

# Path Integral Control with Rollout Clustering and Dynamic Obstacles

Steven Patrick<sup>1</sup> and Efstathios Bakolas<sup>1</sup>

**Abstract**—Model Predictive Path Integral (MPPI) control has proven to be a powerful tool for the control of uncertain systems (such as systems subject to disturbances and systems with unmodeled dynamics). One important limitation of the baseline MPPI algorithm is that it does not utilize simulated trajectories to their fullest extent. For one, it assumes that the average of all trajectories weighted by their performance index will be a safe trajectory. In this paper, multiple examples are shown where the previous assumption does not hold, and a trajectory clustering technique is presented that reduces the chances of the weighted average crossing in an unsafe region. Secondly, MPPI does not account for dynamic obstacles, so the authors put forward a novel cost function that accounts for dynamic obstacles without adding significant computation time to the overall algorithm. The novel contributions proposed in this paper were evaluated with extensive simulations to demonstrate improvements upon the state-of-the-art MPPI techniques.

## I. INTRODUCTION

For autonomous agents to be useful in unstructured environments, motion planning algorithms [1] must be used to ensure avoidance of obstacles (both static and dynamic) at all times. Additionally, the algorithms must be robust to a variety of different disturbances introduced by the real world: actuation noise, measurement error, process noise, and unmodeled dynamics.

A class of trajectory optimization algorithms that has been developed to address a subset of these sources of uncertainty is Model Predictive Path Integral (MPPI) control [2]. The key idea for MPPI control is to simulate multiple sequences of control inputs over a given time horizon and initial condition to determine an optimal control sequence. With each simulation, a realization from a random variable is used to perturb the original input sequence. The resulting trajectory is evaluated with a cost function, and then all of the results from the independent simulations are combined to create a control input sequence robust to the real-world disturbances. The benefits of this algorithm are the speed of planning, robustness to control noise, and versatility to different agent dynamics [2]. MPPI has been applied to several robotics problems, including racing on a dirt track [3], map-less navigation [4], drone flights [5], and manipulation [6] to name a few.

*Literature Review:* Several improvements upon the baseline of MPPI have been proposed. One category of improvements change how the samples are generated. Instead of using a handcrafted static distribution to sample from, other works have used non-static distributions [7] or learned




Fig. 1. Example of standard MPPI failing. Trajectories with low cost, and therefore high value, are separated by a region with high cost trajectory. The resulting weighted average used by MPPI is in the high cost, low value, region.

distributions [8] to generate trajectories. Another improvement is narrowing the search space of control inputs using control barrier functions [9]. A major improvement from the original MPPI is accounting for noise that does not enter the system dynamics through the control input channels. In [10], the authors account for process noise by having a nominal system and a real system. This allows for a larger exploration space of the MPPI algorithm, and it highlights how sampling from distributions other than the standard Gaussian produces better results. In [11], process noise was added to the simulated trajectories and the likelihood of the state disturbance was directly accounted for in the cost function. Even though their system was shown to be more robust than other MPPI variants, the number of trajectories to be simulated was significantly greater. The increase in simulated trajectories means the algorithm can only be applied to systems with a GPU to plan in real-time.

Even with these improvements, there are still cases where MPPI performs poorly. Figure 1 illustrates one such problem where the standard weighted average scheme results in a control input that is in a region of low value for the associated cost function. Another disadvantage of using the baseline MPPI method is it was developed for static environments. When applied to dynamic environments, quick replanning is used to account for moving obstacles. However, this type of planning is likely to fail if the obstacles move quickly or if the agent is close to the moving obstacle. Both cases are likely to result in a collision since the agent is planning under the false assumption that the world is static.

*Contributions:* To account for both of these issues, two novel methods are proposed to augment MPPI (in fact, the proposed improvements can be applied to any state-of-the-art variant of MPPI). The first is based on clustering

\*This research has been supported in part by NSF award ECCS-1924790.

<sup>1</sup>The University of Texas at Austin, Austin, Texas 78712-1221 spatric5@utexas.edu, and bakolas@austin.utexas.edu

trajectories and using MPPI only on clusters of trajectories. The second is sampling trajectories from moving obstacles and incorporating their probability of occurring into the cost of the trajectory. The approach in this paper would be similar to [11] if they augmented the state of the system with dynamic obstacles but with a key difference. In this paper, the cost function is parameterized by simulated obstacle trajectories which reduces the computational load to account for dynamic obstacles.

*Outline:* In Section II, a brief review of the baseline MPPI algorithm is made. After that, the proposed improvements upon the baseline MPPI algorithm in Section III are denoted. Then, the improvements are demonstrated in multiple environments through non-trivial simulations in Section IV. Finally, in Section V the results are discussed and future improvements upon the proposed algorithm are expressed.

## II. BACKGROUND

### A. Model Predictive Path Integral (MPPI)

Consider an agent whose equations of motion can be described by the following discrete-time state space model:

$$\mathbf{x}_{i+1} = f(\mathbf{x}_i, \mathbf{u}_i), \quad (1)$$

where  $\mathbf{x}_i \in \mathbb{R}^n$  is the current state at time step  $i$  and  $\mathbf{u}_i \in \mathbb{R}^m$  is the control input to be applied at the same time step. MPPI uses Eq. (1) to simulate several possible trajectories given a sequence of control inputs over a time horizon,  $N$ . The variable  $U$  is used to denote a sequence of control inputs  $U = \{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}\}$  whose application results in a state sequence,  $X = \{\mathbf{x}_0, \dots, \mathbf{x}_N\}$ , for a given initial condition,  $\mathbf{x}_0$ . In addition,  $\tau = (X, U)$  denotes the combined state and input sequences and is referred to as an agent's trajectory.

In the real world, Eq. (1) is an inaccurate model due to assuming the control inputs are applied without noise. MPPI accounts for this by assuming the disturbance to desired control input is effected by additive noise:  $\mathbf{v}_i = \mathbf{u}_i + \epsilon_i$  where  $\epsilon_i$  is a random normal variable with positive definite variance,  $\Sigma$ , and zero mean. The realization of this control input is determined by the distribution,  $\mathbb{P}$ , referred to as the uncontrolled distribution, whose probability density function (PDF) is given by

$$p(V) = \prod_{i=0}^{N-1} ((2\pi)^m |\Sigma|)^{-1/2} \exp\left(-\frac{1}{2} \mathbf{v}_i^T \Sigma^{-1} \mathbf{v}_i\right) \quad (2)$$

$$\text{where } V = \{\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{N-1}\}. \quad (3)$$

As seen in Eq. (2), the mean is zero because no control input has been applied to the system, which is the reason why the distribution  $\mathbb{P}$  is referred to as the uncontrolled distribution. By contrast, the controlled distribution, which is denoted as  $\mathbb{Q}$ , has the noiseless control input as its mean. The PDF of  $\mathbb{Q}$  is therefore given by

$$q(V) = \prod_{i=0}^{N-1} ((2\pi)^m |\Sigma|)^{-1/2} \exp\left(-\frac{1}{2} (\mathbf{v}_i - \mathbf{u}_i)^T \Sigma^{-1} (\mathbf{v}_i - \mathbf{u}_i)\right). \quad (4)$$

To determine the optimal control input given these potential disturbances, MPPI minimizes the expectation of a cost function over  $\mathbb{P}$ . In particular,

$$U^* = \arg \min_U \mathbb{E}_{\mathbb{P}}(J(V)) \quad (5)$$

$$\begin{aligned} J(V) &= \phi(\mathbf{x}_N) + \sum_{i=0}^{N-1} \psi(\mathbf{x}_i, \mathbf{v}_i) \\ &= \phi(\mathbf{x}_N) + \sum_{i=0}^{N-1} \psi(\mathbf{x}_i, \mathbf{u}_i + \epsilon_i) \end{aligned} \quad (6)$$

where  $\phi$  is the terminal cost function, and  $\psi$  is the running cost function. In [12], it was shown that solving the stochastic optimal control problem given in (5) is equivalent to finding the minimum KL-Divergence between the controlled and uncontrolled distributions. The relationship was used to derive a new distribution,  $\mathbb{Q}^*$ , for importance sampling defined by the following PDF:

$$q^*(V) = \frac{1}{\eta} \exp\left(-\frac{1}{\lambda} J(V)\right) p(V), \quad (7)$$

where  $\eta > 0$  is a normalizing constant such that the integral of the PDF,  $q^*(V)$ , over the sample space is 1. The sensitivity parameter,  $\lambda > 0$ , allows the user to set how important differences in cost between two trajectories are. A large  $\lambda$  results in low sensitivity to cost differences and small  $\lambda$  produce high sensitivity.

This new distribution allows for the control input to be defined as

$$\mathbf{u}_i = \int q^*(V) \mathbf{v}_i dV \quad (8)$$

With Eq. (8), importance sampling is used to optimize the number of samples needed to get a reliable approximation of the optimal control input. To accomplish this, Eq. (8) is multiplied by PDFs that are strictly positive whenever  $q^*(V) \mathbf{v}_i \neq 0$ . Clearly, Eq. (2) and (4) both satisfy this condition and using them for importance sampling results in the following

$$\mathbf{u}_i^* = \int \frac{q^*(V)}{p(V)} \frac{p(V)}{q(V)} q(V) \mathbf{v}_i dV. \quad (9)$$

The next step is to define the so-called weighting function, which is denoted as  $w(V)$  and defined as

$$w(V) = \frac{q^*(V)}{p(V)} \frac{p(V)}{q(V)}. \quad (10)$$

Substituting in the weighting function into Eq. (9) results in the new optimal control input calculated by

$$\mathbf{u}_i^* = \mathbb{E}_{\mathbb{Q}}[w(V) \mathbf{v}_i]. \quad (11)$$

The ratios  $q^*(V)/p(V)$  and  $p(V)/q(V)$  used to define  $w(V)$  in Eq. (10) are known and given by

$$\frac{q^*(V)}{p(V)} = \frac{1}{\eta} \exp\left(-\frac{1}{\lambda} J(V)\right) \quad (12)$$

$$\frac{p(V)}{q(V)} = \exp\left(\sum_{i=0}^{N-1} \frac{1}{2} \mathbf{u}_i^T \Sigma^{-1} \mathbf{u}_i - \mathbf{v}_i^T \Sigma^{-1} \mathbf{v}_i\right). \quad (13)$$

Substituting Eq. (12) and (13) into Eq. (10) results in

$$Q(V) = \sum_{i=0}^{N-1} \frac{1}{2} \mathbf{u}_i^{(k)T} \Sigma^{-1} \mathbf{u}_i^{(k)} - \mathbf{v}_i^{(k)T} \Sigma^{-1} \mathbf{v}_i^{(k)} \quad (14)$$

$$w(V) = \frac{1}{\eta} \exp \left( -\frac{1}{\lambda} J(V) + Q(V) \right). \quad (15)$$

In Eq. (15),  $\eta$  is the only remaining unknown. However, it is computationally intractable to determine the normalizing constant for the PDF in Eq. (7). Instead, a Monte-Carlo approach [9] is used wherein  $K$  realizations of control disturbances are sampled.

$$\eta \approx \sum_{k=1}^K \exp \left( -\frac{1}{\lambda} J(V_k) + Q(V_k) \right). \quad (16)$$

The  $k$ -th sample,  $\mathcal{E}_k$ , is used to generate a sequence of noisy control inputs  $V_k$  from a noiseless control reference  $U_k$ :

$$\mathcal{E}_k = \{\epsilon_0^{(k)}, \dots, \epsilon_{N-1}^{(k)}\} \quad (17)$$

$$U_k = \{\mathbf{u}_0^{(k)}, \dots, \mathbf{u}_{N-1}^{(k)}\} \quad (18)$$

$$V_k = U_k + \mathcal{E}_k. \quad (19)$$

In theory,  $U_k$ , can be sampled arbitrarily from the allowed bounds of the control inputs. In practice, it is more efficient to use a single reference control input  $U_0$  and weight the perturbations with the following equations

$$R(\mathcal{E}_k) = \sum_{i=0}^N \frac{1}{2} \mathbf{u}_i^{(0)T} \Sigma^{-1} (\mathbf{u}_i^{(0)} + 2\epsilon_i^{(k)}) \quad (20)$$

$$w(\mathcal{E}_k) = \frac{1}{\eta} \exp \left( -\frac{1}{\lambda} J(U_0 + \mathcal{E}_k) - R(\mathcal{E}_k) \right) \quad (21)$$

$$\eta = \sum_{k=1}^K \exp \left( -\frac{1}{\lambda} J(U_0 + \mathcal{E}_k) - R(\mathcal{E}_k) \right) \quad (22)$$

$$u_i^* = \mathbf{u}_i^{(0)} + \sum_{k=1}^K w(\mathcal{E}_k) \epsilon_i^{(k)} \quad (23)$$

It can be seen that Eq. (14)-(16) is equivalent to Eq. (20)-(22) when all  $U_k$  are equal. The use of Eq. (23) to determine the final control input sequence can be seen in Algorithm 1. In this algorithm, there is an extra variable  $\rho$  not seen in Eq. (20)-(23). The purpose of this variable is to render the algorithm more stable by preventing rounding errors in calculating the exponent of large negative numbers [10].

---

**Algorithm 1** MPPI Control Algorithm

---

**Require:**  $U_0, \{\mathcal{E}_1, \dots, \mathcal{E}_K\}, \{S_1, \dots, S_K\}, \lambda \triangleright$  Reference

Input, Perturbations, Cost of Trajectories, Sensitivity

- 1:  $\rho \leftarrow \min(S_1, \dots, S_K)$
  - 2:  $\eta \leftarrow \sum_{k=1}^K \exp \left( -\frac{1}{\lambda} (S_k - \rho) \right)$
  - 3: **for**  $k \leftarrow 1$  to  $K$  **do**
  - 4:    $w_k = \frac{1}{\eta} \exp \left( -\frac{1}{\lambda} (S_k - \rho) \right)$
  - 5: **end for**
  - 6:  $U \leftarrow U_0 + \sum_{k=1}^K w_k \mathcal{E}_k$
- 

### III. METHODOLOGY

#### A. Rollout Clustering

As mentioned in Section I, MPPI can perform poorly when the value function has multiple local maxima. To address this limitation of MPPI, clustering rollouts into groups that do not contain sharp discontinuities in the cost function is proposed. The authors' chosen clustering algorithm is a density-based method [13] called DBSCAN [14]. The algorithm is a natural choice for this problem because perturbations that are close together and produce similar costs will be clustered together, but sharp changes in the cost will cause a new cluster boundary to form. Clustering in this manner prevents undesirable valleys separating high-points of the value function from affecting the expected value computation. Since the agent does not know the number of valleys or high points along the value function, this method was also chosen because it allows for a dynamic number of clusters unlike K-means [15].

DBSCAN works by sequentially building clusters as the data comes in. For a new data point, all previous clusters are checked against a ball of radius  $\epsilon$  centered at the new data point. If there are any intersections, the point is either added to an existing cluster, or multiple clusters are merged that intersect with the ball. If no intersections are formed, then the data-point is its own cluster. This process is repeated until all the data points have been processed. The data point,  $\mathbf{d}_k$ , for the  $k$ -th rollout are the perturbations from the control input,  $\mathcal{E}$ , and their associated trajectory cost,  $J(V)$ , from Eq. (6):  $\mathbf{d}_k = \{\epsilon_{N-1}^{(k)}, \dots, \epsilon_0^{(k)}, J(V_k)\}$ . With the clustered data-points, a new distribution can be chosen to derive the importance sampling weights. This distribution should have a high probability for trajectories within a cluster but a low probability for trajectories outside of the cluster. However, this distribution has to satisfy two properties to be a valid distribution for importance sampling. The first being it has to have a valid probability density function. The second is that the probability density function must be strictly positive where the original PDF is non-zero.

To satisfy the two conditions and get the desired attributes, a truncated Gaussian is used as a distribution for importance sampling. The Gaussian from Eq. (7) is truncated such that all noise realizations within a cluster evaluate to non-zero values much greater than realizations outside of the cluster. The resulting PDF is defined as

$$y(V) = \begin{cases} v \exp \left( -\frac{1}{\lambda} J(V) \right) p(V), & V \in \Omega \\ \sigma \exp \left( -\frac{1}{\lambda} J(V) \right) p(V), & V \in \Omega^c \end{cases} \quad (24)$$

where  $\Omega$  is a simply connected compact region that encompasses all of the samples of a given cluster, and  $\Omega^c$  is its complement. Additionally,  $v \gg \sigma > 0$  are normalizing constants for  $y(V)$  to be a valid probability density function.

#### Proposition 1

The truncated normal distribution,  $y(V)$ , is a valid importance sampling distribution.

*Proof:* Found in Appendix. ■




Fig. 2. Example of standard MPPI (left) producing an undesirable result and Clustered MPPI (right) producing multiple valid solutions. The blue solid line is the cost function, the X marks are sampled points, and the filled in stars are the result of the weighted average.

*Remark 1:* It can be seen that the weighting function for the truncated Gaussian distribution can be calculated using the weights from the original Gaussian distribution. Therefore, no additional evaluations of the cost function or extra simulations are required to find the importance sampling weights of a subset of trajectories. It is therefore unnecessary to calculate these parameters as well. Note that the clustered MPPI approach adds little computational burden to the baseline MPPI algorithm.

*Remark 2:* Just like in the original MPPI, the normalizing constant,  $v$  must be estimated with Monte-Carlo methods similar to Eq. (16). This is done using the following equation

$$\eta \approx \sum_{V_k \in \Omega} \exp\left(-\frac{1}{\lambda}J(V_k) + Q(V_k)\right) \quad (25)$$

The other normalizing constant,  $\sigma$ , is not estimated since it is chosen to be arbitrarily close to 0.

---

#### Algorithm 2 Clustered MPPI

---

**Require:**  $f, \Sigma$   $\triangleright$  Difference Equation, Noise Covariance  
**Require:**  $\mathbf{x}_0, U_0$   $\triangleright$  Initial Condition, Control Reference  
**Require:**  $K, N$   $\triangleright$  # of Samples, # of Time Steps  
**Require:**  $\lambda$   $\triangleright$  Cost Sensitivity  
**Require:**  $\psi, \phi$   $\triangleright$  Running Cost, Terminal Cost

```

1: for  $k \leftarrow 1$  to  $K$  do  $\triangleright$  In Parallel
2:    $\tilde{\mathbf{x}}_0^{(k)} \leftarrow \mathbf{x}_0$ 
3:   for  $i \leftarrow 0$  to  $N - 1$  do
4:      $\epsilon_i^{(k)} \leftarrow \text{sample } \mathcal{N}(\mathbf{0}, \Sigma)$ 
5:      $\tilde{\mathbf{x}}_{i+1}^{(k)} \leftarrow f(\tilde{\mathbf{x}}_i, \mathbf{u}_i + \epsilon_i^{(k)})$ 
6:      $S_k \leftarrow \psi(\tilde{\mathbf{x}}_i^{(k)}) + \lambda \mathbf{u}_i^T \Sigma^{-1} \epsilon_i^{(k)}$ 
7:   end for
8:    $S_k \leftarrow \phi(\tilde{\mathbf{x}}_N^{(k)})$ 
9:    $X_k, \mathcal{E}_k \leftarrow \{\tilde{\mathbf{x}}_0^{(k)}, \dots, \tilde{\mathbf{x}}_N^{(k)}\}, \{\epsilon_0^{(k)}, \dots, \epsilon_{N-1}^{(k)}\}$ 
10: end for
11:  $C_1, \dots, C_M \leftarrow \text{dbscan}(\{\mathcal{E}_0, S_0\}, \dots, \{\mathcal{E}_K, S_K\})$ 
12: for  $m \leftarrow 1$  to  $M$  do
13:    $U^{(m)} \leftarrow \text{MPPI}(U_0, \mathcal{E}_{C(m)}, S_{C(m)}, \lambda)$   $\triangleright$  Alg. 1
14: end for
15:  $U \leftarrow \arg \min_{U^{(m)}} J(U^{(m)})$   $\triangleright$  Eq. (6)
```

---

The pseudo-code for Clustered MPPI is given in Algo-

rithm 2. Starting in Line 1, the **for** loop simulates trajectories by sampling perturbations to the reference control input and then storing the cost, trajectory, and associated noise. After the **for** loop in Line 11, the trajectory cost is concatenated with its associated noise vector and the DBSCAN algorithm is performed over the  $K$  data points. The result is an index set,  $C$ , defining  $M$  clusters from the original samples. The **for** loop in Line 14 calculates the new control input for each cluster using Algorithm 1. Finally, in Line 15, the cluster which produces the minimum cost using Eq. (6) is selected as the control input.

An illustrative example of the difference in output between the standard MPPI algorithm versus a clustered approach can be seen in Figure 2. The  $y$ -axis is the value, also known as the negative exponential of the cost  $J(V)$ , and the  $x$ -axis is a deviation from a reference input. Troughs in the value function can be encountered when the reference trajectory has the agent going straight through an obstacle. In this example, standard MPPI fails because the weighted average between two high areas lies within the valley. In contrast, Clustered MPPI separates these two high areas and the low area into separate clusters and performs the weighted average only including the points in the cluster. This produces two valid selections of control input and one invalid.

#### B. Dynamic Obstacles

To account for dynamic obstacles, terminal and running cost functions that are parameterized by realizations of obstacle trajectories are proposed. This is in contrast to methods like [11] where each control input deviation has an independent sampling of obstacle trajectories. This leads to a multiplicative increase in computation time for  $N$  perturbations to the control input and  $P$  realizations of obstacle trajectories:  $O(PN)$ . The proposed approach is additive where the  $P$  realizations are only done once which results in an additive increase in computation time:  $O(P + N)$ .

To reiterate, the  $l$ -th obstacle has  $P$  simulated trajectories collected in the set  $\mathcal{O}_l$ . The  $p$ -th trajectory of the  $l$ -th obstacle,  $\tau_p^l$ , consists of the obstacles state over the time

horizon.

$$\tau_p^l = \{\mathbf{o}_0^{l,p}, \dots, \mathbf{o}_N^{l,p}\} \quad (26)$$

$$\mathcal{O}_l = \{\tau_1^l, \dots, \tau_P^l\} \quad (27)$$

The collection of  $L$  dynamic obstacles trajectories,  $\mathcal{O}$ , is used to create a subset of the configuration space at each time step where the controlled agent is considered in collision.

$$\mathcal{O} = \{\mathcal{O}_1, \dots, \mathcal{O}_L\} \quad (28)$$

If a collision occurs between the controlled agent and the  $p$ -th trajectory of the  $l$ -th obstacle, then the binary function  $\mathbf{1}_p^l$  equals 1 and 0 otherwise. The associated trajectory also has a function,  $\theta$ , that outputs the probability of the trajectory occurring. For example, this probability can be derived from Extended Kalman Filter estimate of the obstacle's position and heading of a differential drive robot. The new running and terminal cost functions are defined as

$$\psi(\mathbf{x}_i|\mathcal{O}) := \psi(\mathbf{x}_i) + \beta \sum_{l=1}^L \sum_{p=1}^P \theta(\tau_p^l) \mathbf{1}_p^l(\mathbf{x}_i, t_i) \quad (29)$$

$$\phi(\mathbf{x}_i|\mathcal{O}) := \phi(\mathbf{x}_i) + \beta \sum_{l=1}^L \sum_{p=1}^P \theta(\tau_p^l) \mathbf{1}_p^l(\mathbf{x}_i, t_i) \quad (30)$$

In Eq. (29) and (30),  $\beta \in \mathbb{R}^+$  is a scalar weight associated with hitting an obstacle. With this formulation, the cost function  $J(V)$  only depends on the agent's state after the obstacles' trajectories have been sampled. This has the aforementioned benefit of reducing the computational cost of accounting for dynamic obstacles over other state-of-the-art methods. The combination of clustered MPPI with simulating dynamic obstacles can be seen in Algorithm 3.

---

#### Algorithm 3 DC-MPPI

---

**Require:**  $f, \Sigma, \mathbf{x}_0, U_0, K, N, \lambda$   $\triangleright$  Alg. 2 Parameters

**Require:**  $\{\mathbf{o}_0^{(1)}, \dots, \mathbf{o}_0^{(L)}\}$   $\triangleright$  obstacles initial states

**Require:**  $\{U_o^{(1)}, \dots, U_o^{(L)}\}$   $\triangleright$  obstacle reference inputs

**Require:**  $g, \Sigma_o$   $\triangleright$  obstacle difference equation, obstacle input variance

```

1: for  $l \leftarrow 1$  to  $L$  do
2:   for  $p \leftarrow 1$  to  $P$  do
3:      $\{\mathbf{u}_0^{l,p}, \dots, \mathbf{u}_{N-1}^{l,p}\} \leftarrow U_o^{(l)}$ 
4:      $\mathbf{o}_0^{l,p} \leftarrow \mathbf{o}_0^{(l)}$ 
5:     for  $i \leftarrow 0$  to  $N-1$  do
6:        $\epsilon \leftarrow \text{sample } \mathcal{N}(0, \Sigma_o)$ 
7:        $\mathbf{o}_{i+1}^{l,p} \leftarrow g(\mathbf{o}_i^{l,p}, \mathbf{u}_i^{l,p} + \epsilon)$ 
8:     end for
9:      $\tau_p^l \leftarrow \{\mathbf{o}_0^{l,p}, \dots, \mathbf{o}_N^{l,p}\}$ 
10:   end for
11:    $\mathcal{O}_l \leftarrow \{\tau_1^l, \dots, \tau_P^l\}$ 
12: end for
13:  $\phi \leftarrow \phi(\{\mathcal{O}_1, \dots, \mathcal{O}_L\})$   $\triangleright$  Eq. (29)
14:  $\psi \leftarrow \psi(\{\mathcal{O}_1, \dots, \mathcal{O}_L\})$   $\triangleright$  Eq. (30)
15:  $U \leftarrow \text{Clust. MPPI}(f, \Sigma, \mathbf{x}_0, U_0, K, N, \lambda, \phi, \psi)$   $\triangleright$  Alg. 2

```

---

For clarity and compactness, the DC-MPPI pseudo-code is only written for one dynamic obstacle. The **for** loop starting in Line 1 simulates how a dynamic obstacle move over time. Lines 13 and 14 create a terminal and running cost function respectively based on the trajectories of the obstacles that were simulated. The resulting terminal and running cost function only rely on the agent's state. Finally, Line 14 sends all of the necessary data to Algorithm 2 to calculate the final control input.

The improvements of both clustering the trajectories and accounting for dynamic obstacles can be seen in Figure 4. The reference control input is for the agent to go straight forward at 1 m/s. The deviation from the control input is only in the heading. The dynamic obstacle is directly between the agent and the goal causing a valley in the value function as seen in all of the bottom plots of Figure 4. Only the baseline MPPI algorithm chooses a control input within this trough since it does not cluster. The other algorithms are able to choose control inputs that produce safe trajectories if the obstacle was stationary. However, the obstacle is dynamic. Therefore, the region of bad trajectories expands in the DC-MPPI algorithm, but not for the Clustered algorithm. This results in the DC-MPPI algorithm getting out of the way of the obstacle rather than going towards the goal like the clustered algorithm and the baseline.

## IV. EXPERIMENTS

All simulations were performed in MATLAB 2022b on a laptop with Intel i7 processor and 32GB of RAM. MATLAB's implementation of DBSCAN was used for clustering. Furthermore, MATLAB's ode45 was used to rollout trajectories.

### A. Static Obstacles

For the first experiment, an agent with Dubins car dynamics is maneuvering in a 2D plane and avoiding randomly placed obstacles. The state of the agent is its position,  $(x, y)$ , and its direction of motion,  $\theta$ . The dynamics of the system are given by

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} v \cos(\theta) \\ v \sin(\theta) \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \omega. \quad (31)$$

The linear velocity,  $v$ , is constant and positive, and the agent only has control over the angular velocity,  $\omega$ . When choosing  $\omega$ , the agent must always satisfy  $0 < R_{\min} \leq v/\omega$  where  $R_{\min}$  is the minimum turning radius.

The running cost and terminal cost are equal to the distance of, respectively, the current position and terminal position of the system from the goal position  $(x_g, y_g)$  along with a penalty cost for being in collision with a static cost map:  $\mathbf{1}(\mathbf{x})$ . The collision cost is weighted by the hyper-parameter  $\alpha > 0$ .

$$\phi(\mathbf{x}) = \psi(\mathbf{x}) = \sqrt{(x - x_g)^2 + (y - y_g)^2} + \alpha \mathbf{1}(\mathbf{x}) \quad (32)$$




Fig. 3. Random forest environment example. The agent is the magenta dot, the goal is the green star, and the obstacles are light and dark red.

The proposed method was tested along with the standard MPPI formulation and another state-of-the-art method [10] using the authors' own implementation of each algorithm. The start and goal positions were randomly generated for locations that were not in collision of any static obstacles within a randomly generated environment. The obstacles are static and non-uniform in shape, and an example can be seen in Figure 3. The first iteration of MPPI is warm-started with a straightforward initial guess, but all subsequent MPPI calculations are initialized with the previous iteration's solution. 500 trajectories are simulated and the perturbation from the control reference is kept constant for the entire trajectory. The hyper-parameters were set  $\lambda = 1$ ,  $\alpha = 1,000$ ,  $\Sigma = 0.1$ . Lastly, the algorithms were tested with 3 different noise profiles: noiseless, input noise, process and input noise.

The algorithms were evaluated in 1000 different randomly generated environments and report the number of collision occurred, percentage failure which includes crashing and not reaching the goal in a timely manner. The authors define a timely manner as being less than the triple the time it would take the agent to go around the entire map without any obstacles. The results can be seen in Table I. For the static trials, DC-MPPI was not included since it produces the same output as Clustered MPPI when there are no moving obstacles. It can be seen that the proposed method reduces the number of collisions and number of failures for all noise configurations while adding minimal computation time.

### B. Dynamic Obstacles

For the dynamic environment, the agent has the same dynamics as the previous experiment. The state of the agent is its position  $(x, y)$  and heading  $\theta$  and respective velocities:  $v_x, v_y, v_\theta$ . The control inputs are  $u_x, u_y, u_\theta$ . The environment was populated with 100 dynamic obstacles. These dynamic obstacles have Dubins car dynamics from Eq. (31) and they all have a 1 meter collision radius. The starting locations of the obstacles are uniformly distributed over  $[0, 75] \times [0, 50] m^2$  space and the orientations are uniformly distributed over  $[0, 2\pi)$  rad. The linear and angular velocities are similarly distributed with  $v \in [0.5, 1.5]$  m/s and  $\omega \in$

| Test                      | Algorithm      | # Collisions | % Failure Time (ms) | Avg. Comp. |
|---------------------------|----------------|--------------|---------------------|------------|
| No Noise                  | Baseline       | 11           | 1.2                 | 190        |
|                           | Tube-MPPI [10] | 9            | 1.2                 | 364        |
|                           | Clustered      | 6            | 0.9                 | 193        |
| Control Noise             | Baseline       | 59           | 6.0                 | 195        |
|                           | Tube-MPPI [10] | 44           | 4.5                 | 376        |
|                           | Clustered      | 26           | 2.9                 | 199        |
| Control and Process Noise | Baseline       | 48           | 4.9                 | 355        |
|                           | Tube-MPPI [10] | 49           | 5.0                 | 376        |
|                           | Clustered      | 37           | 3.9                 | 360        |

TABLE I  
QUANTITATIVE PERFORMANCE OVER 1000 RUNS.

$[-0.5, 0.5]$  rad/s. The obstacles' velocities are not explicitly known to the agent, but a Gaussian distribution is known of the potential linear and angular velocities of the obstacles is known. This distribution is used to sample potential obstacle trajectories as seen in Figure 4. The first time step of the environment can be seen in Figure 5. Additionally, an example of how trajectories are sampled for the agent and obstacles and the samples effect on the value function of the separate algorithms can be seen in Figure 4.

The same hyper-parameters were used in this experiment as the previous one. For DC-MPPI, 25 samples of each obstacle were used and  $\beta = 10$ . Using the baseline MPPI algorithm resulted in 3 instances of obstacle collisions whereas DC-MPPI had no collisions. Without the dynamic obstacle sampling, the Clustered algorithm performed better than MPPI with 2 collisions. It can be seen in the video<sup>1</sup> that both the baseline and Clustered MPPI algorithms collide with moving obstacles when an obstacle passes through the reference path. It should also be noted that Clustered MPPI has marginal impact on the computational time required with an average 288 milliseconds per time step for Clustered MPPI compared to 283 milliseconds per time step for the baseline algorithm. This is in contrast to DC-MPPI which averaged 454 milliseconds. The increase in computation time is clearly worth the decrease in collisions. Additionally, the a-priori knowledge of the obstacles velocities does not diminish the applicability of this paper's novel algorithm to the real world. When applied to a real system, a simple Extended Kalman Filter [16] would be sufficient to generate a distribution to sample obstacle trajectories from.

### V. CONCLUSION

In this paper, two improvements upon MPPI were presented. The first of which is clustering trajectories so that the resulting control input sequence is prevented from producing unsafe trajectories. This improvement is specifically meant to separate clusters of good trajectories thereby preventing an unsafe average between neighboring clusterings. It was shown that the clustering step adds little computation time to the MPPI algorithm, and demonstrated clear cases where the baseline fails. The efficacy of DC-MPPI depends greatly

<sup>1</sup>Video link: <https://youtu.be/G68DbQ01ouM>






Fig. 4. Dynamic obstacle example where the agent is a magenta circle, the obstacle is a red circle traveling in the negative  $x$  direction, and the goal is the green star. Top left are all the simulated trajectories of the agent. Top middle are the simulated obstacle trajectories. Top right are the resulting paths from each algorithm. The bottom row are the plots of the value functions with respect to the deviation from the reference control input. Different clusters of trajectories are denoted in different colors. The resulting control input deviation is denoted with a yellow star.




Fig. 5. Dynamic environment example. Agent is magenta circle, obstacles are red, green and blue circles with their direction of motion indicated by black arrows. The goal is the green star. The yellow and cyan lines are the first iteration of DC-MPPI with the yellow being the reference input and the cyan being the result.

on the hyper-parameter selection for DBSCAN. For high-dimensional control input spaces, holes can appear in the middle of a cluster. Therefore, hyper-parameters should be selected where clusters stay relatively small ( $< 20\%$  of the trajectories) to prevent this from happening.

The second is creating a cost function that accounts for dynamic obstacles. Adding dynamic obstacles allows for MPPI to prevent collisions where fast replanning is not sufficient. In the experiments, it was shown that the proposed improvements reduced the number of collisions with a slight increase in the overall computation time. With a GPU speedup, this additional time would significantly reduce.

Importantly, the proposed improvements can be added to the majority of MPPI variants like [11] without much retooling. As seen in Algorithms 2 and 3, the novel algorithms can be considered a pre-processing step to the original MPPI algorithm.

## APPENDIX

*Proof:* Let one consider the distribution  $\mathbb{Y}$  whose probability density function is given by Eq. (24). For  $y(v)$  to be a valid PDF, the following must be true

$$1 = v \int_{\Omega} \exp\left(-\frac{1}{\lambda} J(V)\right) p(V) d\Omega + \sigma \int_{\Omega^c} \exp\left(-\frac{1}{\lambda} J(V)\right) p(V) d\Omega^c. \quad (33)$$

Because of the constraint in Eq. (33) on  $v$  and  $\sigma$ , the variables are dependent upon one another. A new importance sampling scheme using  $y(V)$  can be made through the following,

$$\begin{aligned} \mathbf{u}_i^* &= \int y(V) q(V) \mathbf{v}_i dV \\ &= \int \frac{y(V)}{p(V)} \frac{p(V)}{q(V)} q(V) \mathbf{v}_i dV \\ &= v \int_{\Omega} \frac{\exp\left(-\frac{1}{\lambda} J(V) p(V)\right)}{p(V)} \frac{p(V)}{q(V)} q(V) \mathbf{v}_i d\Omega \\ &\quad + \sigma \int_{\Omega^c} \frac{\exp\left(-\frac{1}{\lambda} J(V) p(V)\right)}{p(V)} \frac{p(V)}{q(V)} q(V) \mathbf{v}_i d\Omega^c \\ &= v \int_{\Omega} \exp\left(-\frac{1}{\lambda} J(V)\right) \frac{p(V)}{q(V)} q(V) \mathbf{v}_i d\Omega \\ &\quad + \sigma \int_{\Omega^c} \exp\left(-\frac{1}{\lambda} J(V)\right) \frac{p(V)}{q(V)} q(V) \mathbf{v}_i d\Omega^c \end{aligned} \quad (34)$$

The ratio  $p(V)/q(V)$  was defined in Eq. (13). Therefore, a new weighting function  $w^*(V)$  is used and defined as

$$w^*(V) = \begin{cases} vw(V), & V \in \Omega \\ \sigma w(V), & V \in \Omega^c \end{cases} \quad (35)$$

$$w(V) = \frac{1}{\rho} \exp\left(-\frac{1}{\lambda}J(V) + Q(V)\right) \quad (36)$$

$$\begin{aligned} \rho &= v \int_{\Omega} \exp\left(-\frac{1}{\lambda}J(V) + Q(V)\right) d\Omega \\ &+ \sigma \int_{\Omega^c} \exp\left(-\frac{1}{\lambda}J(V) + Q(V)\right) d\Omega^c \end{aligned} \quad (37)$$

$$Q(V) = \sum_{i=0}^{N-1} \frac{1}{2} \mathbf{u}_i^{(k)T} \Sigma^{-1} \mathbf{u}_i^{(k)} - \mathbf{v}_i^{(k)T} \Sigma^{-1} \mathbf{v}_i^{(k)}. \quad (38)$$

As previously mentioned, there is only 1 degree of freedom for the variables  $v$  and  $\sigma$ . Thus the limit can be taken with respect to  $\sigma$  approaching 0. Taking this limit with Eq. (34) results in

$$\lim_{\sigma \rightarrow 0} (\mathbf{u}_i^*) = v \int_{\Omega} w(V) q(V) \mathbf{v}_i d\Omega \quad (39)$$

To ensure  $y(V)$  is still a valid PDF, Eq. (33) must still hold resulting in

$$\begin{aligned} \frac{1}{v} &= \int_{\Omega} \exp\left(-\frac{1}{\lambda}J(V)\right) p(V) d\Omega \\ &+ \frac{\lim_{\sigma \rightarrow 0}(\sigma)}{v} \int_{\Omega^c} \exp\left(-\frac{1}{\lambda}J(V)\right) p(V) d\Omega^c \\ &= \int_{\Omega} \exp\left(-\frac{1}{\lambda}J(V)\right) p(V) d\Omega. \end{aligned} \quad (40)$$

$$\rho = v \int_{\Omega} \exp\left(-\frac{1}{\lambda}J(V) + Q(V)\right) d\Omega \quad (41)$$

With the normalizing constants derived, the final weighting for importance sampling within  $\Omega$  becomes

$$w^*(V) = \frac{1}{\eta} \exp\left(-\frac{1}{\lambda}J(V) + Q(V)\right) \quad (42)$$

$$\eta = \int_{\Omega} w^*(V) d\Omega \quad (43)$$

and 0 otherwise. Thus it has been shown that  $y(V)$  has the needed properties to be a valid PDF, and it is a PDF that is strictly positive when  $p(V)$  is non-zero. With all conditions satisfied for  $y(V)$  to be a valid PDF for importance sampling with respect to  $p(V)$ , the proof is complete. ■

## REFERENCES

- [1] S. M. LaValle, *Planning algorithms*. Cambridge University Press, 2006.
- [2] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *ICRA*. IEEE, 2016, pp. 1433–1440.
- [3] 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 *ICRA*, 2017, pp. 1714–1721.
- [4] I. S. Mohamed, K. Yin, and L. Liu, "Autonomous navigation of agvs in unknown cluttered environments: LOG-MPPI control strategy," *IEEE RA-L*, vol. 7, no. 4, pp. 10 240–10 247, 2022.
- [5] I. S. Mohamed, G. Allibert, and P. 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)*. IEEE, 2020, pp. 196–203.
- [6] M. Bhardwaj, B. Sundaralingam, A. Mousavian, N. D. Ratliff, D. Fox, F. Ramos, and B. Boots, "Storm: An integrated framework for fast joint-space model-predictive control for reactive manipulation," in *Conference on Robot Learning*. PMLR, 2022, pp. 750–759.
- [7] I. M. Balci, E. Bakolas, B. Vlahov, and E. A. Theodorou, "Constrained covariance steering based tube-MPPI," in *2022 American Control Conference (ACC)*, 2022, pp. 4197–4202.
- [8] J. Sacks and B. Boots, "Learning sampling distributions for model predictive control," in *6th Annual Conference on Robot Learning*, 2023.
- [9] C. Tao, H.-J. Yoon, H. Kim, N. Hovakimyan, and P. Voulgaris, "Path integral methods with stochastic control barrier functions," in *IEEE CDC*, 2022, pp. 1654–1659.
- [10] 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*, 2018.
- [11] J. Yin, Z. Zhang, and P. Tsiotras, "Risk-aware model predictive path integral control using conditional value-at-risk," *arXiv preprint arXiv:2209.12842*, 2022.
- [12] E. A. Theodorou and E. Todorov, "Relative entropy and free energy dualities: Connections to path integral and KL control," in *IEEE CDC*, 2012, pp. 1466–1473.
- [13] R. J. Campello, P. Kröger, J. Sander, and A. Zimek, "Density-based clustering," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 10, no. 2, p. e1343, 2020.
- [14] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *kdd*, vol. 96, 1996, pp. 226–231.
- [15] J. A. Hartigan and M. A. Wong, "Algorithm as 136: A k-means clustering algorithm," *Journal of the Royal Statistical Society. Series C (applied statistics)*, vol. 28, no. 1, pp. 100–108, 1979.
- [16] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.