

# Output-Sampled Model Predictive Path Integral Control (o-MPPI) for Increased Efficiency

Leon (Liangwu) Yan and Santosh Devasia

**Abstract**—The success of the model predictive path integral control (MPPI) approach depends on the appropriate selection of the input distribution used for sampling. However, it can be challenging to select inputs that satisfy output constraints in dynamic environments. The main contribution of this paper is to propose an output-sampling-based MPPI (o-MPPI), which improves the ability of samples to satisfy output constraints and thereby increases MPPI efficiency. Comparative simulations and experiments of dynamic autonomous driving of bots around a track are provided to show that the proposed o-MPPI is more efficient and requires substantially (20-times) less number of rollouts and (4-times) smaller prediction horizon when compared with the standard MPPI for similar success rates.

## I. INTRODUCTION

Sampling-based approaches such as model predictive path integral control (MPPI) [1] have become popular methods to solve optimization problems due to fast computations possible with graphic processing units and parallelized computing. In such methods, models are used to predict the cost for a series of input rollouts, and the final input selection is a cost-weighted average of the rollouts. An advantage of the sample-based approach is that it is derivative-free, does not require approximation of the system dynamics and cost functions, and allows for non-differentiable cost functions [2].

Successful use of the MPPI algorithm requires proper selections of the mean and the covariance of the input samples. If the selected mean leads to a rollout that enters inside an infeasible region (this could be caused by the poor initialization, system disturbance, or sudden environment changes), then most of the sampled rollouts can result in failure [3]. Although a larger covariance can be used to alleviate failures by increasing exploration [4], it also often requires a larger number of samples and increased computation load, and can lead to undesirable input chattering [1], [5]. Therefore, there is substantial ongoing interest in methods to appropriately select the input samples to improve MPPI performance. However, in general, it is challenging to appropriately select a desirable input mean to meet constraints in the output space, especially in dynamic environments. In general, even when the reference input can be selected as desired outputs to a closed-loop system that satisfies constraints, it does not ensure that the output achieves precision tracking of the desired trajectories.

This work was supported by NSF Grant CMMI 1824660

Leon (Liangwu) Yan (liangy00@uw.edu) and Santosh Devasia (IEEE Fellow, devasia@uw.edu, see <http://faculty.washington.edu/devasia/>) are with the Department of Mechanical Engineering, University of Washington, Seattle,

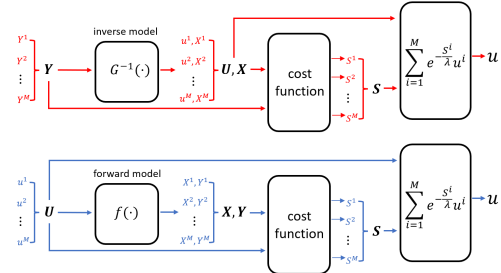


Fig. 1. Comparison between the proposed output-MPPI (o-MPPI) in red (top) and the standard MPPI in blue (bottom). There are  $M$  rollouts and  $f(\cdot)$  maps inputs to states and outputs in the standard MPPI and the inverse  $G^{-1}(\cdot)$  is the inverse dynamics that maps output trajectories to inputs and the states in the proposed o-MPPI.  $S^m$  is the cost for the  $m^{th}$  rollout output-state-input  $(Y^m, X^m, u^m)$ .  $\lambda \in R^+$  is the temperature parameter used in the weighting to obtain the optimized input  $u$ .

With the goal of increasing rollouts that satisfy output constraints, the current paper proposes the sampling of trajectories directly from the output space. Then, the proposed output-sampling-based MPPI (o-MPPI) approach uses the inverse model  $G^{-1}(\cdot)$  to map outputs to inputs rather than the traditional approach in MPPI of using the input-to-output forward model  $f(\cdot)$  as illustrated in Fig. 1(top). The inputs obtained using the inverse model are then weighted to obtain the optimized control sequence, as in standard MPPI.

The main contributions of the paper are the following.

- 1) It proposes an output-sampling-based MPPI (o-MPPI) using inverse models, which improves the ability to select rollouts that satisfy typical output constraints. An advantage of output sampling is that it can leverage well-established trajectory planning algorithms in robotics to select output samples and then use inverse models (physics-based or data-based) to find the associated inputs. This is different from selecting inputs and then finding the associated outputs using forward models.
- 2) Comparative simulations and experiments of autonomous control of bots (TurtleBot3 Burger) in a dynamic autonomous driving around a track, are used to show that the proposed o-MPPI requires substantially (20-times) less number of rollouts and (4-times) smaller prediction horizon when compared with the standard MPPI to achieve a similar success rate.

## II. RELATED WORK

### A. MPPI with adjusted trajectory distribution

Several efforts are made to improve the robustness and sample efficiency of standard MPPI by appropriately adjusting the trajectory distribution using different sampling

strategies [3], [6], [7], [8]. For example, warm-starts are used [1] to improve MPPI-variants. In [6], covariances are adapted to accommodate different actions for better exploration. In [7], the samples are drawn from a normal and log-normal (NLN) distribution instead of the Gaussian distribution, which can yield improved performance in cluttered environments. In [8], the diagonal covariance matrix  $\Sigma_\epsilon$  is changed into a nondiagonal covariance matrix and updated through the adaptive importance sampling procedure. Recently, constraints were added to the terminal state of the prediction horizon in [3] to adjust the input distribution. In all these works sampling is in the input space followed by the use of forward models to determine the outputs. In contrast, the current work proposes sampling in the output space (since constraints might be more easily satisfied in the output space) and then uses an inverse model to find the corresponding inputs. Once the input-output pairs are found, the proposed o-MPPI is similar to standard MPPI and can use the same cost functions and weighting strategies.

### B. MPPI with output-space-informed mean

The output space has been leveraged in the past to improve the input-sampling efficiency of MPPI. Since a better mean for the input distribution can also help improve MPPI, an appropriately selected output can be used to inform the mean selection used in standard MPPI. For example, in [4], the fast rapid-exploring-tree (*RRT\**) algorithm is used to provide guidance about the mean value of the input for improving sampling efficiency in the input space. In [9], trained Conditional Variational Autoencoders that take into account the contextual information are used to better inform the mean values by taking into the control uncertainties. Thus, the above methods utilize output-space information to guide the selection of the mean value for the input sampling. The proposed o-MPPI approach extends this idea and fully samples in the output space and then uses inverse models to find the corresponding input?. It is noted that sampling of the output space has been used in the path-planning community to smooth and potentially optimize feasible trajectories, e.g., [10], [11]. Here, these planned trajectories are inputs to the closed-loop system, and forward models (of the closed-loop system) are used to predict the system response and for optimizing the selected input (final output trajectory) to the closed-loop system. The proposed o-MPPI can be used with any trajectory planning algorithm such as the fast rapid-exploring-tree (*RRT\**) [4] – the main difference is that inverse models are used in o-MPPI to find inputs that track the selected output trajectories.

### C. Inverse models

The proposed o-MPPI leverages the strong history in inverse dynamics to find maps from outputs to inputs. For example, physics-based models can be used to analytically find the inverse input [12] in robotics applications. When the physics-based models are not sufficiently precise, the Gaussian process can be used to learn the discrepancy [13]. The entire inverse model can be approximated with deep

learning models [14] or Gaussian process models [15]. Additionally, data-enabled models can be used to learn the inverse output-to-input map from input-output data [16]. Thus, the proposed o-MPPI can use a wide range of inverse modeling tools to find the inputs associated with the sampled output trajectories.

## III. PROPOSED FRAMEWORK

The proposed output-sampling-based MPPI (o-MPPI) approach is summarized in Algorithm 1. Essentially, there are four steps (detailed below): (i) selection of the cost function; (ii) sampling in the output space; (iii) inversion to find the associated inputs; and (iv) weighted selection of the input.

**Algorithm 1** o-MPPI (red is the difference from standard MPPI)

- 
- 1: **Given:** Number of rollouts & time steps  $M, N$ ; Cost function; Temperature parameter  $\lambda$ ; **Inverse model**  $G^{-1}$ ;
  - 2: **while** Task is not done **do**
  - 3:    $X_k \leftarrow$  state estimate
  - 4:   **for**  $m \rightarrow 0$  to  $M - 1$  in parallel **do**
  - 5:      $Y_0^m \leftarrow h(X_k), S^m \leftarrow 0$
  - 6:     **[Sampling]** **Sample the  $m^{\text{th}}$  trajectory rollout  $Y^m$**
  - 7:     **[Inverse]** **Compute the corresponding inverse input  $u^m$  and states  $X^m$  from the inverse model  $G^{-1}$**
  - 8:     Calculate the trajectory cost  $S^m$  based on the cost function and the  $m^{\text{th}}$  rollout ( $u^m, X^m, Y^m$ )
  - 9:   **end for**
  - 10: **[Weighting]** For  $m = 1, 2, \dots, M$ , compute the normalized weights  $\{w_m\}$  based on the trajectory costs  $\{S^m\}$  and the selected temperature parameter  $\lambda$
  - 11: Obtain the weighted average  $u = \sum_{m=0}^M w_m u^m$
  - 12: Apply the first entry of  $u$  to the system
  - 13:    $k \leftarrow k + 1$
  - 14: **end while**
- 

As in standard MPPI, a cost function is used to specify the desirability of a specific input-state-output rollout. In particular, consider the optimization

$$\min_u J(u) = \phi(X_{N-1}) + \sum_{k=0}^{N-1} q(X_k) + \frac{1}{2} u_k^T R u_k \quad (1)$$

subject to system dynamics

$$X_{k+1} = f(X_k, u_k), Y_k = h(X_k), \quad (2)$$

where  $f$  represents the forward dynamics,  $h$  maps the state  $X_k$  to the output  $Y_k$  at time step  $k$ , and the output trajectory is given by  $\mathbf{Y} \triangleq [h(X_0) \ h(X_1) \ \dots \ h(X_{N-1})]$ . In the above optimization,  $q(\cdot)$  is the running cost,  $R$  is the weight matrix of the input energy cost, and the terminal cost is  $\phi(\cdot)$  in Eq. (1). Additionally, there are constraints on the output trajectory to lie in an acceptable region, i.e.,  $\mathbf{Y} \in \mathcal{Y}$ .

The output sampling can be generated from any trajectory planning method, e.g., parametric ones like spline or bezier curves by specifying the waypoints as in Fig. 2, using optimization methods such as  $k^{\text{th}}$ -order constrained

path optimization (KOMO) [17], via-point-based Stochastic Trajectory Optimization (VP-STO) [18], or from a data-based planner [19], [20], [21], [22]. For each sampled output, inverse maps are used to find the corresponding input that yields tracking of the output.




Fig. 2. A generic way of sampling trajectory rollouts based on fitted smooth curves specified by waypoints that are selected from the regions of interest (dashed ellipsoids).

The optimized input is obtained as the weighted sum of all the input rollouts as in standard MPPI, i.e.,

$$\sum_{m=0}^M w_m u^m = u_{\text{mean}} + \sum_{m=0}^M w_m \epsilon^m, \quad (3)$$

where  $u^m = u_{\text{mean}} + \epsilon^m$  and  $\{w_m\}$  for  $m = 1, 2, \dots, M$  are normalized weights, i.e.,  $\sum_{m=1}^M w_m = 1$ .

#### IV. APPLICATION OF O-MPPI TO EXPERIMENTAL SETUP

The system for evaluating the proposed o-MPPI is set up to mimic a dynamic autonomous driving scenario with moving obstacles.

##### A. System description

In the experiment, the goal is for a fast bot (green) to move as close as possible to its desired speed of 20 cm/s while maneuvering around slower bots that can be considered as moving obstacles at the same time. The ability for the fast bot to maintain its desired speed and overtake the slower bots, as illustrated in Fig. 3, are used to quantitatively compare the performance of the standard MPPI and the proposed o-MPPI. An example simulation run is shown in the left plot in Fig. 3. The slower bots are moving at constant speeds of 10 cm/s (blue) and 12 cm/s (red). All bots are driving in the counter-clockwise direction in ellipsoidal tracks with dimensions tabulated in Table I and illustrated in Fig. 3, similar to the shape of the track in [23].

TABLE I

TURTLEBOT, AND SPECIFICATIONS OF TRACK SHOWN IN FIG. 3.

bot turning circle radius (cm)	10.5	lane width $lw$ (cm)	30
max vel. $\bar{v}$ (cm/s)	22	straight lines $sl$ (cm)	150
max angular vel. $\bar{\omega}$ (rad/s)	2.8	inner radius $ir$ (cm)	40
collision area width $cw$ (cm)	30	collision area length $cl$ (cm)	63

##### B. Standard MPPI

1) *Forward model  $F$* : The bot (TurtleBot3 Burger) state  $X$  at time step  $k$  is defined as  $X_k \triangleq [x_k \ y_k \ \theta_k \ v_k \ \omega_k]$ .  $(x_k, y_k)$  is the position pair and  $\theta_k$  is the orientation with respect to the Turtle Track coordinates, as in Fig. 3.  $v_k$  is the forward velocity and  $\omega_k$  is the angular velocity. The input  $u$  to the system is  $u_k \triangleq [v_k^{\text{des}} \ \omega_k^{\text{des}}]$  where  $v_k^{\text{des}}$  is the desired forward velocity and  $\omega_k^{\text{des}}$  is the desired angular velocity, and the system output is  $Y_k \triangleq [x_k \ y_k]$ .

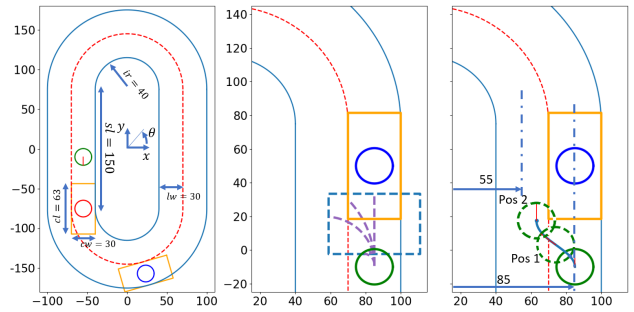


Fig. 3. The left plot is the schematic drawing of the track. The middle plot illustrates an example region of interest (dashed blue rectangle) from which the output waypoint is sampled. The right plot shows an example output rollout where the solid circles indicate the current position and dashed circles indicate future positions in the rollout. The bright yellow rectangles indicate the collision regions.

The forward bot dynamics  $\dot{X} = f(X, u)$  is approximated by the forward Euler discretization  $X_{k+1} = \mathbf{F}(X_k, u_k)$ ,

$$x_{k+1} = x_k + v_k \cos(\theta_k) * \Delta t \quad (4)$$

$$y_{k+1} = y_k + v_k \sin(\theta_k) * \Delta t \quad (5)$$

$$\theta_{k+1} = \theta_k + \omega_k * \Delta t \quad (6)$$

$$v_{k+1} = \text{sat}_{\bar{v}}(v_k + \alpha(v_k^{\text{des}} - v_k) * \Delta t) \quad (7)$$

$$\omega_{k+1} = \text{sat}_{\bar{\omega}}(\omega_k + \alpha(\omega_k^{\text{des}} - \omega_k) * \Delta t), \quad (8)$$

where the sampling time is  $\Delta t = 0.04$  s,  $\text{sat}_{\beta}(x) = \beta \text{sign}(x)$  if  $\text{abs}(x) > \beta$  otherwise  $\text{sat}_{\beta}(x) = x$  is a saturation function used to bound with the magnitudes of the speed and angular velocity with  $\bar{v}$  and  $\bar{\omega}$  defined in Table I, and the settling-time parameter  $\alpha = 4/0.35$  is estimated from the measured step response.

2) *Cost function selection*: The cost function  $J$  in Eq. (1) is designed to have only the running cost function  $q(\cdot)$ , with zero input energy cost ( $R=0$ ), similar to the cost functions used in autonomous driving in [23], [24]. The running cost function is designed to have three components, i.e.,  $q(\cdot) \triangleq c_l(\cdot) + c_v(\cdot) + c_c(\cdot)$ . The first cost  $c_l(\cdot)$ , designed for lane keeping, is the product of the squared values of the bot's distance to the center lines of the two lanes, as shown in the right plot in Fig. 3. Formally,

$$c_l(X_k) = w_0(r_k - 55)^2(r_k - 85)^2 + w_1\gamma_{OT}(x_k, y_k), \quad (9)$$

where  $w_0 = 0.001$ ,  $w_1 = 600$ ,  $\gamma_{OT}(x_k, y_k) = 1$  if  $(x_k, y_k)$  is outside of the designed track shown in Fig. 3 and zero otherwise,  $r_k = \sqrt{x_k^2 + z_k^2}$ , and  $z_k = 0$  if  $\text{abs}(y_k) < 75.0$  and  $z_k = y_k - 75.0 * \text{sign}(y_k)$  otherwise. The second cost  $c_v(\cdot)$ , designed for driving the bot at a speed close to the desired speed of 20 cm/s, is the squared of the speed difference between the current speed  $v_k$  and the desired one,

$$c_v(X_k) = w_1(v_k - 20)^2, \quad w_1 = 0.4. \quad (10)$$

The third cost  $c_c(\cdot)$ , designed to avoid collisions with the slower bots, is set to a large value 500 if  $(x_k, y_k)$  is inside the collision region at time step  $k$

$$c_c(X_k, \text{agent}_k) = \begin{cases} 500 & \text{proj}_f < (0.5cl + cr) \\ & \text{and proj}_l < 0.5cw, \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

where  $cr$  is the turning circle radius, and  $cl, cw$  are the collision region length and width, respectively, as seen in Table I. The computation of the projection distance  $\text{proj}_f$  and  $\text{proj}_l$  are described below

$$\text{proj}_f = \text{abs}(\mathbf{1}_f[\text{dis}_x, \text{dis}_y]^T), \text{proj}_l = \text{abs}(\mathbf{1}_l[\text{dis}_x, \text{dis}_y]^T)$$

where  $\text{dis}_x = x_{\text{agent},k} - x_k$ ,  $\text{dis}_y = y_{\text{agent},k} - y_k$ , and

$$\mathbf{1}_f = [\cos(\theta_{\text{agent},k}), \sin(\theta_{\text{agent},k})],$$

$$\mathbf{1}_l = [\cos(\theta_{\text{agent},k} - \frac{\pi}{2}), \sin(\theta_{\text{agent},k} - \frac{\pi}{2})].$$

3) *MPPI algorithm*: The standard MPPI [23] uses the temperature parameter selected as  $\lambda = 2.0$  and the covariance matrix is selected as  $\Sigma_\epsilon = \text{diag}([4.0^2, 1.0^2])$  to enable sufficient exploration. In general, an additional cost can be added to the cost function for input deviation as  $c(\cdot) = \sum_{k=0}^{N-1} \gamma u_k^T \Sigma_\epsilon^{-1} \epsilon_k$  [23]. However, the hyperparameter  $\gamma$  is set to zero to promote exploration in the dynamic environment with moving obstacles. The mean of the input distribution  $u_{\text{mean}}$  in Eq. (3) and the prediction horizon  $T$  are varied to illustrate their impact on MPPI effectiveness in the results and discussion Section V.

4) *Weighting*: Given the trajectory costs  $S^m$  for all the rollouts ( $0 \leq m \leq M-1$ ) based on the cost function, the weight  $w^m$  associated with each rollout ( $0 \leq m \leq M-1$ ) is computed as  $w^m = \frac{1}{\eta}(S^m - \beta)$ , where  $\beta$  is the minimum cost of the rollouts  $\beta = \min_m [S_m]$  and the normalizing factor  $\eta$  is  $\eta = \sum_{m=0}^{M-1} \exp(-\frac{1}{\lambda}(S_m - \beta))$ .

### C. o-MPPI

The proposed o-MPPI (see Algorithm 1) differs from the standard MPPI in the sampling of the output and the use of inverse to find the input – these are described below. The proposed o-MPPI uses the same temperature parameter  $\lambda$ , cost function and weighting to determine the input as the standard MPPI case in Sections IV-B.2 and IV-B.3.

1) *Sampling the trajectory rollout*: The generic way in Fig. 2 is used to sample the trajectory rollouts in this paper with only two waypoints: the initial output and the final output. Specifically, a rectangle-shape region of interest for the endpoint  $Y_{k+N-1} = [x_{k+N-1} \quad y_{k+N-1}]$  is determined based on the allowed magnitudes of the inputs  $u$  and the prediction horizon  $N$ . Then, the  $m^{\text{th}}$  sampled trajectory rollout is generated by fitting cubic splines between the current output and the  $m^{\text{th}}$  final output  $Y_{k+N-1}^m = (x_e^m, y_e^m)$  sampled inside the region of interest, as seen in the middle plot in Fig. 3. Specifically, the  $m^{\text{th}}$  rollout expression  $t \in [0, T]$  (for the output  $x$ ) is selected as a cubic spine as in [25]

$$x^m(t) = a_0 t + a_1 t + a_2 t^2 + a_3 t^3, \quad (12)$$

where the coefficients  $a_i$  for  $i = 1, \dots, 4$  can be determined from the boundary conditions

$$x^m(0) = x_k, \dot{x}^m(0) = v_k \cos(\theta_k),$$

$$x^m(T) = x_e^m, \dot{x}^m(T) = v_e^m \cos(\theta_e^m),$$

where the endpoint orientation  $\theta_e^m$  is designed to be aligned with the road direction and the endpoint forward speed  $v_e^m$  is

the travel distance divided by the prediction horizon  $T$ , i.e.,  $v_e^m = \sqrt{(x_e^m - x_k)^2 + (y_e^m - y_k)^2}/T$ . The output  $y^m(t)$  is generated in a similar manner with the selected final waypoint output  $y_e^m$ .

2) *Inverse model  $G^{-1}$* : Given the  $m^{\text{th}}$  differentiable output trajectory  $x^m(t), y^m(t)$  in the previous subsection with  $t \in [0, T]$  and starting state  $X_k$ , i.e.,  $x^m(0) = x_k$ ,  $\dot{x}^m(0) = v_k \cos(\theta_k)$ ,  $y^m(0) = y_k$ ,  $\dot{y}^m(0) = v_k \sin(\theta_k)$ , the inverse input can be obtained at time steps  $k+j$  with  $j = 0, 1, \dots, N-1$  from the forward dynamics in Eq. (7) to (8), as:

$$v_{k+j}^{\text{des}} = \frac{v_{p,k+j+1} - v_{p,k+j}}{\alpha \Delta t} + v_{p,k+j}$$

$$\omega_{k+j}^{\text{des}} = \frac{\omega_{p,k+j+1} - \omega_{p,k+j}}{\alpha \Delta t} + \omega_{p,k+j}, \quad (13)$$

where

$$v_{p,k+j} = \sqrt{(\dot{x}^m)^2(j\Delta t) + (\dot{y}^m)^2(j\Delta t)},$$

$$\omega_{p,k+j} = \begin{cases} (\theta_{p,k+j} - \theta_{p,k+j-1})/\Delta t & j > 0 \\ \omega_k & j = 0 \end{cases},$$

$$\theta_{p,k+j} = \arctan\left(\frac{\dot{y}^m(j\Delta t)}{\dot{x}^m(j\Delta t)}\right),$$

$v_{p,s}, \omega_{p,s}, \theta_{p,s}$  are the planned speed, angular velocity, and the orientation at time step  $s$ , respectively. Saturation is not considered when computing the inverse input in Eq. (13). However, if the final optimal input is large (for either o-MPPI or standard MPPI), then it is saturated in the experimental system and bounded in simulations by the saturation function in Eqs. (7)- (8).

## V. RESULTS AND DISCUSSION

The ability to avoid and successfully overtake moving obstacles in dynamic environments are comparatively evaluated below for the proposed o-MPPI and the standard MPPI. The evaluation starts with the study of a single dynamic obstacle, which is followed by the case with multiple obstacles.

### A. Case with a single dynamic obstacle

For comparative evaluation with a single obstacle, three MPPI and one o-MPPI cases are simulated with the following conditions.

- **Case 1 o-MPPI**: The proposed o-MPPI with prediction horizon  $T = 2.0$  s.
- **Case 2 Small-horizon standard MPPI**: Standard MPPI with prediction horizon  $T = 2.0$  s, i.e., prediction steps  $N = 50$  with sampling time period  $\Delta t = 0.04$  s. The nominal speed is 15 cm/s, i.e., with initial control sequence  $(u_0, u_1, \dots, u_{N-1}) \equiv [15 \text{ cm/s}, 0 \text{ rad/s}]$ ;
- **Case 3 Large-horizon standard MPPI**: Standard MPPI as in Case 2 except with a larger prediction horizon  $T = 8.0$  s and corresponding prediction steps  $N = 200$ ;
- **Case 4 Slow-initial standard MPPI**: Standard MPPI as in Case 2 except that the initial speed is

lower at 10 cm/s i.e., with initial control sequence  $(u_0, u_1, \dots, u_{N-1}) \equiv [10 \text{ cm/s}, 0 \text{ rad/s}]$ .

The controlled bot (green in Fig. 4) is said to achieve a successful overtake of a moving obstacle (blue circle in Fig. 4) if: (i) the controlled bot is driving in the counter-clockwise direction all the time, (ii) the controlled bot does not run outside of the track or into the collision regions of the moving obstacle, and (iii) within a specified time, the position of the controlled bot is ahead of the slower constant-speed obstacle. as indicated by the red dashed lines in Fig. 5. For all these cases, the initial position of the moving obstacle is  $(x, y) = (85 \text{ cm}, 50 \text{ cm})$  on one corner of the track as shown in Fig. 5 and it has a constant speed of 10 cm/s. The simulation is run till the moving obstacle reaches the other corner of the track as indicated in Fig. 5. The initial position of the controlled bot is behind the moving obstacle at  $X_0 = [x_0 \ y_0 \ \theta_0 \ v_0 \ \omega_0]^T = [85 \text{ cm} \ -10 \text{ cm} \ \frac{\pi}{2} \text{ rad} \ 15 \text{ cm/s} \ 0 \text{ rad/s}]^T$ . The initial trajectory rollouts for o-MPPI and standard MPPI for the four different cases are visualized in Fig. 4 for comparative evaluation. Additionally, 100 repeated simulations were run, and the success rates for the four cases with different numbers of rollouts  $M$  are tabulated in Table II and III.

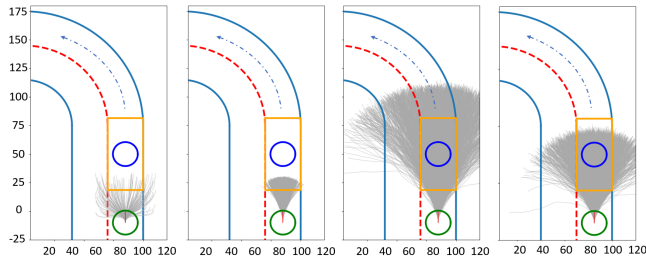


Fig. 4. Initial rollouts for Cases 1-4 (left to right). The green circle represents the controlled bot with the red line indicating its orientation. The blue circle denotes the constant-speed bot of speed 10 cm/s. The bright yellow rectangle depicts the collision region of the constant-speed bot. The cost function component  $c_c$  in Eq. (11) penalizes nonempty intersection of the controlled bot with the collision region. Grey lines demonstrate the trajectory rollouts of **Case studies 1 to 4** in Section V-A.




Fig. 5. Successful versus unsuccessful overtake, as defined in Section V-A. (Left plot) Unsuccessful overtake where the controlled bot runs outside the track, passes through the moving obstacle, or drives in the clockwise direction. (Middle plot) Unsuccessful overtake where the final position is not ahead of the moving obstacle's collision region within specified time. (Right plot) Successful overtake that avoids undesirable situations in the left and the middle plots.

1) *Less number of rollouts (efficient exploration)*: The proposed o-MPPI method requires **20-times** less number of rollouts compared to the standard MPPI, as seen in

TABLE II  
O-MPPI SUCCESS RATE (%) FOR **Case 1** IN SECTION V-A.

M	50	100	200
success rate	100	100	100

Tables III and II, In particular, to achieve 100% success rate of overtaking, the standard MPPI requires  $M = 1000$  rollouts with a prediction horizon  $T = 8.0$  s while the o-MPPI method needs  $M = 50$  rollouts with a prediction horizon  $T = 2.0$  s. Thus, the proposed o-MPPI achieves a more efficient exploration compared to standard MPPI.

2) *Smaller prediction horizon*: The o-MPPI method requires 4-times smaller prediction horizon compared to the standard MPPI as seen in Table III and II. In particular, to achieve 100% success rate in overtaking, the standard MPPI requires a prediction horizon  $T = 8.0$  s while the inversion-based sampling method needs only  $M = 50$  rollouts with a prediction horizon  $T = 2.0$  s.

When the standard MPPI method is restricted to a prediction horizon of 2.0 s, the controlled bot only achieves around 25% success rate, even with a large number of rollouts  $M$ . This is because o-MPPI has rollouts that run into the other lane as seen in the second left plot in Fig. 4. In contrast, rollouts of MPPI (rollout number  $M = 2000$ ) with a prediction horizon  $T = 2.0$  s are mostly restricted to within the original lane. Note that the explored area of the standard MPPI increases substantially with a larger prediction horizon of  $T = 8.0$  s as seen in Fig. 4.

In this sense, the proposed o-MPPI with a smaller prediction horizon requirement is beneficial since requiring a larger prediction horizon might not be feasible due to sensor-range limitations and it can also lead to increased computational load.

3) *MPPI performance is susceptible to the slow initial*: MPPI performance is susceptible to improper initial input sequences, even with sufficient prediction horizon  $T = 8.0$  s and number of rollouts  $M = 2000$ . For **Case 4** with slow-initial speed of 10 cm/s example, the controlled bot has difficulty overtaking the moving obstacle since the explored area has less overlap with the overtake regions, as seen in Fig. 4. Therefore, MPPI performance is affected by the initialization strategy and is susceptible to dynamic environments.

### B. Case with multiple dynamic obstacles

For comparative evaluation with multiple obstacles, two moving obstacles are considered as in the left plot of Fig. 3. The cost function is kept the same as in the previous Section V-A with a single dynamic obstacle.

The first obstacle has a constant speed of 10 cm/s (red circle) and tracks the inner lane with initial position  $(x, y) = (23 \text{ cm}, -156 \text{ cm})$  and the second obstacle has speed 12 cm/s (blue circle) and tracks the outer lane with initial position  $(x, y) = (-55 \text{ cm}, -75 \text{ cm})$ . The initial position of the controlled bot (green) is behind the first moving obstacle and has the initial condition  $X_0 = [x_0 \ y_0 \ \theta_0 \ v_0 \ \omega_0]^T =$

TABLE III  
MPPI-CONTROLLED BOT SUCCESS RATE (%) OF OVERTAKES FOR  
DIFFERENT PREDICTION HORIZON  $T$  AND NUMBER OF ROLLOUTS  $M$   
FOR **Cases 2 to 4** IN SECTION V-A

T	M	succ. rate	T	M	succ. rate
<b>Case 2</b>					
2.0	50	22	4.0	50	4
2.0	500	28	4.0	500	2
2.0	1000	15	4.0	1000	6
2.0	2000	22	4.0	2000	2
<b>Case 3</b>					
8.0	50	49	6.0	50	34
8.0	500	75	6.0	500	52
8.0	1000	100	6.0	1000	89
8.0	2000	100	6.0	2000	97
<b>Case 4</b>					
8.0	2000	1			

$[-55 \text{ cm} \quad -10 \text{ cm} \quad -\frac{\pi}{2} \text{ rad} \quad 15 \text{ cm/s} \quad 0 \text{ rad/s}]^T$  Comparative simulation and experiments of o-MPPI and standard MPPI were performed with settings as in **Case 1** (o-MPPI) and **Case 3** (Larger-horizon standard MPPI) proposed in Section V-A.

1) *Comparison of two obstacle avoidance*: Only the o-MPPI-controlled bot managed to successfully maneuver around both moving obstacles after slowing down because both lanes were blocked by the two slower constant-speed bots as seen in the trace plots in Fig. 6. First, the o-MPPI-controlled bot switched to the inner lane since the inner-lane, constant-speed bot has a relatively faster speed at 12 cm/s compared to the outer-lane, constant-speed bot moving at 10 cm/s and second then the o-MPPI-controlled bot switched back to the outer lane once there is sufficient room to drive at a faster speed that is closer to the desired speed 20 cm/s. In contrast, the standard MPPI-controlled bot got stuck behind the constant-speed bot, as seen from the trace in the upper region of the track in Fig. 6. As shown with **Case 4** (slow-initial standard MPPI), with a smaller speed the nominal mean of the MPPI input distribution become smaller, and can lead to lower success rate for overtaking. It is noted that both the o-MPPI-controlled bot and the MPPI-controlled bot were able to successfully maneuver around a single obstacle when the controlled bot is moving at a higher speed as seen in the accompanying video at <https://youtu.be/snhlZj3l5CE>.

2) *Oscillatory driving versus overtaking*: Comparison of traces of the experimental snapshots in Fig. 6 (left) shows that the MPPI-controlled bot can not only get stuck behind the slower moving obstacle bot (with speed 10 cm/s in the outer lane), but it also tends to have oscillatory driving behavior. This is because the oscillatory motion allows the bot to achieve a relatively faster speed closer to its desired speed in the cost function component  $c_v$  in Eq. (10). In contrast, the trace of the o-MPPI-controlled bot in Fig. 6 (right) show that it is able to maintain higher speed by

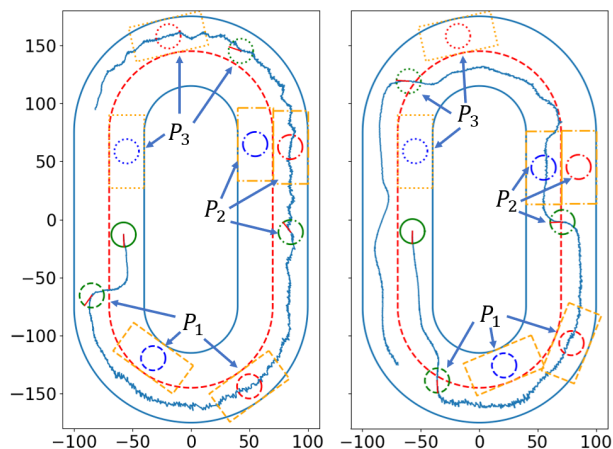


Fig. 6. Experimental traces of the MPPI-controlled bot (left) and o-MPPI-controlled bot (right). The o-MPPI-controlled bot could maneuver multiple times (snapshots of positions are indicated as  $P_1$  (dashed),  $P_2$  (dash-dotted), and  $P_3$  (dotted plots)) to overtake two constant-speed moving bots. However, the MPPI-controlled bot got stuck behind the constant-speed moving bot (blue) although it switched to the outer lane initially (at position  $P_1$  (dashed plot)) to avoid a single moving bot. Note that the MPPI-controlled bot moved in a zig-zag manner, as seen in the top right of the left plot. The full experiments can be seen in the accompanying video at <https://youtu.be/snhlZj3l5CE> and additional details are in [26].

successfully switching to the inner lane since the obstacle bot in the inner lane in the front has a higher speed 12 cm/s compared to the bot in the outer lane with speed 10 cm/s.

## VI. CONCLUSION AND FUTURE WORKS

This paper proposed an output-sampling-based model predictive path integral control (o-MPPI) to improve the efficiency of standard MPPI. An advantage of the proposed output sampling is that it can leverage the substantial work on trajectory planning in robotics to meet constraints that are often posed in the output space, which in turn improves the efficiency of MPPI. Instead of forward models that map from input to output as in standard MPPI, the proposed o-MPPI uses inverse models to map from the sampled output to inputs. The improved efficiency of the o-MPPI was seen in both the simulation and the experimental results — o-MPPI required less number of rollouts and a shorter prediction horizon, compared to the standard MPPI. Future works will explore the use of other trajectory-planning methods for output sampling with o-MPPI, as well as the use of data-enabled deep-learning and Gaussian process models for the inverse map. A blending of the input sampling and output sampling can also be considered depending on the type of constraints. Our ongoing theory work seeks to connect the proposed o-MPPI and stochastic optimal control.

## VII. ACKNOWLEDGEMENT

The authors would like to thank Sarmad Hassan and Digjay Das for implementing the controller for the constant-speed moving bots, Mohamed Safwat and Henry Chang for help with camera system and Robot Operating System (ROS).

## REFERENCES

- [1] Grady Williams, Andrew Aldrich, and Evangelos A Theodorou. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40(2):344–357, 2017.
- [2] Ihab S Mohamed, Guillaume Allibert, and Philippe Martinet. Model predictive path integral control framework for partially observable navigation: A quadrotor case study. In *2020 16th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pages 196–203. IEEE, 2020.
- [3] Ji Yin, Zhiyuan Zhang, Evangelos Theodorou, and Panagiotis Tsiotras. Trajectory distribution control for model predictive path integral control using covariance steering. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 1478–1484. IEEE, 2022.
- [4] Chuyuan Tao, Hunmin Kim, and Naira Hovakimyan. RRT-guided model predictive path integral method. *arXiv preprint arXiv:2301.13143*, 2023.
- [5] Taekyung Kim, Gyuhyun Park, Kiho Kwak, Jihwan Bae, and Wonsuk Lee. Smooth model predictive path integral control without smoothing. *IEEE Robotics and Automation Letters*, 7(4):10406–10413, 2022.
- [6] Mohak Bhardwaj, Balakumar Sundaralingam, Arsalan Mousavian, Nathan D Ratliff, Dieter Fox, Fabio Ramos, and Byron Boots. Storm: An integrated framework for fast joint-space model-predictive control for reactive manipulation. In *Conference on Robot Learning*, pages 750–759. PMLR, 2022.
- [7] Ihab S Mohamed, Kai Yin, and Lantao Liu. Autonomous navigation of agvs in unknown cluttered environments: log-mppi control strategy. *IEEE Robotics and Automation Letters*, 7(4):10240–10247, 2022.
- [8] Dylan M Asmar, Ransalu Senanayake, Shawn Manuel, and Mykel J Kochenderfer. Model predictive optimized path integral strategies. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3182–3188. IEEE, 2023.
- [9] Raphael Kusumoto, Luigi Palmieri, Markus Spies, Akos Csizsar, and Kai O Arras. Informed information theoretic model predictive control. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2047–2053. IEEE, 2019.
- [10] Mrinal Kalakrishnan, Sachin Chitta, Evangelos Theodorou, Peter Pastor, and Stefan Schaal. Stomp: Stochastic trajectory optimization for motion planning. In *2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4569–4574. IEEE, 2011.
- [11] Liang Ma, Jianru Xue, Kuniaki Kawabata, Jihua Zhu, Chao Ma, and Nanning Zheng. Efficient sampling-based motion planning for on-road autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 16(4):1961–1976, 2015.
- [12] Santosh Devasia, Degang Chen, and Brad Paden. Nonlinear inversion-based output tracking. *IEEE Transactions on Automatic Control*, 41(7):930–942, 1996.
- [13] Franziska Meier, Daniel Kappler, Nathan Ratliff, and Stefan Schaal. Towards robust online inverse dynamics learning. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4034–4039. IEEE, 2016.
- [14] Athanasios S Polydoros, Lazaros Nalpantidis, and Volker Krüger. Real-time deep learning of robotic manipulator inverse dynamics. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3442–3448. IEEE, 2015.
- [15] Diego Romeres, Mattia Zorzi, Raffaello Camoriano, Silvio Traversaro, and Alessandro Chiuso. Derivative-free online learning of inverse dynamics models. *IEEE Transactions on Control Systems Technology*, 28(3):816–830, 2019.
- [16] Leon Liangwu Yan and Santosh Devasia. Precision data-enabled koopman-type inverse operators for linear systems. *IFAC-PapersOnLine*, 55(37):181–186, 2022.
- [17] Marc Toussaint. A tutorial on newton methods for constrained trajectory optimization and relations to slam, gaussian process smoothing, optimal control, and probabilistic inference. *Geometric and numerical foundations of movements*, pages 361–392, 2017.
- [18] Julius Jankowski, Lara Bruder Müller, Nick Hawes, and Sylvain Calinon. Vp-sto: Via-point-based stochastic trajectory optimization for reactive robot behavior. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10125–10131. IEEE, 2023.
- [19] Oktay Arslan and Panagiotis Tsiotras. Machine learning guided exploration for sampling-based motion planning algorithms. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2646–2652. IEEE, 2015.
- [20] Ahmed H Qureshi, Anthony Simeonov, Mayur J Bency, and Michael C Yip. Motion planning networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2118–2124. IEEE, 2019.
- [21] Jacob J Johnson, Linjun Li, Fei Liu, Ahmed H Qureshi, and Michael C Yip. Dynamically constrained motion planning networks for non-holonomic robots. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6937–6943. IEEE, 2020.
- [22] Jiankun Wang, Wenzheng Chi, Chenming Li, Chaqun Wang, and Max Q-H Meng. Neural RRT\*: Learning-based optimal path planning. *IEEE Transactions on Automation Science and Engineering*, 17(4):1748–1758, 2020.
- [23] Grady Williams, Paul Drews, Brian Goldfain, James M Rehg, and Evangelos A Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.
- [24] Ji Yin, Zhiyuan Zhang, and Panagiotis Tsiotras. Risk-aware model predictive path integral control using conditional value-at-risk. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7937–7943. IEEE, 2023.
- [25] Zhimin Zhang, John Tomlinson, and Clyde Martin. Splines and linear control theory. *Acta Applicandae Mathematica*, 49:1–34, 1997.
- [26] Leon Yan. Inverse models for trajectory control aided by data, machine learning models, and GPUs. *PhD Thesis, U. of Washington*, 2024.