

Trajectory Distribution Control for Model Predictive Path Integral Control using Covariance Steering

Ji Yin, Zhiyuan Zhang, Evangelos Theodorou, Panagiotis Tsiotras

Abstract—This paper presents a novel control approach for autonomous systems operating under uncertainty. We combine Model Predictive Path Integral (MPPI) control with Covariance Steering (CS) theory to obtain a robust controller for general nonlinear systems. The proposed Covariance-Controlled Model Predictive Path Integral (CC-MPPI) controller addresses the performance degradation observed in some MPPI implementations owing to unexpected disturbances and uncertainties. Namely, in cases where the environment changes too fast or the simulated dynamics during the MPPI rollouts do not capture the noise and uncertainty in the actual dynamics, the baseline MPPI implementation may lead to divergence. The proposed CC-MPPI controller avoids divergence by controlling the dispersion of the rollout trajectories at the end of the prediction horizon. Furthermore, the CC-MPPI has adjustable trajectory sampling distributions that can be changed according to the environment to achieve efficient sampling. Numerical examples using a ground vehicle navigating in challenging environments demonstrate the proposed approach.

I. INTRODUCTION

As autonomous vehicles and other robots become increasingly popular, one of the major concerns is whether humans can trust the robots' ability to complete their assigned tasks safely. For autonomous vehicles, for example, neglecting or misinterpreting the disturbances during driving can lead to serious consequences within milliseconds. The demand for safety has led many researchers develop robust control and planning algorithms for autonomous robotic systems. For example, sampling-based planning algorithms that consider uncertainty and collision probability in the vertex selection or evaluation processes have been proposed in [1], [2], [3], [4], and optimization-based planning algorithms that consider systematic disturbances and chance constraints explicitly by solving optimization problems have been developed in [5], [6], [7], [8].

In this paper, we propose an MPC-based robust trajectory planning approach that deals with environmental and plant uncertainty, while providing guarantees on the dispersion of the *closed-loop* system future trajectories. Model Predictive Control (MPC) is an algorithmic, optimization-based control design [9] that has gained popularity for autonomous vehicle control over the past years [10], [11]. The traditional deterministic MPC approaches are model-based and generate trajectories assuming there are no uncertainties in the dynamics. As a result, MPC controllers are typically not robust to model parameter variations. To improve the performance of MPC controllers by taking system uncertainty into account,

Robust MPC (RMPC) controllers have been proposed to handle deterministic uncertainties residing in a given compact set. RMPC generates control commands by considering worst-case scenarios, thus the resulting trajectories can be conservative. Reference [12] provides an extensive review summarizing all types of RMPC controllers. To achieve more aggressive planning, Stochastic MPC (SMPC) utilizes the probabilistic nature of the system uncertainty to account for the most likely disturbances, instead of considering only the worst-case disturbance, as with the RMPC [13], [14]. There are two classes of SMPC approaches in the literature. The first one is based on the analytical solutions of some optimization problem, such as [15], [16], [17], while the second approach relies on randomization to solve optimization problems, such as [18], [19], [20]. The proposed CC-MPPI controller is somewhere in-between these two, as it analytically computes a controlled dynamics by considering the model uncertainty and then generates the optimal control using randomized roll-outs of the controlled dynamics. This is discussed in greater detail in Section IV.

Most of current MPC implementations assume linear system dynamics and formulate the resulting MPC task as a quadratic optimization problem, which helps MPC meet the strict real-time requirements required for safe control. However, these approaches depend on simplified linear models that may not capture accurately the dynamics of the real system. Model Predictive Path Integral (MPPI) control [21] is a type of nonlinear MPC algorithm that solves repeatedly finite-horizon optimal control tasks while utilizing nonlinear dynamics and general cost functions. Specifically, MPPI is a simulation-based algorithm that samples thousands of trajectories around some mean control sequence in real-time, by taking advantage of the parallel computing capabilities of modern Graphic Processing Units (GPUs). It then produces an optimal trajectory along with its corresponding control sequence by calculating the weighted average of the cost of the ensuing sampled trajectories, where the weights are determined by the cost of each trajectory rollout. One of the advantages of the MPPI approach over more traditional MPC controllers is that it does not restrict the form of the cost function of the optimization problem [22], which can be non-quadratic and even discontinuous.

Despite its appealing characteristics, the MPPI algorithm may encounter problems when implemented in practice. In particular, when the mean control sequence lies inside an infeasible region, all the resulting MPPI sampled trajectories are concentrated within the same region, as illustrated in Figure 1, and this may lead to a situation where the trajectories violate the constraints. Two cases this may happen are: first,

The authors are with the D. Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, GA. E-mail: {jyin81, nickzhang, evangelos.theodorou, tsiotras}@gatech.edu




Fig. 1. MPPI Divergence; from [23].

when the MPPI algorithm diverges because the environment changes too fast; and, second, when the algorithm fails because the predicted dynamics do not capture the noise and uncertainty of the actual dynamics. The reason the MPPI algorithm may perform poorly under the previous two cases is because it fails to take into account the disturbances (either from the dynamics or from the environment) so that all sampled trajectories end up violating the constraints. Figure 1 shows the influence of the noise on MPPI sampled trajectories. In this figure, the gray curves are the MPPI sampled trajectories, the red curves show the boundaries of the trajectory sampling distribution, and the green curve represents the simulated trajectory of the robot following the optimal control sequence given the current distribution. In Figure 1(a) the autonomous vehicle has sampling distribution mostly inside the track initially. In Figure 1(b), the vehicle ends up in an unexpected pose due to unforeseen disturbances after it executes the control command. This further leads to the situation depicted in Figure 1(c), where the algorithm diverges because all of the sampled trajectories violate the constraints.

To mitigate the previous shortcomings of the MPPI algorithm, prior works apply a controller to track the output of the MPPI controller in order to keep the actual trajectory as close as possible to the predicted nominal trajectory. These approaches separate the planning and control tasks so that MPPI acts similarly to a path planner. For example, in [23] an iterative Linear Quadratic Gaussian (iLQG) controller was used to track the planned trajectory provided by MPPI. In [24] the authors propose a method that utilizes a tracking controller with L_1 augmentation to compensate for the mismatch between the nominal dynamics and the true dynamics. However, these methods do not improve the performance of the MPPI algorithm if there are significant changes in the environment within a short interval of time. The proposed CC-MPPI algorithm tries to address some of these shortcomings by improving the performance of the MPPI algorithm under the scenarios mentioned above. This is achieved by introducing adjustable trajectory sampling distributions, and by directly controlling the evolution of these trajectory distributions to avoid an uncontrolled dispersion at the end of the control horizon.

II. PROBLEM FORMULATION

The goal of the proposed Covariance-Controlled MPPI (CC-MPPI) controller is to make the distributions of the sampled trajectories more flexible than the ones generated by MPPI, such that the CC-MPPI algorithm samples more

efficiently and with a smaller probability to be trapped in local minima when the optimal trajectory from the previous time step lies inside some high-cost region, as illustrated in Figure 1(c). To this end, we introduce a desired terminal state covariance Σ_f for the states of the dynamics (1b) at the final time step N as a hyperparameter for the CC-MPPI controller. The key idea is that the distribution of the sampled trajectories can be adjusted by a suitable choice of Σ_f together with the control disturbance variance Σ_ϵ . Denoting by $\mathbf{v} = [v_0, \dots, v_{N-1}]^\top \in \mathbb{R}^{n_u N}$, the CC-MPPI controller solves the following optimization problem,

$$\min_{\mathbf{v}} J(\mathbf{v}) = \mathbb{E} \left[\phi(x_N) + \sum_{k=0}^{N-1} \left(q(x_k) + \frac{1}{2} v_k^\top R v_k \right) \right], \quad (1a)$$

subject to,

$$x_{k+1} = F(x_k, v_k + \epsilon_k), \quad (1b)$$

$$x(0) = x_0, \quad \epsilon_k \sim \mathcal{N}(0, \Sigma_\epsilon), \quad (1c)$$

$$\mathbb{E}[x_N x_N^\top] - \mathbb{E}[x_N] \mathbb{E}[x_N]^\top \preceq \Sigma_f, \quad (1d)$$

at each iteration, where the state terminal cost $\phi(x_N)$ and the state portion of the running cost $q(x_k)$ can be arbitrary functions. The objective function (1a) minimizes the expectation of the state and control costs with x_k being a random vector subject to the dynamics (1b).

III. MPPI ALGORITHM REVIEW

The MPPI controller, as described in [22], minimizes (1a) subject to (1b) and (1c). As in Problem (1), the terminal cost $\phi(x_N)$ and the state portion of the running cost $q(x_k)$ of the MPPI can be arbitrary functions.

As in an MPC setting, the MPPI algorithm samples M trajectories during each optimization iteration. Let $\mathbf{v} = [v_0, \dots, v_{N-1}]^\top \in \mathbb{R}^{n_u N}$ be the mean control sequence, $\mathbf{u}^{(m)} = [u_0^{(m)}, \dots, u_{N-1}^{(m)}]^\top \in \mathbb{R}^{n_u N}$ be the actual control sequence, and let $\epsilon^{(m)} = [\epsilon_0^{(m)}, \dots, \epsilon_{N-1}^{(m)}]^\top \in \mathbb{R}^{n_u N}$, $\epsilon_k^{(m)} \sim \mathcal{N}(0, \Sigma_\epsilon)$, be the control disturbance sequence corresponding to the m^{th} sampled trajectory $[x_0^{(m)}, \dots, x_{N-1}^{(m)}]$ at the current iteration, such that $\mathbf{u}^{(m)} = \mathbf{v} + \epsilon^{(m)}$, where $m = 1, \dots, M$. The cost for the m^{th} sampled trajectory is given by [20]

$$S_m = \phi(x_N^{(m)}) + \sum_{k=0}^{N-1} \tilde{q}(x_k^{(m)}, v_k, \epsilon_k^{(m)}), \quad (2)$$

where $\tilde{q}(x_k^{(m)}, v_k, \epsilon_k^{(m)})$ is the cost for the m^{th} sampled trajectory at the k^{th} step, given by

$$\tilde{q}(x_k^{(m)}, v_k, \epsilon_k^{(m)}) = q(x_k^{(m)}) + \frac{1 - \nu^{-1}}{2} \epsilon_k^{(m)\top} R \epsilon_k^{(m)} + v_k^{(m)\top} R \epsilon_k^{(m)} + \frac{1}{2} v_k^{(m)\top} R v_k^{(m)}, \quad (3)$$

where ν is the ratio between the covariance of the injected disturbance $\epsilon^{(m)}$ and the covariance of the disturbance of the original dynamics. The term $\frac{1}{2} v_k^{(m)\top} R v_k^{(m)}$ in (3) is the cost for the disturbance-free portion of the control input, and

both $\frac{1-\nu^{-1}}{2}\epsilon_k^{(m)\top} R\epsilon_k^{(m)}$ and $v_k^{(m)\top} R\epsilon_k^{(m)}$ in (3) penalize large control disturbances and smooth out the resulting control signal. The weights ω_m ($m = 1, \dots, M$) of the m^{th} sampled trajectory are chosen as [25]

$$\omega_m = \exp\left(-\frac{1}{\lambda}(S_m - \beta)\right), \quad (4)$$

where, $\beta = \min_{m=1, \dots, M} S_m$, and where λ determines how selective is the weighted average of the sampled trajectories. Note that the constant β does not influence the solution, and it is introduced to prevent numerical instability of the algorithm. The MPPI algorithm generates the optimal control sequence and the mean sequence of the next iteration \mathbf{v} using the following equations,

$$\mathbf{v} = \sum_{m=1}^M \omega_m \mathbf{u}^{(m)} / \sum_{m=1}^M \omega_m. \quad (5)$$

In Section IV we discuss how the CC-MPPI controller satisfies the terminal constraint (1d), and then present the complete CC-MPPI algorithm.

IV. COVARIANCE-CONTROLLED MPPI

In this section, we introduce the proposed CC-MPPI controller. Section IV-A discusses the linearization of the dynamics (1b), and Section IV-B uses this linearization to achieve the terminal constraint (1d). Section V presents the proposed CC-MPPI algorithm that solves the optimization Problem (1).

A. Linearized Model

We start by linearizing system (1b) along some reference trajectory using the approach outlined in [26]. The reference trajectory of the first optimization iteration can be an arbitrary trajectory; starting with the second iteration, the reference trajectory of the current iteration is the trajectory generated by the optimal control sequence from the previous iteration. To this end, let $\mathbf{w}^r = [w_0^r, \dots, w_{N-1}^r]^\top$ be the reference control sequence at the current iteration, and let $\mathbf{x}^r = [x_0^r, \dots, x_N^r]^\top$ be the corresponding reference state sequence, such that

$$x_{k+1}^r = F(x_k^r, w_k^r), \quad k = 0, \dots, N-1. \quad (6)$$

The dynamical system in (6) can then be approximated in the vicinity of $(\mathbf{x}^r, \mathbf{w}^r)$ with a discrete-time, linear time-varying (LTV) system as follows

$$x_{k+1} = A_k x_k + B_k u_k + d_k, \quad (7)$$

where x_k and u_k are the state and control input respectively to the LTV system at step k , and

$$A_k = \left. \frac{\partial F}{\partial x_k} \right|_{x_k=x_k^r, u_k=w_k^r}, \quad B_k = \left. \frac{\partial F}{\partial u_k} \right|_{x_k=x_k^r, u_k=w_k^r}, \quad (8)$$

$$d_k = F(x_k^r, w_k^r) - A_k x_k^r - B_k w_k^r, \quad (9)$$

where A_k and B_k are the system matrices, and d_k is the residual term of the linearization.

B. Covariance-Controlled Trajectory Sampling

As with the baseline MPPI algorithm, the CC-MPPI algorithm simulates M trajectories during each iteration. Let $\mathbf{u} = [u_0, \dots, u_{N-1}]^\top$ be the control sequence of the m^{th} CC-MPPI sampled trajectory during the current iteration of the algorithm where we drop the superscript m for simplicity. The optimal control sequence from the previous iteration which is also the reference control sequence of the current iteration $\mathbf{w} = [w_0, \dots, w_{N-1}]^\top$ is injected with artificial noise $\epsilon = [\epsilon_0, \dots, \epsilon_{N-1}]^\top$ and a feedback term $K_k y_k$ is added, such that

$$u_k = w_k + \epsilon_k + K_k y_k, \quad \epsilon_k \sim \mathcal{N}(0, \Sigma_\epsilon), \quad (10)$$

where y_k follows the dynamics,

$$y_{k+1} = A_k y_k + B_k \epsilon_k, \quad y_0 = 0, \quad (11)$$

where $y_0 = 0$ since we assume perfect observation of the initial state [17]. Substituting (10) into (7) yields,

$$x_{k+1} = A_k x_k + B_k (w_k + K_k y_k) + d_k + B_k \epsilon_k, \quad (12)$$

where x_k is the state of the m^{th} CC-MPPI sampled trajectory at step k , and d_k is the residual term of the linearization at step k as defined in (9). Let x_0 be the state at the beginning of the current iteration. We can then rewrite the system in (12) in the compact form,

$$\mathbf{x} = \mathcal{A}x_0 + \mathcal{B}(\mathbf{w} + \mathbf{K}\mathbf{y}) + \mathcal{C}\mathbf{d} + \mathcal{B}\epsilon, \quad (13)$$

where $\mathbf{x} = [x_0, \dots, x_N]^\top \in \mathbb{R}^{n_x(N+1)}$, $\mathbf{w} = [w_0, \dots, w_{N-1}] \in \mathbb{R}^{n_u N}$, $\mathbf{y} = [y_0, \dots, y_N]^\top \in \mathbb{R}^{n_x(N+1)}$, $\mathbf{d} = [d_0, \dots, d_{N-1}]^\top \in \mathbb{R}^{n_x N}$, $\epsilon = [\epsilon_0, \dots, \epsilon_{N-1}]^\top \in \mathbb{R}^{n_x N}$ and the augmented system matrices \mathcal{A} , \mathcal{B} , \mathcal{C} , \mathbf{K} are defined similarly as in [7]. In order to compute \mathbf{K} to satisfy the terminal covariance constraint (1d), the CC-MPPI considers a constrained LQG problem (14) at each optimization iteration as follows

$$\begin{aligned} \min_{\mathbf{K}} J(\mathbf{x}, \mathbf{w}, \mathbf{K}) = \\ \mathbb{E} [\mathbf{x}^\top \bar{Q} \mathbf{x} + (\mathbf{w} + \mathbf{K}\mathbf{y})^\top \bar{R} (\mathbf{w} + \mathbf{K}\mathbf{y})], \end{aligned} \quad (14a)$$

subject to,

$$\mathbf{x} = \mathcal{A}x_0 + \mathcal{B}(\mathbf{w} + \mathbf{K}\mathbf{y}) + \mathcal{C}\mathbf{d} + \mathcal{B}\epsilon, \quad (14b)$$

$$E_N(\mathbb{E}[\mathbf{x}\mathbf{x}^\top] - \mathbb{E}[\mathbf{x}]\mathbb{E}[\mathbf{x}]^\top) E_N^\top \preceq \Sigma_f, \quad (14c)$$

where $E_N = [0, \dots, 0, I]^\top \in \mathbb{R}^{n_x \times n_x(N+1)}$, the augmented cost parameter matrices $\bar{Q} = \text{blkdiag}(Q, \dots, Q, Q_f) \in \mathbb{R}^{n_x(N+1) \times n_x(N+1)}$ and $\bar{R} = \text{blkdiag}(R, \dots, R) \in \mathbb{R}^{n_u N \times n_u N}$. Since $\mathbb{E}[\epsilon] = 0$ and $y_0 = 0$, we have $\mathbb{E}[\mathbf{w} + \mathbf{K}\mathbf{y}] = \mathbf{w}$. It follows from (11) and (13) that

$$\bar{\mathbf{x}} = \mathbb{E}[\mathbf{x}] = \mathcal{A}x_0 + \mathcal{B}\mathbf{w} + \mathcal{C}\mathbf{d}, \quad (15)$$

and,

$$\tilde{\mathbf{x}} = \mathbf{x} - \mathbb{E}[\mathbf{x}] = \mathcal{B}\mathbf{K}\mathbf{y} + \mathcal{B}\epsilon = (I + \mathcal{B}\mathbf{K})\mathcal{B}\epsilon. \quad (16)$$

The cost function $J(\mathbf{x}, \mathbf{w}, \mathbf{K})$ in (14) can then be converted to the following equivalent form [8]

$$\begin{aligned} J(\bar{\mathbf{x}}, \tilde{\mathbf{x}}, \mathbf{w}, \mathbf{K}) = \text{tr}(\bar{Q}\mathbb{E}[\tilde{\mathbf{x}}\tilde{\mathbf{x}}^\top]) \\ + \bar{\mathbf{x}}^\top \bar{Q} \bar{\mathbf{x}} + \text{tr}(\bar{R}\mathbb{E}[\mathbf{K}\mathbf{y}\mathbf{y}^\top \mathbf{K}^\top]) + \mathbf{w}^\top \bar{R} \mathbf{w}. \end{aligned} \quad (17)$$

The reference control sequence \mathbf{w} is fixed and is given by the optimal control sequence from the previous CC-MPPI iteration, which implies that $\bar{\mathbf{x}}$ is fixed and is given by (15). For the optimization problem in (14), we can then drop the terms representing constant values in (17) and obtain the cost

$$J_{\Sigma}(\bar{\mathbf{x}}, K) = \text{tr}(\bar{Q}\mathbb{E}[\bar{\mathbf{x}}\bar{\mathbf{x}}^{\top}]) + \text{tr}(\bar{R}\mathbb{E}[K\mathbf{y}\mathbf{y}^{\top}K^{\top}]). \quad (18)$$

Substituting (16) into (18), and using the fact that $\mathbb{E}[\epsilon\epsilon^{\top}] = \Sigma_{\epsilon}$, yields,

$$J_{\Sigma}(K) = \text{tr}(((I + BK)^{\top}\bar{Q}(I + BK) + K^{\top}\bar{R}K) B\bar{\Sigma}_{\epsilon}B^{\top}), \quad (19)$$

where $\bar{\Sigma}_{\epsilon} = \text{blkdiag}(\Sigma_{\epsilon}, \dots, \Sigma_{\epsilon}) \in \mathbb{R}^{n_u N \times n_u N}$. Substituting (13) into (14c), we obtain,

$$\Sigma_f \succeq E_N(I + BK)B\bar{\Sigma}_{\epsilon}B^{\top}(I + BK)^{\top}E_N. \quad (20)$$

Finally, Problem (14) can be converted into the following convex optimization problem,

$$\min_K J_{\Sigma}(K) \quad \text{subject to} \quad (20). \quad (21)$$

The problem (21) can be easily solved by a convex optimization solver such as Mosek [27] to obtain K . Note that the last diagonal term Q_f in \bar{Q} can be viewed as a soft terminal state constraint that can be tuned to control the terminal covariance in situations where a hard constraint (20) is not preferred. It follows from (10) that the control sequence of the m^{th} sampled trajectory is $\mathbf{u} = \mathbf{w} + \epsilon + K\mathbf{y}$. We then rollout the sampled trajectories using \mathbf{u} and the dynamical model (1b). The complete CC-MPPI algorithm is detailed in Section V.

V. THE CC-MPPI ALGORITHM

The CC-MPPI algorithm is given in Algorithm 1. Line 2 obtains the current estimate of the state x_0 at the beginning of the current optimization iteration. Lines 3 to 5 rollout the reference trajectory \mathbf{x} using the discrete-time nonlinear dynamical model F . Line 6 linearizes the model F along \mathbf{x} and its corresponding control sequence \mathbf{w} as described in (6), (7), (8), (9), and calculates the augmented dynamical model matrices \mathcal{A} , \mathcal{B} , \mathcal{C} along with the linearization residual term \mathbf{d} . Line 7 computes the feedback gain K for the closed-loop system in (13) by solving the convex optimization problem (21). Lines 8 to 23 sample the control sequences, perform the rollouts and evaluate the sampled trajectories with the running cost (3). Specifically, lines 12 to 18 introduce sample trajectories of the close-loop dynamics and sample trajectories of zero-mean input, so that the algorithm can balance between smoothness of trajectories and low control cost [22]. Line 24 computes the optimal control sequence \mathbf{w} following (4) and (5). Line 25 sends the first control command w_0 of the optimal control sequence to the actuators. Line 26 removes w_0 , and duplicates w_{N-1} at the end of the horizon for \mathbf{w} .

VI. RESULTS

In this section, we show via a series of numerical examples that the proposed CC-MPPI algorithm outperforms the baseline MPPI algorithm in critical situations described

Algorithm 1: CC-MPPI Algorithm

Given: $\gamma, \lambda, \Sigma_{\epsilon}, \Sigma_f, \alpha$: CC-MPPI Parameters;
Input: \mathbf{w} : Initial control sequence

```

1 while task not complete do
2    $x_0 \leftarrow \text{GetStateEstimate}()$ ;
3   for  $k \leftarrow 0$  to  $N - 1$  do
4      $x_{k+1} \leftarrow F(x_k, w_k)$ ;
5   end
6    $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathbf{d} \leftarrow \text{DynamicsLinearization}(F, \mathbf{x}, \mathbf{w})$ ;
7    $K \leftarrow \text{CovarianceControl}(\Sigma_f, \Sigma_{\epsilon}, \mathcal{B}, \bar{Q}, \bar{R})$ ;
8   for  $m \leftarrow 1$  to  $M$  do
9      $x_0^{(m)} \leftarrow x_0, y_0^{(m)} \leftarrow 0$ ;
10    Sample  $\epsilon^{(m)} \leftarrow [\epsilon_0^{(m)}, \dots, \epsilon_{N-1}^{(m)}]^{\top}$ ;
11    for  $k \leftarrow 0$  to  $N - 1$  do
12      if  $m < (1 - \alpha)M$  then
13         $v_k^{(m)} \leftarrow w_k + K_k y_k^{(m)}$ ;
14      else
15         $v_k^{(m)} \leftarrow 0$ 
16      end
17       $u_k^{(m)} \leftarrow v_k^{(m)} + \epsilon_k^{(m)}$ ;
18       $x_{k+1}^{(m)} \leftarrow F(x_k^{(m)}, u_k^{(m)})$ ;
19       $y_{k+1}^{(m)} \leftarrow A_k y_k^{(m)} + B_k \epsilon_k^{(m)}$ ;
20       $S_m \leftarrow S_m + \tilde{q}(x_k^{(m)}, v_k^{(m)}, \epsilon_k^{(m)})$ ;
21    end
22     $S_m \leftarrow S_m + \phi(x_N^{(m)})$ ;
23  end
24   $\mathbf{w} = \text{CalculateOptimalControl}(S_m, \mathbf{u}^{(m)})$ ;
25   $\text{ExecuteCommand}(w_0)$ ;
26   $\text{Initialize}(\mathbf{w})$ ;
27 end

```

in Section II. The terminal covariance Σ_f in (20) for the CC-MPPI should be determined based on the environment. Please also refer to the accompanied video¹ for a demonstration of more simulation examples.

A. Vehicle Model

We assume that the injected artificial noise $\epsilon^{(m)}$ in CC-MPPI and MPPI algorithms is significantly greater than the inherent noise of the vehicle model, such that the model noise is negligible. We model the vehicle using a single-track bicycle model

$$\dot{x} = v \cos(\beta + \phi), \quad \dot{y} = v \sin(\beta + \phi), \quad (22a)$$

$$\dot{\phi} = \frac{v}{l_r \sin \beta}, \quad \dot{v} = \delta_T, \quad (22b)$$

where $\tan \beta = l_r / (l_f + l_r) \tan \delta_s$ and the parameters l_r, l_f are distances from the CoM to the rear and front wheels, respectively. The x, y are position coordinates of a fixed world coordinate frame. The ϕ is the vehicle yaw angle, and v is velocity at CoM with respect to the world coordinate frame. The δ_T and δ_s are throttle and steering inputs to

¹<https://www.youtube.com/watch?v=cZq4tnBTIqc>



Fig. 2. Track schematic

the model, respectively. We discretize the system (22) with the Euler method, $x_{k+1} = F(x_k, u_k) = x_k + f(x_k, u_k)\Delta t$, where $k = 0, \dots, N - 1$ and time step $\Delta t = 0.02$ s.

B. Controller Setup

Assuming that the model noise is much smaller than the injected noise $\epsilon^{(m)}$, it follows that $\nu^{-1} \approx 0$ in (3). It then follows from (3) that the running cost $\tilde{q}(x_k^{(m)}, v_k^{(m)}, \epsilon_k^{(m)})$ for the k^{th} step of the m^{th} sampled trajectory takes the form,

$$\begin{aligned} \tilde{q}(x_k^{(m)}, v_k^{(m)}, \epsilon_k^{(m)}) &= q(x_k^{(m)}) + \\ &\frac{1}{2}\epsilon_k^{(m)\top} R \epsilon_k^{(m)} + v_k^{(m)\top} R \epsilon_k^{(m)} + \frac{1}{2}v_k^{(m)\top} R v_k^{(m)}, \end{aligned} \quad (23)$$

for $k = 0, \dots, N - 1$, $m = 1, \dots, M$, where we take the state-dependent cost $q(x_k^{(m)})$ as,

$$q(x) = \mu_{\text{bdry}}(x) + c_1 \mu_{\text{obs}}(x). \quad (24)$$

The term $\mu_{\text{bdry}}(x)$ in (24) is the boundary cost which prevents the vehicle from leaving the track, and it is given by

$$\mu_{\text{bdry}}(x) = \begin{cases} 0, & \text{if } \text{WithinBoundary}(x), \\ 2000, & \text{Otherwise.} \end{cases} \quad (25)$$

The term $c_1 \mu_{\text{obs}}(x)$ in (24) penalizes collisions with obstacles, where c_1 is a weighting coefficient. We choose two different forms of $\mu_{\text{obs}}(x)$ in our simulations. The first is discontinuous on the obstacles' edges,

$$\mu_{\text{obs}}(x) = \begin{cases} 0, & \text{if } r_i > d_i(x), \\ 10, & \text{Otherwise,} \end{cases} \quad (26)$$

and the second is continuous on the obstacles' edges,

$$\mu_{\text{obs}}(x) = \sum_{i=1}^P \max(r_i - d_i(x), 0), \quad (27)$$

where $d_i(x)$ describes the distance from the vehicle's CoM to the center of the i^{th} circular obstacle, and r_i is the radius of the i^{th} obstacle. We take $r_i = 0.1$ m for all of the P circular obstacles in this section. In our simulations, the terminal cost $\phi(x_N^{(m)})$ for the MPPI and CC-MPPI controllers has the form,

$$\phi(x) = c_2(1 - s(x)) + 500e^2(x). \quad (28)$$

The first term in (28) is the progress cost, where c_2 is a weighting coefficient and $s(x)$ represents the distance between the current vehicle state and the terminal state of the sample trajectory along the track centerline. The term $e(x)$ in (28) is the vehicle's lateral deviation from the track centerline (see Figure 2). For both the MPPI and the CC-MPPI controllers, we set the control horizon to $N = 15$, the value for $\lambda = 1$, the number of sampled trajectories to

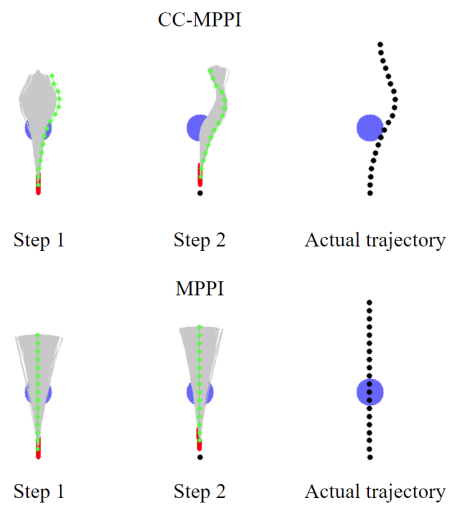


Fig. 3. Responses of the MPPI and CC-MPPI to an unpredictable obstacle. The grey curves are the sampled trajectories, the green curves represent the predicted optimal trajectories generated by the controllers, and the black points show the actual trajectories taken by the vehicle.

$M = 4096$, the portion of uncontrolled sampled trajectories to $\alpha = 0.2$, and the control cost matrix to $R = 0.01I$. The parameter values discussed here are shared by all the controller setups in the simulations of this section.

C. Planning in Fast-changing Environment: Unpredictable Obstacles

This experiment tests the CC-MPPI controller's ability to respond to emergencies owing to unpredictable appearance of obstacles. We test the CC-MPPI against a baseline MPPI controller in an environment where an obstacle suddenly appears in the traveling direction of an autonomous vehicle. In this simulation, the CC-MPPI and MPPI controllers were injected with noises having the same covariance $\Sigma_\epsilon = \text{diag}(0.0049, 0.0012)$, and the same weighting coefficients $c_1 = 712.5$ and $c_2 = 3.3$ in their trajectory costs. Figure 3 demonstrates that the MPPI controller fails to find a feasible solution and results in a collision with the obstacle. Figure 3 further shows that the CC-MPPI has a more effective trajectory sampling distribution strategy, which leads the vehicle to take a feasible trajectory that avoids collision.

D. Aggressive Driving in a Cluttered Environment

To further examine the performance of the CC-MPPI controller in more complicated environments, we run simulations using a CC-MPPI controller and an MPPI controller on a race track environment with obstacles densely scattered on the track. The track has a constant width of 0.6 m, the centerline has a length of 10.9 m and each turn of the centerline has radius 0.3 m. Each obstacle has radius 0.1 m and uses the continuous obstacle cost (27). Both controllers are set to achieve minimum lap time while avoiding collisions with cluttered obstacles, and have the same covariance $\Sigma_\epsilon = \text{diag}(0.49, 0.12)$ for their injected noises. We perform a grid search by varying the cost weight c_1 in (24) which corresponds to avoiding collisions with obstacles, and the

TABLE I
CC-MPPI AND MPPI GRID SEARCH PARAMETER VALUES

Cost Parameter	Min	Max	Interval
c_1	75	450	37.5
c_2	1.65	2.97	0.33




Fig. 4. MPPI and CC-MPPI trajectories on a race track. The trajectories in red are generated by the CC-MPPI, and the trajectories in green are generated by the MPPI using the same injected noise covariance as the CC-MPPI.

cost weight c_2 in (28) for optimizing the vehicle velocity along the track centerline. Table I shows the grid search parameters. Figure 5 presents the results of the grid search in a scatter plot showing the distribution of lap times and average number of collisions. Table II summarizes the grid search and shows the statistics over all the simulations.

We define a failure to be the situation when the vehicle comes to a complete stop, or when the vehicle is too far away from the track centerline (> 1 m). If the vehicle finishes 20 laps without a failure, the simulation is considered a success. Figure 5 shows that the data points corresponding to CC-MPPI occupy the bottom part of the scatter plot, which indicates that the CC-MPPI generates trajectories that are significantly faster than those by the MPPI controller. Table II shows that the CC-MPPI achieves 34.78% smaller average lap time, 6.34% fewer collisions and 218.97% higher success rate than the MPPI in simulations. Moreover, the two data points in the red circles in Figure 5 are produced by MPPI and CC-MPPI with the same set of c_1 and c_2 values, and Figure 4 visualizes the trajectories that correspond to these two data points. We see that the CC-MPPI generates a more aggressive driving maneuver than the MPPI, which helps explain why the CC-MPPI achieves a significantly smaller average lap time.

The performance of CC-MPPI, however, comes with an increased computational overhead. Using our implementation, the CC-MPPI controller runs at 13Hz, while the MPPI controller runs at 97Hz. All simulations were done on a

TABLE II
CC-MPPI Vs. MPPI OF DIFFERENT SETTINGS

Controller	Avg. laptime(s)	No. collision/lap	Success rate
CC-MPPI	4.20	2.68	98.18%
MPPI	6.44	2.86	30.78%

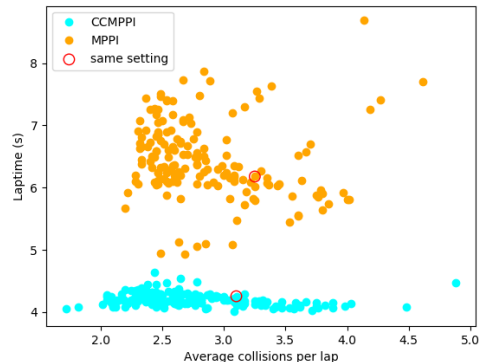


Fig. 5. Lap time and number of collisions distribution. Each point in this figure shows the average lap time and average number of collisions over 20 laps in a simulation using one pair of c_1 and c_2 from Table I. The orange points are produced by the MPPI controller and the cyan points are by the CC-MPPI controller.

desktop computer equipped with an i9 3.5GHz CPU, and an RTX3090 GPU. The main computational bottleneck of CC-MPPI is the computation of the feedback gain K at each iteration. Possible remedies include updating the feedback gain asynchronously, computing the feedback gains off-line and storing them in a lookup table, or using a faster, dedicated convex optimization solver that is more suitable for real-time implementation [28], [29], [30].

VII. CONCLUSIONS AND FUTURE WORK

We have proposed the Covariance-Controlled Model Predictive Path Integral (CC-MPPI) algorithm that incorporates covariance steering within the MPPI algorithm. The CC-MPPI algorithm has adjustable trajectory sampling distributions which can be tuned by changing the terminal covariance constraint Σ_f in (1d) and the covariance of the injected noise Σ_ϵ in (1c), which makes it more flexible and robust than the MPPI algorithm. In the simulations, we showed that the CC-MPPI explores the environment and samples trajectories more efficiently than MPPI for the same level of exploration noise (Σ_ϵ). This results in the vehicle responding faster to unpredictable obstacles and avoid collisions better in a cluttered environment than MPPI. The CC-MPPI performance can be further improved if Σ_ϵ and Σ_f are tuned based on the information of the surrounding environment.

In the future, we plan to design policies to choose judiciously the terminal covariance constraint Σ_f and the injected noise covariance Σ_ϵ on-the-fly. These policies should evaluate the environment and assign Σ_f , Σ_ϵ for the CC-MPPI controller, such that the trajectory sampling distribution of the controller can be tailored to carry out informed and efficient sampling in any environments.

REFERENCES

- [1] R. Pepy and A. Lambert, "Safe path planning in an uncertain-configuration space using RRT," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov. 2006, pp. 5376 – 5381.
- [2] N. A. Melchior and R. Simmons, "Particle RRT for path planning with uncertainty," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 1617–1624.
- [3] Y. Huang and K. Gupta, "Collision-probability constrained prm for a manipulator with base pose uncertainty," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 1426–1432.
- [4] G. Kewlani, G. Ishigami, and K. Iagnemma, "Stochastic mobility-based path planning in uncertain environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 1183–1189.
- [5] I. M. Balci and E. Bakolas, "Covariance steering of discrete-time stochastic linear systems based on Wasserstein distance terminal cost," *IEEE Control Systems Letters*, vol. 5, no. 6, pp. 2000–2005, 2021.
- [6] M. Goldshtein and P. Tsiotras, "Finite-horizon covariance control of linear time-varying systems," in *56th IEEE Conference on Decision and Control (CDC)*, 2017, pp. 3606–3611.
- [7] K. Okamoto and P. Tsiotras, "Optimal stochastic vehicle path planning using covariance steering," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2276–2281, 2019.
- [8] K. Okamoto, M. Goldshtein, and P. Tsiotras, "Optimal covariance control for stochastic systems under chance constraints," *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 266–271, 2018.
- [9] P. Tsiotras and M. Mesbahi, "Toward an algorithmic control theory," *Journal of Guidance, Control, and Dynamics*, vol. 40, pp. 1–3, 02 2017.
- [10] U. Rosolia and F. Borrelli, "Learning how to autonomously race a car: A predictive control approach," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2713–2719, 2020.
- [11] A. Liniger, X. Zhang, P. Aeschbach, A. Georghiou, and J. Lygeros, "Racing miniature cars: Enhancing performance using stochastic MPC and disturbance feedback," in *American Control Conference*, Seattle, WA, 2017, pp. 5642–5647.
- [12] A. A. Jalali and V. Nadimi, "A survey on robust model predictive control from 1999-2006," in *International Conference on Computational Intelligence for Modelling Control and Automation and International Conference on Intelligent Agents Web Technologies and International Commerce (CIMCA'06)*, 2006, pp. 207–207.
- [13] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.
- [14] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear model predictive control with chance constraints – a review," *Journal of Process Control*, vol. 44, pp. 53–67, 08 2016.
- [15] M. Cannon, B. Kouvaritakis, and D. Ng, "Probabilistic tubes in linear stochastic model predictive control," *Systems & Control Letters*, vol. 58, pp. 747–753, 2009.
- [16] M. Cannon, B. Kouvaritakis, and X. Wu, "Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty," *IEEE Transactions on Automatic Control*, vol. 54, no. 7, pp. 1626–1632, 2009.
- [17] K. Okamoto and P. Tsiotras, "Stochastic model predictive control for constrained linear systems using optimal covariance steering," 2019, arXiv:1905.13296.
- [18] D. Bernardini and A. Bemporad, "Scenario-based model predictive control of stochastic constrained linear systems," in *Proceedings of the 48th IEEE Conference on Decision and Control*, Shanghai, China, 2009, pp. 6333–6338, held jointly with the 28th Chinese Control Conference.
- [19] G. C. Calafiore and L. Fagiano, "Robust model predictive control via scenario optimization," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 219–224, 2013.
- [20] 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.
- [21] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic MPC for model-based reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, 2017, pp. 1714–1721.
- [22] 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.
- [23] 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, PA, 2018, pp. 42–51.
- [24] J. Pravitra, K. A. Ackerman, C. Cao, N. Hovakimyan, and E. A. Theodorou, "L1-adaptive MPPI architecture for robust and agile control of multirotors," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 7661–7666.
- [25] G. Williams, B. Goldfain, P. Drews, J. M. Rehg, and E. A. Theodorou, "Autonomous racing with autarally vehicles and differential games," 2017, arXiv:1707.04540.
- [26] P. Falcone, M. Tufo, F. Borrelli, J. Asgari, and H. E. Tseng, "A linear time varying model predictive control approach to the integrated vehicle dynamics control problem in autonomous systems," *IEEE Conference on Decision and Control*, pp. 2980–2985, 2007.
- [27] *MOSEK ApS, The MOSEK Optimization Toolbox for MATLAB Manual. Version 8.1.*, 2017. [Online]. Available: <http://docs.mosek.com>
- [28] Y. Mao, M. Szmuk, and B. Açikmeşe, "A tutorial on real-time convex optimization based guidance and control for aerospace applications," in *Annual American Control Conference (ACC)*, 2018, pp. 2410–2416.
- [29] J. Mattingley and S. Boyd, "CVXGEN: a code generator for embedded convex optimization," *Optimization and Engineering*, vol. 13, no. 1, pp. 1–27, 2012.
- [30] D. Dueri, J. Zhang, and B. Açikmeşe, "Automated custom code generation for embedded, real-time second order cone programming," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 1605–1612, 2014.