

Adaptive Bayesian Optimization for High-Precision Motion Systems

Christopher König¹, Raamadaas Krishnadas¹, Efe C. Balta¹, and Alisa Rupenyan²

Abstract—Controller tuning and parameter optimization are crucial in system design to improve closed-loop system performance. Bayesian optimization has been established as an efficient model-free controller tuning and adaptation method. However, Bayesian optimization methods are computationally expensive and therefore difficult to use in real-time critical scenarios. In this work, we propose a real-time purely data-driven, model-free approach for adaptive control, by online tuning low-level controller parameters. We base our algorithm on GoOSE, an algorithm for safe and sample-efficient Bayesian optimization, for handling performance and stability criteria. We introduce multiple computational and algorithmic modifications for computational efficiency and parallelization of optimization steps. We further evaluate the algorithm's performance on a real precision-motion system utilized in semiconductor industry applications by modifying the payload and reference stepsize and comparing it to an interpolated constrained optimization-based baseline approach.

Note to Practitioners—This work is motivated by developing a comprehensive control and optimization framework for high-precision motion systems. Precision motion is an integral application of advanced mechatronics and a cornerstone technology for high-value industrial processes such as semiconductor manufacturing. The proposed method framework relies on data-driven optimization methods that can be designed by prescribing desired system performance. By using a method based on Bayesian Optimization and safe exploration, our method optimizes desired parameters based on the prescribed system performance. A key benefit is the incorporation of input and output constraints, which are satisfied throughout the optimization procedure. Therefore, the method is suitable for use in practical systems where safety or operational constraints are of concern. Our method includes a variable to incorporate contextual information, which we name the task parameter. Using this variable, users can input external changes, such as changing step sizes for a motion system, or changing weight on top of the motion system. We provide two parallel implementation variants of our framework to make it suitable for run-time operation under context changes and applicable for continuous operation in industrial systems. We demonstrate the optimization framework on simulated examples and experiment on an industrial motion system to showcase its applicability in practice.

I. INTRODUCTION

High-performance mechanical positioning systems used in high-yield industries such as semiconductor manufacturing rely on precision motion systems. Due to the level of precision

required in such applications, closed-loop control is necessary, especially in the context of run-time disturbance, mechanical nonlinearities, and context shifts in practice. Consequently, precision motion control has been a well-studied topic in mechatronics systems [1].

Adaptive control approaches present a compelling alternative to robust controllers, particularly in high-performance precision motion applications, where tracking performance is crucial. The investigation into addressing uncertainties inherent in dynamics and facilitating adaptation has been studied in classical adaptive control literature via methods such as Model Reference Adaptive Control (MRAC) [2], [3]. Gaussian processes (GP) have also found utility in the context of modeling the outputs of nonlinear systems within a dual controller paradigm [4], where the coupling of system states and inputs is encapsulated within the covariance function of the GP model. Learning system dynamics using a $\mathcal{L}_1$ -adaptive control framework has been additionally demonstrated [5], [6].

Rather than pursuing the modeling or learning of system dynamics, an alternate approach involves characterizing the system through its observed performance, directly extracted from measurement data. Subsequently, optimization of the low-level controller parameters can be undertaken to meet the desired performance criteria, as demonstrated in the context of motion systems [7]–[11]. However, the application of such model-free methodologies to continuous adaptive control remains limited, primarily due to challenges associated with maintaining stability and safety under the influence of disturbances and system uncertainties, along with the computational complexity of such methods.

Improving the performance of motion systems requires maintaining good positioning accuracy during motion while ensuring a short settling phase, which is possible by eliminating run-time vibrations and other disturbances. Common methods for motion control include gain scheduling, LPV control approaches, and tuning of low-level controller parameters. However, gain scheduling and LPV models rely on a static view of the system's performance, and are usually tuned to fixed step sizes. Similarly, parameter tuning, which is often implemented via automated routines, suffers from the same problem of lacking adaptivity and flexibility to changes in the system. In the recent literature, run-to-run adaptation of the control parameters in mechanical systems has been demonstrated using an efficient Bayesian optimization-based approach, which prevents instabilities while maximizing tracking performance [12]. However, the approach is limited to slow updates and simple systems, due to its associated computational load.

This project was funded by the Swiss Innovation Agency, grant Nr. 46716, and by the Swiss National Science Foundation under NCCR Automation, grant Nr. 180545.

¹ Control and Automation Group, Inspire AG, Zürich, Switzerland.

² ZHAW Centre for Artificial Intelligence, Zürich University for Applied Sciences, Zürich, Switzerland.

Corresponding author: E. C. Balta: efe.balta@inspire.ch

A. Contributions

This paper proposes an efficient data-driven algorithm for run-to-run control parameter adaptation, suitable for real-time applications, and showcases its efficacy for high-precision motion systems. We summarize our main contributions below.

- 1) We provide a safe and adaptive control algorithm for real-time applications based on parallelization of the optimization via multiprocessing. We accomplish this by separating the optimization into an active- and passive phase, between which the algorithm switches, based on event-triggered criteria.
- 2) Building on the previously introduced GOOSE algorithm for data-driven optimization, we present a computationally efficient implementation, by stripping the algorithm of its discretization of the domain.
- 3) We show the effectiveness of our approach on a real nanometer precision motion system, a crucial instrument in the semiconductor industry.

B. Related work and paper organization

Adaptive control parameter tuning for motion systems.

An adaptive control method with performance guarantees in the presence of model uncertainty is Model Reference Adaptive Control (MRAC) [13], which can be enhanced by modeling the uncertainties in the reference model, as demonstrated in [3] for a quadrotor system. Additionally, adaptive robust control methods for precision motion have been developed for dealing with parametric uncertainties [14]. However, such adaptive controllers require run-time computation and possible changes to the controller architecture, which may not be practical in certain high-speed applications.

Linear parameter varying gain scheduling methods to adapt to sudden changes in the system dynamics of motion systems have previously been designed and studied by [15] for unmanned surface vehicles. The repetitive nature of motion in mechatronics systems has been exploited in various iterative learning control-based approaches (ILC), for example, combined with mismatch modeling to reduce the tracking error [16], or with feed-forward compensation of acceleration, jerk, and snap terms as demonstrated in [17]. While ILC approaches provide an excellent error reduction for repetitive motion, they require additional development to deal with changing trajectories and disturbances. Depending on the system, this can be overcome by describing parametrically the reference and the plant's dynamics [18] and connecting this parametric description to the ILC update law.

BO for controller tuning. Bayesian Optimization (BO) has found applications in refining a variety of controller types designed for high-speed precision motion systems. BO-driven tuning aimed at performance enhancement has been effectively showcased in the context of cascade controllers governing linear axis drives [9], [19]. This approach involves leveraging data-derived performance metrics to deliberately enhance traversal time and tracking accuracy while reducing vibrations. In addition, models describing the noise in the system can be added to improve performance and safety [20].

In the realm of model predictive control (MPC), an alternative strategy emerges by selecting the prediction model to optimize closed-loop performance. Selecting the prediction model is in contrast to robust control design methods of designing controllers for worst-case scenarios. The controller optimization involves tuning parameters within the MPC optimization objective to achieve optimal performance [19]. While BO is efficient for finding good controller parameters before deploying a system in operation, it has been less explored for adaptive tuning during operation. The reasons for the lack of such studies are twofold. Firstly, it requires simplifying assumptions and prior knowledge on the nature of the expected changes in a system [21], and, secondly, it requires a mechanism for continuously reducing data from past time instances [12] which might reduce the predictive capabilities of the approach.

Safe BO. To maintain optimal closed-loop performance for systems subject to variations, controllers must dynamically adapt their parameters online, all while avoiding configurations that might lead to instability. The need for online adaptation is particularly pronounced for high-precision motion systems, where even a single iteration marked by excessive vibrations is deemed unacceptable during continuous operation. Consequently learning the effect of the cascaded controller parameters under safety and stability constraints is necessary to ensure that only safe parameters are explored [22]. One approach to address this challenge involves the utilization of the SafeOpt algorithm introduced in [23]. Notably, while this algorithm is effective in ensuring safety, its exploration strategy may be inefficient in practice [24]. To improve the computational and sample efficiency of SafeOpt, optimization-based search approaches have been recently demonstrated [25].

An alternative was explored in [26], where the safe set is not actively expanded. While the method in [26] might entail a trade-off with optimality, it aligns well with the considered application. Additionally, a distinct alternative for controller tuning involves the incorporation of safety-related log barriers within the cost function, as exemplified by the work presented in [11]. This innovation aims to curtail constraint violations, as compared to approaches grounded in Bayesian optimization with inequality constraints, such as the one outlined in [27].

GoOSE. An efficient strategy for facilitating safe exploration is introduced in the work of GoOSE [28], where the overarching objective is to ensure that all inputs introduced to the system adhere to an unknown yet observable constraint. In the context of controller tuning, GoOSE unifies time-varying Gaussian process bandit optimization [29], multi-task Gaussian processes [30], and an efficient safe set search mechanism grounded in particle swarm optimization [31]. The result is a cohesive framework that proves instrumental in the adaptive tuning of motion systems [12]. However, a noteworthy limitation emerges due to the discretization of the input space of the expanders. This limitation renders the algorithm impractical for scenarios characterized by high-dimensional input and task spaces, or instances involving substantial amounts of data within real-time online applications.

The rest of the paper is organized as follows. Section II presents the proposed method with implementation details.




Fig. 1: The BO-based run-to-run controller tuning setting considered in this work.

Section IV presents numerical and experimental studies to demonstrate the effectiveness of the proposed method. Section V provides concluding remarks and future directions.

II. PROBLEM SETTING AND PRELIMINARIES

Consider a closed-loop control system of a possibly nonlinear plant. Let $x_{\text{opt}} \in \mathcal{X}_{\text{opt}}$ be a vector of controller parameters, such as PID controller gains or feedforward gains in the closed loop. Furthermore, let $x_{\tau} \in \mathcal{X}_{x_{\tau}}$ be the context parameters of the system influencing the behavior of the system (e.g. temperature, payload, reference parameters). The complete input can be written as $x = [x_{\text{opt}}, x_{\tau}]^T \in \mathcal{X}$. The parameters x_{opt} should be continuously optimized online and in real-time, i.e., while the system is operating, such that the overall closed-loop performance of the system is optimized for the current context $x_{\tau,t}$, while stability is ensured at any time. This can be expressed by the constrained optimization problem:

$$\min_{x \in \mathcal{X}} \{f(x) \mid q(x) \leq c \ \& \ x_{\tau} = x_{\tau,t}\}, \quad (1)$$

under the feasibility assumption $\{x \mid q(x) \leq c\} \neq \emptyset$, for a given $c \in \mathbb{R}$ at any possible context $x_{\tau} \in \mathcal{X}_{x_{\tau}}$. The unknown cost function $f(x)$ and the constraint or safety metric $q(x)$ here are expensive-to-evaluate functions, extracted from experimental data, that should be learned on-the-fly.

The setting of the adaptive controller tuning framework presented in this paper is given in Fig. 1. Traditional BO approaches in the literature aim to optimize the controller parameters in an offline optimization step, which is then deployed for online usage. In this work, we continuously update the controller parameters in a run-to-run fashion using the measurement in each run with the goal of solving (1). To ensure desirable transient operation and safety constraints, we base our framework on the GOOSE method, presented below. The main elements of the algorithm enabling continuous safe optimization are built on GOOSE developed in [12] and multi-task optimization.

A. Probabilistic models

To reason under uncertainty about the unknown objective $f(x)$ and safety metrics $q(x)$, we model each as a gaussian process - GP^f and GP^q , characterised by kernel functions κ^f and κ^q as well as priors μ_0^f and μ_0^q , respectively. Note

that in the following we have omitted the f and q indices for ease of notation as the expressions are identical for both. Posterior GP mean and variance denoted by $\mu_t(\cdot)$ and $\sigma_t^2(\cdot)$, respectively, are computed based on the previous measurements $y_{1:t} = [y_1, \dots, y_t]^T$ and a given kernel $\kappa(\cdot, \cdot)$, where t is the number of observations :

$$\begin{aligned} \mu_t(x) &= \kappa_t(x)^T (K_t + \sigma_n^2 I)^{-1} (y_{1:t} - \mu_0(x)), \\ \sigma_t^2(x) &= \kappa(x, x) - \kappa_t(x)^T (K_t + \sigma_n^2 I)^{-1} \kappa_t(x), \end{aligned} \quad (2)$$

where $I \in \mathbb{R}^{t \times t}$ is the identity matrix, σ_n is the noise variance of the GP, K_t is the covariance matrix with $(K_t)_{i,j} = \kappa(x_i, x_j)$, $\kappa_t(\cdot)^T = [\kappa(x_1, \cdot), \dots, \kappa(x_t, \cdot)]^T$. We follow [12], assuming that the functions $f(x)$, $q(x)$ belong to a reproducing kernel Hilbert space (RKHS) associated with their respective kernels, and have bounded RKHS-norms. This assumption enables building well-calibrated confidence intervals over the two functions. To build the confidence intervals, we define the lower and upper confidence bounds $\text{lcb}_t^q(x)$, $\text{ucb}_t^q(x)$ as follows:

$$\text{lcb}_t(x) = \mu_t(x) - \beta^{f/q} \sigma_t(x), \quad (3)$$

$$\text{ucb}_t(x) = \mu_t(x) + \beta^{f/q} \sigma_t(x). \quad (4)$$

The corresponding parameters β^f , β^q can be individually adjusted to ensure the validity of confidence bounds [32] and balance exploration vs. exploitation (further referenced collectively as β parameters).

B. Online adaptation under safety constraints

In this work the underlying assumption is that for each measurement the task x_{τ} can be quantified. The online adaptation is enabled by using multi-task Gaussian Processes, where each measurement can be associated with a context and its corresponding task parameter.

The kernel of a multi-task GP is of the form $k_{\text{multi}}((x_{\text{opt}}, x_{\tau}), (x'_{\text{opt}}, x'_{\tau})) = k_{\tau}(x_{\tau}, x'_{\tau}) \otimes k(x_{\text{opt}}, x'_{\text{opt}})$, where $\otimes$ denotes the Kronecker product. This kernel decouples the correlations in objective values along the input dimensions, captured by k , from those across tasks, captured by k_{τ} [33]. Before the initialization of BO, GOOSE assumes a known non-empty safe seed $\mathcal{X}_0$ [12]. For online adaption using different tasks/contexts, $\mathcal{X}_0$ needs to be safe and fulfill the constraint(s), i.e., $q(\mathcal{X}_0) < c$ for all tasks. The safe seed $\mathcal{X}_0$ is also used as a backup solution for the BO if the set of safe points within the current task x_{τ} is empty, e.g. if the task changes drastically and no data is available to the GPs within a neighborhood of the new task parameters. In this case external prior knowledge of the operator is used.

In [12], GOOSE divides the optimization domain into two subsets: the *pessimistic safe set* $\mathcal{X}_t$ and the *optimistic safe set* $\mathcal{X}_t^{\text{opt}}$. $\mathcal{X}_t$ contains the set of points classified as safe while $\mathcal{X}_t^{\text{opt}}$ contains the points that potentially could be safe at iteration t . Formally,

$$\mathcal{X}_t = \{x \in \mathcal{X} \mid \text{ucb}_t^q(x) \leq c\}, \quad (5)$$

$$\mathcal{X}_t^{\text{opt}} = \{x \in \mathcal{X} \mid \exists \bar{x} \in W_t, \text{ s.t. } g_{\epsilon}^t(\bar{x}, x) > 0\}, \quad (6)$$

where the subscript t indicates the iteration of the Bayesian optimization, and the superscript q marks the constraint

$q(x)$. The set of expanders $W_t \subset \mathcal{X}_t$ corresponds to the periphery of the safe set, such that $\forall x \in W_t$ it holds true that $|\text{ucb}_t^q(x) - \text{lcb}_t^q(x)| > \epsilon$. Furthermore, $g_\epsilon^t(\bar{x}, x)$ is the noisy expansion operator taking values of 0 or 1, defined as

$$g_\epsilon^t(\bar{x}, x) = \mathbb{I} [\text{lcb}_t^q(\bar{x}) + \|\mu_t^\nabla(\bar{x})\|_\infty d(\bar{x}, x) + \epsilon \leq c], \quad (7)$$

for some $\epsilon > 0$ uncertainty threshold relating to the observation uncertainty of $q(x)$. $\|\mu_t^\nabla(\bar{x})\|_\infty$ is the mean of the posterior over the gradient of the constraint function, and $d(\cdot, \cdot)$ is the distance metric associated with Lipschitz continuity of $q(x)$.

In the original GOOSE formulation in [28], the sets $\mathcal{X}_t$ and $\mathcal{X}_t^{\text{opt}}$ need to be explicitly calculated during each BO iteration, following discretization of the whole optimization domain. The proposed definitions in (6) reduce the explicit calculation of the sets $\mathcal{X}_t$ and $\mathcal{X}_t^{\text{opt}}$ to criteria that can be checked for each candidate x . However, discretizing the domain to build the set of expanders W_t cannot be avoided, slowing down the algorithm for high dimensional domains or upon continuous optimization.

III. Modified GOOSE: SAFE CONTINUOUS GOAL-ORIENTED BO

To enable fast online optimization, we remove the expander search from the GOOSE algorithm of [12], without losing its safe expansive behaviour due to thoughtfully chosen hyperparameters, β parameters and prior settings coming from prior knowledge of the objective function $f(x)$ and the constraint $q(x)$. In this work, we use the squared exponential kernel

$$\kappa(x, x') = \sigma_v^2 \exp \left\{ -\frac{(x - x')^T (x - x')}{2L} \right\}, \quad (8)$$

where σ_v is the prior standard deviation of the GP and $L = \text{diag}(l_1, \dots, l_m)$ is the diagonal matrix of the length-scales for the corresponding dimensions of the input x . The modified GOOSE algorithm works with various other kernels e.g. Matérn kernels. As in [12], the acquisition function for selecting the next point is the lower confidence bound search

$$x_t = \underset{x \in \mathcal{X}_t}{\text{argmin}} \text{lcb}_t^f(x). \quad (9)$$

We introduce a method to make the optimizers of the modified GoOSE algorithm expansive and safe and therefore loose the need of explicit expanders. For this purpose we extend the algorithm of [12] by the additional assumption on $f(x)$, lower-bounding the cost by a real constant value:

$$\{\exists \xi \in \mathbb{R} \mid f(x) \geq \xi, \forall x \in \mathcal{X}\}. \quad (10)$$

This assumption can be easily satisfied for many practical applications, e.g. using quadratic or other non negative costs. In Section IV, we demonstrate a non negative $f(x)$, which fits well with our optimization objective. For safe expansive behaviour of the optimizers, the priors $\mu_0^f(x)$ and $\mu_0^q(x)$, the hyperparameters of the prior standard deviation σ_v^f and σ_v^q , and the β parameters β^f and β^q have to be tuned such that

$$\begin{aligned} \mu_0^q(x) + \beta^q \sigma_v^q &> c \quad \forall x \in \mathcal{X}, \\ \mu_0^f(x) - \beta^f \sigma_v^f &\leq \xi \quad \forall x \in \mathcal{X}, \end{aligned} \quad (11)$$

meaning that prior to adding any data to the GPs at $t = 0$, all $x \in \mathcal{X}$ are outside of the safe set $\mathcal{X}_t$, while the lower bound for unknown points $x \in X$ is less or equal to the absolute possible minimum of the cost function $f(x)$ (following from (10) and (11)). Intuitively the first condition ensures safety in the presence of uncertainty and the second condition ensures expansion in the presence of uncertainty, by providing a candidate $x \in \mathcal{X}$ for the acquisition function (see (9)), that is always better than the current best observation $\min(y_{1:t}) \geq \xi$.

A. Implementation of the Modified GOOSE

Algorithm 1 Modified GoOSE

Input: Safe seed $\mathcal{X}_0$, $f \sim \mathcal{GP}^f(\mu_0^f, \kappa^f; \sigma_n^f, \beta^f)$, $q \sim \mathcal{GP}^q(\mu_0^q, \kappa^q; \sigma_n^q, \beta^q)$, $x_\tau = x_{\tau 0}$
 $\mu_0^f, \mu_0^q, \sigma_n^f, \sigma_n^q, \beta^f, \beta^q$ tuned according and 10 and 11

- 1: **while** machine is running **do**
- 2: **while** termination criteria is false **do** ▷ Active phase
- 3: $\mathbf{x}_w \leftarrow (x_{\text{opt}, 1:t-1}, x_{\tau, t})$
- 4: $\mathcal{X}_{tw} \leftarrow \{x \in \mathbf{x}_w : \text{ucb}_t^q(x) \leq c\}$
- 5: $x_t^* \leftarrow \text{PSO}(\text{init}: \mathcal{X}_{tw}, \text{acq}: \text{lcb}_t^f, \text{s.t.}: \text{ucb}_t^q \leq c)$
- 6: Evaluate $f(x^*), q(x^*)$
- 7: **if** $\text{length}(\mathbf{x}_w) \geq \text{data}_{\text{limit}}$ **then**
 remove oldest data point from $\mathcal{GP}^f$ and $\mathcal{GP}^q$
- 8: Update $\mathcal{GP}_t^f$ and $\mathcal{GP}_t^q$ with $x^*, f(x^*), g(x^*)$
- 9: Update $x_{\tau, t}$
- 10: **while** restart criteria is false **do** ▷ Passive phase
- 11: $\mathbf{x}_w \leftarrow (x_{\text{opt}, 1:t-1}, x_{\tau, t})$
- 12: $\mathcal{X}_{tw} \leftarrow \{x \in \mathbf{x}_w : \text{ucb}_t^q(x) \leq c\}$
- 13: $x_t^* \leftarrow \text{PSO}(\text{init}: \mathcal{X}_{tw}, \text{acq}: \text{ucb}_t^f, \text{s.t.}: \text{ucb}_t^q \leq c)$
- 14: Evaluate $f(x^*), q(x^*)$
- 15: update $x_{\tau, t}$

The modified GOOSE is presented in Algorithm 1. It consists of an active and a passive phase. In the active phase, the GPs are updated with data obtained by the measurement. The model is updated until it has sufficiently learned the system performance and is terminated with a termination criterion. The termination criterion should represent the convergence of the optimization. In the experiments of Section IV we use a set of 30 additional data points added to the GPs. At each iteration t , it is checked which of the evaluated $x_{\text{opt}, 1:t-1}$ are safe for the applied task $x_{\tau, t}$, line 4. This is the safe set which is given as an initialization seed for the particle swarm optimization (PSO) of the acquisition function, line 5. In case this set is void, the predefined safe seed is used as the safe set. In the active phase, the acquisition function is set to equal to lcb. The PSO provides an optimizer x^* which is applied to the system. The obtained measurement is used for updating the GPs, line 8. In the passive phase, the learned models $\mathcal{GP}^f$ and $\mathcal{GP}^q$ are used to find the optimum input for the specific task by minimizing the ucb, line 13. The passive phase operates until the restart criterion is triggered, which is the case when the models are not accurate enough and need to be updated. In our experiments in Section IV, the restart criterion is in case of a constraint violation or new task detection.




Fig. 2: Parallel computation scheme to run the modified GoOSE. 1. ParaGoOSE: A horizon of the next m tasks is sampled and fed to the queue. LookupGoOSE: Corresponding to the next task, a neighborhood of tasks within the grid is selected and fed to the queue. 2. The optimization processes pick the entries in the task queue and calculate the corresponding optimizers/optima according to algorithm 1. 3. The optimizers/optima are returned to the manager thread to be executed. 4. The $\mathcal{GP}^f$ and $\mathcal{GP}^q$ are updated with $f(x^*)$ and $q(x^*)$.

B. Parallel computation schemes for fast optimization

Depending on the use case, one or more task/context parameters may change rapidly. Serially calculating the optimizers might not be sufficient to keep pace with the changing context. To be able to anticipate the next optimization step in (near) real-time, we need to develop strategies to move the whole frontier of points given by all possible contexts (or a large set of them) and to have predictions for the upcoming iterations for the whole frontier.

In Fig. 2, we present the overall parallel computation scheme proposed in this work. In the first method, (ParaGoOSE) (see Fig. 2), the optimization algorithm of Algorithm 1) is calculated m times on m different processes. A receding horizon of the next m observations is selected/predicted by the manager thread and fed to the queue of tasks. The prediction requires access in advance to the next tasks. The m processes are selected, removed, and optimized for the tasks in the queue until it is empty in a first come, first serve manner. Note that the number of processes can be less than the horizon length. Synchronization of the processes makes sure that the optimizations on all processes are in the same phase and have the same amount of data points added to the GPs. If the optimization is too slow and a measurement happens before the process is updated, the measurement gets ignored. Should there be no corresponding optimizer/optimum for a given task, due to the optimization being too slow, the predefined safe input $\mathcal{X}_0$ is used.

The second method (LookupGoOSE) uses a discretized set of task parameter combinations $\mathcal{X}_{x_\tau}$. For each task in $\mathcal{X}_{x_\tau}$, the

optimum/optimizer is calculated and saved in a lookup table of optima/optimizers, which is initialized with the predefined safe input $\mathcal{X}_0$. Multiple optimization processes update the optima/optimizer corresponding neighborhood of the applied task $x_{\tau,t}$ of the new data within the discretized set $\mathcal{X}_{x_\tau}$ added to a queue. The neighborhood is defined as:

$$x_\tau \in \mathcal{X}_{x_\tau} : |x_\tau - x_{\tau,t}| < k\Delta x_\tau, k > 0, \quad (12)$$

where Δx_τ is the discretization parameter of $\mathcal{X}_{x_\tau}$. When the phase switches the complete grid of optima/optimizers has to be updated. Equivalent to ParaGoOSE, the optimization processes are synchronized to be in the same state and the queue of tasks is emptied in a first come, first served manner.

One advantage of ParaGoOSE is that the task for which the optimizer/optimum was calculated, corresponds exactly to the task predicted n iterations ago, compared to LookupGoOSE, resulting in more precise optimizers/optima given that the prediction of the task is correct. Due to only a single task update for all processes every iteration the method is also computationally less expensive than LookupGoOSE and/or converges faster due to a lower ignore rate of the measurements. The disadvantage of ParaGoOSE compared to LookupGoOSE is that a n step horizon of the task needs to be computed. Unexpected changes in the task could lead to optima/optimizers calculated for the wrong context, which could lead to suboptimal and unsafe optimizers. Furthermore, each optimization result is calculated on the information level of n iterations in the past.

LookupGoOSE requires no knowledge of the next x_τ . Furthermore, its asynchronous update of the optima/optimizer makes it less dependent on computational resources. On the downside, new data can only be added to the GPs once the update of $\mathcal{X}_{x_\tau}$ is finished. Additionally, data acquired within an update step is ignored. Although the same strategy is applied to ParaGoOSE, this happens only occasionally due to the much lower computational cost of having only a single new task to optimize for each measurement. Therefore the adaptation rate of LookupGoOSE is slower than that of ParaGoOSE. When switching from Active Phase to Passive Phase or vice versa (see Algorithm 1) LookupGoOSE needs to optimize for the whole task grid $\mathcal{X}_{x_\tau}$, due to the change of the acquisition function. This update is computationally expensive and another disadvantage of LookupGoOSE compared to ParaGoOSE.

IV. CASE STUDY

In this section, we demonstrate the proposed approach, first on a numerical simulation of a precision motion system, and then, on the real system.

A. Implementation

The hyperparameters for the GPs in the experiments are given in Table I. For optimization of the acquisition function, all presented implementations of GoOSE use PSO [34] with initialization of a set of 50 particles across $\mathcal{X}_t$ with a fixed initial velocity of the particles v_0 , calculated from the $\mathcal{GP}^f$ lengthscales (see (8)) in random directions. In the case of the presented GoOSE implementation of [12], which we from

TABLE I: GP Hyperparameters for experiments in Section IV. $c(x_\tau)$ is the constraint limit based on the task, see (1).

Numerical experiment	f	q
$l = [P_{kp}, V_{kp}, V_{ki}, A_{ff}, \text{stepsize}]^T$ ($\sigma_n, \mu_0, \sigma_v$)	$[0.4, 0.8, 1.2, 0.3, 0.3]^T$ ($2.5\text{e-}3, 1.8, 0.36$)	$[0.4, 0.8, 1.2, 0.3, 0.3]^T$ ($0.12, c(x_\tau), 1$)
Real experiment	f	q
$l = [P_{kp}, V_{kp}, V_{ki}, A_{ff}, \text{stepsize}, \text{payload}]^T$ ($\sigma_n, \mu_0, \sigma_v$)	$[50, 100, 200, 0.5, 0.3, 1]^T$ ($1\text{e-}2, 1.8, 0.36$)	$[50, 100, 200, 0.5, 0.3, 1]^T$ ($9\text{e-}2, c(x_\tau), 1$)

now on term as *baseline* GOOSE, the particles expand into the optimistic safe set $\mathcal{X}_t^{opt}$ through the PSO update rules, avoiding explicit calculation. In the case of the presented modified GOOSE implementation, the optimistic safe set $\mathcal{X}_t^{opt}$ does not have to be considered, since no expanders are used.

The parameter $\beta = 3$ was fixed for every GP of the implementation, resulting in a 99% confidence interval. For the baseline GOOSE the exploration threshold, determining which parameters of the safe set periphery are defined as expanders, was set to $\epsilon = 6\sigma_q$ for all experiments.

For the calculation of the final optimum after the experiment, the acquisition function was set to

$$x^* = \underset{x \in \mathcal{X}_t}{\operatorname{argmin}} \operatorname{ucb}_t^f(x), \quad (13)$$

resulting in the pessimistic optimum from the passive phase.





Fig. 3: Top panel: Simplified controller architecture in the experimental study. Top panel. The cyan blocks contain control parameters optimized in this work. Bottom panel: Positioning system by Schneeberger Linear Technology AG.

B. System, controller, and optimization setup

The system used for experiments of the modified GOOSE algorithm is a linear axis within a high-precision positioning system manufactured by Schneeberger Linear Technology, as illustrated in Fig. 3. This axis is mobilized by a permanent magnet AC motor, outfitted with encoders offering nanometer-level precision for both position and speed monitoring. The

axis is guided along two rails using monorail bearings, resulting in positioning accuracy of $\leq 10, \mu\text{m}$, a repeatability margin of $\leq 0.7, \mu\text{m}$, and a 3σ stability of $\leq 1, \text{nm}$. The controller's sampling time, along with the data acquisition rate utilized for this system, stands at $f_s = 20\text{kHz}$. Such systems find routine application in the semiconductor industry, biomedical engineering, the fields of photonics, and solar technologies, predominantly for production and quality control purposes.

The control system under consideration is a three-level cascade controller, as depicted in Fig. 3. The hierarchical control structure consists of an outermost loop responsible for position control, employing a P-controller represented as $C_p(s) = P_{kp}$, and a middle loop dedicated to velocity control, employing a PI-controller denoted by $C_v(s) = V_{kp} + V_{ki}/s$. It's worth noting that the inner loop, which governs the current of the drive, is considered well-tuned and remains unchanged throughout the tuning and adaptation process. To expedite the system's response, feedforward structures are employed. The gain of the velocity feedforward, denoted as V_{ff} , is also considered well-calibrated and remains unaltered during the retuning and adaptation procedure. Conversely, the gain associated with acceleration feedforward, A_{ff} , is subject to adjustment via an algorithm during real system experiments and is subsequently set to $A_{ff} = 0$ for the simulation experiments.

The cost $f(x)$ and constraint $q(x)$ are defined as follows:

$$\begin{aligned} f(x) &= \frac{1}{n_P - n_s} \sum_{i=n_s}^{n_P} |\xi(i, n_s) p_e(t_i)|, \\ q(x) &= \max |\mathcal{F}[\xi(i, n_s) v_e(t_i)](x, f = f_1)|, \\ f_1 &= [140\text{Hz}, 1250\text{Hz}], \\ \xi(i, n_s) &= 1 - \frac{1}{1 + \exp(-(i - n_s - 150)/10)}, \end{aligned} \quad (14)$$

where $p_e(t_i)$ and $v_e(t_i)$ are the position and velocity error from the reference at time instance t_i and the settling phase extends from the time index n_s to the time index n_P . The function $\xi(i, n_s)$ is a left-sided sigmoid function used as a filter for the raw data (in this case, the velocity error). The constraint $q(x)$ is calculated from the Fourier transform of the filtered velocity error. For more details on the performance metrics' definition, we refer the reader to [20]. The parameters x that are adapted are the controller gains P_{kp} , V_{kp} , and V_{ki} , as well as the acceleration feedforward gain A_{ff} , and the adaptation is performed according to the $\log_{10}(\text{stepsize})$ of the motion and (on the experiments on the real system) the payload added to the system, which are

also provided as components of the vector x serving as task parameters (thus, they serve as inputs to the corresponding GPs, but are not optimization variables). The maximum speed, acceleration, and jerk of the reference movement are set to $[v_{\max}, a_{\max}, j_{\max}]^T = [0.9 \frac{m}{s}, 20 \frac{m}{s^2}, 200 \frac{m}{s^3}]^T$.

C. Numerical study

In a simulation, we demonstrate the adaptability and optimality of the modified GOOSE algorithm. To achieve this, an artificial drift term is added to the output position p in Fig. 3. This drift term is treated as a task parameter in the algorithm, alongside the stepsize. The artificial drift term is increased linearly over the iterations to test the adaptability of the algorithm.




Fig. 4: Simulation: Comparison of the modified GOOSE with one accounting for the drift term as a task parameter and the other without. The plots show the achieved cost and constraint value with respect to stepsize after convergence of the algorithm. It is evident that the GOOSE without the drift term has difficulty achieving optimality due to the presence of the artificial drift.

In Fig. 4, we compare two modified GOOSE algorithms, one with a task dimension for the drift term and one without. The cost values show that the modified GOOSE with the drift

term can account for the artificial drift and keep the cost at the optimum. Furthermore, Fig. 5 demonstrates how the modified GOOSE algorithm keeps the cost at the minimum despite the drift term by constantly modifying the gains. The results show successful adaptation of the proposed method to shifts due to external disturbances. Such adaptation scenarios have practical relevance in cases such as changing die sizes or weights in a semiconductor manufacturing scenario, or changing cutting tools with different weights in a laser cutting systems.





Fig. 5: Simulation: The achieved optimum cost over the iterations is depicted for the modified GOOSE with the drift term. It is able to compensate for the drift by constantly modifying the controller gains and thus to keep the cost at the optimum.


D. Experimental study

1) *Comparisons to the baseline GOOSE:* In a first experiment on the real system, we compare the baseline GOOSE with the modified GOOSE algorithm described in Section II. For this, we run both optimizations for $T = 135$ iterations with a fixed stepsize of 10mm and a payload that is changed every 15 iterations between 0.4kg and 2kg. Both algorithms start with an identical initial set of data $x_0 = [P_{kp}, V_{kp}, V_{ki}, A_{ff}, \log_{10}(\text{stepsize}), \text{payload}]^T = [200, 600, 1000, [0, 1, 2], 1, 0.4]^T$. The predefined safe set is set to $\mathcal{X}_0 = [P_{kp}, V_{kp}, V_{ki}, A_{ff}]^T = [200, 600, 1000, 0]^T$ for both algorithms.

Figure 6a shows the convergence of both algorithms after alternation of the payload is comparably fast, while the modified GOOSE performs slightly better, by comparing the rolling optimum calculated during the passive phase. Furthermore, it is shown that for both algorithms each further alternation of the payload requires fewer iterations to find the global optimum for the given task. After 3 alternations of the payload, the search for the optimum finishes almost instantaneously. Comparing the iteration time of both algorithms we can see that the modified GOOSE requires approximately the same computational time for each of the $T = 135$ iterations, while the original GOOSE computational time grows with each new



(a) Comparison of the rolling optimum. The solid line shows the prediction of the optimum at every iteration the lighter areas indicate the confidence intervals.



(b) Comparison of the total iteration time including the machine movement of around 2.4s.

Fig. 6: Comparison of the baseline GOOSE with the modified GOOSE tuning executed on the Argus system. The stepsize is fixed to a 10mm movement, while the payload is altered every 15 iterations.

data point, see Fig. 6b. The figure shows the complete iteration time of the experiment including the machine's movement of approximately 2.4s per iteration (10mm back and forth movement).

We conduct a second experiment on the real system to obtain the cost and constraint optima of the modified GOOSE (see Fig. 7). We use the initial samples

$$\begin{aligned} x_0 &= [P_{kp}, V_{kp}, V_{ki}, A_{ff}, \log_{10}(\text{stepsize}), \text{payload}]^T, \\ &= [200, 600, 1000, [0, 1, 2], [0, 2], 0.4]^T, \end{aligned}$$

predefined safe set $\mathcal{X}_0 = [P_{kp}, V_{kp}, V_{ki}, A_{ff}]^T = [200, 600, 1000, 0]^T$, $T = 300$ iterations with random step and alteration of the payload every 33 iterations between 0.4kg and 2kg, is compared with the optima of the built

heuristic non adaptive autotuning of the machine and an adaptive interpolated grid optimization LPV, where a grid of 1900 candidates $x \in X$ is recorded (all permutations of $x_{\text{opt}} = [[200, 300, 400], [600, 800, 1000], [1000, 1250, 1500]]^T$ are combined with all permutations of $x_{\tau} = [\log_{10}([1, 2, 5, 10, 20, 50, 100]), [0.4, 1.2]]^T$.

The optimum for all recorded tasks fulfilling the constraint is calculated. For a task in between or beyond known tasks, a linear interpolation strategy is chosen. It is shown that the modified GOOSE shows better performance than the heuristics-based autotuning of the controller. Both modified GOOSE and autotuning never violate the constraint and keep the system stable for any task, while the interpolated grid optimization LPV often violates the stability constraint and also drives the system once into an unstable position error for the 6mm step, where the automatic emergency stop was triggered, even though the interpolated grid optimization LPV was specifically trained for the shown payload of 0.4kg, and the other two approaches are adaptive also with respect to payload (thus, not specifically tuned to this payload). This can be explained by two effects. 1) The 6mm step was not directly tuned during the grid optimization but interpolated. 2) The interpolated grid optimization LPV method is not robust enough w.r.t. disturbances in the system, while GOOSE handles the disturbances due to the noise model in the GP model (see σ_n in (2) and Table I).

2) *Comparison of parallel schemes:* In a final experiment on the real system the two parallel schemes for the adaptive control modified GOOSE algorithm, described in Section II are compared. We create a similar experiment as used for Fig. 7. Here the tuning contains random stepsizes and 3 different payload settings altering between 0.4kg, 1.2kg, and 2.0kg. The initial samples x_0 and predefined safe set $\mathcal{X}_0$ are equivalent to the previous experiment. Both parallel adaptive control algorithms use 4 parallel processes for optimization. While using ParaGOOSE the predefined safe set $\mathcal{X}_0$ is used on the machine, if the optimization is not fast enough to deliver the corresponding optimizer/optimum (e.g. after the start of the algorithm). Similarly while using LookupGOOSE the grid of optimizers/optima is initialized using the predefined safe set $\mathcal{X}_0$ as the optimizer for every entry in the task grid. The task grid is selected as all permutations of the task

$$\begin{aligned} \mathcal{X}_{x_{\tau}} &= [\log_{10}(\text{stepsize}), \text{payload}]^T, \\ &= [\log_{10}([1, 2, \dots, 10, 20, \dots, 100]), [0.4, 0.6, \dots, 2.0]]^T. \end{aligned}$$

As discretization parameter of the neighborhood $\Delta x_{\tau} = [0.3, 0.2]^T$ is chosen. Fig. 8 summarizes the performance of both parallel schemes. The advantage of the ParaGOOSE is evident by looking at the number of added evaluations to the GPs since all the future tasks were known in advance for this experiment. Otherwise, the LookupGOOSE can handle the unknown future tasks with the bottleneck of its speed due to having to calculate the optima for a neighborhood of the applied task.

V. CONCLUSION

In this work, we present two run-time online Bayesian optimization algorithms that safely perform constrained op-




Fig. 7: Comparison of the cost and constraint of the identified optima calculated by the modified GOOSE algorithm, interpolated grid optimization LPV, and the built-in autotuning heuristic of the machine. In the case of GOOSE and interpolated grid optimization LPV, the optimum depends on the two adaptive task parameters defined as the stepsize and the payload. For illustration, the payload in the figure is fixed to 0.4kg and the x-axis displays the stepsize.

timization. To achieve continuous adaptation, we parallelize optimization iterations and further reduce the computation time of the GOOSE algorithm, by removing the expensive expander calculation of prior GOOSE implementations, while maintaining performance and safety. We demonstrate the proposed framework on a controller tuning and run-time add-on adaptive control application where high performance and safe evaluations are crucial to ensure physical safety without disruption of the operation. We present results in both simulations and on a real high-precision motion system to show the improved performance over state-of-the-art tuning and adaptive control methods that can handle arbitrary controller structures.

Our work was motivated by the adaptation of fast controllers in time-critical applications, however, the presented framework can be effective on a variety of control tasks with various controller structures. Future work will study the implementation of the proposed method in safety-critical industrial robotics applications.




Fig. 8: ParaGoOSE vs LookupGoOSE: First active phase of both algorithms, both algorithms are stopped after the first active phase is completed. The termination criterion is set to adding a total of 16 new data points to the GPs. Both algorithms work independently of the machine scheduling and do not prohibit the machine from moving, due to parallelization. It can be seen that ParaGOOSE is significantly faster at acquiring the data points, due to much fewer measurements being ignored. Furthermore, in the first 75 seconds of the LookupGoOSE, no data is added to the GPs, because the optimization processes have to update the whole task grid.

REFERENCES

- [1] K. K. Tan, T. H. Lee, and S. Huang, *Precision motion control: design and implementation*. Springer Science & Business Media, 2007.
- [2] G. Chowdhary, H. A. Kingravi, J. P. How, and P. A. Vela, "Bayesian nonparametric adaptive control using Gaussian processes," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 3, pp. 537–550, 2015.
- [3] R. C. Grande, G. Chowdhary, and J. P. How, "Experimental validation of Bayesian nonparametric adaptive control using Gaussian processes," *Journal of Aerospace Information Systems*, vol. 11, no. 9, pp. 565–578, 2014.
- [4] D. Sbarbaro and R. Murray-Smith, *Self-tuning Control of Non-linear Systems Using Gaussian Process Prior Models*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 140–157.
- [5] A. Gahlawat, P. Zhao, A. Patterson, N. Hovakimyan, and E. Theodorou, "L1-gp: L1 adaptive control with Bayesian learning," in *Learning for Dynamics and Control*, ser. Proceedings of Machine Learning Research, vol. 120. The Cloud: PMLR, 10–11 Jun 2020, pp. 826–837.
- [6] D. D. Fan, J. Nguyen, R. Thakker, N. Alatur, A.-a. Agha-mohammadi, and E. A. Theodorou, "Bayesian learning-based adaptive control for safety critical systems," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4093–4099.
- [7] F. Berkenkamp, A. P. Schoellig, and A. Krause, "Safe controller optimization for quadrotors with gaussian processes," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 491–496.
- [8] R. R. Duivenvoorden, F. Berkenkamp, N. Carion, A. Krause, and A. P. Schoellig, "Constrained Bayesian optimization with particle swarms for safe adaptive controller tuning," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 11 800 – 11 807, 2017, 20th IFAC World Congress.
- [9] M. Khosravi, V. Behrunani, R. S. Smith, A. Rupenyan, and J. Lygeros, "Cascade control: Data-driven tuning approach based on bayesian optimization," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 382–387, 2020.
- [10] M. Khosravi, V. Behrunani, P. Myszkowski, R. S. Smith, A. Rupenyan, and J. Lygeros, "Performance-driven cascade controller tuning with

- bayesian optimization,” *IEEE Transactions on Industrial Electronics*, p. 1–1, 2021.
- [11] M. Khosravi, C. Koenig, M. Maier, R. S. Smith, J. Lygeros, and A. Rupenyan, “Safety-aware cascade controller tuning using constrained bayesian optimization,” *IEEE Transactions on Industrial Electronics*, 2022.
 - [12] C. König, M. Turchetta, J. Lygeros, A. Rupenyan, and A. Krause, “Safe and efficient model-free adaptive control via bayesian optimization,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 9782–9788.
 - [13] D. Zhang and B. Wei, “A review on model reference adaptive control of robotic manipulators,” *Annual Reviews in Control*, vol. 43, pp. 188–198, 2017.
 - [14] L. Xu and B. Yao, “Adaptive robust precision motion control of linear motors with negligible electrical dynamics: theory and experiments,” *IEEE/ASME transactions on mechatronics*, vol. 6, no. 4, pp. 444–452, 2001.
 - [15] Z. Liu, C. Yuan, and Y. Zhang, “Linear parameter varying adaptive control of an unmanned surface vehicle,” *IFAC-PapersOnLine*, vol. 48, no. 16, pp. 140–145, 2015, 10th IFAC Conference on Manoeuvring and Control of Marine Craft MCMC 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S240589631502162X>
 - [16] E. C. Balta, K. Barton, D. M. Tilbury, A. Rupenyan, and J. Lygeros, “Learning-based repetitive precision motion control with mismatch compensation,” in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 3605–3610.
 - [17] L. Dai, X. Li, Y. Zhu, and M. Zhang, “Feedforward tuning by fitting iterative learning control signal for precision motion systems,” *IEEE Transactions on Industrial Electronics*, vol. 68, no. 9, pp. 8412–8421, 2021.
 - [18] T. Hayashi, H. Fujimoto, Y. Isaoka, and Y. Terada, “Projection-based iterative learning control for ball-screw-driven stage using basis function and data-based friction model,” in *IECON 2019 - 45th Annual Conference of the IEEE Industrial Electronics Society*, vol. 1, 2019, pp. 3293–3298.
 - [19] A. Rupenyan, M. Khosravi, and J. Lygeros, “Performance-based trajectory optimization for path following control using bayesian optimization,” in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 2116–2121.
 - [20] C. König, M. Ozols, A. Makarova, E. C. Balta, A. Krause, and A. Rupenyan, “Safe risk-averse bayesian optimization for controller tuning,” *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8208–8215, 2023.
 - [21] P. Brunzema, A. von Rohr, and S. Trimpe, “On controller tuning with time-varying bayesian optimization,” *arXiv preprint arXiv:2207.11120*, 2022.
 - [22] C. König, M. Khosravi, M. Maier, R. S. Smith, A. Rupenyan, and J. Lygeros, “Safety-Aware Cascade Controller Tuning Using Constrained Bayesian Optimization,” *arXiv e-prints*, p. arXiv:2010.15211, Oct. 2020.
 - [23] Y. Sui, A. Gotovos, J. W. Burdick, and A. Krause, “Safe exploration for optimization with Gaussian processes,” in *Proceedings of International Conference on Machine Learning (ICML)*, 2015.
 - [24] F. Berkenkamp, A. Krause, and A. P. Schoellig, “Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics,” *Machine Learning*, pp. 1–35, 2021.
 - [25] M. Zagorowska, C. König, H. Yu, E. C. Balta, A. Rupenyan, and J. Lygeros, “Efficient safe learning for controller tuning with experimental validation,” *arXiv preprint arXiv:2310.17431*, 2023.
 - [26] M. Fiducioso, S. Curi, B. Schumacher, M. Gwerder, and A. Krause, “Safe contextual bayesian optimization for sustainable room temperature pid control tuning,” in *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. AAAI Press, December 2019, pp. 5850–5856.
 - [27] J. Gardner, M. Kusner, Zhixiang, K. Weinberger, and J. Cunningham, “Bayesian optimization with inequality constraints,” in *Proceedings of International Conference on Machine Learning (ICML)*, ser. Proceedings of Machine Learning Research, vol. 32, no. 2. Beijing, China: PMLR, 22–24 Jun 2014, pp. 937–945.
 - [28] M. Turchetta, F. Berkenkamp, and A. Krause, “Safe exploration for interactive machine learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019, pp. 2891–2901.
 - [29] I. Bogunovic, J. Scarlett, and V. Cevher, “Time-varying gaussian process bandit optimization,” in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 314–323.
 - [30] K. Swersky, J. Snoek, and R. P. Adams, “Multi-task bayesian optimization,” *Advances in neural information processing systems*, vol. 26, 2013.
 - [31] D. Wang, D. Tan, and L. Liu, “Particle swarm optimization algorithm: an overview,” *Soft computing*, vol. 22, pp. 387–408, 2018.
 - [32] J. Kirschner and A. Krause, “Information directed sampling and bandits with heteroscedastic noise,” in *Conference On Learning Theory*. PMLR, 2018, pp. 358–384.
 - [33] K. Swersky, J. Snoek, and R. Adams, “Multi-task Bayesian optimization,” *Advances in Neural Information Processing Systems (NeurIPS)*, 01 2013.
 - [34] J. Kennedy and R. Eberhart, “Particle swarm optimization,” vol. 4, pp. 1942–1948 vol.4, 1995.