

An Efficient Risk-aware Branch MPC for Automated Driving that is Robust to Uncertain Vehicle Behaviors

Zhang, Luyao; Pantazis, Georgios; Han, Shaohang; Grammatico, Sergio

DOI

[10.1109/CDC56724.2024.10886383](https://doi.org/10.1109/CDC56724.2024.10886383)

Publication date

2025

Document Version

Final published version

Published in

Proceedings of the IEEE 63rd Conference on Decision and Control, CDC 2024

Citation (APA)

Zhang, L., Pantazis, G., Han, S., & Grammatico, S. (2025). An Efficient Risk-aware Branch MPC for Automated Driving that is Robust to Uncertain Vehicle Behaviors. In *Proceedings of the IEEE 63rd Conference on Decision and Control, CDC 2024* (pp. 8207-8212). (Proceedings of the IEEE Conference on Decision and Control). IEEE. <https://doi.org/10.1109/CDC56724.2024.10886383>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

An Efficient Risk-aware Branch MPC for Automated Driving that is Robust to Uncertain Vehicle Behaviors

Luyao Zhang¹, Georgios Pantazis¹, Shaohang Han² and Sergio Grammatico¹

Abstract—One of the critical challenges in automated driving is ensuring safety of automated vehicles despite the unknown behavior of the other vehicles. Although motion prediction modules are able to generate a probability distribution associated with various behavior modes, their probabilistic estimates are often inaccurate, thus leading to a possibly unsafe motion plan. To overcome this challenge, we propose an Efficient Risk-Aware Branch MPC (EraBMPC) that appropriately accounts for the ambiguity in the estimated probability distribution. We formulate the risk-aware motion planning problem as a min-max optimization problem and develop an efficient iterative method by incorporating a regularization term in the probability update step. Via extensive numerical studies, we validate the convergence of our method and demonstrate its advantages compared to the state-of-the-art approaches.

I. INTRODUCTION

After two decades of development, automated vehicles have been able to navigate a variety of traffic scenarios successfully. Nevertheless, operating in highly interactive environments, such as intersections without traffic lights, still remains a significant challenge. One main reason is that an automated vehicle struggles to properly account for the different behavior modes of the surrounding vehicles. Specifically, in the classic planning framework [1], the motion planner receives the most likely predicted trajectories of the surrounding vehicles from a prediction module, and then generates a motion plan without considering multi-modal behaviors. One drawback of this framework is that the generated motion plan might be overly aggressive or overly conservative, thus resulting in collisions or traffic jams, respectively.

To properly consider multi-modal behaviors, the so-called branch model predictive control (BMPC) [2] has been proposed, also referred to as contingency planning in [3], [4], or trajectory tree motion planning in [5]–[7]. An advantage of BMPC planners is their ability to utilize the multi-modal trajectory prediction provided by the latest motion prediction module [8]. In contrast with traditional robust motion planners, which aim to generate a motion plan that accommodates all predicted trajectories, BMPC planners construct a trajectory tree with multiple branches corresponding to (possibly) different behavior modes. Thus, BMPC planners

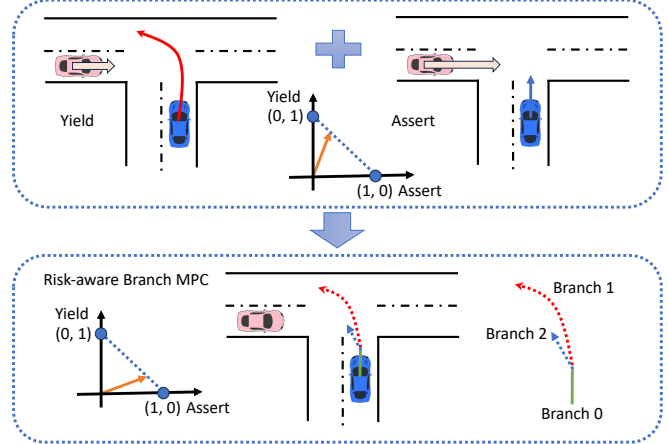


Fig. 1: Unsignalized intersection crossing. The other vehicle (in pink) has two potential behavior modes: “Yield” and “Assert”. The motion predictor indicates that the other vehicle is more likely to “Yield”. However, the behavior mode “Assert” can result in a potential collision. To avoid unsafe motion, the risk-aware branch MPC planner generates a trajectory tree that handles different behavior modes by considering their associated ambiguity. In this example, it focuses more on the behavior mode “Assert” by appropriately reshaping the probability distribution.

can avoid overly conservative motion plans since only the shared branch needs to adapt to all potential predicted scenarios.

However, a significant challenge arising with multi-modal behaviors is that estimates on the probability of each tree branch might be inaccurate [9]. To address this issue, it is necessary to deal with ambiguity in the probability estimates by leveraging tools from risk-aware stochastic optimization. In [2], [10], [11], the authors propose a risk-aware MPC framework, where they employ the so-called Conditional Value at Risk (CVaR), a risk measure that accounts for unlikely scenarios at the tail of the probability distribution. They then develop a risk-aware motion planner, leveraging the dual form of CVaR [12] to obtain a min-max optimization problem with a nonconvex-concave structure.

The works [2], [10] address the min-max reformulation by converting the inner maximization problem into a minimization problem via the dualization technique. However, such a formulation introduces additional auxiliary decision variables and hard constraints. Furthermore, the nonconvexity of such constraints due to the collision cost renders this problem computationally challenging to solve. Another popular category of algorithms for min-max optimization would solve the maximization problem and perform a gradient descent step for the minimization problem. Particularly, in [13], the maximization problem is solved approximately by employing multiple gradient ascent steps. In contrast, in single-loop

*This work is partially supported by NWO under project AMaDeuS and by the ERC under project COSMOS.

¹Luyao Zhang, Georgios Pantazis and Sergio Grammatico are with the Delft Center for Systems and Control, Delft University of Technology, The Netherlands. {l.zhang-7, g.pantazis, s.grammatico}@tudelft.nl.

²Shaohang Han is with the Division of Robotics, Perception and Learning, KTH Royal Institute of Technology, Sweden. shaohang@kth.se.

algorithms such as the gradient descent-ascent (GDA), the update for one variable only occurs once before another variable is updated. However, the GDA method fails to converge even for some simple bilinear problems [14]. To improve convergence, Lu et al. [15] propose the so-called Hybrid Block Successive Approximation (HiBSA), where a regularization term is introduced to the maximization problem, resulting in a perturbed version of the original projected gradient ascent step.

To iteratively solve the outer minimization problem, a projected gradient descent (PGD) method is employed in [13]. However, as the feasible set in the motion planning problem is often non-convex, the projection to such a set is computationally challenging. To compute the update more efficiently, Li et al. [11] replace the gradient descent with iterative LQR (iLQR), an efficient numerical optimization method that exploits the sparse structure inherent in optimal control problems [16, Chapter 8]. iLQR is initially applied to unconstrained motion planning in [17], while in [18], a projected quasi-Newton method is employed to consider input constraints. To handle more general inequality constraints, a barrier function and an augmented Lagrangian function are integrated into the framework of iLQR in [19] and [20], respectively. The authors in [11], [21] have recently developed iLQR-based solvers for the BMPC problem.

In this work, we propose an efficient risk-aware motion planner for applications in autonomous driving. Our contributions with respect to the related literature can be summarized as follows: (i) To circumvent the computational challenges of the formulations in [2], [10], our work builds upon [11] and proposes an iterative algorithm based on the augmented Lagrangian iterative linear quadratic regulator (AL-iLQR) to efficiently solve the minimization problem in the min-max reformulation of the risk-aware motion planning problem. (ii) We combine the method above with a variant of Hybrid Block Successive Approximation (HiBSA) [15]; specifically, to address scenarios of oscillatory behavior observed by implementing the numerical method in [11], we introduce a regularization term to the inner maximization problem. We empirically validate the effectiveness of such an additional term via extensive numerical simulations. (iii) We compare our method with other state-of-the-art methods and demonstrate its applicability in real-time motion planning by considering different intersection-crossing scenarios.

II. PROBLEM FORMULATION

A. Unsignalized intersection crossing

We consider an unsignalized intersection-crossing scenario in which vehicles operate without direct communication. Fig. 1 illustrates such an example, where the ego vehicle makes a left turn while the other vehicle follows a straight path. To turn left successfully, the ego vehicle should take into account the different behavior modes of the other vehicle. In this example, two behavior modes can be identified for the other vehicle. The first behavior mode is called “Yield” mode, i.e., the other vehicle decides to slow down and wait for the ego vehicle to turn left. In this case, the

ego vehicle can maintain the current speed (assuming it is sufficiently slow to make a left turn) and proceed through the intersection. The second behavior mode is the “Assert” mode, where the other vehicle does not decelerate and continues following the straight path. Consequently, the ego vehicle should slow down to avoid a potential collision. The traffic scenario in Fig. 1 illustrates how different behavior modes can have a significant effect on motion planning. To account for this, we introduce a motion planner based on the branch MPC framework.

B. Nominal Branch MPC

A branch MPC planner generates a so-called trajectory tree, i.e., a decision tree whose branches correspond to distinct behavior modes of the surrounding vehicles. An example of the trajectory tree is illustrated in Fig. 1, which comprises one shared branch and a branching point that leads to two possible outcomes. We denote the time step when the shared branch splits into multiple individual ones by T_s and the length of the entire planning horizon by T .

In the trajectory tree, the initial control inputs are constrained to remain consistent over the interval $[0, T_s - 1]$ to allow for adaptation to all potential predicted scenarios. We denote the tree of control inputs by $\bar{\mathbf{u}} := (\bar{\mathbf{u}}^0, \bar{\mathbf{u}}^i)_{i \in [1, d]}$, where d is the number of branches, $\bar{\mathbf{u}}^0 := (u_t^0)_{t \in [0, T_s - 1]}$ represents the shared control input sequence from time step 0 to $T_s - 1$, and $\bar{\mathbf{u}}^i := (u_t^i)_{t \in [T_s, T - 1]}$ denotes the control input sequence for branch i from time step T_s to $T - 1$. We consider a state $x_t^i \in \mathbb{R}^{n_x}$ corresponding to branch i at time step t , which evolves according to the dynamics $x_{t+1}^i = f(x_t^i, u_t^i)$, where $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ is a nonlinear function. The initial state for the shared branch is denoted by x_0 , while x_{T_s} represents the state at the start of the branching. Note that the state evolution can be expressed as a function of the initial condition x_0 and the input $\bar{\mathbf{u}}$. Consider $J^0(x_0, \bar{\mathbf{u}}^0) = \sum_{t=0}^{T_s-1} \ell^0(x_t^0, u_t^0)$ and $J^i(x_{T_s}, \bar{\mathbf{u}}^i) = \sum_{t=T_s}^{T-1} \ell^i(x_t^i, u_t^i) + \ell_T^i(x_T^i)$, where $\ell^i : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ and $\ell_T^i : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$ denote the stage and final cost, respectively, and $p = (p^i)_{i \in [1, d]}$ is the collection of probability estimates for all branches, which originates from a prediction module or behavior planner and takes values in the probability simplex. We define the cost function J of the branch MPC as the weighted sum of the branch costs:

$$J(x_0, \bar{\mathbf{u}}) := J^0(x_0, \bar{\mathbf{u}}^0) + \sum_{i=1}^d p^i J^i(x_{T_s}, \bar{\mathbf{u}}^i). \quad (1)$$

In practice, however, the probability estimate can often be inaccurate, resulting in possibly unsafe motion plans.

C. Risk-aware branch MPC

Revisiting the unsignalized intersection-crossing problem in Fig. 1, the motion predictor assesses the likelihood of each behavior mode and indicates that mode “Yield” has a larger probability of occurrence. However, the behavior mode “Assert” could lead to a potential collision with the ego vehicle, which implies that ambiguity in the probability estimates carries a significant risk for the safety of both

Algorithm 1: Risk-aware AL-iLQR tree

Input: $x_0, \bar{x}_{\text{ref}}, \bar{u}_{\text{ref}}, p$
Output: \bar{u}^*, q
Initialize $\bar{\lambda}_0, \mu_0, \bar{u}_0 \leftarrow \bar{u}_{\text{ref}}, k \leftarrow 0$
Perform rollout using \bar{u}_{ref} to obtain \bar{x}_0
while *stopping criterion not satisfied* **do**
 $\text{minmax_iLQR_tree}(x_0, \bar{x}_{\text{ref}}, \bar{u}_k, p, \bar{\lambda}_k, \mu_k)$;
 Update $\bar{\lambda}_k, \mu_k$;
 $k \leftarrow k + 1$;
end

vehicles. Therefore, the ego vehicle should be more aware of the risk associated with potential misinterpretation of the other vehicle's intentions. To generate a robust motion plan for such intricate cases, we develop a risk-aware branch MPC planner, which focuses more on the potentially dangerous scenarios by minimizing a risk measure of the cost function.

1) *Preliminaries:* Let us first introduce the concept of risk measures. A popular class of risk measures used in stochastic optimization is the so-called coherent risk measures, which satisfy certain properties, including convexity, monotonicity, translation equivariance, and positive homogeneity; see [12] for technical details. While widely used as a coherent risk measure, the expectation fails to effectively account for events that lie at the tail of a given distribution. To circumvent this issue, Conditional Value at Risk (CVaR) [12], [22] is a popular means for improved risk assessment. For computational efficiency, the dual representation of a coherent risk measure [12, Eq. 6.40] is commonly employed: $\rho(X) := \sup_{Q \in \mathcal{A}} \mathbb{E}_Q[X]$, where X is a random variable and the ambiguity set \mathcal{A} is closed and convex. The dual representation indicates that a coherent risk measure can be viewed as the worst-case expectation with respect to all probability distributions in the ambiguity set. In this work, we consider that uncertainty arises from the unknown discrete behavior modes of the surrounding vehicles. This motivates the study of discrete probability distributions $q = (q^i)_{i \in [1, d]}$, where the considered ambiguity set of CVaR_α for $\alpha \in [0, 1]$, is the intersection of the probability simplex and the set of boxes, based on the nominal probability vector p :

$$\mathcal{A}_\alpha(p) = \left\{ q \in \mathbb{R}^d \mid \sum_{i=1}^d q^i = 1, q^i \geq 0, \alpha q^i \leq p^i \right\}. \quad (2)$$

Selecting $\alpha = 0$ implies a lack of confidence in the nominal probability distribution, resulting in the ambiguity set being equivalent to the entire probability simplex. For $\alpha = 1$, the decision maker has more confidence in the nominal probability distribution.

2) *Risk-aware formulation:* By adopting CVaR as a risk measure and leveraging its dual form, the original cost in (1) has the following risk-aware counterpart:

$$J_R(x_0, \bar{u}) := J^0(x_0, \bar{u}^0) + \max_{q \in \mathcal{A}_\alpha(p)} \sum_{i=1}^d q^i J^i(x_{T_s}, \bar{u}^i), \quad (3)$$

where the ambiguity set $\mathcal{A}_\alpha(p)$ is given in (2). Informally speaking, the solution to this maximization problem tends to assign a larger probability to the branch with a higher cost. With the cost function in (3), we can formulate the risk-aware motion planning problem (RAMP) as follows:

$$\text{RAMP} \begin{cases} \min_{\bar{u}} J_R(x_0, \bar{u}) \\ \text{s.t. } x_{t+1}^0 = f(x_t^0, u_t^0), \forall t \in [0, T_s - 1], \\ h^0(x_t^0, u_t^0) \leq 0, \forall t \in [0, T_s - 1], \\ x_{T_s}^i = x_{T_s}^j, \forall i, j \in [1, d], \\ x_{t+1}^i = f(x_t^i, u_t^i), \forall t \in [T_s, T - 1], \forall i \in [1, d], \\ h^i(x_t^i, u_t^i) \leq 0, \forall t \in [T_s, T], \forall i \in [1, d], \end{cases} \quad (4)$$

where $h^i(\cdot)$ encapsulates general inequality constraints for different branches, including control input bounds, collision avoidance constraints, and road boundary constraints. The motion planning problem (RAMP) can be alternatively viewed as a two-player zero-sum game, where the first player seeks a comfortable and collision-free trajectory, while the second player acts as an adversary wishing to increase the total cost by redistributing the probability of each branch.

III. EFFICIENT RISK-AWARE MOTION PLANNING

In the subsequent developments, we draw inspiration from the methods in [11], [15] to obtain a solution to RAMP. Following [11], we keep the probability vector q fixed for each step and solve an approximated optimal control problem of RAMP via the iLQR scheme. We combine this method with a variant of HiBSA proposed in [15]. In particular, to improve convergence, we introduce a regularization term to the maximization problem, which perturbs the original projected gradient ascent step. As in [11], we replace the projected gradient descent with a Newton-like step to further aid in convergence. Our proposed algorithmic scheme is detailed in Algorithm 1 and 2. The vectors \bar{x}_{ref} and \bar{u}_{ref} denote the reference state and input sequences.

A. Augmented Lagrangian iLQR tree

To solve RAMP efficiently, we exploit the inherent sparse structure of the optimal control problem by leveraging ideas from dynamic programming [16]. This leads to an iLQR-based method, where the following subproblem is solved at each time step t :

$$\text{iLQR}_t^i \begin{cases} \min_u Q_t^i(x, u, x') = V_t^i(x) \\ \text{s.t. } x' = f(x, u), h_t^i(x, u) \leq 0, \end{cases} \quad (5)$$

where the Q -function $Q_t^i(\cdot)$ describes the cost incurred after applying the control input u for branch i at time step t , and $h_t^i(\cdot)$ represents general state and control input constraints. Based on dynamic programming, we compute the Q -functions for the trajectory tree as follows:

$$Q_t^0(x, u) = \ell_t^0(x, u) + V_{t+1}^0(f(x, u)), \quad t \in [0, T_s - 2], \quad (6)$$

$$Q_{T_s-1}^0(x, u) = \ell_{T_s-1}^0(x, u) + \sum_{i=1}^d V_{T_s}^i(f(x, u)), \quad (7)$$

$$Q_t^i(x, u) = \ell_t^i(x, u) + V_{t+1}^i(f(x, u)), \quad t \in [T_s, T - 1], \quad (8)$$

Algorithm 2: minmax_iLQR_tree

Input: $x_0, \bar{x}_{\text{ref}}, \bar{u}_0, p, \bar{\lambda}, \mu$
Output: \bar{u}^*, q
Perform rollout using \bar{u}_0 to obtain \bar{x}_0
while *stopping criterion not satisfied* **do**
 Approximate Q -functions;
 Backward pass through Riccati equation (10);
 Forward pass, including rollout and line search;
 if *line search failed* **then**
 Add regularization terms to Hessian matrices;
 else
 Update the probability vector q using (12);
 end
end

where (6) and (8) are associated with the Q -functions of the shared branch and individual branches, respectively. Equation (7) represents the Q -function of the branching node, where the value functions for all branches at the next time step are summed up. Next, the hard constraints are incorporated through the Lagrangian as follows:

$$Q_{a,t}^i(x, u, \lambda, \mu) = Q_t^i(x, u) + \lambda^\top h_t^i(x, u) + \frac{1}{2} h_t^i(x, u)^\top I_{\mu,t}^i h_t^i(x, u), \quad (9)$$

where λ is a vector of Lagrange multipliers, and μ is a penalty weight. $I_{\mu,t}^i$ denotes a diagonal matrix that selects the active constraints:

$$I_{\mu,t,m}^i = 0 \text{ if } h_{t,m}^i(x, u) < 0 \text{ \& } \lambda_m = 0, \text{ otherwise } \mu,$$

where m is the index of the m -th constraint. By using augmented penalty terms, we now convert iLQR_tⁱ into an unconstrained optimization problem. We then linearize the dynamics and approximate the Q -functions using second-order Taylor expansions. We employ the generalized Gauss-Newton Hessian approximation due to its ease of computation and the theoretical guarantee that the approximated Hessian is always positive semi-definite [23]. The Newton descent direction at time step t can then be obtained by minimizing the approximated version $\hat{Q}_{a,t}^i$ of the augmented Lagrangian Q -function:

$$\min_{\delta u} \hat{Q}_{a,t}^i(\delta x, \delta u). \quad (10)$$

Since $\hat{Q}_{a,t}^i$ is in quadratic form, we can derive an affine control law of the form $\delta u = K_t^i \delta x + d_t^i$ [20]. The backward pass involves solving (10) from the leaf tree nodes to the root node. It is worth noting that the computation from the leaf nodes to the branching node can be performed in parallel, thus improving the computation speed significantly. After the backward pass, we conduct a forward rollout using the nonlinear dynamics to obtain the updated trajectory tree. A backtracking line search [24] is performed to achieve a sufficient decrease in the trajectory cost. We introduce a regularization term to aid the Newton direction towards the gradient descent direction in case the line search fails [18].

B. Projected gradient ascent with regularization

We note that the risk-aware motion planning in (4) is a nonconvex-concave problem. For this class of problems, applying the GDA method directly might result in oscillations [14]. Therefore, motivated by [15], we render the original cost function (3) strongly concave with respect to q by introducing a quadratic regularization term:

$$J_R(x_0, \bar{u}) := J^0(x_0, \bar{u}^0) + \max_{q \in \mathcal{A}_\alpha(p)} \sum_{i=1}^d q^i J^i(x_{T_s}, \bar{u}^i) - \frac{\rho}{2} q^i{}^2, \quad (11)$$

where $\rho > 0$ is the regularization weight. A large value of ρ drives the vector q towards the centroid of the simplex. To mitigate the impact of the regularization term on the saddle point, we adopt a diminishing regularization weight $\rho_k = \rho_0 / \sqrt{k+1}$. We then perform a projected gradient ascent to approximately solve (11) at iteration k :

$$q_{k+1} \leftarrow \text{proj}_{\mathcal{A}_\alpha(p)} \left((1 - \gamma \rho_k) q_k + \gamma \bar{J}(x_{T_s}, \bar{u}) \right), \quad (12)$$

where $\bar{J}(\cdot) = (J^i(\cdot))_{i \in [1,d]}$ is the collection of all branch costs except for the shared branch, $\text{proj}(\cdot)$ denotes the projection operator and $\gamma > 0$ is the step size. We compute the projection $\text{proj}_{\mathcal{A}_\alpha(p)}(q) = \text{proj}_{\text{box}[0, \frac{p}{\alpha}]}(q - \phi^* \mathbf{1}_d)$, where $\text{box}[0, \frac{p}{\alpha}] := \{q \in \mathbb{R}^d \mid 0 \leq q^i \leq \frac{p}{\alpha}, \forall i \in [1,d]\}$, $\mathbf{1}_d = (1, 1, \dots, 1) \in \mathbb{R}^d$ represents a vector with all elements being 1, and ϕ^* is a root of the equation $m(\phi) := \mathbf{1}_d^\top \text{proj}_{\text{box}[0, \frac{p}{\alpha}]}(q - \phi \mathbf{1}_d) - 1 = 0$. Since $m(\phi)$ is non-increasing, its root can be efficiently computed through the bisection method [25, Chapter 6.4.3].

C. Detailed formulation

1) *Vehicle modeling*: We model the vehicle as a kinematic bicycle with the state vector $\tilde{x} := (p_x, p_y, \theta, v)$, the control input vector $u := (a, \delta)$, where (p_x, p_y) , θ , and v represent the rear-axle position, heading angle, and speed of the vehicle, respectively; a and δ are the acceleration and steering angle. In (4), the stage cost only depends on the current state and control input. Therefore, to penalize the rate of change of the input, we augment the state by concatenating it with the previous control input, resulting in $x_k := [\tilde{x}_k^\top, u_{k-1}^\top]^\top$.

2) *Cost*: The motion planner is designed to track the reference trajectory tree, maximize the driving comfort level, and keep a safe distance from surrounding vehicles. To account for the driving comfort, we penalize the control input and its rate of change. Moreover, we design the safety cost between the ego vehicle and the other vehicle as $\ell_{\text{saf}} = \max(D_{\text{prox}} - D, 0)^2$, where D is the Euclidean distance between two vehicles, and D_{prox} denotes the threshold distance.

3) *Safety constraints*: To account for the lack of smoothness of the distance between two vehicles, we use instead an overapproximation of the shape of each vehicle, which comprises the union of a collection of linked circles [6], as shown in Fig. 2a. Additionally, following [26], we approximate the road boundary constraints by constructing a safe driving corridor along the reference trajectory tree provided by the behavior planner.




Fig. 2: Test scenarios. The ego vehicle (in blue) intends to turn left. (a) Both surrounding vehicles have two potential behavior modes: “Yield” and “Assert”, represented by the red and orange arrows, respectively. (b) The upper vehicle exhibits different behavior modes: “TurnLeft” and “GoStraight”.

IV. NUMERICAL SIMULATIONS

We test our method (EraBMPC) in two distinct unsignalized intersection-crossing scenarios, as illustrated in Fig. 2. In both scenarios, the ego vehicle (in blue) intends to turn left, while two surrounding vehicles are present with different driving intentions per scenario. Since two possible behavior modes are considered for each surrounding vehicle, we have in total four different combinations of behavior modes.

Note that the choice of the initial guess of the trajectory tree can significantly influence the convergence speed or in some cases hinder the convergence of the motion planner. As such, the initial guess should be appropriately selected. In our setup, we adopt a simple sampling-based behavior planner [27], [28]. Specifically, we control the longitudinal and lateral motion of the ego vehicle via a PD controller and a pure pursuit controller, respectively, and forward simulate its motion under different desired speeds to obtain a set of trajectories. For each joint behavior mode, we select the best trajectory from the trajectory set based on certain user-defined criteria. Our motion planner operates at 10 Hz with a discretization step of 0.1 s and a planning horizon of 5 s. The number of shared nodes T_s is set to 5. We conduct all simulations on a laptop with a 2.30 GHz Intel Core i7-11800H processor and 16 GB RAM. The code is publicly available at <https://github.com/Luyao787/Era-BMPC>.

We run open-loop Monte Carlo simulations to empirically validate the convergence of the proposed motion planner under 500 different initial states of the ego vehicle. To obtain these states, we perturb the nominal initial state by ± 3 m for the longitudinal position, ± 1 m for the lateral position, $\pm 10\%$ for the longitudinal speed. Additionally, we set $p^i = 0.25$ for all $i \in [1, d]$ and $\alpha = 0.6$ when constructing the ambiguity set. We compare our motion planner with MARC [11] and Dual MPC [2]. We implement Dual MPC using the IPOPT [29] interface provided by CasADi [30] with MA57 [31] as the linear solver for enhanced performance. The statistical results on convergence and average computation time are presented in Table I. Our method achieves successful convergence in all the cases. However, when employing MARC, we observe that in certain cases the values of the probability vector q oscillate between two vertices of the ambiguity set. A potential reason is that the solver of the linear program often outputs a vertex as the optimal solution, but the saddle point might be a point on a facet of the

TABLE I: Statistical Results

Test Scenario	Metric	EraBMPC	MARC [11]	Dual MPC [2]
TS1	Convergence	100 %	96.8 %	100 %
	Runtime (ms)	4.45	10.50	1300.8
TS2	Convergence	100 %	100 %	100 %
	Runtime (ms)	14.22	14.52	1520.7

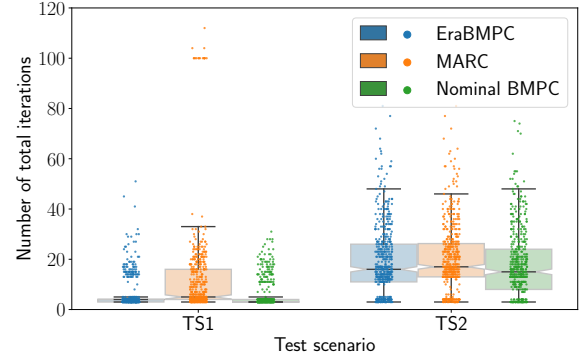


Fig. 3: Box-plots of the total number of iterations obtained from 500 Monte Carlo simulations.




Fig. 4: Velocity profiles for the risk-aware motion planners with $\alpha = 0.1$ and $\alpha = 0.8$ and the nominal motion planner. The nominal probability vector is set to $(0.25, 0.25, 0.25, 0.25)$.

ambiguity set that is not necessarily a vertex. Dual MPC has the highest computational cost since additional decision variables are introduced and the original nonconvex cost becomes a part of hard constraints after dualization.

Fig. 3 illustrates that the number of total iterations required to solve a risk-aware branch MPC problem is slightly larger than that required for the nominal branch MPC problem. Simulation studies indicate that the difference in the number of required iterations is due to the presence of gradually decaying oscillations in the probability update. Such an update is not present in the case of the nominal branch MPC. Such extra computational overheads are still acceptable from a practical point of view, since the projection in (12) can be efficiently computed in just a few microseconds, and the average computation time of our method is below 100 ms. Thus, the proposed planner is well-suited for real-time motion planning after code optimization.

We now compare the closed-loop performance between the risk-aware branch MPC and the nominal branch MPC for the scenario (TS1). We assume that the ego vehicle does not fully know the behavior modes of the surrounding vehicles before

$t_a = 1$ s. In practice, the ego vehicle can infer the intentions of surrounding vehicles based on historical observations after a certain period. Considering this, we assume that the ego vehicle has complete knowledge of these intentions after $t_a = 1$ s. Our simulation studies indicate that, as expected, an increase in t_a , results in gradually more conservative behavior. As the risk-aware planner accounts for the risk against different behavior mode realizations before t_a , its closed-loop trajectory differs from that of the nominal motion planner. Note that even after t_a , when the ego vehicle knows the intentions of the surrounding vehicles, the trajectory of the risk-aware and nominal planners are still different since the initial differences in the trajectory affect their future evolution. Fig. 4 shows the velocity profile of the ego vehicle for the nominal motion planner and the risk-aware motion planner for two distinct values of the risk parameter $\alpha \in [0, 1]$. We note that smaller values of α imply the presence of more ambiguity concerning the probabilistic estimates of the behavior modes. We observe that for up to $t = 3$ s, the velocity of the risk-aware motion is slower than the nominal one, as the risk-aware approach considers the risk of the SV1 not yielding. After $t = 3$ s the velocity increases compared to the nominal motion planner, and the profiles coincide after $t = 7$ s. Finally, we note that for $\alpha = 0.1$, the motion planner is, as expected, more risk-aware of the SV1's behavior, while for $\alpha = 0.8$, its behavior is closer to the nominal one, since ambiguity is smaller.

V. CONCLUSION

Risk-aware branch model predictive control is applicable to motion planning for automated vehicles subject to behavioral uncertainty. However, general-purpose numerical solvers are currently not fast enough to solve the motion planning problem within the given sampling time. Our method shows high potential in closing this computational gap and paves the way towards more efficient real-time motion planning. In future work, we aim at designing an efficient interaction-aware behavior planner for intersection-crossing scenarios. Our final goal is to develop a general motion planning framework by integrating an advanced behavior planner with the proposed risk-aware motion planner.

REFERENCES

- [1] H. Fan, F. Zhu, C. Liu, *et al.*, "Baidu Apollo EM Motion Planner," *arXiv:1807.08048 [cs]*, Jul. 2018, arXiv: 1807.08048.
- [2] Y. Chen, U. Rosolia, W. Ubellacker, N. Csomay-Shanklin, and A. D. Ames, "Interactive multi-modal motion planning with branch model predictive control," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5365–5372, 2022.
- [3] J. Hardy and M. Campbell, "Contingency planning over probabilistic obstacle predictions for autonomous road vehicles," *IEEE Transactions on Robotics*, vol. 29, no. 4, pp. 913–929, 2013.
- [4] J. P. Alsterda, M. Brown, and J. C. Gerdes, "Contingency model predictive control for automated vehicles," in *2019 American Control Conference (ACC)*, 2019, pp. 717–722.
- [5] H. Hu and J. F. Fisac, "Active Uncertainty Reduction for Human-Robot Interaction: An Implicit Dual Control Approach," in *Algorithmic Foundations of Robotics XV*, S. M. LaValle, J. M. O'Kane, M. Otte, D. Sadigh, and P. Tokekar, Eds., ser. Springer Proceedings in Advanced Robotics, Springer International Publishing, 2023, pp. 385–401.
- [6] R. Wang, M. Schuurmans, and P. Patrino, "Interaction-aware model predictive control for autonomous driving," in *2023 European Control Conference (ECC)*, 2023, pp. 1–6.
- [7] I. Batkovic, U. Rosolia, M. Zanon, and P. Falcone, "A robust scenario mpc approach for uncertain multi-modal obstacles," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 947–952, 2021.
- [8] J. Gu, C. Sun, and H. Zhao, "DenseTNT: End-to-end Trajectory Prediction from Dense Goal Sets," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Montreal, QC, Canada: IEEE, Oct. 2021, pp. 15 283–15 292.
- [9] Y. Chen, S. Veer, P. Karkus, and M. Pavone, "Interactive joint planning for autonomous vehicles," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 987–994, 2024.
- [10] P. Sopasakis, D. Herceg, A. Bemporad, and P. Patrino, "Risk-averse model predictive control," *Automatica*, vol. 100, pp. 281–288, 2019.
- [11] T. Li, L. Zhang, S. Liu, and S. Shen, "Marc: Multipolicy and risk-aware contingency planning for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6587–6594, 2023.
- [12] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory, Third Edition*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2021.
- [13] M. Nouiehed, M. Sanjabi, T. Huang, J. D. Lee, and M. Razaviyayn, *Solving a Class of Non-Convex Min-Max Games Using Iterative First Order Methods*, arXiv:1902.08297 [cs, math, stat], Oct. 2019.
- [14] M. Razaviyayn, T. Huang, S. Lu, M. Nouiehed, M. Sanjabi, and M. Hong, "Nonconvex min-max optimization: Applications, challenges, and recent theoretical advances," *IEEE Signal Processing Magazine*, vol. 37, no. 5, pp. 55–66, 2020.
- [15] S. Lu, I. Tsaknakis, M. Hong, and Y. Chen, "Hybrid block successive approximation for one-sided non-convex min-max problems: Algorithms and applications," *IEEE Transactions on Signal Processing*, vol. 68, pp. 3676–3691, 2020.
- [16] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model predictive control: theory, computation, and design*, en, 2nd edition. Madison, Wisconsin: Nob Hill Publishing, 2017, ISBN: 978-0-9759377-3-0.
- [17] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," in *Proceedings of the First International Conference on Informatics in Control, Automation and Robotics*, Setúbal, Portugal, 2004, pp. 222–229.
- [18] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 1168–1175.
- [19] J. Chen, W. Zhan, and M. Tomizuka, "Autonomous driving motion planning with constrained iterative lqr," *IEEE Transactions on Intelligent Vehicles*, vol. 4, no. 2, pp. 244–254, 2019.
- [20] T. A. Howell, B. E. Jackson, and Z. Manchester, "Altro: A fast solver for constrained trajectory optimization," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 7674–7679.
- [21] F. Da, "Comprehensive reactive safety: No need for a trajectory if you have a strategy," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 2903–2910.
- [22] R. T. Rockafellar, S. P. Uryasev, and M. Zabarankin, *Deviation Measures in Risk Analysis and Optimization*, en, SSRN Scholarly Paper, Rochester, NY, Dec. 2002.
- [23] N. N. Schraudolph, "Fast Curvature Matrix-Vector Products for Second-Order Gradient Descent," *Neural Computation*, vol. 14, no. 7, pp. 1723–1738, Jul. 2002.
- [24] J. Nocedal and S. J. Wright, *Numerical Optimization* (Springer Series in Operations Research and Financial Engineering), en. Springer New York, 2006, ISBN: 978-0-387-30303-1.
- [25] A. Beck, *First-Order Methods in Optimization*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2017.
- [26] A. Liniger, A. Domahidi, and M. Morari, "Optimization-Based Autonomous Racing of 1:43 Scale RC Cars," *Optimal Control Applications and Methods*, vol. 36, no. 5, pp. 628–647, Sep. 2015, arXiv:1711.07300 [cs, math].
- [27] L. Zhang, S. Han, and S. Grammatico, "An efficient game-theoretic planner for automated lane merging with multi-modal behavior understanding," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 2023, pp. 3085–3090.
- [28] W. Ding, L. Zhang, J. Chen, and S. Shen, "Epsilon: An efficient planning system for automated vehicles in highly interactive environments," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 1118–1138, 2022.
- [29] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," en, *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, Mar. 2006.
- [30] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi: A software framework for nonlinear optimization and optimal control," en, *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, Mar. 2019.
- [31] I. S. Duff, "MA57—a code for the solution of sparse symmetric definite and indefinite systems," *ACM Transactions on Mathematical Software*, vol. 30, no. 2, pp. 118–144, Jun. 2004.