

Closed-loop Model Selection for Kernel-based Models using Bayesian Optimization

Thomas Beckers¹, Somil Bansal², Claire J. Tomlin² and Sandra Hirche¹

Abstract—Kernel-based nonparametric models have become very attractive for model-based control approaches for nonlinear systems. However, the selection of the kernel and its hyperparameters strongly influences the quality of the learned model. Classically, these hyperparameters are optimized to minimize the prediction error of the model but this process totally neglects its later usage in the control loop. In this work, we present a framework to optimize the kernel and hyperparameters of a kernel-based model directly with respect to the closed-loop performance of the model. Our framework uses Bayesian optimization to iteratively refine the kernel-based model using the observed performance on the actual system until a desired performance is achieved. We demonstrate the proposed approach in a simulation and on a 3-DoF robotic arm.

I. INTRODUCTION

Given a dynamic model, control mechanisms such as model predictive control and feedback linearization can be used to effectively control nonlinear systems. However, when an accurate mathematical model of the system is not available, machine learning offers powerful tools for the modeling of dynamical systems. A special class of models that has obtained a lot of attention recently is kernel-based models, such as Support Vector Machines (SVM) and Gaussian Processes (GP). In contrast to parametric models, kernel-based models require only minimal prior knowledge about the system dynamics, and have been successfully used to model complex, nonlinear systems [1]. Using the kernel-based approach for modeling a system requires the selection of an appropriate kernel function and a set of hyperparameters for that function. Typically, these selections are data-based, e.g. through minimizing a loss function that is often a trade-off between the prediction error and the complexity of the model. However, the full complex and accurate dynamics model might not even be required depending on the task. Moreover, this procedure neglects the fact that the learned model is used for the control of the actual system, which can result in reduced controller performance [2]–[4].

In this work, we propose a Bayesian Optimization (BO)-based active learning framework to optimize the kernel and its hyperparameters directly with respect to the performance of the closed-loop rather than the prediction error, see Fig. 1. This optimization is performed in a sequential fashion where at each step of the optimization, BO takes into account all the past data points and proposes the most promising kernel

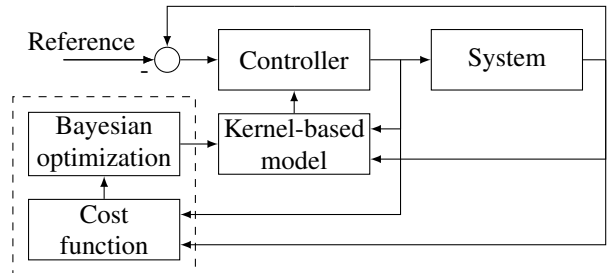


Fig. 1. Closed-loop model selection for kernel-based models. BO is used to optimize the kernel and its hyperparameters directly based on the evaluation of a cost function.

and hyperparameters for the next evaluation. The outcome is used to define a kernel-based model that is utilized by a given controller. The obtained model-based controller is then applied to the actual system in a closed-loop fashion to evaluate its performance. This information is then used by BO to optimize the next evaluation. Consequently, multiple evaluations on the actual system must be performed, which is often feasible such as for systems with repetitive trajectories. BO thus does not aim to obtain the most accurate dynamics model of the system, but rather to optimize the performance of the closed-loop system.

Typically, system identification approaches aim to obtain an open-loop dynamics model of the system by minimizing the state prediction error. This problem has been well studied in literature for both linear systems, e.g. [5], as well as for nonlinear systems using the function approximators such as GP [6]–[8] and neural networks (NN) [9], [10]. However, a model obtained using this open-loop procedure can result in a reduced controller performance on the actual system [4]. To overcome these challenges, adaptive control mechanisms and iterative learning control have been studied where the system dynamics or control parameters are optimized based on the performance on the actual system, e.g. [11]–[13]. However, these approaches are mostly limited to linear systems and controllers or assume at least a parametric system model. Recently, learning-based controller tuning mechanisms have also been proposed [14], [15], but such methods might be highly data-inefficient for general nonlinear systems as they typically completely disregard the underlying dynamics [16].

To overcome the challenges of open-loop system identification, closed-loop system identification methods have been studied that lead to more robust control performance on the actual system [2]–[4]. A similar approach is presented in [17], wherein the authors also propose a goal-driven dynamics learning approach via BO. However, the authors aim to identify a linear dynamics model from *scratch* which might

¹ are with the Chair of Information-oriented Control (ITR), Department of Electrical and Computer Engineering, Technical University of Munich, 80333 Munich, Germany, {t.beckers, hirche}@tum.de

² are with the Department of Electrical Engineering and Computer Sciences, UC, Berkeley, USA, {somil, tomlin}@eecs.berkeley.edu

be a) unnecessary, as often an approximate dynamics model of the system is available and b) insufficient for general nonlinear systems. Moreover, stability of the closed-loop system where the controller is based on the linear dynamics model cannot be guaranteed, whereas our approach explicitly allows to preserve the convergence properties of the initial closed-loop system. To summarize, our key contributions are: a) we present a BO-based framework to optimize the kernel function and its hyperparameters of a kernel-based model to maximize the resultant control performance on the actual system; b) through numerical examples and an experiment on a real 3-DoF robot, we demonstrate the advantages of the proposed approach over classical model selection methods.

Notation: Vectors \mathbf{a} are denoted with bold characters. Matrices A are described with capital letters. The term $A_{i,:}$ denotes the i -th row of the matrix A . The expression $\mathcal{N}(\mu, \Sigma)$ describes a normal distribution with mean μ and covariance Σ . The set $\mathbb{R}_{>0}$ denotes the set of positive real numbers.

II. PROBLEM SETTING

Consider a discrete-time, potentially nonlinear system

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k), \quad k = \{0, \dots, n-1\}, n \in \mathbb{N} \\ \mathbf{y}_k &= \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k) \end{aligned} \quad (1)$$

in which \mathbf{f} , \mathbf{g} are unknown functions of the state $\mathbf{x}_k \in \mathbb{R}^{n_x}$ and input $\mathbf{u}_k \in \mathbb{R}^{n_u}$. For the following, we assume that the state mapping $\mathbf{f}: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ and the output mapping $\mathbf{g}: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_y}$ are such that there exist a unique state and output trajectory for all $\mathbf{u}_k \in \mathbb{R}^{n_u}$ and \mathbf{x}_0 , $k \geq 0$. We assume that a control law $\mathbf{h}: \mathbb{R}^{n_y} \times \mathbb{R}^{n_m} \rightarrow \mathbb{R}^{n_u}$

$$\mathbf{u}_k = \mathbf{h}(\mathbf{y}_k - \mathbf{r}_k, \mathbf{m}_k) \quad (2)$$

is given for the system (1). The reference $\mathbf{r}_k \in \mathbb{R}^{n_y}$ is assumed to be zero but the framework is also applicable for a varying signal. In addition to the reference, the control law also depends on the output $\mathbf{m}_k \in \mathbb{R}^{n_m}$ of a kernel-based model, a regression technique that uses a kernel to perform the regression in a higher-dimensional feature space. The output of a kernel-based model, \mathbf{m}_k , depends on the kernel function \mathfrak{K} , its hyperparameters $\varphi \in \mathbb{R}^{n_\varphi}$ and system input and output, i.e. $\mathbf{m}_k = \mathcal{M}(\mathbf{u}_{0:k-1}, \mathbf{y}_{0:k}, \mathfrak{K}, \varphi)$, where the function \mathcal{M} depends on the class of the kernel-based model, such as GP or SVM, used for the prediction.

Remark 1 For example, the output \mathbf{m}_k can be the prediction of the next state or output of the system based on the current state and input, using the mean and probably the variance of a GP model. This information can then be used by the controller to compute an appropriate system input \mathbf{u}_k .

The control law \mathbf{h} might be an output tracking controller designed based on the predicted model output. For possible control laws for different classes of systems, we refer to [6]–[8], [18], [19]. The goal of this work is to optimize the model kernel and its hyperparameters such that the corresponding model output \mathbf{m}_k minimizes the following cost functional

$$C(\mathbf{y}_{0:k}, \mathbf{u}_{0:k}) = \sum_{k=0}^{n-1} c(\mathbf{y}_k, \mathbf{u}_k), \quad (3)$$

where $c(\mathbf{y}_k, \mathbf{u}_k): \mathbb{R}^{n_y} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ represents the cost incurred for the control input \mathbf{u}_k and the system output \mathbf{y}_k . The cost function here might represent the requirements concerning the closed-loop, e.g. an accurate tracking behavior or a minimized power consumption. Note that the cost functional in (3) implicitly depends on the kernel-based model \mathcal{M} through \mathbf{u}_k , see (2). The optimization of (3) is challenging since the system dynamics in (1) are unknown and the kernel-based model output \mathbf{m}_k indirectly influences the cost. To overcome this challenge, we use BO to optimize the kernel and the hyperparameters based on the direct evaluation of the control law in (2) on the system (1) to find those that minimize the cost functional in (3).

III. PRELIMINARIES

A. Kernel-based models

The prediction of parametric models is based on a parameter vector $\mathbf{w} \in \mathbb{R}^{n_a}$ which is typically learned using a set of training data points. In contrast, nonparametric models typically maintain a subset of the training data points in memory in order to make predictions for new data points. Many linear models can be transformed into a dual representation where the prediction is based on a linear combination of kernel functions. The idea is to transform the data points of a model to an often high-dimensional feature space where a linear regression can be applied to predict the model output. For a nonlinear feature map $\phi: \mathbb{R}^{n_a} \rightarrow \mathbb{R}^{n_\phi}$ with $n_\phi \in \mathbb{N} \cup \{\infty\}$, the kernel function is given by the inner product $\mathfrak{K}(\mathbf{a}, \mathbf{a}') = \langle \phi(\mathbf{a}), \phi(\mathbf{a}') \rangle, \forall \mathbf{a}, \mathbf{a}' \in \mathbb{R}^{n_a}$.

Thus, the kernel implicitly encodes the way the data points are transformed into a higher dimensional space. The formulation as inner product in a feature space allows to extend many standard regression methods. A drawback of kernel-based models is that the selection of the kernel and its hyperparameters heavily influences the interpretation of the data and thus, the quality of the model. Commonly, the kernel and hyperparameters are determined based on the optimization of a loss function such as cross-validation or the likelihood function. In our work, the kernel and its hyperparameters are optimized with respect to performance of the closed-loop system.

B. Gaussian process

Extending the concept of kernel functions to probabilistic models leads to the framework of Gaussian process regression (GPR). In particular, GPR is a supervised learning technique which combines several advantages. As probabilistic kernel techniques, GPs provides not only a mean function but also a measure for the uncertainty of the regression. In this work, we use GPR in BO to model the unknown *closed-loop* objective function, as well as for the kernel-based dynamics model \mathcal{M} in the experiment. The GPR can be derived using a standard linear regression model

$$q(\mathbf{a}) = \mathbf{a}^\top \mathbf{w}, \quad b = q(\mathbf{a}) + \epsilon$$

where $\mathbf{a} \in \mathbb{R}^{n_a}$ is the input vector, \mathbf{w} the vector of weights and $q: \mathbb{R}^{n_a} \rightarrow \mathbb{R}$ the function value. The observed value $b \in \mathbb{R}$

is corrupted by Gaussian noise $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$. Using the feature map $\phi(\mathbf{a})$ instead of \mathbf{a} , leads to $f(\mathbf{a}) = \phi(\mathbf{a})^\top \mathbf{w}$ with $f: \mathbb{R}^{n_a} \rightarrow \mathbb{R}$. The analysis of this model is analogous to the standard linear regression, i.e. we put a prior on the weights such that $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \Sigma_p)$ with $\Sigma_p \in \mathbb{R}^{n_\phi \times n_\phi}$. The mean function is usually defined to be zero, see [20]. Based on m collected training data points $A = [\mathbf{a}_1, \dots, \mathbf{a}_m]$ and $B = [b_1, \dots, b_m]^\top$, the prediction $q_* \in \mathbb{R}$ for a new test point $\mathbf{a}_* \in \mathbb{R}^{n_a}$ can be computed using the Bayes' rule. In particular, it is given by

$$q_* | \mathbf{a}_*, A, B \sim \mathcal{N}(\mathbf{k}_*^\top K_{**}^{-1} B, \mathbf{k}_*^\top K_{**}^{-1} \mathbf{k}_*),$$

where $\mathfrak{K}(\mathbf{a}, \mathbf{a}') = \phi(\mathbf{a})^\top \Sigma_p \phi(\mathbf{a}')$, $\mathbf{k}_{**} = \mathfrak{K}(\mathbf{a}_*, \mathbf{a}_*)$ and $\mathbf{k}_* = [\mathfrak{K}(\mathbf{a}_*, A_{1,:}), \dots, \mathfrak{K}(\mathbf{a}_*, A_{m,:})]^\top$. The covariance matrix $K_{**} = (K + \sigma_n^2 I)$ is defined by $K_{i,j} = \mathfrak{K}(\mathbf{a}_i, \mathbf{a}_j)$. Thus, based on the training data A, B , the estimation of the function value q_* follows a normal distribution where the mean and the variance depend on the test input \mathbf{a}_* . Following Remark 1, the mean and variance can be used for state estimation in the control law (2). The choice of the kernel and hyperparameters $\varphi \in \mathbb{R}^{n_\varphi}$ can be seen as degrees of freedom of the regression. A popular kernel choice in GPR is the squared exponential kernel, see [20]. One possibility for estimating the hyperparameters φ is by means of the likelihood function, thus by maximizing the probability of

$$\varphi^* = \frac{1}{2} (B^\top K_{**}^{-1} B + \log |K_{**}| + m \log 2\pi) \quad (4)$$

which results in an automatic trade-off between the data-fit $B^\top K_{**}^{-1} B$ and model complexity $\log |K_{**}|$, see [20].

C. Bayesian Optimization (BO)

Bayesian Optimization is an approach to minimize an unknown objective function based on (only a few) evaluated samples. We use BO to optimize the cost function (3) based on the kernel-based model as this is in general a non-convex optimization problem with unknown objective function (because the system dynamics are unknown), and probably multiple local extrema. BO is well-suited for this optimization as the task evaluations can be expensive and noisy [21]. Furthermore, BO is a gradient-free optimization method which only requires that the objective function can be evaluated for any given input. Since the objective function is unknown, the Bayesian strategy is to treat it as a random function with a prior, often as Gaussian process. Note that this GP here is used for the closed-loop cost functional approximation in BO and is not related to the kernel-based model for the controller (2) as stated in Remark 1. The prior captures the beliefs about the behaviour of the function, e.g. continuity or boundedness. After gathering the cost (3) of the task evaluation, the prior is updated to form the posterior distribution over the objective function. The posterior distribution is used to construct an acquisition function that determines the most promising kernel/hyperparameters for the next evaluation to minimize the cost. Different acquisition functions are used in literature to trade off between exploration of unseen kernel/hyperparameters and exploitation of promising

combinations during the optimization process. Common acquisition functions are expected improvement, entropy search, and upper confidence bound [22]. To escape a local objective function minimum, the authors of [23] propose a method to modify the acquisition function when they seem to over-exploit an area, namely expected-improvement-plus. That results in a more comprehensive and also partially random exploration of the area and, thus it is probably faster in finding the global minimum. We also use this acquisition function for BO in our simulation and the experiment.

IV. CLOSED-LOOP MODEL SELECTION

Our goal is to optimize the model's kernel and its hyperparameters with respect to the cost functional $C(\mathbf{y}_{0:k}, \mathbf{u}_{0:k})$. Thus, in contrast to the classical kernel selection problem, where the kernel is selected to minimize the state prediction error, our goal here is not to get the most accurate model but the one that achieves the best closed-loop behavior. We now describe the proposed overall procedure for the kernel selection to optimize the closed-loop behavior; we then describe each step in detail.

We start with an initial kernel \mathfrak{K} with hyperparameters φ , and obtain the control law for the system (1) using (2) with the model output $\mathbf{m}_k = \mathcal{M}(\mathbf{u}_{0:k-1}, \mathbf{y}_{0:k}, \mathfrak{K}, \varphi)$. This control law is then applied to the actual system, and the cost function (3) is evaluated after performing the control task. Depending on the obtained cost value, BO suggests a new kernel and corresponding hyperparameters for the kernel-based model \mathcal{M} in order to minimize the cost function on the actual system. With this model, the control task is repeated and, based on the cost evaluation, BO suggests the next kernel and hyperparameters. This procedure is continued until a maximum number of task evaluations is reached or the user rates the closed-loop performance as sufficient enough. We now describe the above three steps, i.e. initialization, evaluation and optimization, in detail.

A. Initialization

We define a set $\mathcal{K} = \{\mathfrak{K}^1, \dots, \mathfrak{K}^{n_{\mathfrak{K}}}\}$ of $n_{\mathfrak{K}} \in \mathbb{N}$ kernel candidates \mathfrak{K}^j that we want to choose the kernel from for our kernel-based model. BO will be used to select the kernel with the best closed-loop performance in this set.

Remark 2 *The selection of possible kernels can be based on prior knowledge about the system, e.g. smoothness with the Matérn kernel or number of equilibria using a polynomial kernel, see [24] and [1] for general properties, respectively.*

In addition, each kernel depends on a set of hyperparameters. Since the number of hyperparameters could be different for each kernel, we define a set of sets $\mathcal{P} = \{\Phi^1, \dots, \Phi^{n_{\mathfrak{K}}}\}$ such that $\Phi^j \subset \mathbb{R}^{n_{\Phi^j}}$ is a closed set. Here, n_{Φ^j} represents the number of hyperparameters for the kernel \mathfrak{K}^j . Moreover, we assume that Φ^j is a valid *hyperparameter set*.

Definition 1 *The set Φ is called a hyperparameter set for a kernel function \mathfrak{K} iff the set Φ is a domain for the hyperparameters of \mathfrak{K} .*

For the first evaluation of the closed-loop, the kernel-based model function \mathcal{M} is created with an initial kernel \mathfrak{K}^j of the set \mathcal{K} and hyperparameters $\varphi^j \in \Phi^j$ with $j \in \{1, \dots, n_{\mathfrak{K}}\}$.

Remark 3 *One potential way to select the initial kernel and hyperparameters is to set them equal to the kernel and hyperparameters of a prediction model that is optimized with respect to a loss function, e.g., using cross-validation or maximization of the likelihood function [1].*

B. Task Evaluation

For the i -th task evaluation, BO determines an index value $j \in \{1, \dots, n_{\mathfrak{K}}\}$ and a $\varphi^j \in \Phi^j$. The control law (2) for the kernel-based model \mathcal{M} , with the determined kernel \mathfrak{K}^j and hyperparameters φ^j , is applied to the system (1)

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{f}(\mathbf{x}_k, \mathbf{h}(\mathbf{y}_k, \mathcal{M}(\mathbf{u}_{0:k-1}, \mathbf{y}_{0:k}, \mathfrak{K}^j, \varphi^j))) \\ \mathbf{y}_k &= \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k) \text{ for } k = \{0, \dots, n-1\} \end{aligned}$$

with fixed $\mathbf{x}_0 \in \mathbb{R}^{n_x}$.

Remark 4 *We focus here on a single, fixed initial state \mathbf{x}_0 . However, multiple (close by) initial states can be considered by using the expected cost across all initial states.*

The corresponding input and output sequences $\mathbf{u}_{0:k}$ and $\mathbf{y}_{0:k}$, respectively, are recorded. Afterwards, the cost function given by $C(\mathbf{y}_{0:k}, \mathbf{u}_{0:k})$ is evaluated.

C. Model Optimization

In the next step, we use BO to minimize the cost function with respect to the kernel and its hyperparameters, i.e.

$$[\mathfrak{K}^j, \varphi^j] = \underset{j \in \{1, \dots, n_{\mathfrak{K}}\}, \varphi^j \in \Phi^j}{\operatorname{argmin}} C(\mathbf{y}_{0:k}, \mathbf{u}_{0:k}). \quad (5)$$

Thus, this problem involves continuous and discrete variables in the optimization task whereas classical BO assumes continuous variables only. To overcome this restriction, a modified version of BO is used where the kernel function is transformed in a way such that integer-valued inputs are properly included [25].

Based on previous evaluations of the cost function, BO updates the prior and minimizes the acquisition function. The result is a kernel \mathfrak{K}^j and hyperparameters φ^j which are used in the model function $\mathcal{M}(\mathbf{u}_{0:k-1}, \mathbf{y}_{0:k}, \mathfrak{K}^j, \varphi^j)$. Then, the corresponding control law is evaluated again on the system and the procedure is repeated until a maximum number of task evaluations has been reached or a sufficient performance level has been achieved.

D. Theoretical Analysis

In this section, we show that, under some additional assumptions, the stability of the closed-loop is preserved during the task evaluation process and that BO converges to the minimum of the closed-loop cost function. Here, we focus on stationary kernels

$$k(\mathbf{x}, \mathbf{x}') = k((\mathbf{x} - \mathbf{x}')^\top \Sigma^{-1} (\mathbf{x} - \mathbf{x}')), \quad \mathbf{x}, \mathbf{x}' \in \mathbb{R}^{n_x}$$

with lengthscales $\varphi \in \mathbb{R}_{>0}^{n_\varphi}$ and $\Sigma = \operatorname{diag}(\varphi_1, \dots, \varphi_{n_x})$. Stationary kernels can always be expressed as a function of the difference between their inputs and they are a common choice for kernel-based models [1].

Assumption 1 *Let $\|\mathbf{f}\|_{\mathfrak{K}^*} < \infty$ and the selected control law (2), based on the model \mathcal{M} with stationary kernel \mathfrak{K}^* and hyperparameters $\varphi^* \in \mathbb{R}_{>0}^{n_\varphi}$, guarantees that $\|\mathbf{y}_k\| \leq r_y \in \mathbb{R}_{>0}$ for the given system (1) for $k > n_1 \in \mathbb{N}$.*

The first part of the assumption, i.e. the bounded reproducing kernel Hilbert space (RKHS) norm, is a measure of smoothness of the function with respect to the kernel \mathfrak{K} with hyperparameters $\varphi^* \in \mathbb{R}_{>0}^{n_\varphi}$. It is a common assumption for stabilizing controllers using kernel-based methods and is discussed in more detail in [19]. Controllers that satisfy this property for nonlinear, unknown systems are given, e.g. by [6], [19], [26]. The focus on stationary kernels is barely restrictive as many successful applied kernels for nonlinear control are stationary.

Lemma 1 *With Assumption 1, there exists a non-empty set \mathcal{K} and a hyperparameter set $\Phi^1 \supset \{\varphi^*\}$ such that $\forall \mathfrak{K}^j \in \mathcal{K}$, for all $\varphi^j \in \Phi^j$ the boundedness $\|\mathbf{y}_k\| \leq r_y$ of the system (1) for $k > n_1$ is preserved.*

This lemma guarantees that there exists a kernel set \mathcal{K} and a set \mathcal{P} of hyperparameters that contains the stabilizing kernel \mathfrak{K}^* and the hyperparameter φ^* of Assumption 1. Thus, the proposed method can be applied to existing kernel-based control methods without losing achieved guarantees. Before we start with the proof, the following lemma is recalled.

Lemma 2 ([23, Lemma 4]) *If $\mathbf{f} \in \mathcal{H}_{\mathfrak{K}^*}$ then $\mathbf{f} \in \mathcal{H}_{\mathfrak{K}^{\varphi'}}$ holds for all $0 < \varphi'_i \leq \varphi_i, \forall i \in \{1, \dots, n_\varphi\}$, and*

$$\|\mathbf{f}\|_{\mathfrak{K}^{\varphi'}}^2 \leq \left(\prod_{i=1}^{n_\varphi} \frac{\varphi_i}{\varphi'_i} \right) \|\mathbf{f}\|_{\mathfrak{K}^*}^2.$$

Proof: [Lemma 1] Assumption 1 inherently guarantees that at least one kernel $\mathfrak{K}^1 = \mathfrak{K}^*$ exist that preserves the boundedness of the system such that we define $\mathcal{K} = \{\mathfrak{K}^1\}$. Since Assumption 1 ensures that $\|\mathbf{f}\|_{\mathfrak{K}^*}$ is bounded and with Lemma 2, the mapping $\mathbf{f} \in \mathcal{H}_{\mathfrak{K}^1}$ and, thus $\|\mathbf{f}\|_{\mathfrak{K}^1}$ is bounded for $\varphi_i \in \mathbb{R}_{>0}, \forall i$ where $\varphi_i < \varphi_i^*, \forall i$. For an upper bound, there exist $\bar{\varphi}_i \in \mathbb{R}_{>0}, \forall i$ such that $\varphi_i^* < \bar{\varphi}_i$ and $\mathbf{f} \in \mathcal{H}_{\mathfrak{K}^{\bar{\varphi}}}$, following Lemma 2. Thus, we define the set

$$\Phi^1 = \{\varphi: \varphi_i \leq \varphi_i \leq \bar{\varphi}_i, \forall i\} \quad (6)$$

as proper superset of φ^* . Based on this set, $\|\mathbf{f}\|_{\mathfrak{K}^1} < \infty$ for all $\varphi \in \Phi^1$ that guarantees the boundedness. ■

Consequently, with Assumption 1, the stability of the control loop is preserved during the task evaluation. Furthermore, the minimum cost is not worse than the initial cost after BO as stated in the following.

Corollary 1 *Let C_{cl} be the minimum cost (3) after BO (5) with $\mathcal{K} = \{\mathfrak{K}^1 = \mathfrak{K}^*\}$ and Φ^1 of (6). Let C_{ol} be the initial cost based on the control with kernel \mathfrak{K}^* and hyperparameter φ^* then $C_{cl} \leq C_{ol}$ holds.*

Proof: Since C_{cl} is the minimum cost after BO that starts with the initial, data-based selected kernel \mathfrak{K}^* and hyperparameter φ^* , it clearly follows that $C_{cl} \leq C_{ol}$ because of $\mathfrak{K}^* \in \mathcal{K}$ and $\varphi^* \in \Phi^1$. ■

We now show that BO can converge to the global minimum of the cost function C under specific conditions starting with the following assumption.

Assumption 2 *The RKHS norm of the cost function is bounded, i.e. $\|C\|_{\mathfrak{K}} \leq r \in \mathbb{R}_{>0}$ with respect to the kernel \mathfrak{K} of the GP (4) that is used as prior $C \sim \mathcal{GP}(0, \mathfrak{K})$ of the Bayesian optimization (5).*

Intuitively, Assumption 2 states that the kernel of the GP for BO is selected such that the GP can properly approximate the cost function. This sounds paradoxical since the cost function is unknown because of the unknown system behavior. However, there exist some kernels, so called universal kernels, which can approximate at least any continuous function arbitrarily precisely [27, Lemma 4.55].

Lemma 3 ([28]) *With Assumption 2, BO in (5) with upper confidence bound acquisition function [28, Eq.(6)] converge with high probability to the global minimum of C .*

V. EVALUATION

In this section, we present a simple illustrative example that highlights our closed-loop model selection approach for kernel-based models. In addition, an example with a 3-DOF robot demonstrates the applicability of the proposed approach to hardware testbeds. BO is used with the expected-improvement-plus as acquisition function because of its satisfactory performance in practical applications, see [23], using a GP as prior.

A. Simulation

Consider the following one-dimensional system

$$x_{k+1} = \exp\left(-\frac{1}{100}x_k^2\right)\sin(x_k) + \frac{1}{3}x_k + u_k, y_k = x_k \quad (7)$$

with state x_k and control u_k at time k . For the purpose of this example, we assume that the system dynamics in (7) are unknown yet we wish to avoid a high-gain control approach due to its unfavorable properties [29], and use the proposed closed-loop model selection framework to optimize the control performance. As control law, a feedback linearization

$$u_k = -\hat{f}(x_k|\mathcal{M}, \mathcal{D}) + \frac{1}{2}x_k \quad (8)$$

is applied with the prediction \hat{f} of a Support vector machine model \mathcal{M} . The data set \mathcal{D} consists of 11 homogeneously distributed training pairs $\{x_k^j, x_{k+1}^j\}_{j=1}^{11}$ of the system (7) in the interval $x_k \in [-10, 10]$ with $u_k = 0$. The linear, polynomial (cubic) and the Gaussian kernel are selected as possible kernel candidates, see Table I for details. The Gaussian kernel possesses one hyperparameter φ_1 which is a scaling factor for the data. In addition, the regression of the SVM depends on a hyperparameter φ_2 that defines the smoothness of the prediction and affects the number of support vectors, see [30].

First, we evaluate a classical, data-based procedure which optimizes the kernel and the hyperparameters with respect to the cross-validation loss function [27] based on the training data only. Using BO, a minimum loss of 0.9127 is found

using the linear kernel with $\varphi_2 = 0.0336$, Table II. Using this linear model in the control loop with the nonlinear system (7) and control law (8) for $x_0 = 3$, the control error remains above zero, see Fig. 2. With the cost function $C = \sum_{k=0}^9 kx_k^2$, the trajectory generates a cost of 204.4769. In comparison, the hyperparameters and the kernel are optimized with the proposed method. For this purpose, we evaluate the performance of the closed-loop system and use BO to compute the next promising kernel and hyperparameter combination. Figure 3 shows the mean and standard deviation of 20 repetitions over 50 trials each. The repetitions are run since BO exploration of the cost is also affected by randomness. The cost is reduced to a mean value of $C = 16.410$ and the loss is 2.491. Figure 2 shows that the regression is more accurate which results in a reduced control error. Table II also presents the results for adding the collected data of all the 50 trials to the existing data to redefine the model (Data-based AT). Even with more training data, the data-based optimization favors the linear kernel.

1) *Discussion:* The example demonstrates that the optimization based on the training data only can lead to a reduced performance of the closed-loop system. Table II clearly shows that the data-based optimization results in a smaller loss with the linear kernel but generates a higher cost of the closed-loop system. In comparison, the closed-loop optimization finds a set of hyperparameters with the Gaussian kernel that significantly reduced the control error even if the loss of the model is higher. Thus, especially in the case of sparse data, the data-based optimization can misinterpret the data which can be avoided with the closed-loop model selection. We observe that at the beginning of the closed-loop optimization, BO switches a lot between the kernels and towards the end, it focus on the hyperparameters. Using the data which is obtained during the 50 trials to refine the model in data-based manner only slightly improves the performance but heavily increases the computational time of the kernel-based model due to the larger training data set.

B. Experiment

1) *Setup:* For the experimental evaluation, we use the 3-dof SCARA robot CARBO as pictured in Fig. 4. The links between the joints have a length of 0.3m and a

TABLE I
KERNEL CANDIDATES

Kernel	Formula
Linear	$\mathfrak{K}(\mathbf{x}, \mathbf{x}') = \mathbf{x}^\top \mathbf{x}'$
Polynomial (cubic)	$\mathfrak{K}(\mathbf{x}, \mathbf{x}') = (1 + \mathbf{x}^\top \mathbf{x}')^3$
Gaussian	$\mathfrak{K}(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\ \mathbf{x} - \mathbf{x}'\ ^2}{\varphi_1^2}\right), \varphi_1 \in \mathbb{R}$

TABLE II
COMPARISON BETWEEN DATA-BASED, DATA-BASED WITH ADDITIONAL TRAINING DATA AND CLOSED-LOOP OPTIMIZATION

Method	Selected kernel	φ_1, φ_2	Loss	Cost
Data-based	Linear	$-, 0.034$	0.913	204.477
Data-based AT	Linear	$-, 0.301$	0.09	199.634
Closed-loop	Gaussian	$2.333, 0.013$	2.491	16.410




Fig. 2. Control error using closed-loop model optimization for 20 repetitions with mean and 5σ deviation (blue) and data-based model selection (red).




Fig. 3. Minimum of the cost function over the number of trials for 20 repetitions for the closed-loop model selection algorithm.

spoon is attached at the end effector of the robot. The goal is to follow a given trajectory as precisely as possible without using high feedback gains, which might result in several practical disadvantages, see [31]. Therefore, a precise model of the system’s dynamics is necessary. Since the modeling of the nonlinear fluid dynamics with a parametric model would be very time consuming, we use a computed torque control method based on a GP model which allows high performance tracking control while also being able to guarantee the stability of the control loop [26]. Underlying, a low level PD-controller enforces the generated torque by regulating the voltage based on a measurement of the current. The controller is implemented in MATLAB/Simulink on a Linux real-time system with a sample rate of 1 ms. For the implementation of the GP model, we use the GPML toolbox¹. The desired trajectory follows a circular stirring movement through the fluid with a frequency of 0.5 Hz.

Modeling: Here, we use a Gaussian process model \mathcal{M} as kernel-based model technique based on 223 collected training points. The data is collected around the desired trajectory using a high gain controller. The placement of the training points heavily influences the control performance. However, the proposed approach focuses on improving the performance based on existing data. Each data pair consists of the position and velocity of all joints $[\mathbf{q}, \dot{\mathbf{q}}]^\top$ and the corresponding torque for the i -th joint, τ_i . Since the GP produces one-dimensional outputs only, 3 GPs are used in total for the modeling of the robot’s dynamics. Each GP $i = 1, \dots, 3$ uses a squared exponential kernel

$$\mathcal{K}(\mathbf{x}, \mathbf{x}') = \varphi_i^2 \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|^2}{\varphi_{i+3}^2}\right), \varphi_i \in \mathbb{R} \setminus \{0\}$$

¹<http://www.gaussianprocess.org/gpml/code>




Fig. 4. Stirring with the 3-dof SCARA robot CARBO.

that can approximate any continuous function arbitrarily exactly. With $\varphi = [\varphi_1, \dots, \varphi_6]$ and the signal noise $\sigma_n \in \mathbb{R}^3$, see [20], a total number of 9 parameters must be optimized. In contrast to the simulation, the kernel is fixed to reduce the optimization space and thus, the number of task evaluations. **Control law:** The control input, i.e. the torque $\tau(\dot{\mathbf{q}}, \mathbf{q})$ for all joints, is generated based on an estimated parametric model and the mean prediction μ of the GP model as feed-forward component and a low gain PD-feedback part

$$\tau_d = \hat{H}\ddot{\mathbf{q}}_d + \hat{C}\dot{\mathbf{q}}_d + \hat{g} + \mu(\dot{\mathbf{q}}, \mathbf{q}|\mathcal{M}) - K_d\dot{e} - K_p e.$$

Here, the desired trajectory is given by \mathbf{q}_d , $\dot{\mathbf{q}}_d$ and $\ddot{\mathbf{q}}_d$ with the error $\dot{e} = \dot{\mathbf{q}}_d - \dot{\mathbf{q}}$, $e = \mathbf{q}_d - \mathbf{q}$. The feedback matrices are given by $K_p = \text{diag}([60, 40, 10])$ and $K_d = \text{diag}([1, 1, 0.4])$. The estimated parametric model is derived from a mathematical model where the parameters are physically measured. For the discretization of the control input, a zero-order method is used. For more details see [26].

2) *Evaluation:* The evaluation of the performance of the closed-loop is based on the cost function

$$C = \frac{1}{2000} \sum_{k=0}^{2000} e(kT)^\top e(kT)$$

with $T = 1$ ms. Therefore, the cost function is a measure for the tracking accuracy of the stirring movement. We consider as kernel candidate the squared exponential kernel, such that only the hyperparameters σ_n, φ are optimized. Table III shows the comparison between the data-based and the closed-loop optimization. In the data-based case, the hyperparameters are optimized based on a gradient method to minimize the log likelihood function (in this case, BO is used to minimize the tracking error in the closed-loop optimization. The initial values of the hyperparameters are set to the values of the data-based optimization. The bounds are defined as the 0.5 and 2 times of the initial values. The evolution of the minimum cost over the trials, where each trial is a single stirring movement, is shown in Fig. 5.

TABLE III
COMPARISON: DATA-BASED AND CLOSED-LOOP OPTIMIZATION

Value	Data-based	Closed-loop
σ_n	[0.10, 3e-3, 6e-4]	[0.20, 4e-3, 3e-4]
$\varphi_{1,2,3}$	[3.49, 1.42, 2.87]	[2.61, 1.68, 5.70]
$\varphi_{4,5,6}$	[1.21, 0.25, 0.27]	[0.80, 0.27, 0.29]
Log. likelihood	[89, -121, -176]	[115, -113, -136]
Cost (Tracking error)	1.49	1.05




Fig. 5. Minimum of the cost function over the number of trials.




Fig. 6. Comparison of the root square position error of all joints.

The comparison of the joint position error for the data-based and closed-loop optimization is shown in Fig. 6.

3) *Discussion:* After 100 trials, the tracking error is decreased by 30% through the optimization of the Gaussian process model only. Even if the resulting hyperparameters are sub-optimal with respect to the likelihood function, see Table III, the performance of the closed-loop is significantly improved. In comparison to collecting more training data to improve the model, the proposed method does not increase the computational burden of the Gaussian process prediction which is often critical in real-time applications. Since only the model is adapted, the properties of the closed-loop control architecture are also preserved.

CONCLUSION

In this paper, we present a framework for the model selection for kernel-based models to directly optimize the overall closed-loop control performance. For this purpose, the kernel and its hyperparameters are optimized using Bayesian optimization with respect to a cost function that evaluates the performance of the closed-loop. It is shown that this approach allows to preserve the control architecture properties as only the model is adapted. Simulations and hardware experiments demonstrate the advantages of the proposed approach to data-based model selection techniques.

ACKNOWLEDGMENTS

The research has received funding from the ERC Starting Grant “Con-humo” n°337654 and BaCaTec grant 9-[2018/1].

REFERENCES

- [1] C. M. Bishop *et al.*, *Pattern recognition and machine learning*, vol. 4. Springer New York, 2006.
- [2] P. Abbeel, M. Quigley, and A. Y. Ng, “Using inaccurate models in reinforcement learning,” in *International Conference on Machine Learning*, 2006.
- [3] M. Gevers, “Identification for control: From the early achievements to the revival of experiment design,” *European journal of control*, vol. 11, pp. 1–18, 2005.
- [4] H. Hjalmarsson, M. Gevers, and F. De Bruyne, “For model-based control design, closed-loop identification gives better performance,” *Automatica*, vol. 32, no. 12, pp. 1659–1673, 1996.
- [5] L. Ljung, “System identification,” in *Signal analysis and prediction*, pp. 163–173, Springer, 1998.

- [6] G. Chowdhary, J. How, and H. Kingravi, “Model reference adaptive control using nonparametric adaptive elements,” in *Conference on Guidance Navigation and Control*, 2012.
- [7] J. Umlauf, T. Beckers, M. Kimmel, and S. Hirche, “Feedback linearization using Gaussian processes,” in *IEEE Conference on Decision and Control*, 2017.
- [8] T. Beckers, J. Umlauf, and S. Hirche, “Stable model-based control with Gaussian process regression for robot manipulators,” in *IFAC World Congress*, 2017.
- [9] K. S. Narendra and K. Parthasarathy, “Identification and control of dynamical systems using neural networks,” *IEEE Transactions on neural networks*, vol. 1, no. 1, pp. 4–27, 1990.
- [10] S. Bansal, A. K. Akametalu, F. J. Jiang, F. Laine, and C. J. Tomlin, “Learning quadrotor dynamics using neural network for flight control,” in *IEEE Conference on Decision and Control*, 2016.
- [11] K. J. Åström and B. Wittenmark, *Adaptive control*. Courier Corporation, 2013.
- [12] D. Clarke, P. Kanjilal, and C. Mohtadi, “A generalized LQG approach to self-tuning control part i. aspects of design,” *International Journal of Control*, vol. 41, no. 6, pp. 1509–1523, 1985.
- [13] D. A. Bristow, M. Tharayil, and A. G. Alleyne, “A survey of iterative learning control,” *IEEE control systems magazine*, vol. 26, no. 3, pp. 96–114, 2006.
- [14] F. L. Lewis and D. Liu, *Reinforcement learning and approximate dynamic programming for feedback control*, vol. 17. John Wiley & Sons, 2013.
- [15] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, “Bayesian optimization for learning gaits under uncertainty,” *Annals of Mathematics and Artificial Intelligence*, vol. 76, no. 1, pp. 5–23, 2015.
- [16] B. Recht, “A tour of reinforcement learning: The view from continuous control,” *Annual Review of Control, Robotics, and Autonomous Systems*, 2018.
- [17] S. Bansal, R. Calandra, T. Xiao, S. Levine, and C. J. Tomlin, “Goal-driven dynamics learning via Bayesian optimization,” in *IEEE Conference on Decision and Control*, 2018.
- [18] J. A. Suykens, J. Vandewalle, and B. De Moor, “Optimal control by least squares support vector machines,” *Neural networks*, vol. 14, no. 1, pp. 23–35, 2001.
- [19] F. Berkenkamp, R. Moriconi, A. P. Schoellig, and A. Krause, “Safe learning of regions of attraction for uncertain, nonlinear systems with Gaussian processes,” in *IEEE Conference on Decision and Control*, 2016.
- [20] C. E. Rasmussen and C. K. Williams, *Gaussian processes for machine learning*, vol. 1. MIT press Cambridge, 2006.
- [21] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, “Taking the human out of the loop: A review of Bayesian optimization,” *IEEE*, vol. 104, no. 1, pp. 148–175, 2016.
- [22] J. Mockus, *Bayesian approach to global optimization: theory and applications*, vol. 37. Springer Science & Business Media, 2012.
- [23] A. D. Bull, “Convergence rates of efficient global optimization algorithms,” *Journal of Machine Learning Research*, vol. 12, no. Oct, pp. 2879–2904, 2011.
- [24] T. Beckers and S. Hirche, “Equilibrium distributions and stability analysis of Gaussian process state space models,” in *IEEE Conference on Decision and Control*, 2016.
- [25] E. C. Garrido-Merchán and D. Hernández-Lobato, “Dealing with integer-valued variables in Bayesian optimization with Gaussian processes,” *arXiv preprint arXiv:1706.03673*, 2017.
- [26] T. Beckers, D. Kulić, and S. Hirche, “Stable Gaussian process based tracking control of Euler-Lagrange systems,” *Automatica*, no. 103, pp. 390–397, 2019.
- [27] I. Steinwart and A. Christmann, *Support vector machines*. Springer Science & Business Media, 2008.
- [28] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, “Information-theoretic regret bounds for Gaussian process optimization in the bandit setting,” *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.
- [29] A. Isidori, *Nonlinear control systems*. Springer Science & Business Media, 2013.
- [30] V. Kecman, *Learning and soft computing: support vector machines, neural networks, and fuzzy logic models*. MIT press, 2001.
- [31] D. Nguyen-Tuong, M. Seeger, and J. Peters, “Computed torque control with nonparametric regression models,” in *American Control Conference*, 2008.