

Towards a Safe Real-Time Motion Planning Framework for Autonomous Driving Systems: A Model Predictive Path Integral Approach

Mehdi Testouri
University of Luxembourg
Interdisciplinary Centre for,
Security, Reliability and Trust (SnT)
mehdi.testouri@uni.lu

Gamal Elghazaly
University of Luxembourg
Interdisciplinary Centre for,
Security, Reliability and Trust (SnT)
gamal.elghazaly@uni.lu

Raphael Frank
University of Luxembourg
Interdisciplinary Centre for,
Security, Reliability and Trust (SnT)
raphael.frank@uni.lu

Abstract—Planning safe trajectories in Autonomous Driving Systems (ADS) is a complex problem to solve in real-time. The main challenge to solve this problem arises from the various conditions and constraints imposed by road geometry, semantics and traffic rules, as well as the presence of dynamic agents. Recently, Model Predictive Path Integral (MPPI) has shown to be an effective framework for optimal motion planning and control in robot navigation in unstructured and highly uncertain environments. In this paper, we formulate the motion planning problem in ADS as a nonlinear stochastic dynamic optimization problem that can be solved using an MPPI strategy. The main technical contribution of this work is a method to handle obstacles within the MPPI formulation safely. In this method, obstacles are approximated by circles that can be easily integrated into the MPPI cost formulation while considering safety margins. The proposed MPPI framework has been efficiently implemented in our autonomous vehicle and experimentally validated using three different primitive scenarios. Experimental results show that generated trajectories are safe, feasible and perfectly achieve the planning objective. The video results as well as the open-source implementation are available at <https://gitlab.uni.lu/360lab-public/mppi>.

Index Terms—Autonomous Driving, Motion Planning, Stochastic Optimization

I. INTRODUCTION

Motion planning is a complex and challenging problem in robotics and autonomous systems [1]. A motion planner has to consider system dynamics, constraints, sensor uncertainties, and real-time constraints to ensure that the system safely and efficiently navigates a given environment. The ultimate goal of a motion planner is to generate a feasible, collision-free, comfortable, and fail-safe trajectory, while strictly satisfying a set of spatio-temporal constraints. Typically these constraints allow generated trajectories to (a) avoid collision with obstacles surrounding the vehicle, (b) ensure dynamic and kinematic feasibility as well as vehicle actuation capabilities (e.g. limits of steering angles, steering rate, maximum acceleration and braking torques), and (c) take into account the complex geometry and topology of the road environment. Violating any of these constraints may lead to unpredictable and risky behaviour of the autonomous vehicle.



Fig. 1: Validation of the MPPI-based motion planner on a closed track using a modified KIA Soul EV equipped for autonomous driving. The figure depicts the object avoidance scenario whereby a virtual object was placed on the path.

Motion planning for robots and autonomous systems somehow amounts to solving a complex optimization problem. The main challenge to solve a motion planning problem derives mainly from the non-convexity of the underlying optimization problem. Furthermore, the spatio-temporal constraints of dynamic obstacles hinder most of the motion planning approaches to reach a feasible trajectory within predefined real-time constraints [2]. Ensuring these constraints are respected is fundamental in safety-critical situations, where the autonomous vehicle must be able to anticipate in a timely manner to prevent a potential collision.

Motion planning is quite an established research topic in the vast state-of-the-art of autonomous vehicles [1], [3]. Several approaches have been developed in the literature to solve the above-described trajectory planning challenge. Recently there has been an increasing interest in using learning-based methods, especially reinforcement learning [4], [5]. Nevertheless, solutions based on these methods are unlikely to meet the safety requirements of autonomous vehicles and it becomes difficult to verify and interpret, thus putting their reliability in question, especially for safety-critical situations

[6]. Numerical optimization has also been a widely accepted approach to solve motion planning problems [7], [8]. Although efficient numerical optimization algorithms exist to converge to a solution, one limitation of these algorithms is that they may become stuck in a local minimum due to non-convexity. Decoupling longitudinal and lateral motions [2] and estimating the driving corridor [9] is one way around reformulating the problem as a convex one. This approximation may, however, result in slightly infeasible trajectories.

Methods from control theory have gained significant interest in motion planning. These methods formulate the motion problem as an optimal control problem subject to constraints on both the states and inputs of the systems [1]. Model Predictive Control (MPC) is one of the efficient methods to solve constrained optimal control problems using the receding horizon concept. MPC iteratively solves an optimal control sequence along a finite-time horizon while minimising the cost function of a planning objective [10]. A major limitation of MPC in motion planning lies in the high computational cost it needs, especially with high-dimensional nonlinear systems and the difficulty to consider obstacles in its formulation. Recently, Model Predictive Path Integral (MPPI) has shown to be an effective framework for optimal motion planning and control for robots and autonomous systems [11], [12]. The key idea of MPPI is to transform the optimal control problem into an expectation over the integral of all possible trajectories [13]–[16]. This enables to solve the optimal control problem using Monte Carlo simulations by forward sampling a large number of trajectories that can be evaluated efficiently using parallel processing, e.g. using GPUs. The way most MPPI frameworks consider avoiding obstacles in their formulation is still in its infancy with obstacles considered as cost maps [17]. Furthermore, most of these frameworks are tailored for robot navigation in unstructured and highly uncertain environments [11], [18], [19]. In the context of Autonomous Driving Systems (ADS), only very few results are available [17], [20]. This paper takes a step forward by developing an MPPI framework that systematically considers obstacles and safety in an autonomous driving setting. Obstacles provided by the perception system are approximated by a finite set of circles while considering a safety margin. This way of representing obstacles makes it possible to penalize trajectories at risk of collision more easily. The proposed MPPI framework is efficiently implemented in our autonomous vehicle and experimentally validated in three different primitive scenarios. The first scenario is to allow the vehicle to merge a lane starting from a lane neighbourhood while the second is to perform a turning manoeuvre to avoid an obstacle. The third scenario is to have the autonomous vehicle follows an obstacle in the same lane. Although these scenarios are individually simple, they indeed meet the requirements of generated trajectories for complex scenarios if combined together. Experimental results show that generated trajectories are safe, feasible and perfectly achieve the planning objective.

The rest of this paper is organized as follows. Section II presents some mathematical preliminaries, the modelling of

the vehicle as well as the formulation of the motion planning as a constrained stochastic nonlinear dynamic optimization problem. The proposed MPPI motion planning framework is presented in Section III. The experimental results are presented in Section IV. Finally, we draw conclusions and discuss future works in Section V.

II. PRELIMINARIES AND PROBLEM STATEMENT

The objective of this paper is to plan feasible safe trajectories for autonomous vehicles to navigate their environment. This section presents a preliminary mathematical background for vehicle modelling and the foundations of MPPI used to develop the approach of this work.

A. Problem Formulation

The motion planning problem in autonomous driving systems can be defined in terms of four main components $\mathcal{P} : \langle \mathcal{V}, \mathcal{S}, \mathcal{B}, \mathcal{J} \rangle$, namely a vehicle motion model \mathcal{V} , a scene representation model \mathcal{S} , a boundaries model \mathcal{B} , and an objective \mathcal{J} . Let us consider the following stochastic nonlinear discrete-time dynamical system to represent the motion model of the vehicle

$$\mathcal{V} : \mathbf{x}_{t+1} = \mathcal{F}(\mathbf{x}_t, \mathbf{u}_t + \epsilon_t) + \mathbf{w}_t, \quad (1a)$$

$$\text{s.t. } \mathbf{h}(\mathbf{x}_t, \mathbf{u}_t) \leq 0 \quad (1b)$$

with $\mathbf{x}_t \in \mathcal{X} \subset \mathbb{R}^n$ is the system state, $\mathbf{u}_t \in \mathcal{U} \subset \mathbb{R}^m$ is the control input, $\epsilon_t \sim \mathcal{N}(0, \Sigma)$ is an additive zero-mean Gaussian noise to control inputs with covariance Σ , $\mathbf{w}_t \in \mathbb{R}^n$ represents unknown but bounded state disturbances, and $\mathcal{F} : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$ is a state transition function. The function $\mathbf{h}(\mathbf{x}_t, \mathbf{u}_t)$ encapsulates all constraints imposed by the system itself, e.g. state and control lower and upper limits as well as kinematic constraints. The scene representation \mathcal{S} defines a model of obstacles to be avoided by the vehicle whether these obstacles are static or dynamic. Let $\mathcal{O}_i^s \subset \mathcal{X}$, $\forall i \in \{1, 2, \dots, n_s\}$ be a set of static obstacles in vehicle's environment and $\mathcal{O}_j^d(t) \subset \mathcal{X}$, $\forall j \in \{1, 2, \dots, n_d\}$. Dynamic obstacles impose additional constraints on the free space available for planning the motion of the vehicle as depicted in Fig. 2. These additional constraints account for collision-free conditions along a given planning horizon $T = t_f - t_0$, where t_0 and t_f are the start and end times of the planning cycle, respectively. The subsets of the state space occupied by obstacles must be avoided while planning vehicle motion. One can represent the scene model as

$$\mathcal{S} : \mathbf{x}_t \notin \cup_{i=1}^{n_s} \mathcal{O}_i^s, \quad (2a)$$

$$\mathbf{x}_t \notin \cup_{j=1}^{n_d} \mathcal{O}_j^d(t), \quad \forall t \in [t_0, t_f] \quad (2b)$$

Unlike robot motion planning in an unstructured environment, autonomous vehicles have to respect the constraints posed by the geometry of a highly structured road environment, often represented by a high-definition map [21].

Basically, this condition can be fulfilled if the vehicle state trajectory is allowed to follow a given reference state $\mathbf{x}_t^r \in [t_0, t_f]$ along the planning horizon. This objective can be satisfied by adding a function $\phi(\mathbf{x}_t, \mathbf{x}_t^r)$ to the cost function to be minimized

$$\mathcal{B} : \phi(\mathbf{x}_t, \mathbf{x}_t^r) \in \mathbb{R}^+ \forall t \in [t_0, t_f] \quad (3)$$

A motion planner seeks to generate a feasible, smooth and collision-free trajectory. More precisely, the planner must handle collisions with surrounding obstacles in the vehicle's environment, safely as defined by (2). Furthermore, it must ensure dynamic and kinematic feasibility as well as vehicle actuation capabilities (e.g. limits of steering angles, steering rate, maximum acceleration and braking torques) as defined in (1). Finally, it must take into account the complex geometry and topology of the road environment defined in (3). The motion planning problem \mathcal{P} can be expressed mathematically by combining (1), (2), and (3) into the following optimization problem

$$\begin{aligned} \min_{\mathbf{u}} \quad & \mathbb{E} \left[\pi(\mathbf{x}_f, \mathbf{x}_f^r) + \sum_{t=t_0}^{t_f-1} \left(\phi(\mathbf{x}_t, \mathbf{x}_t^r) + \frac{1}{2} \mathbf{u}_t^\top \mathbf{R} \mathbf{u}_t \right) \right], \\ \text{s.t.} \quad & \mathbf{x}_{t+1} = \mathcal{F}(\mathbf{x}_t, \mathbf{u}_t + \epsilon_t), \quad \epsilon_t \sim \mathcal{N}(\mathbf{0}, \Sigma), \\ & \mathbf{h}(\mathbf{x}_t, \mathbf{u}_t) \leq 0, \\ & \mathbf{x}_t \notin \cup_{i=1}^{n_s} \mathcal{O}_i^s, \quad \mathbf{x}_t \notin \cup_{j=1}^{n_d} \mathcal{O}_j^d(t), \\ & \mathbf{x}_t, \mathbf{x}_t^r \in \mathcal{X}, \quad \mathbf{u}_t \in \mathcal{U}, \end{aligned} \quad (4)$$

B. Vehicle Motion Model

For simplicity and to reduce computational complexity, this work adopts a bicycle model to represent the motion model of the vehicle as shown in Fig. 3. Let x and y be the coordinates of a vehicle position with respect to a given reference map coordinate system $\mathcal{R}_m : \{O, X, Y\}$, θ be the orientation of the vehicle with respect to \mathcal{R}_m and δ is steering angle. Let r be the radius of the circle centred at the instantaneous centre of rotation (ICR), and the linear velocity of the vehicle v is related to angular velocity about ICR as $v = \dot{\theta}r$. From the triangle formed by the vehicle and the ICR, we have $\tan(\delta) = l/r$, where l is the wheelbase of the vehicle. Thus, one can express the motion model of the vehicle as

$$\frac{d}{dt} \begin{pmatrix} x \\ y \\ \theta \\ v \\ \delta \end{pmatrix} = \begin{pmatrix} v \cos(\theta) \\ v \sin(\theta) \\ v \tan(\delta)/l \\ a \\ \omega \end{pmatrix} \quad (5)$$

where a is the vehicle linear acceleration and ω is the steering rate. One can easily rewrite the vehicle motion model given by (5) in the form of (1a) with the state vector $\mathbf{x} = (x, y, \theta, v, \delta)^T$ and control input $\mathbf{u} = (a, \omega)^T$.

III. MPPI MOTION PLANNING FRAMEWORK

The complex optimization problem in (4) is challenging to solve mainly because of the non-convexity of the search space.




Fig. 2: Definition of the motion planning problem. Given an initial state of the vehicle \mathbf{x}_0 , a reference state \mathbf{x}_t^r , a set of static and \mathcal{O}_i^s and $\mathcal{O}_j^d(t)$. The objective is to generate a collision-free sequence of state $\mathbf{x}_t \in [t_0, t_f]$ while maintaining state feasibility




Fig. 3: Schematic of vehicle motion model

Furthermore, the spatio-temporal constraints of dynamic obstacles hinder most of the motion planning approaches to reach a feasible trajectory to meet real-time constraints. Ensuring these constraints are respected is fundamental in safety-critical situations, where the autonomous vehicle must be able to anticipate in a timely manner to prevent a potential collision. MPPI offers an efficient way to solve (4). MPPI can be used to solve optimization problems subject to nonlinear dynamics and non-convex constraints. The main idea of MPPI is to randomly sample a huge number of M trajectories generated by Monte Carlo simulations of the vehicle state model (1a) in real-time. These trajectories are evaluated in *parallel* by computing the cost-to-go from t_0 to t_f along each trajectory. For a given sample period Δt , the horizon $t \in [t_0, t_f]$ is discretized into T samples. Let $S(\tau_i)$ denotes the cost of trajectory τ_i , $\forall i = 1, 2, \dots, M$.

$$S(\tau_i) = \pi(\mathbf{x}_T, \mathbf{x}_T^r) + \sum_{t=0}^{T-1} \tilde{q}(\mathbf{x}_t, \mathbf{x}_t^r, \mathbf{u}_t, \epsilon_t), \quad (6)$$

where $\pi(\mathbf{x}_T, \mathbf{x}_T^r)$ is referred to the terminal cost with the objective to force the final trajectory state \mathbf{x}_T to be as close




Fig. 4: A perception obstacle \mathcal{O}_i approximated by circles $\mathcal{C}_k = \{x_k, y_k, r_k\}, \forall k = 1, 2, 3$.

as to a selected target point \mathbf{x}_T^r . The term $\tilde{q}(\mathbf{x}_t, \mathbf{x}_T^r, \mathbf{u}_t, \epsilon_t)$ is referred to as the running cost with the goal to satisfy the constraints of the planning problem as defined in (4). For the sake of simplicity, one can decouple the running cost into two components

$$\tilde{q}(\mathbf{x}_t, \mathbf{x}_T^r, \mathbf{u}_t, \epsilon_t) = \tilde{q}_x(\mathbf{x}_t, \mathbf{x}_T^r) + \tilde{q}_u(\mathbf{u}_t, \epsilon_t) \quad (7)$$

where $\tilde{q}_x(\mathbf{x}_t, \mathbf{x}_T^r)$ considers all costs that depend on the system state \mathbf{x}_t and reference state \mathbf{x}_T^r trajectories, and $\tilde{q}_u(\mathbf{u}_t, \epsilon_t)$ is typically quadratic cost to optimize the input sequence

$$\tilde{q}_u(\mathbf{u}_t, \epsilon_t) = \alpha \epsilon_t^\top R \mathbf{u}_t + \mathbf{u}_t^\top R \epsilon_t + \frac{1}{2} \mathbf{u}_t^\top R \mathbf{u}_t \quad (8)$$

with $R \in \mathbb{R}^{m \times m}$ is a symmetric positive semi-definite input weight matrix and α is a parameter that is used to prioritize noise against optimizing inputs, thus controlling the way the state space is explored. It is typically defined as

$$\alpha = \frac{\gamma - 1}{2\gamma}, \quad \gamma \in \mathbb{R}^+ \quad (9)$$

with γ being a positive hyperparameter which is given as an input to the MPPI algorithm. On the other hand, the state-dependent cost $\tilde{q}_x(\mathbf{x}_t, \mathbf{x}_T^r)$ encapsulate all costs needed to force the MPPI to respect the constraints of obstacles and follow the reference trajectory as stated in (2) and (3). For the generated trajectory \mathbf{x}_t to be as close as possible to the reference state trajectory \mathbf{x}_T^r , a weighted quadratic cost of the error between \mathbf{x}_t and \mathbf{x}_T^r is considered. We may then express \tilde{q}_x as

$$\tilde{q}_x = \frac{1}{2} (\mathbf{x}_t - \mathbf{x}_T^r)^\top H (\mathbf{x}_t - \mathbf{x}_T^r) + \tilde{q}_o \quad (10)$$

where $H \in \mathbb{R}^{n \times n}$ is a symmetric positive definite weight matrix and \tilde{q}_o is the cost to avoid obstacles. Our methodology to consider obstacles in MPPI is to approximate the occupied area by each obstacle \mathcal{O}_i by a finite set of circles $\mathcal{C}_k = \{x_k, y_k, r_k\}$ as illustrated in Fig. 4, where x_k, y_k are the coordinates of centre and r_k is the radius. The cost \tilde{q}_o is considered such that distance between (x_k, y_k) and (x_i, y_i) is greater than the sum the radius r_k of a safety margin d_k . This safety margin is used to anticipate the velocity of the vehicle. More details are given in Section IV.

A. Algorithm

Algorithm 1 sketches a pseudocode of the MPPI motion planner. The algorithm starts by reading the vehicle state and the list of perception objects from the localization and perception modules of the ADS, respectively. Perception objects are then approximated by circles. Our MPPI implementation checks the generated noise to guarantee that a set of constraints are never exceeded, namely: $\omega_{max}, a_{max}, a_{min}$, and the target speed v_G . The constraints provide safety and feasibility to the trajectory. Given the cost-to-go $S(\tau_i)$ and an initial input sequence $\mathbf{U} = [\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{T-1}]^\top$, one can update the input sequence iteratively using [13], [18] taking into account all individual costs contributing to the cost-to-go the cost-to-go $S(\tau_i)$

$$\mathbf{u}_t \leftarrow \mathbf{u}_t + \frac{\sum_{m=0}^{M-1} \exp\left(\frac{-1}{\lambda} [\tilde{S}(\tau_{t,m}) - \tilde{S}_{\min}]\right) \epsilon_{t,m}}{\sum_{m=0}^{M-1} \exp\left(\frac{-1}{\lambda} [\tilde{S}(\tau_{t,m}) - \tilde{S}_{\min}]\right)} \quad (11)$$

where \tilde{S}_{\min} is mainly used to guarantee the numerical stability of the algorithm while keeping the optimality conditions of the algorithm. Practically the term \tilde{S}_{\min} is considered as the cost of the trajectory with the minimum cost-to-go value among the set of sampled trajectories [11], [12]. A final step in the MPPI algorithm is to smooth the resulting input sequence using a Savitzky–Golay filter [22] with coefficients Z in a similar way as [11], [12]. Unlike methods that use MPPI to control the dynamic of a system directly by applying the first input in the optimal sequence to the system, for trajectory planning, we concatenate the optimal input sequence together with their corresponding optimal state to constitute the generated trajectory $\{\mathbf{x}_t^*, \mathbf{u}_t^*\}$, which subsequently is fed to the control module of the ADS.

IV. EXPERIMENTAL RESULTS

The MPPI-based motion planner has been implemented and tested on an existing autonomous driving platform [23]. Three driving scenarios were tested on a real track using virtual objects, namely a lane merge, an object avoidance and a vehicle following. Those scenarios were selected as they form primitive functions to a more elaborate motion planner.

A. Implementation Details

The implementation has been done in C++ and integrated into an existing in-house autonomous driving stack. The MPPI being a method relying heavily on parallel computing, the implementation uses OpenCL, which provides the flexibility of using any compatible GPU. In the context of this paper, the MPPI was run on an Intel integrated GPU hinting an efficient resource usage by relieving the main GPU for other tasks that could be involved in autonomous driving.

B. Driving Scenarios

Three driving scenarios were assessed. Those scenarios were deemed relevant as they represent building blocks of what could be a higher-level motion planner.

Algorithm 1 Real-Time MPPI Trajectory Generation

Given:

M : Size of trajectory rollouts
 T : Horizon timesteps
 \mathbf{U} : Initial input sequence
 \mathcal{F} : Vehicle dynamic model
 \mathbf{x}_t^r : Reference state
 π, ϕ : Cost functions
 $\lambda, \gamma, \Sigma, R$: hyper-parameters
 D : safety distance parameters
 Z : Savitzky–Golay filter parameters

```
1: while not done do
2:    $\mathbf{x}_0 \leftarrow \text{GetInitialState}(), \mathbf{x}_0 \in \mathbb{R}^n$ 
3:    $\mathcal{O}_i^s, \mathcal{O}_j^d \leftarrow \text{GetObstacles}(), \mathcal{O}_i^s, \mathcal{O}_j^d \subset \mathcal{X}, \forall i, j$ 
4:    $\mathbf{C}_k \leftarrow \text{ApproximateObstacles}(\mathcal{O}_i^s, \mathcal{O}_j^d), \mathbf{C}_k \in \mathcal{C}, \forall k,$ 
5:    $\epsilon \leftarrow \text{GenerateRandomNoise}(\Sigma), \epsilon \in \mathbb{R}^{M \times T}$ 
6:    $\tilde{S}(\tau_m) \leftarrow \text{GetInitialTrajectoryCost}(), \tilde{S}(\tau_m) \in \mathbb{R}^M$ 
7:   for  $m \leftarrow 0$  to  $M - 1$  in parallel do
8:      $\mathbf{x} \leftarrow \mathbf{x}_0$ 
9:     for  $t \leftarrow 0$  to  $T - 1$  do
10:       $\mathbf{x}_{t+1} \leftarrow \mathcal{F}(\mathbf{x}_t, \mathbf{u}_t + \epsilon_t)$ 
11:       $\tilde{q} \leftarrow \tilde{q}(\mathbf{x}_t, \mathbf{x}_t^r, \mathbf{u}_t, \epsilon_t, \mathbf{C}_k, D, R, \gamma)$ 
12:       $\tilde{S}(\tau_{t+1, m}) \leftarrow \tilde{S}(\tau_{t, m}) + \tilde{q}$ 
13:    end for
14:     $\tilde{S}(\tau_m) \leftarrow \tilde{S}(\tau_{t+1, m}) + \phi(\mathbf{x}_T)$ 
15:  end for
16:   $\tilde{S}_{\min} \leftarrow \min_m[\tilde{S}(\tau_m)]$ 
17:  for  $t \leftarrow 0$  to  $T - 1$  do
18:     $\mathbf{u}_t \leftarrow \mathbf{u}_t + \frac{\sum_{m=0}^{M-1} \exp\left(\frac{-1}{\lambda} [\tilde{S}(\tau_{t, m}) - \tilde{S}_{\min}]\right) \epsilon_{t, M}}{\sum_{m=0}^{M-1} \exp\left(\frac{-1}{\lambda} [\tilde{S}(\tau_{t, m}) - \tilde{S}_{\min}]\right)}$ 
19:  end for
20:   $\mathbf{u}^* \leftarrow \text{SavitzkyGolayFilter}(Z, \mathbf{u})$ 
21:  for  $t \leftarrow 0$  to  $T - 1$  do
22:     $\mathbf{x}_{t+1}^* \leftarrow \mathcal{F}(\mathbf{x}_t^*, \mathbf{u}_t^*)$ 
23:  end for
24:  for  $t \leftarrow 0$  to  $T - 2$  do
25:     $\mathbf{u}_t = \mathbf{u}_{t+1}^*$ 
26:  end for
27:  if completed return optimal trajectory  $\{\mathbf{x}_t^*, \mathbf{u}_t^*\}$ 
28: end while
```

1) *Lane Merge*: This scenario assesses the ability of the vehicle to switch lanes smoothly, a recurrent situation in everyday driving. The vehicle starts from a lane that is parallel to the target lane and has to properly merge into it.

2) *Object Avoidance*: The goal of this scenario is to assess the ability of the motion planner to generate a trajectory that avoids an object or an obstacle obstructing the current lane. A typical real-world occurrence would be a construction site or a stopped vehicle.

3) *Vehicle Following*: This scenario is similar to object avoidance with the added constraint that the vehicle cannot leave the current lane thus forcing it to adjust its speed to keep a safe distance or even stop if the obstacle is not moving.

C. Cost Formulation

The cost function is broken down into five parts each of them pertaining to a component of the autonomous driving task. Let $P_t = (x_t, y_t)$ be the position of the vehicle and θ_t and v_t be respectively the yaw and speed of the vehicle at timestep t . Let $W_t = (x_{W_t}, y_{W_t})$ be the closest waypoint and θ_t^r and v_t^r be the reference yaw and speed. The c_{dist} cost measures the distance between the vehicle and W_t thus penalising diverging trajectories.

$$c_{dist} = \|(P_t - W_t)\|^2 \quad (12)$$

Let T_P be the target point, the c_{target} cost penalises trajectories that do not get closer to the target point, thus preventing the vehicle from going backwards.

$$c_{target} = \begin{cases} 1 & \text{if } \|(P_t - T_P)\| > \|(P_{t-1} - T_P)\| \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

In order to ensure a smooth lane merge, the vehicle needs to stay sufficiently parallel to the target lane while closing in. The c_{yaw} cost measures the difference between the current yaw and the reference yaw, both wrapped in $[-\pi, \pi]$.

$$c_{yaw} = |\theta_t - \theta_t^r|^2 \quad (14)$$

This cost also proved to be helpful in following a curved path. The c_{speed} cost measures the difference between the current speed and the reference speed.

$$c_{speed} = |v_t - v_t^r|^2 \quad (15)$$

MPPI objects are represented by circles of centers O_t^k and radius r_t^k . Let d_{obj} be the minimal object distance and d_{safe} the desired safe distance:

$$\begin{cases} d_{obj} = \min_{\forall k} (\|P_t - O_t^k\| - r_t^k) \\ d_{safe} = d_{safe_c} v_t + d_{safe_0} \end{cases} \quad (16)$$

where d_{safe_c} is the safe distance coefficient and d_{safe_0} is the minimal safe distance. The c_{safe_dist} cost is formulated as follow:

$$c_{safe_dist} = \max(d_{safe} - d_{obj}, 0)^2 \quad (17)$$

In the case of the object avoidance scenario, the cost is only considered if there is a collision i.e $d_{obj} \leq 0$, otherwise c_{safe_dist} is set to 0. The total step cost is a weighted sum of all sub costs.

$$c_{total} = 15.0c_{min_dist} + 7.0c_{target} + 120.0c_{yaw} + 5.0c_{speed} + 25.0c_{safe_dist} \quad (18)$$

D. Experimental Setup

A total of $M = 2560$ trajectories were generated in parallel. The trajectory horizon is set to 4 seconds divided into $T = 16$ steps of $\Delta t = 0.25$ seconds. λ was set to 150 while the noise variance was $\Sigma = \text{Diag}(\sigma_\omega, \sigma_a) = \text{Diag}(0.05, 0.85)$. These parameters were obtained through experiments. The maximum steering rate was $\omega_{max} = 0.11\text{rad/s}$ while the maximum and minimum accelerations were $a_{max} = 1.1\text{m/s}^2$ and $a_{min} = -2.5\text{m/s}^2$. The constraints on ω and a help to

Parameter	Value	Unit
λ	150	
M	2560	
T	16	
Δt	0.25	s
σ_ω	0.05	rad/s
σ_a	0.85	m/s ²
ω_{max}	0.11	rad/s
a_{max}	1.1	m/s ²
a_{min}	-2.5	m/s ²
d_{safe_c}	1.36	
d_{safe_0}	11	m
v_G	30	km/h

TABLE I: MPPI parameters

ensure a safe and feasible trajectory. A target speed $v_G = 30\text{km/h}$ was set along the path. The safe distance coefficient was set to $d_{safe_c} = 1.36$ and the minimum safe distance to $d_{safe_0} = 11$. This corresponds approximately to a safe distance of 18 meters at 30km/h accounting for the length of the vehicle. The parameters are summarized on Table I. A Savitzky-Golay filter was also applied on the generated input \mathbf{u}_t and is shown in (19) where \mathbf{u}_t^* is the filtered input.

$$\mathbf{u}_t^* = \frac{1}{35}(-3\mathbf{u}_{t-2} + 12\mathbf{u}_{t-1} + 17\mathbf{u}_t + 12\mathbf{u}_{t+1} - 3\mathbf{u}_{t+2}) \quad (19)$$

The MPPI was run at 20Hz on an Intel i7-8700T and its integrated UHD Graphics 630. The tests were conducted on a real track using a modified KIA Soul EV equipped for autonomous driving as shown on Figure 1. The objects required by the scenarios were obtained by the simulator module of the ADS running the MPPI. This module provides programmable virtual objects that can be tailored for specific use cases. The use of virtual objects was motivated by the need of a safe and reliable mean to assess the motion planner. The objects dimensions were set to the ones of an average car for both the object avoidance and object following scenarios. In the case of the object avoidance, a relatively tight margin of 0.7m was prescribed between the vehicle and the object while overtaking.

E. Results

For each scenario, the generated trajectory and the vehicle state were recorded along the path. Snapshots of the results are depicted on Figures 5, 6 and 7. The path is represented by the dotted line, the ego vehicle and the virtual object are represented by a red and blue box respectively. The green line represents the generated MPPI trajectory at the given vehicle position. In addition, the realised steering, speed and acceleration are provided for each scenario as to better portray the actual behaviour of the vehicle over the course of the tests. The produced trajectories are smooth and they respect the prescribed acceleration and speed as this was guaranteed by this MPPI implementation. As it can be seen on the realised speed of Figures 5, 6 and 7, the vehicle converged toward the target speed of 30km/h and never exceeded it. The realised steering also shows a smooth operation for all three

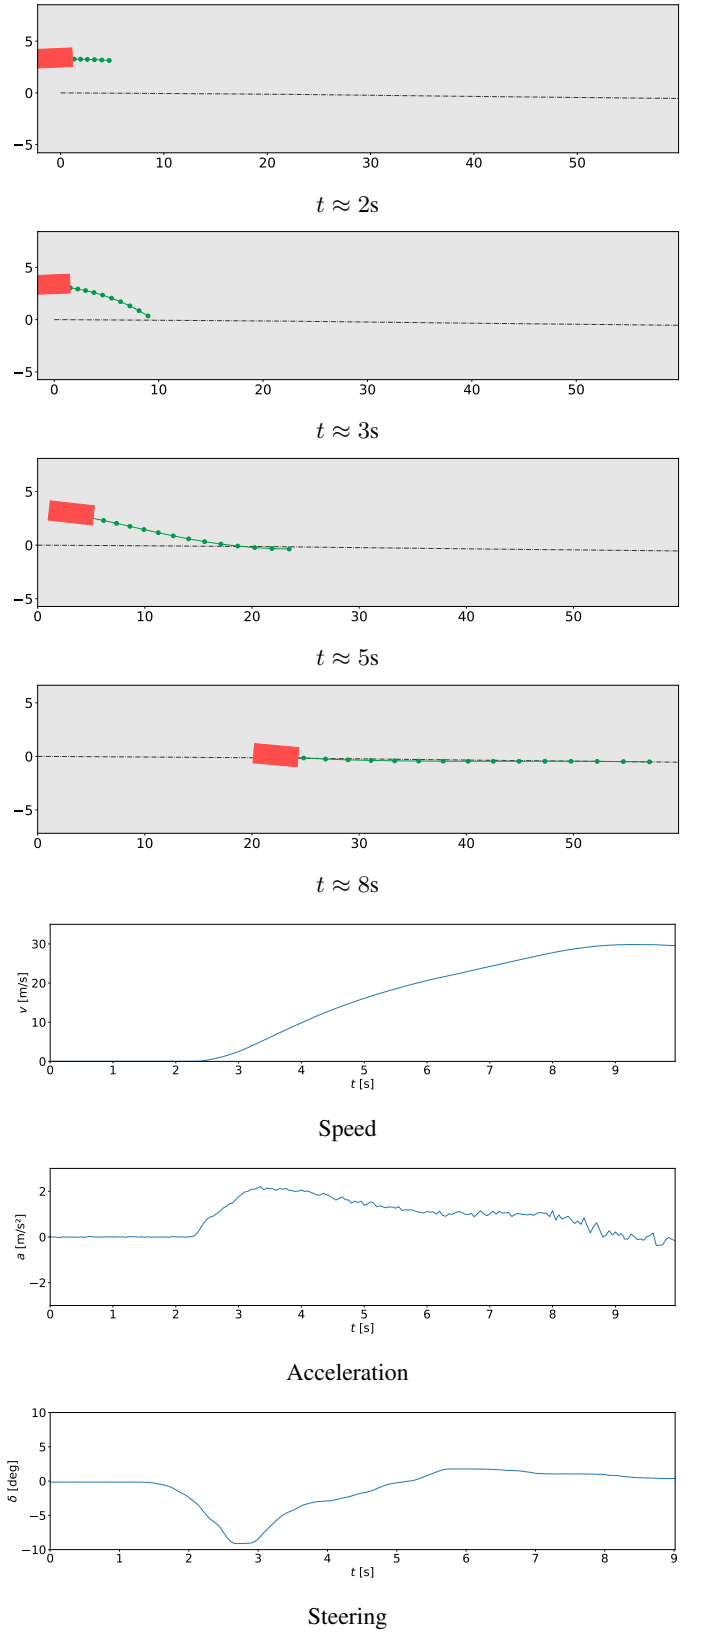


Fig. 5: Lane merge scenario.

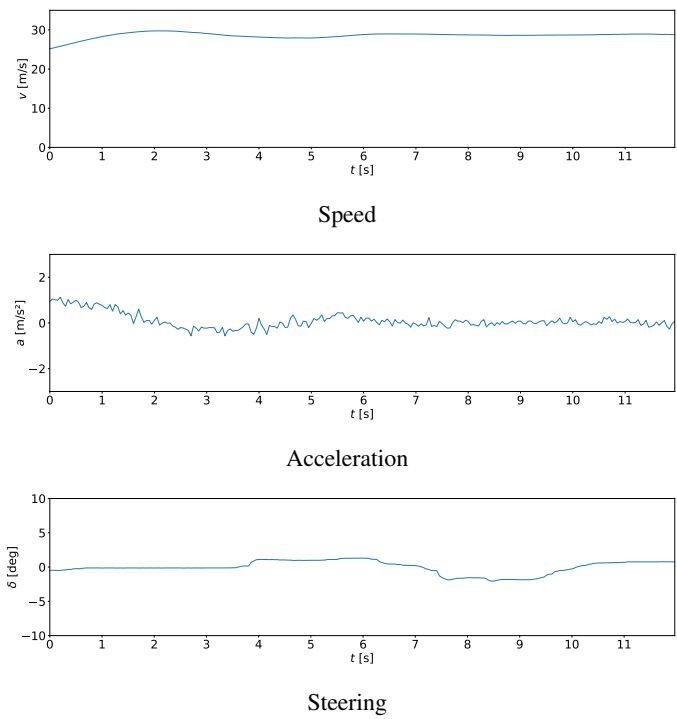



Fig. 6: Object avoidance scenario.

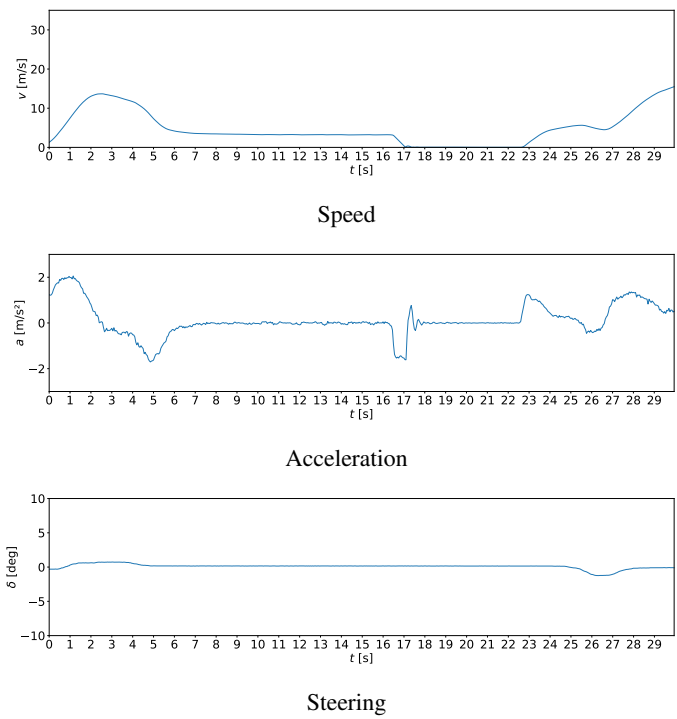
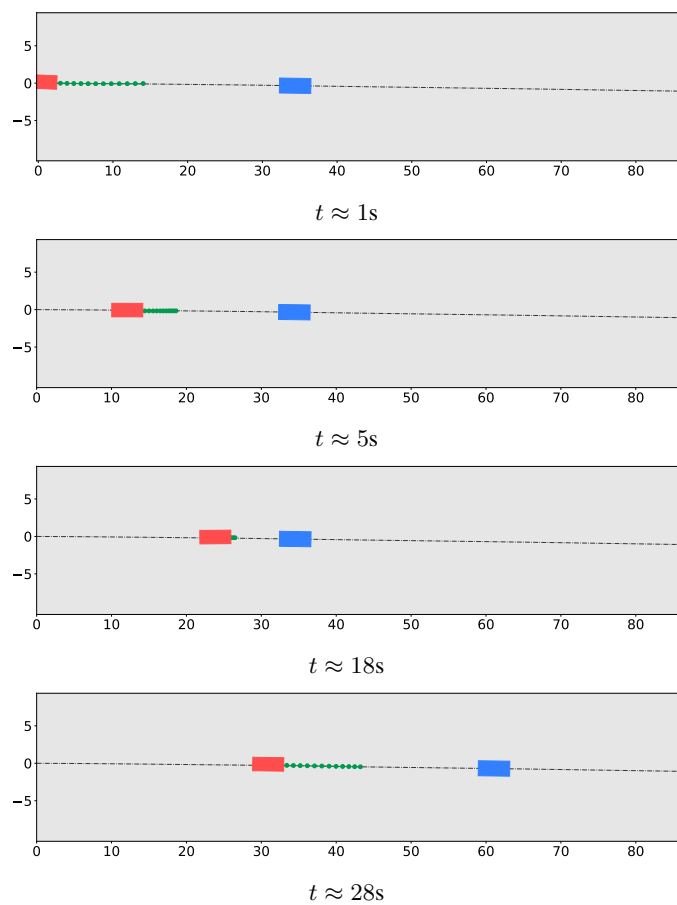


Fig. 7: Object following scenario.

scenarios and never exceeded 10 degrees. The trajectories are proven safe as they not only avoid objects but also maintain a safe distance, and no collision occurred during the tests. It is important to note, however, that while the constraints on acceleration and steering rate ensure that the generated trajectory remains fairly easy to follow, and thus feasible, this cannot be guaranteed as the bicycle model used constitutes an estimation of the vehicle dynamics. The scope of this paper was to use the MPPI as a motion planner and therefore another module was in charge of the control of the vehicle. This module also has imperfections which could translate into a realised trajectory deviating from the target one. This happened during the tests and by looking at the realised acceleration on Figures 5 and 7, it can be seen that the initial acceleration from a standstill exceeded the maximum value of a_{max} by reaching almost 2. The minimum acceleration a_{min} , however, was never exceeded.

V. CONCLUSION AND FUTURE WORKS

This paper introduces a real-time and safe MPPI-based motion planner for autonomous driving. The method can handle obstacles and guarantees bounds for speed, acceleration and steering rate. Three primitive driving scenarios were proposed and tested on a real track at a maximum of speed of 30km/h. The results show that the method can perfectly handle those base scenarios and it achieved a smooth and accurate ride overall but also avoided collisions and kept a safe distance from obstacles. A slight deviation between the MPPI trajectories and the actual trajectories, especially for the acceleration, was observed and the underlying control module could be at fault. Future works should focus on more complex and realistic scenarios involving a more dynamic environment such as a lane change between moving vehicle. More work should also be done to better model the dynamics of the vehicle in the MPPI motion model.

ACKNOWLEDGMENT

The authors would like to thank the Centre de Formation pour Conducteurs, Luxembourg for providing us access their training track. The authors would also like to thank Ihab S. Mohamed, Vehicle Autonomy and Intelligence Lab, Indiana University for the valuable discussion about MPPI.

REFERENCES

- [1] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on intelligent vehicles*, vol. 1, no. 1, pp. 33–55, 2016.
- [2] S. Manzi, C. Pek, and M. Althoff, "Using reachable sets for trajectory planning of automated vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 6, no. 2, pp. 232–248, 2020.
- [3] D. González, J. Pérez, V. Milanés, and F. Nashashibi, "A review of motion planning techniques for automated vehicles," *IEEE Transactions on intelligent transportation systems*, vol. 17, no. 4, pp. 1135–1145, 2015.
- [4] L. Zhang, R. Zhang, T. Wu, R. Weng, M. Han, and Y. Zhao, "Safe reinforcement learning with stability guarantee for motion planning of autonomous vehicles," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 12, pp. 5435–5444, 2021.
- [5] K. Rezaee, P. Yadmellat, and S. Chamorro, "Motion planning for autonomous vehicles in the presence of uncertainty using reinforcement learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3506–3511, IEEE, 2021.
- [6] M. Xu, Z. Liu, P. Huang, W. Ding, Z. Cen, B. Li, and D. Zhao, "Trustworthy reinforcement learning against intrinsic vulnerabilities: Robustness, safety, and generalizability," *arXiv preprint arXiv:2209.08025*, 2022.
- [7] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-based collision avoidance," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 3, pp. 972–983, 2020.
- [8] J. Ziegler, P. Bender, T. Dang, and C. Stiller, "Trajectory planning for bertha—a local, continuous method," in *2014 IEEE intelligent vehicles symposium proceedings*, pp. 450–457, IEEE, 2014.
- [9] W. Ding, L. Zhang, J. Chen, and S. Shen, "Safe trajectory generation for complex urban environments using spatio-temporal semantic corridor," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2997–3004, 2019.
- [10] Y. Yoon, J. Shin, H. J. Kim, Y. Park, and S. Sastry, "Model-predictive active steering and obstacle avoidance for autonomous ground vehicles," *Control Engineering Practice*, vol. 17, no. 7, pp. 741–750, 2009.
- [11] I. S. Mohamed, J. Xu, G. Sukhatme, and L. Liu, "Towards efficient MPPI trajectory generation with unscented guidance: U-MPPI control strategy," *arXiv preprint arXiv:2306.12369*, 2023.
- [12] I. S. Mohamed, K. Yin, and L. Liu, "Autonomous navigation of AGVs in unknown cluttered environments: log-MPPI control strategy," *IEEE Robot. and Automat. Lett.*, vol. 7, no. 4, pp. 10240–10247, 2022.
- [13] G. Williams, A. Aldrich, and E. A. Theodorou, "Model predictive path integral control: From theory to parallel computation," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [14] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [15] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *IEEE Int. Conf. on Robot. and Automat. (ICRA)*, pp. 1433–1440, 2016.
- [16] G. Williams, B. Goldfain, P. Drews, K. Saigol, J. M. Rehg, and E. A. Theodorou, "Robust sampling based model predictive control with sparse objective information," in *Robotics: Science and Systems*, (Pittsburgh, Pennsylvania, USA), pp. 42–51, 2018.
- [17] A. Buyval, A. Gabdullin, K. Sozkyin, and A. Klimchik, "Model predictive path integral control for car driving with autogenerated cost map based on prior map and camera image," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 2109–2114, IEEE, 2019.
- [18] I. S. Mohamed, G. Allibert, and P. Martinet, "Model predictive path integral control framework for partially observable navigation: A quadrotor case study," in *16th Int. Conf. on Control, Automation, Robotics and Vision (ICARCV)*, (Shenzhen, China), pp. 196–203, Dec. 2020.
- [19] I. S. Mohamed, G. Allibert, and P. Martinet, "Sampling-based MPC for constrained vision based control," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pp. 3753–3758, 2021.
- [20] L. A. S. Guardini, A. Spalanzani, P. Martinet, C. Laugier, T. Genevois, and A.-L. Do, "Minimal injury risk motion planning using active mitigation and sampling model predictive control," in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1262–1267, IEEE, 2022.
- [21] G. Elghazaly, R. Frank, S. Harvey, and S. Saffo, "High-definition maps: Comprehensive survey, challenges and future perspectives," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 4, pp. 527–550, 2023.
- [22] A. Savitzky and M. J. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.
- [23] G. Varisteas and R. Frank, "Poster: Junior, a research platform for connected and automated driving," in *2019 IEEE Vehicular Networking Conference (VNC)*, pp. 1–2, 2019.