

Dual Online Stein Variational Inference for Control and Dynamics


Lucas Barcelos^{*‡}, Alexander Lambert[†], Rafael Oliveira^{*}, Paulo Borges[‡], Byron Boots^{§¶}, and Fabio Ramos^{*¶}
^{*}The University of Sydney, [†]Georgia Institute of Technology, [‡]CSIRO, [§]University of Washington, [¶]NVIDIA

Abstract—Model predictive control (MPC) schemes have a proven track record for delivering aggressive and robust performance in many challenging control tasks, coping with nonlinear system dynamics, constraints, and observational noise. Despite their success, these methods often rely on simple control distributions, which can limit their performance in highly uncertain and complex environments. MPC frameworks must be able to accommodate changing distributions over system parameters, based on the most recent measurements. In this paper, we devise an implicit variational inference algorithm able to estimate distributions over model parameters and control inputs on-the-fly. The method incorporates Stein Variational gradient descent to approximate the target distributions as a collection of particles, and performs updates based on a Bayesian formulation. This enables the approximation of complex multi-modal posterior distributions, typically occurring in challenging and realistic robot navigation tasks. We demonstrate our approach on both simulated and real-world experiments requiring real-time execution in the face of dynamically changing environments.


I. INTRODUCTION

Real robotics applications are invariably subjected to uncertainty arising from either unknown model parameters or stochastic environments. To address the robustness of control strategies, many frameworks have been proposed of which model predictive control is one of the most successful and popular [6]. MPC has become a primary control method for handling nonlinear system dynamics and constraints on input, output and state, taking into account performance criteria. It originally gained popularity in chemical and processes control [12], being more recently adapted to various fields, such as agricultural machinery [10], automotive systems [9], and robotics [17, 42]. In its essence, MPC relies on different optimisation strategies to find a sequence of actions over a given control horizon that minimises an optimality criteria defined by a cost function.

Despite their success in practical applications, traditional dynamic-programming approaches to MPC (such as iLQR and DDP [32]) rely on a differentiable cost function and dynamics model. Stochastic Optimal Control variants, such as iLQG [33] and PDDP [23], can accommodate stochastic dynamics, but only under simplifying assumptions such as additive Gaussian noise. These approaches are generally less effective in addressing complex distributions over actions, and it is unclear how these methods should incorporate model uncertainty, if any. In contrast, sampling-based control schemes have gained increasing popularity for their general robustness



(a) Wombot AGV



(b) Posterior distribution over x_{ICR}

Figure 1: Online parameter estimation for autonomous ground vehicles. Distributions over system parameters such as the inertial center of rotation (ICR), are adapted in real-time. (a) The custom built skid-steer robot platform used in experiments. (b) Distribution over x_{ICR} at different time steps. The mass load on the robot is suddenly increased during system execution. The parameter distribution estimate quickly changes to include a second mode that better explains the new dynamics. Our particle-based control scheme can accommodate such multi-modal uncertainty and adapt to dynamically changing environments.

to model uncertainty, ease of implementation, and ability to contend with sparse cost functions [40].

To address some of these challenges, several approaches have been proposed. In sampling based Stochastic Optimal Control (SOC), the optimisation problem is replaced by a sampling-based algorithm that tries to approximate an optimal distribution over actions [39]. This has shown promising results in several applications and addresses some of the former disadvantages, however it may inadequately capture the complexity of the true posterior.

In recent work [19, 22], the MPC problem has been formulated as a Bayesian inference task whose goal is to estimate the posterior distribution over control parameters given the state and observed costs. To make such problem tractable in real-world applications, variational inference has been used to approximate the optimal distribution. These approaches are better suited at handling the multi-modality of the distribution over actions, but do not attempt to dynamically adapt to changes in the environment.

Conversely, previous work has demonstrated that incorporating uncertainty in the evaluation of SOC estimates can improve performance [3], particularly when this uncertainty is periodically re-estimated [26]. Although this method is more robust to model mismatch and help address the sim-to-real gap,

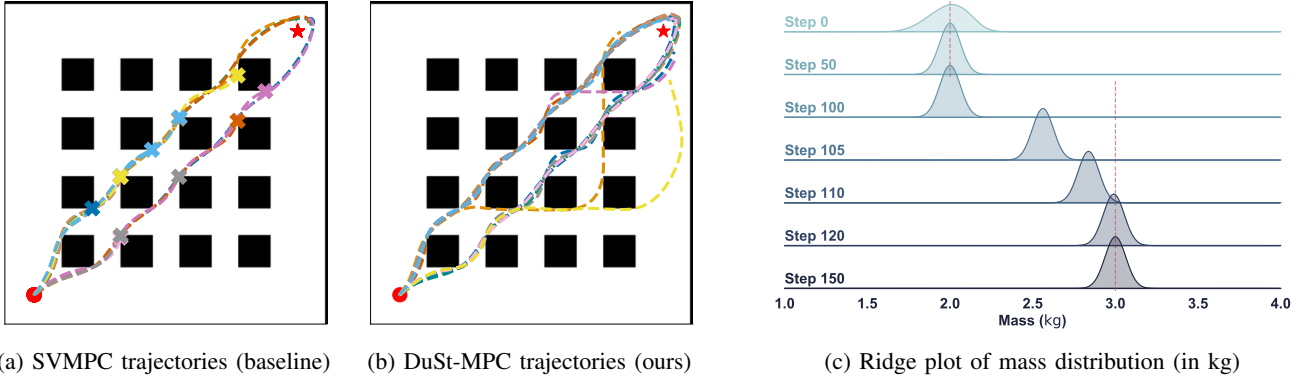


Figure 2: **Point-mass navigation task.** The plots shows trajectories from the start position (red dot) towards the goal (red star). **Left:** Trajectories executed by SVMPC. Note that, as the mass of the robot changes, the model mismatch causes many of the episodes to crash (x markers). **Centre:** Trajectories executed by DuSt-MPC. Depending on the state of the system when the mass change occurs, a few trajectories deviate from the centre path to avoid collisions. A few trajectories are truncated due to the fixed episode length. **Right:** Ridge plot of the distribution over mass along several steps of the simulation. The vertical dashed line denotes the true mass. Mass is initially set at 2 kg, and changed to 3 kg at step 100.

the strategy relies on applying moment-matching techniques to propagate the uncertainty through the dynamical system which is approximated by a Gaussian distribution. This diminishes the effectiveness of the method under settings where multi-modality is prominent.

In this work, we aim to leverage recent developments in variational inference with MPC for decision-making under complex multi-modal uncertainty over actions while simultaneously estimating the uncertainty over model parameters. We propose a Stein variational stochastic gradient solution that models the posterior distribution over actions and model parameters as a collection of particles, representing an implicit variational distribution. These particles are updated sequentially, online, in parallel, and can capture complex multi-modal distributions. We call this method Dual Stein Variational Inference MPC, or DuSt-MPC for conciseness.¹

Specifically, the main contributions of this paper are:

- We propose a principled Bayesian solution of introducing uncertainty over model parameters in Stein variational MPC and empirically demonstrate how this can be leveraged to improve the control policy robustness;
- We introduce a novel method to extend the inference problem and simultaneously optimise the control policy while refining our knowledge of the environment as new observations are gathered. Crucially, by leveraging recent advancements in sequential Monte Carlo with kernel embedding, we perform online, sequential updates to the distribution over model parameters which scales to large datasets;
- By capturing the uncertainty on true dynamic systems in distributions over a parametric model, we are able to incorporate domain knowledge and physics principles

¹Note, that *dual* is used in the sense of a doubly stochastic problem, i.e. inferring a distribution over policies and system dynamics, and not as a duality problem in optimisation.

while still allowing for a highly representative model. This simplifies the inference problem and drastically reduces the number of interactions with the environment to characterise the model.

We implement the algorithm on a real autonomous ground vehicle (AGV) (fig. 1a), illustrating the applicability of the method in real time. Experiments show how the control and parameter inference are leveraged to adapt the behaviour of the robot under varying conditions, such as changes in mass. We also present simulation results on an inverted pendulum and an 2D obstacle grid, see fig. 2, demonstrating an effective adaptation to dynamic changes in model parameters.

This paper is organised as follows. In Section II we review related work, contrasting the proposed method to the existing literature. In Section III we provide background on stochastic MPC and Stein variational gradient descent as a foundation to the proposed method, which is presented in Section IV. In Section V we present a number of real and simulated experiments, followed by relevant conclusions in Section VI.

II. RELATED WORK

Sampling-based approaches for stochastic MPC have shown to be suitable for a range of control problems in robotics [38, 36]. At each iteration, these methods perform an approximate evaluation by rolling-out a stochastic policy with modelled system dynamics over a finite-length horizon. The optimisation step proceeds to update the policy parameters in the direction that minimises the expected cost and a statistical distance to a reference policy or prior [36]. This can equivalently be interpreted as a statistical inference procedure, where policy parameters are updated in order to match an optimal posterior distribution distribution [29, 38, 19]. This connection has motivated the use of common approximate inference procedures for SOC. Model Predictive Path Integral Control (MPPI) [39, 38] and the Cross Entropy Method (CEM) [5], for instance, use importance sampling to match a Gaussian proposal

distribution to moments of the posterior. Variational inference (VI) approaches have also been examined for addressing control problems exhibiting highly non-Gaussian or multi-modal posterior distributions. This family of Bayesian inference methods extends the modelling capacity of control distribution to minimise a Kullback-Leibler divergence with the target distribution. Traditional VI methods such as Expectation-Maximisation have been examined for control problems [37, 22], where the model class of the probability distribution is assumed to be restricted to a parametric family (typically Gaussian Mixture Models). More recently, the authors in [19] proposed to adapt the Stein Variational Gradient Descent (SVGD) [21] method for model predictive control. Here, a distribution of control sequences is represented as a collection of particles. The resulting framework, Stein-Variational Model Predictive Control (SVMPC), adapts the particle distribution in an online fashion. The non-parametric representation makes the approach particularly suitable to systems exhibiting multi-modal posteriors. In the present work, we build on the approach in [19] to develop an MPC framework which leverages particle-based variational inference to optimise controls *and* explicitly address model uncertainty. Our method *simultaneously* adapts a distribution over dynamics parameters, which we demonstrate improves robustness and performance on a real system.

Model predictive control is a reactive control scheme, and can accommodate modelling errors to a limited degree. However, its performance is largely affected by the accuracy of long-range predictions. Modelling errors can compound over the planning horizon, affecting the expected outcome of a given control action. This can be mitigated by accounting for model uncertainty, leading to better estimates of expected cost. This has been demonstrated to improve performance in stochastic optimal control methods and model-based reinforcement learning [23, 8]. Integrating uncertainty has typically been achieved by learning probabilistic dynamics models from collected state-transition data in an episodic setting, where the model is updated in between trajectory-length system executions [7, 22, 35, 28]. A variety of modelling representations have been explored, including Gaussian processes [8], neural network ensembles [7], Bayesian regression, and meta-learning [16]. Alternatively, the authors in [28, 26] estimate posterior distributions of physical parameters for black-box simulators, given real-world observations.

Recent efforts have examined the online setting, where a learned probabilistic model is updated based on observations made *during execution* [24, 1, 13, 14]. The benefits of this paradigm are clear: incorporating new observations and adapting the dynamics *in situ* will allow for better predictions, improved control, and recovery from sudden changes to the environment. However, real-time requirements dictate that model adaptation must be done quickly and efficiently, and accommodate the operational timescale of the controller. This typically comes at the cost of modelling accuracy, and limits the application of computationally-burdensome representations, such as neural networks and vanilla GPs. Previous work has included the use of sparse-spectrum GPs and efficient

factorization to incrementally update the dynamics model [24, 25]. In [16], the authors use a meta-learning approach to train a network model offline, which is adapted to new observations using Bayesian linear regression operating on the last layer. However, these approaches are restricted to Gaussian predictive distributions, and may lack sufficient modelling power for predicting complex, multi-modal distributions.

Perhaps most closely related to our modeling approach is the work by Abraham et al. [1]. The authors propose to track a distribution over simulation parameters using a sequential Monte Carlo method akin to a particle filter. The set of possible environments resulting from the parameter distribution is used by an MPPI controller to generate control samples. Each simulated trajectory rollout is then weighted according to the weight of the corresponding environment parameter. Although such an approach can model multi-modal posterior distributions, we should expect similar drawbacks to particle filters, which require clever re-sampling schemes to avoid mode collapse and particle depletion. Our method also leverages a particle-based representation of parameter distributions, but performs deterministic updates based on new information and is more sample efficient than MC sampling techniques.

III. BACKGROUND

A. Stochastic Model Predictive Control

We consider the problem of controlling a discrete-time nonlinear system compactly represented as:

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

where f is the transition map, $\mathbf{x}_t \in \mathcal{X}$ denotes the system state, $\mathbf{u}_t \in \mathcal{U}$ the control input, and \mathcal{X} and \mathcal{U} are respectively appropriate Euclidean spaces for the states and controls. More specifically, we are interested in the problem where the real transition function $f(\mathbf{x}_t, \mathbf{u}_t)$ is unknown and approximated by a transition function with *parametric uncertainty*, such that:

$$f(\mathbf{x}_t, \mathbf{u}_t) \approx \hat{f}(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\xi}) =: \hat{f}_{\boldsymbol{\xi}}(\mathbf{x}_t, \mathbf{u}_t) \quad (2)$$

where $\boldsymbol{\xi} \in \Xi$ are the simulator parameters with a prior probability distribution $p(\boldsymbol{\xi})$ and the non-linear forward model $\hat{f}(\mathbf{x}, \mathbf{u}, \boldsymbol{\xi})$, is represented as $\hat{f}_{\boldsymbol{\xi}}$ for compactness.

Based on the steps in [3], we can then define a fixed length control sequence $U_t = \{\mathbf{u}_t\}_{t < t+H}$ over a fixed control horizon H , onto which we apply a receding horizon control (RHC) strategy. Moreover, we define a mapping operator, \mathcal{H} , from input sequences U_t to their resulting states by recursively applying $\hat{f}_{\boldsymbol{\xi}}$ given \mathbf{x}_t ,

$$\begin{aligned} X_t &= \mathcal{H}(U_t; \mathbf{x}_t) \\ &= \left[\mathbf{x}_t, \hat{f}_{\boldsymbol{\xi}}(\mathbf{x}_t, \mathbf{u}_t), \hat{f}_{\boldsymbol{\xi}}(\hat{f}_{\boldsymbol{\xi}}(\mathbf{x}_t, \mathbf{u}_t), \mathbf{u}_{t+1}), \dots \right], \end{aligned} \quad (3)$$

noting that X_t and U_t define estimates of future states and controls. Finally, we can define a *trajectory* as the joint sequence of states and actions, namely:

$$\boldsymbol{\tau} = (X_t, U_t). \quad (4)$$

In MPC, we want to minimise a task dependent cost functional generally described as:

$$C(X_t, U_t) = c_{term}(\hat{\mathbf{x}}_{t+H}) + \sum_{h=0}^{H-1} c(\hat{\mathbf{x}}_{t+h}, \hat{\mathbf{u}}_{t+h}), \quad (5)$$

where $\hat{\mathbf{x}}$ and $\hat{\mathbf{u}}$ are the estimated states and controls, and $c(\cdot)$ and $c_{term}(\cdot)$ are respectively arbitrary instantaneous and terminal cost functions. To do so, our goal is to find an optimal policy $\pi(\mathbf{x}_t)$ to generate the control sequence U_t at each time-step. As in [36], we define such feedback policy as a parameterised probability distribution $\pi_{\theta_t} = p(\mathbf{u}_t | \mathbf{x}_t; \theta_t)$ from which the actions at time t are sampled, i.e. $U_t \sim \pi_{\theta_t}$. The control problem can then be defined as:

$$\min_{\theta_t \in \Theta} \hat{J}(\pi_{\theta_t}; \mathbf{x}_t), \quad (6)$$

where \hat{J} is an estimator for a statistic of interest, typically $\mathbb{E}_{\pi_{\theta_t}, \hat{\mathbf{x}}_{\xi}} [C(X_t, U_t)]$. Once π_{θ_t} is determined, we can use it to sample U_t , extract the first control \mathbf{u}_t and apply it to the system.

One of the main advantages of stochastic MPC is that \hat{J} can be computed even when the cost function is non-differentiable w.r.t. to the policy parameters by using importance sampling [39].

B. Stein Variational Gradient Descent

Variational inference poses posterior estimation as an optimisation task where a candidate distribution $q^*(\mathbf{x})$ within a distribution family \mathcal{Q} is chosen to best approximate the target distribution $p(\mathbf{x})$. This is typically obtained by minimising the Kullback-Leibler (KL) divergence:

$$q^*(\mathbf{x}) = \operatorname{argmin}_{q \in \mathcal{Q}} D_{\text{KL}}(q(\mathbf{x}) || p(\mathbf{x})). \quad (7)$$

The solution also maximises the Evidence Lower Bound (ELBO), as expressed by the following objective:

$$q^*(\mathbf{x}) = \operatorname{argmax}_{q \in \mathcal{Q}} \mathbb{E}_q[\log p(\mathbf{x})] - D_{\text{KL}}(q(\mathbf{x}) || p(\mathbf{x})) \quad (8)$$

In order to circumvent the challenge of determining an appropriate \mathcal{Q} , while also addressing eq. (7), we develop an algorithm based on Stein variational gradient descent (SVGD) for Bayesian inference [21]. The non-parametric nature of SVGD is advantageous as it removes the need for assumptions on restricted parametric families for $q(\mathbf{x})$. This approach approximates a posterior $p(\mathbf{x})$ with a set of particles $\{\mathbf{x}^i\}_i^{N_p}$, $\mathbf{x} \in \mathbb{R}^p$. These particles are iteratively updated according to:

$$\mathbf{x}^i \leftarrow \mathbf{x}^i + \epsilon \phi^*(\mathbf{x}^i), \quad (9)$$

given a step size ϵ . The function $\phi(\cdot)$ is known as score function and defines the velocity field that maximally decreases the KL-divergence.

$$\phi^* = \operatorname{argmax}_{\phi \in \mathcal{S}} \{-\nabla_{\epsilon} D_{\text{KL}}(q_{[\epsilon\phi]} || p(\mathbf{x})), \text{ s.t. } \|\phi\|_{\mathcal{S}} \leq 1\} \quad (10)$$

where \mathcal{S} is a Reproducing Kernel Hilbert Space (RKHS) induced by the kernel function used and $q_{[\epsilon\phi]}$ indicates the

particle distribution resulting from taking an update step as in eq. (9). In [21] this has been shown to yield a closed-form solution which can be interpreted as a functional gradient in \mathcal{S} and approximated with the set of particles:

$$\phi^*(\mathbf{x}) = \frac{1}{N_p} \sum_{j=1}^{N_p} [k(\mathbf{x}^j, \mathbf{x}) \nabla_{\mathbf{x}^j} \log p(\mathbf{x}^j) + \nabla_{\mathbf{x}^j} k(\mathbf{x}^j, \mathbf{x})] \quad (11)$$

IV. METHOD

In this section, we present our approach for joint inference over control and model parameters for MPC. We begin by formulating optimal control as an inference problem and address how to optimise policies in section IV-C. Later, in section IV-D, we extend the inference to also include the system dynamics. A complete overview of the method is described in algorithm in appendix A.

A. MPC as Bayesian Inference

MPC can be framed as a Bayesian inference problem where we estimate the posterior distribution of policies, parameterised by θ_t , given an optimality criterion. Let $\mathcal{O} : \mathcal{T} \rightarrow \{0, 1\}$ be an *optimality* indicator for a trajectory $\tau \in \mathcal{T}$ such that $\mathcal{O}[\tau] = 1$ indicates that the trajectory is optimal. Now we can apply Bayes' Rule and frame our problem as estimating:

$$p(\theta_t | \mathcal{O}) \propto p(\mathcal{O} | \theta_t) p(\theta_t) = p(\mathcal{O}, \theta_t). \quad (12)$$

A reasonable way to quantify $\mathcal{O}[\tau]$ is to model it as a Bernoulli random variable conditioned on the trajectory τ , allowing us to define the *likelihood* $p(\mathcal{O} | \theta_t)$ as:

$$\begin{aligned} p(\mathcal{O} | \theta_t) &= \mathbb{E}_{\tau \sim p(\tau | \theta_t)} [p(\mathcal{O}[\tau] = 1 | \tau)] \\ &\approx \frac{1}{n} \sum_{i=1}^n \exp(-\alpha C[\tau_i]), \end{aligned} \quad (13)$$

where $\tau_i \stackrel{i.i.d.}{\sim} p(\tau | \theta_t)$, $i \in \{1, 2, \dots, n\}$, $p(\mathcal{O}[\tau] = 1 | \tau) := \exp(-\lambda C[\tau])$, and we overload $p(\mathcal{O}[\tau] = 1 | \theta_t)$ to simplify the notation. The variable α is known as the inverse temperature parameter and controls the amount of exploration of the policy. A lower α results in smaller separation between each cost functional and encourages more trajectories to contribute in the gradient step. Conversely, higher values of α increase the cost functional separation of each rollout and promotes a more exploitative policy. Now, assuming a prior $p(\theta_t)$ for θ_t , the posterior over θ_t is given by:

$$p(\theta_t | \mathcal{O}) \propto p(\mathcal{O} | \theta_t) p(\theta_t) = \int_{\mathcal{T}} p(\mathcal{O} | \tau) p(\tau | \theta_t) p(\theta_t) d\tau. \quad (14)$$

This posterior corresponds to the probability (density) of a given parameter setting θ_t conditioned on the hypothesis that (implicitly observed) trajectories generated by θ_t are *optimal*. Alternatively, one may say that $p(\theta_t | \mathcal{O})$ tells us the probability

of θ_t being the generator of the optimal set \mathcal{T}^* . Lastly, note that the trajectories conditional $p(\tau|\theta_t)$ factorises as:

$$p(\tau|\theta_t) = \prod_{h=0}^{H-1} p_{\xi}(\mathbf{x}_{t+h+1}|\mathbf{x}_{t+h}, \mathbf{u}_{t+h})\pi_{\theta_t}(\mathbf{x}_{t+h}). \quad (15)$$

B. Joint Inference for Policy and Dynamics with Stein MPC

In this section, we generalise the framework presented in [19] to simultaneously refine our knowledge of the dynamical system, parameterised by ξ , while estimating optimal policy parameters θ_t . Before we proceed, however, it is important to notice that the optimality measure defined in section IV-A stems from *simulated* rollouts sampled according to eq. (13). Hence, these trajectories are not actual observations of the agent's environment and are not suitable for inferring the parameters of the system dynamics. In other words, to perform inference over the parameters ξ we need to collect *real observations* from the environment.

Moreover, we argue that the two inference problems can be naturally factorised. From the perspective of the plant, i.e. the physical system, the dynamics inference depends solely on previously observed data, which despite including past control actions, is independent of the current policy. On the other hand, the policy inference assumes that the distribution over the system dynamics at a given time is given and invariant and relies exclusively on simulated future rollouts. With that in mind, the problem statement in eq. (12) can be rewritten as:

$$p(\theta_t, \xi|\mathcal{O}, \mathcal{D}_{1:t}) = p(\theta_t|\mathcal{O}, \xi)p(\xi|\mathcal{D}_{1:t}), \quad (16)$$

where $\mathcal{D}_{1:t} := \{(\mathbf{x}_t^r, \mathbf{u}_{t-1}^r, \mathbf{x}_{t-1}^r)\}_{t=1}^{N_D}$ represents the dataset of collected environment observations. Note that the distribution over the parameters of the plant, the dynamical system, is independent from the policy optimality, and therefore the conditioning on \mathcal{O} has been dropped. Similarly, once the distribution over ξ is defined, the policy parameters θ are independent from the previous environment observations $\mathcal{D}_{1:t}$, and again we omit the conditioning.

If one were to solve a joint inference problem by defining a new random variable $\Theta = \{\theta, \xi\}$, and following the steps outlined in [19], careful consideration should be taken to prevent the partial derivative of ξ from including terms dependent on the policy optimality (see appendix D for further discussion). Furthermore, by factorising the problem we can perform each inference step separately and adjust the computational effort and hyper-parameters based on the idiosyncrasies of the problem at hand.

C. Policy Inference for Bayesian MPC

Having formulated the inference problem as in eq. (16), we can proceed by solving each of the factors separately. We shall start with optimising π_{θ_t} according to $p(\theta_t|\mathcal{O}, \xi)$. It is evident that we need to consider the dynamics parameters when optimising the policy, but at this stage we may simply assume the inference over ξ has been solved and leverage our

knowledge of $p(\xi|\mathcal{D}_{1:t})$. More concretely, let us rewrite the factor on the RHS of eq. (16) by marginalising over ξ so that:

$$p(\theta_t|\mathcal{O}, \mathcal{D}_{1:t}) \propto \int_{\Xi} \ell_{\pi}(\theta_t)p(\theta_t)p(\xi|\mathcal{D}_{1:t})d\xi, \quad (17)$$

where, as in eq. (13), the likelihood $\ell_{\pi}(\theta_t)$ is defined as:

$$\begin{aligned} \ell_{\pi}(\theta_t) &= p(\mathcal{O}|\theta_t, \xi) := P(\mathcal{T}^*|\theta_t, \xi) \\ &= \int_{\mathcal{T}} p(\mathcal{O}|\tau)p(\tau|\theta_t, \xi)d\tau \\ &\approx \frac{1}{N_{\tau}} \sum_{i=1}^{N_{\tau}} \exp(-\alpha C[\tau_i]), \end{aligned} \quad (18)$$

with $\tau_i \stackrel{i.i.d.}{\sim} p(\tau|\theta_t, \xi)$, $i \in \{1, 2, \dots, N_{\tau}\}$, and $p(\xi|\mathcal{D}_{1:t})$ defines the inference problem of updating the posterior distribution of the dynamics parameters given all observations gathered from the environment, which we shall discuss in the next section. Careful consideration of eq. (18) tells us that unlike the case of a deterministic transition function or even maximum likelihood point estimation of ξ , the optimality of a given trajectory now depends on its *expected* cost over the distribution $p(\xi)$.

Hence, given a prior $p(\theta_t)$ and $p(\xi|\mathcal{D}_{1:t})$, we can generate samples from the likelihood in eq. (18) and use an stochastic gradient method as in section III-B to infer the posterior distribution over θ_t . Following the steps in [19], we approximate the prior $p(\theta_t)$ by a set of particles $q(\theta_t) = \{\theta_t\}_{i=1}^{N_{\tau}}$ and take sequential SVGD updates:

$$\theta_t^i \leftarrow \theta_t^i + \epsilon \phi^*(\theta_t^i), \quad (19)$$

to derive the posterior distribution. Where, again, ϕ^* is computed as in eq. (11) for each intermediate step and ϵ is a predetermined step size. One ingredient missing to compute the score function is the gradient of the log-posterior of $p(\theta_t|\mathcal{O}, \xi)$, which can be factorised into:

$$\nabla_{\theta_t^i} \log p(\theta_t^i|\mathcal{O}, \xi) = \nabla_{\theta_t^i} \log \ell(\theta_t^i) + \nabla_{\theta_t^i} \log q(\theta_t^i). \quad (20)$$

In practice we typically assume that the $C[\cdot]$ functional used as surrogate for optimality is not differentiable w.r.t. θ_t , but the gradient of the log-likelihood can be usually approximated via Monte-Carlo sampling.

Most notably, however, is the fact that unlike the original formulation in SVGD, the policy distribution we are trying to infer is time-varying and depends on the actual state of the system. The measure $P(\mathcal{T}^*|\theta_t; \mathbf{x}_t)$ depends on the current state, as that is the initial condition for all trajectories evaluated in the cost functional $C[\tau]$.

Theoretically, one could choose an uninformative prior with broad support over the \mathcal{U} at each time-step and employ standard SVGD to compute the posterior policy. However, due to the sequential nature of the MPC algorithm, it is likely that the prior policy $q(\theta_{t-1})$ is close to a region of low cost and hence is a good candidate to bootstrap the posterior inference of the subsequent step. This procedure is akin to the prediction step

commonly found in Sequential Bayesian Estimation [11]. More concretely,

$$q(\boldsymbol{\theta}_t) = \int p(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1})q(\boldsymbol{\theta}_{t-1})d\boldsymbol{\theta}_{t-1}, \quad (21)$$

where $p(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1})$ can be an arbitrary transitional probability distribution. More commonly, however, is to use a probabilistic version of the shift operator as defined in [36]. For a brief discussion on action selection, please refer to appendix B. For more details on this section in general the reader is encouraged to refer to [19, Sec. 5.3].

D. Real-time Dynamics Inference

We now focus on the problem of updating the posterior over the simulator parameters. Note that, due to the independence of each inference problem, the frequency in which we update $p(\boldsymbol{\xi})$ can be different from the policy update.

In contrast, we are interested in the case where $p(\boldsymbol{\xi}|\mathcal{D}_{1:t})$ can be updated in *real-time* adjusting to changes in the environment. For that end, we need a more efficient way of updating our posterior distribution. The inference problem at a given time-step can then be written as:

$$p(\boldsymbol{\xi}|\mathcal{D}_{1:t}) \propto p(\mathcal{D}_t|\boldsymbol{\xi}, \mathcal{D}_{1:t-1})p(\boldsymbol{\xi}|\mathcal{D}_{1:t-1}). \quad (22)$$

Note that in this formulation, $\boldsymbol{\xi}$ is considered time-invariant. This based on the implicit assumption that the frequency in which we gather new observations is significantly larger than the covariate shift to which $p(\boldsymbol{\xi}|\mathcal{D}_{1:t})$ is subject to as we traverse the environment. In appendix C, we discuss the implications of changes in the the latent parameter over time.

In general, we do not have access to direct measurements of $\boldsymbol{\xi}$, only to the system state. Therefore, in order to perform inference over the dynamics parameters, we rely on a generative model, i.e. the simulator $\hat{f}_{\boldsymbol{\xi}}$, to generate samples in the state space \mathcal{X} for different values of $\boldsymbol{\xi}$. However, unlike in the policy inference step, for the dynamics parameter estimation we are not computing deterministic simulated rollouts, but rather trying to find the explanatory parameter for each observed transition in the environment. Namely, we have:

$$\mathbf{x}_t^r = f(\mathbf{x}_{t-1}^r, \mathbf{u}_{t-1}) = \hat{f}_{\boldsymbol{\xi}}(\mathbf{x}_{t-1}^r, \mathbf{u}_{t-1}) + \eta_t, \quad (23)$$

where \mathbf{x}_t^r denotes the true system state and η_t is a time-dependent random variable closely related to the *reality gap* in the sim-to-real literature [34] and incorporates all the complexities of the real system not captured in simulation, such as model mismatch, unmodelled dynamics, etc. As a result, the distribution of η_t is unknown, correlated over time and hard to estimate.

In practice, for the feasibility of the inference problem, we make the standard assumption that the noise is distributed according to a time-invariant normal distribution $\eta_t \sim \mathcal{N}(0, \Sigma_{\text{obs}})$, with an empirically chosen covariance matrix. More concretely, this allows us to define the likelihood term

in eq. (22) as:

$$\begin{aligned} \ell(\boldsymbol{\xi}|\mathcal{D}_{1:t}) &:= p(\mathcal{D}_t|\boldsymbol{\xi}, \mathcal{D}_{1:t-1}) \\ &= p(\mathbf{x}_t^r|\boldsymbol{\xi}, \mathcal{D}_{1:t-1}) \\ &= \mathcal{N}(\mathbf{x}_t^r; \hat{f}_{\boldsymbol{\xi}}(\mathbf{x}_{t-1}^r, \mathbf{u}_{t-1}), \Sigma_{\text{obs}}), \end{aligned} \quad (24)$$

where we leverage the symmetry of the Gaussian distribution to centre the uncertainty around \mathbf{x}_t^r . It follows that, because the state transition is Markovian, the likelihood of $\ell(\boldsymbol{\xi}|\mathcal{D}_{1:t})$ depends only on the current observation tuple given by \mathcal{D}_t , and we can drop the conditioning on previously observed data. In other words, we now have a way to quantify how likely is a given realisation of $\boldsymbol{\xi}$ based on the data we have collected from the environment. Furthermore, let us define a single observation $\mathcal{D}_t = (\mathbf{x}_t^r, \mathbf{u}_{t-1}^r, \mathbf{x}_{t-1}^r)$ as the tuple of last applied control action and observed state transition. Inferring exclusively over the current state is useful whenever frequent observations are received and prevents us from having to store information over the entire observation dataset. This approach is also followed by ensemble Kalman filters and particle flow filters [20].

Equipped with eq. (24) and assuming that an initial prior $p(\boldsymbol{\xi})$ is available, we can proceed as discussed in section III-B by approximating each prior at time t with a set of particles $\{\boldsymbol{\xi}^i\}_{i=1}^{N_{\boldsymbol{\xi}}}$ following $q(\boldsymbol{\xi}|\mathcal{D}_{1:t-1})$, so that our posterior over $\boldsymbol{\xi}$ at a given time t can then be rewritten as:

$$p(\boldsymbol{\xi}|\mathcal{D}_{1:t}) \approx q(\boldsymbol{\xi}|\mathcal{D}_{1:t}) \propto \ell(\boldsymbol{\xi}|\mathcal{D}_t)q(\boldsymbol{\xi}|\mathcal{D}_{1:t-1}), \quad (25)$$

and we can make recursive updates to $q(\boldsymbol{\xi}|\mathcal{D}_{1:t})$ by employing it as the prior distribution for the following step. Namely, we can iteratively update $q(\boldsymbol{\xi}|\mathcal{D}_{1:t})$ a number of steps L by applying the update rule with a functional gradient computed as in eq. (11), where:

$$\nabla_{\boldsymbol{\xi}} \log p_t(\boldsymbol{\xi}|\mathcal{D}_{1:t}) \approx \nabla_{\boldsymbol{\xi}} \log p(\mathcal{D}_t|\boldsymbol{\xi}) + \nabla_{\boldsymbol{\xi}} \log q_t(\boldsymbol{\xi}|\mathcal{D}_{1:t}). \quad (26)$$

An element needed to evaluate eq. (26) is an expression for the gradient of the posterior density. An issue in sequential Bayesian inference is that there is no exact expression for the posterior density [27]. Namely, we know the likelihood function, but the prior density is only represented by a set of particles, not the density itself.

One could forge an empirical distribution $q(\boldsymbol{\xi}|\mathcal{D}_{1:t}) = \frac{1}{N_{\boldsymbol{\xi}}} \sum_{i=1}^{N_{\boldsymbol{\xi}}} \delta(\boldsymbol{\xi}^i)$ by assigning Dirac functions at each particle location, but we would still be unable to differentiate the posterior. In practice, we need to apply an efficient density estimation method, as we need to compute the density at each optimisation step. We choose to approximate the posterior density with an equal-weight Gaussian Mixture Model (GMM) with a fixed diagonal covariance matrix:

$$q(\boldsymbol{\xi}|\mathcal{D}_{1:t}) = \frac{1}{N_{\boldsymbol{\xi}}} \sum_{i=1}^{N_{\boldsymbol{\xi}}} \mathcal{N}(\boldsymbol{\xi}; \boldsymbol{\xi}^i, \Sigma_s), \quad (27)$$

where the covariance matrix Σ_s can be predetermined or computed from data. One option, for example, is to use a Kernel Density bandwidth estimation heuristic, such as Improved

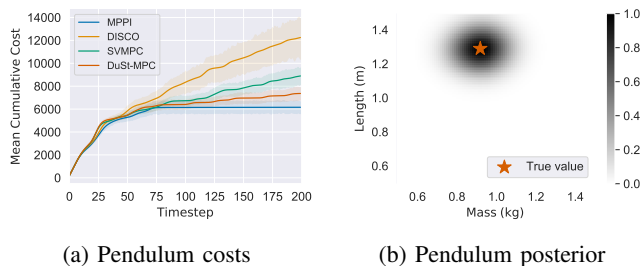


Figure 3: **Inverted pendulum.** (a) The image shows the mean cumulative cost over 10 episodes. The shaded region represents the 50% confidence interval. The high variance is expected since each scenario has parameters sampled from a uniform distribution. (b) Plot of the posterior distribution over the pendulum pole-mass at the final step of one of the episodes. The true latent value is shown by the red star marker.

Sheather Jones [4], Silverman’s [31] or Scott’s [30] rule, to determine the standard deviation σ and set $\Sigma_s = \sigma^2 I$.

One final consideration is that of the support for the prior distribution. Given the discussion above, it is important that the density approximation of $q(\xi|\mathcal{D}_t)$ offers support on regions of interest in the parameter space. In our formulation, that can be controlled by adjusting Σ_s . Additionally, precautions have to be taken to make sure the parameter space is specified correctly, such as using log-transformations for strictly positive parameters for instance.

V. EXPERIMENTS

In the following section we present experiments, both in simulation and with a physical autonomous ground vehicle (AGV), to demonstrate the correctness and applicability of our method. Although these experiments have low-dimensional control spaces it is important to note that in sampling-based MPC each control action in a sequential plan is independent. Therefore, the inference problem involves effectively solving a distribution whose dimensionality is the size of the control space *times* the number of control steps. As such, this demonstrates that the method is suitable in large-dimension spaces and that the control horizon can be adjusted according to the desired computational complexity.

For each experiment, the hyperparameters α , ϵ , H , Σ , Σ_a where optimised by performing Bayesian Optimisation within a given search space [2]. The optimal values found were subsequently rounded-off to fewer significant digits and are listed in appendix E.

A. Inverted pendulum with uncertain parameters

We first investigate the performance of DuSt-MPC in the classic inverted pendulum control problem. As usual, the pendulum is composed of a rigid pole-mass system controlled at one end by a 1-degree-of-freedom torque actuator. The task is to balance the point-mass upright, which, as the controller is typically under-actuated, requires a controlled swing motion to overcome gravity. Contrary to the typical case, however, in our experiments the mass and length of the pole-mass are

unknown and equally likely within a range of 0.6 kg to 1.3 kg and 0.6 m to 1.3 m, respectively.

At each episode, a set of latent model parameters is sampled and used in the simulated environment. Each method is then deployed utilising this same parameter set. MPPI is used as a baseline and has *perfect knowledge* of the latent parameters. This provides a measure of the task difficulty and achievable results. As discussed in section II, we compare against DISCO and SVMPC as additional baselines. We argue that, although these methods perform no online update of their knowledge of the world, they offer a good underpinning for comparison since the former tries to leverage the model uncertainty to create more robust policies, whereas the latter shares the same variational inference principles as our method. DISCO is implemented in its unscented transform variant applied to the uninformative prior used to generate the random environments. SVMPC uses the mean values for mass and length as point-estimates for its fixed parametric model. For more details on the hyperparameters used, refer to appendix E.

Figure 3a presents the average cumulative costs over 10 episodes. Although the results show great variance, due to the randomised environment, it is clear that DuSt-MPC outperforms both baselines. Careful consideration will show that the improvement is more noticeable as the episode progresses, which is expected as the posterior distribution over the model parameters being used by DuSt-MPC gets more informative. The final distribution over mass and length for one of the episodes is shown in fig. 3b. Finally, a summary of the experimental results is presented in table I.

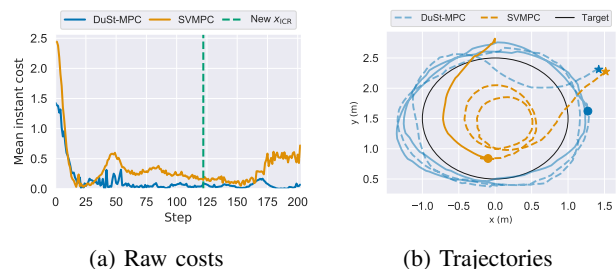


Figure 4: **AGV trajectory tracking.** (a) Raw cost over time. Amount of steps before and after the change of mass are normalised for proper comparison. (b) Trajectories executed by each method. Line style changes when mass changes. Markers denote initial and change of mass position.

B. Point-mass navigation on an obstacle grid

Here, we reproduce and extend the planar navigation task presented in [19]. We construct a scenario in which an holonomic point-mass robot must reach a target location while avoiding obstacles. As in [19], colliding with obstacles not only incurs a high cost penalty to the controller, but prevents all future movement, simulating a crash. The non-differentiable cost function makes this a challenging problem, well-suited for sampling-based approaches. Obstacles lie in an equally spaced 4-by-4 grid, yielding several multi-modal solutions. This is

	Point-mass		Pendulum	
	Cost ($\mu \pm \sigma$)	Succ. [†]	Cost ($\mu \pm \sigma$)	Succ. [‡]
MPPI [§]	—	—	30.8 \pm 12.6	100%
DISCO	250.8 \pm 29.9	20%	61.3 \pm 40.0	70%
SVMPC	191.7 \pm 56.5	25%	44.5 \pm 17.9	70%
DuSt-MPC	118.3 \pm 07.9	100%	36.8 \pm 14.0	80%

Table I: **Simulation results.** Summary of results for simulation experiments. The mean episode cost is given by the sum of the instant costs over the episode length. Values shown do not include the crash penalty for a more comparable baseline. [§]Not used in the navigation task; has perfect knowledge in the pendulum task. [†]Successes are episodes with no crashes. [‡]Successes are episodes whose last five steps have a instant cost below 4 ($\approx 10^\circ$ from the upright position).

depicted in fig. 2. Additionally, we include barriers at the boundaries of the simulated space to prevent the robot from easily circumventing obstacles.

The system dynamics is represented as a double integrator model with non-unitary mass m , s.t. the particle acceleration is given by $\ddot{\mathbf{x}} = m^{-1}\mathbf{u}$ and the control signal is the force applied to the point-mass. In order to demonstrate the inference over dynamics, we forcibly change the mass of the robot at a fixed step of the experiment, adding extra weight. This has a direct parallel to several tasks in reality, such as collecting a payload or passengers while executing a task. Assuming the goal position is denoted by \mathbf{x}_g , the cost function that defines the task is given by:

$$c(\mathbf{x}_t, \mathbf{u}_t) = 0.5\mathbf{e}_t^\top \mathbf{e}_t + 0.25\dot{\mathbf{x}}_t^\top \dot{\mathbf{x}}_t + 0.2\mathbf{u}_t^\top \mathbf{u}_t + p \cdot \mathbf{1}\{\text{col.}\}$$

$$c_{term}(\mathbf{x}_t, \mathbf{u}_t) = 1000\mathbf{e}_t^\top \mathbf{e}_t + 0.1\dot{\mathbf{x}}_t^\top \dot{\mathbf{x}}_t,$$

where $\mathbf{e}_t = \mathbf{x}_t - \mathbf{x}_g$ is the instantaneous position error and $p = 10^6$ is the penalty when a collision happens. A detailed account of the hyper-parameters used in the experiment is presented in appendix E.

As a baseline, we once more compare against DISCO and SVMPC. In fig. 2 we present an overlay of the trajectories for SVMPC and DuSt-MPC over 20 independent episodes and we choose to omit trajectories of DISCO for conciseness. Collisions to obstacles are denoted by a \mathbf{x} marker. Note that in a third of the episodes SVMPC is unable to avoid obstacles due to the high model mismatch while DuSt-MPC is able to avoid collisions by quickly adjusting to the new model configuration online. A typical sequential plot of the posterior distribution induced by fitting a GMM as in eq. (27) is shown on fig. 2c for one episode. There is little variation between episodes and the distribution remains stable in the intermediate steps not depicted.

C. Trajectory tracking with autonomous ground vehicle

We now present experimental results with a physical autonomous ground robot equipped with a skid-steering drive mechanism. The kinematics of the robot are based on a modified unicycle model, which accounts for skidding via an additional parameter [18]. The parameters of interest in this model are

the robot’s wheel radius r_w , axial distance a_w , i.e. the distance between the wheels, and the displacement of the robot’s ICR from the robot’s centre x_{ICR} . A non-zero value on the latter affects turning by sliding the robot sideways. The robot is velocity controlled and, although it possess four-wheel drive, the controls is restricted to two-degrees of freedom, left and right wheel speed. Individual wheel speeds are regulated by a low-level proportional-integral controller.

The robot is equipped with a 2D Hokuyo LIDAR and operates in an indoor environment in our experiments. Prior to the tests, the area is pre-mapped using the gmapping package [15] and the robot is localised against this pre-built map. Similar to the experiment in section V-B, we simulate a change in the environment that could be captured by our parametric model of the robot to explain the real trajectories. However, we are only applying a relatively simple kinematic model in which the effects of the dynamics and ground-wheel interactions are not accounted for. Therefore, friction and mass are not feasible inference choices. Hence, out of the available parameters, we opted for inferring x_{ICR} , the robot’s centre of rotation. Since measuring x_{ICR} involves a laborious process, requiring different weight measurements or many trajectories from the physical hardware [41], this also makes the experiment more realistic. To circumvent the difficulties of ascertaining x_{ICR} , we use the posterior distribution estimated in [3], and bootstrap our experiment with $x_{ICR} \sim \mathcal{N}(0.5, 0.2^2)$.

To reduce the influence of external effects, such as localisation, we defined a simple control task of following a circular path at a constant tangential speed. Costs were set to make the robot follow a circle of 1 m radius with $c(\mathbf{x}_t) = \sqrt{d_t^2 + 10(s_t - s_0)^2}$, where d_t represents the robot’s distance to the edge of the circle and $s_0 = 0.2 \text{ m s}^{-1}$ is a reference linear speed.

The initial particles needed by DuSt-MPC in eq. (26) for the estimation of x_{ICR} are sampled from the bootstrapping distribution, whereas for SVMPC we set $x_{ICR} = 0.5 \text{ m}$, the distribution mean. Again, we want to capture whether our method is capable of adjusting to environmental changes. To this end, approximately halfway through the experiment, we add an extra load of approximately 5.3 kg at the rear of the robot in order to alter its centre of mass. These moments are indicated on the trajectories shown in fig. 4b. In fig. 4a we plot the instant costs for a fixed number of steps before and after the change of mass. For the complete experiment parameters refer to the appendix E.

We observe that considering the uncertainty over x_{ICR} and, crucially, refining our knowledge over it (see fig. 1b in appendix E), allows DuSt-MPC to significantly outperform the SVMPC baseline. Focusing on the trajectories from SVMPC we note that our estimation of x_{ICR} is probably not accurate. As the cost function emphasises the tangential speed over the cross-track error, this results in circles correctly centred, but of smaller radius. Crucially though, the algorithm cannot overcome this poor initial estimation. DuSt-MPC initially appears to find the same solution, but quickly adapts, overshooting the target trajectory and eventually converging to a better result. This

behaviour can be observed both prior to and after the change in the robot’s mass. Conversely, with the addition of mass, the trajectory of SVMPC diverged and eventually led the robot to a halt.

VI. CONCLUSIONS

We present a method capable of simultaneously estimating model parameters and controls. The method expands previous results in control as implicit variational inference and provides the theoretical framework to formally incorporate uncertainty over simulator parameters. By encapsulating the uncertainty on dynamic systems as distributions over a parametric model, we are able to incorporate domain knowledge and physics principles while still allowing for a highly representative model. Crucially, we perform an *online* refinement step where the agent leverages system feedback in a sequential way to efficiently update its beliefs regardless of the size of observation dataset.

Simulated experiments are presented for a randomised inverted pendulum environment and obstacle grid with step changes in the dynamical parameters. Additionally, a trajectory tracking experiment utilising a custom built AGV demonstrates the feasibility of the method for online control. The results illustrate how the simplifications on dynamic inference are effective and allow for a quick adjustment of the posterior belief resulting in a *de facto* adaptive controller. Consider, for instance, the inverted pendulum task. In such periodic and non-linear system, untangling mass and length can pose a difficult challenge as there are many plausible solutions [28]. Nonetheless, the results obtained are quite encouraging, even with very few mapping steps per control loop.

Finally, we demonstrated how, by incorporating uncertainty over the model parameters, DuSt-MPC produces more robust policies to dynamically changing environments, even when the posterior estimation of the model parameters is rather uncertain. This can be seen, for instance, in the adjustments made to trajectories in the obstacle grid.

REFERENCES

- [1] I. Abraham et al. “Model-Based Generalization Under Parameter Uncertainty Using Path Integral Control”. In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 2864–2871.
- [2] T. Akiba et al. “Optuna: A Next-generation Hyperparameter Optimization Framework”. In: *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2019.
- [3] L. Barcelos et al. “DISCO: Double Likelihood-Free Inference Stochastic Control”. In: *Proceedings of the 2020 IEEE International Conference on Robotics and Automation*. ICRA. Paris, France: IEEE Robotics and Automation Society, May 31, 2020, p. 7. DOI: [978-1-7281-7395-5/20](https://doi.org/10.1109/7281-7395-5/20).
- [4] Z. I. Botev, J. F. Grotowski, and D. P. Kroese. “Kernel Density Estimation via Diffusion”. In: *The Annals of Statistics* 38.5 (Oct. 2010), pp. 2916–2957. ISSN: 0090-5364, 2168-8966. DOI: [10.1214/10-AOS799](https://doi.org/10.1214/10-AOS799).
- [5] Z. I. Botev et al. “The cross-entropy method for optimization”. In: *Handbook of statistics*. Vol. 31. Elsevier, 2013, pp. 35–59.
- [6] E. F. Camacho and C. B. Alba. *Model Predictive Control*. 2nd ed. Advanced Textbooks in Control and Signal Processing. London: Springer-Verlag, 2013. ISBN: 978-0-85729-398-5.
- [7] K. Chua et al. “Deep Reinforcement Learning in a Handful of Trials Using Probabilistic Dynamics Models”. In: *Advances in Neural Information Processing Systems*. Ed. by S. Bengio et al. Vol. 31. Curran Associates, Inc., 2018.
- [8] M. Deisenroth and C. Rasmussen. “PILCO: A Model-Based and Data-Efficient Approach to Policy Search”. In: *Proceedings of the 28th International Conference on Machine Learning (ICML)*. Omnipress, 2011, pp. 465–472.
- [9] S. Di Cairano and I. V. Kolmanovsky. “Automotive applications of model predictive control”. In: *Handbook of Model Predictive Control*. Springer, 2019, pp. 493–527.
- [10] Y. Ding et al. “Model predictive control and its application in agriculture: A review”. In: *Computers and Electronics in Agriculture* 151 (2018), pp. 104–117.
- [11] A. Doucet. *Sequential Monte Carlo Methods in Practice*. 2001. ISBN: 978-1-4757-3437-9.
- [12] J. W. Eaton and J. B. Rawlings. “Model-predictive control of chemical processes”. In: *Chemical Engineering Science* 47.4 (1992), pp. 705–720.
- [13] D. Fan, A. Agha, and E. Theodorou. “Deep Learning Tubes for Tube MPC”. In: *Robotics: Science and Systems XVI*. Robotics: Science and Systems 2020. Robotics: Science and Systems Foundation, July 12, 2020. ISBN: 978-0-9923747-6-1. DOI: [10.15607/RSS.2020.XVI.087](https://doi.org/10.15607/RSS.2020.XVI.087).
- [14] J. F. Fisac et al. “A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems”. In: *IEEE Transactions on Automatic Control* 64.7 (July 2019), pp. 2737–2752. ISSN: 0018-9286, 1558-2523, 2334-3303. DOI: [10.1109/TAC.2018.2876389](https://doi.org/10.1109/TAC.2018.2876389).
- [15] G. Grisetti, C. Stachniss, and W. Burgard. “Improved techniques for grid mapping with rao-blackwellized particle filters”. In: *IEEE transactions on Robotics* 23.1 (2007), pp. 34–46.
- [16] J. Harrison, A. Sharma, and M. Pavone. “Meta-learning priors for efficient online bayesian regression”. In: *International Workshop on the Algorithmic Foundations of Robotics*. Springer. 2018, pp. 318–337.
- [17] M. Kamel et al. “Model predictive control for trajectory tracking of unmanned aerial vehicles using robot operating system”. In: *Robot operating system (ROS)*. Springer, 2017, pp. 3–39.
- [18] K. Kozłowski and D. Pazderski. “Modeling and Control of a 4-wheel Skid-steering Mobile Robot”. In: *Int. J. Appl. Math. Comput. Sci.* 14.4 (2004), pp. 477–496.

- [19] A. Lambert et al. “Stein Variational Model Predictive Control”. In: *Proceedings of the 4th Annual Conference on Robot Learning*. CoRL. 2020.
- [20] P. J. Leeuwen et al. “Particle Filters for High-dimensional Geoscience Applications: A Review”. In: *Q.J.R. Meteorol. Soc.* 145.723 (July 2019), pp. 2335–2365. ISSN: 0035-9009, 1477-870X. DOI: [10.1002/qj.3551](https://doi.org/10.1002/qj.3551).
- [21] Q. Liu and D. Wang. “Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm”. In: *Advances in Neural Information Processing Systems*. Ed. by D. Lee et al. Vol. 29. Curran Associates, Inc., 2016.
- [22] M. Okada and T. Taniguchi. “Variational Inference MPC for Bayesian Model-Based Reinforcement Learning”. In: *Proceedings of the Conference on Robot Learning (CoRL)*. Ed. by L. P. Kaelbling, D. Kragic, and K. Sugiura. Vol. 100. Proceedings of Machine Learning Research. PMLR, Oct. 2020, pp. 258–272.
- [23] Y. Pan and E. Theodorou. “Probabilistic differential dynamic programming”. In: *Advances in Neural Information Processing Systems 27* (2014), pp. 1907–1915.
- [24] Y. Pan et al. “Adaptive probabilistic trajectory optimization via efficient approximate inference”. In: *29th Conference on Neural Information Processing System* (2016).
- [25] Y. Pan et al. “Prediction under uncertainty in sparse spectrum Gaussian processes with applications to filtering and control”. In: *Proceedings of the 34th international conference on machine learning*. Ed. by D. Precup and Y. W. Teh. Vol. 70. Proceedings of machine learning research. tex.pdf: <http://proceedings.mlr.press/v70/pan17a/pan17a.pdf>. International Convention Centre, Sydney, Australia: PMLR, Aug. 6, 2017, pp. 2760–2768.
- [26] R. Possas et al. “Online BayesSim for Combined Simulator Parameter Inference and Policy Improvement”. In: International Conference on Intelligent Robots and Systems (IROS). 2020, p. 8.
- [27] M. Pulido and P. J. van Leeuwen. “Sequential Monte Carlo with kernel embedded mappings: The mapping particle filter”. In: *Journal of Computational Physics* 396 (Nov. 2019), pp. 400–415. ISSN: 00219991. DOI: [10.1016/j.jcp.2019.06.060](https://doi.org/10.1016/j.jcp.2019.06.060).
- [28] F. Ramos, R. Possas, and D. Fox. “BayesSim: Adaptive Domain Randomization Via Probabilistic Inference for Robotics Simulators”. In: *Proceedings of Robotics: Science and Systems*. FreiburgimBreisgau, Germany, June 2019. DOI: [10.15607/RSS.2019.XV.029](https://doi.org/10.15607/RSS.2019.XV.029).
- [29] K. Rawlik, M. Toussaint, and S. Vijayakumar. “On stochastic optimal control and reinforcement learning by approximate inference”. In: *Proceedings of Robotics: Science and Systems VIII* (2012).
- [30] D. W. Scott. *Multivariate Density Estimation: Theory, Practice, and Visualization*. 1st ed. Wiley Series in Probability and Statistics. Wiley, Aug. 17, 1992. ISBN: 978-0-470-31684-9. DOI: [10.1002/9780470316849](https://doi.org/10.1002/9780470316849).
- [31] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Boston, MA: Springer US, 1986. ISBN: 978-1-4899-3324-9. DOI: [10.1007/978-1-4899-3324-9](https://doi.org/10.1007/978-1-4899-3324-9).
- [32] Y. Tassa, N. Mansard, and E. Todorov. “Control-limited differential dynamic programming”. In: *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2014, pp. 1168–1175.
- [33] E. Todorov and W. Li. “A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems”. In: *Proceedings of the 2005, American Control Conference, 2005*. IEEE. 2005, pp. 300–306.
- [34] E. Valassakis, Z. Ding, and E. Johns. “Crossing The Gap: A Deep Dive into Zero-Shot Sim-to-Real Transfer for Dynamics”. In: International Conference on Intelligent Robots and Systems (IROS). 2020. arXiv: [2008.06686 \[cs\]](https://arxiv.org/abs/2008.06686).
- [35] K. P. Wabersich and M. Zeilinger. “Bayesian model predictive control: Efficient model exploration and regret bounds using posterior sampling”. In: *Learning for Dynamics and Control*. PMLR. 2020, pp. 455–464.
- [36] N. Wagener et al. “An Online Learning Approach to Model Predictive Control”. In: *Proceedings of Robotics: Science and Systems (RSS)*. 2019.
- [37] J. Watson, H. Abdulsamad, and J. Peters. “Stochastic optimal control as approximate input inference”. In: *Conference on Robot Learning*. PMLR. 2020, pp. 697–716.
- [38] G. Williams et al. “Information-Theoretic Model Predictive Control: Theory and Applications to Autonomous Driving”. In: *IEEE Transactions on Robotics* 34.6 (Dec. 2018), pp. 1603–1622. DOI: [10.1109/TRO.2018.2865891](https://doi.org/10.1109/TRO.2018.2865891).
- [39] G. Williams, A. Aldrich, and E. A. Theodorou. “Model Predictive Path Integral Control: From Theory to Parallel Computation”. In: *Journal of Guidance, Control, and Dynamics* 40.2 (Feb. 2017), pp. 344–357. ISSN: 0731-5090, 1533-3884. DOI: [10.2514/1.G001921](https://doi.org/10.2514/1.G001921).
- [40] G. Williams et al. “Robust Sampling Based Model Predictive Control with Sparse Objective Information.” In: *Robotics: Science and Systems*. 2018.
- [41] J. Yi et al. “Kinematic Modeling and Analysis of Skid-Steered Mobile Robots With Applications to Low-Cost Inertial-Measurement-Unit-Based Motion Estimation”. In: *IEEE Transactions on Robotics* 25.5 (Oct. 2009). Conference Name: IEEE Transactions on Robotics, pp. 1087–1097. ISSN: 1941-0468. DOI: [10.1109/TRO.2009.2026506](https://doi.org/10.1109/TRO.2009.2026506).
- [42] M. Zanon et al. “Model predictive control of autonomous vehicles”. In: *Optimization and optimal control in automotive systems*. Springer, 2014, pp. 41–57.

APPENDIX A
ALGORITHM

Algorithm 1: Sim-to-Real in the loop SVMPC

```

1 Sample  $\{\theta_{t_0}^i\}_{i=1}^{N_\pi} \sim q(\theta_{t_0})$ 
2 foreach policy  $i \in N_\pi$  do  $\pi_{\theta^i} \leftarrow \mathcal{N}(\theta_{t_0}^i, \Sigma_a)$ 
3 while task not complete do
4    $\mathbf{x}_t^r \leftarrow \text{GetStateEstimate}()$ 
5    $\mathcal{D}_{1:t} \leftarrow \mathcal{D}_{1:t-1} \cup \{\mathbf{x}_t^r, \mathbf{x}_{t-1}^r, \mathbf{u}_{t-1}^r\}$ 
6   for  $l \leftarrow 1$  to  $L$  do // Dynamics inference loop
7     for  $m \leftarrow 1$  to  $N_\xi$  do
8        $\ell(\xi|\mathcal{D}_{1:t}) \leftarrow \mathcal{N}(\mathbf{x}_t^r; \hat{f}_\xi(\mathbf{x}_{t-1}^r, \mathbf{u}_{t-1}^r), \Sigma_{\text{obs}})$  // Condition likelihood, eq. (24)
9        $\nabla_\xi \log p(\xi^m|\mathcal{D}_{1:t}) \approx \nabla_\xi \log \ell(\xi^m|\mathcal{D}_{1:t}) + \nabla_\xi \log q(\xi^m|\mathcal{D}_{1:t-1})$  // eq. (26)
10       $\phi(\xi^m) \leftarrow \frac{1}{N_\xi} \sum_{j=1}^{N_\xi} k(\xi^j, \xi^m) \nabla_\xi \log p_t(\xi^m|\mathcal{D}_{1:t}) + \nabla_\xi k(\xi^j, \xi^m)$  // Stein gradient, eq. (11)
11       $\xi^m \leftarrow \xi^m + \epsilon \phi(\xi^m)$ 
12       $q(\xi|\mathcal{D}_{1:t}) = \frac{1}{N_\xi} \sum_{m=1}^{N_\xi} \mathcal{N}(\xi^m, \Sigma_s)$  // Update posterior, eq. (27)
13    end
14  end
15  Sample  $\{\xi^m\}_{m=1}^{N_\xi} \sim q(\xi|\mathcal{D}_{1:t})$ 
16  for  $i \leftarrow 1$  to  $N_\pi$  do // Policies inference loop
17    Sample  $\{U_n^i\}_{n=1}^{N_a} \sim \pi_{\theta^i}$ 
18    foreach  $n \in N_a$  and  $m \in N_s$  do  $C_{n,m}^i \leftarrow \text{GetRolloutCosts}(U_n^i, \xi^m, \mathbf{x}_t^r)$  // eqs. (3) and (5)
19     $\nabla_\theta \log p(\theta_t^i|\mathcal{O}, \xi) \approx \nabla_\theta \log q(\theta_{t-1}^i) + \nabla_\theta \log \left( \frac{1}{N_a N_s} \sum_{n=1}^{N_a} \sum_{m=1}^{N_s} \exp(-\alpha C_{n,m}^i) \right)$  // eq. (20)
20     $\phi(\theta_t^i) \leftarrow \frac{1}{N_\pi} \sum_{j=1}^{N_\pi} k(\theta_{t-1}^j, \theta_t^i) \nabla_\theta \log p(\theta_t^j|\mathcal{O}, \xi) + \nabla_\theta k(\theta_t^j, \theta_t^i)$  // Stein gradient
21     $\theta_t^i \leftarrow \theta_t^i + \epsilon \phi(\theta_t^i)$ 
22     $\omega_t^i \leftarrow \frac{q(\theta_{t-1}^i)}{N_a N_s} \sum_{n=1}^{N_a} \sum_{m=1}^{N_s} \exp(-\alpha C_{n,m}^i)$  // Compute weights, eq. (28)
23  end
24  foreach policy  $i \in N_\pi$  do  $\omega_t^i \leftarrow \frac{\omega_t^i}{\sum_{i=1}^{N_\pi} \omega_t^i}$ 
25   $i^* = \text{argmax}_i \omega_t^i$ 
26   $\text{SendToActuators}(U_{t^*}^{i^*} = \theta_t^{i^*})$  // Applies first action of the control sequence
27   $\theta_t \leftarrow \text{RollPolicies}(\theta_t)$  // Shift one step ahead, adds noise to last step
28   $q(\theta_t) = \sum_{i=1}^{N_\pi} \omega_t^i \mathcal{N}(\theta_t^i, \Sigma)$  // Update policies prior, eq. (21)
29  foreach policy  $i \in N_\pi$  do  $\pi_{\theta^i} \leftarrow \mathcal{N}(\theta_t^i, \Sigma_a)$ 
30   $t \leftarrow t + 1$ 
31 end

```

APPENDIX B
CONSIDERATIONS ON ACTION SELECTION

Once policy parameters have been updated according to eq. (19), we can update the policy for the current step, π_{θ_t} . However, defining the updated policy is not enough, as we still need to determine which immediate control action should be sent to the system. There are many options which could be considered at this stage. One alternative would be to take the expected value at each time-step, although that could compromise the multi-modality of the solution found. Other options would be to consider the modes π_{θ_t} of at each horizon step or sample the actions directly. Finally, we adopt the choice of computing the probabilistic weight of each particle and choosing the highest weighted particle as the action sequence for the current step. More formally:

$$\omega_i = \frac{p(\mathcal{O}|\theta_t^i, \xi)q(\theta_{t-1}^i)}{\sum_{j=1}^m p(\mathcal{O}|\theta_t^j, \xi)q(\theta_{t-1}^j)} \approx p(\theta_t^i|\mathcal{O}, \xi). \quad (28)$$

And finally, the action deployed to the system would be given by $U_t^{i^*} = \theta_t^{i^*}$, where $i^* = \text{argmax}_i \omega_t^i$.

APPENDIX C
CONSIDERATIONS ON COVARIATE SHIFT

The proposed inference method assumes that the parameters of the dynamical model are fixed over time. Note that, even if the distribution over parameters changes over time, this remains a plausible consideration, given that the control loop will likely have a relatively high frequency when compared to the environment covariate shift. This also ensures that trajectories generated in simulation are consistent and that the resulting changes are being governed by the variations in the control actions and not the environment.

However, it is also clear that the method is intrinsically adaptable to changes in the environment, as long as there is a minimum probability of the latent parameter being feasible under the prior distribution $q(\xi|\mathcal{D}_t)$. To see this, consider the case where there is an abrupt change in the environment (e.g. the agent picks-up some load or the type of terrain changes). In this situation, $q(\xi|\mathcal{D}_{t-1})$ would behave as if a poorly specified prior, meaning that as long the probability density around the true distribution $p(\xi)$ is non-zero, we would still converge to the true distribution, albeit requiring further gradient steps.

In practice, the more data we gather to corroborate a given parameter set, the more concentrated the distribution would be around a given location in the parameter space and the longer it would take to transport the probability mass to other regions of the parameter space. This could be controlled by including heuristic weight terms to the likelihood and prior in eq. (22). However, we deliberately choose not to include extra hyper-parameters based on the hypothesis that the control loop is significantly faster than the changes in the environment, which in general allows the system to gather a few dozens or possibly more observations before converging to a good estimate of ξ .

APPENDIX D
BIAS ON JOINT INFERENCE OF POLICY AND DYNAMICS

Note that, if we take the gradient of eq. (16) w.r.t. ξ , we get:

$$\begin{aligned} \nabla_{\xi} p(\theta_t, \xi | \mathcal{O}, \mathcal{D}_{1:t}) &= \nabla_{\xi} [p(\theta_t | \mathcal{O}, \xi) p(\xi | \mathcal{D}_{1:t})] \\ &= p(\xi | \mathcal{D}_{1:t}) \nabla_{\xi} p(\theta_t | \mathcal{O}, \xi) + p(\theta_t | \mathcal{O}, \xi) \nabla_{\xi} p(\xi | \mathcal{D}_{1:t}). \end{aligned} \quad (29)$$

The first term on the RHS of the equation above indicates that the optimality gradient of *simulated* trajectories would contribute to the update of ξ . This implies that the estimation of $p(\xi)$ would be driven to regions that would yield lower cost policies, i.e. possibly due to easier to control dynamics. Naturally, it is important that the likelihood of ξ be taken into account during the policy updates, but we don't want the inference of the *physical* parameters to be biased by our optimality measure. In other words, the distribution over ξ shouldn't conform to whatever would benefit the optimality policy, but the other way around.

APPENDIX E
FURTHER DETAILS ON PERFORMED EXPERIMENTS

Parameter	Inverted Pendulum	Point-mass Navigation	AGV Traj. Tracking
Initial state, \mathbf{x}_0	[3 rad, 0 rad/sec]	—	—
Environment maximum velocity	8 rad/sec	—	—
Environment maximum acceleration/torque	2 N m	—	—
Policy samples, N_a	32	64	50
Dynamics samples, N_s	8	4	4
Cost Likelihood inverse temperature, α	1.0	1.0	1.0
Control authority, Σ	2.0^2	5.0^2	0.1^2
Control horizon, H	20	40	20
Number of policies, N_{π}	3	6	2
Policy Kernel, $k_{\pi}(\cdot, \cdot)$		Radial Basis Function	
Policy Kernel bandwidth selection		Silverman's rule	
Policy prior covariance, Σ_a	2.0^2	5.0^2	1.0^2
Policy step size, ϵ	2.0	100.0	0.02
Dynamics prior distribution	$m = \mathcal{U}(0.6, 1.3), l = \mathcal{U}(0.6, 1.3)$	$\mathcal{N}(2, 0.1^2)$	$\mathcal{N}(0.5, 0.2^2)$
Dynamics number of particles, N_{ξ}	50	50	50
Dynamics Kernel, $k_{\xi}(\cdot, \cdot)$		Radial Basis Function	
Dynamics GMM covariance, Σ_s	Improved Sheather Jones	0.25^2	0.0625^2
Dynamics likelihood covariance, Σ_{obs}	0.1^2	0.1^2	0.1^2
Dynamics update steps, L	20	20	5
Dynamics step size, ϵ	0.001	0.01	0.05
Dynamics in log space	No	Yes	No
Unscented Transform spread [3], α	0.5	—	—

Table II: Hyperparameters used in the experiments.