

Sample-Efficient and Smooth Cross-Entropy Method Model Predictive Control Using Deterministic Samples

Markus Walker, Daniel Frisch, and Uwe D. Hanebeck

Abstract—Cross-entropy method model predictive control (CEM-MPC) is a powerful gradient-free technique for nonlinear optimal control, but its performance is often limited by the reliance on random sampling. This conventional approach can lead to inefficient exploration of the solution space and non-smooth control inputs, requiring a large number of samples to achieve satisfactory results. To address these limitations, we propose deterministic sampling CEM (dsCEM), a novel framework that replaces the random sampling step with deterministic samples derived from localized cumulative distributions (LCDs). Our approach introduces modular schemes to generate and adapt these sample sets, incorporating temporal correlations to ensure smooth control trajectories. This method can be used as a drop-in replacement for the sampling step in existing CEM-based controllers. Experimental evaluations on two nonlinear control tasks demonstrate that dsCEM consistently outperforms state-of-the-art iCEM in terms of cumulative cost and control input smoothness, particularly in the critical low-sample regime.

Index Terms—Model predictive control, cross-entropy method, deterministic sampling, localized cumulative distribution, sampling efficiency.

I. INTRODUCTION

Optimal control plays a crucial role in achieving desired performance in many applications, such as autonomous driving and robotics. One widely used method for solving these types of problems is model predictive control (MPC), which operates in a receding horizon fashion. At each time step, MPC solves an optimization problem to determine the optimal control inputs over a finite horizon. It then applies only the first input to the system [1]. However, this optimization can be challenging for nonlinear systems, non-differentiable dynamics, and non-convex cost functions. In such cases, gradient-free optimization methods are often preferred because they can navigate complex cost landscapes and avoid local minima [2].

The cross-entropy method (CEM) [3] is a popular gradient-free technique that has been successfully applied to various optimal control problems [2], [4]. It iteratively refines a proposal distribution over control inputs by sampling, evaluating performance, and updating the distribution based on the best-performing samples. CEM-MPC integrates this into the MPC framework, using a discrete-time stochastic process (e.g., a Gaussian process, see Fig. 1) to generate control input sequences. The parameters of this process are updated based

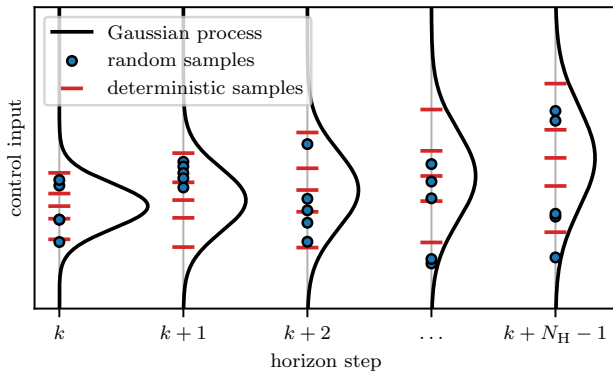


Fig. 1: Schematic showing control input sampling over a finite horizon using either deterministic or random samples. As can be seen, deterministic samples cover the stochastic process without large gaps or clusters. For simplicity, time correlations are neglected.

on the cost of each sampled trajectory, which focuses the search on promising regions of the control space.

A key challenge in standard CEM-MPC is that control inputs are sampled independently at each time step, which can result in non-smooth control sequences and requires a large number of samples to adequately explore the control space. The improved CEM (iCEM) [2] introduces smoothness by temporal correlations, allowing for more informed sampling strategies that lead to better performance. However, reliance on random sampling, even with temporal correlations, can still produce non-smooth control trajectories. This is often undesirable in practice because it can damage actuators and cause erratic system behavior.

To overcome these challenges, we propose replacing random sampling with a deterministic sampling strategy. Our approach, which we term *deterministic sampling Cross-Entropy Method* (dsCEM), leverages precomputed optimal sample sets based on localized cumulative distributions (LCDs) [5], [6]. These deterministic samples are structured to cover the solution space more efficiently and with fewer discrepancies than random samples (see Fig. 1). Inspired by iCEM, we incorporate temporal correlations into the sampling process with the goal of generating control sequences that are both more effective and significantly smoother. The proposed sampling schemes are designed as modular drop-in replacements for the sampling step in CEM-based controllers, making this approach broadly applicable.

This work is part of the German Research Foundation (DFG) AI Research Unit 5339 regarding the combination of physics-based simulation with AI-based methodologies for the fast maturation of manufacturing processes.

Markus Walker, Daniel Frisch and Uwe D. Hanebeck are with the Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Germany (e-mail: {markus.walker, daniel.frisch, uwe.hanebeck}@kit.edu).

A. Contribution And Outline

We first state the optimal control problem and summarize the CEM–MPC approach in Secs. II and III. Then, in Sec. IV, we introduce the concept of deterministic samples based on LCDs and describe how to integrate them into the CEM–MPC framework. Finally, we evaluate the proposed approach on two benchmark tasks, (i) the mountain car and (ii) the cart-pole swing-up task, in Sec. V, and compare it against the widely used iCEM method [2].

B. Notation

In this paper, underlined letters, e.g., \underline{x} , denote vectors, boldface letters, such as $\underline{\mathbf{x}}$, represent random variables, while boldface capital letters, such as \mathbf{A} , indicate matrices. Sets are denoted by calligraphic letters, e.g., \mathcal{E} . The mean of a random variable is denoted by $\hat{\cdot}$, e.g., $\hat{\underline{x}}$, and covariance matrices are denoted by \mathbf{C} . Diagonal matrices are represented by $\text{diag}(\cdot)$. The indicator function is denoted by $\mathbb{1}$, e.g., $\mathbb{1}_A$ is 1 if event A is true and 0 otherwise.

II. OPTIMAL CONTROL PROBLEM

We consider discrete-time, finite-horizon deterministic optimal control problems (OCPs) with cumulative cost

$$J_k = g_{N_H}(\underline{x}_{k+N_H}) + \sum_{n=k}^{k+N_H-1} g_n(\underline{x}_n, \underline{u}_n), \quad (1)$$

where k is the current time step, N_H the prediction horizon, $\underline{x}_k \in \mathbb{R}^{d_x}$ the system state, $\underline{u}_n \in \mathbb{R}^{d_u}$ the control input, $g_n(\underline{x}_n, \underline{u}_n): \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}$ the stage cost at time step n , and $g_{N_H}(\underline{x}_{k+N_H}): \mathbb{R}^{d_x} \rightarrow \mathbb{R}$ the terminal cost. The discrete-time system dynamics are given by

$$\underline{x}_{n+1} = \underline{a}_n(\underline{x}_n, \underline{u}_n) \quad \text{for } n = k, \dots, k + N_H - 1, \quad (2)$$

with initial state \underline{x}_k at time step k and system function $\underline{a}_n(\underline{x}_n, \underline{u}_n): \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d_x}$. The OCP is to find the optimal control input sequence $\underline{u}_k^*, \dots, \underline{u}_{k+N_H-1}^*$ that minimizes the cumulative cost J_k while satisfying system dynamics and additional constraints on states and controls. The OCP can be solved in a receding horizon fashion using MPC, where at each time step k , the first control input \underline{u}_k^* of the optimal sequence is applied to the system, and the process is repeated at the next time step.

Typically, the OCP is solved using gradient-based approaches such as direct single- or multiple-shooting methods [7], differential dynamic programming [8], iterative linear quadratic regulator [9], or gradient-free methods such as model predictive path integrals [10], or the CEM [2], [4]. In this work, we focus on CEM–MPC, which is described in more detail below.

III. CROSS-ENTROPY METHOD MPC

Before introducing the CEM–MPC, we will briefly explain the concept of the CEM as an optimizer, based on [3], [11]. Note that, to prevent confusion with the control input \underline{u} and state \underline{x} , we denote the d_ξ -dimensional optimization variable for the CEM as $\xi \in \mathbb{R}^{d_\xi}$.

A. Cross-Entropy Method

The CEM aims to find a density $\xi \sim f(\xi; \theta)$, parameterized by θ , whose samples $\xi^{(i)}$ tend to achieve small costs $J(\xi^{(i)})$. By setting a performance threshold γ and defining the event $A_\gamma = \{\xi \mid J(\xi) \leq \gamma\}$, i.e., the event of achieving low-cost solutions, the corresponding rare-event probability can be expressed as [3]

$$l = \mathbb{P}_{f(\xi; \theta)}(J(\xi) \leq \gamma) = \mathbb{E}_{f(\xi; \theta)}\{\mathbb{1}_{J(\xi) \leq \gamma}\}, \quad (3)$$

where $\mathbb{1}_{J(\xi) \leq \gamma}$ is the indicator function. By multiplying $\tilde{f}(\xi)/\bar{f}(\xi)$ within the expectation (3), rearranging yields

$$l = \mathbb{E}_{f(\xi; \theta)}\left\{\mathbb{1}_{J(\xi) \leq \gamma} \frac{\tilde{f}(\xi)}{\bar{f}(\xi)}\right\} = \mathbb{E}_{\tilde{f}(\xi)}\left\{\mathbb{1}_{J(\xi) \leq \gamma} \frac{f(\xi; \theta)}{\tilde{f}(\xi)}\right\},$$

where $\tilde{f}(\xi)$ is an importance sampling probability density function (PDF). Note that the expectation is now w.r.t. the importance PDF $\tilde{f}(\xi)$. According to importance sampling theory, the optimal choice of $\tilde{f}^*(\xi)$ that minimizes the variance of the estimator is [3]

$$\tilde{f}^*(\xi) = \frac{\mathbb{1}_{J(\xi) \leq \gamma} f(\xi; \theta)}{l}.$$

However, this is infeasible because it depends on l , which is the quantity we want to estimate using $\tilde{f}(\xi)$, and is therefore a-priori unknown.

To overcome this infeasibility, the central idea of CEM is to *iteratively approximate* the optimal PDF $\tilde{f}^*(\xi)$. At each iteration j , we start with a known proposal distribution $\tilde{f}(\xi; \theta_j)$ and aim to find a better set of parameters θ_{j+1} for the next iteration. This is achieved by finding the parameters θ_{j+1} that minimize the Kullback–Leibler (KL) divergence between the optimal PDF derived from the current step,

$$\tilde{f}_j^*(\xi) = \frac{\mathbb{1}_{J(\xi) \leq \gamma_j} \tilde{f}(\xi; \theta_j)}{l_j},$$

and the next proposal distribution $\tilde{f}(\xi; \theta_{j+1})$. Formally, the next parameters are given by

$$\begin{aligned} \theta_{j+1} &= \arg \min_{\theta'} D_{\text{KL}}\left(\tilde{f}_j^*(\xi) \parallel \tilde{f}(\xi; \theta')\right) \\ &= \arg \min_{\theta'} H\left(\tilde{f}_j^*(\xi), \tilde{f}(\xi; \theta')\right) - H\left(\tilde{f}_j^*(\xi)\right) \end{aligned}$$

where $H(\cdot, \cdot)$ is the cross-entropy between both distributions, and $H(\tilde{f}_j^*(\xi))$ is the entropy of $\tilde{f}_j^*(\xi)$. Since entropy $H(\cdot)$ does not depend on θ' , minimizing the KL divergence is equivalent to minimizing the cross-entropy

$$\begin{aligned} \theta_{j+1} &= \arg \min_{\theta'} - \int_{\mathbb{R}^{d_\xi}} \tilde{f}_j^*(\xi) \log \tilde{f}(\xi; \theta') d\xi \\ &= \arg \min_{\theta'} - \int_{\mathbb{R}^{d_\xi}} \frac{\mathbb{1}_{J(\xi) \leq \gamma_j} \tilde{f}(\xi; \theta_j)}{l_j} \log \tilde{f}(\xi; \theta') d\xi, \end{aligned}$$

Algorithm 1: Cross-Entropy Method

Input: Initial parameters θ_0 , elite set size N_e , number of samples N_{CEM} , number of iterations j_{max}

- 1 **for** $j \leftarrow 0$ **to** $j_{\text{max}} - 1$ **do**
- 2 Sample $\{\xi^{(i)}\}_{i=1}^{N_{\text{CEM}}} \sim \tilde{f}(\xi; \theta_j)$
- 3 Evaluate costs $\{J(\xi^{(i)})\}_{i=1}^{N_{\text{CEM}}}$
- 4 Select elite set \mathcal{E}_j of the N_e best samples
- 5 Update parameters θ_{j+1} using \mathcal{E}_j
- 6 **return** best $\xi^{(i)}$ from final elite set $\mathcal{E}_{j_{\text{max}}}$

where the unknown constant l_j can be dropped, as it does not depend on θ' . By rewriting the integral as an expectation with respect to the *known* current PDF $\tilde{f}(\xi; \theta_j)$, we get

$$\theta_{j+1} = \arg \min_{\theta'} -\mathbb{E}_{\tilde{f}(\xi; \theta_j)} \{ \mathbb{1}_{J(\xi) \leq \gamma_j} \log \tilde{f}(\xi; \theta') \} .$$

This expectation can be approximated by samples $\{\xi^{(i)}\}_{i=1}^N$ from $\tilde{f}(\xi; \theta_j)$. This yields the objective for the next set of parameters

$$\theta_{j+1} = \arg \min_{\theta'} -\frac{1}{N} \sum_{i=1}^N \mathbb{1}_{J(\xi^{(i)}) \leq \gamma_j} \log \tilde{f}(\xi; \theta') .$$

Therefore, the new parameters θ' are fitted to the subset of best samples $\mathcal{E}_j = \{\xi^{(i)} \mid J(\xi^{(i)}) \leq \gamma_j\}$, i.e., the *elite set*.

For common choices of \tilde{f} , the minimization has simple closed-form solutions, e.g., for a Gaussian $\tilde{f}(\xi; \theta) = \mathcal{N}(\xi; \hat{\xi}, \mathbf{C})$, where the solution is the maximum likelihood estimate on the elite set

$$\begin{aligned} \hat{\xi}_{j+1} &= \frac{1}{|\mathcal{E}_j|} \sum_{\xi^{(i)} \in \mathcal{E}_j} \xi^{(i)} , \\ \mathbf{C}_{j+1} &= \frac{1}{|\mathcal{E}_j|} \sum_{\xi^{(i)} \in \mathcal{E}_j} (\xi^{(i)} - \hat{\xi}_{j+1})(\xi^{(i)} - \hat{\xi}_{j+1})^\top , \end{aligned} \quad (4)$$

with $\hat{\xi}_{j+1}$ and \mathbf{C}_{j+1} being the mean and covariance matrix of the Gaussian proposal summarized in parameter vector $\theta_{j+1} = (\hat{\xi}_{j+1}, \mathbf{C}_{j+1})$ and the elite set size $|\mathcal{E}_j|$.

In each iteration, γ_j is reduced to focus on better-performing samples (i.e., more rare events). In practice, γ_j is usually not specified directly but rather, the N_e -th best cost among the samples is used, e.g., the best 10 samples out of 100. After the last iteration j_{max} , the best sample from the final elite set $\mathcal{E}_{j_{\text{max}}}$ is returned as the solution. Note that in some variants, the mean $\hat{\xi}_{j_{\text{max}}}$ of the final proposal distribution is returned instead.

In summary, the CEM is an iterative method that refines the proposal distribution over the solution space, concentrating on regions with lower costs, and effectively guiding the search toward optimal or near-optimal solutions. The complete CEM procedure is shown in Alg. 1.

Algorithm 2: Cross-Entropy Method MPC Step

Input: State x_k , initial parameters $\theta_0 = (\hat{\xi}_0, \mathbf{C}_0)$

- 1 **for** $j \leftarrow 0$ **to** $j_{\text{max}} - 1$ **do**
- 2 Sample $\mathbf{u}_{k:k+N_H-1}^{(1)}, \dots, \mathbf{u}_{k:k+N_H-1}^{(N_{\text{CEM}})} \sim \tilde{f}(\cdot; \theta_j)$
- 3 Trajectory shooting using (2) // parallel
- 4 Evaluate costs $\{J_k(\mathbf{u}_{k:k+N_H-1}^{(i)})\}_{i=1}^{N_{\text{CEM}}}$ using (1) // parallel
- 5 Select elite set \mathcal{E}_j of the N_e best samples
- 6 Update parameters θ_{j+1} using \mathcal{E}_j
- 7 **return** first control \mathbf{u}_k^* from the best sequence in $\mathcal{E}_{j_{\text{max}}}$

B. Application to MPC

To apply the CEM to the OCP from Sec. II, the entire control sequence over the prediction horizon $\mathbf{u}_{k:k+N_H-1} = \mathbf{u}_k, \dots, \mathbf{u}_{k+N_H-1}$, is treated (in flattened form) as the random vector to be optimized at each time step k . I.e., the optimization variable is the random vector $\xi = [\mathbf{u}_k^\top, \dots, \mathbf{u}_{k+N_H-1}^\top]^\top$, whose sample space is $\mathbb{R}^{d_u \cdot N_H}$, and the goal is to find a realization that minimizes the cost function J_k (1).

The sequence of control inputs can be viewed as a discrete-time stochastic process. A natural choice for the proposal distribution $\tilde{f}(\cdot)$ is therefore a multivariate Gaussian distribution over the flattened control sequence $\tilde{f}(\xi; \theta_j) = \mathcal{N}(\xi; \hat{\xi}_j, \mathbf{C}_j)$, where the parameters $\theta_j = (\hat{\xi}_j, \mathbf{C}_j)$ consist of the mean control sequence $\hat{\xi}_j = [\hat{\mathbf{u}}_k^\top, \dots, \hat{\mathbf{u}}_{k+N_H-1}^\top]^\top$ and the covariance matrix $\mathbf{C}_j \in \mathbb{R}^{(d_u \cdot N_H) \times (d_u \cdot N_H)}$. This formulation is powerful because the covariance matrix \mathbf{C}_j can model temporal correlations between control inputs at different time steps, effectively treating the control sequence as a Gaussian Process. As with standard CEM, the parameters are iteratively updated based on the elite set of best-performing control sequences. Notably, trajectory shooting and cost evaluation can be *efficiently parallelized* across all sampled control sequences on modern hardware. The overall procedure for a single CEM-MPC step is summarized in Alg. 2.

C. Practical Improvements and Related Work

A major challenge in this basic approach is the high dimensionality of the optimization problem. While adapting a full covariance matrix can capture correlations, it is prone to estimation errors unless a large number of samples is used. To address this, a common simplification is to assume a diagonal covariance matrix \mathbf{C}_j , e.g., as done in [12], which reduces the number of parameters to be estimated and simplifies the sampling process, i.e., $\mathbf{C}_j = \text{diag}(\sigma_{j,1}^2, \dots, \sigma_{j,d_u \cdot N_H}^2)$, where $\sigma_{j,i}^2$ is the variance of the i -th control input in the flattened sequence. However, this assumption ignores temporal correlations between control inputs at different time steps, which can be crucial for generating smooth control sequences.

A standard technique to improve convergence smoothness is to introduce a momentum term in the parameter updates [11], which smooths the updates across iterations. The momentum update is given by $\hat{\xi}_{j+1} = \alpha \hat{\xi}_j + (1 - \alpha) \hat{\xi}_{j,e}$, where $\hat{\xi}_{j,e}$ is

the mean from the previous iteration, $\hat{\xi}_{j,e}$ is the mean of the current elite set, and $\alpha \in [0, 1)$ is a momentum factor. Analogously, the momentum can also be applied to the covariance matrix update [11].

A further standard improvement, used in various MPC approaches [1], is to warm-start the optimization at each time step by shifting the mean control sequence from the previous time step. This is achieved by setting the initial mean to $\hat{\xi}_0 = [u_{k-1,1}^\top, \dots, u_{k-1,N_H-1}^\top, u_{\text{init}}^\top]^\top$, where $u_{k-1,i}$ is the i -th control input from the optimized sequence at the previous time step and u_{init} is an initial guess for the last control input, often set to zero or the last optimized input.

A notable improvement is the iCEM algorithm [2], which introduces temporal correlations in a structured way. Instead of adopting a full covariance matrix, iCEM samples control actions from colored noise, imposing a smoothness prior on the control sequences. This is achieved by generating samples from a stationary distribution whose power spectral density follows a power law, $PSD(f) \propto 1/f^\beta$, where f is frequency and β controls the noise color. To further increase sample efficiency, iCEM retains a fraction of the elite set from the previous MPC time step, shifting the sequences in time and adding them to the current sample pool. Other improvements in iCEM include clipping control samples to ensure feasibility and decaying the number of samples over iterations via $N_{\text{CEM},j} = \max(N_{\text{CEM}}/\eta^j, 2N_e)$, where $\eta \geq 1$ is a decay factor. This reduces computational cost as the optimizer converges. An overview of the iCEM algorithm is provided in [2, Alg. 1].

Besides the improvements mentioned above, further enhancements focus on learning the sampling distribution itself, for instance, by using Gaussian processes [13] or normalizing flows [14].

However, in this work, we focus on the sample generation process itself by introducing deterministic sampling to CEM-MPC. This approach is orthogonal to the learning-based improvements and can be combined with them.

IV. CROSS-ENTROPY METHOD MPC USING DETERMINISTIC SAMPLES

As an alternative to random sampling, we propose to use deterministic samples based on the LCD. In this section, we briefly summarize how to obtain such optimal samples and then describe how to integrate them into CEM-MPC.


A. Deterministic Samples Using Localized Cumulative Distributions

The LCD is a multivariate generalization [5] of the univariate cumulative density function. It is defined as

$$F(\underline{m}, b) = \int_{\mathbb{R}^{d_\xi}} f(\xi) K(\xi, \underline{m}, b) d\xi ,$$

where $f(\xi)$ is a PDF over \mathbb{R}^{d_ξ} , and $K(\cdot, \underline{m}, b)$ is a kernel function centered at $\underline{m} \in \mathbb{R}^{d_\xi}$ with bandwidth $b \in \mathbb{R}_+$. The kernel function is typically chosen as a Gaussian kernel [5]

$$K(\xi, \underline{m}, b) = \exp\left(-\frac{\|\xi - \underline{m}\|_2^2}{2b^2}\right) .$$



(a) Standard normal samples (b) Transformed samples
Fig. 2: Example of 25 two-dimensional deterministic samples, where the background color indicates the PDF.

We then obtain a general distance measure between two PDFs by comparing their respective LCDs with a modified Cramér-von Mises (CvM) distance [5]

$$D_{\text{CvM}} = \int_0^\infty w(b) \int_{\mathbb{R}^{d_\xi}} \left(\tilde{F}(\underline{m}, b) - F(\underline{m}, b) \right)^2 d\underline{m} db ,$$

where $\tilde{F}(\underline{m}, b)$ and $F(\underline{m}, b)$ are the LCDs of $\tilde{f}(\xi)$ and $f(\xi)$, respectively, and $w(b)$ is a weighting function, typically $w(b) = b^{1-d_\xi}$. Here we use this distance to compare a Gaussian PDF $\tilde{f}(\xi)$ with a Dirac mixture $f(\xi) = \frac{1}{N} \sum_{i=1}^N \delta(\xi - \xi^{(i)})$. By minimizing D_{CvM} w.r.t. the sample locations $\xi^{(i)}$ we obtain an optimal Dirac mixture approximation of the Gaussian $\tilde{f}(\xi)$ [6]. An example of such optimal samples in two dimensions is shown in Fig. 2.

B. Integration of Deterministic Samples into CEM-MPC

The core idea is to use deterministic samples rather than random ones in the CEM. To speed up, we use pre-computed deterministic samples since online computation is prohibitive. Optimal samples $\{\tilde{\xi}^{(i)}\}$ are generated offline for the isotropic standard Gaussian $\mathcal{N}(\xi; \mathbf{0}, \mathbf{I})$ [15], [16]. These samples are then transformed at runtime (Alg. 2, line 2) to match the current proposal distribution $\mathcal{N}(\xi; \hat{\xi}_j, \mathbf{C}_j)$ of the CEM optimizer via [16]

$$\xi^{(i)} = \hat{\xi}_j + \mathbf{L}_j \tilde{\xi}^{(i)} , \quad (5)$$

where \mathbf{L}_j is the matrix square root of the Gaussian process covariance matrix, i.e., $\mathbf{C}_j = \mathbf{L}_j \mathbf{L}_j^\top$. While this transformation does not preserve the samples' optimality in the CvM sense [16] for a non-isotropic target distribution, it provides a practical and efficient way to generate structured samples with low discrepancy.

A naive implementation would use the same transformed set of deterministic samples in every iteration of the optimizer. However, we found this limits exploration and can lead to premature convergence, especially with a small number of samples. To address this, we propose three schemes for generating varied sample sets across iterations.

a) *Sample Set Variability Schemes:* To enhance exploration, we introduce variability into the sampling process either across iterations or across time steps.

(V1) **Random Rotation:** In each iteration j of the CEM, the pre-computed isotropic standard Gaussian samples $\{\tilde{\xi}^{(i)}\}$

are rotated by a random rotation matrix \mathbf{R}_j before the transformation. The matrix is drawn from the special orthogonal group $SO(d_\xi)$, which consists of all $d_\xi \times d_\xi$ orthogonal matrices with a determinant of +1 [17]. Therefore the transformation in (5) changes to $\xi^{(i)} = \hat{\xi}_j + \mathbf{L}_j \mathbf{R}_j \tilde{\xi}^{(i)}$. Note that rotating the samples of an isotropic distribution results in another valid sample set without losing their optimality in the CvM sense. This introduces stochasticity, improving exploration at the cost of losing the fully deterministic nature of the algorithm.

(V2) **Deterministic Joint Density Sampling:** To maintain a fully deterministic algorithm, we pre-compute optimal deterministic samples from a higher-dimensional random vector that encompasses all CEM iterations, i.e., we generate deterministic samples in dimension $d_\xi \cdot j_{\max}$. Each high-dimensional sample $[\tilde{\xi}_1^{(i)\top}, \dots, \tilde{\xi}_{j_{\max}}^{(i)\top}]^\top$ is then partitioned into j_{\max} separate samples of dimension d_ξ , providing a unique, deterministic sample set for each iteration, which are then transformed using (5).

(V3) **Combined Approach:** This hybrid scheme balances determinism within the optimization loop with stochastic exploration across time steps. A single set of high-dimensional samples is pre-computed as in (V2). At the beginning of each MPC time step k , this set is rotated by a single random matrix $\mathbf{R}_k \in SO(d_\xi)$. The resulting samples are then partitioned and used deterministically for all subsequent CEM iterations within that time step.

b) *Covariance Matrix Structures and Adaptation:* Given the sample set by one of the above schemes, the samples are transformed to match the current proposal distribution of the CEM optimizer using (5). The use of deterministic samples is compatible with, e.g., maintaining a diagonal covariance matrix, or by adapting a full covariance matrix. We focus on two approaches that incorporate temporal correlations, which are known to be beneficial for generating smooth control sequences [2]. The initial correlation structure is derived from the power spectral density of colored noise (e.g., pink noise, $\beta = 1$) via the Wiener–Khinchin theorem, resulting in a Toeplitz structured correlation matrix \mathbf{C}_ρ [18, pp. 576–578].

(M1) **Fixed Correlation with Adaptive Variance:** In this approach, the correlation structure of the proposal distribution is fixed throughout the optimization. The covariance matrix is constructed as $\mathbf{C}_j = \text{diag}(\sigma_j) \mathbf{C}_\rho \text{diag}(\sigma_j)$, where \mathbf{C}_ρ is the fixed time-correlation matrix and σ_j is the vector of marginal standard deviations. In each iteration, only the marginal variances are updated based on the elite set, which simplifies (4) to

$$\sigma_{j+1}^2 = \frac{1}{|\mathcal{E}_j|} \sum_{\xi^{(i)} \in \mathcal{E}_j} (\xi^{(i)} - \hat{\xi}_{j+1})^2.$$

This reduces the number of parameters to be estimated and is less demanding on the number of elite samples; a minimum of two non-identical elite samples is sufficient. The transformation in (5) then uses $\mathbf{L}_j = \text{diag}(\sigma_j) \mathbf{A}_\rho$, where \mathbf{A}_ρ is the matrix square root of \mathbf{C}_ρ .

(M2) **Adaptive Full Covariance:** This strategy allows for greater flexibility by adapting the full covariance matrix

TABLE I: Task-specific parameters for evaluation.

Parameter	Mountain Car	Cart-Pole Swing-Up
Control input limits	$u \in [-1, 1]$	$u \in [-20, 20] \text{N}$
Goal state \underline{x}_g ($\hat{\underline{x}}_g$)	$[\pi/2, 0]^\top$	$[0, 0, 1, 0, 0]^\top$
State weights \mathbf{Q}	$\text{diag}([1, 1])$	$\text{diag}([0.1, 0.1, 1, 0.1, 0.1])$
Control weight r	0.1	10^{-4}
Terminal weights \mathbf{Q}_{N_H}	$\text{diag}([1, 1])$	$\text{diag}([10, 0.1, 10, 0.1, 0.1])$
Process noise \mathbf{C}^w	$\text{diag}([0, 10^{-7}])$	$\text{diag}([0, 10^{-8}, 0, 10^{-8}])$
Time discretization Δt	3 s	0.02 s
Total time steps T	150	300
Noise color for CEM β	0.25	1.0
Initial $\hat{\xi}_0$ for CEM	0	0
Initial $\hat{\sigma}_0$ for CEM	$1.5 \cdot \mathbf{1}$	$10 \cdot \mathbf{1}$

\mathbf{C}_j in each iteration. The optimization is initialized with a structured covariance matrix based on colored noise as in the previous scheme. Subsequently, the entire matrix is updated using the maximum likelihood estimate from the elite set. This allows the optimizer to adapt the temporal correlations online. However, estimating a full $d_\xi \times d_\xi$ covariance matrix requires the elite set to contain at least $d_\xi + 1$ non-collinear samples [19].

The sample set variation schemes (V1–V3) and covariance adaptation methods (M1–M2) can be combined arbitrarily. For instance, random rotations (V1) can be used with either a fixed correlation structure with adaptive variances (M1) or a full adaptive covariance matrix (M2). This modularity allows the proposed deterministic sampling strategies to serve as drop-in replacements for the random sampling step in various CEM–MPC algorithms.

V. EXPERIMENTS

To evaluate the proposed CEM–MPC with deterministic samples, we conduct experiments on two tasks: (i) the mountain car [20] and (ii) the cart-pole swing-up task [21].

In contrast to standard simulation environments, such as Gymnasium [20], where the system dynamics are deterministic, we add additive Gaussian process noise when the optimized control inputs are applied to the system during simulation. The simulated dynamics are then given by

$$\underline{x}_{k+1} = \underline{a}(\underline{x}_k, u_k) + \underline{w}_k,$$

where \underline{w}_k is sampled from a zero-mean Gaussian PDF $\mathcal{N}(\underline{w}; \underline{0}, \mathbf{C}^w)$ with process noise covariance \mathbf{C}^w . However, the controller does not have access to the process noise or its PDF and assumes deterministic dynamics $\underline{a}(\cdot, \cdot)$ for prediction, and therefore has to deal with random disturbances. For both tasks, quadratic stage cost functions of the form $g_n(\underline{x}_n, u_n) = (\underline{x}_n - \underline{x}_g)^\top \mathbf{Q}(\underline{x}_n - \underline{x}_g) + r \cdot u_n^2$ are used, where \underline{x}_g is the goal state, and \mathbf{Q} and r are state and control weights, respectively. The terminal cost function is set to $g_{N_H}(\underline{x}_{k+N_H}) = (\underline{x}_{k+N_H} - \underline{x}_g)^\top \mathbf{Q}_{N_H}(\underline{x}_{k+N_H} - \underline{x}_g)$. The task-specific parameters are summarized in Tab. I.

For the evaluation, we denote the use of our deterministic sampling strategies as *deterministic sampling CEM* (dsCEM). Specifically, we refer to the adaptation methods from Sec. IV-B as *dsCEM-Var* (M1) and *dsCEM-Cov* (M2), and for the variability schemes we add *V1*, *V2*, and *V3*, when the specific

scheme is used. dsCEM-Var is evaluated with all three variability schemes, while dsCEM-Cov is only evaluated with scheme V3, to keep visualizations clear.

We compare our approach against the iCEM method [2], where the only difference is that we swapped the random sampling step with our proposed deterministic sampling strategies, i.e., both methods have the same computational budget. For a fair comparison, both methods share hyperparameters as used in [2]: a horizon $N_H = 30$, $j_{\max} = 3$ iterations, momentum $\alpha = 0.1$, and shift-initialization warm-starting (see Sec. III-C). The elite set size is $N_e = 10$ (or $N_e = 40$ for dsCEM-Cov), and a fraction of the top 0.3 of the elite set is carried over to the next time step (rounded down). The initial correlation for dsCEM is derived from the noise color β used in iCEM as described in Sec. IV-B. To isolate the impact of the sampling strategy, iCEM’s sample decay is disabled, ensuring an equal number of shot trajectories for both methods.

We evaluate all methods over a sample size N_{CEM} ranging from 20 to 300. Each configuration is repeated for 100 runs with different random seeds. The range for dsCEM-Cov starts at 40 samples due to its higher elite sample requirement. Additionally, we establish a performance baseline by running iCEM with an extensive sample size of 10^4 .

We evaluate controller performance using cumulative cost and control input smoothness. The cumulative cost is the sum of stage costs $g_k(\underline{x}_k, \underline{u}_k)$ over the entire simulation. The smoothness is measured using [14]

$$S = \sum_{k=1}^{T-1} \|\underline{u}_k - \underline{u}_{k-1}\|^2,$$

where a lower value of S indicates a smoother trajectory by penalizing large changes between consecutive control inputs.

A. Mountain Car Task

The dynamics of the mountain car task [20] are given by

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -0.0025 \cos(3x_1) \end{bmatrix} + \begin{bmatrix} 0 \\ 0.0015 \end{bmatrix} u,$$

where x_1 is the position and x_2 is the velocity. The system is discretized using Runge–Kutta 4th order (RK4) integration. The main challenge is that the underpowered car cannot drive directly up the hill and must perform a swing-up maneuver. Unlike the standard task, our objective is to reach the goal state at the top of the hill *and stop*, which is more difficult than reaching the top position with an arbitrary final velocity. The state is initialized with position x_1 drawn uniformly from $[-0.7, -0.3]$, and with zero initial velocity.

The results for the mountain car task are presented in Fig. 3. In terms of cumulative cost (Fig. 3a), all proposed dsCEM methods outperform iCEM, particularly at small sample sizes. While the performance of all methods converges for larger sample sizes, they do not reach the baseline established by iCEM with 10^4 samples. Regarding smoothness (Fig. 3b), all dsCEM variants, except for dsCEM-Cov V3, produce smoother control trajectories than iCEM. Notably, dsCEM-Var V2 consistently achieves the highest degree of smoothness (lowest values) across most sample sizes. An interesting trend

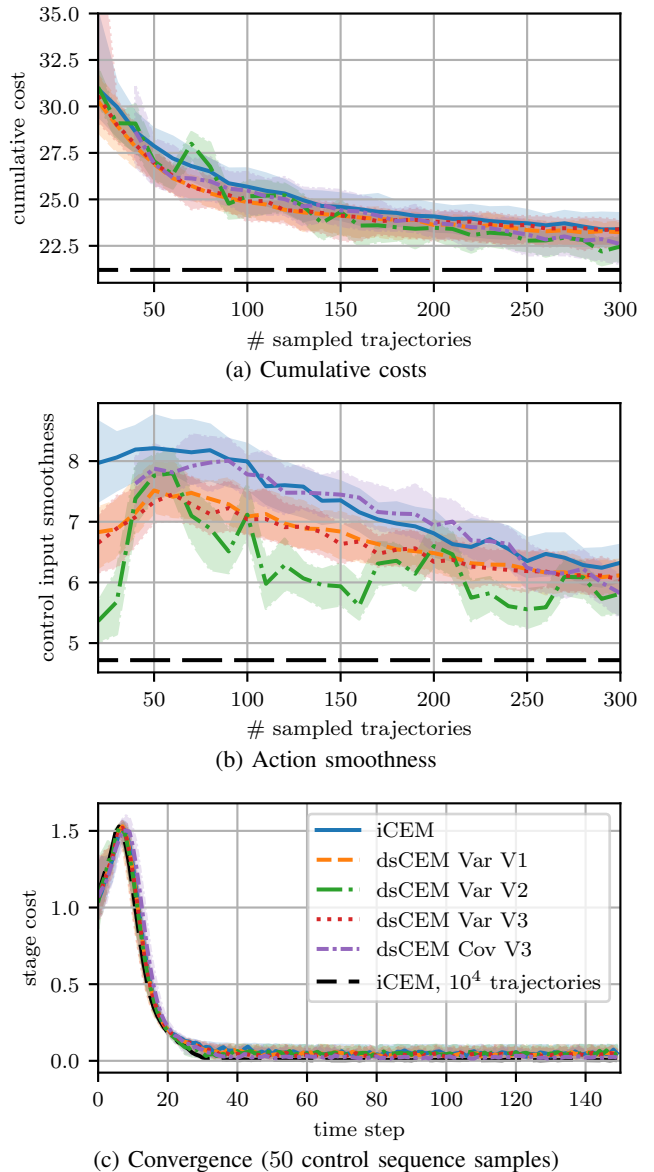


Fig. 3: The results for the Mountain Car Task are given for (a) cumulative cost and (b) control input smoothness over sample size, as well as for (c) convergence behavior for a fixed sample size of 50. All plots show the median (line) and the interquartile range (shaded area) across 100 runs. The different methods’ colors are consistent across all plots.

is visible in the smoothness measure, where the score for all methods initially increases, peaking around a sample size of 50, before decreasing again. This peak likely indicates the sample budget required to discover aggressive swing-up maneuvers that are effective for reducing costs. With fewer samples, the optimizer settles for smoother but less optimal trajectories, while larger budgets allow it to find and subsequently refine these more dynamic strategies. The convergence behavior for a sample size of 50 is depicted in Fig. 3c. While the median stage cost evolves similarly for all methods, our proposed dsCEM variants demonstrate

slightly faster convergence, particularly around time step 30, and their performance closely approaches the extensive iCEM baseline. As the costs approach zero, the differences between the methods become less significant because all controllers primarily focus on counteracting disturbances.

B. Cart-Pole Swing-Up Task

The dynamics of the frictionless cart-pole swing-up task are given by [21]

$$\ddot{\phi} = \frac{g \sin(\phi) - \cos(\phi) \cdot \frac{u + m_p l \dot{\phi}^2 \sin(\phi)}{m_p + m_c}}{l \left(\frac{4}{3} - \frac{m_p \cos^2(\phi)}{m_p + m_c} \right)},$$

$$\ddot{x} = \frac{u + m_p l \left(\dot{\phi}^2 \sin(\phi) - \ddot{\phi} \cos(\phi) \right)}{m_p + m_c},$$

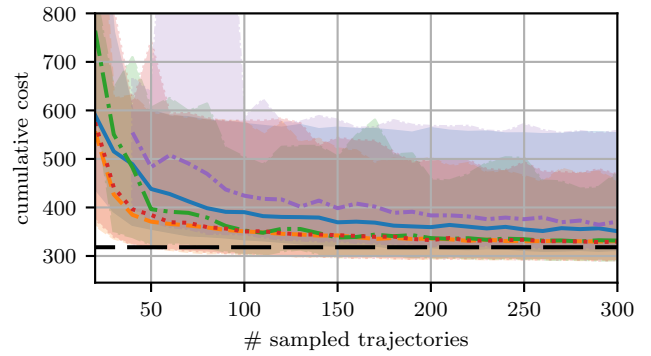
where the state is $\mathbf{x} = [x, \dot{x}, \phi, \dot{\phi}]^\top$, with cart position x , pole angle ϕ (where $\phi = 0$ is the upright position), and their respective (angular) velocities. The physical parameters are cart mass $m_c = 1$ kg, pole mass $m_p = 0.1$ kg, pole length $l = 0.5$ m, and gravity $g = 9.81$ m s⁻². The system is discretized using RK4. The initial state is set to zero for the cart's position and both velocities, while the pole's angle ϕ is uniformly sampled from $[145^\circ, 215^\circ]$. The costs are evaluated using an augmented state vector $\tilde{\mathbf{x}} = [x_1, x_2, \cos(x_3), \sin(x_3), x_4]$ to account for the periodicity of the angle x_3 .

In terms of cumulative cost (Fig. 4a), most proposed dsCEM methods outperform iCEM and closely approach the performance of the extensive iCEM baseline. However, the submethod based on the full covariance (dsCEM Cov V3) struggled to match the performance of our variance-based methods (dsCEM Var V1–V3). This suggests that using a high-dimensional covariance matrix may not be suitable for this task with small sample sizes. Regarding control input smoothness (Fig. 4b), all dsCEM methods demonstrate superior performance over iCEM, with dsCEM-Var V2 being the smoothest across all sample sizes. Notably, with more than 50 samples, dsCEM-Var V2 even surpasses the smoothness of the extensive iCEM baseline. This is further illustrated in Fig. 4d, which shows that the control input trajectories of dsCEM-Var V2 are significantly less noisy than those of iCEM with the same sample budget, and even smoother than the extensive baseline. The convergence plot for a sample size of 50 (Fig. 4c) confirms that most dsCEM variants converge faster than iCEM. However, the wide interquartile range for dsCEM-Var V3 suggests that this specific variant is less robust.

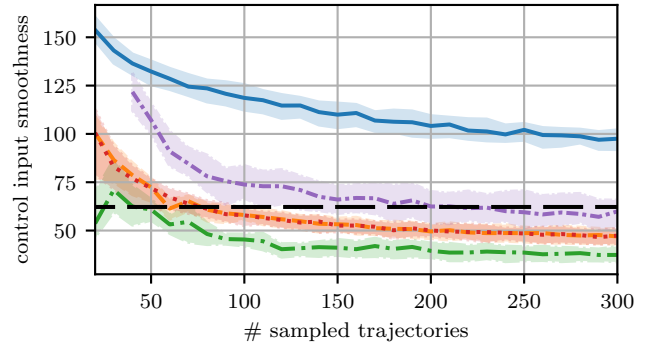
C. Discussion

The experimental results demonstrate that replacing standard random sampling in CEM with our proposed deterministic sampling strategies, termed dsCEM, leads to significant improvements in performance and sample efficiency. Across both the mountain car and cart-pole swing-up tasks, the dsCEM variants consistently outperform the baseline iCEM method, particularly when the number of samples is limited.

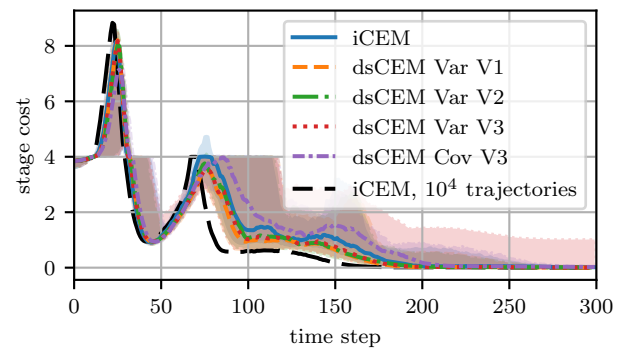
A key finding is the superior control input smoothness achieved by the dsCEM methods. While smooth controls can



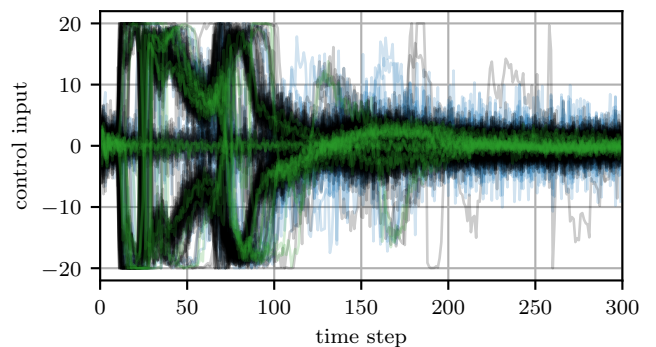
(a) Cumulative costs



(b) Action smoothness



(c) Convergence (50 control input samples)



(d) Applied control inputs for 20 randomly selected runs, with each method distinguished by color. The optimization for each run was performed using 300 samples.

Fig. 4: Results for the cart-pole task. The different methods' colors are consistent across all plots. To give intuition about runtime, for 100 samples iCEM and dsCEM both take 166 ms per MPC step on one Intel Xeon Platinum 8358 core, indicating no measurable overhead from deterministic sampling and variation schemes; parallelization can further reduce runtime.

also be encouraged by adding explicit smoothness penalties, this introduces additional hyperparameters to trade off control performance and smoothness; in our view, the optimizer should produce smooth controls without such extra tuning. This is especially pronounced for the fully deterministic variant dsCEM-Var V2, which consistently yields the smoothest control trajectories. This outcome is expected, as deterministic sample sets provide low-discrepancy coverage of the sampling space, mitigating the clustering and gaps inherent to random sampling. Notably, in the cart-pole swing-up task, dsCEM produces smoother control inputs than even the extensive iCEM baseline, despite using a fraction of the samples. This highlights the ability of deterministic sampling to find high-quality, smooth solutions with remarkable efficiency.

In terms of cumulative cost, dsCEM shows a clear advantage at lower sample sizes, indicating faster convergence to effective control strategies. Although all methods tend to converge as sample sizes increase, the gains from dsCEM in the low-sample regime are critical for practical applications.

The enhanced sample efficiency of dsCEM has profound implications for the deployment of MPC. Reducing the required number of samples lowers computational demand, a critical factor for real-time performance. This reduction impacts four key stages of the CEM algorithm: (i) the creation of random or deterministic samples; (ii) the determination of the elite set; (iii) the computationally expensive trajectory shooting, especially for complex system dynamics; and (iv) the evaluation of cost functions for each trajectory and all its stages. Furthermore, as hardware generally has limited parallelization capabilities, even a slight reduction in sample size can prevent substantial computational overhead. For example, on such hardware, an algorithm requiring one more sample than the hardware's capacity can *double the execution time*. By achieving superior performance with fewer samples, dsCEM promises to apply MPC to systems with more complex models and longer prediction horizons without sacrificing real-time capability.

VI. CONCLUSION

In this paper, we addressed the limitations of random sampling in CEM-MPC by introducing dsCEM, a novel framework that leverages deterministic samples based on LCDs. We proposed several schemes for integrating these samples, demonstrating that this approach is a modular, drop-in replacement for the conventional sampling step. Our experimental evaluation on two nonlinear control benchmarks showed that dsCEM consistently outperforms the state-of-the-art iCEM method in terms of both cumulative cost and control input smoothness, especially in the critical low-sample regime. These findings highlight the potential of deterministic sampling to significantly improve the efficiency of CEM-MPC, making it a promising option for real-time control on computationally constrained hardware.

As our approach is orthogonal to learning-based improvements, a viable option is to combine dsCEM with these techniques, e.g., using a learned policy to warm-start the optimization process and apply it to robotic systems.

REFERENCES

- [1] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed. Madison, WI, USA: Nob Hill Publishing, 2017.
- [2] C. Pinneri, S. Sawant, S. Blaes, J. Achterhold, J. Stueckler, M. Rolinek, and G. Martius, "Sample-efficient cross-entropy method for real-time planning," in *Proceedings of the 2020 Conference on Robot Learning*, vol. 155, Nov. 2021, pp. 1049–1065.
- [3] R. Rubinstein, "The cross-entropy method for combinatorial and continuous optimization," *Methodology And Computing In Applied Probability*, vol. 1, no. 2, pp. 127–190, 1999.
- [4] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 4759–4770.
- [5] U. D. Hanebeck and V. Klumpp, "Localized cumulative distributions and a multivariate generalization of the Cramér-von Mises distance," in *Proceedings of the 2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2008)*, Seoul, Republic of Korea, August 2008, p. 33–39.
- [6] U. D. Hanebeck, M. F. Huber, and V. Klumpp, "Dirac mixture approximation of multivariate Gaussian densities," in *Proceedings of the 2009 IEEE Conference on Decision and Control (CDC 2009)*, Shanghai, China, December 2009.
- [7] M. Diehl, H. Bock, H. Diedam, and P.-B. Wieber, "Fast direct multiple shooting algorithms for optimal robot control," in *Fast Motions in Biomechanics and Robotics: Optimization and Feedback Control*. Berlin, Heidelberg: Springer, 2006, pp. 65–93.
- [8] D. H. Jacobson and D. Q. Mayne, *Differential Dynamic Programming*, ser. Modern Analytic and Computational Methods in Science and Mathematics. New York, NY: American Elsevier Publ, 1970, no. 24.
- [9] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," in *Proceedings of the First International Conference on Informatics in Control, Automation and Robotics*. Setúbal, Portugal, 2004, p. 222–229.
- [10] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.
- [11] P.-T. De Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein, "A tutorial on the cross-entropy method," *Annals of Operations Research*, vol. 134, no. 1, pp. 19–67, Feb. 2005.
- [12] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, "Learning latent dynamics for planning from pixels," in *Proceedings of the 36th International Conference on Machine Learning*, vol. 97, Jun. 2019, pp. 2555–2565.
- [13] J. Watson and J. Peters, "Inferring smooth control: Monte Carlo posterior policy iteration with Gaussian processes," in *Proceedings of the 6th Conference on Robot Learning*, vol. 205, Dec. 2023, pp. 67–79.
- [14] T. Power and D. Berenson, "Learning a generalizable trajectory sampling distribution for model predictive control," *IEEE Transactions on Robotics*, vol. 40, pp. 2111–2127, 2024.
- [15] J. Steinbring and U. D. Hanebeck, "S2KF: The smart sampling Kalman filter," in *Proceedings of the 16th International Conference on Information Fusion (Fusion 2013)*, Istanbul, Turkey, July 2013.
- [16] —, "LRKF revisited: The smart sampling Kalman filter (S2KF)," *Journal of Advances in Information Fusion*, vol. 9, no. 2, pp. 106–123, December 2014.
- [17] C. A. León, J.-C. Massé, and L.-P. Rivest, "A statistical model for random rotations," *Journal of Multivariate Analysis*, vol. 97, no. 2, pp. 412–430, 2006.
- [18] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice-Hall, Inc., 1993.
- [19] S. Julier and J. Uhlmann, "Reduced sigma point filters for the propagation of means and covariances through nonlinear transformations," in *Proceedings of the 2002 American Control Conference*, Anchorage, AK, USA, 2002, pp. 887–892.
- [20] M. Towers, A. Kwiatkowski, J. Terry, Balis *et al.*, "Gymnasium: a standard interface for reinforcement learning environments," *arXiv preprint:2407.17032*, 2024.
- [21] A. G. Barto, R. S. Sutton, and C. W. Anderson, "Neuronlike adaptive elements that can solve difficult learning control problems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, no. 5, pp. 834–846, 1983.