

Stein Variational Model Predictive Control

Alexander Lambert^{1,2} * Adam Fishman^{2,3} Dieter Fox^{2,3}

Byron Boots^{2,3} Fabio Ramos^{2,4}

¹ Georgia Tech ² NVIDIA ³ University of Washington ⁴ The University of Sydney

Abstract: Decision making under uncertainty is critical to real-world, autonomous systems. Model Predictive Control (MPC) methods have demonstrated favorable performance in practice, but remain limited when dealing with complex probability distributions. In this paper, we propose a generalization of MPC that represents a multitude of solutions as posterior distributions. By casting MPC as a Bayesian inference problem, we employ variational methods for posterior computation, naturally encoding the complexity and multi-modality of the decision making problem. We propose a Stein variational gradient descent method to estimate the posterior over control parameters, given a cost function and a sequence of state observations. We show that this framework leads to successful planning in challenging, non-convex optimal control problems.

Keywords: Model Predictive Control, Variational Inference

1 Introduction

Model predictive control (MPC) is a powerful framework for sequential decision making in robotics [1, 2]. This success can be largely attributed to its simplicity and scalability in dealing with stochastic, non-stationary optimization problems encountered on real systems. MPC has been applied effectively in different areas, including autonomous driving [1–3], humanoid locomotion [4], and dextrous manipulation [5]. However, common approaches to MPC often fall short in their ability to adequately contend with complex, multi-modal distributions over possible actions. For instance, such distributions may arise from non-convexity of constraints, such as obstacles [6, 7] or from multiple goal locations [8]. Although common sampling-based SOC algorithms have been shown to exhibit symmetry-breaking in the presence of sudden disturbances [9] and multiple optima [10], the sampling scheme may inadequately resolve the true posterior [6]. We require a new class of MPC algorithms that can effectively contend with complex, non-Gaussian distributions.

In the following, we formulate MPC as a Bayesian inference problem, where the target posterior is defined directly over control policy parameters or control inputs, as opposed to joint probabilities over states and actions [11, 12]. By taking this perspective, we can construct a relative-entropy minimization problem to approximate the posterior, and leverage recent advances in variational inference [13–15] to derive the optimal distribution over parameters. Specifically, we use Stein variational gradient descent [16], to infer a set of solutions which constitute a nonparametric approximation to the posterior distribution over decision parameters, given state observations and a defined cost function. The generality of this approach offers flexibility in designing appropriate MPC algorithms, and is directly related to common MPC approaches. Furthermore, we show how this framework can be extended to general trajectory optimization problems.

Related Work The duality between probabilistic inference and optimization for stochastic optimal control (SOC) has been examined extensively in previous work [17–20]. A wide range of approximate inference methods have been proposed, including Expectation-Propagation (EP) and moment-matching approaches to control and trajectory optimization [3, 19]. These methods typically assume a restricted form of the target posterior distribution over controls (usually in the exponential family), which, under simplifying assumptions, return the optimal control input. However, these distributions are often insufficiently expressive for many SOC problems, where non-convexity may arise from nonlinear dynamics or non-convex cost functions, for instance.

*E-mail: alambert6@gatech.edu

Our interpretation of MPC as an approximate inference problem is perhaps most closely related to the formulation presented in [11, 18]. Here, an iterative KL-minimization problem for finite-horizon problems is presented, where the prior is defined to be the control distribution obtained from the previous iteration. However, this does not include a strategy for contending with non-stationary distributions, where the posterior may change between iterations, and the analysis is restricted to exact inference of the log-partition function.

Expectation-Maximization (EM) approaches to SOC have also been considered [21–23], which iteratively optimize a variational lower-bound. However, these approaches share the limitations inherent with common EM strategies, in that the form of the posterior distribution is assumed to be known and tractable. Additionally, representations based on mixture models must contend with mode-collapse and poor local minima, especially prevalent in higher dimensions [24, 25]. These issues can be mitigated, for example, by introducing entropy regularization heuristics [22].

The Path Integral (PI) formulation of stochastic optimal control [26–28] bears close resemblance to the open-loop-controls characterization of our proposed approach, where a particular choice of the marginal log-likelihood is assumed to contain an exponentiated cost. In fact, PI attempts to minimize a variational lower-bound with respect to the controlled stochastic dynamics. This follows from the perspective presented here. This comparison applies equally to KL-control [20], the discrete-time counterpart to PI-control. A more detailed discussion is included in Appendix E.

2 Model Predictive Control

We consider the discrete-time stochastic dynamical system: $\mathbf{x}_{t+1} \sim f(\mathbf{x}_t, \mathbf{u}_t)$, where at time t , the system state is denoted by $\mathbf{x}_t \in \mathbb{R}^n$ and the control input as $\mathbf{u}_t \in \mathbb{R}^d$. The stochastic transition map $f: \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^n$ randomly produces the subsequent state \mathbf{x}_{t+1} and this state is accompanied by an instantaneous cost $c(\mathbf{x}_t, \mathbf{u}_t)$.

Over a time horizon H , we define a control trajectory as a sequence of control inputs beginning at time $t: U_t \triangleq (\mathbf{u}_t, \mathbf{u}_{t+1}, \dots, \mathbf{u}_{t+H-1})$. Similarly, we define the state trajectory: $X_t \triangleq (\mathbf{x}_t, \mathbf{x}_{t+1}, \dots, \mathbf{x}_{t+H-1}, \mathbf{x}_{t+H})$. The total cost incurred over H timesteps can then be specified as

$$C(X_t, U_t) = c_{\text{term}}(\mathbf{x}_{t+H}) + \sum_{h=0}^{H-1} c(\mathbf{x}_{t+h}, \mathbf{u}_{t+h}), \quad (1)$$

where $c_{\text{term}}(\cdot)$ is the terminal cost. As in [2], we define an instantaneous feedback policy, $\pi_{\theta_t}(\mathbf{x}_t)$, as a parameterized probability distribution $p(\mathbf{u}_t|\mathbf{x}_t; \theta_t)$ used to generate a control input at time t given the \mathbf{x}_t , *i.e.*, $\mathbf{u}_t \sim \pi_{\theta_t}(\mathbf{x}_t)$, where $\theta_t \in \Theta$, is the set of feasible parameter values. MPC describes the process of finding the optimal, time-indexed sequence of policy parameters $\boldsymbol{\theta}_t \triangleq (\theta_t, \theta_{t+1}, \dots, \theta_{t+H-1})$, which determine the sequence of instantaneous feedback policies $\boldsymbol{\pi}_{\theta_t} \triangleq (\pi_{\theta_t}, \pi_{\theta_{t+1}}, \dots, \pi_{\theta_{t+H-1}})$. At each time step, we must find $\boldsymbol{\theta}_t$, the parameters that define the optimal policy. We can do this by defining a statistic $J(\cdot)$ on cost $C(X_t, U_t)$ where the minimal $J(\cdot)$ occurs at the optimal $\boldsymbol{\theta}_t$. In real-world situations, the true dynamics function f is often unavailable, and is commonly estimated using a parameterized function \hat{f}_ξ with parameters ξ . As such, we define the surrogate loss function $\hat{J}(\boldsymbol{\pi}_\theta; \mathbf{x}_t) = \mathbb{E}_{\boldsymbol{\pi}_\theta, \hat{f}_\xi} [C(X_t, U_t)]$. For each MPC-step, the optimal decision is defined as $\boldsymbol{\theta}_t = \arg \min_{\boldsymbol{\theta}} \hat{J}(\boldsymbol{\pi}_\theta; \mathbf{x}_t)$ which parameterizes the optimal policy π_{θ_t} , from which we can sample a new control value for the first timestep: $\mathbf{u}_t \sim \pi_{\theta_t}(\mathbf{x}_t) = p(\mathbf{u}_t|\mathbf{x}_t; \theta_t)$. This is then executed on the physical system to generate the next state value: $\mathbf{x}_{t+1} \sim f(\mathbf{x}_t, \mathbf{u}_t)$.

3 Bayesian Model Predictive Control

MPC as Bayesian Inference Optimal control can be framed as Bayesian inference by considering the distribution over parameters $\boldsymbol{\theta}$. Similarly to [12, 18], we introduce an auxiliary binary random variable $\mathcal{O}_\tau = \{0, 1\}$ to indicate optimality of the state-action trajectory $\tau = (X_t, U_t)$ with respect to the cost function $C(\cdot)$. Using Bayes' rule, the distribution of parameters $\boldsymbol{\theta}$ conditioned on the requirement for optimal trajectories ($\mathcal{O}_\tau = 1$) and the current state \mathbf{x}_t can be expressed as

$$p_t(\boldsymbol{\theta}|\mathcal{O}_\tau = 1; \xi, \mathbf{x}_t) = \frac{p_t(\mathcal{O}_\tau = 1|\boldsymbol{\theta}; \xi, \mathbf{x}_t) p_t(\boldsymbol{\theta}; \mathbf{x}_t)}{\int p_t(\mathcal{O}_\tau = 1|\boldsymbol{\theta}; \xi, \mathbf{x}_t) p_t(\boldsymbol{\theta}; \mathbf{x}_t) d\boldsymbol{\theta}}, \quad (2)$$

where explicit dependence on state \mathbf{x}_t is included for generality. In the remaining discussion, we will denote $\mathcal{O}_\tau = 1$ simply as \mathcal{O}_τ without ambiguity. The likelihood $p_t(\mathcal{O}_\tau|\boldsymbol{\theta}; \xi, \mathbf{x}_t)$ is defined as the marginal probability over all possible control and state trajectories:

$$p_t(\mathcal{O}_\tau|\boldsymbol{\theta}; \xi, \mathbf{x}_t) = \int \int p(\mathcal{O}_\tau|X_t, U_t) p(X_t, U_t|\boldsymbol{\theta}; \xi, \mathbf{x}_t) dU_t dX_t \quad (3)$$

where $p(\mathcal{O}_\tau|X_t, U_t)$ is the probability of optimality given the observed trajectory $\tau = (X_t, U_t)$, and $p(X_t, U_t|\boldsymbol{\theta}; \xi, \mathbf{x}_t)$ is the joint probability of state-control trajectories, conditioned on parameters $\boldsymbol{\theta}$ and assumed dynamics model $\hat{f}_\xi = p_\xi(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$. In the discrete-time case, the joint probability can be factorized as

$$p(X_t, U_t|\boldsymbol{\theta}; \xi, \mathbf{x}_t) = \prod_{h=0}^{H-1} p_\xi(\mathbf{x}_{t+h+1}|\mathbf{x}_{t+h}, \mathbf{u}_{t+h}) \pi_{\theta_h}(\mathbf{x}_{t+h}), \quad (4)$$

where the current state \mathbf{x}_t is observed. If dynamics parameters must also be inferred, the equations can be extended to include ξ as an inference variable, where the posterior is defined over both parameters: $p_t(\boldsymbol{\theta}, \xi|\mathcal{O}_\tau; \mathbf{x}_t) \propto p_t(\mathcal{O}_\tau|\boldsymbol{\theta}, \xi; \mathbf{x}_t) p_t(\boldsymbol{\theta}|\xi; \mathbf{x}_t) p(\xi)$, suggesting the definition of an alternative latent random variable, such as $\boldsymbol{\Theta} = (\boldsymbol{\theta}, \xi)$.

We can further model $p(\mathcal{O}_\tau|X_t, U_t)$ using a non-negative function $\mathcal{L}(\tau) \propto p(\mathcal{O}_\tau|X_t, U_t)$, which we refer to as the ‘‘cost-likelihood’’. This is defined to be the composition $\mathcal{L} = g \circ C$ of the cost function $C(\cdot)$ with a monotonically decreasing function $g(\cdot)$. The likelihood in (3) then takes the form

$$p_t(\mathcal{O}_\tau|\boldsymbol{\theta}; \xi, \mathbf{x}_t) \propto \int \mathcal{L}(\tau) p(\tau|\boldsymbol{\theta}; \xi, \mathbf{x}_t) d\tau = \mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi}} [\mathcal{L}(\tau)]. \quad (5)$$

Nonparametric Bayesian MPC Instead of performing inference over policy parameters, the Bayesian MPC formulation can be used for inference over control input sequences. In this case, the inference variable $\boldsymbol{\theta}$ is defined as the sequence of open-loop controls: $\boldsymbol{\theta} \triangleq (\mathbf{u}_t, \mathbf{u}_{t+1}, \dots, \mathbf{u}_{t+H-1})$. This can be interpreted as a nonparametric version of Bayesian-MPC, as no assumption is made on the existence of a parametrized policy π . Although the relationship in (2) still holds, the likelihood function $p_t(\mathcal{O}_\tau|\boldsymbol{\theta}; \xi, \mathbf{x}_t)$ must then be re-defined as

$$p_t(\mathcal{O}_\tau|\boldsymbol{\theta}; \xi, \mathbf{x}_t) = \int p(\mathcal{O}_\tau|X_t, \boldsymbol{\theta}) p(X_t|\boldsymbol{\theta}; \xi, \mathbf{x}_t) dX_t \propto \mathbb{E}_{\hat{f}_\xi} [\mathcal{L}(X_t, \boldsymbol{\theta})], \quad (6)$$

where $p(X_t|\boldsymbol{\theta}; \xi, \mathbf{x}_t)$ is the probability of state trajectories, conditioned on decisions $\boldsymbol{\theta}$ and assumed dynamics model $\hat{f}_\xi = p_\xi(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t)$. In the discrete-time case, this can be written as the product of state transition probabilities along a trajectory:

$$p(X_t|\boldsymbol{\theta}; \xi, \mathbf{x}_t) = \prod_{h=0}^{H-1} p_\xi(\mathbf{x}_{t+h+1}|\mathbf{x}_{t+h}, \boldsymbol{\theta}_{t+h}). \quad (7)$$

Furthermore, if we assume a fixed current state ($\mathbf{x}_t = \text{const.}$), we can apply this model to trajectory optimization problems by performing approximate inference on the posterior $p(\boldsymbol{\theta}|\mathcal{O}_\tau; \xi)$, and taking the *maximum a posteriori* estimate, $\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}} p(\boldsymbol{\theta}|\mathcal{O}_\tau; \xi)$.

4 Inference for Bayesian MPC

Variational Inference Variational inference poses posterior estimation as an optimization task where a candidate distribution $q^*(\boldsymbol{\theta})$ within a distribution family $\mathcal{Q} = \{q(\boldsymbol{\theta})\}$ is chosen to best approximate the target distribution $p(\boldsymbol{\theta}|\mathcal{O}_\tau)$. This is typically obtained by minimizing the Kullback-Leibler (KL) divergence:

$$q^* = \arg \min_{q \in \mathcal{Q}} D_{KL}(q(\boldsymbol{\theta})||p(\boldsymbol{\theta}|\mathcal{O}_\tau)). \quad (8)$$

The solution also maximizes the Evidence Lower Bound (ELBO), as expressed by the following objective (details in Appendix D):

$$q^* = \arg \min_{q \in \mathcal{Q}} -\mathbb{E}_q [\log p(\mathcal{O}_\tau|\boldsymbol{\theta})] + D_{KL}(q(\boldsymbol{\theta}) || p(\boldsymbol{\theta})). \quad (9)$$

This optimization seeks to maximize the log-likelihood of the observations with the first term while penalizing for differences between the target and the prior with the second term. For a high-capacity model space \mathcal{Q} that includes the target distribution, the second term becomes increasingly small. Selecting a model space \mathcal{Q} with both high capacity and computational efficiency is critical to variational inference, but remains a challenging problem.

Stein Variational Gradient Descent In order to circumvent the challenge of determining an appropriate \mathcal{Q} , while also addressing (8), we develop an algorithm based on Stein variational gradient descent (SVGD) for Bayesian inference. The nonparametric nature of SVGD is advantageous as it removes the need for assumptions on restricted parametric families for q . This approach approximates a posterior $p(\boldsymbol{\theta}|x)$ with a set of particles $\{\boldsymbol{\theta}^i\}_{i=1}^m$, $\boldsymbol{\theta}^i \in \mathbb{R}^p$. The particles are iteratively updated according to

$$\boldsymbol{\theta}^i \leftarrow \boldsymbol{\theta}^i + \epsilon \boldsymbol{\phi}^*(\boldsymbol{\theta}^i) \quad (10)$$

given a step-size ϵ . The function $\boldsymbol{\phi}^*(\cdot)$ lies in the unit-ball of an \mathbb{R}^p -valued reproducing kernel Hilbert space (RKHS) of the form $\mathcal{H} = \mathcal{H}_0 \times \dots \mathcal{H}_0$, where \mathcal{H}_0 is a scalar-valued RKHS with kernel $k(\boldsymbol{\theta}', \boldsymbol{\theta})$. It characterizes the optimal perturbation or velocity field (i.e. gradient direction) which maximally decreases the KL-divergence:

$$\boldsymbol{\phi}^* = \arg \max_{\boldsymbol{\phi} \in \mathcal{H}} \left\{ -\nabla_{\epsilon} D_{KL}(q_{[\epsilon \boldsymbol{\phi}]}) \parallel p(\boldsymbol{\theta}|x) \right\} \text{ s.t. } \|\boldsymbol{\phi}\|_{\mathcal{H}} \leq 1, \quad (11)$$

where $q_{[\epsilon \boldsymbol{\phi}]}$ indicates the particle distribution resulting from taking an update step $\boldsymbol{\theta} = \boldsymbol{\theta} + \epsilon \boldsymbol{\phi}(\boldsymbol{\theta})$. This has been shown to yield a closed-form solution [16] which can be interpreted as a functional gradient in RKHS and approximated with the set of particles:

$$\hat{\boldsymbol{\phi}}^*(\boldsymbol{\theta}) = \frac{1}{m} \sum_{j=1}^m \left[k(\boldsymbol{\theta}^j, \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}^j} \log p(\boldsymbol{\theta}^j|x) + \nabla_{\boldsymbol{\theta}^j} k(\boldsymbol{\theta}^j, \boldsymbol{\theta}) \right]. \quad (12)$$

Eq. 12 has two terms that control different aspects of the algorithm. The first term is essentially a scaled gradient of the log-likelihood over the posterior's particle approximation. The second term is known as the *repulsive force*. Intuitively, it pushes particles apart when they get too close to each other and prevents them from collapsing into a single mode. This allows the method to approximate complex, possibly multi-modal posteriors in MPC. When there is only a single particle, the method reduces to a standard optimization of the log-likelihood or a MAP estimate of the posterior as the repulsive force term vanishes, i.e. $\nabla_{\boldsymbol{\theta}} k(\boldsymbol{\theta}, \boldsymbol{\theta}) = 0$. SVGD's optimization structure empirically provides better particle efficiency than other popular sampling procedures, such as Markov Chain Monte Carlo [29].

5 Stein Variational MPC

In this section we present our novel method for Stein inference, specifically designed around MPC requirements. The full algorithm can be found in Appendix B.

Posterior Sequential Updates As described in Section 3, the Bayesian interpretation of MPC seeks to find the posterior distribution over decision parameters at time t . Recalling Eq. 2:

$$p_t(\boldsymbol{\theta}|\mathcal{O}_\tau; \xi, \mathbf{x}_t) = \frac{p_t(\mathcal{O}_\tau|\boldsymbol{\theta}; \xi, \mathbf{x}_t) \tilde{q}_t(\boldsymbol{\theta}; \mathbf{x}_t)}{\int p_t(\mathcal{O}_\tau|\boldsymbol{\theta}^j; \xi, \mathbf{x}_t) \tilde{q}_t(\boldsymbol{\theta}^j; \mathbf{x}_t) d\boldsymbol{\theta}} \quad (13)$$

with a prior $\tilde{q}_t(\boldsymbol{\theta}; \mathbf{x}_t)$. Our approach approximates the posterior over decision parameters using a *weighted* set of particles $\{\boldsymbol{\theta}^i\}_{i=1}^m$, where the proposal q is defined as the empirical distribution $q(\boldsymbol{\theta}) = \sum_{i=1}^m w^i \delta(\boldsymbol{\theta}^i)$ with weights evaluated according to:

$$w^i = \frac{p_t(\mathcal{O}_\tau|\boldsymbol{\theta}^i; \xi, \mathbf{x}_t) \tilde{q}_t(\boldsymbol{\theta}^i; \mathbf{x}_t)}{\sum_{j=1}^m p_t(\mathcal{O}_\tau|\boldsymbol{\theta}^j; \xi, \mathbf{x}_t) \tilde{q}_t(\boldsymbol{\theta}^j; \mathbf{x}_t)} \quad (14)$$

such that $\sum_{i=1}^m w^i = 1$. Following the procedure outlined in (10)-(12), an SVGD update can be computed for the individual particles by computing the functional gradient

$$\hat{\boldsymbol{\phi}}^*(\boldsymbol{\theta}^i) = \frac{1}{m} \sum_{j=1}^m \left[k(\boldsymbol{\theta}^j, \boldsymbol{\theta}^i) \nabla_{\boldsymbol{\theta}^j} \log p_t(\boldsymbol{\theta}^j|\mathcal{O}_\tau; \xi, \mathbf{x}_t) + \nabla_{\boldsymbol{\theta}^j} k(\boldsymbol{\theta}^j, \boldsymbol{\theta}^i) \right] \quad (15)$$

and performing the gradient step: $\theta^i \leftarrow \theta^i + \epsilon \hat{\phi}^*(\theta^i)$. The evaluation of (15) requires computation of the log-posterior gradient, which can be written as the sum of the gradients of both the log-prior and log-likelihood:

$$\nabla_{\theta^i} \log p_t(\theta^i | \mathcal{O}_\tau; \xi, \mathbf{x}_t) = \nabla_{\theta^i} \log p_t(\mathcal{O}_\tau | \theta^i; \xi, \mathbf{x}_t) + \nabla_{\theta^i} \log \tilde{q}_t(\theta^i; \mathbf{x}_t) \quad (16)$$

$$= \nabla_{\theta^i} \log \mathbb{E}_{\pi_{\theta^i, \hat{f}_\xi}} [\mathcal{L}(\tau)] + \nabla_{\theta^i} \log \tilde{q}_t(\theta^i; \mathbf{x}_t) \quad (17)$$

The first RHS term requires that we define the cost-likelihood function \mathcal{L} . The SV-MPC framework allows for different possible definitions for \mathcal{L} . We will consider two in particular:

- Exponentiated Utility (EU): $\mathcal{L}(\tau) = \exp(-\alpha C(X_t, U_t))$, where $\alpha > 0$ (18)

- Probability of Low Cost (PLC): $\mathcal{L}(\tau) = \mathbb{1}_{C \leq C_{\max}}(C(X_t, U_t))$ (19)

It is generally assumed that the cost function $C(X_t, U_t)$ is non-differentiable with respect to the decision parameter θ , and that resulting expectations are difficult to evaluate analytically. As such, the gradients can be estimated via Monte-Carlo sampling, where a set of N control and state trajectory samples are drawn from the policy and the modeled dynamics: $\{\tau^s\}_{s=1}^N \sim p(X_t, U_t | \theta; \xi, \mathbf{x}_t)$, where $\tau^s = (X_t^s, U_t^s)$. This leads to the following approximation:

$$\nabla_{\theta^i} \log \mathbb{E}_{\pi_{\theta^i, \hat{f}_\xi}} [\mathcal{L}(\tau)] = \frac{\mathbb{E}_{\pi_{\theta^i, \hat{f}_\xi}} [\mathcal{L}(\tau) \nabla_{\theta^i} \log \pi_{\theta^i}]}{\mathbb{E}_{\pi_{\theta^i, \hat{f}_\xi}} [\mathcal{L}(\tau)]} \approx \frac{\sum_{s=1}^N \mathcal{L}(\tau^s) \nabla_{\theta^i} \log \pi_{\theta^i}(U_t^s)}{\sum_{s=1}^N \mathcal{L}(\tau^s)} \quad (20)$$

For a single particle ($m = 1$), the full gradient in (15) reduces to: $\phi^*(\theta) = \nabla p_t(\theta | \mathcal{O}_\tau; \xi, \mathbf{x}_t)$. SVGD then produces a local MAP estimate of the posterior distribution over θ . As a result, the SV-MPC update step exhibits strong similarity to common MPC algorithms (such as MPPI [1] and CEM [30]) depending on the chosen likelihood function, exhibiting equivalence under parameter values and choice of prior. Discussion of these comparisons can be found in Appendix F.

A significant advantage of the SV-MPC formulation is the robustness to highly-peaked posterior distributions. If particles are initialized poorly, or if the target posterior changes significantly between time-steps, many particles may find themselves in regions of low-probability. Indeed, this may occur frequently for likelihoods $p(\mathcal{O}_\tau | \theta^i; \xi, \mathbf{x}_t)$ with exponentiated cost (see Appendix F, for example). However, the shared gradient terms in (15) allow these particles to overcome this degeneracy quickly, while avoiding collapse due to the repulsive term (the reader may refer to Figure 1. in [31] for an intuitive illustration of this phenomenon.) As a consequence, SV-MPC avoids the problem of particle depletion often encountered in Sequential Monte Carlo methods [32].

Kernels for trajectories High-dimensional inference problems pose significant challenges for SVGD as the repulsive force given by the derivative of the kernel with respect to the inputs diminishes as the dimensionality increases [33]. Inspired by probabilistic graphical models and the conditional independence assumptions encoded in Markov random fields, we tackle this issue by devising a kernel that factorizes a high-dimensional input into a sum of kernels defined over cliques of dimensions. This allows the exploitation of the Markov structure of the trajectories to address the curse of dimensionality. For example, assume that the posterior over the parameters θ satisfies the conditional independence relations encoded in a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with vertices \mathcal{V} , and edges \mathcal{E} such that $p(\theta) \propto \prod_{d \in \mathcal{E}} \psi_d(\theta_d) \prod_{(dt) \in \mathcal{E}} \psi_{dt}(\theta_d, \theta_t)$, where $\psi_d(\theta_d)$ and $\psi_{dt}(\theta_d, \theta_t)$ are unary and pairwise potential functions respectively. We define the kernel over parameters as,

$$k(\theta, \theta') = \sum_{d \in \mathcal{V}} k(\theta_d, \theta'_d) + \sum_{(d,t) \in \mathcal{E}} k(\theta_{(d,t)}, \theta'_{(d,t)}) \quad (21)$$

The kernel is a sum of positive semi-definite kernels, so the result is a valid reproducing kernel Hilbert space [34] but less sensitive to the curse of dimensionality. In this paper we adopt the smooth RBF kernel; $k(\theta, \theta') = \exp\{-\|\theta - \theta'\|_2^2/h\}$, where h is evaluated using the median heuristic on the set of particles: $h = \text{med}(\{\|\theta^i\|_2\}^2) / \log m$.

Action Selection The variational inference procedure results in an approximation of the posterior distribution over θ . Following this step, a decision must be made on θ_t , such that a control action can be generated from the resulting policy π_{θ_t} and executed on the real system. Here we outline two possible methods for choosing an appropriate θ_t .

We can first consider the relative probabilistic weight of the particles as an approximation to their posterior probabilities:

$$w_i = \frac{p(\mathcal{O}_\tau | \theta^i; \xi, \mathbf{x}_t) \tilde{q}_{t-1}(\theta^i)}{\sum_{j=1}^m p(\mathcal{O}_\tau | \theta^j; \xi, \mathbf{x}_t) \tilde{q}_{t-1}(\theta^j)} \quad (22)$$

$$\approx p(\theta^i | \mathcal{O}_\tau; \xi, \mathbf{x}_t). \quad (23)$$

One strategy to selecting θ_t is to pick the highest-weighted particle $\theta_t = \theta_{i^*}$, which corresponds to the approximate MAP solution:

$$i^* = \arg \max_i [w_i] \quad (24)$$

$$\approx \arg \max_i p(\theta^i | \mathcal{O}_\tau; \xi, \mathbf{x}_t). \quad (25)$$

An alternative approach is to randomly sample from the posterior distribution, which can be approximated by sampling from the set of particles according to their weight w_i . This can be performed by drawing from the categorical distribution over particle weights: $i^* \sim \text{Cat}(w_i)$.

Shifting the distribution The prior $\tilde{q}_t(\theta; \mathbf{x}_t)$ is defined as the marginalized transition over the approximate posterior distribution obtained from the previous iteration: $\tilde{q}_t(\theta; \mathbf{x}_t) = \int p_t(\theta | \theta_{t-1}; \mathbf{x}_t) q_{t-1}(\theta_{t-1}) d\theta_{t-1}$, with the transition probability $p_t(\theta | \theta_{t-1}; \mathbf{x}_t)$ including a dependence on the currently observed state \mathbf{x}_t for generality. This operation is akin to the prediction step commonly found in Bayes filtering and sequential Monte-Carlo methods, and can be interpreted as a probabilistic version of the shift operator defined in [2] which serves to bootstrap the previous MPC solution to initialize the current iteration. Further details can be found in Appendix C.

Non-parametric SV-MPC In the setting where inference must be performed over open-loop controls $\theta_t = U_t$, we can define the cost-likelihood as a function of state: $\mathcal{L} = \mathcal{L}(X_t)$. A cost on control can be replaced by defining an additional prior factor $p(\theta)$, which can be combined with the transition probability above. For further details, refer to Appendix G.

Trajectory optimization The variational inference framework defined by SV-MPC can be modified to accommodate general motion planning problems common to many robotics applications. This special case can be considered by (1) using the non-parametric formulation, (2) assuming a stationary posterior distribution (no shifting), and (3) defining a deterministic dynamics model \tilde{f} such that state trajectory probabilities can be represented as $p(X | \theta) = \delta(X - \tilde{f}(\theta, x_0))$. A prior over sequences $p(\theta)$ can be defined to encourage desired behavior such as smoothness [35]. Because individual particles are not guaranteed to produce a MAP estimate once the SVGD optimization has converged, they can be subsequently refined by applying stochastic gradient descent (SGD) using posterior gradients without kernelization. Selection of a feasible and optimal plan can be generated by simply picking the best particle. The resulting algorithm, SV-TrajOpt, is described in Algorithm 2, with the likelihood gradient derived in Appendix H.

6 Experiments

In the following section, apply SV-MPC to common robotics problems: navigation, manipulation and locomotion. We end with an example of the SV-TrajOpt algorithm applied to a planar motion-planning scenario. All control algorithms were implemented in PyTorch, with batched gradient computation across particles and parallel generation of rollouts using either simulated or analytic dynamical models. Additional experimental details and results can be found in the Appendix.

Planar Navigation We construct a 2D robot navigation task, where a holonomic point-robot must reach a target location while avoiding obstacles (Figure 2). Hitting an obstacle will cause the agent to "crash" and prevent any further movement. This added non-differentiability makes the problem particularly challenging, but also suitable for sampling-based control schemes. The system exhibits

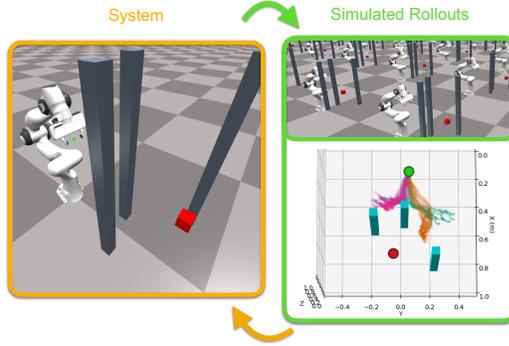


Figure 1: **7-DOF manipulation task.** The SV-MPC framework is capable of reasoning over multi-modal distributions of trajectories in high-dimensional spaces. Here, the controller iteratively explores the posterior over joint-velocities by simulating trajectories in parallel (green frame) in order to guide the system (orange frame). Each particle-generated distribution is shown by a unique coloring over the generated state trajectories, as seen from a top-down view of the workspace. The robot arm manages to reach the goal (red), while avoiding poor local minima.

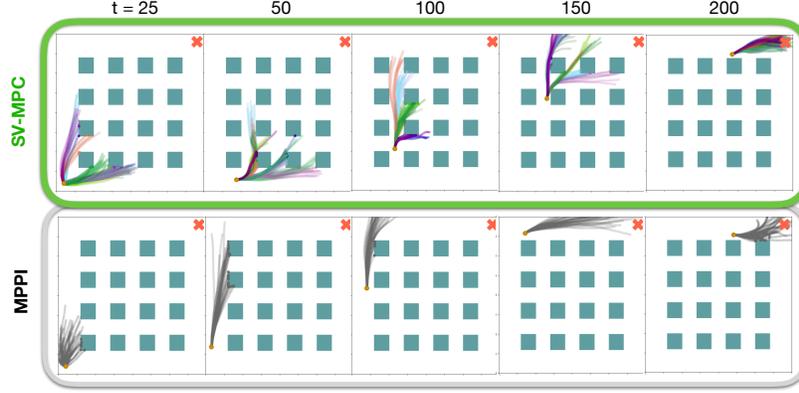


Figure 2: **Planar Navigation Task.** The robot (orange dot) attempts to reach the goal location (red cross) while avoiding obstacles. Each frame depicts the environment state at a particular time-step, along with the distributions of sampled state-trajectory rollouts generated by the MPC controllers using the modeled dynamics. Each trajectory color is associated with a single particle from SV-MPC. The multi-modal distribution of SV-MPC is able to explore passages between obstacles and find shorter paths to the goal.

stochastic dynamics, and is defined as a double-integrator model with additive Gaussian noise. The SV-MPC controller with exponentiated-utility (EU) is compared against MPPI (Figure 2) and CEM. We provide additional quantitative results in (Appendix J).

Manipulation We demonstrate SV-MPC on a 7-DOF reaching task (Figure 1). Velocity-based control commands are generated in the configuration space of a simulated robot manipulator, which must reach a stationary goal in its work-space. We leverage the GPU-accelerated Isaac-Gym library [36] for parallel computation of trajectory rollouts in simulation during MPC iterations.

A sampling-based SV-MPC controller is compared against MPPI and CEM for an open-loop, constant-covariance Gaussian control distribution: $\pi_{\theta} = \mathcal{N}(U_t; \theta, \Sigma)$. The obstacles and the target are placed in order to demonstrate a local-minima trap: given a finite H , the optimal control solution is to move left. However, the opening between the obstacles in this direction is too narrow for the robot to move through. To avoid getting stuck, a sampling-based control scheme must generate a sequence which will move the robot in the other direction. For a uni-modal distribution, this may have a very low-probability, and recovery will not occur. Increasing the sampling covariance would mitigate this, but would require a larger amount of samples to reduce variance.

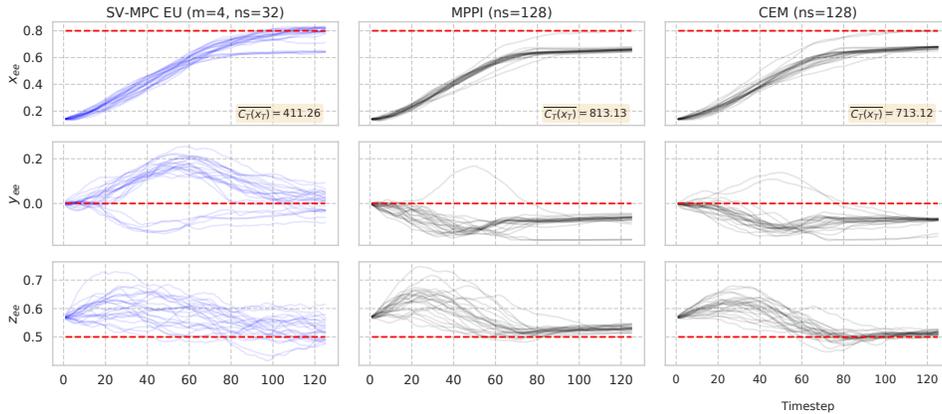


Figure 3: **Manipulation.** Examples of end-effector Cartesian trajectories resulting from application of different MPC algorithms on the Franka reaching experiment. The relative distance to the fixed target location is plotted over the length of each episode. The dashed red line indicates the coordinates of the target. The sample-averaged terminal cost for the final state $C_T(\mathbf{x}_T)$ is evaluated over the 24 independent trials. With four particles ($m = 4$), SV-MPC with Exponentiated utility likelihood manages to avoid bad local minima, despite higher-variance gradients due to fewer samples used to evaluate gradients ($ns = 32$ vs. $ns = 128$).

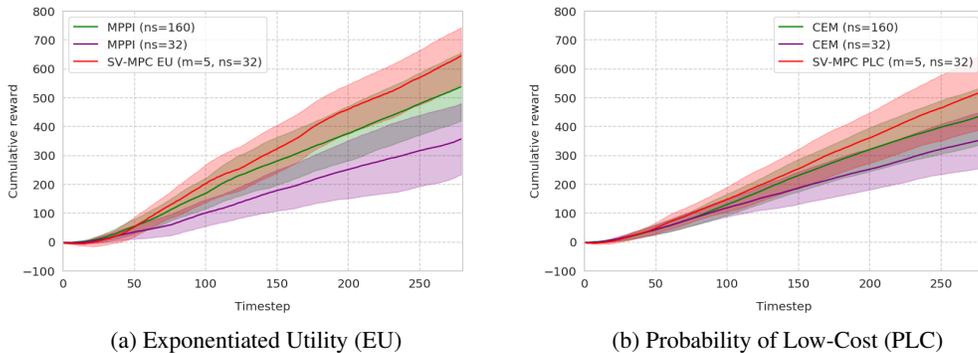


Figure 4: **Stochastic HalfCheetah**. Comparisons of cumulative-reward distributions over 16 independent trials, with mean and standard deviations shown. SV-MPC is capable of finding high-reward trajectories, using the same total amount of samples as MPPI and CEM.

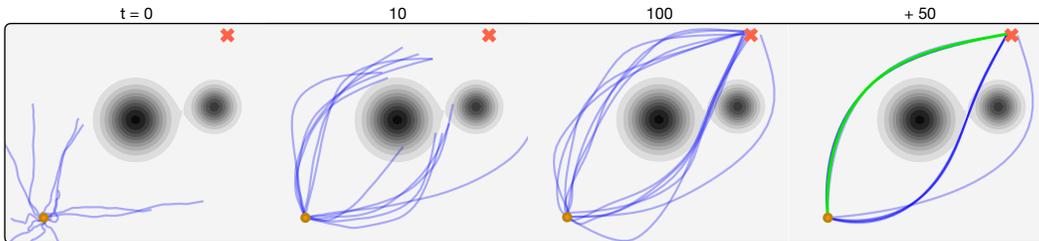


Figure 5: **Motion Planning**. SV-TrajOpt is applied to a velocity-controlled planar robot with fully-differentiable costs. Each blue state-trajectory results from a single particle control-sequence. Particles are randomly initialized from the prior ($t = 0$), and are optimized until convergence ($t = 10, 100$). Independent local MAP approximations are generated after 50 SGD iterations, with the lowest-cost particle shown in green.

Examples of trajectory executions are shown in Figure 3. The total number of generated control samples is held constant across algorithms. Both MPPI and CEM tend towards the sub-optimal local minimum, leading the robot to get stuck between two obstacles. Using a particle-based representation of the posterior, SV-MPC can resolve multiple optima simultaneously, switching to lower-cost modes around obstacles and successfully reaching the goal.

Stochastic Half-Cheetah To test our approach on a complex nonlinear system with discontinuous dynamics, we consider an environment common in many Reinforcement Learning benchmarks: the Half-Cheetah [37]. We use a stochastic version of the dynamics with additive noise in the control space. The cumulative rewards over multiple trials are plotted in Figure 4.

Motion Planning We use the SV-TrajOpt algorithm to infer a distribution over optimal control sequences for a planar motion planning problem on a point robot. The robot must find a velocity-based control sequence which results in a low-cost, feasible path around a set of obstacles to reach the goal. The occupancy map is fully-differentiable, allowing a gradient on obstacle cost to be computed numerically. A ‘smoothness’ prior is defined over velocities as a multi-variate Gaussian with a tri-diagonal precision matrix, which favors low-acceleration trajectories.

7 Conclusion

A novel formulation for Bayesian model predictive control is presented, where inference is performed directly over the control parameters and inputs. An algorithm for approximate inference is then proposed, where the posterior is represented as a set of particles and is optimized via SVGD. In contrast to pure Monte-Carlo sampling methods, gradient-based information can be exploited to improve particle efficiency, where many computationally-intensive operations can be run in parallel using effective GPU implementation. The flexibility of the approach can accommodate different cost transformations to modulate risk-seeking behavior, and can naturally be extended to trajectory optimization problems. We compare against common MPC baselines, demonstrating improved performance on a variety of control tasks.

References

- [1] 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*, 34(6):1603–1622, 2018.
- [2] N. Wagener, C.-A. Cheng, J. Sacks, and B. Boots. An online learning approach to model predictive control. *arXiv preprint arXiv:1902.08967*, 2019.
- [3] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Aggressive driving with model predictive path integral control. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1433–1440. IEEE, 2016.
- [4] M. Kasaei, N. Lau, and A. Pereira. A robust closed-loop biped locomotion planner based on time varying model predictive control. *arXiv preprint arXiv:1909.06873*, 2019.
- [5] V. Kumar, E. Todorov, and S. Levine. Optimal control with learned local models: Application to dexterous manipulation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 378–383. IEEE, 2016.
- [6] H. J. Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11):P11011, 2005.
- [7] V. Gómez, S. Thijssen, A. Symington, S. Hailes, and H. J. Kappen. Real-time stochastic optimal control for multi-agent quadrotor systems. In *Twenty-Sixth International Conference on Automated Planning and Scheduling*, 2016.
- [8] W. Wiegierinck, B. v. d. Broek, and H. Kappen. Stochastic optimal control in continuous space-time multi-agent systems. *arXiv preprint arXiv:1206.6866*, 2012.
- [9] G. Williams, B. Goldfain, P. Drews, K. Saigol, J. M. Rehg, and E. A. Theodorou. Robust sampling based model predictive control with sparse objective information. In *Robotics: Science and Systems*, 2018.
- [10] H. J. Kappen. Linear theory for control of nonlinear stochastic systems. *Physical review letters*, 95(20):200201, 2005.
- [11] K. Rawlik, M. Toussaint, and S. Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.
- [12] S. Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- [13] M. I. Jordan, Z. Ghahramani, and et al. An introduction to variational methods for graphical models. In *Machine Learning*, pages 183–233. MIT Press, 1999.
- [14] M. J. Beal. Variational algorithms for approximate bayesian inference. Technical report, University College London, 2003.
- [15] M. Wainwright and M. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.
- [16] Q. Liu and D. Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2378–2386. Curran Associates, Inc., 2016.
- [17] E. Todorov. General duality between optimal control and estimation. In *2008 47th IEEE Conference on Decision and Control*, pages 4286–4292. IEEE, 2008.
- [18] K. Rawlik, M. Toussaint, and S. Vijayakumar. Approximate inference and stochastic optimal control. *arXiv preprint arXiv:1009.3958*, 2010.

- [19] M. Toussaint. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pages 1049–1056. ACM, 2009.
- [20] H. J. Kappen, V. Gómez, and M. Opper. Optimal control as a graphical model inference problem. *Machine learning*, 87(2):159–182, 2012.
- [21] M. Toussaint and A. Storkey. Probabilistic inference for solving discrete and continuous state markov decision processes. In *Proceedings of the 23rd international conference on Machine learning*, pages 945–952, 2006.
- [22] M. Okada and T. Taniguchi. Variational inference mpc for bayesian model-based reinforcement learning. In *Conference on Robotics Learning (CoRL)*, 2019.
- [23] J. Watson, H. Abdulsamad, and J. Peters. Stochastic optimal control as approximate input inference. In *Conference on Robot Learning*, pages 697–716, 2020.
- [24] C. M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [25] D. Ormoneit and V. Tresp. Averaging, maximum penalized likelihood and bayesian estimation for improving gaussian mixture probability density estimates. *IEEE Transactions on Neural Networks*, 9(4):639–650, 1998.
- [26] E. Theodorou, J. Buchli, and S. Schaal. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11:3137–3181, 2010.
- [27] E. A. Theodorou and E. Todorov. Relative entropy and free energy dualities: Connections to path integral and kl control. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 1466–1473. IEEE, 2012.
- [28] E. A. Theodorou. Nonlinear stochastic control and information theoretic dualities: Connections, interdependencies and thermodynamic interpretations. *Entropy*, 17(5):3352–3375, 2015.
- [29] W. Y. Chen, A. Barp, F.-X. Briol, J. Gorham, M. Girolami, L. Mackey, and C. J. Oates. Stein point markov chain monte carlo. In *ICML*, 2019.
- [30] Z. I. Botev, D. P. Kroese, R. Y. Rubinstein, and P. L’Ecuyer. The cross-entropy method for optimization. In *Handbook of statistics*, volume 31, pages 35–59. Elsevier, 2013.
- [31] Q. Liu. Stein variational gradient descent: Theory and applications.
- [32] N. Kantas, J. Maciejowski, and A. Lecchini-Visintini. Sequential monte carlo for model predictive control. In *Nonlinear model predictive control*, pages 263–273. Springer, 2009.
- [33] J. Zhuo, C. Liu, J. Shi, J. Zhu, N. Chen, and B. Zhang. Message passing Stein variational gradient descent. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 6018–6027, Stockholm, Sweden, 10–15 Jul 2018. PMLR.
- [34] B. Schölkopf and A. Smola. *Learning with Kernels - Support Vector Machines, Regularization, Optimization and Beyond*. MIT Press, Cambridge, MA, 2001.
- [35] 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*, 37(11):1319–1340, 2018.
- [36] M. Macklin, K. Erleben, M. Müller, N. Chentanez, S. Jeschke, and V. Makoviychuk. Non-smooth newton methods for deformable multi-body dynamics. *ACM Transactions on Graphics (TOG)*, 38(5):1–20, 2019.
- [37] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [38] S. S. Varadhan. *Large deviations and applications*. SIAM, 1984.
- [39] C. Hartmann, L. Richter, C. Schütte, and W. Zhang. Variational characterization of free energy: Theory and algorithms. *Entropy*, 19(11):626, 2017.

A Notation

The table below summarizes notation used throughout the paper.

Table 1: Notations and definitions

Notation	Description
\mathbf{x}_t	System state at time t
X_t	State trajectory: sequence of length $H + 1$, from time t
\mathbf{u}_t	Control input at time t
U_t	Control trajectory: sequence of length H , from time t
τ	State-control trajectory tuple, $\tau = (X_t, U_t)$
$\hat{f}_\xi(\cdot, \cdot)$	Modeled system dynamics, parameterized by ξ
$C(\cdot, \cdot)$	Total cost for horizon H
$\pi_{\theta_t}(\cdot)$	Policy at time t parameterized by θ_t
$\boldsymbol{\theta}_t$	Sequence of policy parameters <i>or</i> control inputs
$\boldsymbol{\pi}_{\theta_t}(\cdot)$	Sequence of control policies
$\mathcal{L}(\cdot)$	Cost-likelihood function
$q(\cdot)$	Approximate posterior distribution
$\phi(\cdot)$	Perturbation direction or velocity field that decreases the KL divergence between the particle distribution and target distribution
$\hat{\phi}(\cdot)$	Approximated perturbation direction over particles
$k(\cdot, \cdot)$	Kernel specifying a reproducing kernel Hilbert space

B Algorithm

Algorithm 1: Stein Variational MPC

Input: Initial state \mathbf{x}_0 , dynamics \hat{f}_ξ , policy $\boldsymbol{\pi}$, cost-likelihood \mathcal{L} , prior $p_0(\boldsymbol{\theta}; \cdot)$, kernel k

Initialize $\tilde{q}_0(\boldsymbol{\theta}) = p_0(\boldsymbol{\theta}; \mathbf{x}_0)$

Sample $\{\boldsymbol{\theta}^i\}_{i=1}^m \sim \tilde{q}_0(\boldsymbol{\theta})$

for $t = 0, 1, \dots, T - 1$ **do**

for $i = 1, 2, \dots, m$ **do in parallel**

$$\nabla_{\boldsymbol{\theta}^i} \log p_t(\boldsymbol{\theta}^i | \mathcal{O}_\tau; \xi, \mathbf{x}_t) = \nabla_{\boldsymbol{\theta}^i} \log \mathbb{E}_{\boldsymbol{\pi}_{\boldsymbol{\theta}^i}, \hat{f}_\xi} [\mathcal{L}(\tau)] + \nabla_{\boldsymbol{\theta}^i} \log \tilde{q}_t(\boldsymbol{\theta}^i)$$

end

for $i = 1, 2, \dots, m$ **do**

$$\Delta \boldsymbol{\theta}^i \leftarrow \frac{1}{m} \sum_{j=1}^m k(\boldsymbol{\theta}^j, \boldsymbol{\theta}^i) \nabla_{\boldsymbol{\theta}^j} \log p_t(\boldsymbol{\theta}^j | \mathcal{O}_\tau; \xi, \mathbf{x}_t) + \nabla_{\boldsymbol{\theta}^j} k(\boldsymbol{\theta}^j, \boldsymbol{\theta}^i)$$

$$\boldsymbol{\theta}^i \leftarrow \boldsymbol{\theta}^i + \epsilon \Delta \boldsymbol{\theta}^i$$

end

$$w^i \leftarrow p(\mathcal{O}_\tau | \boldsymbol{\theta}^i; \mathbf{x}_t) \tilde{q}_t(\boldsymbol{\theta}^i)$$

$$w^i \leftarrow \frac{w^i}{\sum_{j=1}^m w^j}$$

 Pick $\boldsymbol{\theta}^*$ (using (22) or (24), for example)

 Sample control: $\mathbf{u}_t \sim \boldsymbol{\pi}_{\boldsymbol{\theta}^*}(\mathbf{x}_t)$

 Sample true dynamics: $\mathbf{x}_{t+1} \sim f(\mathbf{x}_t, \mathbf{u}_t)$

 Shift particles: $\tilde{\boldsymbol{\theta}}^i = \Phi(\boldsymbol{\theta}^i)$

 Update prior: $\tilde{q}_{t+1}(\boldsymbol{\theta}; \mathbf{x}_{t+1}) = \sum_{i=1}^m w^i p_t(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}}^i; \mathbf{x}_{t+1})$

end

Algorithm 2: SV-TrajOpt

Input: Initial state \mathbf{x}_0 , dynamics f , cost function C ,
prior $p_0(\boldsymbol{\theta})$, kernel k , term. condition $\text{Done}(\cdot)$

Sample $\{\boldsymbol{\theta}^i\}_{i=1}^m \sim p_0(\boldsymbol{\theta})$

Set $\Delta\boldsymbol{\theta}^i = \text{Inf} \forall i \in 1 : m$

while $\text{Done}(\{\Delta\boldsymbol{\theta}^i\})$ is *False* **do**

for $i = 1, 2, \dots, m$ **do in parallel**

 Forward dynamics : $X^i = X(f, \boldsymbol{\theta}^i, \mathbf{x}_0)$

$\nabla_{\boldsymbol{\theta}^i} \log p(\boldsymbol{\theta}^i | \mathcal{O}_\tau; \xi) = -\alpha \nabla_{\boldsymbol{\theta}^i} C(X^i) + \nabla_{\boldsymbol{\theta}^i} \log p_0(\boldsymbol{\theta}^i)$

end

for $i = 1, 2, \dots, m$ **do**

$\Delta\boldsymbol{\theta}^i \leftarrow \frac{1}{m} \sum_{j=1}^m k(\boldsymbol{\theta}^j, \boldsymbol{\theta}^i) \nabla_{\boldsymbol{\theta}^j} \log p(\boldsymbol{\theta}^j | \mathcal{O}_\tau; \xi) + \nabla_{\boldsymbol{\theta}^j} k(\boldsymbol{\theta}^j, \boldsymbol{\theta}^i)$

$\boldsymbol{\theta}^i \leftarrow \boldsymbol{\theta}^i + \epsilon \Delta\boldsymbol{\theta}^i$

end

end

(Optional: SGD refinement)

for $iter = 1, \dots, N$ **do**

$\Delta\boldsymbol{\theta}^i \leftarrow \nabla_{\boldsymbol{\theta}^i} \log p(\boldsymbol{\theta}^i | \mathcal{O}_\tau; \xi) \quad \forall i \in 1 : m$

$\boldsymbol{\theta}^i \leftarrow \epsilon_r \Delta\boldsymbol{\theta}^i \quad \forall i \in 1 : m$

end

$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta}^i} \log p(\boldsymbol{\theta}^i | \mathcal{O}_\tau; \xi)$

C Shift operation

As described in 5, the proposal distribution q is propagated after each round of MPC by marginalizing over the previous solution:

$$\tilde{q}_{t+1}(\boldsymbol{\theta}; \mathbf{x}_{t+1}) = \int p_t(\boldsymbol{\theta} | \boldsymbol{\theta}_t; \mathbf{x}_{t+1}) q_t(\boldsymbol{\theta}_t) d\boldsymbol{\theta}_t \quad (26)$$

given a transition probability, $p_t(\boldsymbol{\theta} | \boldsymbol{\theta}_t; \mathbf{x}_t)$. This operation serves to approximate the prior at the new iteration. Given the empirical distribution q_t , we can simplify the expression above:

$$\tilde{q}_{t+1}(\boldsymbol{\theta}; \mathbf{x}_{t+1}) = \sum_{i=1}^m \int p_t(\boldsymbol{\theta} | \boldsymbol{\theta}_t; \mathbf{x}_{t+1}) w^i \delta_{\boldsymbol{\theta}_t}(\boldsymbol{\theta}_t) d\boldsymbol{\theta}_t = \sum_{i=1}^m w^i p_t(\boldsymbol{\theta} | \boldsymbol{\theta}^i; \mathbf{x}_{t+1}) \quad (27)$$

resulting in a mixture of conditional probabilities.

In many open-loop MPC implementations [1, 2], it is assumed that the solution does not change significantly between each round, given an accurate dynamics model and sufficiently high controller frequency to resolve the dynamics and possible perturbations. This motivates a common heuristic used to reduce the computational burden between timesteps, which is to shift the control distribution $\boldsymbol{\pi}_{\boldsymbol{\theta}_t}$, forward-in-time by one step. That is, given an initial parameter sequence $\boldsymbol{\theta}_t = (\theta_t, \theta_{t+1}, \dots, \theta_{t+H-1})$:

$$\tilde{\boldsymbol{\theta}}_{t+1} = \Phi(\boldsymbol{\theta}_t) \quad (28)$$

$$= (\theta_{t+1}, \theta_{t+2}, \dots, \theta_{t+H-1}, \tilde{\theta}_{t+H-1}) \quad (29)$$

where the new parameter $\tilde{\theta}_{t+H-1}$ is chosen to reflect the expected final action. In the implementation, we adopt a similar heuristic, where the empirical distribution is first shifted deterministically according to:

$$\tilde{\boldsymbol{\theta}}^i = \Phi(\boldsymbol{\theta}^i) \quad \forall i \in 1 : m \quad (30)$$

and setting the resulting distribution as $\tilde{q}_t(\boldsymbol{\theta}) = q_t(\boldsymbol{\theta})$. The shifted particles, $\tilde{\boldsymbol{\theta}}^i$, are then used in the following iteration to approximate the posterior, and the prior is updated according to (27). The shift operation can then be summarized by the following two sub-steps:

1. Shift particles: $\tilde{\boldsymbol{\theta}}^i = \Phi(\boldsymbol{\theta}^i)$
2. Update prior: $\tilde{q}_{t+1}(\boldsymbol{\theta}; \mathbf{x}_{t+1}) = \sum_{i=1}^m w^i p_t(\boldsymbol{\theta}|\tilde{\boldsymbol{\theta}}^i; \mathbf{x}_{t+1})$

D Variational Inference

In variational inference, a target distribution $p(\boldsymbol{\theta}|\mathcal{O}_\tau)$ is approximated by a candidate $q^*(\boldsymbol{\theta})$, belonging to a specified family of distributions $\mathcal{Q} = \{q(\boldsymbol{\theta})\}$, by minimizing the KL-divergence:

$$q^* = \arg \min_{q \in \mathcal{Q}} D_{KL}(q(\boldsymbol{\theta})||p(\boldsymbol{\theta}|\mathcal{O}_\tau)) \quad (31)$$

The optimal distribution q^* is also the solution to the following objective:

$$q^* = \arg \max_{q \in \mathcal{Q}} \mathbb{E}_q \left[\log p(\mathcal{O}_\tau|\boldsymbol{\theta}) \right] - D_{KL}(q(\boldsymbol{\theta}) || p(\boldsymbol{\theta})) \quad (32)$$

Proof 1:

$$q^* = \arg \min_{q \in \mathcal{Q}} D_{KL}(q(\boldsymbol{\theta})||p(\boldsymbol{\theta}|\mathcal{O}_\tau)) \quad (33)$$

$$= \arg \min_{q \in \mathcal{Q}} \int \log q(\boldsymbol{\theta}) dq(\boldsymbol{\theta}) - \int \log p(\boldsymbol{\theta}|\mathcal{O}_\tau) dq(\boldsymbol{\theta}) \quad (34)$$

$$= \arg \min_{q \in \mathcal{Q}} \int \log q(\boldsymbol{\theta}) dq(\boldsymbol{\theta}) - \int \log p(\mathcal{O}_\tau|\boldsymbol{\theta}) dq(\boldsymbol{\theta}) - \int \log p(\boldsymbol{\theta}) dq(\boldsymbol{\theta}) + \log Z \quad (35)$$

$$= \arg \min_{q \in \mathcal{Q}} \int \log q(\boldsymbol{\theta}) dq(\boldsymbol{\theta}) - \int \log p(\mathcal{O}_\tau|\boldsymbol{\theta}) dq(\boldsymbol{\theta}) - \int \log p(\boldsymbol{\theta}) dq(\boldsymbol{\theta}) \quad (36)$$

$$= \arg \min_{q \in \mathcal{Q}} \mathbb{E}_q \left[\log q(\boldsymbol{\theta}) \right] - \mathbb{E}_q \left[\log p(\mathcal{O}_\tau|\boldsymbol{\theta}) + \log p(\boldsymbol{\theta}) \right] \quad (37)$$

$$= \arg \max_{q \in \mathcal{Q}} \mathbb{E}_q \left[\log p(\mathcal{O}_\tau|\boldsymbol{\theta}) \right] - D_{KL}(q(\boldsymbol{\theta}) || p(\boldsymbol{\theta})) \quad (38)$$

where Z is the normalizing constant. The functional in the objective is also known as the variational free energy $F_q(\boldsymbol{\theta}) := \mathbb{E}_q \left[\log p(\mathcal{O}_\tau|\boldsymbol{\theta}) \right] - D_{KL}(q(\boldsymbol{\theta}) || p(\boldsymbol{\theta}))$.

An alternative derivation, which is relevant to the discussion in Appendix E, can be found by deriving the lower-bound to the log marginal-likelihood:

Proof 2:

$$\log \mathbb{E}_p \left[p(\mathcal{O}_\tau|\boldsymbol{\theta}) \right] = \log \int p(\mathcal{O}_\tau|\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (39)$$

$$= \log \int p(\mathcal{O}_\tau|\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (40)$$

$$= \log \int p(\mathcal{O}_\tau|\boldsymbol{\theta}) \frac{p(\boldsymbol{\theta})}{q(\boldsymbol{\theta})} q(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (41)$$

$$= \log \mathbb{E}_q \left[p(\mathcal{O}_\tau|\boldsymbol{\theta}) \frac{p(\boldsymbol{\theta})}{q(\boldsymbol{\theta})} \right] \quad (42)$$

$$\geq \mathbb{E}_q \left[\log \frac{p(\mathcal{O}_\tau|\boldsymbol{\theta}) p(\boldsymbol{\theta})}{q(\boldsymbol{\theta})} \right] \quad (43)$$

$$= \mathbb{E}_q \left[\log p(\mathcal{O}_\tau|\boldsymbol{\theta}) \right] - \mathbb{E}_q \left[\log \frac{p(\boldsymbol{\theta})}{q(\boldsymbol{\theta})} \right] \quad (44)$$

$$= \mathbb{E}_q \left[\log p(\mathcal{O}_\tau|\boldsymbol{\theta}) \right] - D_{KL}(q(\boldsymbol{\theta}) || p(\boldsymbol{\theta})) \quad (45)$$

where (43) results from log-concavity and application of Jensen's inequality. Equality is obtained when q matches the posterior probability. In other words, provided that $p(\boldsymbol{\theta}|\mathcal{O}_\tau) \in \mathcal{Q}$, we have:

$$q^*(\boldsymbol{\theta}) = \frac{p(\mathcal{O}_\tau|\boldsymbol{\theta}) p(\boldsymbol{\theta})}{\int p(\mathcal{O}_\tau|\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}} = p(\boldsymbol{\theta}|\mathcal{O}_\tau) \quad (46)$$

As such, we can define the minimization:

$$-\log \mathbb{E}_p [p(\mathcal{O}_\tau|\boldsymbol{\theta})] = \min_{q \in \mathcal{Q}} -\mathbb{E}_q [\log p(\mathcal{O}_\tau|\boldsymbol{\theta})] + D_{KL}(q(\boldsymbol{\theta}) || p(\boldsymbol{\theta})) \quad (47)$$

E Connection to Path Integral Control

The equation derived in (47) is well-known in statistical thermodynamics, where the random variable under consideration is the energy $C(x) : \Omega_x \mapsto \mathbb{R}$, a non-negative real-valued measurable property, and $x \in \mathbb{R}^n$ is the state of the system. The Helmholtz free energy of C , with respect to probability density p , is defined as the function :

$$-\frac{1}{\alpha} \log \mathbb{E}_p [\exp(-\alpha C(x))], \quad (48)$$

where $\alpha > 0$ is the (inverse) temperature. The corresponding variational inequality, known as the Donsker-Varadhan principle [38, 39], relates the free energy as the Legendre transform of the entropy:

$$-\frac{1}{\alpha} \log \mathbb{E}_p [\exp(-\alpha C(x))] = \min_q -\frac{1}{\alpha} \mathbb{E}_q [\log \exp(-\alpha C(x))] + \frac{1}{\alpha} D_{KL}(q(x) || p(x)) \quad (49)$$

$$= \min_q \mathbb{E}_q [C(x)] + \frac{1}{\alpha} D_{KL}(q(x) || p(x)) \quad (50)$$

Following the same derivation presented in (39)-(47), the solution is then found to be the Gibbs distribution:

$$q^* = \frac{\exp(-\alpha C(x)) p(x)}{\int \exp(-\alpha C(x)) p(x) dx}. \quad (51)$$

In the context of stochastic dynamics, the above is also applicable to random paths generated by a Markov diffusion process [39]. This can then be extended to optimal control, by addressing the KL-minimization problem between controlled and uncontrolled stochastic systems. The connection was developed and explored for continuous-time dynamics in previous work, such as [27, 28]. Here, the nonlinear-affine dynamics under consideration are subject to Brownian motion:

$$\mathbf{x}_{t+1} = \bar{f}(\mathbf{x}_t) + G(\mathbf{x}_t) \mathbf{u}_t + B(\mathbf{x}_t) \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(0, \Sigma) \quad (52)$$

$$= f(\mathbf{x}_t, \mathbf{u}_t) \quad (53)$$

for the discrete-time case. For a SOC expected-cost objective:

$$\min_{U_t} J(X_t) = \min_{U_t = (\mathbf{u}_0, \dots, \mathbf{u}_T)} \mathbb{E}_f \left[c_{\text{term}}(\mathbf{x}_T) + \sum_{t=0}^T c_t(\mathbf{x}_t) + \frac{1}{2} \mathbf{u}_t^\top R \mathbf{u}_t \right], \quad (54)$$

the optimal control distribution could be derived by a change of measure using Girsanov's theorem, which was shown to satisfy the Hamilton-Jacobi-Bellman (HJB) equations for optimality. To solve the HJB equations, an exponentiated form of the value function, the desirability-function, was defined: $\Psi(\mathbf{x}_t) = \exp(-\alpha V(\mathbf{x}_t))$. By application of the Feynman-Kac lemma, the transformed HJB partial differential equation can be solved as:

$$\Psi(\mathbf{x}_t) = \int \exp(-\alpha C(X_t)) p_f(X_t; \mathbf{x}_t) dX_t \quad (55)$$

$$= \mathbb{E}_{p_f} [\exp(-\alpha C(X_t))] \quad (56)$$

where p_f denotes the passive dynamics *i.e.* the probability density of trajectories X_t , resulting from the stochastic dynamics $f(\mathbf{x}_t, \mathbf{u}_t)$ with $\mathbf{u}_t = 0$. Under the assumption that $R = \frac{1}{\alpha}\Sigma^{-1}$, this can then be used to then evaluate the optimal control law:

$$\mathbf{u}_t^* = -\alpha\Sigma G(\mathbf{x}_t)\nabla_{\mathbf{x}}\mathbf{V}(\mathbf{x}_t) \quad (57)$$

$$= \Sigma G(\mathbf{x}_t)\nabla_{\mathbf{x}}\log\Psi(\mathbf{x}_t) \quad (58)$$

$$= \Sigma G(\mathbf{x}_t)\frac{\nabla_{\mathbf{x}}\Psi}{\Psi}. \quad (59)$$

For receding-horizon control, we can borrow from [1] and consider a discrete-time case of Path-Integral control for a nonlinear stochastic dynamical system:

$$\mathbf{x}_{t+1} = \bar{f}(\mathbf{x}_t, \mathbf{v}_t), \quad \mathbf{v}_t \sim p(\mathbf{v}_t|\mathbf{u}_t) \quad (60)$$

with nominal deterministic dynamics \bar{f} , commanded control input \mathbf{u}_t , and stochastic perturbations \mathbf{v}_t which are exhibited in the control input channel. The uncontrolled system is then realized when $\mathbf{u}_t = 0$, and is controlled otherwise. Given a sequence of perturbations: $V_t = (\mathbf{v}_t, \mathbf{v}_{t+1}, \dots, \mathbf{v}_{t+H-1})$, we can write the resulting state trajectory as $X_t = F(\mathbf{x}_t, V_t)$, where F performs consecutive application of the dynamics \bar{f} given \mathbf{x}_t and a sequence V_t . We can then consider probability distributions directly over V_t , with $p(V_t)$ for uncontrolled dynamics and $q(V_t|U_t)$ for the controlled system, with the sequence of control inputs given by $U_t = (\mathbf{u}_t, \mathbf{u}_{t+1}, \dots, \mathbf{u}_{t+H-1})$. A cost on state $C(X_t)$ is mapped to the random variable V_t by the convolution $S = C \circ F$. Similarly to equations (47) and (50) above, we then seek a solution to minimize the variational objective:

$$q^* = \arg \min_q D_{KL}(q || q^*(V_t)) \quad (61)$$

$$= \arg \min_q -\frac{1}{\alpha}\mathbb{E}_q \left[\log \exp(-\alpha S(V_t)) \right] + \frac{1}{\alpha}D_{KL}(q || p) \quad (62)$$

$$= \arg \min_q \mathbb{E}_q [S(V_t)] + \frac{1}{\alpha}D_{KL}(q || p) \quad (63)$$

where the optimal distribution is then known to be

$$q^*(V_t) = \frac{\exp(-\alpha S(V_t))p(V_t)}{\int \exp(-\alpha S(V_t))p(V_t)dV_t}. \quad (64)$$

Although we have the form of the optimal control distribution, an analytic solution is generally intractable due to the partition function. The algorithm for Model Predictive Path Integral Control (MPPI) [1, 3] proceeds by defining a surrogate cross-entropy minimization problem:

$$\min_{q \in \mathcal{Q}} D_{KL}(q^*(V_t) || q) \quad (65)$$

for a tractable family of distributions \mathcal{Q} , typically fixed-covariance Gaussians with mean parameters $\mu_t = U_t$ and an equivalent assumption on the controlled dynamics distribution.

In the SV-MPC framework, the original variational objective of Path Integral control can be addressed, where

$$\min_{q \in \mathcal{Q}} D_{KL}(q || q^*(V_t)). \quad (66)$$

We can apply the non-parametric Bayesian-MPC formulation (Section 3) to solve for $q^*(V_t)$ directly over V_t , setting the inference parameter to be $\boldsymbol{\theta} = V_t$ with a prior given by the passive dynamics $p(\boldsymbol{\theta}; \mathbf{x}_t) = p(V_t)$. The likelihood term $p(\mathcal{O}_\tau | \boldsymbol{\theta}; \xi, \mathbf{x}_t)$ can be derived by using the exponentiated-utility $\mathcal{L}(X_t) = \exp(-\alpha C(X_t))$ and defining a dirac measure over state trajectories to denote the probability density of X_t given $\boldsymbol{\theta}$:

$$p(\mathcal{O}_\tau | \boldsymbol{\theta}; \xi, \mathbf{x}_t) \propto \int \exp(-\alpha C(X_t)) \delta(X_t - F(\mathbf{x}_t, \boldsymbol{\theta})) dX_t \quad (67)$$

$$= \exp(-\alpha C(F(\mathbf{x}_t, \boldsymbol{\theta}))) \quad (68)$$

$$= \exp(-\alpha S(\boldsymbol{\theta})). \quad (69)$$

where we drop ξ -notation for the dynamics parameters for simplicity.

We then recover the Bayesian formulation:

$$q^*(\boldsymbol{\theta}) = \frac{p(\mathcal{O}_\tau | \boldsymbol{\theta}; \xi, \mathbf{x}_t) p(\boldsymbol{\theta}; \mathbf{x}_t)}{\int p(\mathcal{O}_\tau | \boldsymbol{\theta}; \xi, \mathbf{x}_t) p(\boldsymbol{\theta}; \mathbf{x}_t) d\boldsymbol{\theta}} \quad (70)$$

$$= \frac{\exp(-\alpha S(\boldsymbol{\theta})) p(\boldsymbol{\theta}; \mathbf{x}_t)}{\int \exp(-\alpha S(\boldsymbol{\theta})) p(\boldsymbol{\theta}; \mathbf{x}_t) d\boldsymbol{\theta}} \quad (71)$$

$$= p(\boldsymbol{\theta} | \mathcal{O}_\tau; \xi, \mathbf{x}_t) \quad (72)$$

Approximate inference on the posterior $p(\boldsymbol{\theta} | \mathcal{O}_\tau; \xi, \mathbf{x}_t)$ can then be performed using [Algorithm 1](#), where the optimal control action is approximated by:

$$\mathbf{u}_t^* = \Sigma G(\mathbf{x}_t) \frac{\nabla_{\mathbf{x}} \Psi}{\Psi} = \Sigma G(\mathbf{x}_t) \frac{\mathbb{E}_p \left[\exp(-\alpha S(V_t)) V_t \right]}{\mathbb{E}_p \left[\exp(-\alpha S(V_t)) \right]} \quad (73)$$

$$\approx \Sigma G(\mathbf{x}_t) \frac{\sum_{i=1}^m \exp(-\alpha S(\boldsymbol{\theta}_i)) \boldsymbol{\theta}_i}{\sum_{i=1}^m \exp(-\alpha S(\boldsymbol{\theta}_i))} \quad (74)$$

F Likelihood functions

Exponentiated Utility (EU) By picking the cost-likelihood to be: $\mathcal{L}(\tau) = \exp(-\alpha C(X_t, U_t))$, we obtain the likelihood function

$$\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\mathcal{L}(\tau) \right] = \mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \right]. \quad (75)$$

This is otherwise known as the Free-Energy of the cost function $C(\mathbf{x}_t, \mathbf{u}_t)$ [27], as well as the "soft maximum" or "risk-aware" loss. This yields the likelihood-gradient:

$$\nabla_{\boldsymbol{\theta}} \log \mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\mathcal{L}(\tau) \right] = \nabla_{\boldsymbol{\theta}} \log \mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \right] \quad (76)$$

$$= \frac{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}} \right]}{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \right]} \quad (77)$$

As $\alpha \rightarrow \infty$, regions of high-cost are assigned lower probability, making the distribution of resulting policies risk-averse. Conversely, as $\alpha \rightarrow 0$, high-cost regions have higher likelihood, making the policy distributions more risk-seeking.

With this form of the gradient, we can choose a control policy as a sequence of Gaussian distributions over open-loop controls with fixed covariance: given that $\pi_{\boldsymbol{\theta}_t} = (\pi_{\theta_t}, \pi_{\theta_{t+1}}, \dots, \pi_{\theta_{t+H-1}})$, we have the instantaneous policy $\pi_{\theta_t} = \mathcal{N}(\boldsymbol{\mu}_t, \Sigma_t)$, our decision parameter is thus $\boldsymbol{\theta}_t = \boldsymbol{\mu}_t = (\mu_t, \mu_{t+1}, \dots, \mu_{t+H-1})$, where $\mu_t \in \mathbb{R}^d$, $\Sigma_t \in \mathbb{R}^{d \times d}$. Considering the SV-MPC step in (15) for a single particle, we can derive the update for a parameter element $h \in (t, t+1, \dots, t+H-1)$:

$$\theta'_h = \theta_h + \epsilon \left(\nabla_{\theta_h} \log p_t(\mathcal{O}_\tau | \boldsymbol{\theta}; \xi, \mathbf{x}_t) + \nabla_{\theta_h} \log \tilde{q}_t(\boldsymbol{\theta}) \right) \quad (78)$$

$$= \theta_h + \epsilon \frac{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \nabla_{\theta_h} \log \pi_{\theta_h} \right]}{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \right]} + \epsilon \nabla_{\theta_h} \log \tilde{q}_t(\boldsymbol{\theta}) \quad (79)$$

$$= \theta_h + \epsilon \Sigma_t^{-1} \frac{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) (\mathbf{u}_h - \theta_h) \right]}{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \right]} + \epsilon \nabla_{\theta_h} \log \tilde{q}_t(\boldsymbol{\theta}) \quad (80)$$

$$= (I - \epsilon \Sigma_t^{-1}) \theta_h + \epsilon \Sigma_t^{-1} \frac{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \mathbf{u}_h \right]}{\mathbb{E}_{\pi_{\boldsymbol{\theta}, \hat{f}_\xi} \left[\exp(-\alpha C(X_t, U_t)) \right]} + \epsilon \nabla_{\theta_h} \log \tilde{q}_t(\boldsymbol{\theta}) \quad (81)$$

where the gradient of the log-prior will depend on the particular choice of transition probability, $p_t(\boldsymbol{\theta} | \tilde{\boldsymbol{\theta}}; \mathbf{x}_{t+1})$, used in the shift operation (see [Appendix C](#)). The update reduces to that found in

MPPI [1] if we consider the control distribution to be uncorrelated across dimensions (as is often done in practice): $\Sigma_t = \sigma_t^2 I$. Setting the step-size to $\epsilon = \sigma^2$, and assuming a uniform prior on controls, $p_t(\theta|\hat{\theta}; \mathbf{x}_{t+1}) = \mathcal{U}(\theta_{min}, \theta_{max})$:

$$\theta'_h = (I - \sigma_t^2 \sigma_t^{-2} I) \theta_h + \sigma_t^2 \sigma_t^{-2} I \frac{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\exp(-\alpha C(X_t, U_t)) \mathbf{u}_h]}]{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\exp(-\alpha C(X_t, U_t))]}]} + 0 \quad (82)$$

$$= \frac{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\exp(-\alpha C(X_t, U_t)) \mathbf{u}_h]}]{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\exp(-\alpha C(X_t, U_t))]}}, \quad (83)$$

recovering the MPPI update rule.

Probability of Low Cost (PLC) Similarly to [2], we can incorporate a threshold-utility to indicate preference for costs below a given threshold, using the indicator function: $\mathcal{L}(\tau) = \mathbb{1}_{C \leq C_{t, \max}}(C(X_t, U_t))$. The likelihood then takes the form

$$\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\mathcal{L}(\tau)]} = \mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\mathbb{1}_{C \leq C_{t, \max}}(C(X_t, U_t))]} \quad (84)$$

with the resulting gradient:

$$\nabla_{\theta} \log \mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\mathcal{L}(\tau)]} = \frac{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\mathbb{1}_{C \leq C_{t, \max}}(C(X_t, U_t)) \nabla_{\theta} \log \pi_{\theta}]}]{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\mathbb{1}_{C \leq C_{t, \max}}(C(X_t, U_t))]}]} \quad (85)$$

The threshold parameter $C_{t, \max}$ is set adaptively as the largest cost of the top member in the elite fraction of sampled trajectories. Using the same derivation for the case of Exponentiated Utility likelihood, the choice of a Gaussian policy with a threshold-utility reduces to the update rule for the Cross Entropy Method [30]:

$$\theta'_h = \frac{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\mathbb{1}_{C \leq C_{t, \max}}(C(X_t, U_t)) \mathbf{u}_h]}]{\mathbb{E}_{\pi_{\theta, \hat{f}_\xi} [\mathbb{1}_{C \leq C_{t, \max}}(C(X_t, U_t))]}]} \quad (86)$$

G Nonparametric Bayesian MPC

Inference can be performed directly over the posterior of control input sequences $\theta \triangleq (\mathbf{u}_t, \mathbf{u}_{t+1}, \dots, \mathbf{u}_{t+H-1})$. Assuming the likelihood gradient in the Bayesian MPC setting can be derived for differentiable dynamics:

$$\nabla_{\theta} \log \mathbb{E}_{\hat{f}_\xi} [\mathcal{L}(X_t)] = \frac{\mathbb{E}_{\hat{f}_\xi} [\mathcal{L}(X_t) \nabla_{\theta} \log p(X_t | \theta; \xi, \mathbf{x}_t)]}{\mathbb{E}_{\hat{f}_\xi} [\mathcal{L}(X_t)]} \quad (87)$$

Similarly to the parametric formulation, the form of the cost-likelihood \mathcal{L} will result in a particular update rule. The gradient can generally be evaluated by approximating the expectations with Monte Carlo sampling of state trajectories given the controlled stochastic dynamics., and evaluating the gradients on sampled trajectories. Under certain conditions, however, it may be evaluated in closed form. Such is the case for a Linear-Quadratic-Gaussian (LQG) system, for example [27].

H Motion Planning

A special case can be considered for deterministic dynamics. We can define the stochastic dynamics using a dirac measure on the space of trajectories: $p(X_t | \theta; \xi, \mathbf{x}_t) = \delta(X_t - \bar{X}_t)$, where $\bar{X}_t = F(\theta, \xi, \mathbf{x}_t)$, and F performs consecutive application of the deterministic dynamics f . Setting

$\mathcal{L}(X_t, \theta_t) = \exp(-\alpha C(X_t, \theta_t))$, the gradient then reduces to :

$$\nabla_{\theta} \log \mathbb{E}_{\hat{f}_{\xi}} \left[\mathcal{L}(X_t, \theta) \right] = \frac{\nabla_{\theta} \int \exp(-\alpha C(X_t, \theta)) \delta(X_t - \bar{X}_t) dX_t}{\int \exp(-\alpha C(X_t, \theta)) \delta(X_t - \bar{X}_t) dX_t} \quad (88)$$

$$= \frac{\nabla_{\theta} \exp(-\alpha C(\bar{X}_t, \theta))}{\exp(-\alpha C(\bar{X}_t, \theta))} \quad (89)$$

$$= \frac{-\alpha \exp(-\alpha C(\bar{X}_t, \theta)) \nabla_{\theta} C(\bar{X}_t, \theta)}{\exp(-\alpha C(\bar{X}_t, \theta))} \quad (90)$$

$$= -\alpha \nabla_{\theta} C(\bar{X}_t, \theta) \quad (91)$$

The gradient can then be evaluated in a straightforward manner via back-propagation on the cost function, through the dynamics.

I Complexity

The implementation of SV-MPC requires the computation of the kernel Gram-matrix $K(\theta^i, \theta^j) \forall i, j \in (1, \dots, m)$ for the SVGD update. This requires an inner-product operation for all particle pairs, resulting in a computational complexity of $\mathcal{O}(m^2hd)$ (where m : number of particles, h : horizon, d : control dimension). However, by exploiting structured kernels for trajectories, the scaling with respect to the horizon can be removed by parallel computation of kernel factors. This results in a $\mathcal{O}(m^2 + k)$ operation, where k is the cost of a kernel factor evaluation. In practice a relatively low order of particles is required: $m \ll 1 \times 10^3$. The core bottleneck is typically the generation of rollout trajectories from control samples during each iteration of MPC, which can be performed in parallel but is linear in time with respect to the horizon h . As with most MPC applications, the horizon length, number of samples, etc. can be varied to balance accuracy with runtime complexity, depending on the constraints of the system.

J Additional Results

Planar Navigation Table 6, we include quantitative summaries of performance across control types. The performance of the SV-MPC controller improves with increasing particle number.

Figure 6: Statistics for planar navigation task over 25 trials (4x4 obstacle grid)

Controller	Num. of particles	Avg. cost of success ($\times 10^3$)	Success rate (%)
SV-MPC	32	20.7	96
	12	21.8	96
	6	25.6	84
MPPI	—	26.5	64
CEM	—	25.4	64

K Experimental Details

The environment and controller parameters used in the experiments are summarized in tables 2 and 3, respectively. Note for Table 3: the utility function parameter, such as α and *elite fraction*, share the same values as those used by SV-MPC for either EU or PLC likelihood.

Planar navigation The SV-MPC controller with exponentiated-utility (EU) is constructed with a set of 6 to 32 particles, where the gradient of each particle is estimated via Monte-Carlo sampling of control and state trajectories. At each round, the best-performing particle is chosen to generate the action using (24).

For state $\mathbf{x}_t = (x_t, \dot{x}_t)$ and control $\mathbf{u}_t = \ddot{x}_t$, where $x_t, \dot{x}_t, \ddot{x}_t$ are the 2D position, velocity and acceleration, respectively, and the x_{goal} the target 2D goal position, we define the instantaneous and terminal costs (with respect to (1)):

$$c(\mathbf{x}_t, \mathbf{u}_t) = 0.5(x_t - x_{goal})^\top (x_t - x_{goal}) + 0.25\dot{x}_t^\top \dot{x}_t + 0.2\ddot{x}_t^\top \ddot{x}_t \quad (92)$$

$$c_{\text{term}}(\mathbf{x}_t) = 1000(x_t - x_{goal})^\top (x_t - x_{goal}) + 0.1\dot{x}_t^\top \dot{x}_t \quad (93)$$

Manipulation This provides the ability to efficiently compute highly-resolved geometric constraints between the robot and obstacles, eliminating the need for coarsely-defined heuristics for collision detection (such as signed-distance fields [35]).

Although the dynamics are deterministic, the problem remains challenging since the posterior probability distribution implies a nonlinear mapping (via inverse kinematics) from the work-space to the sampling space.

The cost-function consists of a cost on cartesian distance-to-goal from the end-effector, as well as a penalty on control. For state $\mathbf{x}_t = (e_t, \dot{e}_t)$ and control $\mathbf{u}_t = \dot{q}_t$, where e_t, \dot{e}_t , are the cartesian end-effector positions and velocities, respectively, \dot{q} the joint velocities, and the e_{goal} the target 3D goal position, we define the instantaneous and terminal costs (with respect to (1)):

$$c(\mathbf{x}_t, \mathbf{u}_t) = 1(e_t - e_{goal})^\top (e_t - e_{goal}) + 0.25\dot{e}_t^\top \dot{e}_t + 0.1\dot{q}_t^\top \dot{q}_t \quad (94)$$

$$c_{\text{term}}(\mathbf{x}_t) = 5000(e_t - e_{goal})^\top (e_t - e_{goal}) + 0.1\dot{e}_t^\top \dot{e}_t \quad (95)$$

Stochastic HalfCheetah We modify the cost function to reward forward velocity only if the agent is forward-facing, as done in [22], along with a control penalty. Without this alteration, progress can be made fairly easily by applying torque commands in a single direction, and ‘cart-wheeling’ the system. For instantaneous forward velocity v_t and body angle β ,

$$c(\mathbf{x}_t, \mathbf{u}_t) = 0.1\mathbf{u}_t^\top \mathbf{u}_t - \frac{v_t}{2}(1 + \text{sgn}(\cos \beta)) \quad (96)$$

Motion Planning The dynamics consist of a deterministic, velocity controlled single-integrator model on the 2D position: $\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{u}_t\Delta t$. The initial set of particles is drawn from a multi-variate Gaussian prior: $\theta^i \sim \mathcal{N}(0, \Sigma)$. A smooth obstacle cost-map is generated using a bi-modal mixture of Gaussians, with the probability of collision is given by $p_{\text{obs}}(x_t)$. The cost function is then given by:

$$c(\mathbf{x}_t) = 1 \times 10^5 p_{\text{obs}}(x_t) \quad (97)$$

$$c_{\text{term}}(\mathbf{x}_t) = 1000(x_t - x_{goal})^\top (x_t - x_{goal}) \quad (98)$$

Table 2: Environment Parameters

Experiment	Episode length	Trials	state space ($\mathbf{x}_t \in$)	control space ($\mathbf{u}_t \in$)	control limits	time-step (Δt)	dyn. noise σ_{dyn}^2
Planar Nav.	300	25	\mathbb{R}^4	\mathbb{R}^2	$[-50, 50]$	0.015	0.1
Manipulation	150	24	\mathbb{R}^{14}	\mathbb{R}^7	$[-1, 1]$	0.067	—
Stoch. HalfCheetah	300	16	\mathbb{R}^{18}	\mathbb{R}^6	$[-1, 1]$	0.05	0.25

Table 3: Experiment Controller Settings.

		Planar Nav.	Manipulation	Stoch. HalfCheetah
Planning Horizon (H)		64	64	30
Warm-start iterations		30	25	20
Optimization iterations per timestep		1	1	5
Control variance (σ^2)		100	0.25	1
SV-MPC	num. particles (m)	6, 12, 32	4	5
	num. ctrl. samples per particle	8	32	32
	step-size (ϵ)	10	0.1	1
	cost-likelihood ($\mathcal{L}(\cdot)$)	EU	EU	EU, PLC
MPPI	num. control samples	32	128	160
	inverse-temperature (α)	1×10^{-3}	0.1	1
CEM	num. control samples	32	128	160
	elite fraction	0.1	0.1	0.1