

Localized active learning of Gaussian process state space models

Alexandre Capone

Gerrit Noske

Jonas Umlauf

Thomas Beckers

Armin Lederer

Sandra Hirche

Chair of Information-oriented Control

Department of Electrical and Computer Engineering

Technical University of Munich

D-80333 Munich, Germany

ALEXANDRE.CAPONE@TUM.DE

GA24COQ@MYTUM.DE

JONAS.UMLAUFT@TUM.DE

T.BECKERS@TUM.DE

ARMIN.LEDERER@TUM.DE

HIRCHE@TUM.DE

Editors: A. Bayen, A. Jadbabaie, G. J. Pappas, P. Parrilo, B. Recht, C. Tomlin, M. Zeilinger

Abstract

While most dynamic system exploration techniques aim to achieve a globally accurate model, this is generally unsuited for systems with unbounded state spaces. Furthermore, many applications do not require a globally accurate model, e.g., local stabilization tasks. In this paper, we propose an active learning strategy for Gaussian process state space models that aims to obtain an accurate model on a bounded subset of the state-action space. Our approach aims to maximize the mutual information of the exploration trajectories with respect to a discretization of the region of interest. By employing model predictive control, the proposed technique integrates information collected during exploration and adaptively improves its exploration strategy. To enable computational tractability, we decouple the choice of most informative data points from the model predictive control optimization step. This yields two optimization problems that can be solved in parallel. We apply the proposed method to explore the state space of various dynamical systems and compare our approach to a commonly used entropy-based exploration strategy. In all experiments, our method yields a better model within the region of interest than the entropy-based method.

Keywords: exploration, Bayesian inference, data-driven control, model predictive control

1. Introduction

Autonomous systems often need to operate in complex environments, of which a model is difficult or even impossible to derive from first principles. Learning-based techniques have become a promising paradigm to address these issues (Pillonetto et al., 2014). In particular, Gaussian processes (GPs) have been increasingly employed for system identification and control (Umlauf et al., 2018; Capone and Hirche, 2019; Berkenkamp and Schoellig, 2015; Deisenroth and Rasmussen, 2011). GPs possess very good generalization properties (Rasmussen and Williams, 2006), which can be leveraged to obtain data-efficient learning-based approaches (Deisenroth and Rasmussen, 2011; Kamthe and Deisenroth, 2018). By employing a Bayesian framework, GPs provide an automatic trade-off between model smoothness and data fitness. Moreover, GPs provide an explicit estimate of the model uncertainty that is used to derive probabilistic bounds in control settings (Capone and Hirche, 2019; Beckers et al., 2019; Umlauf and Hirche, 2020).

A performance-determining factor of data-driven techniques is the quality of the available data. In settings where data is insufficient to achieve accurate predictions, new data needs to be gathered via exploration (Umlauf and Hirche, 2020). In classical reinforcement learning settings, exploration is often enforced by randomly selecting a control action with a predetermined probability that tends to zero over time (Dayan and Sejnowski, 1996). However, this is generally inefficient, as regions of low uncertainty are potentially revisited in multiple iterations. These issues have been addressed by techniques that choose the most informative exploration trajectories (Alpcan and Shames, 2015; Ay et al., 2008; Burgard et al., 2005; Schreiter et al., 2015). The goal of these methods is to obtain a globally accurate model. While this is reasonable for systems with a bounded state-action space, it is unsuited for systems with unbounded ones, particularly if a non-parametric model is used. This is because a potentially infinite number of points is required to achieve a globally accurate model. Furthermore, in practice a model often only needs to be accurate locally, e.g., for stabilization tasks.

In this paper, we propose a model predictive control-based exploration approach that steers the system towards the most informative points within a bounded subset of the state-action space. By modeling the system with a Gaussian process, we are able to quantify the information inherent in each data point. Our approach chooses actions by approximating the mutual information of the system trajectory with respect to a discretization of the region of interest. This is achieved by first selecting the single most informative data point within the region of interest, then steering the system towards that point using model predictive control. Through this approximation, the solution approach is rendered computationally tractable.

The remainder of this paper is structured as follows. Section 2 describes the general problem. Section 3 discusses how GPs are employed for modeling and exploration. Section 4 presents the LocAL algorithm, and is followed by a numerical simulation example, in Section 5.

2. Problem Statement

We consider the problem of exploring the state and control space of a discrete-time nonlinear system with Markovian dynamics of the form

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

where $t \in \mathbb{N}_0$, $\mathbf{x}_t \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$ and $\mathbf{u}_t \in \mathcal{U} \subseteq \mathbb{R}^{d_u}$ are the system's state vector and control vector at the t -th time step, respectively. The system is disturbed by multivariate Gaussian process noise $\mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_w)$ with $\boldsymbol{\Sigma}_w = \text{diag}(\sigma_{w,1}^2, \dots, \sigma_{w,d_x}^2)$, $\sigma_{w,i}^2 \in \mathbb{R}_{+,0}$. The concatenation $\tilde{\mathbf{x}}_t := (\mathbf{x}_t^\top \quad \mathbf{u}_t^\top)^\top \in \tilde{\mathcal{X}}$, where $\tilde{\mathcal{X}} := \mathcal{X} \times \mathcal{U}$ is employed for simplicity of exposition. The nonlinear function $\mathbf{f} : \tilde{\mathcal{X}} \rightarrow \mathcal{X}$ represents the known component of the system dynamics, e.g., a model obtained using first principles, while $\mathbf{g} : \tilde{\mathcal{X}} \rightarrow \mathcal{X}$ corresponds to the unknown component of the system dynamics.

We aim to obtain an approximation of the function $\mathbf{g}(\cdot)$, denoted $\hat{\mathbf{g}}(\cdot)$, which provides an accurate estimate of $\mathbf{g}(\cdot)$ on a predefined bounded subset of the augmented state space $\tilde{\mathcal{X}}_B \subset \tilde{\mathcal{X}}$. This is often required in practice, e.g., for local stabilization tasks.

3. Gaussian Processes

In order to faithfully capture the stochastic behavior of (1), we model the system as a *Gaussian process* (GP), where we employ measurements of the augmented state vector $\tilde{\mathbf{x}}_t$ as training inputs, and the differences $\mathbf{x}_{t+1} - \mathbf{f}(\tilde{\mathbf{x}}_t) = \mathbf{g}(\tilde{\mathbf{x}}_t) + \mathbf{w}_t$ as training targets.

A GP is a collection of dependent random variables, for which any finite subset is jointly normally distributed (Rasmussen and Williams, 2006). It is specified by a mean function $m : \mathcal{X} \rightarrow \mathbb{R}$ and a positive definite covariance function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, also known as kernel. In this paper, we set $m \equiv 0$ without loss of generality, as all prior knowledge is already encoded in $f(\cdot)$. The kernel $k(\cdot, \cdot)$ is a similarity measure for evaluations of $\mathbf{g}(\cdot)$, and encodes function properties such as smoothness and periodicity.

In the case where the state is a scalar, i.e., $d_x = 1$, given n training input samples $\tilde{\mathbf{X}} = \{\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n\} \subset \mathcal{X}$ and training outputs $\mathbf{y}_{\tilde{\mathbf{X}}} = (g(\tilde{\mathbf{x}}_1) + w_1 \dots g(\tilde{\mathbf{x}}_n) + w_n)^\top$, the posterior mean and variance of the GP corresponds to a one-step transition model. Starting at a point $\tilde{\mathbf{x}}_t$, the difference between the subsequent state and the known component is normally distributed, i.e.,

$$\mathbf{x}_{t+1} - \mathbf{f}(\tilde{\mathbf{x}}_t) \sim \mathcal{N}(\mu_n(\tilde{\mathbf{x}}_t), \sigma_n^2(\tilde{\mathbf{x}}_t)), \quad (2)$$

with mean and variance given by

$$\mu_n(\tilde{\mathbf{x}}_t) := \mathbf{k}^\top(\tilde{\mathbf{x}}_t) (\mathbf{K} + \sigma_w^2 \mathbf{I})^{-1} \mathbf{y}_{\tilde{\mathbf{X}}}, \quad \sigma_n^2(\tilde{\mathbf{x}}_t) := k(\tilde{\mathbf{x}}_t, \tilde{\mathbf{x}}_t) - \mathbf{k}^\top(\tilde{\mathbf{x}}_t) (\mathbf{K} + \sigma_w^2 \mathbf{I})^{-1} \mathbf{k}(\tilde{\mathbf{x}}_t),$$

respectively, where $\mathbf{k}(\cdot) = (k(\tilde{\mathbf{x}}_1, \cdot) \dots k(\tilde{\mathbf{x}}_n, \cdot))^\top$, and the entries of the covariance matrix \mathbf{K} are computed as $K_{ij} = k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$, $i, j = 1, \dots, n$.

In the case where the state is multidimensional, we model dimension of the state transition function using a separate GP. This corresponds to the assumption that the state transition function entries are conditionally independent. For simplicity of exposition, unless stated otherwise, we henceforth assume $d_x = 1$. However, the methods presented in this paper extend straightforwardly to the multivariate case.

3.1. Performing multi-step ahead predictions

The GP model presented in the previous section serves as a one-step predictor given a known test input $\tilde{\mathbf{x}}_t$. However, if only a distribution $p(\tilde{\mathbf{x}}_t)$ is available, the successor states' distribution generally cannot be computed analytically. Hence, the distributions of future states cannot be computed exactly, but only approximated, e.g., using Monte Carlo methods (Candela et al., 2003). Alternatively, approximate computations exist that enable to propagate the GP uncertainty over multiple time steps, such as moment-matching and GP linearization (Deisenroth et al., 2015). In this paper, we employ the GP mean to perform multi-step ahead predictions, without propagating uncertainty, i.e., $\mathbf{x}_{t+1} = \mathbf{f}(\tilde{\mathbf{x}}_t) + \mu_n(\tilde{\mathbf{x}}_t)$, $t \in \mathbb{N}$. However, the proposed method is also applicable using models that propagate uncertainty, e.g., moment-matching or linearization-based methods (Deisenroth et al., 2015).

3.2. Quantifying utility of data

In order to steer the system along informative trajectories, we need to quantify the utility of data points in the augmented state space $\tilde{\mathcal{X}}$. To this end, we consider the mutual information between observations $\mathbf{y}_{\tilde{\mathbf{X}}}$ at training inputs $\tilde{\mathbf{X}}$ and evaluations $\mathbf{y}_{\tilde{\mathbf{X}}_{\text{ref}}}$ at reference points $\tilde{\mathbf{X}}_{\text{ref}}$. Here $\tilde{\mathbf{X}}_{\text{ref}} \subset \tilde{\mathcal{X}}_{\text{ref}}$

is a discretization of the bounded subset $\tilde{\mathcal{X}}_{\text{ref}}$. Formally, the mutual information between $\mathbf{y}_{\tilde{\mathcal{X}}}$ and $\mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}}$ is given by

$$I(\mathbf{y}_{\tilde{\mathcal{X}}}, \mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}}) = \int_{\mathcal{X}^{|\tilde{\mathcal{X}}_{\text{ref}}|} \times \mathcal{X}^{|\tilde{\mathcal{X}}|} p(\mathbf{y}_{\tilde{\mathcal{X}}}, \mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}}) \log \left(\frac{p(\mathbf{y}_{\tilde{\mathcal{X}}}, \mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}})}{p(\mathbf{y}_{\tilde{\mathcal{X}}}) p(\mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}})} \right) d\mathbf{y}_{\tilde{\mathcal{X}}} d\mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}} \quad (3)$$

respectively denote the differential entropy of $\mathbf{y}_{\tilde{\mathcal{X}}}$ and the conditional differential entropy of $\mathbf{y}_{\tilde{\mathcal{X}}}$ given $\mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}}$. In practice, computing (3) for a multi-step GP prediction is intractable. However, we can obtain the single most informative data point $\xi^* \in \tilde{\mathcal{X}}$ with respect to $\mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}}$ by computing the unconstrained minimum of

$$I(\mathbf{y}_{\xi}, \mathbf{y}_{\tilde{\mathcal{X}}_{\text{ref}}}) = \frac{1}{2} \log \left(\frac{(|k(\xi, \xi) + \sigma_w)| |\mathbf{K}_{\tilde{\mathcal{X}}_{\text{ref}}} + \sigma_w \mathbf{I}|)}{|\mathbf{K}_{\tilde{\mathcal{X}}_{\text{ref}}} \cup \xi + \sigma_w \mathbf{I}|} \right), \quad (4)$$

where $|\cdot|$ denotes the determinant of a square matrix. In settings with unconstrained decision spaces, sequentially computing a minimizer of (4) has been shown to yield a solution that corresponds to at least 63% of the optimal value (Krause et al., 2008).

4. The LocAL algorithm

The system dynamics (1) considerably limit the decision space at every time step t . Furthermore, after a data point is collected, both the GP model and mutual information change. Hence, we employ a model predictive control (MPC)-based approach to steer the system towards areas of high information. Ideally, at every MPC-step t , we would like to minimize (3) with respect to a series of N_H inputs $\mathbf{U} := \{\mathbf{u}_t, \dots, \mathbf{u}_{t+N_H-1}\}$. However, this is generally infeasible, limiting its applicability in an MPC setting. Hence, we consider an approximate solution approach that sequentially computes the most informative data point by minimizing (4) separately from the MPC optimization. This is achieved as follows. At every time step t , an unconstrained minimizer ξ^* of (4) is computed. Afterwards, the MPC computes the approximate optimal inputs \mathbf{U}^* by minimizing a constrained optimization problem that penalizes the weighted distance to the reference point ξ using a positive semi-definite weight matrix \mathbf{Q} . The ensuing state is then measured, the GP model is updated, and the procedure is repeated. These steps yield the Localized Active Learning (LocAL) algorithm, which is presented in Algorithm 1.

The square weight matrix $\mathbf{Q} \in \mathbb{R}^{d_x+d_u} \times \mathbb{R}^{d_x+d_u}$ should be chosen such that the MPC steers the system as close to ξ as possible. This represents a system-dependent task. Alternatively, \mathbf{Q} can be chosen such that the MPC cost function corresponds to a quadratic approximation of the mutual information, e.g., such that $I(\mathbf{y}_{\tilde{\mathbf{x}}_t}, \mathbf{y}_{\xi}) \approx \sum_{t=1}^{N_H} (\xi - \tilde{\mathbf{x}}_t)^T \mathbf{Q} (\xi - \tilde{\mathbf{x}}_t)$ holds for $\tilde{\mathbf{x}}_t \approx \xi$.

4.1. Sensitivity analysis

We now provide a sensitivity analysis of (4). Specifically, we show that the difference in mutual information at ξ^* and $\tilde{\mathbf{x}}_t$ is lower bounded. To this end, we require the following assumption.

Assumption 1 *The kernel $k(\cdot, \cdot)$ is Lipschitz continuous and upper bounded by $k_{\max} > 0$.*

This assumption holds for many commonly used kernels, e.g., the squared exponential kernel.

Algorithm 1 LocAL (Localized Active Learning)

Input: $x_0, f(\cdot), k(\cdot, \cdot)$ 1: **for** $t = 0, 1, 2, 3, \dots$ **do**

2: Solve

$$\xi^* = \arg \max_{\xi \in \tilde{\mathcal{X}}} \frac{1}{2} \log \left(\frac{(k(\xi, \xi) + \sigma_w) |\mathbf{K}_{\tilde{\mathcal{X}}_{\text{ref}}} + \sigma_w \mathbf{I}|}{|\mathbf{K}_{\tilde{\mathcal{X}}_{\text{ref}}} \cup \xi + \sigma_w \mathbf{I}|} \right),$$

3: Solve

$$\mathbf{U}^* = \arg \min_{\mathbf{U} \in \mathcal{U}^{N_H}} \sum_{t=1}^{N_H} (\xi^* - \tilde{x}_t)^T \mathbf{Q} (\xi^* - \tilde{x}_t)$$

$$\text{s.t. } x_{t+\tau+1} = f(\tilde{x}_{t+\tau}) + \mu_t(\tilde{x}_{t+\tau}), \mathbf{u}_{t+\tau} \in \mathcal{U}, \quad \forall \tau \in \{0, \dots, N_H - 1\}$$

4: Apply \mathbf{u}_{t+1}^* to system5: Measure \mathbf{x}_{t+1} , set $\tilde{\mathbf{X}} = \tilde{\mathbf{X}} \cup \mathbf{x}_{t+1}$, and update GP model $\mu_t(\cdot), \sigma_t^2(\cdot)$ 6: **end for**

Theorem 1 Let $\Delta \xi_t^* := \xi^* - \tilde{x}_t$ be the difference between the augmented state and the most informative data point ξ^* at time step t . Moreover, let $\Delta I_t^* := I(\mathbf{y}_{\xi^*}, \mathbf{y}_{\tilde{\mathbf{X}}_{\text{ref}}}) - I(\mathbf{y}_{\tilde{x}_t}, \mathbf{y}_{\tilde{\mathbf{X}}_{\text{ref}}})$ denote the corresponding difference in mutual information, and let Assumption 1 hold. Then, there exists a constant $L \geq 0$, such that $\Delta I_t^* \leq |\tilde{\mathbf{X}}_{\text{ref}}| \log(1 + C(\Delta \mathbf{x}_t^*))$ holds, where $C(\Delta \mathbf{x}_t^*) := \sigma_w^{-2} \min\{k_{\max}, L(t + |\tilde{\mathbf{X}}_{\text{ref}}|)^{-1/2} \|\Delta \mathbf{x}_t^*\|_2\}$.

Here we only provide a sketch of the proof. For the full proof, the reader is referred to [Capone et al. \(2020\)](#).

Proof Sketch By using the Lipschitz continuity of $k(\cdot, \cdot)$, we are able to show that the entries of the matrix \mathbf{K} are Lipschitz continuous with respect to the data. We then employ this result to show that the posterior covariance $\sigma_t(\cdot)$ is Lipschitz continuous with respect to the data. By employing a result from [Srinivas et al. \(2012\)](#), we reformulate the mutual information as a sum of logarithmic expressions that contain the GP posterior variances up to time step t . The resulting expression can then be bounded using the aforementioned results, which yields Theorem 1.

In particular, Theorem 1 implies that the mutual information with respect to ξ^* can be approximated arbitrarily accurately by reducing the difference $\xi^* - \tilde{x}_t$. In many control applications, this can be achieved in spite of the model error, e.g., by increasing control gains [Capone and Hirche \(2019\)](#); [Umlauf and Hirche \(2020\)](#). Hence, Theorem 1 can be potentially employed to guarantee a gradual improvement in model accuracy, or even to derive the worst-case number of iterations required to learn the system.

5. Numerical Experiments

In this section, we apply the proposed algorithm to four different dynamical systems. We begin with a toy example, with which we can easily illustrate the explored portions of the state space. Afterwards, we apply the proposed approach to a pendulum, a cart-pole, and a synthetic model that generalizes the mountain car problem. The exploration is repeated 50 times for each system using different starting points x_0 sampled from a normal distribution. To quantify the performance of each approach, we compute the root mean square model error (RMSE) on 500 points sampled from a uniform distribution on the region of interest $\tilde{\mathcal{X}}_{\text{ref}}$.

We employ a squared-exponential kernel in all examples, and train the hyperparameters online using gradient-based log likelihood maximization (Rasmussen and Williams, 2006). We employ an MPC horizon of $N_H = 10$, and choose weight matrix for the MPC optimization step as

$$\mathbf{Q} = \sum_{d=1}^{d_x} \sigma_{g_d} \text{diag}(l_{1,g_d}^{-2}, \dots, l_{d_x,g_d}^{-2}) \quad \forall i = 1, \dots, N,$$

where σ_{g_d} denotes the standard deviation of the GP kernel corresponding to the d -th dimension, and $l_{1,g_d}^{-2}, \dots, l_{d_x,g_d}^{-2}$ denote the corresponding lengthscales. In order to ease the computational burden, we apply the first 7 inputs computed by the LocAL algorithm before computing a new solution.

We additionally explore each system using a greedy entropy-based cost function, as suggested in Koller et al. (2018) and Schreiter et al. (2015), and compare the results. In all three cases, the LocAL algorithm yields a better model in the regions of interest than the entropy-based algorithm.

5.1. Toy Problem

Consider the continuous-time nonlinear dynamical system

$$\dot{x} = 10(\sin(x) + \arctan(x) + u), \quad (5)$$

with state space $\mathcal{X} = \mathbb{R}$ and input space $\mathcal{U} = [-5, 5]$. We are interested in obtaining an accurate dynamical model within the region $\tilde{\mathcal{X}}_{\text{ref}} = \{[x \ u]^T \in \tilde{\mathcal{X}} \mid x \in [-\pi, \pi], u \in [-1, 1]\}$. To obtain a discrete-time system in the form of (1), we discretize (5) with a discretization step of $\Delta t = 0.1$ and set the prior model to $f(x_t, u_t) = x$. The results are displayed in Figure 1.

The LocAL algorithm yields a substantial improvement in model accuracy in every run. This is because the system stays close to the region of interest $\tilde{\mathcal{X}}_{\text{ref}}$ during the whole simulation. By contrast, the greedy entropy-based method covers a considerably more extensive portion of the state space. This comes at the cost of a poorer model on $\tilde{\mathcal{X}}_{\text{ref}}$, as indicated by the respective RMSE.

5.2. Surface exploration

We apply the LocAL algorithm to the dynamical system given by

$$\dot{x}_1 = 3u_1 + 10 \cos(5x_1) \cos(5x_2), \quad \dot{x}_2 = 3u_2 + 10 \sin(5x_1) \sin(5x_2). \quad (6)$$

This setting can be seen as a surface exploration problem, i.e., an agent navigates a surface to learn its curvature. We aim to obtain an accurate model of the dynamics within the region given by $\tilde{\mathcal{X}}_{\text{ref}} = \{[\mathbf{x}^T \ \mathbf{u}^T]^T \in \tilde{\mathcal{X}} \mid \mathbf{x} \in [-\pi/4, \pi/4]^2, \mathbf{u} \in [-1, 1]^2\}$. To run the LocAL algorithm, we employ a discretization step of $\Delta t = 0.02$ s and set the prior model to $\mathbf{f}(\tilde{\mathbf{x}}) = \mathbf{x} + \Delta t \mathbf{u}$. The results are shown in Figure 2.

The LocAL algorithm manages to significantly improve its model after 400 time steps, while the entropy-based strategy does not yield any improvement. This is because every variable of the state space is unbounded, i.e., the state space can be explored for a potentially infinite amount of time without ever reaching the region of interest $\tilde{\mathcal{X}}_{\text{ref}}$.

5.3. Pendulum

We now consider a two-dimensional pendulum, whose state $\mathbf{x} = [\vartheta, \dot{\vartheta}]$ is given by the angle ϑ and angular velocity $\dot{\vartheta}$. The input torque u is constrained to the interval $\mathcal{U} = [-10, 10]$. Our goal is to

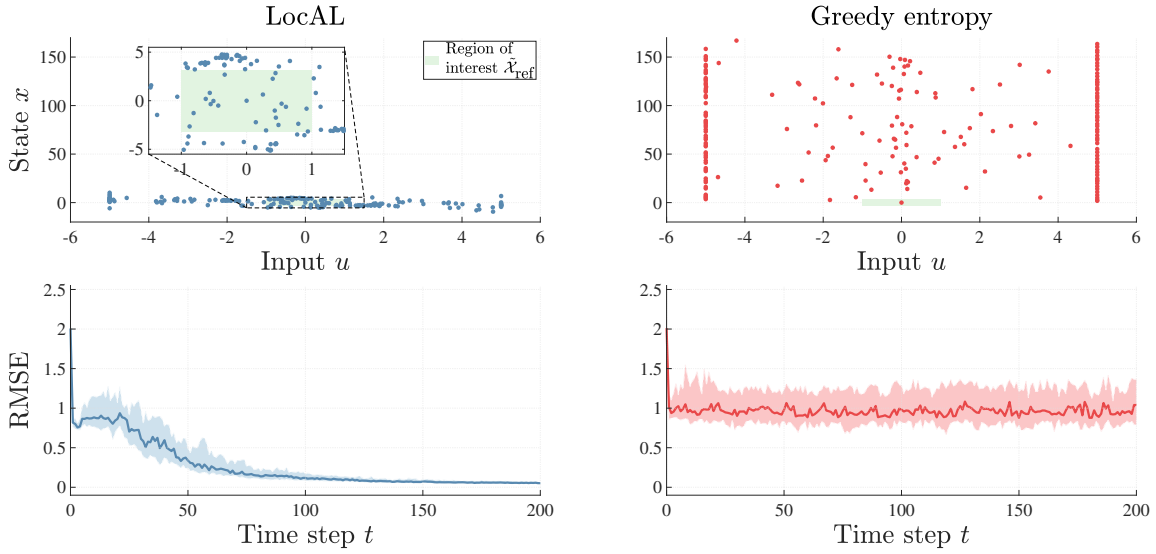


Figure 1: Toy problem results. Collected data in augmented state space $\mathcal{X} \times \mathcal{U}$ after 200 times steps (top). Median, lower and upper quartile of RMSE on region of interest (bottom). The LocAL algorithm explores the region of interest more thoroughly than the entropy-based approach




Figure 2: Surface exploration results. Median, lower and upper quartile of RMSE on region of interest.

obtain a suitable model within $\tilde{\mathcal{X}}_{\text{ref}} = \{[\mathbf{x}^T u]^T \in \tilde{\mathcal{X}} \mid x_1 \in [\pi/2, 3/2\pi], x_2 \in [-5, 5], u \in [-3, 3]\}$. Obtaining a precise model around this region is particularly useful for the commonly considered task of stabilizing the pendulum around the upward position $\vartheta = \dot{\vartheta} = 0$. The results are depicted in Figure 3.

The RMSE indicates that the LocAL algorithm yields a similar model improvement in every run. The model obtained with the entropy-based strategy, by contrast, exhibits a significantly stronger variance.

5.4. Cart-pole

We apply the LocAL algorithm to the cart-pole system (Barto et al., 1983). In this example, the state space is given by $\mathbf{x} = [v, \vartheta, \dot{\vartheta}]$, where v is the cart velocity, ϑ is the pendulum angle, and $\dot{\vartheta}$ is the angular velocity of the pendulum. Here we ignore the cart position without loss of generality, as it has no influence on the system dynamics. The region of interest is given by




Figure 3: Pendulum results. Median, lower and upper quartile of RMSE on region of interest.




Figure 4: Cart-pole results. Median, lower and upper quartile of RMSE on region of interest.

$\tilde{\mathcal{X}}_{\text{ref}} = \left\{ [\mathbf{x}^T \mathbf{u}^T]^T \in \tilde{\mathcal{X}} \mid x_1 \in [-2, 2], x_2 \in [-\pi/4, \pi/4], u_1, u_2 \in [-5, 5] \right\}$. Similarly to the pendulum case, obtaining an accurate model on this region is useful to address the balancing task. The discretization step is set to $\Delta t = 0.05$ s, the prior model is $\mathbf{f}(\tilde{\mathbf{x}}) = \mathbf{x}$. The results are shown in Figure 4.

Similarly to the pendulum case, the model obtained with the LocAL algorithm exhibits low variance compared to the one obtained with the entropy-based approach. This is because the region of interest is explored more thoroughly with our approach.

6. Conclusion

A technique for efficiently exploring bounded subsets of the state-action space of a system has been presented. The proposed technique aims to minimize the mutual information of the system trajectories with respect to a discretization of the region of interest. It employs Gaussian processes both to model the unknown system dynamics and to quantify the informativeness of potentially collected data points. In numerical simulations of four different dynamical systems, we have demonstrated that the proposed approach yields a better model after a limited amount of time steps than a greedy entropy-based approach.

Acknowledgements

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 .

References

- T. Alpcan and I. Shames. An information-based learning approach to dual control. *IEEE transactions on neural networks and learning systems*, 26(11):2736–2748, 2015.
- N. Ay, N. Bertschinger, R. Der, F. Güttler, and E. Olbrich. Predictive information and explorative behavior of autonomous robots. *The European Physical Journal B*, 63(3):329–339, 2008.
- A. G. Barto, R. S. Sutton, and C. W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5): 834–846, 1983.
- T. Beckers, D. Kulić, and S. Hirche. Stable Gaussian process based tracking control of euler-lagrange systems. *Automatica*, (103):390–397, 2019.
- F. Berkenkamp and A. P. Schoellig. Safe and robust learning control with Gaussian processes. In *Control Conference (ECC), 2015 European*, pages 2496–2501. IEEE, 2015.
- W. Burgard, D. Fox, and S. Thrun. Probabilistic robotics. *The MIT Press*, 2005.
- J. Q. Candela, A. Girard, J. Larsen, and C. E. Rasmussen. Propagation of uncertainty in Bayesian kernel models - application to multiple-step ahead forecasting. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03).*, volume 2, pages II–701, April 2003.
- A. Capone and S. Hirche. Backstepping for partially unknown nonlinear systems using Gaussian processes. *IEEE Control Systems Letters*, 3:416–421, 2019.
- A. Capone, J. Umlauf, T. Beckers, A. Lederer, and S. Hirche. Localized active learning of Gaussian process state space models. 2020. URL <https://arxiv.org/submit/3157874/view>.
- P. Dayan and T. J. Sejnowski. Exploration bonuses and dual control. *Machine Learning*, 25(1): 5–22, 1996.
- M. P. Deisenroth, D. Fox, and C. E. Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2): 408–423, 2015.
- Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472, 2011.
- S. Kamthe and M. P. Deisenroth. Data-efficient reinforcement learning with probabilistic model predictive control. In *AISTATS*, volume 84, pages 1701–1710. Proceedings of Machine Learning (PMLR), 2018.
- Torsten Koller, Felix Berkenkamp, Matteo Turchetta, and Andreas Krause. Learning-based model predictive control for safe exploration. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 6059–6066, 2018.

- A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9 (Feb):235–284, 2008.
- G. Pillonetto, F. Dinuzzo, T. Chen, G. De Nicolao, and L. Ljung. Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, 50(3):657–682, 2014.
- C. E. Rasmussen and C. K. Williams. Gaussian processes for machine learning. 2006. *The MIT Press, Cambridge, MA, USA*, 38:715–719, 2006.
- F. Schreiter, D. Nguyen-Tuong, M. Eberts, B. Bischoff, H. Markert, and M. Toussaint. Safe exploration for active learning with Gaussian processes. In *Machine Learning and Knowledge Discovery in Databases*, pages 133–149, Cham, 2015. Springer International Publishing.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias W Seeger. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.
- J. Umlauft and S. Hirche. Feedback linearization based on Gaussian processes with event-triggered online learning. *IEEE Transactions on Automatic Control*, 2020.
- J. Umlauft, L. Pöhler, and S. Hirche. An uncertainty-based control Lyapunov approach for control-affine systems modeled by Gaussian process. *IEEE Control Systems Letters*, 2018.