Proceedings of Machine Learning Research vol 144:1–13, 2021

Safe Bayesian Optimisation for Controller Design by Utilising the Parameter Space Approach

Lorenz Dörschel* L.DOERSCHEL@IRT.RWTH-AACHEN.DE
David Stenger* D.STENGER@IRT.RWTH-AACHEN.DE
Dirk Abel D.ABEL@IRT.RWTH-AACHEN.DE
Institut für Regelungstechnik, RWTH Aachen University, Germany

Abstract

As control systems become more and more complex, the optimal tuning of control parameters using Bayesian Optimisation gained an increased interest of research in recent years. Safe Bayesian Optimisation, tries to prevent sampling of unsafe parametrizations and therefore allow parameter tuning in real world experiments. Usually this is achieved by approximating a safe set using probabilistic Gaussian Process Regression predictions. In contrast in this work, analytical knowledge about robustly stable parameter configurations is gained by the parameter space approach and then incorporated within the optimisation as constraint. Simulation results on linear and non-linear systems with uncertain parameters show a significant performance gain compared to standard approaches.

Keywords: Optimisation of controller parameters, Bayesian Optimisation, parameter space approach, performance index ITAE

1. Introduction

In the past decades multiple efficient control algorithms have been proposed, which are able to control complex systems. Nevertheless, controllers are not plug-and-play but have to be tuned properly to fully exploit their possibilities. Hence, the task of a control engineer has often become to invest a lot of resources in time consuming tuning of these control parameters. To support engineers in this regard optimisation techniques have been proposed, which optimise the control parameters with respect to arbitrarily chosen objective functions using simulation of the closed loop behaviour.

The **controller optimisation problem** is challenging for several reasons: Firstly, we have a black-box optimisation problem. A closed analytical form of the objective function is not available and information can only be obtained through sampling of the experiment or simulation with different parameter combinations. Secondly, simulations and experiments might be time consuming. Thirdly, the objective function might be corrupted by noise and fourthly, sampling in unsafe areas of the controller parameter space must be prevented for real world system. These challenges rule out optimisation algorithms relying on gradients or relaxations. Although they have been used in this regard, meta-heuristics such as Genetic Algorithms or Particle Swarm Optimisation are considered to be not sample-efficient since they discard some previously obtained evaluations.

Bayesian Optimisation (BO) is a common method used for sample-efficient noisy optimisation. It can be attributed to the *Micro Data RL* branch of Reinforcement Learning, see Chatzilygeroudis et al. (2020). Surrogate models of black-box responses are learned by fitting fast-to-evaluate probabilistic regression models using all data obtained through previous sampling of the black-box

* These authors contributed equally and are listed in alphabetical order.

during optimisation. These surrogate models are used to find the next promising sample point. Recent examples from the control and robotics community include learning gaits under uncertainty in Calandra et al. (2016), local linear dynamics learning in Bansal et al. (2017), MPC tuning in Stenger et al. (2020) and tuning for a linear quadratic regulator in Marco et al. (2016).

In **Safe Bayesian Optimisation** the goal is to prevent the algorithm from sampling in unsafe controller parameter regions and therefore allow safe experimental parameter tuning. The algorithm SafeOpt originally proposed by Sui et al. (2015) approximates the safe controller parameter set by leveraging the probabilistic GPR-predictions. This method was applied by Khosravi et al. (2019) to a heat pump. Recently, constrained Bayesian Optimisation was compared with SafeOpt for safety critical real world controller tuning by König et al. (2020). Although, in the paper it was reported that no unsafe evaluations were conducted, sampling in unsafe regions is still possible for the combination of expected improvement and the probability of feasibility (see, e.g., Stenger et al. (2020)). It was also shown that with hyperparameter optimisation, SafeOpt resulted in frequent unsafe samples. This is in agreement, with Berkenkamp et al. (2020) as incorporation of domain knowledge by manual setting of bounds on the GP-hyperparameters was reported to be necessary.

With this contribution we seek to achieve safe controller tuning by incorporating analytical safety boundaries obtained from classical control engineering results. For linear control systems, the area of all **stabilising controller parameters** can be calculated symbolically, while also uncertain parameters can be included within the calculation scheme. First approaches can be traced back to Vyshnegradsky (1876) and nowadays several algorithms are available, which fulfil the same task but have important differences. Famous criteria like the Routh-Hurwitz criterion lead to expensive calculations if conducted analytically with multiple unknown parameters, as roots of polynomials of high order have to be found. This typically occurs if controllers (with parameters of choice) have to be robustly tuned for a system with parametric uncertainties. More efficient formulations are delivered by methods based on the classical **parameter space approach (PSA)** introduced by Ackermann et al. (2002), which leads to a decomposition of the parameter space into several regions, such that the number of unstable poles is identical for every point in each single region. The classical parameter space approach has certain drawbacks, as it is only applicable for a restricted system class and the use of frequency based techniques often leads to frequency sweeping, which makes a discretization of the parameter space necessary. The most recent and most efficient formulation is a parameter space approach based on a stability boundary mapping incorporating Lyapunov stability formulation introduced by the authors in Schrödel et al. (2015). It has been shown, e.g. in Pyta et al. (2017), that already for systems with medium complexity this approach outperforms other formulation with respect to accuracy and computational time.

Within the framework of the parameter space approach, not only stability but also performance criteria can be incorporated directly. Here, a variety of different methods is available, where the method of choice depends on the performance criterion to evaluate. Methods reach from frequency domain requirements in Bünte (2000), Odenthal and Blue (2000) to positioning of the eigenvalues of the closed loop systems in Siljak (1964). Those methods also have been extended for the Lyapunov based formulation in Voßwinkel et al. (2019) but still the performance mapping only covers performance criteria, which are typical for linear systems. More complex performance requirements like the ITAE-criterion, which are also used in practice, cannot be addressed by means of any parameter space technique nowadays.

The main idea of the paper is to design a new safe controller tuning algorithm for uncertain systems and arbitrary performance criteria by combining Bayesian Optimisation and the parameter

space approach. Both itself cannot solve this problem: The parameter space approach can find stable parameter regions and deal with uncertainties but is restricted to certain performance criteria. Bayesian Optimisation on the other hand can deal with objective functions, which can be noisy (uncertain) and also highly non-linear. However Bayesian Optimisation cannot reliably achieve safe sampling if it is not carefully tuned for the respective application. Hence, the structure of the paper is as follows. First, the fundamentals of Bayesian Optimisation are recapped in Section 2. The same follows for the parameter space approach techniques in Section 3. Section 4 explains and presents the new algorithm, whereas the developed algorithm is compared with approaches from the literature on an exemplary system in Section 5. A conclusion summarises the paper.

2. Fundamentals of Bayesian Optimisation

The aim of this section is to explain the fundamentals of noisy Bayesian Optimisation. For a detailed introduction to Bayesian Optimisation, we refer to Shahriari et al. (2016). For simplicity here, we consider a noisy minimisation problem with box constraints:

$$\min_{x} \quad \mathbb{E}\left[y(x)\right] \quad s.t. \quad x_{\min} \leq x \leq x_{\max}, \quad x \in \mathbb{R}^m \tag{1}$$

The decision variable of the problem is a m-dimensional parameter vector x with box constraints $x_{\min}$ and $x_{\max}$. It is assumed that information about the unknown cost function can only be obtained through noisy sampling of an expensive-to-evaluate black-box simulation or experiment. Each time the black-box is evaluated with parameters x_i, a corresponding noisy sample y_i is obtained at iteration i. Algorithm 1 shows the procedure of Bayesian Optimisation. The main idea is to use all samples obtained so far ($\mathcal{X}_i = [x_1, \ldots, x_i], \mathcal{Y}_i = [y_1, \ldots, y_i]$) to construct a fast-to-evaluate surrogate model of $y(x)$ at each iteration (Line 3) and use that model to search for the next promising sample point (Line 4). This way the surrogate model is iteratively refined in promising regions.

Algorithm 1 Bayesian Optimisation

1: Initial sampling of $\mathcal{X}_1$ and $\mathcal{Y}_1$:
2: **for** i = 1; 2; . . . ; **do**
3: update probabilistic GPR surrogate models using $\mathcal{X}_i$ and $\mathcal{Y}_i$
4: select x_i by optimising an acquisition function: $x_i = \arg\max_x(\alpha(x))$
5: query objective function with parameters x_i to obtain response y_i
6: augment data: $\mathcal{X}_{i+1} = \{\mathcal{X}_i, x_i\}, \mathcal{Y}_{i+1} = \{\mathcal{Y}_i, y_i\}$
7: **end for**

In this work Gaussian Process Regression (GPR) is used as the surrogate model. For a detailed introduction to GPR the reader is referred to Rasmussen and Williams (2006). GPR is a non-parametric regression and interpolation model which provides a probabilistic prediction of the objective function for parametrisations which have not been evaluated yet. In this work the model is defined by a constant mean function, a squared exponential kernel with automated relevance detection and a homoscedastic Gaussian observation model (to account for the noisy samples). In this case GPR yields normally distributed predictions $\tilde{y}$ with mean $\bar{y}$ and standard deviation σ_y:

$$\tilde{y}(\boldsymbol{x}) \sim \mathcal{N}(\bar{y}(\boldsymbol{x}|\mathcal{X}_i, \mathcal{Y}_i, \theta_i), \sigma_y^2(\boldsymbol{x}|\mathcal{X}_i, \mathcal{Y}_i, \theta_i)) \tag{2}$$

Hyperparameters θ_i are optimised at each iteration by maximising the marginal log-likelihood. Hyperpriors are placed on the hyperparameters to avoid potential over fitting. Based on the GPR predictions, an acquisition function $\alpha(\boldsymbol{x}) = f(\tilde{y}(\boldsymbol{x}))$ is used to balance between exploitation and exploration when searching for the next sample point (Line 4). Popular acquisition functions are for example *Probability of Improvement*, *Expected Improvement (EI)* and *Entropy search*. The *SafeOpt* algorithm can also be seen as a special acquisition function. After evaluation of the new sample point (Line 5), the data set is augmented (Line 6).

3. Fundamentals of the Parameter Space Approach

In Bayesian Optimisation for controller tuning, information about stable and unstable configurations would be helpful to avoid sampling of unstable parameter combinations. An analytical solution of calculating all stabilising parameter configurations for a linear system is provided by the parameter space approach, where the formulation from Voßwinkel et al. (2019) will be used in the following.

Consider the linear state space description with constant parameters k and p

$$\dot{z} = A(k,p)\, z, \quad A \in \mathbb{R}^{n\times n}, \quad z \in \mathbb{R}^n, \quad k \in \mathbb{R}^{l_k}\,,\; p \in \mathbb{R}^{l_p}. \tag{3}$$

The system (3) may either describe the open control loop with p being unknown or uncertain parameters or the closed loop system with the control parameters k or a combination of both. Applying the well known Lyapunov equation to (3) we get

$$A^*(k,p)P(k,p) + P(k,p)A(k,p) = -Q, \tag{4}$$

with $A^*(k,p)$ being the Hermitian transpose of $A(k,p)$. The connection between this equation and the stability of (3) is given by the fact, that the system is asymptotically stable, if for every positive definite matrix $Q \in \mathbb{R}^{n\times n}$ there exists a unique positive definite Hermitian matrix $P(k,p)$, which fulfills (4). Thus, the set of all stabilizing parameters is given by those parameters k and p, for which (4) can be solved with such P. Equation (4) can be rearranged in vector form

$$\left(I \otimes A^*(k,p) + A^T(k,p) \otimes I\right) \operatorname{vec}(P(k,p)) = -\operatorname{vec}(Q), \tag{5}$$

with I being the $n \times n$ identity matrix and $\otimes$ is representing the Kronecker product. The operator $\operatorname{vec}(\cdot)$ transforms matrices into column vectors by rearranging column-wise. The matrix P in (5) can be computed with the linear set of equations $\operatorname{vec}(P(k,p)) = -M^{-1}\operatorname{vec}(Q)$, with

$$M(k,p) = I \otimes A^*(k,p) + A^T(k,p) \otimes I \in \mathbb{R}^{n^2\times n^2}. \tag{6}$$

Thus the following theorem holds (Schrödel et al., 2015):

Theorem 1 *Let $M(k,p) = I \otimes A^*(k,p) + A^T(k,p) \otimes I$. If the matrix $A(k,p)$ is at its stability boundary, then the determinant of $M(k,p)$ equals 0 or ∞.*

Note, that this condition is necessary but not sufficient. Also pairs of eigenvalues of A which are point reflective about the origin (e. g. 5 and -5) cause the determinant of $M(k,p)$ to become zero. Nevertheless, this does not limit the applicability of Theorem 1, as those boundaries clearly will lie outside of stable regions. Thus, no conservatism is introduced within this method and the final obtained set is necessary and sufficient for system stability. For improving calculation

times, the computation can be simplified by removing doubled elements in M using elimination and duplication matrices from Magnus and Neudecker (1980), such that the dimension of $M(k,p)$ can be reduced from $n^2 \times n^2$ to $(n(n+1)/2) \times (n(n+1)/2)$. The parameter space approach only gives exact results for linear systems. Nevertheless necessary conditions can be given for non-linear system by investigating the linearization of these systems. Further note, that only parametric uncertainties can be incorporated within this framework.

4. Safe Bayesian Optimisation with the Parameter Space Approach (BO-PSA)

We now seek to optimise controller parameters k w.r.t. some arbitrary non-linear performance metric $y(k,p)$ of the closed loop behaviour of the system described in (3). The systems unknown parameters p are assumed to be distributed according to some unknown distribution with known bounded support: $\Pr(p_{\max} < p) = \Pr(p < p_{\min}) = 0$. Whenever the closed-loop evaluation of parameter k is performed some realisation of the parameter p is drawn at random from the unknown distribution. The expectation of the performance metric considering the distribution of unknown parameters (approximated by the GPs Gaussian likelihood model) is minimised:

$$\min_{k} \quad \mathbb{E}_p\left[y(k,p)\right] \quad s.t. \quad k_{\min} \leq k \leq k_{\max}, \quad k \in \mathbb{R}^{l_k} \tag{7}$$

It should be noted that the problem in (7) is closely related to robust and contextual BO. However, unlike in Toscano-Palmerin and Frazier (2018), Groot et al. (2010) and Fröhlich et al. (2020) here we assume the analytical form of $p(x)$ to be unknown. Additionally, we assume the realisation of p to be unknown before the experiment unlike in e.g. Krause and Ong (2011).

Additionally, we require that stability is maintained for the combination of all evaluated controller parametrizations k and possible realizations of p. With the parameter space approach, we can obtain analytical expressions, which are zero, if the system is on its stability boundary.

$$\det(M(k,p)) = \pm \prod_{j} c_j(k,p) = 0 \tag{8}$$

Here, $c_j(k,p)$ are polynomials in k and p. No non-polynomial parts arise in $M(k,p)$ as the closed loop system $A(k,p)$ has the form $A(k,p) = \hat{A}(p) – B(p) \cdot K(k)$ with the open loop system $\hat{A}(p)$. If $\hat{A}(p)$ is non-polynomial in p, one can define surrogate parameters $\hat{p}$, such that it becomes polynomial in $\hat{p}$. Furthermore the sign of each $c_j(k,p)$ can be defined such, that $c_j(k,p) > 0$ has to hold for stability. These calculations are exact if conducted symbolically and the arising expressions $c_j(k,p)$ can be evaluated fast in general. The idea is now, to include these conditions $c_j(k,p)$ in a straight forward manner to Bayesian Optimisation by treating them as constraints in the acquisition function optimisation step of BO. Since the realisation of the parameters p are unknown before the experiment, they have to be included in a worst-case fashion. Thus, instead of maximising the acquisition function over the complete domain (cf. Step 4 of Algorithm 1), the PSA constraint is included as follows in the acquisition function maximisation:

$$k_i = \arg\max_{k} \quad \alpha(k) \tag{9a}$$

$$s.t. \quad k_{\min} \leq k \leq k_{\max}, \quad k \in \mathbb{R}^{l_k} \tag{9b}$$

$$\min_{p,j}(c_j(k,p)) > 0 \tag{9c}$$

Here (9c) ensures stability as for each parametrization p, each of the polynomials j needs to be larger than zero, which corresponds to the necessary (and for linear system also sufficient) conditions of PSA. If infeasibility occurs, it is not possible to robustly stabilize the system. We further use the RI (cf. Forrester et al. (2006)) and EI (cf. Jones et al. (1998)) criterion for noisy and deterministic objective functions respectively.[1] However, the approach is not limited to any specific type of acquisition function. Therefore extensions to BO such as contextual BO or multi-objective BO can be included naturally. The acquisition function (9a) is non-convex. Therefore in this contribution, the optimisation problem (9) is solved using a hybrid approach of random search and MATLAB's fmincon. The starting point of the locally searching fmincon is taken as the best performing feasible random sampling point. For higher dimensional problems, meta-heuristics such as genetic algorithms can be used. Note that the set of parameters fulfilling (9c) is identical over all iterations and could therefore be calculated once at the start of the optimisation for a finite set of candidates.

5. Example

5.1. Problem description

To illustrate the presented methodology, the following linear system is investigated:

$$G(s) = \frac{(s+1)^2(-s+1)(s+p)}{(s+5)^2(s+0.1)^2} \tag{10}$$

Here p is a single unknown parameter, which influences the position of the systems zero and the steady state gain as well. This system shall be controlled with a classical PI-controller $K(s) = K_P + \frac{K_I}{s}$, where $k = \begin{bmatrix} K_P & K_I \end{bmatrix}$ are parameters of choice with $k_{\min} = \begin{bmatrix} -2 & 0 \end{bmatrix}$ and $k_{\max} = \begin{bmatrix} 1.5 & 10 \end{bmatrix}$. As performance criterion, the logarithm of the famous ITAE-criterion

$$y(k,p) = \log\left(\mathrm{ITAE}(k,p)\right) \quad \text{with} \quad \mathrm{ITAE}(k,p) = \int_0^\infty t \cdot \|e(t,k,p)\| \, \mathrm{d}t \tag{11}$$

with control error e and time t shall be minimised. Our approach is general and other performance metrics can be used with the algorithm. Because in practice, we cannot assume to perfectly know the system dynamics in advance, we assume the parameter p to be distributed according to a truncated normal distribution with the mean $\overline{p} = 0.6$, variance $\sigma_p = 1/15$ as well as support $p_{\min} = 0.4$ and $p_{\max} = 0.8$. Note, that no analytical solutions for minimising the ITAE-criterion are unknown in literature. The ITAE of a given controller parametrization and randomly drawn system parameter is evaluated using the ode45 solver of MATLAB of a single reference step simulated for $100\,\mathrm{sec}$.

5.2. Parameter Space Approach

Application of the parameter space approach to the system with respect to the parameters $k = \begin{bmatrix} K_P & K_I \end{bmatrix}$ and p delivers analytically stability conditions for (k,p), which have to be fulfilled. In the present case, these are lengthy polynomials of order 4, which are easy to evaluate but not given here explicitly. The arising stable regions within the parameter space are shown in Fig. 1. For better comprehensibility in a two-dimensional visualisation, the results are shown for equidistant

1. Note that the applicability of EI to the noisy case is limited. Therefore RI is used in the noisy case.

discretized $p \in [0.2, 1]$ with red lines, $\overline{p}$ in black and the robustly stabilising area in green. As clearly visible in Fig. 1, the region of stability is rather complicated and splits into two separated regions if p is high enough. This is especially challenging for the SafeOpt algorithm (Berkenkamp et al. (2016)), since it can only expand the safe set locally. Furthermore one can notice, that the region of stability shrinks with increased value of p, as an higher open loop gain and a shift of the phase gain through the systems zero towards higher frequency makes the system harder to stabilise.

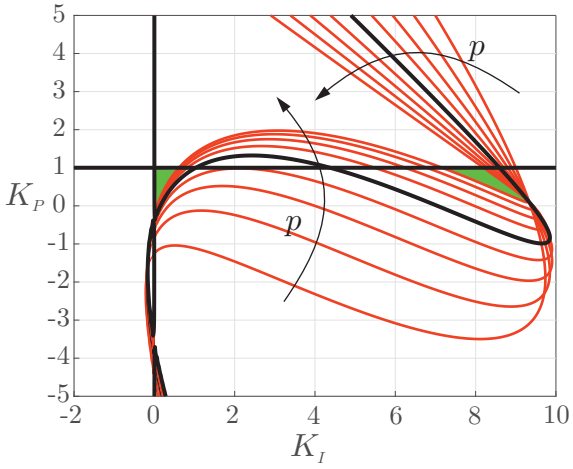

Figure 1: Stability boundaries for discretized p in red, $\overline{p} = 0.6$ in black and the robustly stabilizing area in green

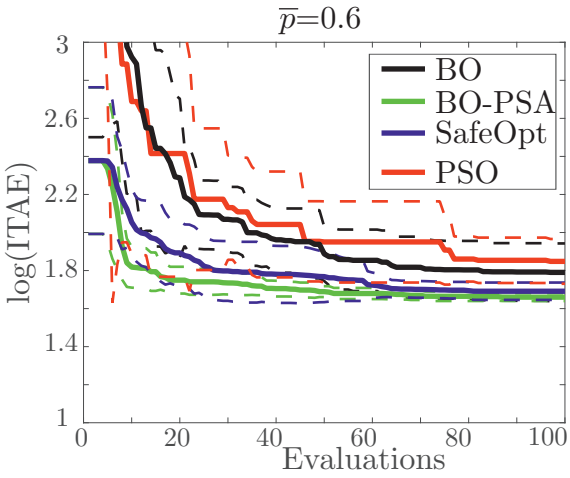

Figure 2: Mean and 1-sigma bound of the best observed performance as a function of the number of evaluations.

5.3. Benchmark Optimisation Algorithms

In order to assess the performance of the algorithm proposed in Section 4, it is compared with other approaches from the literature which have previously been used for optimisation-based PI controller tuning. In addition to standard BO without safety constraints, we use the SafeOpt algorithm proposed by Berkenkamp et al. (2016) and Particle Swarm Optimisation (PSO) without the incorporation of safety constraints. The later was used by Injeti and Divyavathi (2019) for optimisation of the ITAE-criterion. Note that **BO** and **BO-PSA** use identical settings for the GPR[2] model. The lower bounds on the kernel length scales are set such, that on a distance of 0.01 in controller parameter space the respective entry in the covariance matrix is at least 0.1. This resembles the introduction of minor prior knowledge. Additionally, the GP's prior mean is fixed to the mean of the observations. When using **SafeOpt**, we cannot directly use stability as a safety constraint. Therefore the safety boundary i.e. the maximum safe log-ITAE is set to $\log(0.5 \cdot 100^2)$, which corresponds to the log of the open loop ITAE. The safety margin is set to $b_w = 3$. Preliminary results with a GP-parametrization identical to BO and BO-PSA let to frequent unsafe evaluations. Therefore the upper bounds of the length scales are fixed such that on a distance of 0.5 in controller parameter space, the respective entry in the covariance matrix is at maximum 0.1. The constant GP mean function is set to be equal to the safety bound. With these choices prior domain knowledge is brought in the

2. The GPML Toolbox (Rasmussen and Nickisch (2010)) is used to construct the GPR models.

optimisation although it may not be available in the general case. The choice of the mean function means that unless we obtained a stable sample close by, we expect the parametrization to be unsafe with probability of 50%. We limit the extrapolation capabilities of the GP by setting an upper bound for the length scale resulting in a rather conservative setup of the algorithm. In contrast to standard BO a discretization of the parameter space is needed. Here we use a discretization of 0.005 for K_P and K_I, respectively. For **PSO**, the swarm size is set to 10 and all other parameters are kept to MATLAB 2017b default values. It should be noted, that PSO may achieve better results, if its hyperparameters are tuned more carefully, however this is not feasible for real world experiments.

5.4. Optimisation on nominal model

First, the algorithms are applied to two specific realisations of the uncertain parameter: $p = 0.4$ and $p = 0.6$. Each of the algorithm is ran 10 times for 100 evaluations. All BO-based algorithms use four initial random samples with the initial samples for BO-PSA and SafeOpt being identically and taken at random from the stable region of the controller parameter space. Figures 2 shows the convergence behaviour of the presented algorithms whereas Table 1 summarises the quantitative performance criteria. Figure 4 show the optimal parametrizations obtained with BO-PSA for both cases. It can be seen that BO-PSA outperforms the benchmark algorithms in convergence speed, final performance and number of unsafe parameter evaluations. Even in the case where unsafe evaluations are not problematic for example during optimisation on simulations only, standard BO as well as PSO are unable to produce competitive parametrisations after 100 evaluations. In comparison to BO-PSA, SafeOpt converges slower and less consistently. Secondly, although prior domain knowledge is introduced in SafeOpt, it consistently results in a few unsafe evaluations. The results are statistically significant according to the Wilcoxon Rank Sum Test.

Table 1: Performance characteristics of all approaches for p = 0.6 (p = 0.4)

Approach:	BO-PSA	SafeOpt	BO	PSO
Final log(ITAE)	1.660 ± 0.020	1.691 ± 0.045	1.790 ± 0.152	1.848 ± 0.117
	(1.380 ± 0.029)	(1.391 ± 0.017)	(1.447 ± 0.037)	(1.483 ± 0.058)
% unstable evals.	0 ± 0	4 ± 2	53.1 ± 12.8	13.5 ± 0.9
	(0 ± 0)	(3.1 ± 1.3)	(39.7 ± 3.9)	(66.5 ± 3.72)
Calc. time in sec	236.7 ± 28.7	473.0 ± 10.5	222.7 ± 50.0	241.0 ± 87.0
	(221.8 ± 25.3)	(465.1 ± 33.6)	(124.0 ± 2.9)	(14.8 ± 3.9)

Figure 3 showcases the problems of using standard BO and SafeOpt in combination with the ITAE-criterion. We can see, that the ITAE-criterion has a discontinuity at the stability limit at $K_p = 1$, whereas at the stability limit for small K_P the slope remains moderate. In Figure 3 (left) we see the effect of two samples being evaluated at either side of the discontinuity. The discontinuity contradicts the smoothness assumptions imposed by the squared exponential kernel. Therefore the hyper-parameter optimisation of the GP leads towards extremely small length scales resulting in poor extrapolation capabilities. Figure 3 - middle shows the GP-prediction, if no upper bound on the length scale is set. It can be observed that the estimated safe set is substantially larger than the actual safe set. If an upper bound is imposed on the length scale (cf. Figure 3 - right) as described in Section 5.3, the estimated safe set is much more conservative. However naturally the discontinuity at $K_P = 1$ still cannot be predicted from the ITAE criterion alone.

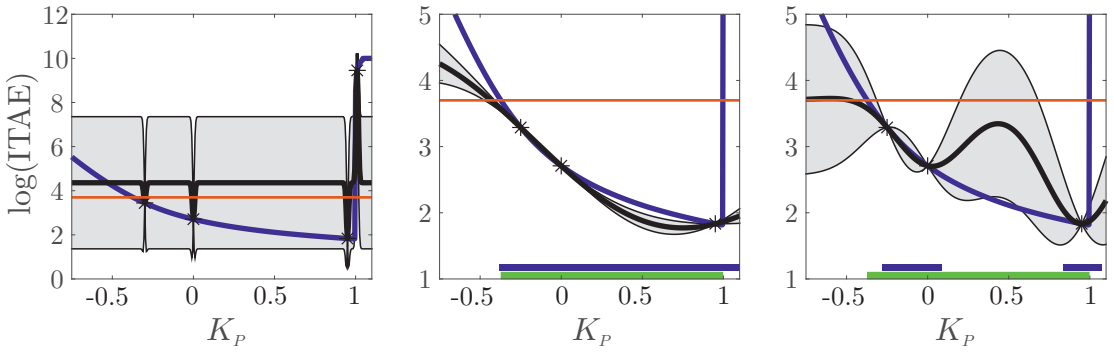

Figure 3: **Left:** Standard BO **Center:** SafeOpt **Right:** SafeOpt with an upper bound on the Kernel length scale. The ground truth (blue), evaluated samples (x), GPR-Predictions with mean (black) and confidence interval (grey), the safety boundary (red) as well as the actual safe set (green) and the estimated safe set (blue) are shown for $K_i = 8$ and $p = 0.6$.

5.5. Optimisation on noisy model

Often times in real world experiments, the system behaviour might differ from one evaluation to an other due to changing operating conditions. Here, we show that the proposed algorithm can also be used in this noisy setting (see Section 5.1). In Figure 5, the validation performance evaluated on 50 draws of p is shown. We can observe that BO-PSA out performs SafeOpt as well as standard BO, with the difference between SafeOpt and BO-PSA being larger than in the nominal case (cf. Fig. 2). The red line in Fig. 5 shows the validation performance of the parametrizations obtained through optimising on the nominal model. Optimising on the nominal model is not sufficient. This can be explained by the optimum parametrizations for the nominal model ($p = 0.6$) not being located within the robustly stable region (cf. Fig. 4), where the optimal parameters are marked with stars, the stability boundaries for the specific value for p as lines in the same colour and the robustly stable area in green. It can be seen, that many of the nominal optimal values are not robustly stable.

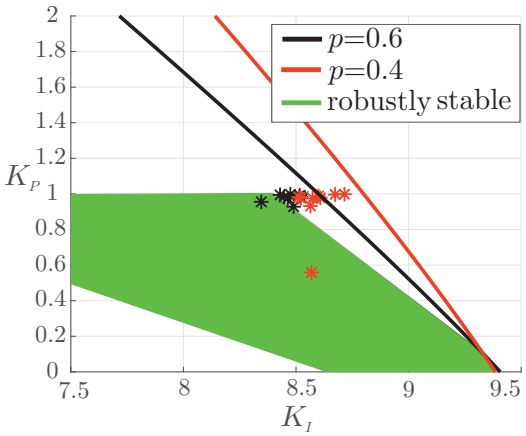

Figure 4: Optimal parameters obtained by BO-PSA for $p = 0.4$ and $p = 0.6$

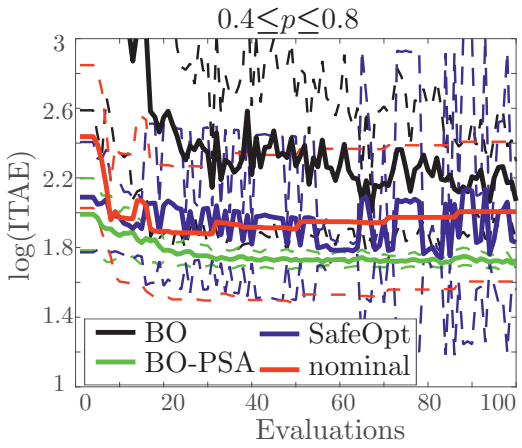

Figure 5: Validation performance evaluated on 50 random draws of p

5.6. Performance in the non-linear case

Sufficient conditions for stability for the non-linear case are not available in general and may depend on the initial position. However, the stability of the linearised system around the target operating point is a necessary condition for the stability of the non-linear system. As a result BO-PSA approach can reduce the parameter search space even in the non-linear case by including that necessary condition. This is demonstrated on a standard cart-Pole system (c.f. Fig. 6). The feedback gains of a state feedback controller are optimised with BO-PSA with the control objectives of balancing the pole and track a cart reference trajectory. Fig. 7 shows the convergence of standard BO, Safe-Opt and BO-PSA[3]. The percentage of parameter combinations leading to an angle greater than $90°$ were 20.6%, 6.7% and 4.6%, respectively. The number of unsafe samples with BO-PSA can be further reduced to 0.5% by introducing uncertainty in the pole length.

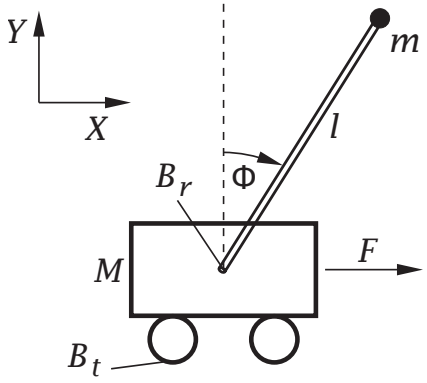

Figure 6: Schematic of the non-linear cart-pole system with friction.

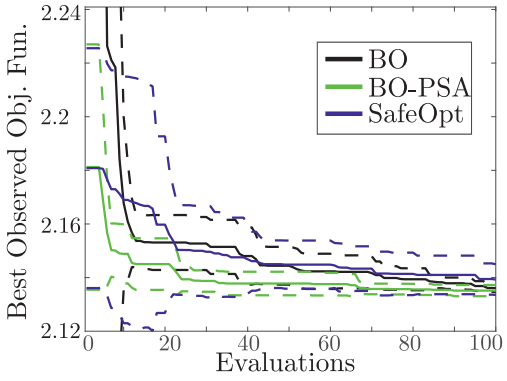

Figure 7: Best observed performance for the non-linear cart-pole example.

6. Conclusion

In this paper a novel algorithm for safe Bayesian Optimisation for tuning controller parameters of uncertain dynamical systems has been presented. Stability criteria from the parameter space approach were included to noisy Bayesian Optimisation. The approach was compared with approaches from the literature on linear and non-linear systems. While in the linear case, the SafeOpt algorithm was not able to prevent sampling in the unstable controller parameter region the BO-PSA approach guaranteed only stable sampling, while achieving faster convergence and better final results. Similar was observed for a non-linear example. Hence, the paper is an example on how incorporating control engineering domain knowledge in data driven techniques can be beneficial.

Acknowledgments

David Stenger gratefully acknowledges the support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation – GRK 1856).

3. The same BO settings as in Section 5.4 were used. Identical hyperpriors are used for all approaches. They were not tuned for the specific application. For SafeOpt a safety constraint is set to the maximum observed pole angle of 45 degrees.

References

J. Ackermann, P. Blue, T. Bünte, L. Guevenc, D. Kaesbauer, M. Kordt, M. Muhler, and D. Odenthal. *Robust Control: The Parameter Space Approach*. Communications and Control Engineering. Springer London, 2002.

Somil Bansal, Roberto Calandra, Ted Xiao, Sergey Levine, and Claire J. Tomiin. Goal-driven dynamics learning via bayesian optimization. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 5168–5173. IEEE, 2017. ISBN 978-1-5090-2873-3. doi: 10.1109/CDC.2017.8264425.

Felix Berkenkamp, Angela P. Schoellig, and Andreas Krause. Safe controller optimization for quadrotors with gaussian processes. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 491–496. IEEE, 2016. ISBN 978-1-4673-8026-3. doi: 10.1109/ICRA.2016.7487170.

Felix Berkenkamp, Andreas Krause, and Angela P. Schoellig. Bayesian optimization with safety constraints: Safe and automatic parameter tuning in robotics, 2020. URL https://arxiv.org/pdf/1602.04450.pdf.

T. Bünte. Mapping of nyquist/popov theta-stability margins into parameter space. In *Proc. of 3rd IFAC Symposium on Robust Control Design*, pages 241–246, 2000.

Roberto Calandra, André Seyfarth, Jan Peters, and Marc Peter Deisenroth. Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1-2): 5–23, 2016. ISSN 1012-2443. doi: 10.1007/s10472-015-9463-9.

Konstantinos Chatzilygeroudis, Vassilis Vassiliades, Freek Stulp, Sylvain Calinon, and Jean-Baptiste Mouret. A survey on policy search algorithms for learning robot controllers in a handful of trials: Ieee transactions on robotics, 1-20. 2020. doi: 10.1109/TRO.2019.2958211.

A. Forrester, A. Keane, and N. Bressloff. Design and analysis of noisy computer experiments. *AIAA Journal*, 44(10):2331–2339, 2006. ISSN 0001-1452. doi: 10.2514/1.20068.

Lukas Fröhlich, Edgar Klenske, Julia Vinogradska, Christian Daniel, and Melanie Zeilinger. Noisy-input entropy search for efficient robust bayesian optimization. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2262–2272. PMLR, 26–28 Aug 2020. URL http://proceedings.mlr.press/v108/frohlich20a.html.

P.C. Groot, A. Birlutiu, and T. Heskes. Bayesian monte carlo for the global optimization of expensive functions. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010), 16-19 August 2010, Lisbon, Portugal*, pages 249–254, Netherlands, 2010. IOS Press. ISBN 978-1-60750-605-8.

Satish Kumar Injeti and M. Divyavathi. Optimal gain scheduling of pid controller for the speed control of pmsm drive using bio-inspired optimization algorithms. *International Journal on Electrical Engineering and Informatics*, 11(2):308–325, 2019.

D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.

M. Khosravi, A. Eichler, N. Schmid, R. S. Smith, and P. Heer. Controller tuning by bayesian optimization an application to a heat pump. In *2019 18th European Control Conference (ECC)*, pages 1467–1472, 2019. doi: 10.23919/ECC.2019.8795801.

Christopher König, Mohammad Khosravi, Markus Maier, Roy S. Smith, Alisa Rupenyan, and John Lygeros. Safety-aware cascade controller tuning using constrained bayesian optimization, 2020. URL https://arxiv.org/pdf/2010.15211.pdf.

Andreas Krause and Cheng Ong. Contextual gaussian process bandit optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL https://proceedings.neurips.cc/paper/2011/file/f3f1b7fc5a8779a9e618e1f23a7b7860-Paper.pdf.

J. R. Magnus and H. Neudecker. The elimination matrix: Some lemmas and applications. *SIAM Journal on Algebraic Discrete Methods*, 1(4):422–449, 1980.

Alonso Marco, Philipp Hennig, Jeannette Bohg, Stefan Schaal, and Sebastian Trimpe. Automatic lqr tuning based on gaussian process global optimization. pages 270–277, 2016. doi: 10.1109/ICRA.2016.7487144. URL http://arxiv.org/pdf/1605.01950v1.

D. Odenthal and P. Blue. Mapping of frequency response performance specifications into parameter space. volume 33, pages 531 – 536, 2000.

L. Pyta, I. Mutlu, F. Schrödel, N. Bajcinca, M.T. Söylemez, and D. Abel. Stability boundary mapping for discrete time linear systems. In *2017 20th IFAC World Congress*, pages 15060–15065, Jul 2017.

Carl Edward Rasmussen and Hannes Nickisch. Gaussian processes for machine learning (gpml) toolbox. *J. Mach. Learn. Res.*, 11:3011–3015, 2010. ISSN 1532-4435.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2006. ISBN 026218253X.

F. Schrödel, E. Almodaresi, A. Stump, N. Bajcinca, and D. Abel. Lyapunov stability bounds in the controller parameter space. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 4632–4636, Dec 2015.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175, 2016. ISSN 0018-9219. doi: 10.1109/JPROC.2015.2494218.

D. D. Siljak. Analysis and synthesis of feedback control systems in the parameter plane i-linear continuous systems. *IEEE Transactions on Applications and Industry*, 83(75):449–458, Nov 1964.

David Stenger, Muzaffer Ay, and Dirk Abel. Robust parametrization of a model predictive controller for a cnc machining center using bayesian optimization. *IFAC-PapersOnLine*, 53(2):10388–10394, 2020. ISSN 2405-8963. doi: https://doi.org/10.1016/j.ifacol.2020.12.2778. URL https://www.sciencedirect.com/science/article/pii/S2405896320335412. 21th IFAC World Congress.

Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 997–1005, Lille, France, 07–09 Jul 2015. PMLR. URL http://proceedings.mlr.press/v37/sui15.html.

Saul Toscano-Palmerin and Peter I. Frazier. Bayesian optimization with expensive integrands, 2018.

Rick Voßwinkel, Lorenz Pyta, Frank Schrödel, İlhan Mutlu, Dinu Mihailescu-Stoica, and Naim Bajcinca. Performance boundary mapping for continuous and discrete time linear systems. *Automatica*, 107:272 – 280, 2019. ISSN 0005-1098. doi: https://doi.org/10.1016/j.automatica.2019.05.055. URL http://www.sciencedirect.com/science/article/pii/S0005109819302833.

I. Vyshnegradsky. Sur la theorie generale des regulateurs. *CR Acad. Sci. Paris*, 83:318–321, 1876.