic controller with level-2 reasoning

Among the three algorithms, the PGC-Level 2 shows the best overall performance. In general, the vehicles under this controller do not exhibit very close car-following behavior, large oscillations, hesitation during overtaking, or long periods of low-speed following caused by reluctance to pass slower vehicles. In addition, the resulting trajectories in the animation are also visually more regular and natural.
Between 01:55 and 02:10, as shown in Figure 5 (a), all five vehicles come into close proximity and form a complex interaction scenario. Nevertheless, the faster vehicles are still able to overtake the slower ones while maintaining safe separation and smooth maneuvers throughout the process. A similar interaction pattern can also be observed between 03:45 and 04:00, as shown in Figure 5 (b).
The quantitative results in Table 2 shows the global minimum static clearance is 0.159 m, which lies between that of the DGC and PGC-Level 1. Even after subtracting the safety buffer, there still remains a considerable safety margin of approximately half a vehicle length, which is sufficient to ensure safe vehicle interactions. At the same time, the vehicle speeds are higher than those under the deterministic algorithm. Therefore, this algorithm is able to maintain a good balance between safety and efficiency.

4 Conclusion

This study investigated lane-free AV control in a curved closed-loop road environment and compared the proposed controller with its level-1 and deterministic counterparts. The results show

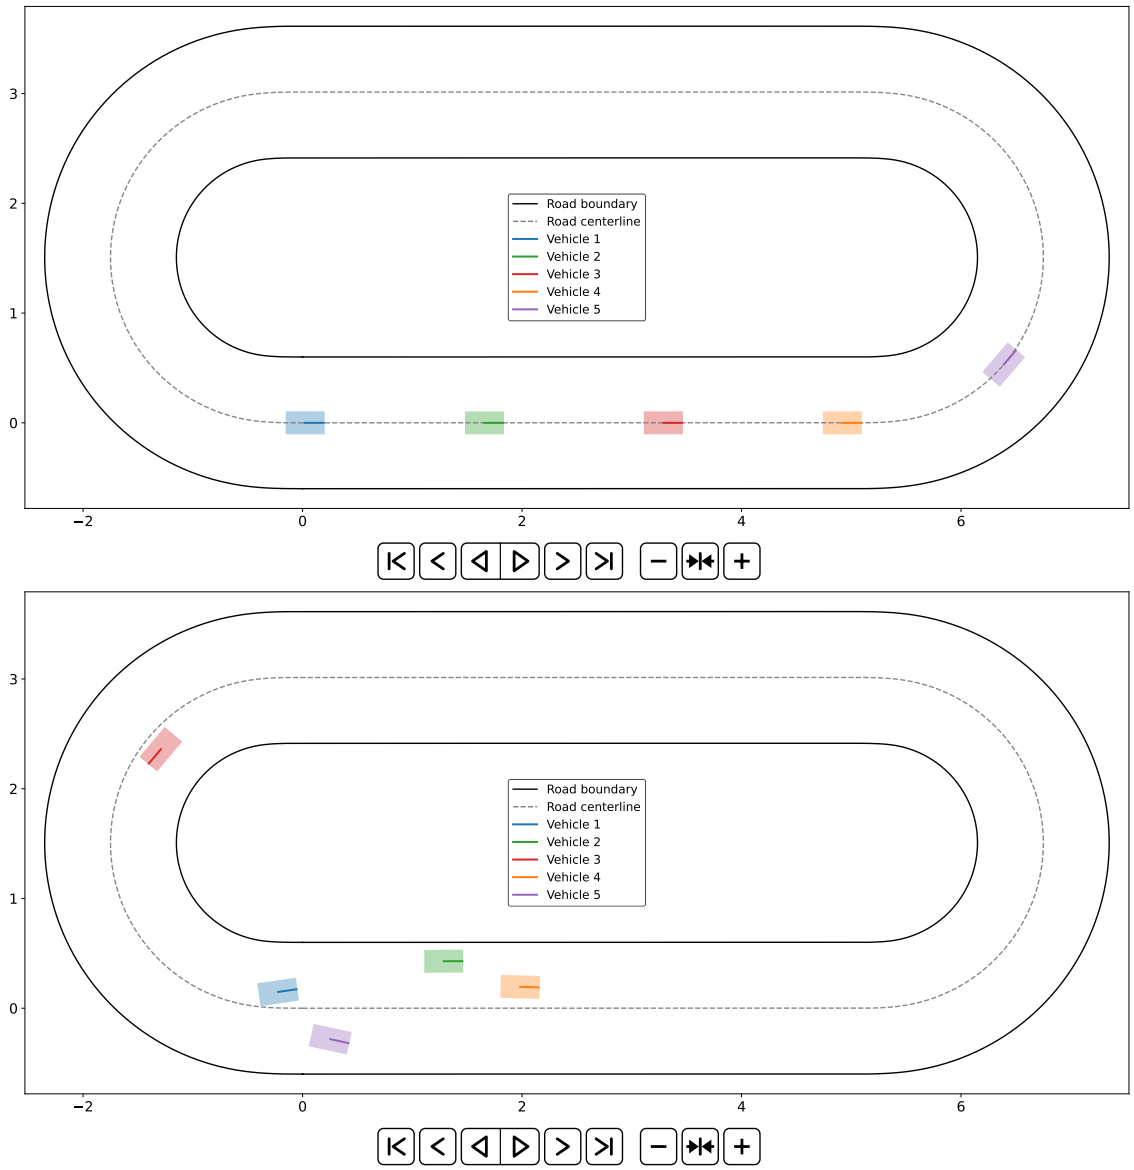


Figure 3: Simulation animation of the DGC: (a) green vehicle cut off by the red vehicle (video time: 00:00 - 00:15), and (b) green vehicle hesitating to overtake (video time: 02:40 - 02:55) [Viewable in Adobe Acrobat].

that, compared with the deterministic controller, the probabilistic game-theoretic controller (PGC) can better reduce overly conservative maneuvers, because the CVaR-based risk and the IRA-based constraint allow the controller to allocate more attention to critical interactions while avoiding unnecessarily cautious responses in less risky situations. Furthermore, compared with the level-1 version, the level-2 version provides a more informative prediction of neighboring vehicles' behaviors and therefore yields smoother, safer, and more efficient maneuvers.
These findings offer potential value for future automated traffic operation in flexible road spaces. In future work, we will increase the number of vehicles to study denser lane-free traffic conditions. We will also transfer the proposed controllers from simulation to our reduced-scale mobile robots for further physical validation.

Acknowledgements

This research was partially funded by the Australian Research Council (ARC) Discovery Project DP220100882.

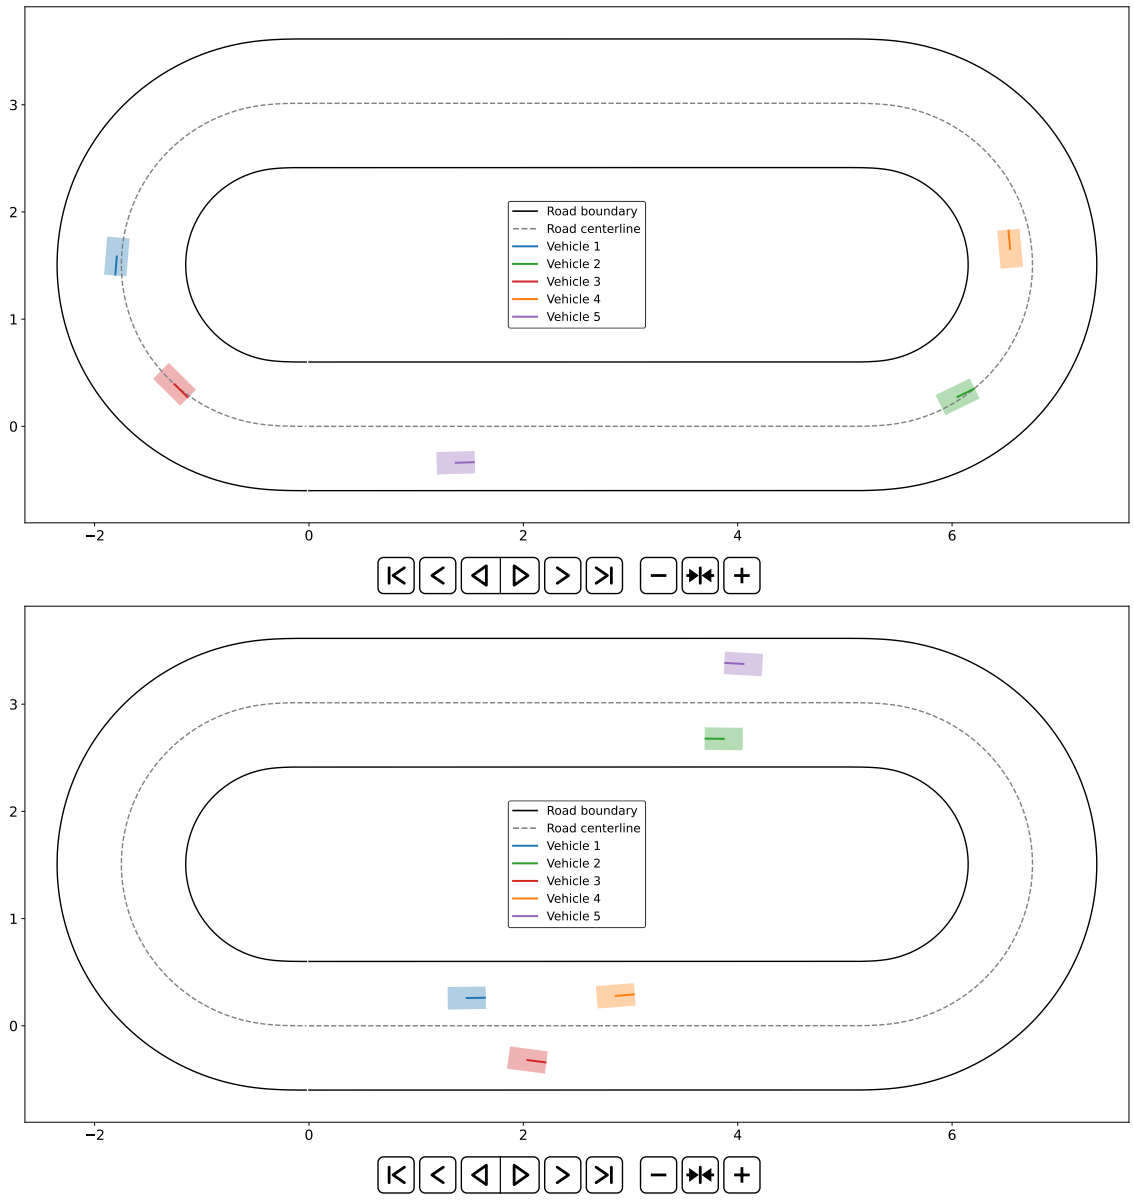
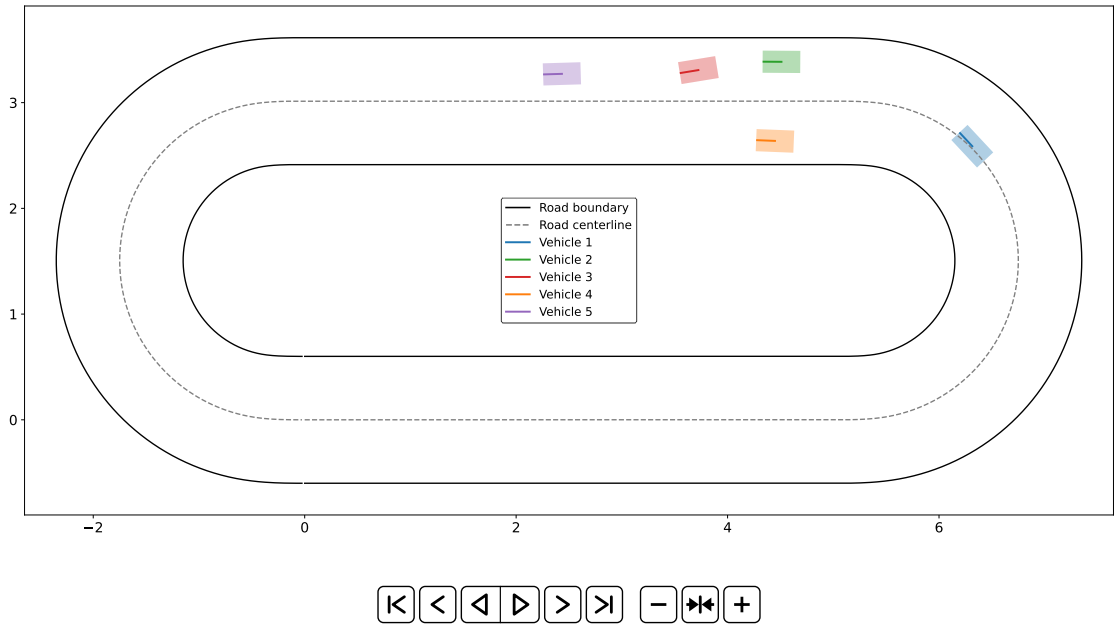


Figure 4: Simulation animation of the PGC-Level 1: (a) blue vehicle passing through a multi-vehicle interaction (video time: 01:00 - 01:10), (b) blue vehicle overtaking orange vehicle with a small gap (video time: 03:15 - 03:25) [Viewable in Adobe Acrobat].



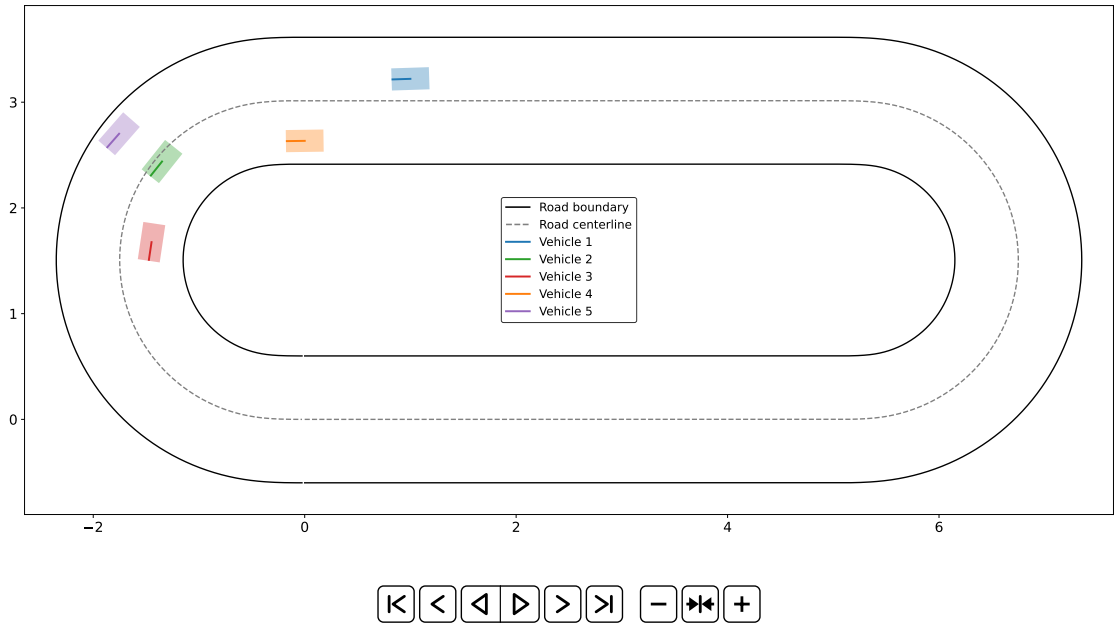


Figure 5: Simulation animation of the PGC-Level 2: (a) smooth overtaking in a dense five-vehicle interaction (video time: 01:55 - 02:10), (b) smooth overtaking in another complex interaction (video time: 03:45 - 04:00) [Viewable in Adobe Acrobat].

References

Berahman, M., Rostmai-Shahrbabaki, M., & Bogenberger, K. (2022). Driving strategy for vehicles in lane-free traffic environment based on deep deterministic policy gradient and artificial forces. *IFAC-PapersOnLine*, *55*(14), 14–21. doi: 10.1016/j.ifacol.2022.07.576

Beza, A. D., Xie, Z., Ramezani, M., & Levinson, D. (2025). From lane-less to lane-free: Implications in the era of automated vehicles. *Transportation Research Part C: Emerging Technologies*, *170*, 104898.

Bonzanini, A. D., Mesbah, A., & Cairano, S. D. (2024). Perception-aware model predictive control for constrained control in unknown environments. *Automatica*, *160*, 111418. doi: 10.1016/j.automatica.2023.111418

Jeong, Y. (2021). Stochastic model-predictive control with uncertainty estimation for autonomous driving at uncontrolled intersections. *Applied Sciences*, *11*(20), 9397. doi: 10.3390/app11209397

Karalakou, A., Troullinos, D., Chalkiadakis, G., & Papageorgiou, M. (2023). Deep reinforcement learning reward function design for autonomous driving in lane-free traffic. *Systems*, *11*(3), 134. doi: 10.3390/systems11030134

Levy, R., & Haddad, J. (2021). Path and trajectory planning for autonomous vehicles on roads without lanes. In *2021 ieee international intelligent transportation systems conference (itsc)* (pp. 3871–3876). doi: 10.1109/ITSC48978.2021.9565023

Levy, R., & Haddad, J. (2022). Cooperative path and trajectory planning for autonomous vehicles on roads without lanes: A laboratory experimental demonstration. *Transportation Research Part C: Emerging Technologies*, *144*, 103813. doi: 10.1016/j.trc.2022.103813

Naderi, M., Papageorgiou, M., Troullinos, D., Karafyllis, I., & Papamichail, I. (2023). Controlling automated vehicles on large lane-free roundabouts. *IEEE Transactions on Intelligent Vehicles*, *9*(1), 3061–3074.

Naderi, M., Typaldos, P., & Papageorgiou, M. (2025). Lane-free signal-free intersection crossing via model predictive control. *Control Engineering Practice*, *154*, 106115. doi: 10.1016/j.conengprac.2024.106115

Papageorgiou, M., Mountakis, K.-S., Karafyllis, I., Papamichail, I., & Wang, Y. (2021). Lane-free artificial-fluid concept for vehicular traffic. *Proceedings of the IEEE*, *109*(2), 114–121. doi: 10.1109/JPROC.2020.3042681

Rostami-Shahrbabaki, M., Weikl, S., Niels, T., & Bogenberger, K. (2023). Modeling vehicle flocking in lane-free automated traffic. *Transportation Research Record*, *2677*(9), 499–512. doi: 10.1177/03611981231159405

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*, 120. doi: 10.1186/s10033-022-00790-5

Yanumula, V. K., Typaldos, P., Troullinos, D., Malekzadeh, M., Papamichail, I., & Papageorgiou, M. (2021). Optimal path planning for connected and automated vehicles in lane-free traffic. In *2021 ieee international intelligent transportation systems conference (itsc)* (pp. 3545–3552). doi: 10.1109/ITSC48978.2021.9564698

Yanumula, V. K., Typaldos, P., Troullinos, D., Malekzadeh, M., Papamichail, I., & Papageorgiou, M. (2023). Optimal trajectory planning for connected and automated vehicles in lane-free traffic with vehicle nudging. *IEEE Transactions on Intelligent Vehicles*, *8*(3), 2385–2399. doi: 10.1109/TIV.2023.3241200

Yin, J., Zhang, Z., & Tsiotras, P. (2023). Risk-aware model predictive path integral control using conditional value-at-risk. In *2023 ieee international conference on robotics and automation (icra)* (pp. 7937–7943). doi: 10.1109/ICRA48891.2023.10161100