

Robust Parametrization of a Model Predictive Controller for a CNC Machining Center Using Bayesian Optimization ^{*}

David Stenger ^{*} Muzaffer Ay ^{*} Dirk Abel ^{*}

^{*} *Institute of Automatic Control, RWTH Aachen University, Aachen,
Germany (e-mail: {d.stenger,m.ay,d.abel}@irt.rwth-aachen.de).*

Abstract: Control algorithms such as model predictive control (MPC) and state estimators rely on a number of different parameters. The performance of the closed loop usually depends on the correct setting of these parameters. Tuning is often done manually by experts based on a simulation model of the system. Two problems arise with this procedure. Firstly, experts need to be skilled and still may not be able to find the optimal parametrization. Secondly, the performance of the simulation model might not be able to be carried over to the real world application due to model inaccuracies within the simulation. With this contribution, we demonstrate on an industrial milling process how Bayesian optimization can automate the tuning process and help to solve the mentioned problems. Robust parametrization is ensured by perturbing the simulation with arbitrarily distributed model plant mismatches. The objective is to minimize the expected integral reference tracking error, guaranteeing acceptable worst case behavior while maintaining real-time capability. These verbal requirements are translated into a constrained stochastic mixed-integer black-box optimization problem. A two stage min-max-type Bayesian optimization procedure is developed and compared to benchmark algorithms in a simulation study of a CNC machining center. It is showcased how the empirical performance model obtained through Bayesian optimization can be used to analyze and visualize the results. Results indicate superior performance over the case where only the nominal model is used for controller synthesis. The optimized parametrization improves the initial hand-tuned parametrization notably.

Keywords: Constrained Bayesian optimization, Outlier detection, Noisy optimization, Model Predictive Control, Automatic parameter tuning, Milling, CNC machining center

1. INTRODUCTION

A major goal of controller synthesis is to achieve best possible closed-loop performance. Besides choosing a suitable controller and possibly state estimator, their correct parametrization can heavily improve closed-loop performance. Analytical design rules are only available for a limited number of controllers and scenarios. Manual tuning of controller parameters is tedious and often suboptimal. Mathematical optimization can lead to a more systematic approach and overall better performance.

The problem of finding the optimal configuration of a controller can in the most simple case be written as:

$$\theta_c^{\text{opt}} = \arg \min_{\theta_c \in \mathcal{D}} C(\theta_c) \quad (1)$$

The value of the loss function to be minimized $C(\theta_c)$ is obtained for one set of controller parameters through simulation or experiment. This results in three challenges for the optimizer used to find the solution to Eq. 1:

Firstly, we have a black-box optimization problem. Secondly, simulations and experiments might be time consuming requiring a sample efficient optimization algorithm. Thirdly, the objective function might be corrupted by noise $\tilde{C}(\theta_c) = C(\theta_c) + \epsilon$. These challenges rule out a number of common optimization algorithms. Algorithms relying on gradients or relaxations of the problem are not suitable. Meta-heuristics such as Genetic Algorithms or Particle Swarm Optimization are considered to be not sample efficient since they discard some of the previously obtained simulation results. Additionally, they do not incorporate stochastic constraints and objective functions inherently.

Bayesian optimization is a common method used for noisy optimization. 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 during optimization. These surrogate models are used to find the next promising sample point.

Bayesian optimization is a tool widely used in algorithmic tuning across different domains such as tuning of optimization algorithms [Hutter et al. (2011)] and Fault detection and isolation (FDI) [Marzat et al. (2011)]. Recent examples from the control and robotics community include learning gaits under uncertainty [Calandra et al. (2016)],

^{*} The presented research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2023 Internet of Production - 390621612.

throttle valve control [Neumann-Brosig et al. (2019)], local linear dynamics learning [Bansal et al. (2017)] and tuning for a linear quadratic regulator for robots [Marco et al. (2016)]. A related line of research is safe Bayesian optimization which tries to prevent the algorithm from sampling in unsafe controller parameter regions and therefore allows experimental parameter tuning [Berkenkamp et al. (2016)]. This method was recently applied by Khosravi et al. (2019).

In this contribution we apply Bayesian optimization to an industrial CNC machining process with machinery sensitive to improper actuation. We cannot expect this process to be modeled accurately with low computational cost a-priori¹. Therefore the focus of this contribution is to find a robust parameterization for the combination of a Kalman Filter (KF) and a MPC with respect to random model plant mismatches in the simulation. Goal is to maximize expected performance while ensuring safe worst case behavior.

Parameter tuning for MPC with Bayesian optimization was recently examined by Piga et al. (2019) and Lucchini et al. (2020). However, the authors do not consider safe worst case behavior or optimization of the state estimator. Andersson et al. (2016) optimize tuning parameters of soft constraints for stochastic collision avoidance with MPC taking into account probabilistic safety constraints.

Ensuring minimal worst case behavior has been previously addressed in the FDI community by Marzat et al. (2011). Ghosh et al. (2018) used Bayesian optimization to search for adversarial examples for controllers. Krause and Ong (2011) considered finding optimal parameters for different application-contexts using Bayesian optimization.

Our main contributions in the context of parameter tuning for MPC with Bayesian optimization are to

- ensure acceptable worst case behavior by explicitly considering model uncertainty while
- simultaneously optimizing the MPC and KF parametrization and
- expand previous work on outlier detection by adding a classification step to prohibit resampling in an area where outliers often occur.

The presented method can be used to find a safe parametrization to initiate experimental manual tuning, tuning with safe Bayesian optimization [Berkenkamp et al. (2016)] or other online learning methods.

The paper is organized as follows: In Section 2, the application system is described in detail and the optimization problem is stated. In Section 3, Bayesian optimization including design choices made for the problem at hand is introduced. Section 4 shows how it is used in a two stage framework to solve the problem described in Section 2. In Section 5, the presented approach is evaluated.

2. PROBLEM STATEMENT

Milling is a fast and flexible machining process, which is highly acclaimed in production due to its productivity.

¹ Accurate modeling is possible via dextral simulations. Computational costs prohibit their usage in controller synthesis.

During milling, a rotating tool is moved against a workpiece to cut material along a predefined trajectory. Thus, a desired geometry can be manufactured. The presented optimization approach is examined for the quality control during milling. The quality of the production is defined by the deviation of the manufactured geometry from the desired geometry. This deviation relates to the deflection of the working tool during the process. Hence, a reproducible quality for the milling process requires the control of the cutting force, which leads to tool deflection. The force in turn relates to the feed velocity of the tool. Therefore, a multi-stage approach is applied, where on the outer loop the trajectory for the feed velocity is optimized with respect to the resulting force and on the inner loop the feed velocity of the tool is controlled. The tool dynamic is described as

$$\ddot{v}(t) + 2D\omega_0\dot{v}(t) + \omega_0^2v(t) = K\omega_0^2u(t - t_d), \quad (2)$$

where $v(t)$ stands for the tool velocity at time t in dependence of the control input u after the delay-time t_d and the model parameters K for gain, D for damping and ω_0 for resonance frequency.

In order to avoid high cutting forces before occurrence, a model predictive control strategy is applied for the quality control during milling. In addition, a Kalman-Filter (KF) is used for state-estimation during the process. The reader may refer to previous work by Stemmler et al. (2019) for further details about the structure of the control approach.

The performance of the overall control strategy depends in part on the parametrization of the MPC and KF. Namely for the MPC, the prediction horizon H_p , control horizon H_u are relevant parameters. Additionally, the ratio $\lambda_{\text{MPC}} \in \mathbb{R}$ between the weighting matrices $\mathbf{Q} = \mathbf{I}$ and $\mathbf{R} = \mathbf{I}10^{\lambda_{\text{MPC}}}$ in the MPC cost function,

$$J = \|\mathbf{e}\|_{\mathbf{Q}} + \|\Delta\mathbf{u}\|_{\mathbf{R}} \quad (3)$$

are considered within optimization. This way, the deviation from the desired trajectory \mathbf{e} and the change of the control input $\Delta\mathbf{u}$ can be weighted differently during optimization. Regarding the KF, the covariance matrix for measurement noise is predefined as $\mathbf{R}_{\text{KF}} = \mathbf{I} \cdot 0.001$ which corresponds to an experimentally determined variance of the sensor. In order to weight between prediction and measurement at correction, the covariance matrix for process noise is set as $\mathbf{Q}_{\text{KF}} = \mathbf{I}10^{\lambda_{\text{KF}}}$ by the ratio $\lambda_{\text{KF}} \in \mathbb{R}$.

Satisfactory tracking performance for the underlying control of the feed velocity was achieved with the hand-tuned default parametrization ($H_u = 15$, $H_p = 15$, $\lambda_{\text{MPC}} = -3$, $\lambda_{\text{KF}} = -1$) on the nominal model. In order to increase the robustness of the parameterization, sources of uncertainty are introduced in the simulation model. Measurement signals are perturbed with zero mean Gaussian noise. Additionally, a model plant mismatch is introduced by modifying the plant models stiffness $\tilde{\omega}_0 = \omega_0 \cdot \theta_{e,1}$ and damping $\tilde{D} = D \cdot \theta_{e,2}$. Note that while the simulated process model is disturbed, the internal MPC model is kept constant. From now on we refer to certain realizations of $\theta_e = \{\theta_{e,1}, \theta_{e,2}\}$ as context. Note that in general other types of environmental conditions can be included in the

contextual variables such as different reference trajectories or operating modes.

Goal now is to find a robust set of parameters $\theta_c = \{H_u, H_p, \lambda_{\text{MPC}}, \lambda_{\text{KF}}\}$ with respect to a known probability distribution of θ_e . The inputs and outputs of the simulation model used during optimization are shown in Fig. 1. Note that even in a fixed context, the simulation is probabilistic because the realization of measurement noise is drawn at random.




Fig. 1. Overview of the simulation environment

The optimization problem is stated in Eq. 5. The objective function (Eq. 5a) is the expected (with respect to the assumed probability distribution of θ_e) value of the integral tracking error, $C_{\text{EITE}}(\theta_c)$. For one simulation it is calculated by comparing the actual velocity v_k at timestep k with the reference velocity $v_{\text{ref},k}$. Here it is assumed that θ_e is distributed according to a truncated normal distribution with no correlation in between components of θ_e as follows:

$$\theta_e \in \mathbb{R}^2 \sim \psi(\mu_e, \sigma_e^2 \mathbf{I}, \theta_e^{\min}, \theta_e^{\max}) \quad (4)$$

where $\psi(\mu_e, \sigma_e^2 \mathbf{I}, \theta_e^{\min}, \theta_e^{\max})$ is a normal distribution with $p(\theta_e < \theta_e^{\min}) = p(\theta_e > \theta_e^{\max}) = 0$.

Note however that the presented approach is in principle applicable for all probability distributions with bounded support. In order to prevent damage to the milling head the maximal overshoot in the worst case context $\Delta h^*(\theta_c)$ is constrained (Eq. 5e). This constraint is not explicitly considered in the original MPC formulation². In addition, to be able to use cost efficient hardware the maximal computation time $T(\theta_e|\theta_c)$ is limited (Eq. 5d). With a probability of $\Phi(z=3) = 0.96$ it is not allowed to exceed a critical value of T^{\max} .

$$\min_{\theta_c} C_{\text{EITE}}(\theta_c) = \mathbb{E} \left[\sum_{k=0}^N (v_{\text{ref},k} - v_k(\theta_e|\theta_c))^2 \right] \quad (5a)$$

$$s. t. \quad \theta_c \in \mathbb{Z}^2 \times \mathbb{R}^2 \quad (5b)$$

$$\theta_c^{\min} \leq \theta_c \leq \theta_c^{\max} \quad (5c)$$

$$p(T(\theta_e|\theta_c) < T^{\max}) > \Phi(z=3) = 0.96 \quad (5d)$$

$$\max_{\theta_e} (\Delta h(\theta_e|\theta_c)) = \Delta h^*(\theta_c) < \Delta h^{\max} \quad (5e)$$

It should be noted that the optimization methodology presented is not limited to the constraints and objective

² This constraint can be integrated in the MPC formulation. This might increase calculation time and its fulfillment cannot be guaranteed if the MPC model is inaccurate as it is in the presented case.

function chosen here. Arbitrary constraints or performance matrices can be used. Preferably those which cannot be explicitly considered in the controller.

3. NOISY CONSTRAINED BAYESIAN OPTIMIZATION WITH OUTLIER DETECTION

The aim of this Section is to explain the fundamentals of noisy Bayesian optimization with outlier detection as well as highlight the design choices made for the problem at hand. Two instances of the algorithm described in this section are used in a hierarchical approach as will be described in Section 4. For notational simplicity, we now consider a simplified version of the optimization problem stated in Eq. 5:

$$\min_x \mathbb{E}[y(x)] \quad (6a)$$

$$s. t. \quad \mathbf{x} \in \mathbb{R}^m \quad (6b)$$

$$p(g(x) < g^{\max}) > \Phi(z) \quad (6c)$$

The objective function $y(x)$ as well as the constrained black-box response $g(x)$ are corrupted by noise. Each time the Black-Box is evaluated with parameters x_i , corresponding noisy samples y_i and g_i are obtained.

Algorithm 1 Bayesian optimization with outlier detection

- 1: Initial sampling of X_1, Y_1 and G_1 ;
 - 2: **for** $k = 1; 2; \dots$; **do**
 - 3: update probabilistic surrogate models using $\tilde{X}_{k+1}, \tilde{Y}_{k+1}$ and \tilde{G}_{k+1}
 - 4: select x_{k+1} by optimizing an acquisition function: $x_{k+1} = \arg \max_x (\alpha(x))$
 - 5: query objective function to obtain y_{k+1} and g_{k+1}
 - 6: augment data $X_{k+1} = \{X_k, x_{k+1}\}$, $Y_{k+1} = \{Y_k, y_{k+1}\}$, $G_{k+1} = \{G_k, g_{k+1}\}$,
 - 7: $\tilde{X}_{k+1}, \tilde{Y}_{k+1}, \tilde{G}_{k+1} \leftarrow \text{OutDetect}(X_{k+1}, Y_{k+1}, G_{k+1})$
 - 8: **end for**
-

Algorithm 1 shows the procedure of Bayesian optimization. A detailed introduction to Bayesian optimization is provided for example by Shahriari et al. (2016). The main idea is to use all sample points obtained so far ($X_k = [x_1, \dots, x_k], Y_k = [y_1, \dots, y_k], G_k = [g_1, \dots, g_k]$) to construct fast-to-evaluate surrogate models of $y(x)$ and $g(x)$ at each iteration (Line 3) and use these models to search for the next promising sample point. This way the surrogate models are iteratively refined in promising regions.

In this work Gaussian process regression (GPR) models are used as surrogate models. A separate GPR model is built for each of the responses. For a detailed introduction on GPR the reader is referred to Rasmussen and Williams (2006). GPR is a nonparametric regression and interpolation model which provides a probabilistic prediction of each of the black-box responses for parametrization which have not been evaluated yet. The model is here defined by a constant a-priori mean, a squared exponential kernel with automated relevant detection and a homoscedastic Gaussian observation model (to account for the noisy samples). Hyperparameters are optimized at each iteration

by maximizing the likelihood. Hyper-priors are placed on the hyperparameters to include expert knowledge in the optimization and avoid potential over fitting. E.g. the lower bound on the kernel length scales in the direction of prediction and control horizon are set to $l_{H_n}^{\min} = l_{H_p}^{\min} = 0.22$ which roughly speaking corresponds to a covariance of the objective function values of at least 0.1 if the horizons are changed by one.

Based on these models, an acquisition function $\alpha(x)$ is used to balance between exploitation (sampling close to the current optimum) and exploration (sampling where the model is uncertain) when searching for the next sample point (Line 4). Here the so called reinterpolation procedure (RI) proposed by Forrester et al. (2006) is used. Using RI prohibits multiple evaluations of the objective function for identical parameters by fitting an intermediate interpolating surrogate model. This is beneficial in the presented case since repeating the second stage of the optimization (cf. Section 4.2) twice for identical parameters would be a waste of computational resources. In addition to maximize the noisy expected improvement of the objective function by maximizing $RI(x_q)$, we want our query point to be feasible $p_{\text{feas}}(x_q)$, to not be an outlier $p_{\text{out}}(x_q)$ with off-the-charts objective function value and we want the simulation not to fail $p_{\text{fail}}(x_q)$. Therefore, the following acquisition function is used:

$$\alpha(x_q) = RI(x_q)p_{\text{feas}}(x_q)(1 - p_{\text{out}}(x_q))(1 - p_{\text{fail}}(x_q)) \quad (7)$$

The probability of feasibility is calculated by using the probabilistic predictions of $g(x)$. At each iteration, the next sample point is chosen by maximizing $\alpha(x_q)$ using particle swarm optimization.

After evaluation of the new sample point (Line 5), outlier detection is performed (Line 7). In Bayesian optimization with GPR, special care needs to be taken of outliers. At the problem at hand, some parametrization may lead to very large objective function values. This is problematic because outliers can severely corrupt the probabilistic surrogate models, leading to unrealistically small length scales. In order to detect outliers we follow the approach presented by Martinez-Cantin et al. (2018). In this paper outliers are detected by first fitting a robust GPR model with a student-t observation model and then detecting the observations with a low likelihood.

In the original paper, outliers are discarded from the set of observations. We take this approach one step further by building a k-nearest-neighbor (knn) classifier³ to estimate the probability $p_{\text{out}}(x_q)$ of observing an outlier at a given location x_q . The training data set of the knn-classifier consists of all sample points within the controller design parameter space labeled according to whether they were identified as outliers or not.

For some parameterizations the control algorithm to be configured may lead the simulation environment to crash due to numerical issues. To prevent repetitive sampling in these areas an additional knn-classifier is built to estimate

³ A squared exponential kernel as distance metric and a squared inverse distance weighting is used .

the probability $p_{\text{fail}}(x_q)$ of a parametrization x_q to lead to a crash.

4. TWO STAGE BAYESIAN OPTIMIZATION FOR MPC AND KF TUNING

The algorithm presented in Section 3 is applied within a two stage optimization approach to solve the problem stated in Eq. 5. An overview is given in Fig. 2. Stage one has the task to find the optimal controller configuration characterized by low expected integral tracking error, acceptable worst case overshoot and real time capable execution time. For each controller configuration, θ'_c , queried by Stage one, the overshoot belonging to the worst possible combination of environmental variables, $\Delta h^*(\theta'_c) = \max_{\theta_e}(\Delta h(\theta_e|\theta'_c))$, is searched for in Stage two. This is achieved by solving a second stage optimization problem and ensures that constraint 5e is satisfied. This procedure is similar to the one presented by Marzat et al. (2011) in the context of FDI.




Fig. 2. Overview of the two stage optimization approach

4.1 Stage one: Optimize controller parameters

The algorithm presented in Section 3 is used within Stage one. Although $p(\theta_e)$ is assumed to be a truncated normal distribution, we cannot make any assumptions about the true structure of the probability density function of a response (i.e. integral tracking error, return time, and maximum overshoot) for a given controller parametrization θ_c because the simulation model is non-linear. By using GPR with Gaussian likelihood (observation model) we approximate the true unknown probability distribution of the responses, by a deterministic function corrupted by Gaussian noise. This way in total 3 GPR models are built - one for each of the relevant black-box responses: $C_{\text{EITE}}(\theta_c)$, $T(\theta_c)$ and $\Delta h^*(\theta_c)$. Latin hypercube sampling is used as initial sampling. As the objective function (Step 5 of Algorithm 1) of stage one, stage two of the optimization procedure is called.

4.2 Stage two: Find worst possible context

The goal of Stage two is to find the maximum overshoot for a given controller parametrization θ'_c queried by Stage one. Note that using RI in the first stage prohibits multiple evaluations of the second stage with identical controller parameters. Noisy Bayesian optimization is used in this Stage, as well. But in this case only box constraints need to be considered. Outlier detection is also not used. Therefore, only one GPR model to approximate $\Delta h(\theta_e|\theta'_c)$ is built and standard RI is used as the acquisition function. Initial sampling is performed by drawing 5 different contexts from the truncated normal distribution of environmental variables. The corresponding 5 different integral

tracking errors and return times of the initial sampling are returned to Stage one. The overshoot is maximized in the successive steps. Stage two is terminated either if an overshoot exceeding the maximum overshoot allowed was observed or if 10 evaluations were performed. Note that the GPR models built in both stages are completely independent from one another. This follows the approach by Marzat et al. (2011). An alternative approach would be to use a joint model using the controller parameter as well as the environmental parameters as explanatory variables for the overshoot as for example explained by Krause and Ong (2011).

5. RESULTS

The algorithm presented in Sections 3 and 4 is applied to the optimization problem stated in Section 2 and was implemented in MATLAB 2017b using the GPML toolbox [Rasmussen and Nickisch (2010)] to create the GPR surrogate models. Additionally, a simplified benchmark test case (Section 5.2) is used to compare the developed approach with three different benchmark approaches (Section 5.1). Parameters are optimized with respect to a single reference velocity trajectory. This is reasonable because in a real world milling application we can expect to know the desired trajectory in advance. However, the presented approach is not limited to use only one single reference trajectory during optimization.

5.1 Benchmark algorithms

The presented algorithm (Algo. IV) is compared with three benchmark algorithms (Algo. I-III):

- I Bayesian optimization on nominal model
- II Random sampling in Stage one
- III Bayesian optimization without outlier detection
- IV Bayesian optimization with outlier detection

The goal is to evaluate the performance impact of the individual proposed optimization steps. Algo. I is used for parameter optimization on the nominal model whereas Algo. (II-IV) are run on randomly varying realizations of the perturbed model using the second Stage presented in Section 4.2 to determine feasibility w.r.t. constraint 5e. Algo. II uses random sampling in Stage one instead of finding the next sample point by maximizing the acquisition function presented in Eq. 7. Algo. III is identical to the presented approach except that the outlier detection and classification scheme is not used.

5.2 Benchmark test case

In order to compare the algorithms quantitatively each algorithm is run 10 times for 3 hours (only one hour for (Algo. I)) respectively on a simplified benchmark test case⁴. The benchmark test case considers only two controller parameters H_u and λ_{MPC} . The prediction horizon H_p is set to the controller horizon H_u and λ_{KF} is kept at its hand-tuned default value. In each of the 10 runs, initial sampling was kept identical for each algorithm. The

⁴ All experiments were performed on a desktop PC with an AMD Ryzen 7 1700 Eight-Core Processor @3GHZ and 16 GB Ram

validation performance of a given parametrization is estimated by running the simulation with 25 different random draws from the distributions of model uncertainties. Table 1 shows the average quality of the final solution of each of the algorithms. Fig. 3 depicts the average validation objective function of the feasible best solution after a given number of simulations during optimization.

Table 1. Quality of the final solution averaged over 10 runs (benchmark test case)

Algorithm:	I	II	III	IV
Feasibility	50%	100%	90%	100%
Obj. fun. validation	0.147	0.182	0.148	0.155
Δ Obj. fun. val. - train	+0.042	+0.001	+0.003	+0.007




Fig. 3. Validation performance after a given number of iterations considering only feasible solutions averaged over 10 runs (benchmark test case)

It becomes apparent that only optimizing the MPC on the nominal model (Algo. I.) is not sufficient. Although it requires the algorithm only ~ 25 simulations to converge, half of the time, the final solution is infeasible due to too much overshoot (Constraint 5e) when model uncertainty is incorporated during validation. Additionally, the expected integral tracking error is underestimated substantially.

In contrast, Algos. II-IV are able to consistently find solutions which are feasible during validation and only slightly underestimate the validation integral tracking error. Random Sampling (Algo. II) is not able to find a competitive solution. Although the outlier detection and classification scheme used in Algo. IV improves the initial convergence (until around 200 simulations), the average performance of the final solution is slightly better in Algo. III.

Preliminary experiments have shown that the surrogate model of the objective function generated with outlier detection is more plausible than without outlier detection⁵. Without outlier detection, negative integral tracking errors are predicted in some regions and the surrogate models are far less smooth and in general uncertainty is larger.

Furthermore, it was observed that with outlier detection, the hyper parameter optimization favored larger length scales which is consistent with the smoother predictions

⁵ Plots are not shown here due to space limitations.

and smaller uncertainty. We hypothesize that the larger length scales hinder exploration in the later stages of optimization due to lower model uncertainty. Therefore interestingly, the more realistic fit obtained by GPR with outlier detection does not automatically mean that optimization performance is increased.

The experiments conducted by Martinez-Cantin et al. (2018) have shown that outlier detection consistently improves the optimization procedure. This somewhat contradicts our findings. One reason is that in comparison to Martinez-Cantin et al. (2018) we have *deterministic* outliers, which consistently occur in one area of the design parameter space instead of *random* outliers.




Fig. 4. Final surrogate model of the benchmark test case with estimated feasible area and objective function.

Fig. 4 depicts the final estimated behavior of the closed loop within the controller design parameter space. In general the results are comprehensible. Calculation time is solely a function of H_u . Due to the presented outlier classification scheme we can observe that deterministic outliers (very large integral tracking errors) appear in areas with small H_u and large λ_{MPC} . In contrast, the worst case overshoot constraint is violated for small values of λ_{MPC} (i.e. low penalty on change of inputs).

Interestingly, the optimum has a smaller H_u than the calculation time constraint would allow. At first glance this may be counterintuitive. But given enough model uncertainty a longer prediction horizon does increase the chance of misprediction and therefore of critical overshoot. The default parametrization is located in the infeasible region. It was therefore considerably improved by the optimization approach.

5.3 Full test case

In the full test case all four parameters are considered for optimization $\theta_c = \{H_u, H_p, \lambda_{MPC}, \lambda_{KF}\}$. Algos. III & IV were run 4 times for 6 hours each.

The best and average validation performance as well as the average feasibility of the final solution are shown in Table 2. Results indicate that the proposed algorithms can find a competitive solution even in a higher dimensional controller parameter space. Yet performances are less consistent than in the benchmark test case. Additionally, we can observe that the default parametrization of $\lambda_{KF} = -1$ and $H_p = H_u$ could not be improved upon.

Table 2. Validation obj. fun. of the final solution averaged over 4 runs (full test case)

	Best	Average	Feasibility
Algo. III	0.143	0.151	75%
Algo. IV	0.141	0.174	75%

Temporarily, an alternative parametrization was considered to be the best feasible solution during the course of optimization. This parametrization is shown with some of its neighboring parameterizations in Table 3. Similar parameterizations were found in 50% of the runs of Algos. III and IV.

Table 3. Parametrization with the best validation performance and neighboring parameterizations.

H_u	H_p	λ_{MPC}	λ_{KF}	Obj. Fun.	Feasibility
1	4	-3.9	2	0.12	✓
1	3	-3.9	2	0.19	✗
1	5	-3.9	2	0.22	✓
1	4	-3.9	-1	0.23	✓
2	4	-3.9	2	0.14	✗

Although this parametrization shows superior validation performance in comparison to the best solution of the simplified test case, its neighboring solutions are either infeasible or the expected integral tracking error is unacceptably high. Therefore from a control engineering perspective this solution cannot be considered robust and would be rejected in practice.

Fortunately, during later stages of the optimization these parameterizations are discarded by the optimization algorithms, because of the worse performing or infeasible neighborhood. This can be explained by the fact that in Bayesian optimization with GPR, black-box responses are assumed to be smooth on the characteristic kernel length scales. By setting a minimal kernel length scale based on domain knowledge as done in the present work, the optimizer implicitly considers the neighborhood of the optimal solution. This can be seen as an additional advantage of Bayesian optimization over other optimizers such as genetic algorithms where only the best solution is considered regardless of its neighborhood. Additionally, this alternative parametrization highlights the non-convexity of the optimization problem and of the relevance of all considered parameters.

6. SUMMARY & CONCLUSION

In this contribution Bayesian optimization was used to simultaneously optimize hyperparameters of a MPC and KF for an industrial CNC machining process. In order to

achieve a robust parameterization, the simulation model was perturbed with model-plant mismatches randomly drawn from a known distribution with bounded support as well as randomly drawn measurement noise. Goal was to minimize expected integral tracking error, ensure worst case safety by constraining the maximum overshoot and enforce real-time capability by limiting the return time. The optimization problem is solved using a two-stage Bayesian optimization procedure relying on GPR with Gaussian observation model, the RI-acquisition function as well as outlier detection and classification. On a simplified benchmark test case, it was shown that optimization on the nominal model does not produce satisfactory parameter combinations. The default parametrization as well as random sampling was outperformed considerably. It was also observed that outlier detection did not consistently improve the convergence although surrogate models are more comprehensible. The empirical performance model obtained through Bayesian optimization allowed to analyze and visualize the results. Optimization on the full test case revealed the relevance of all parameters and the non-convexity of the optimization problem. Furthermore, it was shown how the model assumptions encoded in the hyperprior helps the optimization to avoid narrow and physically incomprehensible local minima.

The presented approach can help control engineers to find an *initial* robust and safe parametrization for controllers and state estimators given a closed loop simulation of the system and a probabilistic assumption over the model plant mismatch. It can empirically enforce constraints or performance metrics which are not or cannot explicitly be considered within the controller. This parametrization can then be further improved online for example by safe Bayesian optimization.

REFERENCES

- Andersson, O., Wzorek, M., Rudol, P., and Doherty, P. (2016). Model-predictive control with stochastic collision avoidance using bayesian policy optimization. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 4597–4604. IEEE. doi:10.1109/ICRA.2016.7487661.
- Bansal, S., Calandra, R., Xiao, T., Levine, S., and Tomiini, C.J. (2017). Goal-driven dynamics learning via bayesian optimization. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 5168–5173. IEEE. doi:10.1109/CDC.2017.8264425.
- Berkenkamp, F., Schoellig, A.P., and Krause, A. (2016). Safe controller optimization for quadrotors with gaussian processes. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 491–496. IEEE. doi:10.1109/ICRA.2016.7487170.
- Calandra, R., Seyfarth, A., Peters, J., and Deisenroth, M.P. (2016). Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1-2), 5–23. doi:10.1007/s10472-015-9463-9.
- Forrester, A., Keane, A., and Bressloff, N. (2006). Design and analysis of noisy computer experiments. *AIAA Journal*, 44(10), 2331–2339. doi:10.2514/1.20068.
- Ghosh, S., Berkenkamp, F., Ranade, G., Qadeer, S., and Kapoor, A. (2018). Verifying controllers against adversarial examples with bayesian optimization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 7306–7313. IEEE. doi:10.1109/ICRA.2018.8460635.
- Hutter, F., Hoos, H.H., and Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. In C.A.C. Coello (ed.), *Learning and Intelligent Optimization*, 507–523. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Khosravi, M., Eichler, A., Schmid, N., Smith, R.S., and Heer, P. (2019). Controller tuning by bayesian optimization an application to a heat pump. In *2019 18th European Control Conference (ECC)*, 1467–1472. IEEE. doi:10.23919/ECC.2019.8795801.
- Krause, A. and Ong, C.S. (2011). Contextual gaussian process bandit optimization. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 24*, 2447–2455. Curran Associates, Inc.
- Lucchini, A., Formentin, S., Corno, M., Piga, D., and Savaresi, S.M. (2020). Torque vectoring for high-performance electric vehicles: an efficient mpc calibration. *IEEE Control Systems Letters*, 1. doi:10.1109/LCSYS.2020.2981895.
- Marco, A., Hennig, P., Bohg, J., Schaal, S., and Trimpe, S. (2016). Automatic lqr tuning based on gaussian process global optimization. 270–277. doi:10.1109/ICRA.2016.7487144.
- Martinez-Cantin, R., Tee, K., and McCourt, M. (2018). Practical bayesian optimization in the presence of outliers. In A. Storkey and F. Perez-Cruz (eds.), *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, 1722–1731. PMLR, Playa Blanca, Lanzarote, Canary Islands.
- Marzat, J., Piet-Lahanier, H., and Walter, E. (2011). Min-max hyperparameter tuning, with application to fault detection. *18th IFAC World Congress*. doi:10.3182/20110828-6-IT-1002.00476.
- Neumann-Brosig, M., Marco, A., Schwarzmann, D., and Trimpe, S. (2019). Data-efficient auto-tuning with bayesian optimization: An industrial control study. *IEEE Transactions on Control Systems Technology*, 1–11. doi:10.1109/TCST.2018.2886159.
- Piga, D., Forgione, M., Formentin, S., and Bemporad, A. (2019). Performance-oriented model learning for data-driven mpc design. *IEEE Control Systems Letters*, 3(3), 577–582. doi:10.1109/LCSYS.2019.2913347.
- Rasmussen, C.E. and Nickisch, H. (2010). Gaussian processes for machine learning (gpml) toolbox. *J. Mach. Learn. Res.*, 11, 3011–3015.
- Rasmussen, C.E. and Williams, C.K.I. (2006). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., and de Freitas, N. (2016). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1), 148–175. doi:10.1109/JPROC.2015.2494218.
- Stemmler, S., Ay, M., Schwenzer, M., Abel, D., and Bergs, T. (2019). Model-based predictive force control in milling. In *2019 18th European Control Conference (ECC)*, 4313–4318. IEEE. doi:10.23919/ECC.2019.8795716.