

Inferring Smooth Control: Monte Carlo Posterior Policy Iteration with Gaussian Processes

Joe Watson¹ Jan Peters¹²³⁴

¹Department of Computer Science, Technical University of Darmstadt

²Centre for Cognitive Science, Technical University of Darmstadt

³German Research Center for AI ⁴Hessian.AI

{joe,jan}@robot-learning.de

Abstract: Monte Carlo methods have become increasingly relevant for control of non-differentiable systems, approximate dynamics models and learning from data. These methods scale to high-dimensional spaces and are effective at the non-convex optimizations often seen in robot learning. We look at sample-based methods from the perspective of inference-based control, specifically posterior policy iteration. From this perspective, we highlight how Gaussian noise priors produce rough control actions that are unsuitable for physical robot deployment. Considering smoother Gaussian process priors, as used in episodic reinforcement learning and motion planning, we demonstrate how smoother model predictive control can be achieved using online sequential inference. This inference is realized through an efficient factorization of the action distribution and a novel means of optimizing the likelihood temperature to improve importance sampling accuracy. We evaluate this approach on several high-dimensional robot control tasks, matching the sample efficiency of prior heuristic methods while also ensuring smoothness. Simulation results can be seen at monte-carlo-ppi.github.io.

Keywords: approximate inference, policy search, model predictive control




Figure 1: High-dimensional, contact-rich tasks such as manipulation (left) can be performed effectively using sample-based model predictive control. While prior work uses correlated actuator noise to improve sample-efficiency and exploration, these methods do not preserve the smoothness in the downstream actuation a , resulting in aggressive control (center). We use smooth Gaussian process priors to infer posterior actions (right), which preserves smoothness while maintaining performance and sample efficiency, as both are using only 32 samples. Rewards r show quartiles over 25 seeds.

1 Introduction

Learning robot control requires optimization to be performed on sampled transitions of the environment [1]. Monte Carlo methods [2] provide a principled means to approach such algorithms, bridging black-box optimization and approximate inference techniques. These methods have been adopted extensively by the community for their impressive simulated [3, 4, 5, 6] and real-world [7, 8, 9, 10, 11, 12, 13, 14, 15] robot learning results. Their appeal includes requiring only function evaluations of the dynamics and objective, so can be applied to complex environments with minimal overhead (Figure 1). Moreover, their stochastic nature also avoids issues with local minima

that occur with gradient-based solvers [16, 17]. Finally, while Monte Carlo sampling is expensive, shooting methods can be effectively parallelized across processes and the advent of simulations on GPUs also provides a means of acceleration [18, 13]. However, some aspects of black-box optimization are open to criticism. Sample-based solvers such as the cross-entropy method (CEM) [19] appear ‘wasteful’, ignoring computation by throwing away the majority of samples, while others enforce high-entropy search distributions to avoid premature convergence [18]. Moreover, many design decisions and hyperparameters are heuristic in nature, which is undesirable from both the user- and research perspective when interpreting, tuning or advancing these methods.

In this work, we consider Monte Carlo optimal control through the broader perspective of inference-based control [20, 21, 22, 23, 24, 25, 26, 27], where optimization is achieved through importance sampling [28]. This approach covers settings such as policy search [29], motion planning [8, 30] and model predictive control (MPC) [18]. From this view point, we highlight two key design decisions: the likelihood temperature and the distribution over action sequences. An adaptive temperature scheme is crucial for controlling the optimization behavior across objectives and distributions, but in many methods this aspect is ignored or opaque. Moreover, correlated action sequences are equally crucial for performing effective exploration and control in practical settings. *Smoothness*, arising from such correlations, is an aspect of human motion [31]. Smooth priors have taken many forms across domains, such as movement primitives [32], smoothed- [11, 12] or coloured noise [4]. We use Gaussian processes [33] as action priors and show how they can be scaled to high-dimensional action spaces through factorization of the covariance. Evaluating on simulated robotic systems, we reproduce prior results on policy search while transferring these ideas to MPC, matching prior performance with respect to sample efficiency while ensuring smooth actuation.

Contribution. First, we present a perspective of episodic inference-based control based on Gibbs posteriors. Using this view, we then present novel Monte Carlo variants that incorporate the approximate inference error due to importance sampling, simplifying the hyperparameter while providing regularization. Thirdly, we demonstrate how richer Gaussian process priors can be combined with these regularized Gibbs posteriors for Monte Carlo MPC using online sequential inference, which achieves greater smoothness and sample efficiency than standard white noise priors. We highlight connections between this approach to MPC and effective prior approaches to episodic policy search.

2 Monte Carlo Methods for Optimal Control

This section outlines the problem setting and introduces variational optimization and posterior policy iteration methods. We consider the standard (stochastic) optimal control setting in discrete-time, with states $\mathbf{s} \in \mathbb{R}^{d_s}$, actions $\mathbf{a} \in \mathbb{R}^{d_a}$. Optimization is framed as maximizing a reward $r: \mathbb{R}^{d_s} \times \mathbb{R}^{d_a} \rightarrow \mathbb{R}$ under dynamics $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ and initial state distribution $p(\mathbf{s}_1)$,

$$\max_{\mathbf{a}_1, \dots, \mathbf{a}_T} \mathbb{E} \left[\sum_{t=1}^T r(\mathbf{s}_t, \mathbf{a}_t) \right] \quad \text{s.t.} \quad \mathbf{s}_{t+1} \sim p(\cdot | \mathbf{s}_t, \mathbf{a}_t), \quad \mathbf{s}_1 \sim p(\mathbf{s}_1). \quad (1)$$

This work focuses on the episodic setting, where optimization is performed after evaluating the current solution over a finite-time horizon T . We frequently use the episodic return R , where $R(\mathbf{S}, \mathbf{A}) = \sum_{t=1}^T r(\mathbf{s}_t, \mathbf{a}_t)$, using upper-case to denote sequences, e.g. $\mathbf{A} := \{\mathbf{a}_1, \dots, \mathbf{a}_T\}$.

2.1 Variational Optimization with Gibbs Posteriors

The optimization outlined above is amenable to gradient-based solvers such as stochastic differential dynamic programming [34]. However, to aid optimization through exploration and regularization, we can consider optimizing a parametric *belief* over action sequences $q \in \mathcal{Q}$. The variational formulation (Equation 3) generalizes Bayes’ rule beyond optimizing likelihoods and resembles many learning algorithms [35, 36]. This work concerns optimizing an open-loop action sequence to maximize an episodic return. Bayesian inference of an action sequence from data, known as input estimation, can be performed using message passing of the appropriate probabilistic graphical model, capturing the sequential structure of the problem and necessary priors [27]. If the measurement log-likelihood is replaced with the control objective, this inference computation can be shown to have precise dualities with dynamic programming-based optimal control [37]. While this switch in objective provides a powerful suite of inference tools for efficient computation, it requires treating the control objective as a Markovian log-likelihood, which is not the case for episodic objectives. The Gibbs likelihood is a general treatment of the objective-as-likelihood (Definition 1) [38, 39].

Definition 1. (*Gibbs likelihoods and posteriors*) For a loss f and prior $p(\mathbf{x})$, the Gibbs posterior q_α for parameter \mathbf{x} is derived by constructing the Gibbs likelihood $\exp(-\alpha f(\mathbf{x}))$ from the loss,

$$q_\alpha(\mathbf{x}) = \frac{1}{Z_\alpha} \exp(-\alpha f(\mathbf{x})) p(\mathbf{x}), \quad Z_\alpha = \int \exp(-\alpha f(\mathbf{x})) p(\mathbf{x}) d\mathbf{x}, \quad \alpha \geq 0. \quad (2)$$

This posterior minimizes the following objective

$$q_\alpha = \arg \min_{q \in \mathcal{Q}} \mathbb{E}_{\mathbf{x} \sim q(\cdot)} [f(\mathbf{x})] + \frac{1}{\alpha} \mathbb{D}_{\text{KL}}[q(\mathbf{x}) \parallel p(\mathbf{x})]. \quad (3)$$

This objective appears in PAC-Bayes methods [38], mirror descent methods [40] and Bayesian inference as the evidence lower bound objective when $f(\mathbf{x})$ is a negative log-likelihood [39].

Augmenting the variational optimization objective with prior regularization (Equation 3), we obtain an expression of the optimal belief in the action sequence (Equation 2). The parameter α has a range of meanings, depending on context. In PAC-Bayes it is the dataset size, in mirror descent it is an update step size and in risk-sensitive control it is the sensitivity [41, 42]. Example 1 in Appendix A examines a tractable linear-quadratic-Gaussian example of this update, demonstrating its relation to Newton-like optimization and highlighting the effect α has on the regularized update.

2.2 Posterior Policy Iteration

The optimal control problem (Equation 1) is ambiguous regarding whether the action sequence or state-action trajectory is the optimization variable. Applying the Gibbs posterior to the optimal control setting recovers Rawlik et al.’s posterior policy iteration [41], which can be implemented using the joint distribution or policy. We consider the following joint state-action distribution, that factorizes in the following Markovian fashion $p(\mathbf{S}, \mathbf{A}) = p(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t) p(\mathbf{a}_t \mid \mathbf{s}_t)$. Posterior policy iteration updates the state-action distribution through the policy, constructing a Gibbs likelihood from the reward, as the dynamics and initial state distribution are constant.

Definition 2. (*Posterior policy iterations (PPI)* [43]) As the initial distribution and dynamics are shared by the prior and posterior joint state-action distribution, the joint Gibbs posterior $q_\alpha(\mathbf{S}, \mathbf{A}) \propto \exp(\alpha R(\mathbf{S}, \mathbf{A})) p(\mathbf{S}, \mathbf{A})$ can be alternatively expressed using the policy posterior update $q_\alpha(\mathbf{A} \mid \mathbf{S}) \propto \exp(\alpha R(\mathbf{S}, \mathbf{A})) p(\mathbf{A} \mid \mathbf{S})$.

Using this update, the key decisions are choosing $p(\mathbf{A} \mid \mathbf{S})$, α and the inference approximation. If p and q_α are Gaussian, then PPI involves iterative refinement of the distribution. In the Monte Carlo setting, q_α takes the form of an importance-weighted empirical distribution. To apply iteratively, p is updated using the M-projection, following the objective (Equation 3), i.e. a weighted maximum likelihood fit of the policy parameters [29]. This approach is a stochastic approximate expectation maximization (SAEM) method [44] and described fully in Algorithm 1 in the Appendix. We argue a key aspect of PPI methods is how to specify the inverse temperature α during optimization (Section 3), as it has a strong influence on the posterior, which is important when fitting rich distributions such as Gaussian processes (Section 4) from samples. Gaussian process action priors can be applied to several control settings, such as policy search and model predictive control (Section 6).

3 Posterior Policy Constraints for Monte Carlo Optimization

The Gibbs posterior in Definition 2 has been adopted widely in control, albeit from a range of different perspectives, such as Bayesian smoothing [23], solutions to the Feynman-Kac equation [45], maximum entropy [26], mirror descent [46] and entropy-regularized reinforcement learning [47]. An open question is how best to set α for Monte Carlo optimization? Relative entropy policy search (Definition 3), provides a principled and effective means of deriving α for stochastic optimization, using the constrained optimization view of entropy-regularized optimal control.

Definition 3. (*Episodic relative entropy policy search (eREPS)* [29]) Maximize the expected return, subject to a hard KL bound ϵ on the policy update,

$$\max_{\theta} \mathbb{E}_{\mathbf{s}_{t+1} \sim p(\cdot \mid \mathbf{s}_t, \mathbf{a}_t), \mathbf{a}_t \sim q_\theta(\cdot \mid \mathbf{s}_t), \mathbf{s}_1 \sim p(\cdot)} [R(\mathbf{s}_t, \mathbf{a}_t)] \quad \text{s.t.} \quad \mathbb{D}_{\text{KL}}[q_\theta(\mathbf{A} \mid \mathbf{S}) \parallel p(\mathbf{A} \mid \mathbf{S})] \leq \epsilon.$$

The posterior policy takes the form $q_\theta(\mathbf{A} \mid \mathbf{S}) \propto \exp(\alpha R) p(\mathbf{A} \mid \mathbf{S})$, where α is derived from Lagrange multiplier calculated by minimizing the empirical dual $\mathcal{G}(\cdot)$ using N samples,

$$\min_{\alpha} \mathcal{G}(\alpha) = \frac{\epsilon}{\alpha} + \frac{1}{\alpha} \log \int p(\mathbf{S}, \mathbf{A}) \exp(\alpha R(\mathbf{S}, \mathbf{A})) d\mathbf{S} d\mathbf{A} \approx \frac{\epsilon}{\alpha} + \frac{1}{\alpha} \log \frac{1}{N} \sum_{n=1}^N \exp(\alpha R_n).$$

While REPS is a principled approach to stochastic optimization, we posit two weaknesses: The hard KL constraint ϵ is difficult to specify, as it depends on the optimization problem, distribution family and dimensionality. Secondly, the Monte Carlo approximation of the dual has no regularization and may poorly adhere to the KL constraint without sufficient samples. Therefore, we desire an alternative approach that resolves these two issues, capturing the Monte Carlo approximation error with a simpler hyperparameter. To tackle this problem, we interpret the REPS update as a pseudo-posterior, where the temperature is calculated using the KL constraint. We make this interpretation concrete by reversing the objective and constraint, switching to an equality constraint for the expectation,

$$\min_{\theta} \mathbb{D}_{\text{KL}}[q_{\theta}(\mathbf{A} \mid \mathbf{S}) \parallel p(\mathbf{A} \mid \mathbf{S})] \quad \text{s.t.} \quad \mathbb{E}_{\mathbf{s}_{t+1} \sim p(\cdot \mid \mathbf{s}_t, \mathbf{a}_t), \mathbf{a}_t \sim q_{\theta}(\cdot \mid \mathbf{s}_t), \mathbf{s}_1 \sim p(\cdot)} [\sum_t r(\mathbf{s}_t, \mathbf{a}_t)] = R^*.$$

This objective is a *minimum relative entropy problem* [48], which yields the same Gibbs posterior as eREPS (Lemma 1, Appendix A). With exact inference, a suitable prior and oracle knowledge of the maximum return, this program computes the optimal policy in a single step by setting R^* to the optimal value. However, in this work, the expectation constraint requires self-normalized importance sampling (SNIS) on sampled returns $R^{(n)}$ using samples from the current policy prior,

$$\mathbb{E}_{\mathbf{s}_{t+1} \sim p(\cdot \mid \mathbf{s}_t, \mathbf{a}_t), \mathbf{a}_t \sim q_{\theta}(\cdot \mid \mathbf{s}_t), \mathbf{s}_1 \sim p(\cdot)} [\sum_t r(\mathbf{s}_t, \mathbf{a}_t)] \approx \sum_n w_{q/p}^{(n)} R^{(n)} = \frac{\sum_n R^{(n)} \exp(\alpha R^{(n)})}{\sum_n \exp(\alpha R^{(n)})} = R^*.$$

Rather than specifying R^* here, we identify that this estimator is fundamentally limited by inference accuracy. We capture this error by applying an IS-derived concentration inequality to this estimate (Theorem 1) [49]. This lower bound can be used as an objective for optimizing α , balancing policy improvement with approximate inference accuracy.

Theorem 1. (*Importance sampling estimator concentration inequality (Theorem 2, [49])*) Let q and p be two probability densities such that $q \ll p$ and $d_2[q \parallel p] < +\infty$. Let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ i.i.d. random variables sampled from p and $f : \mathcal{X} \rightarrow \mathbb{R}$ be a bounded function ($\|f\|_{\infty} < +\infty$). Then, for any $0 < \delta \leq 1$ and $N > 0$ with probability at least $1 - \delta$:

$$\mathbb{E}_{\mathbf{x} \sim q(\cdot)} [f(\mathbf{x})] \geq \frac{1}{N} \sum_{i=1}^N w_{q/p}(\mathbf{x}_i) f(\mathbf{x}_i) - \|f\|_{\infty} \sqrt{\frac{(1 - \delta) d_2[q(\mathbf{x}) \parallel p(\mathbf{x})]}{\delta N}}. \quad (4)$$

The divergence term $d_2[q \parallel p]$ is the exponentiated Rényi-2 divergence, $\exp \mathbb{D}_2[q \parallel p]$. While this is tractable for the multivariate Gaussian, it is otherwise not available in closed form. Fortunately, we can use the effective sample size (ESS) [50] as an approximation, as $\hat{N}_{\alpha} \approx N / d_2[q_{\alpha} \parallel p]$ [49, 51] (Lemma 2, see Section A of the Appendix). Combining Equation 4 with our constraint, instead of setting R^* , we maximize the IS lower bound R_{LB}^* to form an objective for the inverse temperature α which incorporates the inference accuracy due to the sampling given inequality probability $1 - \delta$,

$$\max_{\alpha} R_{\text{LB}}^*(\alpha, \delta) = \mathbb{E}_{q_{\alpha}/p}[R] - \mathcal{E}_R(\delta, \hat{N}_{\alpha}), \quad \mathcal{E}_R(\delta, \hat{N}_{\alpha}) = \|R\|_{\infty} \sqrt{\frac{(1 - \delta)}{\delta}} \frac{1}{\sqrt{\hat{N}_{\alpha}}}. \quad (5)$$

We refer to this approach as *lower-bound policy search* (LBPS). This objective combines the expected performance of q_{α} , based on the IS estimate $\mathbb{E}_{q_{\alpha}/p}[\cdot]$, with regularization \mathcal{E}_R based on the return and inference accuracy. Treating p , N , $\|R\|_{\infty}$ as task-specific hyperparameters, the only algorithm hyperparameter $\delta \in [0, 1)$ defines the probability of the bound. In practice, self-normalized importance sampling is used for PPI, as the normalizing constants of the Gibbs likelihoods are not available. While Metelli et al. also derive an SNIS lower bound [49], we found, as they did, that the IS lower bound with SNIS estimates work better in practice due to the conservatism of the SNIS bound. An interpretation of this approach is that the Rényi-2 regularization constrains the Gibbs posterior to be one that can be estimated from the finite samples, as the divergence is used in evaluating IS sample complexity [52, 53]. Moreover, the role of the ESS for regularization is similar to the ‘elite’ samples in CEM. Connecting these two mechanisms as robust maximum estimators (Section A), we also propose *effective sample size policy search* (ESSPS), which optimizes α to achieve a desired ESS N^* , i.e. a Rényi-2 divergence bound, using the objective $\min_{\alpha} |\hat{N}_{\alpha} - N^*|$. More details regarding PPI (Section A) and temperature selection methods (Table 1) are in the Appendix.

This section introduces two methods, LBPS and ESSPS, for constraining the Gibbs posteriors for Monte Carlo optimization. These methods provide statistical regularization through soft and hard constraints involving the effective sample size, which avoids the pitfall of fitting high-dimensional distributions to a few effective samples. A popular setting for these methods is MPC, which performs episodic optimization over short planning horizons while adapting each time step to the current state. Moreover, for optimal control, we also need to specify a suitable prior over action sequences. To apply PPI to the MPC setting, we must implement online optimization given this prior over actions.




Figure 2: A practical aspect of Monte Carlo control methods for robotics is optimizing smooth action sequence. This example shows a non-smooth optimal sequence $--$, which may be undesirable, though optimal, to fit exactly. Prior methods struggle at providing both effective smooth solutions in the mean $—$ and action samples $—$, as they ultimately fit the action distribution in an independent fashion. Using kernel-derived covariance function provides both. The line $---$ denotes the optimization horizon, beyond which is exploratory actions derived from both the posterior and prior.

4 Online Posterior Policy Iteration & Prior Design

In this section, we derive the online realization of posterior policy iteration that uses and maintains correlated action priors, computing the finite-horizon H future actions given a likelihood on a subset of actions from the past. R represents the return-based Gibbs likelihood term (Definition 1),

$$q_\alpha(\mathbf{a}_{t:t+H} \mid R_{1:\tau}) = \int q_\alpha(\mathbf{a}_{1:t+H} \mid R_{1:\tau}) d\mathbf{a}_{1:t-1} \propto \int p(R_{1:\tau} \mid \mathbf{a}_{1:\tau}) p(\mathbf{a}_{1:t+H}) d\mathbf{a}_{1:t-1}, \quad (6)$$

where $\tau \leq t + H$. As an analogy, this is equivalent to combining forecasting with state estimation, i.e. $p(\mathbf{x}_{t:t+H} \mid \mathbf{y}_{1:t})$ for states \mathbf{x} and measurements \mathbf{y} . For correlated priors on the action space, this computation is tractable if working with Gaussian processes. In fact, a recurring aspect across several many posterior policy iteration-like approaches is the use of Gaussian process policies,

$$p(\mathbf{A} \mid \mathbf{S}) = \begin{cases} \prod_t \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t), & \text{(Independent Gaussian noise, e.g. [18]),} \\ \mathcal{N}(\boldsymbol{\mu}_w^\top \phi(t), \phi(t)^\top \boldsymbol{\Sigma}_w \phi(t)), & \text{(Bayesian linear regression, e.g. PROMP [32]),} \\ \prod_t \mathcal{N}(\mathbf{k}_t + \mathbf{K}_t \mathbf{s}, \boldsymbol{\Sigma}_t), & \text{(time-varying linear Gaussian e.g. [41, 54, 42]),} \\ \mathcal{GP}(\boldsymbol{\mu}(\mathbf{s}), \boldsymbol{\Sigma}(\mathbf{s})), & \text{(non-parametric Gaussian process [55]).} \end{cases}$$

Despite the simplicity of Gaussian action noise, for robotics, more sophisticated noise is often desired for safety and effective exploration [56, 29]. Prior work has proposed first-order smoothing [11, 12]. Using $\mathbf{v}_t^{(n)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\beta \in [0, 1]$ and $\boldsymbol{\Sigma}_t = \mathbf{L}_t \mathbf{L}_t^\top$, actions are sampled using

$$\mathbf{a}_t^{(n)} = \boldsymbol{\mu}_t + \mathbf{L}_t \mathbf{n}_t^{(n)}, \quad \mathbf{n}_t^{(n)} = \beta \mathbf{v}_t^{(n)} + (1 - \beta) \mathbf{n}_{t-1}^{(n)}, \quad \text{or} \quad \mathbf{n}_t^{(n)} = \beta \mathbf{v}_t^{(n)} + \sqrt{(1 - \beta^2)} \mathbf{n}_{t-1}^{(n)}.$$

However, in practice it is also implemented as $\mathbf{a}_t^{(n)} = \beta(\boldsymbol{\mu}_t + \mathbf{L}_t \mathbf{v}_t^{(n)}) + (1 - \beta) \mathbf{a}_{t-1}^{(n)}$ ¹. While this approach directly smooths the actuation, it also introduces a lag, which may deteriorate performance. Other approaches have used colored noise for sampling the noise \mathbf{n} [4]. Contrast these approaches to Gibbs sampling a multivariate Gaussian joint distribution with 1-step cross-correlations [58], which is $\mathbf{a}_{t|t-1}^{(n)} = \boldsymbol{\mu}_{t|t-1}^{(n)} + \mathbf{L}_{t|t-1} \mathbf{v}_t^{(n)}$, where $\boldsymbol{\mu}_{t|t-1}^{(n)} = \boldsymbol{\mu}_t + \boldsymbol{\Sigma}_{t,t-1} \boldsymbol{\Sigma}_{t-1}^{-1} (\mathbf{a}_{t-1}^{(n)} - \boldsymbol{\mu}_{t-1})$, and $\boldsymbol{\Sigma}_{t|t-1} = \boldsymbol{\Sigma}_t - \boldsymbol{\Sigma}_{t,t-1} \boldsymbol{\Sigma}_{t-1}^{-1} \boldsymbol{\Sigma}_{t-1}^\top$. The differences are subtle, but important. The initial proposed sampling scheme essentially adds correlated noise to the mean for exploration, but does not consider the smoothness of the mean itself. The practical implementation incorporates the previous action, but through exponential smoothing, which introduces a fixed lag that potentially degrades the quality of the mean action sequence. Correct sampling of the joint distribution has neither of these issues and naturally extends to correlations over several time steps. We do this in a general fashion by considering the (continuous time) Gaussian process (see Section G, Appendix), so $p(\mathbf{a}_t) = \mathcal{N}(\boldsymbol{\mu}_{t_i:t_j}, \boldsymbol{\Sigma}_{t_i:t_j}) = \mathcal{GP}(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t))$ for a discrete-time sequence $\mathbf{t} = [t_i, \dots, t_j]$. Proposition 1 shows how the time shift for MPC can be implemented in a general fashion when using GPs.

Proposition 1. *Given a Gaussian process prior $\mathcal{GP}(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t))$ and multivariate normal posterior $q_\alpha(\mathbf{a}_{t_1:t_2}) = \mathcal{N}(\boldsymbol{\mu}_{t_1:t_2|R}, \boldsymbol{\Sigma}_{t_1:t_2|R})$ for t_1 to t_2 , the posterior for t_3 to t_4 is expressed as*

$$\boldsymbol{\mu}_{t_3:t_4|R} = \boldsymbol{\mu}_{t_3:t_4} + \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2} \boldsymbol{\nu}_{t_1:t_2}, \quad \boldsymbol{\Sigma}_{t_3:t_4|R} = \boldsymbol{\Sigma}_{t_3:t_4} - \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2} \boldsymbol{\Lambda}_{t_1:t_2}^\top \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2}^{-1}, \quad (7)$$

where $\boldsymbol{\nu}_{t_1:t_2} = \boldsymbol{\Sigma}_{t_1:t_2}^{-1} (\boldsymbol{\mu}_{t_1:t_2|R} - \boldsymbol{\mu}_{t_1:t_2})$ and $\boldsymbol{\Lambda}_{t_1:t_2} = \boldsymbol{\Sigma}_{t_1:t_2}^{-1} (\boldsymbol{\Sigma}_{t_1:t_2} - \boldsymbol{\Sigma}_{t_1:t_2|R}) \boldsymbol{\Sigma}_{t_1:t_2}^{-1}$.

This update combines the new sequence prior from t_3 to t_4 and the previous likelihood used in the update for t_1 to t_2 , obtained from the posterior and prior. Note, the cross-covariance $\boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2}$ is computed using the covariance function of the prior GP. The proof is in Appendix A.

¹See the source code for Nagabandi et al. [11] and MBRL-lib [57].

For a stationary kernel, fixed planning horizon and fixed control frequency, the term $\Sigma_{t_1:t_2}^{-1}$ is $\Sigma_{t:t+H}^{-1}$ and is constant, so can be computed at initialization to avoid repeated inversion. Figure 2 demonstrates how this update lets us combine our prior with previous posterior in a principled fashion. Moreover, its continuous-time construction means that the time resolution can be updated, not just the time window, for planning at different timescales [30].

Compared to the independence assumption, modeling correlations between actions introduces complexity. The full covariance over (flattened) time and action has a complexity $R(T^3 d_a^3)$, which is infeasible to work with. Assuming independence between actions, a GP per action has a complexity of $R(T^3 d_a)$, requiring d_a GPs to be fit, which is not desirable for online methods such as MPC. To avoid the linear scaling w.r.t. d_a , we propose using the *matrix Normal distribution* (Definition 4) for scalability, as it is parameterized into single T and d_a -dimensional covariances,

Definition 4. (*Matrix Normal Distribution (MaVN)* [59]) For a random matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, it follows a matrix normal distribution $\mathbf{X} \sim \mathcal{MN}(\mathbf{M}, \mathbf{K}, \mathbf{\Sigma})$, where $\mathbf{M} \in \mathbb{R}^{n \times p}$, $\mathbf{K} \in S_+^n$ and $\mathbf{\Sigma} \in S_+^p$, if and only if $\text{vec}(\mathbf{X}) \sim \mathcal{N}(\text{vec}(\mathbf{M}), \mathbf{\Sigma} \otimes \mathbf{K})$, where \otimes denotes the Kronecker product.

Using the Kronecker-structured covariance provides a useful decomposition of the time-based covariance \mathbf{K} , that defines correlations between time steps, and an action covariance $\mathbf{\Sigma}$ that captures correlations between actions. Typically we assume actions are independent, but cross-correlations could be learned from experience for richer coordination. While this Kronecker structure does not fully capture the correlations between time and actions, the structure is very useful for MPC on robotic systems, where the actions space could be very high but the planning horizon is sufficiently small for covariance estimation using a reasonable number of Monte Carlo rollouts.

Feature Approximations. Despite the matrix Normal factorization, computing the correlations between actions still requires a dense $H \times H$ covariance matrix \mathbf{K} for planning horizon H . To sparsify this quantity, we consider kernel approximations, such as the canonical basis functions $\sum_n k(\cdot, \mathbf{x}_n)$ and *spectral* approximations using random features $\sum_n \phi_n(\cdot)$ [60], for a Bayesian linear model $\phi_t^\top \mathbf{W}$. Focusing on the squared exponential (SE) kernel, this results in radial basis function (RBF) and random Fourier features (RFF) respectively. Interestingly, RBF features are closely related to probabilistic movement primitives, used extensively in policy search for robotics [32]. For one-dimensional inputs, RFFs are effectively approximated by applying Gauss-Hermite quadrature [61] to the random weights [62]. RBF features and RFFs approximate w.r.t. time and frequency respectively and could be combined [63]. Using these continuous-time features, the optimization is now abstracted from planning horizon and control frequency, providing much greater flexibility. Secondly, due to the features, a factorized weight covariance approximation does not sacrifice smoothness. Moreover, the moment updates described above are not needed, as only ϕ_t is updated.

5 Related Work

Inference-based control. Posterior policy iteration was proposed by Rawlik et al. [41] and covers prior methods developed from Bayesian smoothing [23, 37], expectation maximization [22, 56], entropy regularization [47, 9] and path integral [64] perspectives. For MPC specifically, the path integral-based MPPI was proposed [18], with alternative formulations based on mirror descent [46] and variational inference [5, 65, 66]. Mukadam et al. [25] models the optimal state-action distribution as a sparse Gaussian process and uses linearization for approximate inference. The same approach is used for Gaussian process motion planning [30], which are also optimized using sampling [8]. Gaussian quadrature is also used for inference-based MPC [27]. Concurrent work uses the ESS for a temperature adjusting heuristic for MPPI [15] and also combines policy search with MPC using PPI techniques [67]. See Section B for a more in-depth discussion on these related works.

Policy design and regularization. Smooth actuation is important in robot learning for safety and exploration, having been proposed for Monte Carlo MPC [11, 12, 4] and more broadly incorporated using augmented objectives, parameter sampling and policy design, e.g. [68, 69, 70].

Stochastic search. Probabilistic interpretations of black-box optimization algorithms are well established [71, 72, 73], however prior work did not connect the ESS and elite samples. CEM and extensions have also been adopted widely as a solver for MPC [3, 4, 6].

Gaussian processes for control. This work adopts GPs for correlated action priors. This is distinct from prior work which uses GPs to approximate dynamics or value functions, e.g. [74, 75, 76, 77].

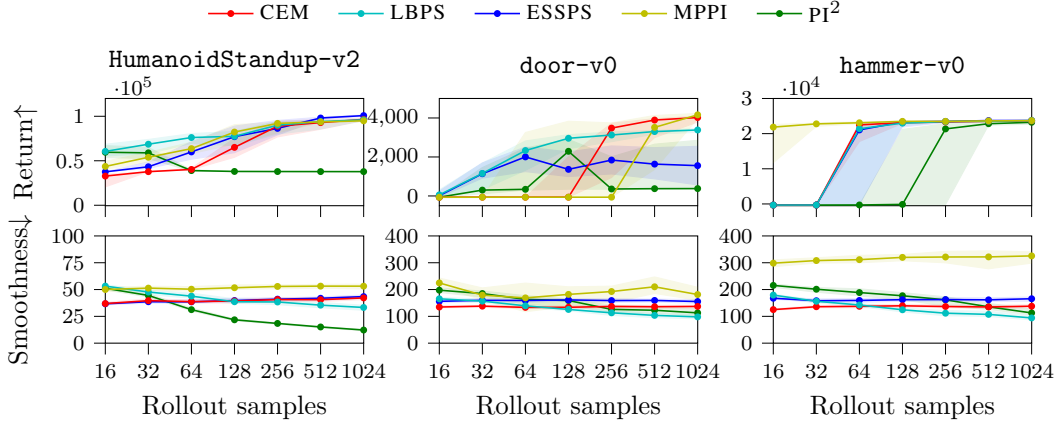


Figure 3: MPC return and smoothness with white noise priors. Displaying quartiles over 50 seeds. These priors require a large number of samples for good performance, across methods and tasks.

6 Experimental Results

We assess the Gibbs posterior methods and policy design empirically across various settings. Black-box optimization (Section 6.1) considers standard benchmarks, while policy search (Section 6.2) optimizes action sequences for a robotic task. For MPC (Section 6.3), we evaluate online PPI approaches with white noise and smooth priors on high-dimensional, contact-rich tasks. For the code, see github.com/JoelWatson/monte-carlo-posterior-policy-iteration.

6.1 Black-box Optimization

To understand the behaviour of the proposed PPI variants, the performance of LBPS and ESSPS on a range of standard black-box optimization functions over a range of hyperparameters, with eREPS and CEM as baselines, are shown in Appendix D.1. Figures 7 – 10 show that ESS is a useful metric for these methods, as each solver exhibits consistent ESS for a given hyperparameter value. However, the uniform weights used by CEM (Figure 7) maintain entropy longer than ESSPS (Figure 8), which can lead to better optima, so the ESS is not sufficient to fully capture the behavior of these solvers.

6.2 Policy Search

As LBPS and ESSPS are closely related to eREPS, we repeat the experiment from prior work performing the ‘ball in a cup’ task using policy search using a Barret WAM [14], which has been shown to transfer to the physical system [78, 14, 79], as a benchmark task. Moreover, we replace PROMPs with Matrix normal RBF and RFF policies. From the kernel perspective, this feature approximation is motivated by the large ($T \simeq 1000$) task horizon. The results in Appendix D.2 confirm that these solvers are all capable of solving the task, based on success rate, where RBF (Figure 11) and RFF features (Figure 12) perform equally well w.r.t. the convergence of the success rate for each approach.

6.3 Model Predictive Control with Oracle Dynamics

We evaluate online PPI across a range of high-dimensional robotic control tasks in MuJoCo [80], including HumanoidStandup-v2 in Gym [81] and door-v0, hammer-v0 from `mj_envs`, using the Adroit hand (Figure 1) [82]. To measure smoothness, we adopt the FFT-based score $\frac{2}{Nf_s} \sum_{i=1}^N a_i f_i$ [68], with sampling frequency f_s and N resolvable frequencies f with amplitudes a . We compute the Euclidean norm of the action sequence over time and apply the smoothness measure to this signal. For the evaluation, we focus on a low computational budget, with 1 or 2 iterations per timestep. To assess the impact of approximate inference, we assess performance over an logarithmic range of sample rollouts, following prior work [4]. Details may be found in Appendix E.2.

White noise priors. Figure 3 shows MPC with white noise priors using LBPS and ESSPS, with MPPI, CEM and PI^2 baselines (see Table 1). While each solver performs comparably for 1024 rollouts, the low sample regime shows greater performance variance. While MPPI seems particularly effective, Figure 13 shows that its average ESS is particularly low, $\simeq 1$ for many cases. Combined with the fixed variances, this suggests optimization is closer to greedy random search than importance sampling. The poor door-v0 performance of ESSPS is due to slow opening, rather than task failure.

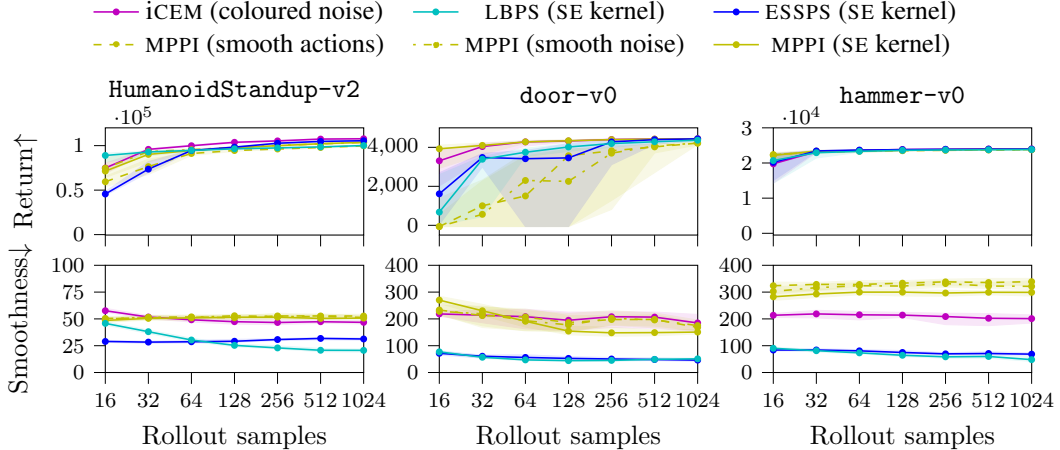


Figure 4: MPC return and **smoothness** with smooth action priors. Displaying quartiles over 50 seeds. Compared to white noise priors, smooth action priors improve sample efficiency dramatically, but only PPI methods (LBPS, ESSPS) preserve this smoothness in the downstream control.

Policy design for smooth control. Figure 4 shows online PPI with action priors. LBPS, ESSPS and MPPI use the SE kernel, with iCEM [4] and MPPI with smooth actions and noise as a baseline. The smooth action distributions greatly improves performance across models, tasks and sample sizes, due to effective exploration. As desired, the richer SE kernel provides much greater smoothness, by up to a factor of 2 compared to baselines, with limited impact to performance. It is unsurprising that the smoothness bias reduces performance if optimal behavior is non-smooth, as illustrated in Figure 2. Appendix D.3.2 shows some of the action sequences from Figure 4, where we see GP smoothness varies with increasing rollout samples and also results in significant actuator amplitude reduction.

Comparing kernel- and feature-based policies. To assess feature approximations for smooth MPC, we replace the SE kernel with RBF and RFF features, while keeping the lengthscale fixed. These policies perform worse given fewer samples, but are comparable to the true kernel with sufficient samples (Figure 15). We attribute this to the compounding errors of kernel approximation and fewer effective samples. In contrast to the policy search task, RFFs appear superior to RBF features.

Learning priors from data. A benefit of using GP priors is the ability to optimize hyperparameters from expert demonstrations through the likelihood or a divergence. Moreover, the matrix normal distribution is useful for analyzing high-dimensional action sequences, as it decomposes temporal- and action correlations into viewable covariance matrices. Section D.3.3 shows the matrix normal distributions of expert demonstrations of the tasks, obtained through human and RL experts. The results show that, surprisingly, the demonstrations are rougher than the smoothness achievable with MPC. We attribute this to control artifacts from demonstration collection and the use of Gaussian noise by RL agents. Applying the same methodology to the demonstrations of the smooth MPC agents proposed here extracts the expected action correlations across tasks. This analysis also raises the question of whether smoothness is an inductive bias we enforce for practicality, or a phenomena we expect to arise from optimality. If the latter, it may be that the simulated environments or objectives considered are lacking components that encourage smoothness, such as energy efficiency.

7 Conclusion

We present a broad perspective on episodic posterior policy iteration method for robotics and new methods for the Monte Carlo setting, based on regularizing the IS approximations. By considering vector-valued Gaussian processes for action priors, we have demonstrated how sample-efficient MPC can be performed as online inference and with greater control over actuator smoothness, connecting Monte Carlo MPC to prior work on policy search. This approach was validated on a set of high-dimensional MPC tasks closely matching baseline performance while achieving greater smoothness.

Limitations. Much of the prior work is motivated by simplicity, minimizing hyperparameter tuning and numerical procedures such as matrix inversion [64]. In contrast, the contributions of this work introduces complexity, i.e. online temperature optimization and the use of dense covariance matrices in order to perform more sophisticated approximate inference. While this additional complexity has an impact on execution time (Table 2, Appendix), we hope the sample-efficiency when combined with accelerations such as GPU-integration should produce real-time algorithms [13].

Acknowledgments

This work built on prior codebases developed by Johannes Silberbauer, Michael Lutter and Hany Abdulsamad. The large-scale experiments and ablations were conducted on the Lichtenberg high performance computer of the TU Darmstadt. The authors wish to thank Hany Abdulsamad, Boris Belousov, Michael Lutter, Fabio Muratore, Pascal Klink, Georgia Chalvatzaki, Kay Hansel, Oleg Arenz and the anonymous conference reviewers for helpful feedback on earlier drafts.

References

- [1] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, 2018.
- [2] A. B. Owen. *Monte Carlo theory, methods and examples*. 2013.
- [3] K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, 2018.
- [4] 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 *Conference on Robot Learning*, 2020.
- [5] M. Okada and T. Taniguchi. Variational inference MPC for Bayesian model-based reinforcement learning. In *Conference on Robot Learning*, 2019.
- [6] M. Lutter, L. Hasenclever, A. Byravan, G. Dulac-Arnold, P. Trochim, N. Heess, J. Merel, and Y. Tassa. Learning dynamics models for model predictive agents. *arXiv preprint arXiv:2109.14311*, 2021.
- [7] F. Stulp, E. Theodorou, J. Buchli, and S. Schaal. Learning to grasp under uncertainty. In *2011 IEEE International Conference on Robotics and Automation*, 2011.
- [8] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal. STOMP: Stochastic trajectory optimization for motion planning. In *2011 IEEE International Conference on Robotics and Automation*, 2011.
- [9] C. Daniel, G. Neumann, O. Kroemer, J. Peters, et al. Hierarchical relative entropy policy search. *Journal of Machine Learning Research*, 2016.
- [10] 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*, 2018.
- [11] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar. Deep dynamics models for learning dexterous manipulation. In *Conference on Robot Learning*, 2019.
- [12] Y. Yang, K. Caluwaerts, A. Iscen, T. Zhang, J. Tan, and V. Sindhwani. Data-efficient reinforcement learning for legged robots. In *Conference on Robot Learning*, 2019.
- [13] 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*, 2021.
- [14] M. Lutter, J. Silberbauer, J. Watson, and J. Peters. Differentiable physics models for real-world offline model-based reinforcement learning. In *IEEE International Conference on Robotics and Automation*, 2021.
- [15] J. Carius, R. Ranftl, F. Farshidian, and M. Hutter. Constrained stochastic optimal control with learned importance sampling: A path integral approach. *The International Journal of Robotics Research*, 2022.
- [16] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber. Natural evolution strategies. *Journal of Machine Learning Research*, 2014.

- [17] A. Abdolmaleki, R. Lioutikov, J. R. Peters, N. Lau, L. Pualo Reis, and G. Neumann. Model-based relative entropy stochastic search. In *Advances in Neural Information Processing Systems*, 2015.
- [18] G. Williams, A. Aldrich, and E. A. Theodorou. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 2017.
- [19] R. Y. Rubinstein and D. P. Kroese. *The Cross Entropy Method: A Unified Approach To Combinatorial Optimization, Monte Carlo Simulation*. 2004.
- [20] P. Dayan and G. E. Hinton. Using expectation-maximization for reinforcement learning. *Neural Computation*, 1997.
- [21] M. Toussaint and A. Storkey. Probabilistic inference for solving discrete and continuous state Markov decision processes. In *International Conference of Machine Learning*, 2006.
- [22] J. Peters and S. Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *International Conference on Machine Learning*, 2007.
- [23] M. Toussaint. Robot trajectory optimization using approximate inference. In *International Conference on Machine Learning*, 2009.
- [24] H. J. Kappen, V. Gómez, and M. Opper. Optimal control as a graphical model inference problem. In *International Conference on Automated Planning and Scheduling*, 2013.
- [25] M. Mukadam, C. Cheng, X. Yan, and B. Boots. Approximately optimal continuous-time motion planning and control via probabilistic inference. In *IEEE International Conference on Robotics and Automation*, 2017.
- [26] S. Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- [27] J. Watson, H. Abdulsamad, R. Findeisen, and J. Peters. Efficient stochastic optimal control through approximate Bayesian input inference. *arXiv preprint arXiv:2105.07693*, 2021.
- [28] H. J. Kappen and H. C. Ruiz. Adaptive importance sampling for control and inference. *Journal of Statistical Physics*, 2016.
- [29] M. P. Deisenroth, G. Neumann, J. Peters, et al. A survey on policy search for robotics. *Foundations and Trends in Robotics*, 2013.
- [30] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots. Continuous-time Gaussian process motion planning via probabilistic inference. *The International Journal of Robotics Research*, 2018.
- [31] T. Flash and N. Hogan. The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience*, 1985.
- [32] A. Paraschos, C. Daniel, J. Peters, G. Neumann, et al. Probabilistic movement primitives. In *Advances in Neural Information Processing Systems*, 2013.
- [33] C. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [34] E. Theodorou, Y. Tassa, and E. Todorov. Stochastic differential dynamic programming. In *American Control Conference*, 2010.
- [35] A. Zellner. Optimal information processing and Bayes’s theorem. *The American Statistician*, 1988.
- [36] M. E. Khan and H. Rue. The Bayesian learning rule. *arXiv preprint arXiv:2107.04562*, 2021.
- [37] J. Watson, H. Abdulsamad, and J. Peters. Stochastic optimal control as approximate input inference. In *Conference on Robot Learning*, 2019.
- [38] B. Guedj. A primer on PAC-Bayesian learning. In *Proceedings of the second congress of the French Mathematical Society*, 2019.

- [39] P. Alquier, J. Ridgway, and N. Chopin. On the properties of variational approximations of Gibbs posteriors. *Journal of Machine Learning Research*, 2016.
- [40] B. Dai, N. He, H. Dai, and L. Song. Provable bayesian inference via particle mirror descent. In *International Conference on Artificial Intelligence and Statistics*, 2016.
- [41] K. Rawlik, M. Toussaint, and S. Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference. In *Robotics: Science and Systems*, 2012.
- [42] J. Watson and J. Peters. Advancing trajectory optimization with approximate inference: Exploration, covariance control and adaptive risk. In *American Control Conference*, 2021.
- [43] K. C. Rawlik. *On probabilistic inference approaches to stochastic optimal control*. PhD thesis, The University of Edinburgh, 2013.
- [44] C. Wirth, J. Fürnkranz, and G. Neumann. Model-free preference-based reinforcement learning. In *Conference on Artificial Intelligence*, 2016.
- [45] S. Satoh, H. J. Kappen, and M. Saeki. An iterative method for nonlinear stochastic optimal control based on path integrals. *IEEE Transactions on Automatic Control*, 2017.
- [46] N. Wagener, C. an Cheng, J. Sacks, and B. Boots. An online learning approach to model predictive control. In *Robotics: Science and Systems*, 2019.
- [47] J. Peters, K. Mülling, and Y. Altün. Relative entropy policy search. In *Conference on Artificial Intelligence*, 2010.
- [48] J. van Campenhout and T. Cover. Maximum entropy and conditional probability. *IEEE Transactions on Information Theory*, 1981.
- [49] A. M. Metelli, M. Papini, N. Montali, and M. Restelli. Importance sampling techniques for policy optimization. *Journal of Machine Learning Research*, 2020.
- [50] A. Kong. A note on importance sampling using standardized weights. *University of Chicago, Dept. of Statistics, Technical Report*, 1992.
- [51] C. Cortes, Y. Mansour, and M. Mohri. Learning bounds for importance weighting. In *Advances in Neural Information Processing systems*, 2010.
- [52] S. Agapiou, O. Papaspiliopoulos, D. Sanz-Alonso, and A. M. Stuart. Importance sampling: Intrinsic dimension and computational cost. *Statistical Science*, 2017.
- [53] J. Hernández-González and J. Cerquides. A robust solution to variational importance sampling of minimum variance. *Entropy*, 2020.
- [54] V. Gómez, H. J. Kappen, J. Peters, and G. Neumann. Policy search for path integral control. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2014.
- [55] H. van Hoof, G. Neumann, and J. Peters. Non-parametric policy search with limited information loss. *Journal of Machine Learning Research*, 2017.
- [56] J. Kober and J. Peters. Policy search for motor primitives in robotics. In *Advances in Neural Information Processing Systems*, 2009.
- [57] L. Pineda, B. Amos, A. Zhang, N. O. Lambert, and R. Calandra. MBRL-Lib: A modular library for model-based reinforcement learning. *arXiv preprint arXiv:2104.10159*, 2021.
- [58] A. Doucet. A note on efficient conditional simulation of Gaussian distributions. *Departments of Computer Science and Statistics, University of British Columbia*, 2010.
- [59] A. P. Dawid. Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika*, 1981.
- [60] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2007.

- [61] F. B. Hildebrand. *Introduction to numerical analysis*. Courier Corporation, 1987.
- [62] M. Mutny and A. Krause. Efficient high-dimensional Bayesian optimization with additivity and quadrature Fourier features. In *Advances in Neural Information Processing Systems*, 2018.
- [63] J. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *International Conference on Machine Learning*, 2020.
- [64] E. Theodorou, J. Buchli, and S. Schaal. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 2010.
- [65] A. Lambert, A. Fishman, D. Fox, B. Boots, and F. Ramos. Stein variational model predictive control. In *Conference on Robot Learning*, 2020.
- [66] 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.
- [67] Y. Song and D. Scaramuzza. Policy search for model predictive control with application to agile drone flight. *IEEE Transactions on Robotics*, 2022.
- [68] S. Mysore, B. Mabsout, R. Mancuso, and K. Saenko. Regularizing action policies for smooth control with reinforcement learning. In *IEEE International Conference on Robotics and Automation*, 2021.
- [69] A. Raffin, J. Kober, and F. Stulp. Smooth exploration for robotic reinforcement learning. In *Conference on Robot Learning*, 2021.
- [70] D. Korenkevych, A. R. Mahmood, G. Vasan, and J. Bergstra. Autoregressive policies for continuous control deep reinforcement learning. In *International Joint Conference on Artificial Intelligence*, 2019.
- [71] F. Stulp and O. Sigaud. Path integral policy improvement with covariance matrix adaptation. In *International Conference on Machine Learning*, 2012.
- [72] Y. Ollivier, L. Arnold, A. Auger, and N. Hansen. Information-geometric optimization algorithms: A unifying picture via invariance principles. *Journal of Machine Learning Research*, 2017.
- [73] A. Abdolmaleki, B. Price, N. Lau, L. P. Reis, and G. Neumann. Deriving and improving CMA-ES with information geometric trust regions. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2017.
- [74] F. Berkenkamp and A. P. Schoellig. Safe and robust learning control with Gaussian processes. In *European Control Conference*, 2015.
- [75] J. Kocijan, R. Murray-Smith, C. Rasmussen, and A. Girard. Gaussian process model based predictive control. In *American Control Conference*, 2004.
- [76] M. Deisenroth, C. Rasmussen, and J. Peters. Gaussian process dynamic programming. *Neuro-computing*, 2009.
- [77] M. Maiworm, D. Limon, J. Maria Manzano, and R. Findeisen. Stability of Gaussian process learning based output feedback model predictive control. In *IFAC Conference on Nonlinear Model Predictive Control*, 2018.
- [78] P. Klink, H. Abdulsamad, B. Belousov, C. D’Eramo, J. Peters, and J. Pajarinen. A probabilistic interpretation of self-paced learning with applications to reinforcement learning. *Journal of Machine Learning Research*, 2021.
- [79] F. Muratore, C. Eilers, M. Gienger, and J. Peters. Data-efficient domain randomization with Bayesian optimization. *IEEE Robotics and Automation Letters*, 2021.
- [80] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In *IEEE International Conference on Intelligent Robots and Systems*, 2012.

- [81] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. OpenAI Gym, 2016.
- [82] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In *Robotics: Science and Systems*, 2018.
- [83] B. D. Anderson and J. B. Moore. *Optimal filtering*. 2012.
- [84] T. D. Barfoot, C. H. Tong, and S. Särkkä. Batch continuous-time trajectory estimation as exactly sparse Gaussian process regression. In *Robotics: Science and Systems*, 2014.
- [85] O. Cappé, S. J. Godsill, and E. Moulines. An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE*, 2007.
- [86] K. Asadi and M. L. Littman. An alternative softmax operator for reinforcement learning. In *International Conference on Machine Learning*, 2017.
- [87] C. Chu, J. Blanchet, and P. Glynn. Probability functional descent: A unifying perspective on GANs, variational inference, and reinforcement learning. In *International Conference on Machine Learning*, 2019.
- [88] A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 2003.
- [89] 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*, 2017.
- [90] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, Í. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. . . Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 2020.
- [91] P. Dutilleul. The MLE algorithm for the matrix normal distribution. *Journal of Statistical Computation and Simulation*, 1999.
- [92] L. Peng, H. Qian, Z. Shen, C. Zhang, and F. Li. Generative actor-critic: An off-policy algorithm using the push-forward model. *arXiv preprint arXiv:2105.03733*, 2021.
- [93] S. Boyd, S. P. Boyd, and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [94] S. Särkkä and A. Solin. *Applied Stochastic Differential Equations*. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2019.




Figure 5: Visualization of Example 1 for three values of α , showing interpolation between no update and towards the Newton update.

A Extended Discussion Points and Results

This section discusses specific details not covered at length in the main section.

Posterior policy iteration as optimization. For clarity, we can ground inference-based optimization clearly by considering Gaussian inference and quadratic optimization (Example 1).

Example 1. (*Newton method via Gaussian inference*). Consider minimizing a function $f(x)$ from a Gaussian prior $p(x) = \mathcal{N}(\mu_i, \Sigma_i)$. Considering a second-order Taylor expansion about the prior mean, $f(x) \approx f(\mu_i) + \nabla_x f(\mu_i)(x - \mu_i) + \frac{1}{2}(x - \mu_i)^\top \nabla_{xx} f(\mu_i)(x - \mu_i)$, the Gaussian pseudo-posterior $q_\alpha(x) = \mathcal{N}(\mu_{i+1}, \Sigma_{i+1}) \propto \exp(-\alpha f(x)) p(x)$ is

$$\mu_{i+1} = \mu_i - \alpha \Sigma_{i+1} \nabla_x f(\mu_i), \quad \Sigma_{i+1} = (\Sigma_i + \alpha \nabla_{xx}^2 f(\mu_i))^{-1}.$$

As $\alpha \rightarrow 0$, $q_\alpha \rightarrow p$, while as $\alpha \rightarrow \infty$, the mean update tends to the Newton step,

$$\mu_{i+1} \rightarrow \mu_i - \nabla_{xx}^2 f(\mu_i)^{-1} \nabla_x f(\mu_i).$$

Therefore, α acts as a learning rate and regularizer, with repeated inference performing regularized Newton decent. This regularization is beneficial for non-convex optimization, as the Hessian may be negative definite.

Minimum relative entropy problems. To motivate REPS as a form of constrained Gibbs posterior, we reverse the objective and constraint, which yields the minimum relative entropy problem (Lemma 1). The posterior of this problem is the same as REPS and the temperature is also a Lagrangian multiplier, but the constraint is now defined by the expectation value rather than the KL divergence.




Figure 6: Visualization of Example 1 for three values of α , using self-normalized importance sampling. High values of α lead to the posterior collapsing around a few samples. Moreover, compared to the Gaussian case, the update is limited to the support of the samples.

Lemma 1. (Minimum relative entropy [48]) Find q that minimizes the relative entropy from p , given a function \mathbf{f} that describes random variable \mathbf{x} , $\min_q \mathbb{D}_{\text{KL}}[q(\mathbf{x}) \parallel p(\mathbf{x})]$ s.t. $\mathbb{E}_{\mathbf{x} \sim q(\cdot)}[\mathbf{f}(\mathbf{x})] = \hat{\mathbf{f}}$. The solution to this problem is $q(\mathbf{x}) \propto \exp(-\boldsymbol{\lambda}^\top \mathbf{f}(\mathbf{x}))p(\mathbf{x})$, where $\boldsymbol{\lambda}$ is the Lagrange multiplier required to satisfy the constraint. When $p(\mathbf{x})$ is uniform, $q(\mathbf{x})$ is a maximum entropy distribution.

Interestingly, in maximum entropy problems, function \mathbf{f} defines the sufficient statistics of the random variable. For our setting, \mathbf{f} is based on the reward function (episodic return), which indicates the reward objective is used to ‘summarize’ the random policy samples.

Rényi-2 divergence and the effective sample size. The LBPS objective uses the effective sample size as a practical estimator for the exponentiated Rényi-2 divergence (Lemma 2).

Lemma 2. (Exponentiated Rényi-2 divergence estimator [51]) The effective sample size diagnostic (\hat{N}) is an estimate of the exponentiated Rényi-2 divergence for finite samples, as $d_2(p \parallel q)/N \approx \hat{N}^{-1}$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \frac{(\sum_n w_{q/p}(\mathbf{x}_n))^2}{\sum_n w_{q/p}(\mathbf{x}_n)^2} = \int \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x} = \exp \mathbb{D}_2(p \parallel q)^{-1} = d_2(p \parallel q)^{-1}.$$

This result connects divergence-based stochastic optimization methods, such as REPS, to ‘elite’-based stochastic solvers such as CEM, as the number of elites is equal to the effective sample size of the weights, as discussed below.

Time Shifting with Gaussian Process We derive the multivariate normal posterior shifting update introduced in Section 4, that is facilitated by continuous-time Gaussian process priors.

Proposition 1. Given a Gaussian process prior $\mathcal{GP}(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t))$ and multivariate normal posterior $q_\alpha(\mathbf{a}_{t_1:t_2}) = \mathcal{N}(\boldsymbol{\mu}_{t_1:t_2|R}, \boldsymbol{\Sigma}_{t_1:t_2|R})$ for t_1 to t_2 , the posterior for t_3 to t_4 is expressed as

$$\begin{aligned} \boldsymbol{\mu}_{t_3:t_4|R} &= \boldsymbol{\mu}_{t_3:t_4} + \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2} \boldsymbol{\nu}_{t_1:t_2}, \quad \boldsymbol{\Sigma}_{t_3:t_4|R} = \boldsymbol{\Sigma}_{t_3:t_4} - \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2} \boldsymbol{\Lambda}_{t_1:t_2} \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2}^\top, \\ \text{where } \boldsymbol{\nu}_{t_1:t_2} &= \boldsymbol{\Sigma}_{t_1:t_2}^{-1} (\boldsymbol{\mu}_{t_1:t_2|R} - \boldsymbol{\mu}_{t_1:t_2}) \text{ and } \boldsymbol{\Lambda}_{t_1:t_2} = \boldsymbol{\Sigma}_{t_1:t_2}^{-1} (\boldsymbol{\Sigma}_{t_1:t_2} - \boldsymbol{\Sigma}_{t_1:t_2|R}) \boldsymbol{\Sigma}_{t_1:t_2}^{-1}. \end{aligned} \quad (7)$$

Proof. For this update, we require the Gaussian likelihood $\mathcal{N}(\boldsymbol{\mu}_{R|t_1:t_2}, \boldsymbol{\Sigma}_{R|t_1:t_2})$ that achieves the posterior $q_\alpha(\mathbf{a}_{t_1:t_2})$ given the prior $p(\mathbf{a}_{t_1:t_2})$. Using $\mathbf{K}_{t_1:t_2} = \boldsymbol{\Sigma}_{t_1:t_2} \boldsymbol{\Sigma}_{R|t_1:t_2}^{-1}$, we write the Gaussian posterior update in the Kalman filter form

$$\begin{aligned} \boldsymbol{\mu}_{t_1:t_2|R} &= \boldsymbol{\mu}_{t_1:t_2} + \mathbf{K}_{t_1:t_2} (\boldsymbol{\mu}_{R|t_1:t_2} - \boldsymbol{\mu}_{t_1:t_2}), \\ \boldsymbol{\Sigma}_{t_1:t_2|R} &= \boldsymbol{\Sigma}_{t_1:t_2} - \mathbf{K}_{t_1:t_2} \boldsymbol{\Sigma}_{R|t_1:t_2} \mathbf{K}_{t_1:t_2}^\top. \end{aligned}$$

We introduce $\boldsymbol{\nu}$ and $\boldsymbol{\Lambda}$ to capture the unknown terms involving $\mathbf{a}_{R|t_1:t_2}$ w.r.t. $\mathbf{a}_{t_1:t_2}$ and $\mathbf{a}_{t_1:t_2|R}$,

$$\begin{aligned} \boldsymbol{\nu}_{t_1:t_2} &= \boldsymbol{\Sigma}_{R|t_1:t_2}^{-1} (\boldsymbol{\mu}_{R|t_1:t_2} - \boldsymbol{\mu}_{t_1:t_2}) = \boldsymbol{\Sigma}_{t_1:t_2}^{-1} (\boldsymbol{\mu}_{t_1:t_2|R} - \boldsymbol{\mu}_{t_1:t_2}), \\ \boldsymbol{\Lambda}_{t_1:t_2} &= \boldsymbol{\Sigma}_{R|t_1:t_2}^{-1} = \boldsymbol{\Sigma}_{t_1:t_2}^{-1} (\boldsymbol{\Sigma}_{t_1:t_2} - \boldsymbol{\Sigma}_{t_1:t_2|R}) \boldsymbol{\Sigma}_{t_1:t_2}^{-1}. \end{aligned}$$

When extrapolating to the new time sequence, we use the Kalman posterior update again, but with the prior for the new time sequence. This step computes the joint over old and new timesteps, conditions on the objective (i.e. Bayes’ rule), and marginalizes out the old timesteps (see Equation 6), so

$$\begin{aligned} \boldsymbol{\mu}_{t_3:t_4|R} &= \boldsymbol{\mu}_{t_3:t_4} + \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2} \boldsymbol{\nu}_{t_1:t_2}, \\ \boldsymbol{\Sigma}_{t_3:t_4|R} &= \boldsymbol{\Sigma}_{t_3:t_4} - \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2} \boldsymbol{\Lambda}_{t_1:t_2} \boldsymbol{\Sigma}_{t_3:t_4, t_1:t_2}^\top. \end{aligned}$$

□

Using the geometric interpretation of the Kalman filter [83], this update can be seen as projecting the solution from the previous sequence into the new time sequence, given the correlation structure defined by the GP prior.

Augmented Cost Design versus Prior Design. Prior work enforces smoothness using an augmented reward function term, e.g. $-\lambda_s \|\mathbf{a}_t - \mathbf{a}_{t+1}\|^2$ per timestep [68]. For PPI, the PPI objective also augments the reward objective with a KL term (Equation 3). For the Monte Carlo setting here, this KL term is $\frac{1}{\alpha} \sum_n w^{(n)} (\log w^{(n)} - \log p(\mathbf{A}^{(n)}))$, where $w^{(n)} \propto \exp(\alpha R^{(n)})$. When $p(\mathbf{A})$ is a

Gaussian process, the log probability has the familiar quadratic form, but this time applies to the whole sequence \mathbf{A} , due to the episodic setting. Independent white noise only applies a quadratic penalty per action, due to its factorized covariance. A first-order Markovian GP would regularize one-step correlations, as its inverse covariance is banded [84]. The squared exponential kernel, used in this work, has infinite order [33] and a dense inverse covariance matrix and therefore regularizes the whole sequence, which is why it was chosen to achieve smoothness.

Episodic or sequential inference? For the Monte Carlo setting, sequential Monte Carlo (SMC) [85] is an obvious choice for inference given its connection to Bayesian smoothing methods. However, we found that an SMC approach provided limited benefits for control. The variance reduction when resampling decreases exploration from the optimization viewpoint.

Moreover, SMC smoothing reweighs the particles sampled in the forward pass, with $\mathcal{O}(N^2T)$ complexity for N particles [85]. Therefore, no exploration occurs during the backward pass, in contrast to Gaussian message passing where smoothing performs Riccati equation updates on the state-action distribution. When considering just sequential importance sampling for trajectories, we arrive at the episodic case, as $w_T^{(n)} = \exp\left(\alpha r_T^{(n)}\right) w_{T-1}^{(n)} = \exp\left(\alpha \sum_{\tau=0}^T r_\tau^{(n)}\right)$.

Connection to rare event simulation and maximum estimators. The cross entropy method is a popular and effective Monte Carlo optimizer [19]. It broadly works by moment matching an exponential family distribution onto ‘elite’ samples. The elite samples are chosen according to $\mathbb{P}_{q(\mathbf{x})}(f(\mathbf{x}) > f^*)$, a rare event simulator, where improving upon f^* is treated as the rare event. In practice, elites are chosen by sorting the top k samples according to the objective. Our lower-bound reflects CEM in two ways. One, the rare event inequality reflects the expected return bound used in the LBPS objective. Secondly, the k in the top- k estimator matches the effective sample size, used in LBPS and ESSPS. One can view CEM and IS approaches through their maximum estimators

$$\text{Max}_k[\mathbf{R}] = \frac{\sum_{n=1}^k R_n}{\sum_{n=1}^k 1}, \quad \text{Max}_\alpha[\mathbf{R}] = \frac{\sum_{n=1}^N \exp(\alpha R_n) R_n}{\sum_{n=1}^N \exp(\alpha R_n)}.$$

We can see that the SNIS expectation is equivalent to the Boltzmann softmax, used in reinforcement learning [86]. Using k or α this estimator is bounded between the true maximum over the samples and the mean. Therefore the estimators behave in a similar fashion when α is chosen such that the ESS is k . The key distinction is the sparse and uniform weights of CEM vs. the posterior weights of PPI methods, which seems to have an equal or more important influence on optimization, based on the black-box optimization results.

Connection to mirror descent. Functional mirror descent for online learning optimizes a distribution q and introduces a Bregman divergence \mathbb{D} for information-geometric regulation, i.e. $q_{i+1} = \arg \min_{q \in \mathcal{Q}} \mathbb{E}_{\mathbf{x} \sim q(\cdot)}[f(\mathbf{x})] + \alpha_i \mathbb{D}[q \parallel q_i]$ [87]. The temperature typically follows a pre-defined schedule that comes with convergence guarantees, e.g. $\alpha_i \propto \sqrt{i}$, such that $\alpha \rightarrow \infty$ as $i \rightarrow \infty$ [88]. Considering Equation 3, the DMD view applies to our setting when using KL regularization, so DMD view opens up alternative Bregman divergences for regularization.

B Extended Related Work

The methods presented in the main section have been investigated extensively in the prior literature. This section provides a more in-depth discussion of the differences in approach and implementation.

Temperature tuning approaches. Across control-as-inference methods, the temperature value plays an important role. By the construction of the policy updates, mis-specification of the temperature leads overly greedy or conservative optimization. Table 1 provides a summary. A popular and often effective approach is to tune a fixed temperature, as done in AICO [23] and MPPI [18]. However, the optimization behavior depends the values of the objective, and nonlinear optimization methods such as Levenberg–Marquardt and mirror descent suggest an adaptive regularization scheme may perform better. Reward-weighted regression (RWR) [22], and later input inference for control (i2C) [27], take a completely probabilistic view of the Gibbs likelihood and used a closed-form update to optimize the temperature using expectation maximization. However, this update is motivated by the probabilistic interpretation of the optimization, not the optimization itself, so there is no guarantee this approach improves optimization, beyond the attractive closed-form update. Follow-up work regulated optimization by normalizing the returns, using the range or standard deviation [64, 56].

Table 1: A review of different adaptive temperature strategies across episodic control-as-inference methods and settings. We propose LBPS and ESSPS as optimization-driven methods with intuitive hyperparameters.

Algorithm	Method	Description
MPPI [18], AICO [23]	Constant	α
PI ² [64], POWER [56]	Normalization	$\bar{\alpha}/(\max[\mathbf{R}] - \min[\mathbf{R}]), \bar{\alpha}/\sigma_R$
RWR [22], I2C [37]	EM	$\alpha_{i+1} = \sum_n \exp(\alpha_i R_i^{(n)}) / \sum_n R_i^{(n)} \exp(\alpha_i R_i^{(n)})$
REPS [47], MORE [17]	KL bound	$\arg \min_{\alpha} \epsilon/\alpha + \frac{1}{\alpha} \log \sum_n \exp(\alpha R^{(n)})$
LBPS (this work)	IS lower-bound	$\arg \max_{\alpha} \mathbb{E}_{q_{\alpha}}[R] - \mathcal{E}_R(\delta, \hat{N}_{\alpha})$
ESSPS (this work)	ESS	$\arg \min_{\alpha} \hat{N}_{\alpha} - N^* $

While this resolved the objective dependency, it is a heuristic and not interpretable from the optimization perspective. Later, REPS framed the update as a constrained optimization problem, subject to a KL constraint [47, 9]. The temperature was then obtained by minimizing the Lagrangian dual function. However, in practice this dual is approximated using Monte Carlo integration which limits the constraint accuracy. Moreover, a hard KL constraint now requires a distribution-dependent hyperparameter. This work introduces LBPS and ESSPS. LBPS optimizes a lower-bound of the importance-sampled expected return, with the hyperparameter controlling greediness via the confidence of the concentration inequality and therefore is objective- and distribution independent. ESSPS bridges LBPS and CEM, optimizing for a desired effective sample size as an analogy to elite samples. Finally, self-paced contextual episodic policy search has previously adopted the minimum KL form of the REPS optimization problem, in order to assess the expected performance per context task [78].

Gaussian message passing methods. Using Bayesian smoothing of the state-action distribution, trajectory optimization can be performed by treating a step-based objective analogously to the observation log likelihood in state estimation [23, 27]. While the control prior defined as temporally independent in prior work, the smoothing of the state-action distribution imbues an inherent smoothness to the solution (e.g. Figure 6, [27]). Approximate inference can be performed using linearization or quadrature, which, while effective are less amenable to parallelization and therefore do not scale so gracefully to high-dimensional state and action spaces [27]. Moreover, these methods require access to the objective function for linearization, rather than just rollout evaluations.

Mirror descent model predictive control. Mirror descent, used for proximal optimization, introducing information geometric regularization through a chosen Bregman divergence \mathbb{D}_{Ψ} . *Dynamic* mirror descent (DMD), motivates an autonomous update Φ to the optimization variable to improve performance for online learning. DMD-MPC [46] incorporates this into an MPC scheme by incorporating an explicit time shift operator Φ , so $\theta_{t+1} = \Phi(\theta)$, and $\theta = \arg \min_{\theta \in \Theta} \nabla l_t(\theta_t)^\top \theta + \mathbb{D}_{\psi}(\theta || \theta_t)$. Unlike mirror descent, DMD-MPC does not consider an adaptive temperature strategy. Moreover, it introduces CEM- and MPPI- like updates by explicitly transforming the objective and performing gradient descent, rather than estimating the posterior.

Path integral control. Path integral theory connects Monte Carlo estimations to partial differential equation (PDE) solutions [28]. Using this, path integral control is used to motivate sample-approximation to the continuous-time Hamiltonian-Jacobian PDE. Crucially, the exponential transform is introduced to the value function term to make the PDE linear, a requirement to path integral theory. Path integral theory is limited to estimating an optimal action trajectory, given an initial state, and requires several additional assumptions on the dynamics and disturbance noise [64]. Later, a divergence-minimization perspective was applied to path integral methods, to produce algorithms for more general settings and discrete time, referred to both as (information theoretic) IT-MPC [10] and also MPPI [89]. This alternative view is closer to PPI and no longer contains the key assumptions that apply to path integral theory.

Variational Inference MPC Variational MPC [5] is the same posterior policy iteration scheme outlined in the main text. The key difference is a posterior entropy bonus, like MORE, added to the objective

$$\min_q \mathbb{E}_q[R] + \frac{1}{\lambda_1} \mathbb{D}_{\text{KL}}[q || p] + \frac{1}{\lambda_2} \mathbb{H}[q].$$

This results, once reparameterized, in the usual posterior update with an annealing coefficient κ on the prior

$$w_n \propto \exp(\alpha r_n) p(\mathbf{a}_n)^{-\kappa}.$$

Due to hyperparameter sensitivity, this entropy regularization required a normalization step itself

$$p(\mathbf{a})^{-\kappa} = \exp(-\kappa \log p(\mathbf{a})) \rightarrow \exp\left(-\bar{\kappa} \frac{\log p - \max \log p}{\min \log p - \max \log p}\right). \quad (8)$$

As a result, the entropy regularization is dynamic in practice. The mechanism of this update is to increase the weight of low-likelihood samples such that the entropy of the posterior increases. In the ablation study, the effect of this regularization was relatively small with limited statistical significance. The work also uses a Gaussian mixture model as the variational family for multi-modal action sequences. As an MPPI baseline, the authors use PI² for MPC with an adaptive temperature through normalization with $\bar{\alpha}$ set to 10.

Variational Inference MPC using Tsallis Divergence Adopting the ‘generalized’ variational inference approach, Wang et al. replace the KL divergence with the Tsallis divergence [66], resulting in discontinuous, rather than exponential, expression for the posterior weights

$$w_n \propto \max(0, 1 + (\gamma - 1) \alpha r_n)^{\frac{1}{\gamma-1}}, \quad \gamma > 0. \quad (9)$$

In practice, like in the CEM, k elites are used to define the r_* threshold

$$w_n \propto \begin{cases} \exp\left(\frac{1}{\gamma-1} \log\left(1 - \frac{r_n}{r_*}\right)\right) & \text{if } r_n < r_*, \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

The resulting parameterization, essentially combines the CEM- and importance sampling approaches, using elites but weighing them through γ . For the implementation, the number of elites and γ were tuned but kept constant during MPC. From the perspective outlined in this work, this approach limits the effective sample size at a maximum value, but allows it to drop if there is sufficient range in the elite returns. As a result, its optimization is greedier than CEM, resulting in better downstream control performance. In the paper, the authors compare to MPPI with a static temperature and do not consider adaptive schemes. In theory, the Tsallis divergence could be combined with LBPS to optimize γ adaptively.

Stein Variational Inference Model Predictive Control Stein variational gradient descent (SVGD) is an approximate inference method for learning multi-modal posteriors parameterized by particles and a kernel function, $q(\cdot) = \sum_n k(\cdot, \mathbf{x}_n)$. The key quality of SVGD is the kernel, which provides a repulsion force during learning to encourage particle diversity.

For control, SVGD can be used to infer multi-modal action sequence which exhibit high entropy for exploration, as done in SV-MPC [65]. Due to the limitation of kernel methods with high-dimensional inputs, the kernel is designed to be factorized in time and action. The kernel design is also Markovian, so it considers the temporally adjacent variables as well, encouraging smoothness. This design decomposes the kernel into the sum of $H(H-1)d_a$ one dimensional kernels for H timesteps. Moreover, extrapolation into the future is realized by copying the last timestep (i.e. a zero-order hold). SVGD requires gradients of the loglikelihood, using backpropagation through the dynamics or Monte Carlo estimates. SV-MPC has an additional learning rate hyperparameter and uses an independent Gaussian action prior across timesteps.

Algorithm 1 Open-loop Episodic Monte Carlo Posterior Policy Iteration

Requires:Markov decision process MDP, initial policy π_1 , posterior strategy GibbsPosterior, initial state s_0 .

- 1: **for** $i \leftarrow 1$ to N **do**
 - 2: Sample action sequences and associated parameters $\theta, \mathbf{A}^{(n)}, \theta^{(n)} \sim \pi_i(\cdot|s_0)$
 - 3: Obtain returns $R^{(n)} \sim \text{MDP}(\mathbf{A}^{(n)})$
 - 4: Compute importance weights, $\mathbf{w} \leftarrow \text{GibbsPosterior}(\mathbf{R})$ to estimate $q_\alpha(\mathbf{A}|\mathcal{O})$
 - 5: Update policy $\pi_{i+1} \leftarrow \text{MProjection}(\mathbf{w}, \theta)$, performing $\min_\pi \mathbb{D}_{\text{KL}}[q_\alpha || \pi]$
 - 6: **end for**
-

C Implementation Details

Algorithm 1 describes the general routine of Monte Carlo PPI. The specific GibbsPosterior update must be chosen, e.g. REPS, PI², MPPI, LBPS or ESSPS. CEM corresponds to uniform importance weights applied to the elite samples in GibbsPosterior. By definition, MPPI has a fixed covariance. For MPC, CEM methods reset their covariance each timestep. For actuator limits, we treat them as properties of the dynamics, and later fit the posterior using the *clipped* actions. This recognizes that the applied policy is not necessarily the same as the sampled one and is observed along with the reward. For feature regression weights, clipping is applied but cannot be incorporated into the fitting, as the model weights are being fit, which may explain its worse performance for MPC.

As done in REPS and MORE, the temperature optimization uses off-the-shelf minimizers, such as those found in `scipy.opt.minimize` [90]. While REPS and MORE use gradient-based solvers such as LBFGS-B, for LBPS and ESSPS, the `brent` method in `scipy.opt.minimize_scalar` was found to be slightly faster, presumably due to its quasi-convex objective (see Section F).

For the matrix normal distribution, a weighted maximum likelihood fit and sampling is very similar to the normal distribution. Sampling a matrix normal $\mathcal{MN}(\mathbf{M}, \mathbf{K}, \Sigma)$ is achieved using

$$\mathbf{X}^{(n)} = \mathbf{M} + \mathbf{A}\mathbf{W}^{(n)}\mathbf{B}, \quad \mathbf{K} = \mathbf{A}\mathbf{A}^\top, \quad \Sigma = \mathbf{B}^\top\mathbf{B}, \quad W_{ij}^{(n)} \sim \mathcal{N}(0, 1).$$

The weighted maximum likelihood fit of the input covariance is computed using [91]

$$\mathbf{K} = \sum_n w_n (\mathbf{X}^{(n)} - \mathbf{M})(\mathbf{X}^{(n)} - \mathbf{M})^\top \Sigma^{-1}, \quad \text{where } \sum_n w_n = 1.$$

For the feature approximation of the squared exponential kernel, we used RBF features and quadrature RFFs (QRFF). RBF features require careful normalization in order to approximate the SE kernel at the limiting case [33]. A d -dimensional RBF feature for a SE kernel approximation with lengthscale l is defined as

$$\phi(t) = \frac{1}{\sqrt{\sqrt{\pi} d \lambda}} \exp\left(-\frac{(t - \mathbf{c})^2}{2\lambda^2}\right),$$

where $\lambda = l/\sqrt{2}$. The d -dimensional centers \mathbf{c} are linearly placed along the task horizon.

Quadrature random Fourier features require $d = (2\nu)^k$ features for order ν and input dimension k . Therefore, for time-series we require 2ν features where

$$\phi_j(t) = \begin{cases} w_j \cos(\omega_j t) & j \leq \nu \\ w_{j-\nu} \sin(\omega_{j-\nu} t) & \nu < j \leq 2\nu \end{cases}$$

Gauss-Hermite quadrature provides points \mathbf{u} and weights \mathbf{v} for a given order, to be used for approximating integrals. In QRFFs, the Monte Carlo integration with Gaussian frequencies [60] is replaced with quadrature. As a result, $w_i = 2v_j/\sqrt{\pi}$. Incorporating the lengthscale, $\omega_i = \sqrt{2}u_i/l$. See Mutný et al. for a more in-depth description of QRFFs, including for higher-dimensional inputs [62].

A practical issue in stochastic search methods is maintaining a sufficiently exploratory search distribution. In accordance with Bayesian methods, PPI methods maintain a belief and do not keep a fixed covariance like MPPI. However, for some tasks it was found that the search distribution had insufficient variance to solve the task effectively, therefore, we introduced an adjustment to Equation 7 to

Algorithm	n_samples		
	16	128	1024
MPPI (white noise)	0.06	0.37	2.63
iCEM (coloured noise)	0.07	0.18	1.08
LBPS (SE kernel)	0.06	0.36	2.74
MPPI (SE kernel)	0.06	0.36	2.74
LBPS (RBF features)	0.06	0.37	2.79

Table 2: Wall-clock time (s) of one MPC calculation for HumanoidStandup-v2. Averaged over 10 timesteps and 5 seeds. Computation was performed on a AMD Ryzen 9 3900X 12-Core Processor @ 3.8GHz, parallelized across 24 processes.

‘anneal’ the covariance, recognizing the the update subtracts a likelihood-based term from the prior. Using

$$\Sigma_{t_3:t_4, t_3:t_4} | \mathcal{O} = \Sigma_{t_3:t_4, t_3:t_4} - \gamma \Sigma_{t_3:t_4, t_1:t_2} \Lambda_{t_1:t_2} \Sigma_{t_3:t_4, t_1:t_2}^\top, \quad 0 \leq \gamma \leq 1,$$

we perform standard PPI for $\gamma=1$, but adopt an MPPI-like approach for $\gamma=0$. For HumanoidStandup-v2, $\gamma=0.5$ was required to transition from standing to stabilization effectively. Appendix D.3.4 provides ablation studies of this aspect.

Investigating the runtime of these methods, comparing iCEM, MPPI and LBPS in Table 2, PPI methods are slower than CEM/iCEM, increasing with sample size. Rather than due to the α optimization or kernel-based prior construction, the bottleneck is using all the samples in the matrix normal MLE step for the covariance, computed using einsum operations on the sampled parameters. Based on this study, one approximation to speed up PPI methods is use the ESS to reduce the number of samples used in the matrix normal MLE step, pruning samples that have negligible weight, like in CEM.

D Extended Experimental Results

D.1 Black-box Optimization

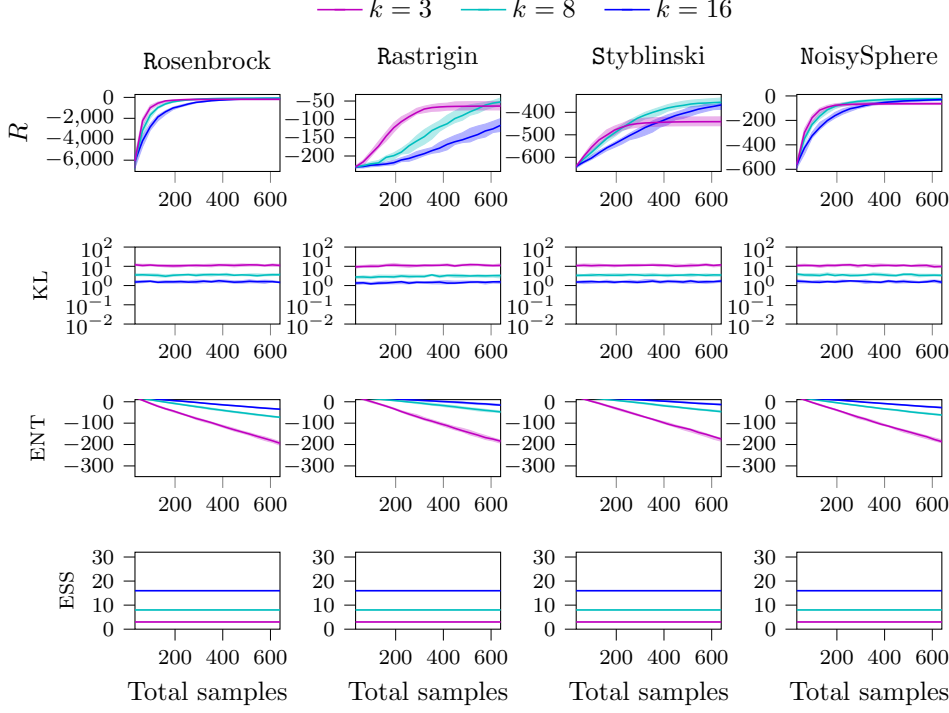


Figure 7: Black-box optimization with CEM and Monte Carlo sampling, with 32 samples over 20 episodes, displaying quartiles over 25 seeds. k is the number of ‘elite’ samples.

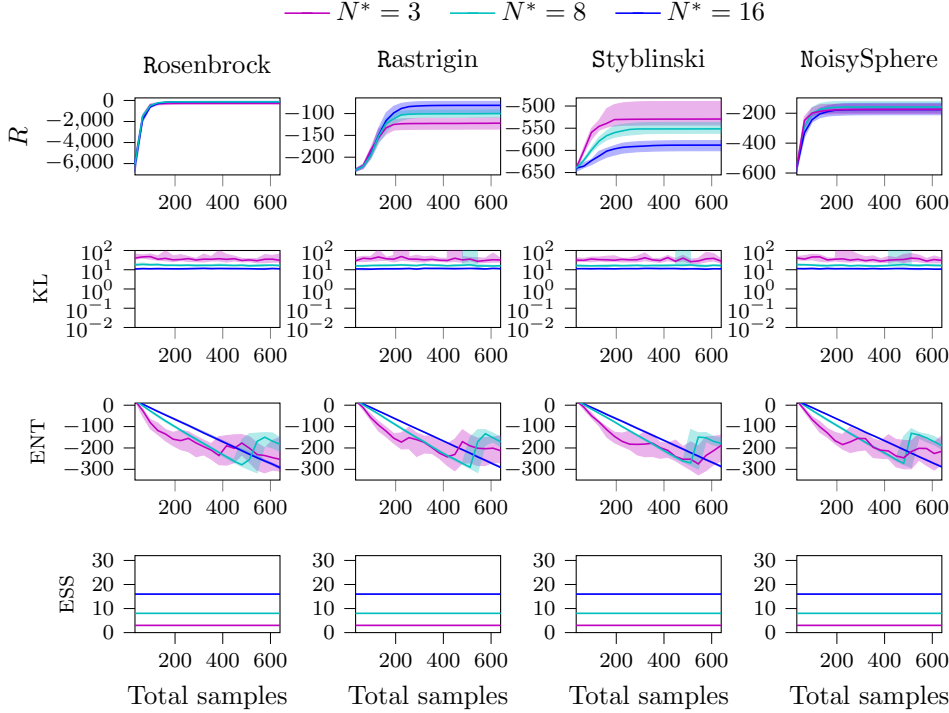


Figure 8: Black-box optimization with ESSPS and Monte Carlo sampling, with 32 samples over 20 episodes, displaying quartiles over 25 seeds. N^* is the desired effective sample size.

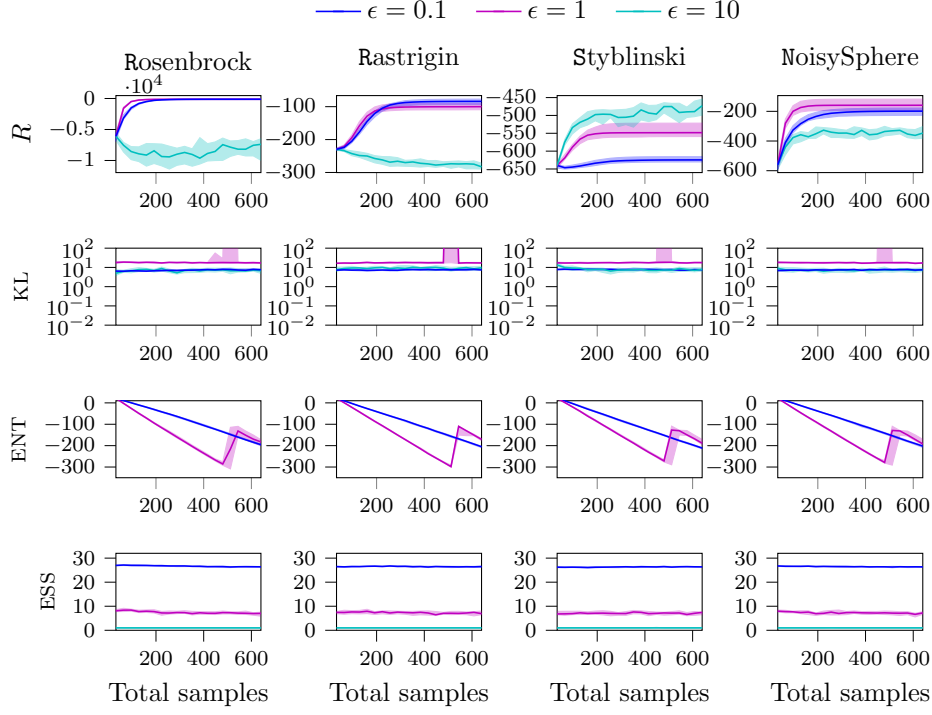


Figure 9: Black-box optimization with REPS and Monte Carlo sampling, with 32 samples over 20 episodes, displaying quartiles over 25 seeds. ϵ is the KL bound.

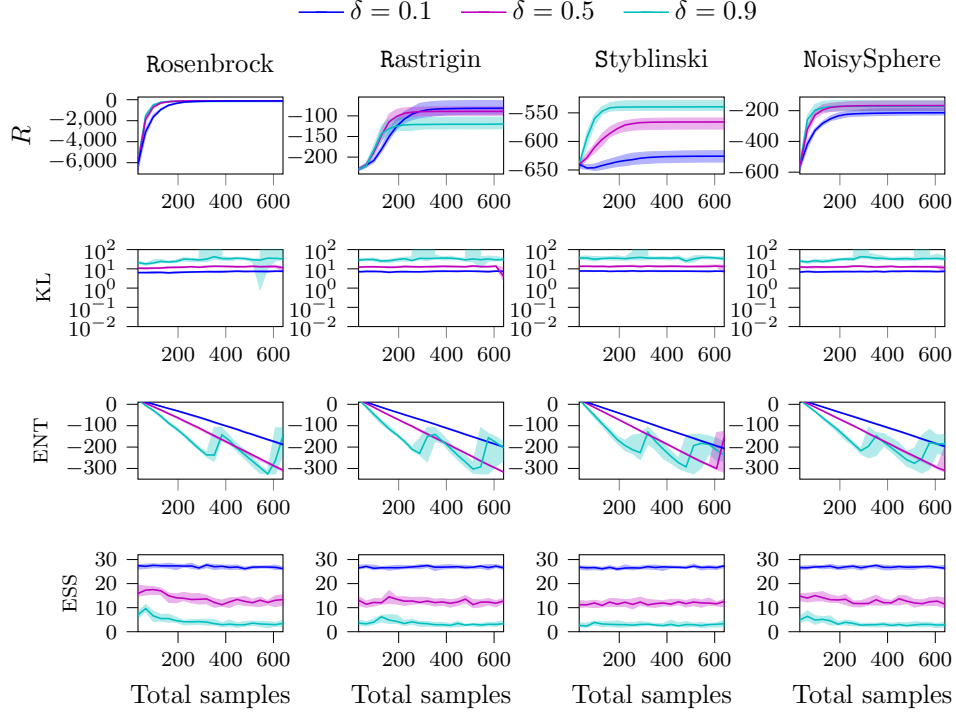


Figure 10: Black-box optimization with LBPS and Monte Carlo sampling, with 32 samples over 20 episodes, displaying quartiles over 25 seeds. δ is the probability of the lower bound.

D.2 Policy Search




Figure 11: Success rate for policy search with RBF features over 20 seeds, using 128 rollout samples per episode. Displaying uncertainty in quartiles




Figure 12: Success rate for policy search with RFF features over 20 seeds, using 128 rollout samples per episode. Displaying uncertainty in quartiles

D.3 Model Predictive Control




Figure 13: Average ESS for Monte Carlo MPC with white noise priors. Displaying quartiles over 50 seeds.

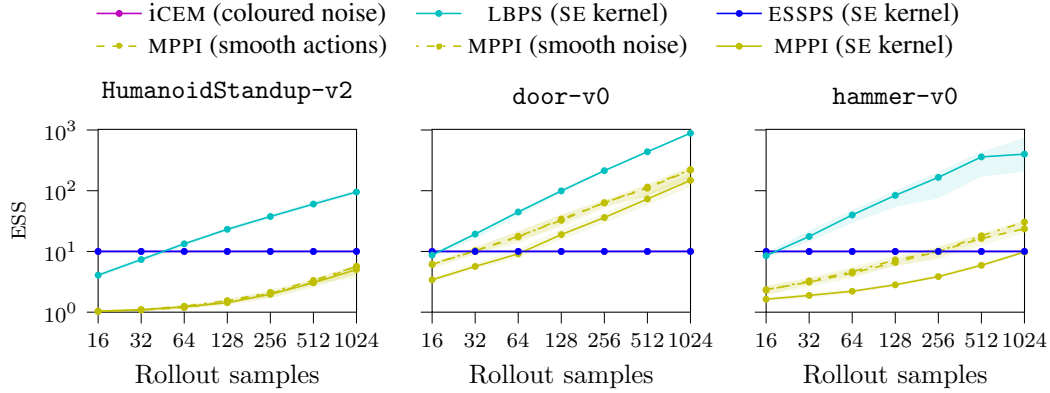


Figure 14: Average ESS for Monte Carlo MPC with smooth noise priors. Displaying quartiles over 50 seeds. iCEM and ESSPS both have an ESS of 10, due to iCEM having 10 elites.

D.3.1 MPC with finite feature approximations

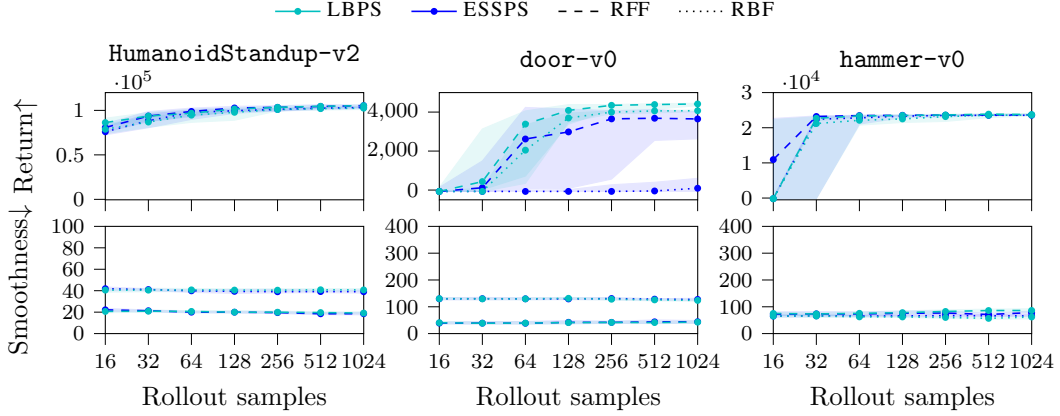


Figure 15: Monte Carlo MPC with with RBF and RFF policies. Displaying quartiles over 50 seeds.




Figure 16: Average ESS for Monte Carlo MPC with RBF and RFF policies. Displaying quartiles over 50 seeds.




Figure 17: Monte Carlo MPC with RBF and RFF policies, using a (shrinking) planning horizon of 250 rather than 30, the full task duration. Displaying quartiles over 20 seeds.

D.3.2 Visualizing actuation profiles

To accompany the smoothness metric used in the MPC evaluation, we showed how the SE kernel with PPI produces smoother and lower amplitude policies than alternative priors. Due to the high-dimensional action spaces, we depict all actions overlapped as they are normalized.




Figure 18: iCEM MPC, with coloured noise, action sequence using 16 rollouts.




Figure 19: iCEM MPC, with coloured noise, action sequence for using 1024 rollouts.




Figure 20: LBPS MPC, using the SE kernel, action sequence using 16 rollouts.




Figure 21: LBPS MPC, using the SE kernel, action sequence using 1024 rollouts.




Figure 22: ESSPS MPC, using the SE kernel, action sequence using 16 rollouts.




Figure 23: ESSPS MPC, using the SE kernel, action sequence using 1024 rollouts.




Figure 24: MPPI MPC, using the SE kernel, action sequence using 16 rollouts.




Figure 25: MPPI MPC, using the SE kernel, action sequence using 1024 rollouts.




Figure 26: MPPI MPC, using smooth action noise, action sequence using 16 rollouts.




Figure 27: MPPI MPC, using smooth action noise, action sequence using 1024 rollouts.




Figure 28: MPPI MPC, using smooth exploration noise, action sequence using 16 rollouts.




Figure 29: MPPI MPC, using smooth exploration noise, action sequence using 1024 rollouts.

Setting	Return	Smoothness	Lengthscale
door-v0 (expert)	[3301, 4001, 4028]	[139 145 150]	1×10^{-5}
door-v0 (LBPS MPC)	[4326, 4370, 4388]	[50, 55, 58]	0.05
hammer-v0 (expert)	[14042, 17054, 21272]	[142, 147, 151]	1×10^{-5}
hammer-v0 (LBPS MPC)	[23810, 23828, 23904]	[51, 53, 62]	0.025
HumanoidStandup-v2 (GAC)	[92429, 92972, 93300]	– [†]	1×10^{-5}
HumanoidStandup-v2 (LBPS MPC)	[99902 100147 101337]	[20 22 23]	0.05

Table 3: Comparing expert demonstrations to PPI MPC performance w.r.t. return, smoothness and lengthscale. LBPS MPC performance is reported for 1024 rollouts. Uncertainty shows quartiles over dataset demonstrations or experiment seeds. See the discussion text for the omitted result [†].

D.3.3 Learning priors from data

A benefit of the Gaussian process view of policy priors is the ability to perform model selection from data, rather the hyperparameter tuning. This data could be expert demonstrations or a partial solution, such as the ‘warm start’ in MPC.

We investigate data driven priors by taking expert demonstrations from human and RL agents, analyzing them using the matrix normal distribution. The covariances of the matrix normal allow us to assess the stationarity of the temporal correlations, as well as the correlations between actions. For example, in Figures 30–31 we can see non-stationarity of door-v0 and hammer-v0, due to the motions before and after completing the task, whereas HumanoidStandUp-v2 appears surprisingly stationary due to the task having ‘stand up’ and ‘stabilize’ components. However, we found the smoothness that benefited MPC agents in Section 6.3 was not present in the demonstrations. For door-v0 and hammer-v0, the action space is desired joint positions and the demonstrations were collected using motion capture technology [82]. With this in mind, the non-smoothness may be the result of motion capture artefacts or the inverse kinematics used. While the video results of the demonstrations do not suggest rough motion, the action sequences in the dataset do appear rough (Figures 30–31). Moreover, there may be unknown components of the control stack that smooth out the desired joint setpoints. However, the issue may lie in the matrix normal factorization, which assumes each action share the same temporal correlations. This coupling may result in missing smooth sequences if many dimensions have a rough actions.

For HumanoidStandUp-v2, we train and use demonstrations from a gac (generative actor critic) agent [92]. To our knowledge, this is the only model-free deep RL algorithm that solves HumanoidStandUp-v2. However, the policy learned by GAC is a bang-bang controller that operates at the action limits. As a result, the smoothness measure introduced in Section 6.3 ‘breaks’, as the norm of this action sequence is *constant*, suggesting a smoothness of 0. Looking per action independently, the smoothness metric varies around 100 to 300. This result suggests that the IID noise used in deep RL exploration may influence its optimal policy estimate towards rough behaviors such as bang-bang control, which limits their usefulness as expert demonstrators.

The second quantity of interest is the action covariance Σ . The demonstrations of door-v0 and hammer-v0 suggest that the independence assumption of Σ is a reasonable one. The action covariance of HumanoidStandUp-v2 depicts much stronger correlations between actions, which could be used to improve exploration through coordination.




Figure 30: door-v0 expert demonstrations, showing matrix normal covariance fit and action sequence. The viridis colourmap is used to express correlations.




Figure 31: hammer-v0 expert demonstrations, showing matrix normal covariance fit and action sequence. The viridis colourmap is used to express correlations.




Figure 32: hammer-v0 expert demonstrations from a MPC solvers (MPPI, LBPS, ESSPS) using the SE kernel, showing matrix normal covariance fit and action sequence. The viridis colourmap is used to express correlations.

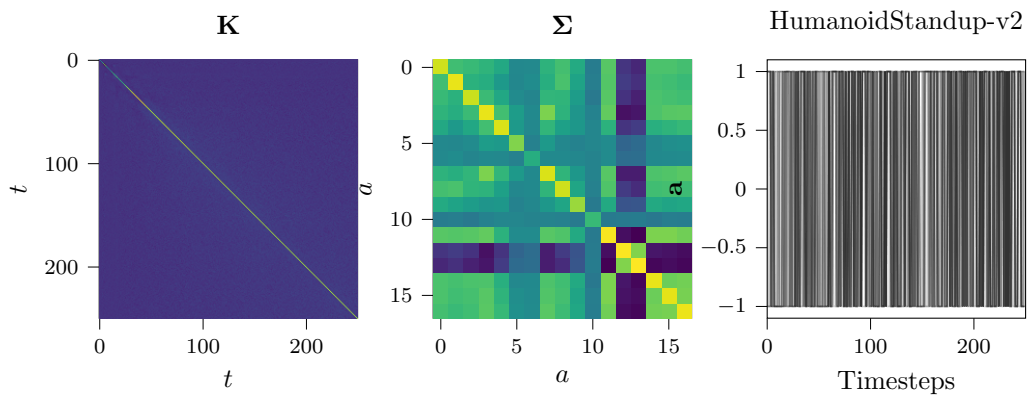


Figure 33: HumanoidStandup-v2 expert demonstrations from a GAC agent, showing matrix normal covariance fit and action sequence. The viridis colourmap is used to express correlations.




Figure 34: HumanoidStandup-v2 expert demonstrations from a MPC solvers (MPPI, LBPS, ES-SPS) using the SE kernel, showing matrix normal covariance fit and action sequence. The viridis colourmap is used to express correlations.

D.3.4 Model predictive control ablation studies

The MPC methods compared in Section 6.3 have subtle variations in their implementation. MPPI has a constant covariance during optimization, while CEM MPC resets the covariance each timestep. To understand the implications of these details, we provide ablations over these design decisions, in comparison to the results in Section 6.3.

LBPS and ESSPS with constant covariances

On the whole the performance drops compared to the updated covariance results (Figure 4), which suggests that a fixed variance requires greedier optimization (i.e. MPPI) or more iterations per timestep for good performance.




Figure 35: Monte Carlo MPC with LBPS and ESSPS with fixed covariance, like MPPI, as opposed to updated each timestep. Displaying quartiles over 20 seeds.

MPPI with covariance updates

Conversely to the result above, MPPI’s greedier optimization does not work effectively with covariance updates, as entropy is quickly lost during the control.




Figure 36: Monte Carlo MPC using MPPI with covariance updates. Displaying quartiles over 20 seeds.

CEM with a SE kernel prior

Compared to the white noise prior in Figure 3, the SE kernel provides an improvement boost for HumanoidStandup-v2 and hammer-v0. There is a smoothness improvement, but slightly smaller than compared to Figure 4. Based on the results of Figure 35, we can attribute this to CEM MPC resetting the policy covariance each timestep.




Figure 37: Monte Carlo MPC using CEM with the SE. Displaying quartiles over 20 seeds. Results closely match ESSPS with the SE kernel and fixed covariance (Figure 35).

E Experimental Details

E.1 Hyperparameter Selection

Parameters were swept with there respective solver for best average performance with 256 sample rollouts, averaged over 5 seeds. Parameters were not swept if prior work had identified effective values.

Table 4: Hyperparameter selection for model predictive control methods

	α (MPPI)	$\bar{\alpha}$ (PI ²) [64]	δ (LBPS)	k, N^* (CEM, iCEM, ESSPS) [4]
Considered	[0.1, 1, 10]	-	[0.1, 0.5, 0.9]	-
HumanoidStandup-v2	10	10	0.9	10
door-v0	10	10	0.5	10
hammer-v0	10	10	0.5	10

Table 5: Hyperparameter selection for model predictive control policies

	β (smooth noise)	β (smooth action)	β (coloured noise) [4]	l (SE kernel)
Considered	[0.5, 0.7, 0.9]	[0.5, 0.7, 0.9]	-	[0.01, 0.025, 0.05, 0.1]
HumanoidStandup-v2	0.5	0.7	2.0	0.05
door-v0	0.9	0.9	2.5	0.05
hammer-v0	0.7	0.7	2.5	0.025

E.2 Experiment Hyperparameters

Black-box Optimization. We evaluated the solvers on the 20-dimensional version of the test functions. The initial search distribution was $\mathcal{N}(\mathbf{1}, 0.5\mathbf{I})$ for all methods. 32 samples were used over 20 iterations.

Policy Search. The experimental setting of prior work [14] was reproduced, with the same MuJoCo simulation and episodic reward function.

Table 6: Hyperparameters for policy search tasks.

	T	d_a	n.iter	n.samples	RBF features	RFF order	lengthscale, l
BallInACup	1000	4	40	128	20	10	$\sqrt{0.03}$

Model Predictive Control. For the MPC tasks, the action mean and variance for each dimension i was set using

$$\mu_i = \frac{a_{\max,i} + a_{\min,i}}{2}, \quad \sigma_i^2 = \frac{(a_{\max,i} - a_{\min,i})^2}{4}.$$

This specification means that $\mu_i \pm \sigma_i$ reaches the actuator limits, ensuring coverage across the actuation range when sampling. μ_i defines a ‘mean function’ applied as a bias to the policy, which had nominally zero mean. σ_i^2 defines the diagonal of the matrix normal action covariance Σ . To match the covariance size of the kernel, 30 features were used for the RBF and RFF approximations, i.e. order $\nu = 15$.

Table 7: Hyperparameters for model predictive control tasks.

	T	H	d_a	n.iters	n.warmstart.iters
HumanoidStandup-v2	250	30	17	2	50
door-v0	250	30	25	1	50
hammer-v0	250	30	25	1	50

F Characterizing the Lower-bound Objective

The temperature objective of REPS is the dual function, which is convex (Definition 3). For the lower bound introduced Section 3, we characterize its nature to understand its suitability for optimization.

First, we express objective in terms of functions f , g and h of α ,

$$\begin{aligned}\max_{\alpha} \mathcal{L}(\alpha) &= \frac{\sum_n \exp(\alpha r_n) r_n}{\sum_n \exp(\alpha r_n)} - \lambda \sqrt{\frac{1}{\hat{N}(\alpha)}}, \quad \hat{N}(\alpha) = \frac{(\sum_n \exp(\alpha r_n))^2}{\sum_n \exp(2\alpha r_n)} \\ &= \frac{\sum_n \exp(\alpha r_n) r_n - \lambda \sqrt{\sum_n \exp(2\alpha r_n)}}{\sum_n \exp(\alpha r_n)} = \frac{f(\alpha) - \lambda g(\alpha)}{h(\alpha)}.\end{aligned}$$

The gradient is available in closed form

$$\frac{d}{d\alpha} \mathcal{L}(\alpha) = \frac{(\sum_n w_n)(\sum_n w_n r_n^2) - \lambda(\sum_n w_n) \sum_n r_n w_n^2 / \sqrt{\sum_n w_n^2} - (\sum_n w_n r_n)^2 + (\sum_n r_n w_n)(\lambda \sqrt{\sum_n w_n^2})}{(\sum_n w_n)^2}.$$

This objective is not concave in α , but is quasi-concave. With $\alpha \in \mathbb{R}_+$, we enforce $r_n \leq 0$ and $\exp(\alpha r_n)$ is convex in α for $r_n \in \mathbb{R}$. We rewrite the objective using in dot product form

$$\min_{\alpha} \mathcal{L}(\alpha) = \frac{\sum_n \exp(\alpha r_n) r_n + \lambda \sqrt{\sum_n \exp(2\alpha r_n)}}{\sum_n \exp(\alpha r_n)} = \frac{\mathbf{w}_{\alpha}^{\top} \mathbf{r} - \lambda \sqrt{\mathbf{w}_{\alpha}^{\top} \mathbf{w}_{\alpha}}}{\mathbf{w}_{\alpha}^{\top} \mathbf{1}},$$

where $\mathbf{w}_{\alpha} = [\exp(\alpha r_1), \dots, \exp(\alpha r_N)]^{\top}$. The term $\mathbf{w}_{\alpha}^{\top} \mathbf{r}$ is concave as it is a negative weighted sum of convex functions. The term $\mathbf{w}_{\alpha}^{\top} \mathbf{1}$ is convex as it is a positive sum of weighted convex functions. The remaining term $g(\alpha)$ can be shown to be convex in α (Lemma 3) using standard techniques [93].

Lemma 3. *The function $g(\alpha) = \sqrt{\sum_n \exp(2\alpha r_n)}$ is convex in α for $r_n \leq 0 \forall n$.*

Proof. For convexity, where $\theta \in [0, 1]$, $\alpha \in \mathbb{R}_+$, $\beta \in \mathbb{R}_+$,

$$\sqrt{\sum_n \exp(2(\theta\alpha + (1-\theta)\beta)r_n)} \leq \theta \sqrt{\sum_n \exp(2\alpha r_n)} + (1-\theta) \sqrt{\sum_n \exp(2\beta r_n)}$$

Starting with the right-hand term, we take interpolation term θ inside

$$\theta \sqrt{\sum_n \exp(2\alpha r_n)} + (1-\theta) \sqrt{\sum_n \exp(2\beta r_n)} = \sqrt{\sum_n (\theta \exp(\alpha r_n))^2} + \sqrt{\sum_n ((1-\theta) \exp(\beta r_n))^2}$$

Apply Jensen's inequality to both exponential terms, where $\exp(xy) \leq x \exp(y)$, as they are inside Euclidean norms so we can use $\sqrt{\sum_i p_i^2} \leq \sqrt{\sum_i q_i^2}$ if $p_i \leq q_i$,

$$\sqrt{\sum_n \exp(\theta\alpha r_n)^2} + \sqrt{\sum_n \exp((1-\theta)\beta r_n)^2} \leq \sqrt{\sum_n (\theta \exp(\alpha r_n))^2} + \sqrt{\sum_n ((1-\theta) \exp(\beta r_n))^2}.$$

Apply Minkowski's inequality to the LHS, where $(\sum_i |x_i + y_i|^p)^{1/p} \leq (\sum_i |x_i|^p)^{1/p} + (\sum_i |y_i|^p)^{1/p}$

$$\sqrt{\sum_n (\exp(\theta\alpha r_n) + \exp((1-\theta)\beta r_n))^2} \leq \sqrt{\sum_n \exp(\theta\alpha r_n)^2} + \sqrt{\sum_n \exp((1-\theta)\beta r_n)^2}$$

Expand the terms of the LHS and remove the (non-negative) squared terms

$$\sqrt{\sum_n 2 \exp((\theta\alpha + (1-\theta)\beta)r_n)} \leq \sqrt{\sum_n \exp(2\theta\alpha r_n) + \exp(2(1-\theta)\beta r_n) + 2 \exp((\theta\alpha + (1-\theta)\beta)r_n)}$$

Apply Jensen's inequality again to recover the initial left hand term

$$\sqrt{\sum_n \exp(2(\theta\alpha + (1-\theta)\beta)r_n)} \leq \sqrt{\sum_n 2 \exp((\theta\alpha + (1-\theta)\beta)r_n)}.$$

□

With the lemma, the negative penalty term is concave as $\lambda \geq 0$.

For quasi-convexity, we require the t -level sets to be convex $\{\alpha \in \mathbb{R}_+ \mid \mathcal{L}(\alpha) \leq t\}$, $t \in \mathbb{R}$. Due to the structure of the objective, we require $f(\alpha) - \lambda g(\alpha) \leq t h(\alpha)$. As $f(\alpha) \leq 0$, $-\lambda g(\alpha) \leq 0$ and $h(\alpha) \geq 0$, for $t > 0$ we have the empty set, which is convex. For $t \leq 0$, we have the sum of two concave function which are both less or equal to zero, so the set is also convex.

G Stochastic processes, Gaussian processes and coloured noise

This section summarizes the connections between stochastic processes, smoothed noise, coloured noise and Gaussian processes to motivate the use of kernels in action priors. For further details, we refer to Section 12.3 of Särkkä et al. [94] and Chapters 4 and Appendix B of Rasmussen et al. [33].

Section 4 introduced smooth Gaussian noise processes of the form

$$n_t^{(n)} = \sum_{i=1}^p a_i n_{t-i}^{(n)} + b_0 v^{(n)}, \quad v^{(n)} \sim \mathcal{N}(0, 1).$$

to sample action sequences. This is known as a discrete-time autoregressive AR(p) process. The ARMA(p, q) process introduces a noise history, such that ARMA($p, 0$) is a AR(p) process

$$n_t^{(n)} = \sum_{i=1}^p a_i n_{t-i}^{(n)} + \sum_{j=1}^q b_j v_{t-j}^{(n)}, \quad v_t^{(n)} \sim \mathcal{N}(0, 1).$$

In continuous-time, the AR(1) is analogous to the Ornstein–Uhlenbeck (OU) process

$$dn(t) = a_0 n(t) dt + v(t) dt.$$

The OU covariance function is $\text{Cov}(t, t') = \sigma^2 \exp(-a_0 |t - t'|)$, which is also known as the exponential kernel. For additional smoothness we can consider higher-order derivatives, which results in the Matérn family of kernels. In stochastic differential equation form, they are written as

$$n(t) = \mathbf{H} \mathbf{f}(t), \quad d\mathbf{f}(t) = \mathbf{A} \mathbf{f}(t) dt + \mathbf{L} v(t) dt,$$

where \mathbf{f} contains $n(t)$ and its derivatives, describing the state. The order ν of the Matérn defines the size of the state space and \mathbf{A} . A first-order kernel reduces to the exponential kernel. These Matérn kernels are Markovian in their state space, following a linear Gaussian dynamical system (LGDS). Therefore, they can be compared to the Gaussian processes used in STOMP [8] and GPMP [30], which are also LGDSs but with different state space models that perform Euler integration, producing priors with Gaussian noise on the velocity, acceleration or jerk. Extending the derivatives for the Matérn kernel, as the order $\nu \rightarrow \infty$, we arrive at the squared exponential kernel $\text{Cov}(t, t') = \sigma^2 \exp(-|t - t'|^2 / 2l^2)$. Comparing to the exponential kernel earlier, we observe the a_0 in the OU process defines the lengthscale of the covariance function.

Considering stationary covariance functions, where $\text{Cov}(t, t+\tau) = \text{Cov}(\tau)$, the power spectral density is defined as the Fourier transform of the covariance function

$$S(\omega) = \int \text{Cov}(\tau) \exp(i\omega\tau) d\tau.$$

From the linear systems perspective, the parameters a and b of an ARMA process describe a linear filter $F(\omega)$ where $N(\omega) = F(\omega)V(\omega)$. In the frequency domain, such is realized as a filter

$$F(\omega) = \frac{|B(\exp i\omega)|^2}{|A(\exp i\omega)|^2}, \quad \text{where} \quad A(\omega) = \sum_{k=1}^p a_k \exp(ik\omega).$$

Conversely, *coloured* noise with parameter β applies a filter $\propto 1/\omega^\beta$ to Gaussian noise. The power spectrum of the ν -order Matérn kernel is $S(\omega) \propto (l^2 + \omega^2)^{\nu/2+1}$ and the squared exponential is $S(\omega) \propto \exp(-l^2\omega^2/2)$. While there is no explicit connection between coloured noise and Gaussian processes, by reasoning about their power spectrums it can be seen that they can produce similar sampled paths.