

# Model Predictive Optimized Path Integral Strategies

Dylan M. Asmar,<sup>1</sup> Ransalu Senanayake,<sup>1</sup> Shawn Manuel,<sup>1</sup> and Mykel J. Kochenderfer<sup>1</sup>

**Abstract**—We generalize the derivation of model predictive path integral control (MPPI) to allow for a single joint distribution across controls in the control sequence. This reformation allows for the implementation of adaptive importance sampling (AIS) algorithms into the original importance sampling step while still maintaining the benefits of MPPI such as working with arbitrary system dynamics and cost functions. The benefit of optimizing the proposal distribution by integrating AIS at each control step is demonstrated in simulated environments including controlling multiple cars around a track. The new algorithm is more sample efficient than MPPI, achieving better performance with fewer samples. This performance disparity grows as the dimension of the action space increases. Results from simulations suggest the new algorithm can be used as an anytime algorithm, increasing the value of control at each iteration versus relying on a large set of samples.

**Repository**—<https://github.com/sisl/MPOPIs>

## I. INTRODUCTION

Many modern control applications benefit from modeling nonlinear dynamics of the system. Examples include vehicle collision avoidance [1], aerial maneuvers [2], and autonomous racing [3]. Model predictive control (MPC) has become a popular method to accommodate the nonlinearity of the dynamics [4]. Some MPC formulations use a quadratic objective function and linear constraints encoding an approximation of the dynamics. Other approaches involve solving the nonlinear program associated with the dynamics by using variants of sequential quadratic programming and interior point methods to approximate solutions [5]. Another technique that has emerged is the use of the Koopman operator [6]. Korda and Mezić [7] show that combining this approach with a linear MPC controller reduces the computational complexity of the underlying optimization problem to one that is comparable to a linear dynamical system of similar state and action space sizes. However, efficiently developing finite-dimensional representations of the linear dynamics is still an area of active research [8].

An alternative MPC-based approach is to use Monte Carlo sampling as part of the optimization of the trajectory. Model predictive path integral control (MPPI) is one such approach, which has been motivated from the perspective of the Hamilton-Jacobi-Bellman equation [9], information theory [10], [11], and stochastic optimization [12]. MPPI can work with arbitrary system dynamics and cost functions, and it uses importance sampling to estimate the optimal control input. Williams *et al.* [13] demonstrated MPPI to work in real time with a neural network representing the system dynamics. The main idea behind MPPI is to sample a set

of control sequences, compute the cost of each sequence, and then calculate a cost informed weighted average of the sampled controls. The first control in the sequence is executed and the remainder of the weighted average is used as the mean of the proposal distributions at the next time step.

The insight that the performance of MPPI is tightly coupled with sampling from a proposal distribution defined from the previous control sequence has led to research in different areas. One focus has been on increasing the robustness when a disturbance changes the state of the system. Such a disturbance often results in samples from the proposal distribution, which is based on the previous control sequence, to not accurately estimate the the optimal control. Williams *et al.* [14] augment MPPI with an ancillary controller which tracks the nominal trajectory forming Tube-MPPI. Gandhi *et al.* [15] further extend this approach and provide bounds on the free energy growth while rejecting unknown disturbances. This methodology still relies on developing a nominal control input from sampling separate distributions at each time step around the previous control sequence.

The formulation of MPPI structured the importance sampling to be from separate distributions at each control step. This formulation limits the form of the proposal distribution. In execution, we are limited to a finite number of samples and the quality of the estimate using importance sampling is influenced by the reference distribution [16]. This limitation often requires large sample sizes to approximate the optimal control and manifests in experiments when the dimension of the action space increases [17].

Kusumoto *et al.* [18] extend MPPI by adding an informed sampling process. They use conditional variational autoencoders to learn distributions that imitate samples from a training data set containing optimized controls. Using these learned distributions, they adjust the mean of the sampling distribution before executing the MPPI process. Okada and Taniguchi [19] use optimization methods originally developed for gradient descent to adjust the control input at each time step, increasing the convergence to the optimal control. Lambert *et al.* [20] introduce a Bayesian MPC approach that allows for multi-modal distributions and approximate the posterior distribution using Stein variational gradient descent. Pravitra *et al.* [21] implement an iteration step within the MPPI framework, but only adjust the mean of the sampling distributions from the weighted average of the previous MPPI iteration. Wang *et al.* [22] use the Tsallis divergence to form a general sampling based algorithm and use a similar iteration procedure to refine the distribution within the control update step. However, the sampling and

<sup>1</sup> Stanford Intelligent Systems Laboratory (SISL), Stanford University  
<asmard, ransalu, sman64, mykel>@stanford.edu

optimization are over separate distributions at each control step.

These works either adjust the control output after the sampling, adjust only the mean of the proposal distribution, or update individual distributions corresponding to the control at each time step. Our work extends these ideas by providing a method to further refine the proposal distribution by learning a more suitable mean and covariance after obtaining feedback from an original set of samples. To accomplish this change, we generalize the derivation of MPPI to allow for a single joint distribution across controls in the control sequence. With this formulation, we add an adaptive importance sampling (AIS) procedure at each time step. The other approaches discussed can be implemented in addition with our proposed reformulation and benefit from the introduction of an AIS step within the MPPI process. Adapting the proposal distribution across control inputs results in an increase in performance with fewer samples, albeit with a reduction in parallel capability. Our work allows for the MPPI algorithm to be implemented with smaller sample sizes and as an anytime algorithm, increasing performance as computation time allows.

## II. MODEL PREDICTIVE OPTIMIZED PATH INTEGRAL CONTROL

In this section, we generalize MPPI to treat the controls across a time horizon as a single input. We then discuss modifications outlined in other work with MPPI, present the model predictive optimized path integral (MPOPI) control algorithm, and then comment on the computational cost.

### A. Generalized Problem Formulation and Derivation

In MPPI, the control sequence is treated as a set of distinct controls at time  $t$  with the inputs to the system as  $\mathbf{v}_t \sim \mathcal{N}(\mathbf{u}_t, \Sigma)$  where  $\mathbf{u}_t$  is the commanded control input at time  $t$  and  $\Sigma$  is the covariance of the Gaussian noise associated with the input. Therefore, when sampling from the control distribution and estimating the optimal control, we are sampling from separate distributions corresponding to each time step. However, the weights derived from importance sampling are dependent on the cost of the whole control sequence. Using separate distributions at each time step presents a problem if we want to improve our estimation of the expectation by integrating techniques that use the calculated trajectory costs and the control sequence as a single input. Our solution is to treat the sequence of controls as one control input and construct a combined covariance matrix and sample from one, larger joint distribution. This formulation allows different optimization techniques to be applied between time steps to obtain a better estimate of the optimal control through adaptive importance sampling.

Mirroring the derivation of Williams *et al.* [10], we consider the same dynamical system,  $\mathbf{x}_{t+1} = \mathbf{F}(\mathbf{x}_t, \mathbf{v}_t)$ , where  $\mathbf{x}_t \in \mathbb{R}^n$  is the state of the system at time  $t$ ,  $\mathbf{v}_t \in \mathbb{R}^m$  is the input to the system, and  $\mathbf{F}$  is the state-transition function of the system. The input,  $\mathbf{v}_t$ , is a random vector generated by a white-noise process with density function

$\mathbf{v}_t \sim \mathcal{N}(\mathbf{u}_t, \Sigma_t)$  where  $\mathbf{u}_t \in \mathbb{R}^m$  is the commanded control input at time  $t$  and  $\Sigma_t$  is the covariance of the Gaussian noise associated with input  $\mathbf{u}_t$ . We will define the vector  $\mathbf{U}$  to be composed of the components of each control input  $\mathbf{U} = [\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{T-1}]$ ,  $\mathbf{U} \in \mathbb{R}^{mT}$  and define the combined covariance matrix as

$$\Sigma = \begin{bmatrix} \Sigma_0 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \Sigma_1 & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \Sigma_{T-1} \end{bmatrix}, \Sigma \in \mathbb{R}^{mT \times mT}. \quad (1)$$

Therefore, the input,  $\mathbf{V}$ , is a random vector with density function  $\mathbf{V} \sim \mathcal{N}(\mathbf{U}, \Sigma)$  and we can express the analytical form of the density function  $q$  corresponding to the distribution  $\mathbb{Q}_{\mathbf{U}, \Sigma}$  as

$$q(\mathbf{V} | \mathbf{U}, \Sigma) = ((2\pi)^{mT} |\Sigma|)^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{V} - \mathbf{U})^\top \Sigma^{-1}(\mathbf{V} - \mathbf{U})\right). \quad (2)$$

Similar to Williams *et al.* [10], we define our cost function as one that can be decomposed into a state dependent cost  $c(\mathbf{X})$  and a quadratic control cost  $\mathcal{L}(\mathbf{X}, \mathbf{U}) = c(\mathbf{X}) + \frac{\lambda}{2}(\mathbf{U}^\top \Sigma^{-1} \mathbf{U} + \beta_{\tilde{\mathbf{U}}}^\top \mathbf{U}) + c_{\tilde{\mathbf{U}}}$  where  $\mathbf{X} = [\mathbf{x}_0, \mathbf{F}(\mathbf{x}_0, \mathbf{v}_0), \mathbf{F}(\mathbf{F}(\mathbf{x}_0, \mathbf{v}_0), \mathbf{v}_1), \dots]$ ,  $\mathbf{X} \in \mathbb{R}^{nT}$ ,  $\mathbf{v}_t \in \mathbb{R}^m$  are the components of  $\mathbf{V} \in \mathbb{R}^{mT}$ , and  $\beta_{\tilde{\mathbf{U}}}$  and  $c_{\tilde{\mathbf{U}}}$  are constants. We can now write our optimal control problem as

$$\mathbf{U}^* = \arg \min_{\mathbf{U} \in \mathcal{U}} \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} [\phi(\mathbf{X}) + \mathcal{L}(\mathbf{X}, \mathbf{U})] \quad (3)$$

where  $\mathcal{U}$  is the set of all valid control vectors and  $\phi(\mathbf{X})$  is a terminal cost function operating on the terminal state within  $\mathbf{X}$ . Based on our definition of  $\mathbf{X}$ , the state cost of a control sequence is  $S(\mathbf{V}) = c(\mathbf{X}) + \phi(\mathbf{X})$  and the free-energy of the control system is then

$$\mathcal{F}(S, p, \mathbf{X}, \lambda) = -\lambda \log \left( \mathbb{E}_{\mathbb{P}} \left[ \exp \left( -\frac{1}{\lambda} S(\mathbf{V}) \right) \right] \right) \quad (4)$$

$$= -\lambda \log \left( \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} \left[ \exp \left( -\frac{1}{\lambda} S(\mathbf{V}) \right) \frac{p(\mathbf{V})}{q(\mathbf{V} | \mathbf{U}, \Sigma)} \right] \right) \quad (5)$$

where  $\lambda$  is the inverse temperature and  $\mathbb{P}$  is the distribution over a set of control vectors with probability density,  $p$ . From the concavity of the logarithm function, we can apply Jensen's inequality and then use the definition of the KL divergence

$$\mathcal{F}(S, p, \mathbf{X}, \lambda) \leq \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} [S(\mathbf{V})] + \lambda \mathbb{D}_{KL}(\mathbb{Q}_{\mathbf{U}, \Sigma} \| \mathbb{P}). \quad (6)$$

This inequality relates the free energy to the state cost of the control problem and the divergence between distribution  $\mathbb{Q}_{\mathbf{U}, \Sigma}$  and  $\mathbb{P}$ .

Suppose the base distribution is  $p(\mathbf{V}) = q(\mathbf{V} | \tilde{\mathbf{U}}, \Sigma)$  where  $\tilde{\mathbf{U}}$  is a nominal input applied to the system. Defining,  $\beta_{\tilde{\mathbf{U}}}^\top = -2\tilde{\mathbf{U}}^\top \Sigma^{-1}$  and  $c_{\tilde{\mathbf{U}}} = \tilde{\mathbf{U}}^\top \Sigma^{-1} \tilde{\mathbf{U}}$ , we have

$$\mathbb{D}_{KL}(\mathbb{Q}_{\mathbf{U}, \Sigma} \| \mathbb{Q}_{\tilde{\mathbf{U}}, \Sigma}) = \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} \left[ \log \left( \frac{q(\mathbf{U} | \Sigma)}{q(\tilde{\mathbf{U}} | \Sigma)} \right) \right] \quad (7)$$

$$= \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} \left[ \frac{1}{2} \left( (\mathbf{U} - \tilde{\mathbf{U}})^\top \Sigma^{-1} (\mathbf{U} - \tilde{\mathbf{U}}) \right) \right] \quad (8)$$

$$= \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} \left[ \frac{1}{2} \left( \mathbf{U}^\top \Sigma^{-1} \mathbf{U} + \beta_{\tilde{\mathbf{U}}}^\top \mathbf{U} + c_{\tilde{\mathbf{U}}} \right) \right]. \quad (9)$$

Using eq. (9) and the state cost, we can rewrite eq. (6) as

$$\begin{aligned} \mathcal{F}(S, p, \mathbf{X}, \lambda) &\leq \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} \left[ \phi(\mathbf{X}) + c(\mathbf{X}) \right. \\ &\quad \left. + \frac{\lambda}{2} \left( \mathbf{U}^\top \Sigma^{-1} \mathbf{U} + \beta_{\tilde{\mathbf{U}}}^\top \mathbf{U} + c_{\tilde{\mathbf{U}}} \right) \right] \quad (10) \\ &\leq \mathbb{E}_{\mathbb{Q}_{\mathbf{U}, \Sigma}} [\phi(\mathbf{X}) + \mathcal{L}(\mathbf{X}, \mathbf{U})] \quad (11) \end{aligned}$$

and we have established that the free energy provides a lower bound on our reformulated optimal control problem. Furthermore, this inequality becomes an equality if we consider the optimal control distribution  $\mathbb{Q}^*$  defined by its density function  $q^*(\mathbf{V}) = \frac{1}{\eta} \exp(-\frac{1}{\lambda} S(\mathbf{V})) p(\mathbf{V})$ . We can now minimize the divergence between our control distribution and the optimal distribution and achieve an optimal control trajectory instead of minimizing eq. (3) directly.

The processes of minimizing the KL-divergence between our control distribution and the optimal distribution also follows closely to the original derivation [10] and results in the quadratic minimization problem

$$\mathbf{U}^* = \arg \min_{\mathbf{U} \in \mathcal{U}} (\mathbb{E}_{\mathbb{Q}^*} [(\mathbf{V} - \mathbf{U})^\top \Sigma^{-1} (\mathbf{V} - \mathbf{U})]) \quad (12)$$

with the optimal solution in the unconstrained case of

$$\mathbf{U}^* = \mathbb{E}_{\mathbb{Q}^*} [\mathbf{V}]. \quad (13)$$

Previous work has addressed incorporating control constraints as part of the dynamical model, therefore we only consider the unconstrained solution and discuss modifications in section II-B.

Using importance sampling to determine the expectation in eq. (13) also follows the original derivation closely. Let  $\mathbb{Q}_{\hat{\mathbf{U}}, \Sigma}$  represent a tentative sampling distribution to perform importance sampling from, then we get our control through

$$\mathbf{U}^* = \mathbb{E}_{\mathbb{Q}_{\hat{\mathbf{U}}, \Sigma}} [w(\mathbf{V}) \mathbf{V}] \quad (14)$$

$$w(\mathbf{V}) = \frac{1}{\eta} \exp \left( -\frac{1}{\lambda} \left( S(\mathbf{V}) + \lambda (\hat{\mathbf{U}} - \tilde{\mathbf{U}})^\top \Sigma^{-1} \mathbf{V} \right) \right) \quad (15)$$

$$\eta = \int \exp \left( -\frac{1}{\lambda} \left( S(\mathbf{V}) + \lambda (\hat{\mathbf{U}} - \tilde{\mathbf{U}})^\top \Sigma^{-1} \mathbf{V} \right) \right) d\mathbf{V}. \quad (16)$$

No assumptions were changed and all of the theoretical results from MPPI still hold such as eq. (14) being a global optimal solution (from the information-theoretic perspective) under the assumption we can evaluate the expectation perfectly. With no other modification, this approach is equivalent to traditional MPPI. When constructing  $\Sigma$  in eq. (1) we only allowed for control input noise in order to mirror MPPI. In the derivation, we made no assumptions on the structure of  $\Sigma$  except that it is a valid covariance matrix and  $\Sigma \in \mathbb{R}^{mT \times mT}$ .

### B. Other Modifications

Information-theoretic model predictive control does present some practical issues such as numerical stability with the trajectory costs. Previous work established methods to address such problems and we extend those concepts to our algorithm. Notice in eq. (15) that  $\lambda$  directly affects the control cost. That is, as  $\lambda$  increases and distributes the weight among trajectories, it also increases the contribution

of the control cost. We can decouple the influence of  $\lambda$  by redefining the base distribution as  $\tilde{p}(\mathbf{V}) = p(\mathbf{V} | \alpha \hat{\mathbf{U}}, \Sigma)$  where  $0 \leq \alpha \leq 1$ . The new base distribution changes our weight to

$$w(\mathbf{V}) = \frac{1}{\eta} \exp \left( -\frac{1}{\lambda} \left( S(\mathbf{V}) + \gamma (\hat{\mathbf{U}} - \tilde{\mathbf{U}})^\top \Sigma^{-1} \mathbf{V} \right) \right) \quad (17)$$

where  $\gamma = \lambda(1 - \alpha)$ . With this change, as  $\alpha$  moves from 0 to 1, the control cost is less of a factor in the overall weight. Additionally, calculating the weight in eq. (17) can result in numerical instability. To avoid those issues a normalizing factor is introduced to set the lowest cost to zero.

The control sequence produced by the sampling methodology can produce inputs that change rapidly across time steps. There are a few ways to address this case. One option is to incorporate a smoothing process into the dynamics model. Since there are no restrictions on the model dynamics in the formulation, we can incorporate different input modifications as part of the dynamics. This idea follows the same concept on how to integrate control constraints allowing us to only consider the unconstrained solution in eq. (13). We can modify our inputs by introducing a function  $g$  that applies the appropriate smoothing and control constraints to our inputs before applying the dynamics and our system becomes  $\mathbf{x}_{t+1} = \mathbf{F}(\mathbf{x}_t, g(\mathbf{v}_t))$ .

### C. Adaptive Importance Sampling

Using importance sampling to estimate eq. (14) is a key aspect of MPPI. However, since we are limited to a finite number of samples, the quality of the estimate is influenced by the reference distribution [16]. One approach to improve our estimate is to use AIS to adapt our proposal distribution iteratively. There are many AIS algorithms with different performance costs and benefits as surveyed by Bugallo *et al.* [23]. Control problems present various challenges and the choice of which AIS algorithm to implement will vary across scenarios. With the generalized formulation, we are not restricted to the choice of an AIS algorithm as  $\mathbf{U}$  and  $\Sigma$  can be updated to achieve a better proposal distribution including correlation terms across control steps. Separating the AIS step from the MPPI algorithm also allows different parameter settings for the control problem and the AIS algorithm (e.g. different  $\lambda$  values to allow exploration during AIS but more selective when calculating the final control). This work focuses on the integration of different AIS algorithms into MPPI and leaves the investigation of specific trade-offs of different algorithms for future research.

### D. Algorithm

Algorithm 1 outlines the MPOPI algorithm. This algorithm is executed at each time step until the task is complete. Bold capital letters (e.g.  $\mathbf{U}, \Sigma, \mathbf{X}$ ) represent the composite vectors/matrices as described in section II-A and bold lowercase letters with subscripts (e.g.  $\mathbf{u}_t, \mathbf{x}_t$ ) represent the components of the composite vectors. The call to *PerformAIS* represents a call to an AIS algorithm to execute at each iteration. One of the key differences from MPPI to note is the calculation

of the control cost component on line 9 and the weighted average to update  $\mathbf{U}$  (line 16). We have to account for the appropriate change in control amount as we change our proposal distribution. If we set  $L = 1$ , then  $\mathbf{U}' - \mathbf{U}$  vanishes and MPOPI is equivalent to MPPI.

---

**Algorithm 1:** Model Predictive Optimized Path Integral Control

---

```

Given:  $\mathbf{F}, g, K, T$ 
 $\mathbf{U} \in \mathbb{R}^{mT}$ 
 $\Sigma \in \mathbb{R}^{mT \times mT}$ 
 $\phi, c, \lambda, \alpha$  /* Cost function/parameters */
 $L$  /* Max number of AIS iterations */
 $\psi_1, \psi_2, \dots, \psi_p$  /* AIS parameters */
1  $\mathbf{x}_0 \leftarrow \text{GetStateEstimate}()$ 
2  $\mathbf{U}' \leftarrow \mathbf{U}$ 
3  $\Sigma' \leftarrow \Sigma$ 
4 for  $\ell \leftarrow 1$  to  $L$  do
5   for  $k \leftarrow 1$  to  $K$  do
6     Sample  $\mathcal{E}_k$  from  $\mathcal{N}(0, \Sigma')$ 
7     for  $t \leftarrow 1$  to  $T$  do
8        $\mathbf{x}_t \leftarrow \mathbf{F}(\mathbf{x}_{t-1}, g(\mathbf{u}'_{t-1} + \boldsymbol{\varepsilon}_{t-1}^k))$ 
9        $s_k \leftarrow c(\mathbf{X}) + \phi(\mathbf{X}) + \lambda(1 - \alpha)\mathbf{U}'^T \Sigma^{-1} (\mathcal{E}_k + \mathbf{U}' - \mathbf{U})$ 
10    if  $\ell < L$  then
11       $\mathbf{U}', \Sigma' \leftarrow \text{PerformAIS}(\mathbf{U}', \Sigma', \mathbf{S}, \psi_1, \psi_2, \dots, \psi_p, \ell)$ 
12     $\rho \leftarrow \min(\mathbf{S})$ 
13     $\eta \leftarrow \sum_{k=1}^K \exp(-\frac{1}{\lambda}(s_k - \rho))$ 
14    for  $k \leftarrow 1$  to  $K$  do
15       $w_k \leftarrow \frac{1}{\eta} \exp(-\frac{1}{\lambda}(s_k - \rho))$ 
16       $\mathbf{U} += w_k (\mathcal{E}_k + \mathbf{U}' - \mathbf{U})$ 
17    SendToActuators( $\mathbf{u}_0$ )
18    for  $t \leftarrow 1$  to  $T - 1$  do
19       $\mathbf{u}_{t-1} \leftarrow \mathbf{u}_t$ 
20     $\mathbf{u}_{T-1} \leftarrow \text{Initialize}(\mathbf{u}_{T-1})$ 

```

---

Figure 1 provides a graphical depiction of how MPOPI can improve samples to better estimate the optimal control. The trajectories were generated from throttle and steering commands for controlling a car. Figure 1a depicts the trajectories of the controls sampled from the initial proposal distribution. After each iteration, the proposal distribution is updated and new samples are generated. Figure 1b and fig. 1c show how the sampled controls evolve, resulting in higher scoring trajectories after each iteration of the AIS algorithm.

### E. Computational Cost

A key enabler for MPPI to work real time is its ability to be implemented in parallel. The computational complexity of MPPI can be summarized as the complexity of obtaining  $KT$  samples and the complexity of propagating the system dynamics  $T$  times for  $K$  samples in parallel. The addition of AIS changes this complexity. The sampling and propagation of the system dynamics must occur  $L$  times, along with the addition of the complexity of the AIS algorithm  $L - 1$  times. Modern sampling and AIS algorithms are efficient. The largest computational burden comes from propagating the dynamics in  $L$  sequential steps. Using AIS will allow the use of fewer samples and still take advantage of parallel operations during the propagation step, but at the burden of reestablishing the parallel computations  $L$  times.

The MPOPI algorithm is more sample efficient and the iterative nature lends itself to an anytime approach where

$L$  can increase with computation time available. However, these benefits come at the cost of a reduction in parallel capability. For example, given  $M = KL$  samples, MPPI would propagate the system dynamics in parallel for all  $M$  samples. However, MPOPI requires propagating  $K$  samples in parallel in  $L$  sequential steps where the execution of the AIS algorithm must occur after each propagation of  $K$  samples. The complexity of establishing parallel computations is hardware, problem, and implementation specific. As problem complexity increases, MPOPI provides an alternative approach to adding more samples.

## III. EXPERIMENTS

We tested the MPOPI algorithm in different virtual environments and compared it to MPPI with the same modifications discussed in section II-B. MPPI has been shown across the literature to perform equitably or outperform numerous other control algorithms and was chosen as the baseline [13], [18], [21]. We integrated five different AIS algorithms with MPOPI. The algorithms have the ability to use multiple proposal distributions, but we chose to use one proposal distribution to keep the comparison closer to MPPI. We implemented a Population Monte Carlo (PMC) algorithm [24], a mean only moment matching AIS algorithm ( $\mu$ -AIS), a mean and covariance moment matching AIS algorithm ( $\mu\Sigma$ -AIS) similar to Mixture-PMC [25], a Cross-Entropy method (CE) [26], and a Covariance Matrix Adaptation Evolutionary Strategy (CMA) [27]. The  $\mu$ -AIS algorithm is similar to the iterative approach introduced by Pravitra *et al.* [21], but we allowed for different values of  $\lambda$  for updating the mean and the final weight calculation of the samples. This separation resulted in better performance versus using a shared  $\lambda$ . We kept parameters constant across similar algorithms to show a more direct comparison. Reference the repository for parameters used and details of the implementation.

We first looked at the continuous version of the MountainCar environment [28] and then performed simulations with a non-linear car model. To increase the action space size and complexity, we modified the car environment to control multiple cars cooperatively and conducted experiments on MuJoCo gym environments. In this section, the term effective samples refers to the total samples an algorithm had available. For MPPI, it is the number of samples and for MPOPI, it is the number of samples multiplied by the number of AIS iterations,  $KL$ . We also discuss rewards when describing objective functions instead of costs. The effective sample sizes were increased by changing the number of samples or AIS iterations with MPOPI or the total number of samples for MPPI. The smallest sample size in the MountainCar and car racing scenarios restricted MPOPI to one AIS iteration to show the equivalence to MPPI and provide a more direct comparison. All simulations used the same parameters for MPPI and MPOPI (e.g.  $\lambda, \alpha, T$ ).

### A. MountainCar Environment

The goal of the MountainCar problem is to get an underpowered car up a hill to a goal location with the dynamics




Fig. 1: Evolution of 150 trajectories from sampled controls using MPOPI and the cross entropy method for the AIS algorithm. Green depicts high valued trajectories and red represents low valued trajectories within each sub-figure. The control is a series of throttle and steering commands. Details of this environment are described in section III-B.

and parameters as specified by Kochenderfer *et al.* [29]. The action space consisted of a single continuous action at each time step. The reward function for this problem was modified to be  $R(x_t, v_t) = -1 + |v_t| + 100000R^*$  where  $R^*$  is an indicator variable for reaching the goal location with  $v_t > 0$ . The environment terminated when the car reached the goal location or at 200 steps. The simulations consisted of 1000 trials for effective samples ranging from 20 to 180 with MPOPI using 20 samples and varying the number of AIS iterations. The results are shown in fig. 2. The algorithms converge to a similar number of steps while also reaching a point where increasing the number of samples provided little return. MPPI was able to reach similar performance but failed to outperform versions of MPOPI except for PMC. The CE and CMA versions of MPOPI were able to reach the plateau with as few as 40 and 60 effective samples respectively. The PMC algorithm uses multinomial resampling and struggles with a small initial sample size.




Fig. 2: MountainCar results. Shaded regions are 95% confidence intervals. Adjusting the proposal distribution allows MPOPI to achieve MPPI level performance in as few as 2 iterations and 40 effective samples.

### B. Car Racing

This scenario simulated a form of car racing and was constructed to be similar to the virtual racing environment

used by Williams [17]. The total length of the track was approximately 1.18 km with multiple curves. We implemented a nonlinear single-track model that captures transient dynamics in coming with a variant of a Fiala brush tire model [30]. The model was a representation of the electric X1, an autonomous experimental car. For details of the model dynamics and parameters of the X1 see Subosits and Gerdes [30] and Brown and Gerdes [31]. The action space for a single car consisted of two continuous inputs per time step: a steering command and a throttle/brake command. The controls were issued at 10 Hz with the dynamics operating at 100 Hz. The reward function for each car was  $R(\mathbf{x}_t) = 2|v_t| - |d| - 5000R_\beta^* - 1000000R_t^*$  where  $v_t$  is the velocity of the car at time step  $t$ ,  $d$  is the distance of the car from the middle of the track lane,  $R_\beta^*$  is an indicator variable for exceeding a certain drift amount  $\beta$ , and  $R_t^*$  is an indicator variable for when the car leaves the track boundary ( $\beta$  is the angle between the velocity vector and the longitudinal axis; the limit was set to  $45^\circ$ ). All runs consisted of 25 trials and each trial concluded when the car completed 2 laps, exceeded the  $\beta$  limit for more than 5 sec, or was outside the track boundary for greater than 1 sec. The number of effective samples varied from 375 to 2250 with MPOPI using iterations of 150 or 375 samples depending on the number of cars and the AIS algorithm. More details of this environment and videos can be seen in the repository.

Results of controlling a single car are shown in fig. 3a. There were no violations of the track boundary or the  $\beta$  limit. All versions of MPOPI except for  $\mu$ -AIS outperformed MPPI and achieved MPPI peak performance with as few as 375 samples and 3 iterations of the AIS algorithm. There was more of a disparity in performance with fewer effective samples. The single car scenario resulted in the car attaining speeds of 39 m/s with  $\beta$  values exceeding  $34^\circ$ .

We modified the car racing environment to cooperatively control multiple cars. For  $N_c$  cars, the reward function was

$$R(\mathbf{x}_t) = \sum_{i=1}^{N_c} R_i - 11000R_c^* - \sum_{i=1}^{N_c-1} \sum_{j=i+1}^{N_c} |d_{i,j}| \quad (18)$$

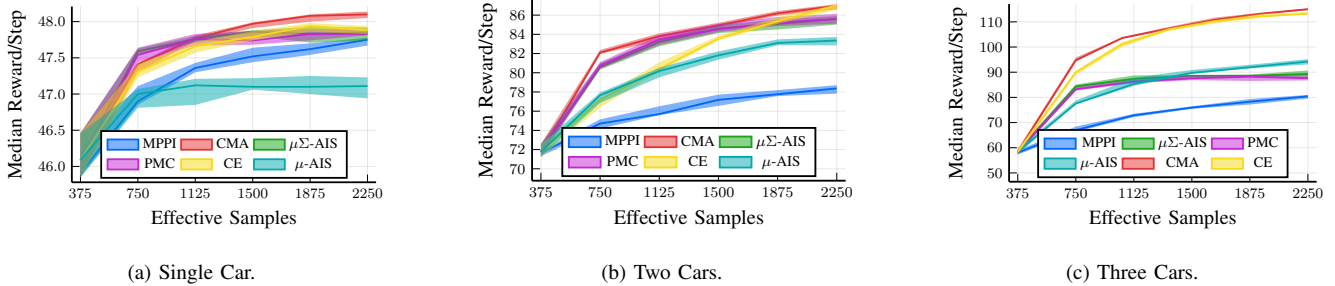


Fig. 3: Car racing results. Shaded regions are 95% confidence intervals. The sample efficiency of MPOPI versus MPPI becomes more pronounced as the problem complexity increases.

TABLE I: 4-Car completion rate with zero violations.

Effective Samples	MPPI	$\mu$ -AIS	$\mu\Sigma$ -AIS	PMC	CE	CMA
750	0.67	0.96	0.88	0.80	1.00	0.96
1125	0.84	1.00	0.88	0.84	1.00	1.00
1500	0.88	1.00	0.96	0.96	1.00	1.00
1875	0.88	1.00	1.00	1.00	1.00	1.00
2250	0.96	1.00	1.00	1.00	1.00	1.00

where  $R_i$  is the individual reward for car  $i$ ,  $R_c^*$  is an indicator for the cars being within 4 m of each other, and  $d_{i,j}$  is the distance between car  $i$  and car  $j$ . The last term was added to increase complexity by emphasizing coordinated actions across the action space.

Results for two and three cars are shown in fig. 3b and fig. 3c. There were no violations for the two and three car scenarios. The success rate of a trial in the 4-car scenario with no violations is shown in table I. For a 5-car scenario, MPPI was only able to complete 3 out of 25 trials with 6000 samples. The CE and CMA versions of MPOPI were able to control up to six cars successfully with as few as 1500 effective samples.

For a given effective sample size, the CE and CMA methods benefited from using fewer samples with more iterations as the number of cars increased. However, as the sample size decreased, the method to approximate the covariance matrix in the CE version of MPOPI became more important. Algorithms specializing in estimating large covariance matrices with few samples resulted in higher performance. Details of the different covariance matrix estimation algorithms tested are provided in the repository. Changing the sample size and iteration count on the single and two car scenarios did not have a noteworthy change in performance.

### C. MuJoCo Gym Environments

To further test MPOPI, we conducted experiments on MuJoCo gym environments [32], [33] using EnvPool [34]. No environmental parameters were changed from the documented implementations. We ran each scenario 10 times and the experiments were conducted for 250 steps. We conducted experiments using only the CE version of MPOPI based on previous results. The average total accumulated reward is shown for two tasks in table II. Algorithm parameters were

not tuned for performance on the different scenarios and were kept consistent across MPPI and MPOPI. Reference the repository wiki page for more details on the parameters of the algorithms used.

HalfCheetah-v4 and Ant-v4 have six and eight dimensional action spaces respectively. These environments were chosen to test MPOPI on different scenarios with complex control spaces. Similar to previous results, MPOPI achieved similar performance as MPPI with fewer effective samples. The MuJoCo gym and multi-car results demonstrate the significance the proposal distribution can have on the quality of the importance sampling estimate of the optimal control. MPOPI increases in performance with each iteration of the AIS algorithm and also increases in performance with a larger initial sample size. These results suggest MPOPI can be used in complex problems as an anytime algorithm where the sample size and iteration count can be adjusted as computation time allows.

TABLE II: MuJoCo environment results. Average total reward for 250 steps on the HalfCheetah-v4 and Ant-v4 environments. The margin shown is the 95% confidence interval.

Effective Samples	HalfCheetah-v4		Ant-v4	
	MPPI	MPOPI-CE	MPPI	MPOPI-CE
250	1554 ± 224	2154 ± 223	1444 ± 92	1711 ± 161
500	1973 ± 162	2771 ± 183	1635 ± 147	1839 ± 93
1000	2071 ± 103	3358 ± 328	1814 ± 100	2197 ± 89
1500	2192 ± 197	3798 ± 140	1946 ± 131	2251 ± 77
3000	2440 ± 108	4737 ± 315	2180 ± 95	2417 ± 114

## IV. CONCLUSION

We increased the sample efficiency of MPPI by implementing adaptive importance sampling into the original importance sampling step of MPPI. The benefit of learning a better proposal distribution from previous samples was demonstrated in simulated environments and became further evident as the dimension of the action space increased. Simulation results suggest MPOPI can be used as an anytime algorithm, increasing the value of control at each iteration versus relying on a single large set of samples. The modification to adjust the proposal distribution to increase the efficiency and quality of the samples used is a promising change to an already high performing control algorithm.

## REFERENCES

- [1] J. Ji, A. Khajepour, W. W. Melek, and Y. Huang, "Path planning and tracking for vehicle collision avoidance based on model predictive control with multiconstraints," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 2, pp. 952–964, 2017.
- [2] Z. Samadikhoshkho, S. Ghorbani, F. Janabi-Sharifi, and K. Zareinia, "Nonlinear control of aerial manipulation systems," *Aerospace Science and Technology*, vol. 104, p. 105945, 2020.
- [3] P. Foehn, D. Brescianini, E. Kaufmann, T. Cieslewski, M. Gehrig, M. Muglikar, and D. Scaramuzza, "Alphapilot: Autonomous drone racing," *Autonomous Robots*, pp. 1–14, 2021.
- [4] M. Schwenzer, M. Ay, T. Bergs, and D. Abel, "Review on model predictive control: An engineering perspective," *International Journal of Advanced Manufacturing Technology*, vol. 117, no. 5, pp. 1327–1349, 2021.
- [5] D. Kouzoupis, G. Frison, A. Zanelli, and M. Diehl, "Recent advances in quadratic programming algorithms for nonlinear model predictive control," *Vietnam Journal of Mathematics*, vol. 46, no. 4, pp. 863–882, 2018.
- [6] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, "A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition," *Journal of Nonlinear Science*, vol. 25, no. 6, pp. 1307–1346, 2015.
- [7] M. Korda and I. Mezić, "Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control," *Automatica*, vol. 93, pp. 149–160, 2018.
- [8] C. A. Johnson and E. Yeung, "A class of logistic functions for approximating state-inclusive Koopman operators," in *American Control Conference (ACC)*, 2018.
- [9] E. A. Theodorou, J. Buchli, and S. Schaal, "A generalized path integral control approach to reinforcement learning," *Journal of Machine Learning Research*, vol. 11, no. 104, pp. 3137–3181, 2010.
- [10] 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)*, 2017.
- [11] E. A. Theodorou and E. Todorov, "Relative entropy and free energy dualities: Connections to path integral and KL control," in *IEEE Conference on Decision and Control (CDC)*, 2012.
- [12] G. I. Boutselis, Z. Wang, and E. A. Theodorou, "Constrained sampling-based trajectory optimization using stochastic approximation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [13] 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.
- [14] 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.
- [15] M. S. Gandhi, B. Vlahov, J. Gibson, G. Williams, and E. A. Theodorou, "Robust model predictive path integral control: Analysis and performance guarantees," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1423–1430, 2021.
- [16] P.-T. D. Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein, "A tutorial on the cross-entropy method," *Annals of Operations Research*, vol. 134, pp. 19–67, 1 2005.
- [17] G. Williams, "Model predictive path integral control: Theoretical foundations and applications to autonomous driving," Ph.D. dissertation, Georgia Institute of Technology, 2019.
- [18] R. Kusumoto, L. Palmieri, M. Spies, A. Csiszar, and K. O. Arras, "Informed information theoretic model predictive control," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [19] M. Okada and T. Taniguchi, "Acceleration of gradient-based path integral method for efficient optimal and inverse optimal control," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [20] A. Lambert, F. Ramos, B. Boots, D. Fox, and A. Fishman, "Stein variational model predictive control," in *Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, 2021.
- [21] J. Pravitra, E. A. Theodorou, and E. N. Johnson, "Flying complex maneuvers with model predictive path integral control," in *AIAA SciTech Forum*, 2021.
- [22] Z. Wang, O. So, J. Gibson, B. Vlahov, M. S. Gandhi, G.-H. Liu, and E. A. Theodorou, "Variational inference MPC using Tsallis divergence," in *Robotics: Science and Systems*, 2021.
- [23] M. F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Míguez, and P. M. Djuric, "Adaptive importance sampling: The past, the present, and the future," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 60–79, 2017.
- [24] O. Cappé, A. Guillin, J.-M. Marin, and C. P. Robert, "Population Monte Carlo," *Journal of Computational and Graphical Statistics*, vol. 13, no. 4, pp. 907–929, 2004.
- [25] O. Cappé, R. Douc, A. Guillin, J.-M. Marin, and C. P. Robert, "Adaptive importance sampling in general mixture classes," *Stat and Comput*, vol. 18, no. 4, pp. 447–459, 2008.
- [26] R. Y. Rubinstein and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Springer-Verlag, 2004, vol. 133.
- [27] Y. El-Laham, V. Elvira, and M. F. Bugallo, "Robust covariance adaptation in adaptive importance sampling," *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 1049–1053, 2018.
- [28] A. W. Moore, "Efficient memory-based learning for robot control," Ph.D. dissertation, University of Cambridge, 1990.
- [29] M. J. Kochenderfer, T. A. Wheeler, and K. H. Wray, *Algorithms for Decision Making*. MIT Press, 2022.
- [30] J. K. Subosits and J. C. Gerdes, "Impacts of model fidelity on trajectory optimization for autonomous vehicles in extreme maneuvers," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 3, pp. 546–558, 2021.
- [31] M. Brown and J. C. Gerdes, "Coordinating tire forces to avoid obstacles using nonlinear model predictive control," *IEEE Transactions on Intelligent Transportation Systems*, vol. 5, no. 1, pp. 21–31, 2020.
- [32] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.
- [33] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, *OpenAI gym*, 2016. eprint: arXiv:1606.01540.
- [34] J. Weng, M. Lin, S. Huang, B. Liu, D. Makoviichuk, V. Makoviychuk, Z. Liu, Y. Song, T. Luo, Y. Jiang, Z. Xu, and S. Yan, "EnvPool: A highly parallel reinforcement learning environment execution engine," *arXiv preprint arXiv:2206.10558*, 2022.