

# DISCO: Double Likelihood-free Inference Stochastic Control

Lucas Barcelos<sup>\*1</sup>, Rafael Oliveira<sup>1</sup>, Rafael Possas<sup>1</sup>, Lionel Ott<sup>1</sup>, and Fabio Ramos<sup>1,2</sup>

**Abstract**—Accurate simulation of complex physical systems enables the development, testing, and certification of control strategies before they are deployed into the real systems. As simulators become more advanced, the analytical tractability of the differential equations and associated numerical solvers incorporated in the simulations diminishes, making them difficult to analyse. A potential solution is the use of probabilistic inference to assess the uncertainty of the simulation parameters given real observations of the system. Unfortunately the likelihood function required for inference is generally expensive to compute or totally intractable. In this paper we propose to leverage the power of modern simulators and recent techniques in Bayesian statistics for likelihood-free inference to design a control framework that is efficient and robust with respect to the uncertainty over simulation parameters. The posterior distribution over simulation parameters is propagated through a potentially non-analytical model of the system with the unscented transform, and a variant of the information theoretical model predictive control. This approach provides a more efficient way to evaluate trajectory roll outs than Monte Carlo sampling, reducing the online computation burden. Experiments show that the controller proposed attained superior performance and robustness on classical control and robotics tasks when compared to models not accounting for the uncertainty over model parameters.

## I. INTRODUCTION

Robustness to model miss-specification and noisy sensor measurements is a critical property for control systems operating in complex robotics applications. The development of powerful and more realistic simulators allows practitioners to analyse and verify the performance of the controller against these variables before the controller is deployed to the real robot. In Model Predictive Control (MPC) one seeks to iteratively find the solution of an optimisation problem for a receding finite time-horizon using an approximate model of the system. When the dynamic model is given by complex simulators that incorporate differential equations and numerical solvers there is little hope the equations can be reversed to reason about the parameters of the simulation to best match the real behaviour of the system. Furthermore, the simulator might abstract away the equations and solver from the user. Effectively, it can be interpreted as a generative model that can be sampled from given a set of parameter values, but not inverted. In this paper we pose the question, can we leverage the power of simulators, treated as generative models, to design control strategies that are robust to parameter uncertainty?

On the other hand, the application of MPC to linear systems has been an active research area for many decades

with extensive deployments to many practical problems [1]–[3]. Notably, the most common setting for linear MPC application are tasks that involve trajectory tracking or stabilisation. However, control tasks in reinforcement learning are usually more complex and therefore less suitable to linearisation, motivating the use of non-linear models [4]. Another motivation for more complex models is the ability to use of more expressive constraints, even if not directly involved in the physical process, such as economic criteria [5]. Despite its vast application in the linear case, the use of MPC in non-linear systems continues to be an increasingly active area of research in control theory [2], [3].

Recent work in the field has led to controllers that are able to incorporate non-linear dynamics without relying on linear or quadratic approximations [6], [7]. However, most MPC controllers still do not consider uncertainty in the parameters of their internal simulator for future trajectories. In addition, estimating parameters for the system’s model usually requires large amounts of data from the real system, which can be infeasible for some applications. Yet, whenever the stochastic system uncertainties can be adequately modelled, it is more natural to explicitly take them into account in the control design method itself. In Stochastic MPC, the uncertainty on the internal system dynamics is intrinsic to the optimal control problem solved at every time step. This allows the controller to trade-off performance and satisfaction of the constraints by regulating the joint probability distribution of the system states and outputs [3].

In this paper we make the following contributions: we develop a Stochastic Non-linear MPC variant which leverages recent advancements in likelihood-free inference to estimate both the uncertainty on the simulator parameters as well as to propagate it throughout the estimated trajectories. We call our method double likelihood-free inference stochastic control, DISCO. The posterior distribution for the parameters of the simulator is estimated by combining simulated data from generative models and observations from the physical system. Using this posterior distribution allows us to take into account the uncertainty about the system’s dynamics in the decision making process during the control task. We proceed to show that the Unscented Transform (UT) [8] provides a computationally efficient alternative when compared to traditional Monte Carlo approaches to propagate the uncertainty from the parameter space to the forward modelling of the trajectory roll outs. In short, DISCO can be seen as a variant of the Information Theoretical MPC (IT-MPC) control algorithm [7] that considers the uncertainty in the system’s parameters in its internal trajectory simulations.

\* Corresponding author: lucas.barcelos@sydney.edu.au

<sup>1</sup> School of Computer Science, The University of Sydney, Australia

<sup>2</sup> NVIDIA, USA

## II. RELATED WORK

The use of MPC in the control of linear systems is very mature and has been widely studied and applied to real systems. However, as seen in [2], Non-linear MPC (NMPC) is still an open-research question, especially for systems where uncertainty over parameters and constraints on controls and state-space are considered. The most common methods for controlling general nonlinear systems are based on Non-linear Programming [9] and Differential Dynamic Programming (DDP) [10]. Both rely on approximations of dynamics and cost functions so that the online optimisation problem becomes tractable. However, these mainstream gradient-based MPC approaches have some shortcomings. In the DDP method, the cost function must be smooth and it is notoriously difficult to include state constraints. Whereas with nonlinear programming constraints may be easily accounted for, but a common issue is what to do when no feasible solution is found.

In [3], a family of Stochastic NMPC (SNMPC) methods are discussed. In Tube-based NMPC the objective of the control policy is to ensure that the forward trajectories will remain inside a desirable tube centred around a given trajectory, however the boundary tube has to be computed offline [11]. A multi-stage NMPC approach has been suggested in which the uncertainty is modelled by a scenario tree approach from stochastic programming. However, the procedure quickly becomes intractable, since the size of the optimisation problem scales exponentially with the time horizon, number of uncertainties and uncertainty levels [12].

Although many of the methods above focus on robustness, they do not incorporate uncertainty over the parameters of the transition function. In [5], this is accounted for by using a SNMPC with an Unscented Kalman Filter to propagate the uncertainty over the state-space. However, this method requires an optimisation with chance constraints to be solved online and, to keep the problem feasible, the variance of the trajectories has to be artificially constrained. The most similar approach is perhaps presented in [13], where the IT-MPC formulation is used in conjunction with a Ensemble of Mixture Density Networks (E-MDN) to approximate the joint probability distribution of states and actions. This is similar to our approach, however as the E-MDN tries to approximate the joint distribution of states and actions, it needs to be retrained entirely on new environments.

In contrast, the variant of IT-MPC proposed in this paper uses the UT to propagate the uncertainty over model parameters. This reduces the dimensionality of the inference problem and results in a controller more adept to generalise to unseen situations. Moreover, unlike the stochastic optimisation strategies, our framework is very amenable to the inclusion of constraints, as the control update law is based on sampled trajectories. As shown in [7] constraints may be applied directly to the control actions. On the other hand, we can apply soft constraints to the state space through the cost function. This is easily achieved as there is no need for the cost function to be differentiable and assures that a feasible

solution will exist.

Additionally, DISCO takes advantage of the BayesSim Likelihood-free Inference (LFI) framework presented in [14] to update the model uncertainty periodically. Hence, given a set of true observations after a specified episode length, we can update our knowledge of the posterior probability density of parameters  $p(\theta|\mathbf{x} = \mathbf{x}^r)$ . This way our model can adapt to variations in the environment, e.g. adjust friction coefficients in case of rain, or intrinsic to the transition function, e.g. change of weight distribution. In contrast to other inference methods, such as Variational Inference or Markov Chain Monte Carlo, where a likelihood function is needed, in LFI we compute an approximated parametric distribution of the true posterior. Furthermore, BayesSim was shown to be more data efficient than other LFI methods, such as Approximate Bayesian Computation [14].

## III. PRELIMINARIES

We consider the problem of controlling a discrete-time stochastic system described by a non-linear set of differential equations of the form:

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

where  $f$  is the transition function,  $\mathbf{x}_t \in \mathbb{R}^n$  denotes the system states, and  $\mathbf{v}_t \sim \mathcal{N}(\mathbf{u}, \Sigma) \in \mathbb{R}^m$  is the control input at a given time  $t$ . We assume a finite time-horizon  $T$ , and that the control frequency is given. Note that there is no direct control over the variable  $\mathbf{v}$ , but we are able to control its mean  $\mathbf{u}$ . This assumption considers not only a multiplicative noise model which is common in robotics, where lower-level actuator controllers are usually present, but also an amount of exploration in our control actions. As such, in practice,  $\Sigma$  is a hyper-parameter of our control system that may need to be artificially increased.

More generally, we are interested in the problem where the real transition function  $f(\mathbf{x}, \mathbf{v})$  is approximated by a parameterised non-linear forward model  $f(\mathbf{x}, \mathbf{v}, \theta)$ , represented as  $f_\theta$  for compactness. Equation (1) may then be rewritten as:

$$\mathbf{x}_{t+1} = f_\theta(\mathbf{x}_t, \mathbf{v}_t). \quad (2)$$

### A. Information Theoretical MPC (IT-MPC)

Following the steps in [7], we can define a fixed length input sequence  $U = (\mathbf{u}_0, \dots, \mathbf{u}_{T-1})$  over a fixed control horizon  $T$ , onto which we apply a Receding Horizon Control strategy. This yields  $V = (\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{T-1}) \in \mathbb{R}^{m \times T}$ , which is itself a random variable. Furthermore, let's denote as  $\mathbb{P}$  the joint probability distribution and  $p$  the corresponding probability density function (pdf) of the uncontrolled system (i.e.  $U \equiv 0$ ). Likewise,  $\mathbb{Q}$  is the joint distribution and  $q$  the corresponding pdf for an open-loop control sequence. The optimal control problem may then be defined as:

$$U^* = \underset{U \in \mathcal{U}}{\operatorname{argmin}} \mathbb{E}_{\mathbb{Q}} \left[ \phi(\mathbf{x}_T) + \sum_{t=0}^{T-1} \mathcal{L}(\mathbf{x}_t, \mathbf{u}_t) \right], \quad (3)$$

where  $\mathcal{U}$  is the set of admissible controls,  $\phi(\mathbf{x}_T)$  is a terminal cost function, and  $\mathcal{L}(\mathbf{x}_t, \mathbf{u}_t)$  is a running cost function of the form:

$$\mathcal{L}(\mathbf{x}_t, \mathbf{u}_t) = c(\mathbf{x}_t) + \frac{\lambda}{2} \left( \mathbf{u}_t^\top \Sigma^{-1} \mathbf{u}_t + \boldsymbol{\beta}_t^\top \mathbf{u}_t \right), \quad (4)$$

where  $\lambda \in \mathbb{R}^+$  is known as the inverse temperature and the affine term  $\boldsymbol{\beta}$  allows the location of the minimum control (rest position) to be different from zero. Noting that the state cost may be considered independent from the control terms, we can define  $C(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T) = \phi(\mathbf{x}_T) + \sum_{t=0}^{T-1} c(\mathbf{x}_t)$ . Moreover, we define a mapping operator,  $\mathcal{H}$ , from input sequences  $V$  to their resulting trajectory by recursively applying  $f_\theta$  given  $\mathbf{x}_0$ ,  $\mathcal{H}(V; \mathbf{x}_0, \boldsymbol{\theta}) = [\mathbf{x}_0, f_\theta(\mathbf{x}_0, \mathbf{v}_0), f_\theta(f_\theta(\mathbf{x}_0, \mathbf{v}_0), \mathbf{v}_1), \dots]$ . This leads to the following state cost function:

$$S(V; \mathbf{x}_0, \boldsymbol{\theta}) = S(V) = C(\mathcal{H}(V)). \quad (5)$$

Finally, IT-MPC relies on the *free-energy* principle to compute a lower bound for the optimal control problem and defines the form of the optimal distribution function  $q^*(V)$  for which this bound is tight and achieves the optimal control  $U^*$ . It can be shown that such distribution is of the form

$$q^*(V) = \frac{1}{\eta^*} \exp\left(-\frac{1}{\lambda} S(V)\right) p(V) \quad (6)$$

$$\eta^* = \int_{\mathbb{R}^{m \times T}} \exp\left(-\frac{1}{\lambda} S(V)\right) p(V) dV, \quad (7)$$

where the base distribution  $p(V)$  has been augmented with the cost of the state trajectory. This results in  $\mathbf{u}_t^* = \mathbb{E}_{Q^*}[\mathbf{v}_t] \forall t \in \{0, 1, \dots, T-1\}$ . Therefore, the optimal open-loop control sequence is the expected value of control trajectories sampled from the optimal distribution. As we cannot sample directly from  $Q^*$ , we can resort to importance sampling [15] to construct an unbiased estimator of the optimal distribution, given the current control distribution, namely

$$\mathbb{E}_Q[\mathbf{v}_t] = \int q^*(V) \mathbf{v}_t dV = \int \omega(V) q(V|\hat{U}, \Sigma) \mathbf{v}_t dV, \quad (8)$$

where  $\omega(V) = q^*(V)/q(V|\hat{U}, \Sigma)$  is the importance sampling weight. Therefore, we can switch the expectation to  $\mathbb{E}_{Q_{\hat{U}}}$ , resulting in  $\mathbb{E}_{Q_{\hat{U}}}[\omega(V) \mathbf{v}_t]$ . We can then use the definition of the optimal distribution w.r.t. the base measure distribution given in [7] to derive the *optimal information-theoretic control law*:

$$\omega(V) = \frac{1}{\eta} \exp\left(-\frac{1}{\lambda} \left( S(V) + \lambda \sum_{t=0}^{T-1} \mathbf{u}_t^\top \Sigma^{-1} \mathbf{v}_t \right)\right) \\ \mathbf{u}_t = \mathbb{E}_{Q_{\hat{U}}}[\omega(V) \mathbf{v}_t], \quad (9)$$

where:

$$\eta = \int \exp\left(-\frac{1}{\lambda} \left( S(V) + \lambda \sum_{t=0}^{T-1} \mathbf{u}_t^\top \Sigma^{-1} \mathbf{v}_t \right)\right) dV \quad (10)$$

and  $\mathbf{u}_t = (\hat{\mathbf{u}}_t - \tilde{\mathbf{u}}_t)$  is the difference between the current control action  $\hat{\mathbf{u}}_t$  and the minimum control  $\tilde{\mathbf{u}}_t$  (adjusted by  $\boldsymbol{\beta}$  and usually zero). Note that in practice, for numerical stability, we multiple the numerator and denominator of  $\omega(V)$  by a factor  $\exp(\frac{1}{\lambda}\rho)$ , where  $\rho$  is defined as the minimum cost.

## B. Likelihood-free parameter estimation

Recent advances in LFI allowed the use of probabilistic inference to learn distributions over simulation parameters [14]. The main idea is that of approximating an intractable posterior  $p(\boldsymbol{\theta}|\mathbf{x} = \mathbf{x}^r)$  using data generated from a generative forward model (or simulator) where trajectories are collected for different simulation configurations. Therefore, one can directly learn a conditional density  $q_\phi(\boldsymbol{\theta}|\mathbf{x})$  where parameters  $\phi$  are learned through an optimisation procedure. The learned model usually takes the form of a mixture of Gaussians where inputs are summary statistics obtained from trajectories and outputs are the parameters of the mixture.

The goal is to maximise the likelihood  $\prod_n q_\phi(\boldsymbol{\theta}_n|\mathbf{x}_n)$ . It has been shown in previous work [14] that  $q_\phi(\boldsymbol{\theta}|\mathbf{x})$  will be proportional to  $\frac{\hat{p}(\boldsymbol{\theta})}{p(\boldsymbol{\theta})} p(\boldsymbol{\theta}|\mathbf{x})$  if the log-likelihood is optimised as follows:

$$\mathcal{L}(\phi) = \frac{1}{N} \sum_n \log q_\phi(\boldsymbol{\theta}_n|\mathbf{x}_n) \quad (11)$$

Consequently, a posterior estimate can be obtained by:

$$\hat{p}(\boldsymbol{\theta}|\mathbf{x} = \mathbf{x}^r) \propto \frac{\hat{p}(\boldsymbol{\theta})}{\hat{p}(\boldsymbol{\theta})} q_\phi(\boldsymbol{\theta}|\mathbf{x} = \mathbf{x}^r). \quad (12)$$

The conditional density  $q_\phi(\boldsymbol{\theta}|\mathbf{x})$  is a mixture of  $K$  Gaussians,

$$q_\phi(\boldsymbol{\theta}|\mathbf{x}) = \sum_{k=1}^K K \alpha_k(\mathbf{x}) \mathcal{N}(\boldsymbol{\theta}|\boldsymbol{\mu}(\mathbf{x})_k, \boldsymbol{\Sigma}_k(\mathbf{x})), \quad (13)$$

where  $\{\alpha_k(\mathbf{x})\}_{k=1}^K$  are mixing functions,  $\{\boldsymbol{\mu}_k(\mathbf{x})\}_{k=1}^K$  are mean functions and  $\{\boldsymbol{\Sigma}_k(\mathbf{x})\}_{k=1}^K$  are covariance functions.

## IV. DISCO

At its core, model-based control relies on an approximated transition function to optimise the control actions over the control horizon. In practice, this transition function is usually defined *a priori* using fundamental physical principles and domain knowledge, or empirically by applying system identification techniques [16] or learning methods from data [16]–[18]. Typically, these methods provide deterministic transition functions that do not incorporate model uncertainty and are invariant over time. As discussed in [6], the closed-loop RHC offers a degree of robustness to model uncertainties, but the compounding error of poor predictions along the control horizon will reduce the stability margins of the system. Using the methods outlined in [section III](#), in this paper we propose a framework to apply the IT-MPC stochastic control formulation to problems where the parameters of the transition function  $f_\theta$  are unknown, but belong to a problem dependent *prior*,  $p(\boldsymbol{\theta})$ . Furthermore, we make use of BayesSim [14] to refine our knowledge of the parameters as we interact with the environment and gather new observations. The intuition behind this approach is that, by refining our knowledge of the parameters of an otherwise well-defined transition function, we will capitalise not only on the application domain knowledge, but also on the adaptability of inference-free learning methods. Since  $\boldsymbol{\theta}$  represents a plausible range of unknown physical parameters, e.g. mass or friction coefficient,

it is straightforward to incorporate domain knowledge to this formulation. Alternatively, an improper uninformative prior may be used when no assumptions are given. On the other hand, by updating our knowledge of  $p(\boldsymbol{\theta}|\mathbf{x} = \mathbf{x}^r)$  given observed data, we are more likely to cope with problems such as *covariate shift* [19] and *reality gap* [14], [20]. The complete method is presented in [algorithm 1](#).

#### A. Problem setup

Given a forward model with parameters  $\boldsymbol{\theta}$  and distributed according to  $p(\boldsymbol{\theta})$ , trajectories can be obtained from it by first sampling  $\boldsymbol{\theta}$  and generating roll outs by propagating the state-action pairs through the transition function. Although the parameters are stochastic, we assume they are invariant throughout the control horizon for a given trajectory. This is a reasonable assumption as the latent parameters are usually stable physical quantities and the update frequency of the control loop is significantly faster than their time constants. In this situation, the optimal distribution given in (6) becomes:

$$q^*(V, \boldsymbol{\theta}) = \frac{1}{\eta} \exp\left(-\frac{1}{\lambda} S_{\boldsymbol{\theta}}(V)\right) p(V|\boldsymbol{\theta}) p(\boldsymbol{\theta}), \quad (14)$$

where we overload the notation to emphasise the dependence of  $S(V)$  on the now stochastic  $\boldsymbol{\theta}$ . However, as  $V$  and  $\boldsymbol{\theta}$  are independent, we can drop the conditioning in  $p(V|\boldsymbol{\theta}) = p(V)$ . As a result, our control law can be expressed as:

$$\mathbf{u}_t = \mathbb{E}_{\mathbb{Q}_{\hat{V}}} [\mathbb{E}_{\mathbb{Q}_{\boldsymbol{\theta}}} [\omega_{\boldsymbol{\theta}}(V) \mathbf{v}_t]] = \mathbb{E}_{\mathbb{Q}_{\hat{V}, \boldsymbol{\theta}}} [w_{\boldsymbol{\theta}}(V) \mathbf{v}_t], \quad (15)$$

where  $\omega_{\boldsymbol{\theta}}$  shows the dependence on  $\boldsymbol{\theta}$  and we applied the law of total expectation to get the resulting equivalence. This means our update rule now has to sample jointly from the distributions of  $V$  and  $\boldsymbol{\theta}$ .

#### B. Propagation of uncertainty over the state-space dynamics

If we sample sufficiently from  $p(V, \boldsymbol{\theta})$ , we are able to reconstruct the joint distribution  $q(V, \boldsymbol{\theta})$  and compute our control updates. However, we note that the increased dimensionality of the sample space requires the number of samples to grow combinatorially. As such, we resort to the unscented transform [8] as more efficient approach to propagate the uncertainty of  $\boldsymbol{\theta}$  throughout the state-space.

In [8], the authors demonstrate how UT is able to reconstruct an approximate  $Y'(x)$  of the random variable  $Y = g(X)$  resulting when an original random variable  $X$  is applied to a non-linear function  $g$ . The premise behind this approach is that it should be easier to approximate a probability distribution than a arbitrary non-linear transformation. The idea is to select a set of *sigma points* able to capture the most important statistical properties of the prior random variable  $X$ . The necessary statistical information captured by the UT is the first and second order moments of  $p(X)$ . The number of sigma-points needed to do this  $L = 2n + 1$ , where  $n$  is the dimension of  $X$ . In [21], it is shown that matching the moments of  $X$  up to the  $n$ th order implies matching the moments of  $Y$  to the same order. By using a larger number of sigma-points, skew and kurtosis can also be captured [22].

In DISCO, we refer to the formulation presented in [21] to compute sigma-points over the distribution  $p(\boldsymbol{\theta})$  of parameters. The expressions to compute the sigma-points and weights for the mean and covariance are presented below:

$$\begin{aligned} \boldsymbol{\chi}_0 &= \bar{\boldsymbol{\theta}} & \boldsymbol{\chi}_i &= \bar{\boldsymbol{\theta}} + (\sqrt{(n+\nu)\Sigma_{\boldsymbol{\theta}}})_{1 \leq i \leq n} \\ \varpi_0^m &= \frac{\nu}{n+\nu} & \boldsymbol{\chi}_i &= \bar{\boldsymbol{\theta}} - (\sqrt{(n+\nu)\Sigma_{\boldsymbol{\theta}}})_{n+1 \leq i \leq 2n} \\ \varpi_0^c &= \varpi_0^m + (1 - \nu^2 + \xi) & \varpi_i^m &= \varpi_i^c = \frac{1}{2(n+\nu)}, \end{aligned} \quad (16)$$

where  $\nu = \alpha^2(n + \kappa) - n$  is the primary scaling factor,  $\kappa$  a secondary scaling (usually 0),  $\alpha$  determines the spread of the sigma points around  $\bar{\boldsymbol{\theta}}$ , and  $\xi$  is a scalar to provide an extra degree freedom. The reader is encouraged to refer to [21] for details on hyperparameter selection. The sigma-points are then applied recursively to the transition function  $f_{\boldsymbol{\theta}}$  to compute the cost  $S_{\boldsymbol{\theta}}(V; \boldsymbol{\theta} = \boldsymbol{\chi}_i)$  for  $i \in \{1, \dots, L\}$ . In practice, the sigma points on the state space are given by  $\boldsymbol{\gamma} = \mathcal{H}(\boldsymbol{\chi})$ . To ensure the trajectory cost of each sigma-point can be summarised using the UT, it is necessary to apply the same action sequence  $V$  sampled i.i.d. to all points. Effectively, this means we need to replicate  $L$  times each action sequence  $V$  during our update step. Finally, the mean trajectory cost is given by

$$S(V) = \sum_{i=0}^L \varpi_i^m S_{\boldsymbol{\theta}}(V; \boldsymbol{\theta} = \boldsymbol{\gamma}_i), \quad (17)$$

and used in (9) for the control law update.

#### C. Updating the parameter prior distribution

At each time step we are computing a new control action  $\mathbf{u}_t$ , applying it to our environment and collecting new observations  $\mathbf{x}_{t+1}^r$ . The pairs of  $[\mathbf{u}_t, \mathbf{x}_t^r]$  represent a trajectory  $\boldsymbol{\tau}$ , up to a specified time-length. This serves as input for the estimate of the posterior probability of  $p(\boldsymbol{\theta}|\boldsymbol{\tau})$ . Once sufficient data has been aggregated in  $\boldsymbol{\tau}$ , we can use the method presented in [14] to refine the posterior estimate.

Finally, the unscented transform requires as an input a mean vector  $\bar{\boldsymbol{\theta}}$  and covariance matrix  $\Sigma_{\boldsymbol{\theta}}$  for the parameters. Therefore, these have to be retrieved from  $q_{\phi}(\boldsymbol{\theta}|\boldsymbol{\tau})$ , or, alternatively, the highest weighted Gaussian may be selected if it is above a specified threshold.

## V. EXPERIMENTAL RESULTS

#### A. Inverted pendulum swing-up task

In this task, the controller has to swing and hold a pendulum upright using a torque command applied directly to the joint of a rigid-arm. We used the simulator in [23], and always set the pendulum initial state to the downright position and at rest. The state cost function used was  $c = 50 \cos(\theta - 1)^2 + \dot{\theta}^2$ , and the terminal cost function  $\phi$  was set to zero. The inverse control temperature  $\lambda$  was set at 10 and the control authority  $\Sigma$  at 1. We have also defined the number of sampled trajectories  $K = 500$  and the control horizon  $T = 30$ . For more details on the experiment parameters refer to the appendix in [24].

---

**Algorithm 1: DISCO**

---

**Control Hyperparameters:**  $\lambda, \Sigma, \beta, c, \phi$ ;**UT Hyperparameters:**  $\nu, \kappa, \alpha, \xi$ ;**Given:**  $f_{\theta}, p(\theta), U_0, T, K, L, \tau$ ;

Update posterior distribution;

 $q_{\phi}(\theta|\tau) \leftarrow \text{BayesSim}(\tau)$ ; $p(\theta) \leftarrow q_{\phi}(\theta|\tau)$ ;**while** task not complete **do**     $\mathbf{x}_0 \leftarrow \text{GetStateEstimate}()$ ;    **for**  $k \leftarrow 0$  **to**  $K - 1$  **do**        Sample  $\mathcal{E}^k = (\epsilon_0^k \dots \epsilon_{T-1}^k), \epsilon_t^k \sim \mathcal{N}(0, \Sigma)$ ;        **for**  $i \leftarrow 1$  **to**  $L$  **do**             $\theta_i \leftarrow \theta \sim p(\theta)$  (MC) or  $\mathcal{X}_i$  (UT);             $\mathbf{x} \leftarrow \mathbf{x}_0$ ;            **for**  $t \leftarrow 1$  **to**  $T$  **do**                 $\mathbf{x} \leftarrow f_{\theta}(\mathbf{x}, \mathbf{v}_t, \theta_i)$ ;                 $S_i^k += c(\mathbf{x}) + \lambda \mathbf{u}_{t-1}^T \Sigma^{-1} (\epsilon_{t-1})$ ;            **end**             $S_i^k += \phi(\mathbf{x})$ ;        **end**         $S^k = \sum_{i=1}^L \varpi_i^m S_i^k$ ;    **end**     $\rho \leftarrow \min(S^k)$ ;     $\eta \leftarrow \sum_{k=1}^K \exp(-\frac{1}{\lambda} (S^k - \rho))$ ;    **for**  $k \leftarrow 1$  **to**  $K$  **do**         $\omega_k \leftarrow \frac{1}{\eta} \exp(-\frac{1}{\lambda} (S^k - \rho))$ ;    **end**    **for**  $t \leftarrow 0$  **to**  $T - 1$  **do**         $\mathbf{u}_t += \sum_{k=0}^{K-1} \omega_k (\mathcal{E}^k) \epsilon_t^k$ ;    **end**    SendToActuators( $\mathbf{u}_0$ );    Append( $\tau, [\mathbf{x}_0, \mathbf{u}_0]$ );    RollControlActions( $\mathbf{u}$ );**end**

---

The results presented in **Figure 1** are the accumulated cost over time for 50 iterations, for a baseline case, DISCO, and Monte Carlo sampling. Note that the oscillatory behaviour of the cost function is expected, as the controller has insufficient authority to swing the pendulum without the swinging action to increase the momentum.

All models used the same hyperparameters described above, with the exception of the  $T$  for the case of MC sampling. Given that we want to compare the performance of the controller when using UT against MC, for the case where MC is used, we increment the amount of trajectories sampled by the number of sigma points  $L$  used by the UT. Effectively, for the MC controller we have  $K = 2500$  trajectories.

The unknown parameters in this example were the length of the arm and the mass of the pendulum. As a prior, we assumed an uniform distribution between 0.1 and 5 for both parameters. The posterior distribution was given by a mixture of Gaussians with 5 components, trained using a reference control policy. Note that in our simulations, both models

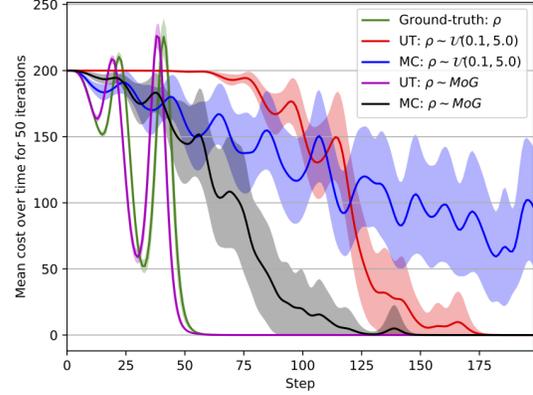


Fig. 1: Mean cost over time for the inverted pendulum experiment. Shaded area represents one  $\sigma$  of standard deviation. Three models were evaluated: a standard IT-MPC with access to the true system parameters (in green); DISCO using unscented transform with a prior distribution over parameters (in red) and with an updated posterior distribution (in magenta); and DISCO using MC sampling with a prior distribution (in blue) and an updated posterior (in black).

shared the same posterior distribution estimate. Once trained and conditioned on the observed data, the resulting mixture had a mean estimate for the length of 0.89 meter and for mass 0.90 kilo. The covariance matrix was diagonal, and the variance was 0.01 for the length estimate and 0.03 for mass. One of the components of the mixture was dominant with a weight of 0.979 and was used as reference for the UT.

DISCO with UT outperforms MC sampling both with an uninformative prior and inferred posterior. Noticeable also, the performance of UT with the posterior distribution is better than the baseline model. This is explainable by the fact that the parameter randomisation introduced by the sigma-points provides more information in the trajectory evaluation. This way, trajectories that are borderline to a higher cost state captured by one of the sigma points get penalised. Effectively, UT works like an automatic calibration of the control temperature, when the prior is broad, many trajectories are considered in the control update average. Conversely, when the posterior gets refined, the controller is more confident to select fewer trajectories.

### B. Skid-steer robot

This section presents experimental results with a physical robot equipped with a skid-steering drive mechanism (**Figure 2**). We modelled the kinematics of the robot based on a modified unicycle model, which accounts for skidding via an additional parameter [25]. The parameters to be estimated via BayesSim 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 instant centre of rotation (ICR) from the robot's centre  $x_{ICR}$ . A non-zero value on the latter affects turning by sliding the robot sideways. To estimate the

parameters, the robot was driven manually around a circle and had its trajectory data recorded. From the trajectory data we computed cross-correlation summary statistics as  $(\bar{x}, \bar{y}, \Delta x, \Delta y)$ , which capture the centre of trajectory and the average linear velocity. In simulation, the wheel speed commands sent to the robot were repeated  $N = 1000$  for different parameter settings sampled from a uniform prior,  $x_{ICR} \sim [0, 0.5]$ ,  $r_w \sim [0, 0.5]$ ,  $a_w \sim [0.1, 0.5]$ .

Figure 2b presents the resulting marginal estimates from BayesSim for each parameter of the robot’s kinematic model. For comparisons, physical measurements indicate a  $r_w$  of around 0.06 m and  $a_w$  of around 0.31 m. Measuring  $x_{ICR}$ , however, involves a laborious process, which would require different weight measurements or many trajectories from the physical hardware [26]. As we are only applying a relatively simple kinematic model of the robot to explain the real trajectories, the effects of the dynamics and ground-wheel interactions are not accounted for. As a result we see that BayesSim automatically estimates some parameters, such as the axial distance to try to compensate for the miss-specifications, which helps in the control task.

The control task was defined as following a circular path at a constant tangential speed. Costs were set to make the robot follow a circle of 0.75 m radius with  $c(\mathbf{x}_t) = \sqrt{d_t^2 + (s_t - s_0)^2}$ , where  $d_t$  represents the robot’s distance to the edge of the circle and  $s_0 = 0.2$  m/s is a reference linear speed. We performed experiments sampling from the uniform prior over the parameters  $p(\theta)$ , sampling from the posterior  $q(\theta|\tau)$ , and using only a point estimate set to  $x_{ICR} = 0.12$ ,  $r_w = 0.06$  and  $a_w = 0.47$ , which was adjusted offline to reduce simulation error. For clarity, instead of the noisy raw costs, we present the mean instant cost, i.e.  $\bar{c}_t = \frac{1}{t} \sum_{i=1}^t c(\mathbf{x}_i)$  and the executed trajectories in Figure 3. For the complete experiment parameters refer to the appendix in [24].

We see that considering parameter uncertainty via DISCO provides significant performance improvements over the baseline MPPI algorithm running with a point estimate. Although, in term of costs, both the prior and posterior estimates offer similar performance, we see the advantages of using the parameter posterior estimates in the trajectories plot, where we see overshooting happening on some portions of the path. The latter can be explained by the prior allowing kinematic parameters candidates that are too far from the true values. Additionally, refining the posterior distribution allows the system to adapt to new configurations or drifts in the model parameters. Lastly, a noisier speed control explains the gap between the baseline MPPI and the DISCO methods, despite the similar performances in terms of path tracking.

## VI. CONCLUSION

This paper is a first step towards incorporating model uncertainty and sophisticated Bayesian inference methods to stochastic model based control. We showed how uncertainty over parameters may be formally incorporated into an stochastic non-linear MPC controller and evaluated methods of propagating the uncertainty into trajectory roll outs. This extension to information theoretical MPC provides

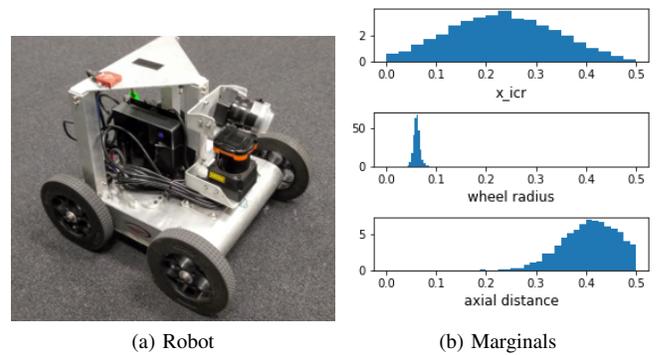


Fig. 2: Skid-steer robot and its parameter estimates

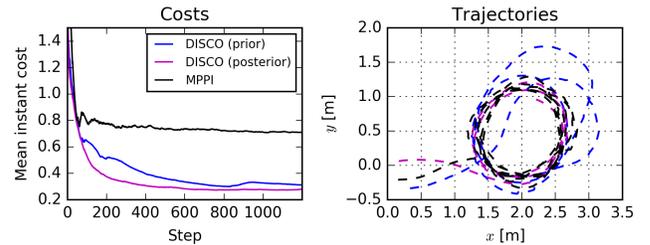


Fig. 3: Results with physical robot

the building blocks of an adaptive controller framework, more resilient to issues arising from reality gap and covariate shift. As shown in the robotic experiments, incorporating uncertainty may lead to a more accurate assessment of the environment and increase the performance.

The unscented transform proved an efficient way to propagate uncertainty, reducing the burden of sampled trajectories. When combined with the ability to impose hard-penalties on the state cost, the result is similar to a chance constraint, where the resulting trajectories from sigma points that violate the soft constraints are heavily penalised. It is worth noticing, that this deterministic method of estimating the moments of the parameter distribution allow the task of sampling actions to be parallelised asynchronously and aggregated when computing the final cost. In future work we intend to explore further possibilities of uncertainty propagation in a principled way.

More importantly, we showed how LFI is a powerful tool to refine the estimation of the posterior distribution. As the inference is based on the same transition function  $f_\theta$  of the controller, it may compensate overly simplified models of the environment. Therefore, we want to explore pathways to efficiently retrain this estimate online so practical experiments with time-variant parameters may be conducted. This is a crucial step towards generalisation of control policies for autonomous robots operating under varying environments and configurations. Crucially, by combining parameter estimation and gradient-free control methods, DISCO may also be used with black-box simulators, such as data-driven function approximators, as long as we are able to sample efficiently from them. This is a promising direction for future research.

## REFERENCES

- [1] L. de Oliveira and E. Camponogara, "Multi-agent model predictive control of signaling split in urban traffic networks," *Transportation Research Part C: Emerging Technologies*, vol. 18, no. 1, 2010.
- [2] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, Dec. 2014.
- [3] A. Mesbah, "Stochastic Model Predictive Control: An Overview and Perspectives for Future Research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, Dec. 2016.
- [4] B. Recht, "A Tour of Reinforcement Learning: Tutorial," *2018 ICML Tutorial*, p. 80, 2018.
- [5] E. Bradford and L. Imsland, "Stochastic Nonlinear Model Predictive Control with State Estimation by Incorporation of the Unscented Kalman Filter," *arXiv:1709.01201 [math]*, Sep. 2017, arXiv: 1709.01201.
- [6] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. Theodorou, "Information theoretic MPC for model-based reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation*. Singapore, Singapore: IEEE, May 2017, pp. 1714–1721.
- [7] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-Theoretic Model Predictive Control: Theory and Applications to Autonomous Driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, Dec. 2018.
- [8] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, Mar. 2004.
- [9] B. Houska, H. J. Ferreau, and M. Diehl, "ACADO toolkit—An open-source framework for automatic control and dynamic optimization," *Optimal Control Applications and Methods*, vol. 32, no. 3, pp. 298–312, 2011. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/oca.939>
- [10] T. Erez, K. Lowrey, Y. Tassa, V. Kumar, S. Koley, and E. Todorov, "An integrated system for real-time model predictive control of humanoid robots," in *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Oct. 2013, pp. 292–299.
- [11] S. V. Rakovic, B. Kouvaritakis, M. Cannon, C. Panos, and R. Findeisen, "Parameterized Tube Model Predictive Control," *IEEE Transactions on Automatic Control*, vol. 57, no. 11, pp. 2746–2761, Nov. 2012.
- [12] S. Thangavel, S. Lucia, R. Paulen, and S. Engell, "Robust nonlinear model predictive control with reduction of uncertainty via dual control," in *2017 21st International Conference on Process Control (PC)*, Jun. 2017, pp. 48–53.
- [13] E. Arruda, M. J. Mathew, M. Kopicki, M. Mistry, M. Azad, and J. L. Wyatt, "Uncertainty averse pushing with model predictive path integral control," in *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*. Birmingham: IEEE, Nov. 2017, pp. 497–502. [Online]. Available: <http://ieeexplore.ieee.org/document/8246918/>
- [14] F. Ramos, R. Possas, and D. Fox, "BayesSim: Adaptive Domain Randomization Via Probabilistic Inference for Robotics Simulators," in *Proceedings of Robotics: Science and Systems*, Jun. 2019, event-place: Freiburg/Breisgau, Germany.
- [15] C. Andrieu, N. De Freitas, A. Doucet, and M. I. Jordan, "An introduction to MCMC for machine learning," *Machine Learning*, 2003, arXiv: 1109.4435v1 ISBN: 0885-6125.
- [16] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, "Learning Without Mixing: Towards A Sharp Analysis of Linear System Identification," *arXiv:1802.08334 [cs, math, stat]*, Feb. 2018, arXiv: 1802.08334. [Online]. Available: <http://arxiv.org/abs/1802.08334>
- [17] S. Schaal, "Learning from Demonstration," in *Advances in Neural Information Processing Systems 9*, M. C. Mozer, M. I. Jordan, and T. Petsche, Eds. MIT Press, 1997, pp. 1040–1046. [Online]. Available: <http://papers.nips.cc/paper/1224-learning-from-demonstration.pdf>
- [18] P. Abbeel, A. Coates, and A. Y. Ng, "Autonomous Helicopter Aerobatics through Apprenticeship Learning," *The International Journal of Robotics Research*, vol. 29, no. 13, pp. 1608–1639, Nov. 2010. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/0278364910371999>
- [19] T. Ganegedara, L. Ott, and F. Ramos, "Online Adaptation of Deep Architectures with Reinforcement Learning," in *Proceedings of the Twenty-second European Conference on Artificial Intelligence*, ser. ECAI'16. IOS Press, 2016, pp. 577–585, tex.acmid: 3305659 tex.numpages: 9 event-place: The Hague, The Netherlands. [Online]. Available: <https://doi.org/10.3233/978-1-61499-672-9-577>
- [20] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox, "Closing the Sim-to-Real Loop: Adapting Simulation Randomization with Real World Experience," *arXiv:1810.05687 [cs]*, Oct. 2018, arXiv: 1810.05687. [Online]. Available: <http://arxiv.org/abs/1810.05687>
- [21] R. van der Merwe, "Sigma-point Kalman filters for probabilistic inference in dynamic state-space models," PhD, Oregon Health & Science University, Beaverton, Oregon, Jan. 2004.
- [22] S. Julier, "The scaled unscented transformation," in *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*. Anchorage, AK, USA: IEEE, 2002, pp. 4555–4559 vol.6. [Online]. Available: <http://ieeexplore.ieee.org/document/1025369/>
- [23] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "OpenAI Gym," *arXiv:1606.01540 [cs]*, Jun. 2016, arXiv: 1606.01540. [Online]. Available: <http://arxiv.org/abs/1606.01540>
- [24] L. Barcelos, R. Oliveira, R. Possas, L. Ott, and F. Ramos, "DISCO: Double Likelihood-free Inference Stochastic Control," *arXiv:2002.07379 [cs]*, Feb. 2020, arXiv: 2002.07379. [Online]. Available: <http://arxiv.org/abs/2002.07379>
- [25] K. Kozłowski and D. Pazderski, "Modeling and Control of a 4-wheel Skid-steering Mobile Robot," *Int. J. Appl. Math. Comput. Sci.*, vol. 14, no. 4, pp. 477–496, 2004.
- [26] J. Yi, H. Wang, J. Zhang, D. Song, S. Jayasuriya, and J. Liu, "Kinematic Modeling and Analysis of Skid-Steered Mobile Robots With Applications to Low-Cost Inertial-Measurement-Unit-Based Motion Estimation," *IEEE Transactions on Robotics*, vol. 25, no. 5, pp. 1087–1097, Oct. 2009, conference Name: IEEE Transactions on Robotics.