

Learning-Based Optimal Control with Performance Guarantees for Unknown Systems with Latent States

Robert Lefringhausen¹, Supitsana Srithasan¹, Armin Lederer², and Sandra Hirche¹

Abstract—As control engineering methods are applied to increasingly complex systems, data-driven approaches for system identification appear as a promising alternative to physics-based modeling. While the Bayesian approaches prevalent for safety-critical applications usually rely on the availability of state measurements, the states of a complex system are often not directly measurable. It may then be necessary to jointly estimate the dynamics and the latent state, making the quantification of uncertainties and the design of controllers with formal performance guarantees considerably more challenging. This paper proposes a novel method for the computation of an optimal input trajectory for unknown nonlinear systems with latent states based on a combination of particle Markov chain Monte Carlo methods and scenario theory. Probabilistic performance guarantees are derived for the resulting input trajectory, and an approach to validate the performance of arbitrary control laws is presented. The effectiveness of the proposed method is demonstrated in a numerical simulation.

I. INTRODUCTION

An accurate mathematical model is fundamental for the model-based control of complex dynamical systems. Since it is often very time-consuming or even impossible to derive physics-based models, e.g., for flexible robotic manipulators or human-robot interaction, data-driven modeling approaches are gaining attention. Accordingly, there are numerous applications in control, with a combination of state-space models (SSMs) and Gaussian processes (GPs) [1] or similar Bayesian approaches being prevalent for safety-critical applications. Their ability to quantify model uncertainties due to limited training data (epistemic uncertainty) and measurement noise (aleatory uncertainty) enables the derivation of probabilistic guarantees for the closed-loop system, even in the presence of noise. Examples can be found, among others, in [2] and [3], where GP-based tracking controllers are proposed and probabilistic bounds on the tracking error are derived.

A serious disadvantage of the aforementioned methods is that full-state measurements are required. In many applications, however, it is unclear which variables represent the

states of the system, or not all of these are measurable. While some learning-based approaches avoid a state-space representation by using past inputs and outputs to make a prediction, e.g., NARX approaches [4], this comes with several drawbacks. First of all, the resulting models cannot separate between noise that permanently changes the trajectory of the system (process noise) and noise that only influences the measurements (measurement noise). They hence provide systematically incorrect open-loop predictions for systems where process and measurement noise differ [5]. Another disadvantage is that it is unclear how prior knowledge can be leveraged since it can be expressed almost exclusively in state-space representation. In addition, the states cannot be used for the control task, for example, to formulate constraints. In many cases, it is thus necessary to overcome the dependence of SSMs on state measurements by jointly estimating the dynamics and the latent state. However, to the best of our knowledge, there are no control approaches for nonlinear systems that rigorously account for uncertainties arising from the joint dynamics and state estimation, thus allowing the derivation of performance guarantees in this setting.

The main contribution of this paper is a novel method for the optimal control of unknown systems based on input-output measurements with probabilistic performance guarantees. In order to quantify uncertainties, which is crucial for deriving formal guarantees, we employ a Bayesian approach and formulate a prior over the unknown dynamics and the system trajectory in state-space representation. Since for practical applicability, the prior must be updated based on input-output measurements, but the corresponding posterior distribution is analytically intractable, we utilize particle Markov chain Monte Carlo (PMCMC) methods [6] to draw samples from this distribution. While similar approaches are proposed, among others, in [7], [8], and [9] for system identification, this fundamental idea has not yet been exploited for control. By employing a scenario point of view [10], the resulting samples from the posterior distribution over future system trajectories allow us to analyze the closed-loop performance under an arbitrary fixed control law. Based on this idea, we propose a scenario optimal control problem (OCP), whose solution we prove to exhibit performance and constraint satisfaction guarantees. The straightforward applicability and flexibility of the proposed method are demonstrated through simulations.

The remainder of this paper is structured as follows. After defining the problem in Sec. II, Sec. III briefly reviews PMCMC methods. Sec. IV describes how the obtained

*This work was supported by the European Research Council (ERC) Consolidator Grant "Safe data-driven control for human-centric systems (CO-MAN)" under grant agreement number 864686, by TUM AGENDA 2030, funded by the Federal Ministry of Education and Research (BMBF) and the Free State of Bavaria under the Excellence Strategy of the Federal Government and the Länder as well as by the Hightech Agenda Bavaria, and by the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40 180545.

¹Chair of Information-oriented Control, TUM School of Computation, Information and Technology, Technical University of Munich, Germany [robert.lefringhausen, supitsana.srithasan, hirche]@tum.de.

²Learning and Adaptive Systems Group, Department of Computer Science, ETH Zurich, Switzerland armin.lederer@inf.ethz.ch.

samples can be used to provide formal guarantees for existing control laws and presents a sample-based OCP formulation and guarantees for its solution. The proposed method is numerically evaluated in Sec. V, followed by some concluding remarks in Sec. VI.

II. PROBLEM FORMULATION

Consider the general nonlinear discrete-time system of the form¹

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

$$\mathbf{y}_t = \mathbf{g}(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}_t, \quad (1b)$$

with state $\mathbf{x}_t \in \mathbb{R}^{n_x \in \mathbb{N}}$, which is not observed explicitly, input $\mathbf{u}_t \in \mathbb{R}^{n_u \in \mathbb{N}}$, output $\mathbf{y}_t \in \mathbb{R}^{n_y \in \mathbb{N}}$, and $t \in \mathbb{Z}$. The system is perturbed by independent and identically distributed (iid) process noise $\mathbf{v}_t \sim \mathcal{V}$ and measurement noise $\mathbf{w}_t \sim \mathcal{W}$, which are independent of the state and the input. The transition and observation functions $\mathbf{f}(\cdot)$ and $\mathbf{g}(\cdot)$ as well as the noise distributions \mathcal{V} and \mathcal{W} are assumed to be unknown. We assume that at time $t = 0$ we have access to the dataset $\mathbb{D} = \{\mathbf{u}_t, \mathbf{y}_t\}_{t=T \in \mathbb{Z}_{<-1:-1}}$ containing the last $|T|$ input-output measurements. In order to enable the theoretically justified inference of a model and to provide formal guarantees for its generalization, we further employ the following assumption.

Assumption 1: The structure of the system (1) is known, i.e., suitable parameterizations $\{\mathbf{f}_\theta(\cdot), \mathbf{g}_\theta(\cdot), \mathcal{V}_\theta, \mathcal{W}_\theta\}$ with a finite number of unknown parameters θ are available. In addition, priors $p(\theta)$ and $p(\mathbf{x}_T)$ for the model parameters and the initial state of the observed trajectory are available². The unknown system and the initial state are samples from these priors.

Knowledge of a reasonable prior is a standard assumption in Bayesian inference. In many cases, models for $\mathbf{f}(\cdot)$ and $\mathbf{g}(\cdot)$ can be constructed based on physical principles, e.g., for cooperative manipulation [11]. Moreover, if such knowledge is unavailable, methods from Gaussian process regression can be used to construct parametric models, which can asymptotically represent all continuous functions with an increasing number of basis functions, such as the approach presented in [12]. Often, the type of noise is also known, and only the parameters of the corresponding distribution must be determined. Therefore, this assumption is not particularly restrictive in practice.

Given a stage cost $c(\mathbf{u}_t, \mathbf{x}_t, \mathbf{y}_t)$, the objective is to find an input trajectory $\mathbf{u}_{0:H}$ that minimizes the accumulated cost

¹**Notation:** Lower/upper case bold symbols denote vectors/matrices, respectively. \mathbb{R} denotes the set of real numbers, \mathbb{Z} the set of integers, \mathbb{N}^0 and \mathbb{N} the set of natural numbers with and without zero, and subscripts $<a$ or $\leq a$ the corresponding subsets whose elements satisfy $<a$ or $\leq a$. $\mathbb{P}(\cdot)$ denotes the probability of an event and $p(\cdot)$ the probability density function (pdf). (Multivariate) Gaussian distributed random variables with mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\Sigma}$ are denoted by $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. $\mathbf{a}_{b \in \mathbb{Z}: c \in \mathbb{Z}}$ is the shorthand notation for $\{\mathbf{a}_b, \dots, \mathbf{a}_c\}$. \mathbf{I}_n denotes the $n \times n$ identity matrix and $\mathbf{0}$ the zero matrix of appropriate dimension. $|\cdot|$ denotes the absolute value of a scalar and $\text{card}(\cdot)$ the cardinality of a set.

²Note that the prior for the model parameters θ and the initial state \mathbf{x}_T also implies a prior for all subsequent states.

over the horizon $H \in \mathbb{N}$

$$J_H = \sum_{t=0}^H c(\mathbf{u}_t, \mathbf{x}_t, \mathbf{y}_t) \quad (2)$$

while satisfying constraints of the form

$$h(\mathbf{u}_{0:H}, \mathbf{x}_{0:H}, \mathbf{y}_{0:H}) \leq 0, \quad (3)$$

where $h(\cdot)$ is an arbitrary deterministic function. Since the unknown parameters and the process and measurement noise may have infinite support, it is generally impossible to guarantee the satisfaction of the constraints (3) in every case, i.e., with probability 1. Moreover, input trajectories with such guarantees tend to be overly conservative and are thus undesirable due to poor performance. Instead, we aim to find an input trajectory that is probabilistically robust, i.e., for which we can derive guarantees of the form $\mathbb{P}(h(\mathbf{u}_{0:H}, \mathbf{x}_{0:H}, \mathbf{y}_{0:H}) \leq 0) \geq \alpha$ when applied to the system. In order to address this problem, we first consider how probabilistic guarantees can be provided for the case that a given control law $\mathbf{u}_t = \boldsymbol{\pi}(\mathbf{u}_{T:t-1}, \mathbf{y}_{T:t-1}, t)$ is applied to the unknown system. Based on this analysis, we consider how the search for an input trajectory $\mathbf{u}_{0:H}$ can be formulated as a tractable OCP for whose solution probabilistic performance and constraint satisfaction guarantees can be inferred directly.

III. PARTICLE MARKOV CHAIN MONTE CARLO METHODS

For practical applicability, the prior presented in the previous section must be updated based on the observations \mathbb{D} , i.e., the posterior distribution over model parameters and latent state trajectories³ $p(\theta, \mathbf{x}_{T:-1} \mid \mathbb{D})$ must be inferred. Without inferring the posterior distribution, the stochastic OCP is usually infeasible since the prior over system trajectories, which implicitly follows from the repeated propagation of $p(\mathbf{x}_T)$ through the prior distribution over the dynamics, exhibits an excessive variance. In general, inference of the posterior distribution $p(\theta, \mathbf{x}_{T:-1} \mid \mathbb{D})$ is analytically intractable, but particle Markov chain Monte Carlo (PMCMC) methods, introduced in [6] and briefly reviewed in this section, can be used to draw samples from it.

PMCMC methods utilize sequential Monte Carlo (SMC) algorithms [13], which numerically approximate the distribution over the latent state trajectory $p(\mathbf{x}_{T:-1} \mid \theta, \mathbb{D})$ by a set of N weighted particles, to generate proposals for Markov chain Monte Carlo (MCMC) samplers, a broad class of algorithms designed for sampling from posterior distributions.

An example is particle Gibbs (PG) sampling, which enables particularly efficient sampling from $p(\theta, \mathbf{x}_{T:-1} \mid \mathbb{D})$ in case samples can be drawn from the conditional distribution $p(\theta \mid \mathbf{x}_{T:-1}, \mathbb{D})$. In PG sampling, a state trajectory $\mathbf{x}_{T:-1}^{[k]}$ is first drawn from the distribution $p(\mathbf{x}_{T:-1}^{[k]} \mid \theta^{[k]}, \mathbb{D})$ using an SMC algorithm. Based on this state trajectory, new model parameters $\theta^{[k+1]}$ are then drawn from $p(\theta^{[k+1]} \mid \mathbf{x}_{T:-1}^{[k]}, \mathbb{D})$. These two steps are repeated until the desired number of samples is reached.

³Note that throughout the work, the dependency on the prior is omitted for notational simplicity.

Remark 1: Note that the distribution $p(\theta \mid \mathbf{x}_{T:-1}, \mathbb{D})$, required for PG sampling, depends on the prior $p(\theta)$. Various combinations of parametric models and priors can be found in the literature that provide a closed-form expression for $p(\theta \mid \mathbf{x}_{T:-1}, \mathbb{D})$ and thus enable efficient PG sampling.

Remark 2: If the conditional distribution $p(\theta \mid \mathbf{x}_{T:-1}, \mathbb{D})$ is not available, other PMCMC methods, e.g., particle marginal Metropolis-Hastings sampling [6], can be applied. The theory presented hereafter is not tied to a specific PMCMC method.

It can be shown that the invariant distribution of PG sampling and other suitable PMCMC methods is $p(\theta, \mathbf{x}_{T:-1} \mid \mathbb{D})$ and that they will asymptotically provide samples from this distribution even for a finite number of particles [6]. Since the first samples depend heavily on the initialization and may not accurately represent the desired distribution, the first K_b samples must be discarded (so-called burn-in period). As for all MCMC methods, no guarantees can be given regarding the length of the burn-in period. The samples are also correlated with samples generated immediately before and after. Since independent samples are desired for the intended application, additional measures must be taken to reduce the correlation. Several approaches are presented in the literature; see, e.g., [14]. A straightforward and comparatively sample-efficient approach is thinning, which reduces the correlation by using only every k_d -th sample. The parameter k_d should be chosen as small as possible so that few samples are discarded but large enough so that the samples are approximately independent. An indicator of an appropriate k_d is a low auto-correlation of successive samples.

Remark 3: How well a PMCMC algorithm can explore the posterior distribution and how small k_d can be chosen also depends significantly on the prior and the parameters of the algorithm. In order to reduce k_d to an acceptable level, scaling the model parameters or adjusting the prior may be necessary.

With a burn-in period of appropriate length and a suitable measure to reduce the correlation, PMCMC methods generate almost uncorrelated samples from the desired distribution, justifying the following assumption.

Assumption 2: The employed PMCMC method provides K independent samples $\{\theta, \mathbf{x}_{T:-1}\}^{[1:K]}$ from the distribution $p(\theta, \mathbf{x}_{T:-1} \mid \mathbb{D})$.

IV. CONTROL WITH GUARANTEES

In this section, we consider how input trajectories with probabilistic performance guarantees can be derived based on the prior over the system dynamics and the latent state trajectory presented in Sec. II. This prior is very general and can accurately represent a wide range of systems due to the separation of process and measurement noise. In addition, the state-space representation makes it very intuitive to incorporate prior knowledge. In order to determine an input trajectory, we propose to utilize PMCMC methods to sample from the posterior distribution over future trajectories of the unknown system, which depends on the control inputs $\mathbf{u}_{0:H}$. Due to the representation as a parametric state-space

Algorithm 1 Scenario generation

Input: Dataset \mathbb{D} , parametric model $\{\mathbf{f}_\theta(\cdot), \mathbf{g}_\theta(\cdot), \mathcal{V}_\theta, \mathcal{W}_\theta\}$, priors $p(\theta)$ and $p(\mathbf{x}_T)$, K , H

Output: Scenarios $\delta^{[1:K]} = \{\theta, \mathbf{x}_0, \mathbf{v}_{0:H}, \mathbf{w}_{0:H}\}^{[1:K]}$

- 1: **for** $k = 1$ to K **do**
 - 2: Sample $\{\theta, \mathbf{x}_{T:-1}\}^{[k]}$ from $p(\theta, \mathbf{x}_{T:-1} \mid \mathbb{D})$ using a PMCMC method.
 - 3: **for** $t = -1$ to H **do**
 - 4: Sample $\mathbf{v}_t^{[k]}$ from $\mathcal{V}_{\theta^{[k]}}$.
 - 5: Sample $\mathbf{w}_t^{[k]}$ from $\mathcal{W}_{\theta^{[k]}}$.
 - 6: **end for**
 - 7: Set $\mathbf{x}_0^{[k]} = \mathbf{f}_{\theta^{[k]}}(\mathbf{x}_{-1}^{[k]}, \mathbf{u}_{-1}) + \mathbf{v}_{-1}^{[k]}$.
 - 8: **end for**
-

model, future state and output trajectories can be expressed as a function of $\mathbf{u}_{0:H}$, which is uniquely characterized by the model parameters θ , the initial state \mathbf{x}_0 , and the realizations of the process and measurement noise $\mathbf{v}_{0:H}$ and $\mathbf{w}_{0:H}$, respectively. Samples δ from the corresponding posterior distribution $p(\theta, \mathbf{x}_0, \mathbf{v}_{0:H}, \mathbf{w}_{0:H} \mid \mathbb{D})$ represent possible future system behavior depending on $\mathbf{u}_{0:H}$ and we refer to these samples as scenarios in the following. Algorithm 1 generates the scenarios based on the observations \mathbb{D} . First, K model parameters θ and state trajectories $\mathbf{x}_{T:-1}$ are drawn from the corresponding posterior distribution using a PMCMC method (line 2). Then, possible realizations of the process and measurement noise are drawn from the distributions parameterized by θ (lines 3-6), and the initial state is determined via forward propagation (line 7).

The scenarios obtained by Algorithm 1 depict not only the aleatory uncertainty but also the epistemic uncertainty. As we will explain in the following, these scenarios are a powerful tool for analyzing and synthesizing control laws. In Sec. IV-A, we first outline how they can be used to provide probabilistic guarantees for existing control laws. Then, in Sec. IV-B, we present a scenario-based OCP formulation and derive probabilistic guarantees for its solution.

A. Analysis of existing control laws

In the following, we show how the scenarios $\delta^{[1:K]}$ can be used to provide probabilistic performance and constraint satisfaction guarantees for arbitrary control laws $\mathbf{u}_t = \pi(\mathbf{u}_{T:t-1}, \mathbf{y}_{T:t-1}, t)$. We thereby use a very natural idea to analyze the closed-loop behavior: since the scenarios represent samples from the distribution over possible future system trajectories depending on the inputs, we can obtain samples $\{\mathbf{u}_{0:H}, \mathbf{x}_{0:H}, \mathbf{y}_{0:H}, J_H\}^{[1:K]}$ from the posterior distribution over future inputs, states, outputs, and costs by simulating the scenarios together with the control law forward as outlined in Algorithm 2. In order to enable a formal performance analysis based on these samples, the following prerequisite must be fulfilled.

Assumption 3: The control law $\mathbf{u}_t = \pi(\cdot)$ is independent of the scenarios $\delta^{[1:K]}$, which are used to analyze it, i.e., it is not designed based on these scenarios.

This assumption is not limiting since additional scenarios can be generated for the analysis if scenario-dependent

Algorithm 2 Forward simulation

Input: Scenarios $\delta^{[1:K]}$, control law $\pi(\mathbf{u}_{T:t-1}, \mathbf{y}_{T:t-1}, t)$, K, H

Output: Samples $\{\mathbf{u}_{0:H}, \mathbf{x}_{0:H}, \mathbf{y}_{0:H}, J_H\}^{[1:K]}$

```
1: for  $k = 1$  to  $K$  do
2:   for  $t = 0$  to  $H$  do
3:     Set  $\mathbf{u}_t^{[k]} = \pi(\mathbf{u}_{T:t-1}^{[k]}, \mathbf{y}_{T:t-1}^{[k]}, t)$ .
4:     Set  $\mathbf{x}_{t+1}^{[k]} = \mathbf{f}_{\theta^{[k]}}(\mathbf{x}_t^{[k]}, \mathbf{u}_t^{[k]}) + \mathbf{v}_t^{[k]}$ .
5:     Set  $\mathbf{y}_t^{[k]} = \mathbf{g}_{\theta^{[k]}}(\mathbf{x}_t^{[k]}, \mathbf{u}_t^{[k]}) + \mathbf{w}_t^{[k]}$ .
6:   end for
7:   Set  $J_H^{[k]} = J_H(\mathbf{u}_{0:H}^{[k]}, \mathbf{x}_{0:H}^{[k]}, \mathbf{y}_{0:H}^{[k]})$ .
8: end for
```

control laws are to be investigated.

In order to derive probabilistic guarantees for the cost incurred when the control law is applied to the unknown system, we consider the maximum of the K samples of the accumulated cost (2) generated by Algorithm 2

$$\overline{J}_H = \max_{k \in \mathbb{N}_{\leq K}} J_H^{[k]}. \quad (4)$$

The following theorem provides a statement on the probability that this worst-case cost \overline{J}_H , which is a random variable itself, is an upper bound for the incurred cost.

Theorem 1: For a given confidence parameter $\beta \in (0, 1)$, under Assumptions 1, 2, and 3, it holds that

$$\mathbb{P}\left(V_J \leq 1 - \sqrt[\kappa]{\beta}\right) \geq 1 - \beta, \quad (5)$$

where V_J denotes the probability that the cost (2), incurred when the control law $\pi(\cdot)$ is applied to the unknown system, exceeds \overline{J}_H , i.e., $V_J = \mathbb{P}(J_H > \overline{J}_H)$.

Proof: Due to Assumptions 1 and 2, the samples $\{\theta, \mathbf{x}_{T:-1}\}^{[1:K]}$ generated by the PMCMC method are independent samples from the posterior distribution over system dynamics and past state trajectories. Due to Assumption 3, the samples $\{\mathbf{u}_{0:H}, \mathbf{x}_{0:H}, \mathbf{y}_{0:H}, J_H\}^{[1:K]}$ generated with Algorithms 1 and 2 are therefore also independent samples from the respective posterior distributions. \overline{J}_H corresponds to the solution of a sample-dependent optimization problem with the single decision variable \overline{J}_H , the objective $\min \overline{J}_H$, which is convex in the decision variable, and the constraint $\max_{k \in \mathbb{N}_{\leq K}} J_H^{[k]} - \overline{J}_H \leq 0$, which is also convex in \overline{J}_H . The result follows from [15, Theorem 1] by simple algebraic manipulations. ■

If the constraints (3) are satisfied for all scenarios, the same line of thought can be used to derive guarantees that the constraints are satisfied when the control law $\pi(\cdot)$ is applied to the unknown system.

Theorem 2: Assume that with a control law $\pi(\cdot)$ the constraints (3) are satisfied for all scenarios, i.e.,

$$h\left(\mathbf{u}_{0:H}^{[k]}, \mathbf{x}_{0:H}^{[k]}, \mathbf{y}_{0:H}^{[k]}\right) \leq 0 \quad \forall k \in \mathbb{N}_{\leq K} \quad (6)$$

where $\{\mathbf{u}_{0:H}, \mathbf{x}_{0:H}, \mathbf{y}_{0:H}\}^{[1:K]}$ are obtained via Algorithm 2. For a given confidence parameter $\beta \in (0, 1)$, under Assumptions 1, 2, and 3, it holds that

$$\mathbb{P}\left(V_h \leq 1 - \sqrt[\kappa]{\beta}\right) \geq 1 - \beta, \quad (7)$$

where V_h denotes the probability that the constraints (3) are violated when the control law $\pi(\cdot)$ is applied to the unknown system.

Proof: This follows from the proof of Theorem 1, if one considers $\max_{k \in \mathbb{N}_{\leq K}} h(\mathbf{u}_{0:H}^{[k]}, \mathbf{x}_{0:H}^{[k]}, \mathbf{y}_{0:H}^{[k]}) \leq 0$ instead of the maximum cost \overline{J}_H . ■

Theorems 1 and 2 thus provide probabilistic performance and constraint satisfaction guarantees when arbitrary control laws are applied to the unknown system. The strength of the guarantees depends on the total number of samples K , and more samples lead to stronger guarantees.

B. Scenario-based optimal control

In the previous section, we have shown that a finite number of scenarios $\delta^{[1:K]}$ can be used to provide probabilistic guarantees for existing control laws. A natural question is whether these scenarios can also be used directly to find a control law or input trajectory that robustly minimizes the cost (2) while satisfying the constraints (3) with high probability. As we will show in the following, it is possible to formulate the search for such an input trajectory⁴ $\mathbf{u}_{0:H}$ as a deterministic OCP for whose solution probabilistic constraint satisfaction guarantees can be directly provided. For this purpose, we make use of the ideas presented in Sec. IV-A and formulate the following OCP:

$$\min_{\mathbf{u}_{0:H}, \overline{J}_H} \overline{J}_H \quad (8a)$$

$$\text{subject to: } \forall k \in \mathbb{N}_{\leq K}, \quad \forall t \in \mathbb{N}_{\leq H}^0$$

$$\mathbf{x}_{t+1}^{[k]} = \mathbf{f}_{\theta^{[k]}}\left(\mathbf{x}_t^{[k]}, \mathbf{u}_t\right) + \mathbf{v}_t^{[k]}, \quad (8b)$$

$$\mathbf{y}_t^{[k]} = \mathbf{g}_{\theta^{[k]}}\left(\mathbf{x}_t^{[k]}, \mathbf{u}_t\right) + \mathbf{w}_t^{[k]}, \quad (8c)$$

$$J_H^{[k]} = J_H\left(\mathbf{u}_{0:H}, \mathbf{x}_{0:H}^{[k]}, \mathbf{y}_{0:H}^{[k]}\right) \leq \overline{J}_H, \quad (8d)$$

$$h\left(\mathbf{u}_{0:H}, \mathbf{x}_{0:H}^{[k]}, \mathbf{y}_{0:H}^{[k]}\right) \leq 0. \quad (8e)$$

In order to ensure the robust minimization of the cost (2), the objective of the OCP is to minimize the worst-case cost \overline{J}_H , i.e., similar to (4), the costs of all scenarios must be less than or equal to \overline{J}_H resulting in (8d). The forward simulation utilized in Sec. IV-A (Algorithm 2) is implicitly included in the optimization problem via the constraints (8b-8c). Similar to the previous section, we also require that the constraints (3) are satisfied for all scenarios, resulting in (8e).

The optimization problem (8) is deterministic, more precisely, a nonlinear programming (NLP) problem in which the uncertainty about the unknown system is incorporated by considering multiple deterministic scenarios rather than in terms of distributions. The problem can thus be solved with well-known methods; see [16] for an overview. Since the problem (8) is generally non-convex, convergence to the global minimum cannot be guaranteed. However, locally optimal solutions are also well-suited for many practical applications.

In the following, we assume that a feasible and locally optimal solution $\{\mathbf{u}_{0:H}^*, \overline{J}_H^*\}$ to the problem (8) is available.

⁴Note that the considerations given here can be readily extended to finding parameters of arbitrary parametric control laws $\pi(\cdot)$.

Note that epistemic and aleatory uncertainties play a crucial role here. In (8e), it is required that the constraints (3) are satisfied for all scenarios. If the sampled scenarios are very different from each other and the uncertainty is thus high, finding a feasible input trajectory may be challenging. However, it is also generally expected that if the uncertainty about the system's behavior is high, it may be difficult or impossible to find an input that satisfies the constraints with high probability - the proposed method is no exception.

Remark 4: In general, the uncertainty about the latent states $\mathbf{x}_{0:H}$ is much higher than about the observable output $\mathbf{y}_{0:H}$ or $\mathbf{g}(\mathbf{x}_{0:H}, \mathbf{u}_{0:H})$. Therefore, in case of high uncertainty (e.g., many unknown parameters and a flat prior), the constraints (3) should depend only on $\mathbf{u}_{0:H}$, $\mathbf{y}_{0:H}$, and $\mathbf{g}(\mathbf{x}_{0:H}, \mathbf{u}_{0:H})$. In practice, the constraints will usually depend only on these variables since, without measurability, it is difficult to formulate meaningful constraints for the states.

Similar to Sec. IV-A, we are interested in deriving probabilistic guarantees for the incurred cost and the satisfaction of constraints when the input trajectory $\mathbf{u}_{0:H}^*$, obtained by solving the OCP (8), is applied to the unknown system. In (8), only a finite number of possible future system trajectories are considered, rather than the corresponding analytically intractable distribution. However, in a safety-critical context, ensuring that the constraints (3) are satisfied for all possible trajectories with high probability is necessary. Theorems 1 and 2 do not apply in this case since the control input depends on the scenarios, and therefore Assumption 3 is not fulfilled. In order to obtain a formal statement about the worst-case cost and the probability of a constraint violation, we instead consider variants of the optimization problem (8) where the constraints (8b-8e) are in effect for only a subset of the scenarios. In order to derive formal guarantees by solving these problems, we need to ensure that the solver used to solve the OCP (8) provides repeatable answers. This is formalized in the following assumption [17].

Assumption 4: The OCP (8) and the employed solver represent a mapping M_K from the scenarios $\delta^{[1:K]}$ to a solution $\{\mathbf{u}_{0:H}^*, \overline{J}_H^*\}$. Let $\mathbb{H}_{\delta^{[i]}}$ denote the set that contains all solutions that are suitable for a scenario $\delta^{[i]}$, i.e., all $\{\mathbf{u}_{0:H}, \overline{J}_H\}$ that satisfy the constraints (8b-8e) for $k = i$. Then, for every $K \in \mathbb{N}$, every $k \in \mathbb{N}$, and for every choice of $\delta^{[1]}, \dots, \delta^{[K]}$ and $\delta^{[K+1]}, \dots, \delta^{[K+k]}$, the following three properties hold:

- (i) If $\delta^{[i_1]}, \dots, \delta^{[i_K]}$ is a permutation of $\delta^{[1]}, \dots, \delta^{[K]}$, then $M_K(\delta^{[1]}, \dots, \delta^{[K]}) = M_K(\delta^{[i_1]}, \dots, \delta^{[i_K]})$.
- (ii) If $M_K(\delta^{[1]}, \dots, \delta^{[K]}) \in \mathbb{H}_{\delta^{[K+i]}}$ for all $i = 1, \dots, k$, then $M_K(\delta^{[1]}, \dots, \delta^{[K]}) = M_{K+k}(\delta^{[1]}, \dots, \delta^{[K+k]})$.
- (iii) If $M_K(\delta^{[1]}, \dots, \delta^{[K]}) \notin \mathbb{H}_{\delta^{[K+i]}}$ for at least one $i = 1, \dots, k$, then $M_K(\delta^{[1]}, \dots, \delta^{[K]}) \neq M_{K+k}(\delta^{[1]}, \dots, \delta^{[K+k]})$.

This assumption is not restrictive, and it can be straightforwardly shown to be satisfied for most deterministic gradient-based solvers, provided that the same initialization is used and the cost function has a unique local minimum, possibly after inserting a regularization term. However, it is

generally not fulfilled for random-based methods such as genetic algorithms, which may provide different solutions when considering the same scenarios due to their reliance on randomness.

Now suppose a set $\mathbb{S} \subset \mathbb{N}_{\leq K}$ with cardinality $s := \text{card}(\mathbb{S}) < K$ is known for which a solver satisfying Assumption 4 yields the same solution as for the original OCP (8) when only the scenarios belonging to this set are considered. That is, the solution of the OCP (8), where the constraints (8b-8e) are in place only for $k \in \mathbb{S}$, is $\{\mathbf{u}_{0:H}^*, \overline{J}_H^*\}$. In the following, we will refer to such a set \mathbb{S} as a support sub-sample. In that case, the following result provides probabilistic guarantees for the maximum cost and the satisfaction of the constraints (3) when $\mathbf{u}_{0:H}^*$ is applied to the unknown system.

Theorem 3: Assume a support sub-sample \mathbb{S} with cardinality s is known. Let V_c denote the probability that the incurred cost (2) exceeds \overline{J}_H^* or that the constraints (3) are violated when the input trajectory $\mathbf{u}_{0:H}^*$ is applied to the unknown system. For a given confidence parameter $\beta \in (0, 1)$, under Assumptions 1, 2 and 4 it holds that

$$\mathbb{P}(V_c \leq 1 - \epsilon(s)) > 1 - \beta, \quad (9)$$

where $\epsilon(s)$ is the unique solution over the interval $(0, 1)$ of the polynomial equation in the v variable

$$\binom{K}{s} (1-v)^{K-s} - \frac{\beta}{K} \sum_{m=s}^{K-1} \binom{m}{s} (1-v)^{m-s} = 0. \quad (10)$$

Proof: Due to Assumptions 1 and 2, the scenarios $\delta^{[1:K]}$ generated by Algorithm 1 are independent samples from the respective distributions. Due to Assumption 4 and after combining (8b-8e) to a single scalar-valued constraint, [17, Theorem 1] is applicable. The result follows since, for a fixed K , the implied function $\epsilon(k)$ monotonically decreases for an increasing k . ■

Thus, we need to find a support sub-sample \mathbb{S} to obtain formal guarantees for the input trajectory $\mathbf{u}_{0:H}^*$. The strength of the guarantees depends on the number of scenarios K and the cardinality s of the support sub-sample. The guarantees become stronger for a fixed K as s decreases. However, determining the support sub-sample with the smallest possible cardinality and, thus, the strongest possible guarantees is a complex combinatorial problem. Therefore, heuristics, e.g., the greedy algorithm presented in [18], are usually employed to determine a possibly suboptimal \mathbb{S} . For a fixed cardinality s , increasing the number of scenarios K leads to stronger guarantees. Therefore, additional scenarios are usually generated if the robustness level is insufficient for the intended application. However, since the problem (8) is generally non-convex due to the nonlinear dynamics, and it is thus not guaranteed that additional scenarios will not also increase s , it is impossible to determine the number of scenarios required for a given robustness level a priori.

With Theorem 3 and a suitable strategy to find a support sub-sample \mathbb{S} , it is possible to provide formal performance and constraint satisfaction guarantees. Due to the dedicated scenario generation process, these guarantee robustness to

aleatory and epistemic uncertainties. The scenario-based formulation thus prevents the solution of the OCP from becoming overconfident in the face of uncertainty, allowing the use of the proposed method for safety-critical applications.

V. SIMULATION

In this section, we demonstrate the effectiveness of the proposed optimal control approach in a simulation⁵. The simulation setup is outlined in Sec. V-A. The proposed control approach is evaluated in Sec. V-B, assuming a known parametric model for the unknown system. In Sec. V-C, we demonstrate that it performs well even if no parametric model for the state transition function is known, using methods from Gaussian process regression.

A. Setup

Consider a system of the form (1) with the unknown state transition function

$$\mathbf{f}(\mathbf{x}, u) = \begin{bmatrix} 0.8x_1 - 0.5x_2 + 0.1 \cos(3x_1)x_2 \\ 0.4x_1 + 0.5x_2 + (1 + 0.3 \sin(2x_2))u \end{bmatrix} \quad (11)$$

and unknown process noise distribution

$$\mathbf{v}_t \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} 0.03 & -0.004 \\ -0.004 & 0.01 \end{bmatrix}\right). \quad (12)$$

In order to simplify the interpretation of the simulation results, we assume that the observation function $g(\mathbf{x}, u) = x_1$ and the distribution of the measurement noise $w_t \sim \mathcal{N}(0, 0.1)$ are known. Note that this assumption is without loss of generality since an unknown observation model can be absorbed into the state transition model when the number of states is increased to $n_x + n_y$ [19, Sec. 3.2.1].

For the evaluation of the proposed optimal control approach, we consider the cost function $J_H = \sum_{t=0}^H u_t^2$ over the horizon $H = 40$. In order to investigate how the control approach reacts to temporarily active constraints, we consider the constraints $2 \leq y_{20:25}$ and $|u| \leq 5$. For the scenario generation, $|T| = 2000$ input-output measurements are obtained by simulating the system forward with a random input trajectory $u_t \sim \mathcal{N}(0, 3)$ starting from a random initial state with known distribution $\mathbf{x}_T \sim \mathcal{N}([2, 2]^T, \mathbf{I}_2)$.

We employ the approach presented in [9] to infer the model parameters and the latent state trajectory. In this approach, it is assumed that the state transition function $\mathbf{f}(\cdot)$ is a linear combination of n_a basis functions $\varphi(\mathbf{x}_t, \mathbf{u}_t)$ and that the process noise is normally distributed with zero mean. This yields the model $\mathbf{x}_{t+1} = \mathbf{A}\varphi(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{v}_t$, $\mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ with unknown parameters $\theta = \{\mathbf{A} \in \mathbb{R}^{n_x \times n_a}, \mathbf{Q} \in \mathbb{R}^{n_x \times n_x}\}$. For \mathbf{Q} , an inverse Wishart (IW) prior with ℓ degrees of freedom and positive definite scale matrix $\mathbf{\Lambda}$ is assumed. The IW distribution is a distribution over real-valued, symmetric, positive definite matrices. For \mathbf{A} a matrix normal (MN) prior with mean matrix $\mathbf{M} \in \mathbb{R}^{n_x \times n_a} = \mathbf{0}$, right covariance matrix $\mathbf{U} \in \mathbb{R}^{n_x \times n_x} = \mathbf{Q}$, and left covariance matrix $\mathbf{V} \in \mathbb{R}^{n_a \times n_a}$, is assumed. This conjugate prior results in a closed-form expression for $p(\theta | \mathbf{x}_{T:-1}, \mathbb{D})$

⁵The code is available at <https://github.com/TUM-ITR/PGopt>

| $\mathbf{\Lambda}, \mathbf{Q}^{\text{init}}$ | ℓ | \mathbf{A}^{init} | K_b | N |
|--|--------|----------------------------|-------|-----|
| $100\mathbf{I}_2$ | 10 | $\mathbf{0}$ | 1000 | 30 |

TABLE I

PARAMETERS OF THE PRIOR AND THE PG SAMPLER.




Fig. 1. Normalized auto-correlation function (ACF) between successive samples of the PG sampler without thinning. The red lines represent the ACF for the 10 different entries of the weight matrix \mathbf{A} , the blue lines represent the ACF for the 4 different entries of the process noise covariance matrix \mathbf{Q} , and the green lines represent the ACF for the 2 different entries of the state \mathbf{x}_{-1} .

[9] and thus the efficient PG sampler presented in [20] can be used to draw samples from $p(\theta, \mathbf{x}_{T:-1} | \mathbb{D})$. The selected parameters are given in Table I.

B. Optimal control with known basis functions

In order to show the applicability of the guarantees derived in this paper, we first assume that the basis functions $\varphi(\mathbf{x}, u) = [x_1, x_2, u, \cos(3x_1)x_2, \sin(2x_2)u]^T$ are known and that only the model parameters $\theta = \{\mathbf{A}, \mathbf{Q}\}$ must be inferred. In order to facilitate the PG sampler's exploration of the posterior distribution, we scale the basis functions with the factors $[0.1, 0.1, 1, 0.01, 0.1]^T$. For the prior for the corresponding weights, we choose $\mathbf{V} = 10\mathbf{I}_5$.

First, we ensure the validity of Assumption 2. For this purpose, we draw $K = 10000$ samples from the posterior distribution over model parameters and latent state trajectories using a PG sampler without thinning. The normalized auto-correlation functions (ACFs) of consecutive samples of the model parameters θ and the state \mathbf{x}_{-1} are shown in Figure 1. The ACF for all 16 parameters decays significantly for a lag of 50. Therefore, a thinning procedure with $k_d = 50$ is used in the following to generate approximately independent samples.

Afterward, $K = 200$ scenarios are generated using Algorithm 1. These scenarios are used to formulate an OCP as described in Sec. IV-B, which is then solved using the trajectory optimizer ALTRO [21]. In order to find a support sub-sample \mathbb{S} , required for the application of Theorem 3, we first sort the scenarios based on their minimum distance to the constraint boundary. We then employ the greedy algorithm presented in [18], which iterates over the scenarios (starting with the scenarios with the largest distance to the constraint boundary) and checks whether removing the constraints associated with a scenario changes the solution. The corresponding scenario is permanently removed from




Fig. 2. Example of the optimal control with known basis functions. The red area shows the output constraints, the gray area encompasses the 200 scenarios that were used to determine the input trajectory, the green line shows the mean prediction, and the blue line shows one realization of the output of the actual system when the input trajectory $\mathbf{u}_{0:H}^*$ is applied from time $t = 0$.

the constraint set if the solution remains unchanged. Based on the cardinality s of the obtained support sub-sample, we compute the guarantees according to Theorem 3 with confidence $1 - \beta = 99\%$, whereby we use the algorithm presented in [10] to solve the polynomial equation (10). Since, in this case, there is no uncertainty about the incurred cost J_H , as it does not depend on the uncertain dynamics, $1 - \epsilon$ directly corresponds to an upper bound on the probability that the constraints are violated. After computing the guarantees, we apply the obtained trajectory $\mathbf{u}_{0:H}^*$ to the actual system for validation.

In order to obtain meaningful results, the process of computing the optimal input trajectory and guarantees is repeated 100 times. Since our approach can also account for epistemic uncertainties, a new training dataset \mathbb{D} is generated for each run as described in Sec. V-A. Of the 100 runs, in 19 cases, a locally optimal input trajectory $\mathbf{u}_{0:H}^*$ is not found within the specified maximum number of iterations of the solver. Possible causes for the solver to fail are that the posed problem is infeasible due to the randomly generated initial state of the OCP or a high uncertainty in this run, or that the solver converges to an infeasible local minimum. The latter problem could be addressed by repeating the optimization with a different initialization. If the optimization problem is infeasible, a potential solution is to reduce the uncertainty by adding more training data or choosing a more specific prior. It may also help to reduce the number of scenarios and, therefore, the strength of the achievable guarantees. Only the 81 runs with a successful optimization are considered in the following.

Figure 2 shows the results of an exemplary run. Initially, the control input is set to zero as this minimizes the stage cost, and the sampled dynamics converge to the stable equilibrium of the unforced system with output $y = 0$. The system must then deviate from the equilibrium to satisfy the constraints $y_{20:25} \geq 2$. The input trajectory is chosen such that some of the scenarios precisely reach the constraint boundary and is thus not overly conservative. After the constraints are no longer present, the control input is again set to zero, and the sampled dynamics converge to the

| | |
|--|--|
| max. constraint violation probability $1 - \epsilon$ | $11.02\% \pm 2.53\%$ |
| number of runs with constraint violations | 4 (= 4.94%) |
| J_H | 26.61 ± 10.93 |
| time to generate scenarios | $900 \text{ s} \pm 56 \text{ s}$ |
| time to compute $\mathbf{u}_{0:H}^*$ | $366 \text{ s} \pm 257 \text{ s}$ |
| time to compute $1 - \epsilon$ | $15\,629 \text{ s} \pm 2118 \text{ s}$ |

TABLE II

RESULTS (MEAN \pm STD) OF THE 81 SUCCESSFUL RUNS WITH KNOWN BASIS FUNCTIONS.

equilibrium again. The greedy algorithm yields a support sub-sample with cardinality $s = 7$ in this run. According to Theorem 3, with confidence $1 - \beta = 99\%$, the probability that the constraints are violated is thus less than $1 - \epsilon = 10.22\%$. The constraints are satisfied when the computed input trajectory is applied to the system.

The results of all 81 runs with a successful optimization are summarized in Table II. The average upper bound for the probability of constraint violation is $1 - \epsilon = 11.02\%$. In the forward simulation, the constraints are violated in 4.94% of the runs. Thus, the theoretical robustness levels presented in this paper appear slightly weaker than the actual robustness. Although the theoretical guarantees in this simulation may not be sufficient for every application, the ability to compute guarantees emphasizes the solid theoretical foundation of our approach. To the best of our knowledge, no other approaches allow the derivation of such guarantees under comparable assumptions. If stronger guarantees are desired, additional scenarios can be considered.

The specified runtimes are obtained on an 2.25 GHz 64-core AMD EPYC 7742 processor, on which the individual runs are executed CPU-parallelized, i.e., they correspond to a single core. As expected, the computational complexity of the approach is comparatively high, especially for the computation of the theoretical guarantees. This is partially due to the fact that solving chance-constrained optimization problems is, in general, computationally intensive. In addition, it should be noted that a very general class of nonlinear dynamical systems is considered and that the approach provides theoretical guarantees even with little prior knowledge and under mild assumptions.

C. Optimal control with generic basis functions

In the following, we show that the proposed optimal control approach can yield good results even if no parametric model is known, which might be the case in practice. Instead of the actual basis functions, we use the reduced-rank GP approximation proposed in [12] to systematically determine the basis functions $\varphi(\mathbf{x}, \mathbf{u})$ and the parameter \mathbf{V} of the prior. We choose a GP with a squared exponential kernel and select the hyperparameters of the GP and the approximation based on the training data. These parameters are given in Table III. Then, $K = 100$ scenarios are generated using a PG sampler with the same parameters as in the previous example, and the resulting OCP is solved as before. Again, the procedure is repeated 100 times, of which the optimizer finds an optimal solution in 77 cases.

Figure 3 shows the results of an exemplary run with

| l | s_f | m_{x_1}, m_{x_2}, m_u | L_{x_1}, L_{x_2} | L_u |
|-----|-------|-------------------------|--------------------|-------|
| 2 | 100 | 5 | 20 | 10 |

TABLE III

PARAMETERS OF THE REDUCED-RANK GP APPROXIMATION.




Fig. 3. Example of the optimal control with generic basis functions. The red area shows the output constraints, the gray area encompasses the 100 scenarios that were used to determine the input trajectory, the green line shows the mean prediction, and the blue line shows one realization of the output of the actual system when the input trajectory $\mathbf{u}_{0:H}^*$ is applied from time $t = 0$.

generic basis functions. The resulting trajectories are similar to the case with known basis functions. Only the uncertainty intervals are larger, which is expected as less specific prior knowledge is brought in. Due to the higher uncertainty, a more conservative input trajectory is chosen, and the distance of the actual output signal to the constraint boundary is larger.

The results of all 77 successful runs with generic basis functions are summarized in Table IV. In this case, no formal guarantees for the constraint satisfaction can be derived since Assumption 1 is not satisfied as the employed basis functions cannot represent the actual dynamics with arbitrary precision. Nevertheless, the resulting input trajectories are robust, and when applied to the actual system, the constraints are violated only in two runs. The higher cost compared to the case with known basis functions is due to the fact that a more conservative input trajectory is chosen due to the higher uncertainty.

As shown in this subsection, the proposed approach is thus able to provide well-functioning and robust input trajectories based on a prior that can be formulated using the well-known GP framework, even without knowledge of a parametric model of the dynamics.

VI. CONCLUSION

This paper presents a novel approach for the optimal control of unknown dynamical systems with latent states. Based on a prior for the dynamics and the latent state trajectory, samples from the posterior distribution over future trajectories of the unknown system are drawn using PMCMC methods. These samples are then used to formulate an optimal control problem that can be solved using well-known methods. Probabilistic performance and constraint satisfaction guarantees are derived for the obtained input trajectory by employing a scenario perspective.

| | |
|---|-------------------------------------|
| number of runs with constraint violations | 2 (= 2.60%) |
| J_H | 81.81 ± 292.5 |
| time to generate scenarios | $14\,138\text{ s} \pm 200\text{ s}$ |
| time to compute $\mathbf{u}_{0:H}^*$ | $749\text{ s} \pm 1108\text{ s}$ |

TABLE IV

RESULTS (MEAN \pm STD) OF THE 77 SUCCESSFUL RUNS WITH GENERIC BASIS FUNCTIONS.

REFERENCES

- [1] C. K. I. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, USA, 2006.
- [2] J. Umlauf and S. Hirche, "Feedback linearization based on Gaussian processes with event-triggered online learning," *IEEE Transactions on Automatic Control*, vol. 65, no. 10, pp. 4154–4169, 2019.
- [3] T. Beckers, D. Kulić, and S. Hirche, "Stable Gaussian process based tracking control of Euler–Lagrange systems," *Automatica*, vol. 103, pp. 390–397, 2019.
- [4] M. Maiworm, D. Limon, and R. Findeisen, "Online learning-based model predictive control with Gaussian process models and stability guarantees," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 18, pp. 8785–8812, 2021.
- [5] J. Schoukens and L. Ljung, "Nonlinear system identification: A user-oriented road map," *IEEE Control Systems Magazine*, vol. 39, no. 6, pp. 28–99, 2019.
- [6] C. Andrieu, A. Doucet, and R. Holenstein, "Particle Markov chain Monte Carlo methods," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 72, no. 3, pp. 269–342, 2010.
- [7] R. Frigola, F. Lindsten, T. B. Schön, and C. E. Rasmussen, "Bayesian inference and learning in Gaussian process state-space models with particle MCMC," *Advances in neural information processing systems*, vol. 26, 2013.
- [8] F. Tobar, P. M. Djurić, and D. P. Mandić, "Unsupervised state-space modeling using reproducing kernels," *IEEE Transactions on Signal Processing*, vol. 63, no. 19, pp. 5210–5221, 2015.
- [9] A. Svensson and T. B. Schön, "A flexible state–space model for learning nonlinear dynamical systems," *Automatica*, vol. 80, pp. 189–199, 2017.
- [10] S. Garatti and M. C. Campi, "Risk and complexity in scenario optimization," *Mathematical Programming*, vol. 191, no. 1, pp. 243–279, 2022.
- [11] P. B. G. Dohmann, A. Lederer, M. Dißemond, and S. Hirche, "Distributed Bayesian online learning for cooperative manipulation," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 2888–2895.
- [12] A. Solin and S. Särkkä, "Hilbert space methods for reduced-rank Gaussian process regression," *Statistics and Computing*, vol. 30, no. 2, pp. 419–446, 2020.
- [13] A. Doucet, A. M. Johansen *et al.*, "A tutorial on particle filtering and smoothing: Fifteen years later," *Handbook of nonlinear filtering*, vol. 12, no. 656–704, p. 3, 2009.
- [14] L. Tierney, "Markov chains for exploring posterior distributions," *the Annals of Statistics*, pp. 1701–1728, 1994.
- [15] M. C. Campi and S. Garatti, "The exact feasibility of randomized solutions of uncertain convex programs," *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1211–1230, 2008.
- [16] J. Nocedal and S. J. Wright, *Numerical optimization*, 2nd ed. New York, NY, USA: Springer, 2006.
- [17] S. Garatti and M. C. Campi, "The risk of making decisions from data through the lens of the scenario approach," *IFAC-PapersOnLine*, vol. 54, no. 7, pp. 607–612, 2021.
- [18] M. C. Campi, S. Garatti, and F. A. Ramponi, "A general scenario theory for nonconvex optimization and decision making," *IEEE Transactions on Automatic Control*, vol. 63, no. 12, pp. 4067–4078, 2018.
- [19] R. Frigola, "Bayesian time series learning with Gaussian processes," Ph.D. dissertation, University of Cambridge, 2015.
- [20] F. Lindsten, T. B. Schön, and M. Jordan, "Ancestor sampling for particle Gibbs," *Advances in Neural Information Processing Systems*, vol. 25, 2012.
- [21] T. A. Howell, B. E. Jackson, and Z. Manchester, "ALTRO: A fast solver for constrained trajectory optimization," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 7674–7679.