

Violation-Aware Contextual Bayesian Optimization for Controller Performance Optimization with Unmodeled Constraints

Wenjie Xu^{a,b}, Colin N Jones^a, Bratislav Svetozarevic^b, Christopher R. Laughman^c and Ankush Chakrabarty^{c,*}

^aLaboratoire d'Automatique, Swiss Federal Institute of Technology Lausanne (EPFL), Lausanne, 1015, Switzerland

^bUrban Energy Systems Laboratory, Swiss Federal Laboratories for Materials Science and Technology (Empa), Dübendorf, Zurich, 8600, Switzerland

^cMitsubishi Electric Research Laboratories, 201 Broadway, 8th Floor, Cambridge, MA 02139, USA

ARTICLE INFO

Keywords:

Violation-Awareness
Bayesian Optimization
Controlled System
Performance Optimization

ABSTRACT

We study the problem of performance optimization of closed-loop control systems with unmodeled dynamics. Bayesian optimization (BO) has been demonstrated to be effective for improving closed-loop performance by automatically tuning controller gains or reference setpoints in a model-free manner. However, BO methods have rarely been tested on dynamical systems with unmodeled constraints and time-varying ambient conditions. In this paper, we propose a violation-aware contextual BO algorithm (VACBO) that optimizes closed-loop performance while simultaneously learning constraint-feasible solutions under time-varying ambient conditions. Unlike classical constrained BO methods which allow unlimited constraint violations, or 'safe' BO algorithms that are conservative and try to operate with near-zero violations, we allow budgeted constraint violations to improve constraint learning and accelerate optimization. We demonstrate the effectiveness of our proposed VACBO method for energy minimization of industrial vapor compression systems under time-varying ambient temperature and humidity.

1. INTRODUCTION

Closed-loop systems can often be optimized after deployment by altering controller gains or reference inputs guided by the performance observed through operational data. Manually tuning these control parameters often requires care and effort along with considerable task-specific expertise. Algorithms that can automatically adjust these control parameters to achieve optimal performance are therefore invaluable for saving manual effort, time, and expenditure.

The optimal performance of a control system is generally defined via domain-specific performance functions whose arguments are outputs measured from the closed-loop system. While the map from measurements to performance may be clearly defined, the map from control parameters (that can actually be tuned) to performance is often unmodeled or unknown, since closed-form system dynamics may not be available during tuning [9]. It is thus common to treat the control parameters-to-performance map as a black-box, and design a data-driven tuning algorithm, where data is collected by experiments or simulations. However, since both experimentation and high-fidelity software simulations are expensive, tuning algorithms must be designed to assign a near-optimal set of control parameters with

as few experiments/simulations (equivalently, performance function evaluations) as possible. Therefore, existing data-driven methods that need a large number of samples, such as genetic algorithms [34], can be impractical.

It is precisely for this reason that Bayesian optimization (BO)¹ has received widespread attention in the context of closed-loop performance optimization. BO is a sample-efficient derivative-free global optimization method [19, 43] that utilizes probabilistic machine learning to intelligently search through parameter spaces. [12] gives a detailed survey of Bayesian Optimization. In recent work, BO has demonstrated potential in controller gain tuning. For example, BO has been applied to the tuning of the PI controller of a heat pump [20] and the tuning of PID cascade controller gains [21]. BO has also been applied to the performance optimization of model predictive control. For example, BO was applied to optimize the nominal linear model of a predictive controller [3], to tune the parameters of MPC to optimize the closed-loop performance [32], and to generate candidate parameters for data-driven scenario optimization [31]. BO has also been proposed to select closed loop kernel based model [4]. BO was used in many other various real-world control applications, such as wind energy systems parameter tuning [1, 2], engine calibration [29], and space cooling system optimization [9].

A challenge that has garnered recent interest is that of *safe* Bayesian optimization; that is, BO in the presence of safety-critical system constraints. These constraints may also be unmodeled ('black-box'), as a mathematical representation of the constraint with respect to the control parameters is not always known or straightforward to represent.

*Corresponding author

✉ wenjie.xu@epfl.ch (W. Xu); colin.jones@epfl.ch (C.N. Jones); bratislav.svetozarevic@empa.ch (B. Svetozarevic); laughman@merl.com (C.R. Laughman); achakrabarty@ieee.org (A. Chakrabarty)

🌐 <https://jackiexuw.github.io/> (W. Xu); (C.N. Jones); (B. Svetozarevic); (C.R. Laughman); (A. Chakrabarty)

ORCID(s):

1..

¹Also known as efficient global optimization (e.g., in [19]) or kriging (e.g., in [18]) in optimization and engineering literature.




Figure 1: A motivating example comparing our violation-aware BO to existing state-of-the-art methods. Safe BO gets stuck in a local minimum and fails to identify the global optimum, while generic constrained BO identifies the global optimum but can incur large constraint violations during the sampling process. In contrast, our method can simultaneously identify the global minimum and manages the constraint violation well.

To handle these constraints, safe BO methods have been recently proposed in [38], improved in [37], and extended to a more general setting in [40]. These methods either operate on the principle of not allowing any constraint violations during optimization, or leverage partial model knowledge to ensure safety via Lyapunov arguments [8]. In either case, safe BO learns optima without violating unmodeled constraints during the sampling process, or risks their violation with a predefined small probability. Often, this conservativeness results in obtaining local minima, slow convergence speeds, and reduced data efficiency. Conversely, generic constrained optimization with BO learns constraints without paying heed to the amount of constraint violation during the exploration phase [13, 14]. More recently, a group of works [42, 44] propose an optimistic constrained optimization approach, with applications to control system tuning. These methods are mostly agnostic to the deleterious consequences of constraint violation, such as long-term damage to expensive hardware caused by large violations, rendering them impractical for many industrial applications. Another direction of BO research proposes the use of budgets on the cost of samples (e.g., neural network training time [36], wall clock time [24], sample number [23] and system failures [26]). However, in these existing BO settings, the budget considered is usually related to the effort or failure risk for performance function evaluation, and does not provide a way to manage the magnitude of constraint violations.

For many industrial systems, small constraint violations over a short period are often acceptable if that exploration improves the convergence rate of optimization, but large violations are strongly discouraged. For example, in vapor compression systems (VCSs), it is imperative that constraint violation on variables such as compressor discharge temperature are limited to short time periods. We aim to find a set of near-optimal parameters within as few samples as possible since performance evaluation is time-consuming and available tuning time can be limited. Therefore, it may be desirable to systematically trade tolerable constraint violations for faster convergence and potentially skipping local minima. In other words, for VCSs, the benefits of

accelerated global convergence outweigh the cost of short-term constraint violations.

The performance of control systems is also influenced by exogenous signals, such as variations in the environment. We refer to such signals as *context* variables. It is often necessary to adapt control policies to maintain the performance and feasibility of the control system despite changes in the context variables. For instance, a controller designed for both performance and safety may adapt to prioritize safety over performance if the context variable predicts an unsafe event about to occur. In order to systematically incorporate contextual information, contextual Bayesian optimization approaches have been proposed in [22], where the inputs to the learner include the context variables augmented with the optimization variables. Contextual Bayesian optimization approaches have recently been tested on controller design applications: for example, safe Bayesian optimization has been extended to the contextual case in [11] to tune a room temperature controller via PID gain tuning. Furthermore, in [30], contextual Bayesian optimization is applied to cooperative wind farm control to maximize the efficiency of generated power. However, the incorporation of context in our setting, that is: how to simultaneously tune the controlled system efficiently and manage the constraint violations under a tolerable level with time-varying contexts, has not been studied previously, to the best of our knowledge.

Beyond the application of Bayesian optimization (BO) in this area, some classical works in process engineering consider the optimal process design with flexibility under uncertainty. For example, [17, 16] consider the optimal design problem of chemical process under parametric uncertainty of some exogenous factors that influence the process. However, these classical works consider the offline optimization of process design but do not consider the online adaptation of control parameters in response to the time-varying exogenous variables, which is the problem considered in this paper.

Along the methodological research line, robust optimization is a widely adopted method to handle uncertainty in the optimization problem [5]. It is also widely applied in process control [33, 39]. However, the goal of robust

optimization is to find a solution that is feasible and optimal for the worst case. However, we consider an online scenario where we can gradually reduce the uncertainty of the objective/constraints through sequential evaluations.

In this paper, we propose a novel violation-aware contextual Bayesian optimization approach (VACBO²) that exhibits accelerated convergence compared to safe BO, while ensuring the violation cost is within a prescribed budget under time-varying exogenous variables, which are typically referred to as contextual variables in Bayesian optimization literature. We demonstrate that our VACBO algorithm is less conservative than ‘safe BO’ algorithms that tend to be sample-inefficient and can get stuck in a local minimum because they cannot allow any constraint violation. The VACBO algorithm is also more cautious than constrained Bayesian optimization, which is agnostic to constraint violations and thus, is likely to incur large violation costs. Our VACBO algorithm is based on the principle of encouraging performance function evaluation at combinations of control parameters that greatly assist the optimization process, as long as it does not incur high constraint violations likely to result in system failure or irreversible damage.

Our VACBO algorithm is an extension of the VABO algorithm [45] to the contextual case. We augment the input space with exogenous variables and design a tractable auxiliary acquisition optimization problem specific to the contextual setting. More specifically, the VABO algorithm can not directly incorporate the impact of exogenous variables, which can be significant in many applications. For example, the impact of ambient temperature and ambient humidity can be significant in vapor compression system set-point tuning. To incorporate the exogenous variables, we augment the input variables, which we can control, with exogenous variables that are measured from the environment. With this augmentation, our method can learn the joint impact of the input variables and exogenous variables based on Gaussian process learning. Furthermore, the existing constrained expected improvement acquisition function used in VABO [45] and other constrained BO papers [14, 13] are not readily applicable to the contextual case. To address this issue, we extend the commonly used Constrained Expected Improvement (CEI) acquisition function to the contextual case and propose the Constrained Proxy Expected Improvement (CPEI) acquisition function.

Our **contributions** include:

1. We propose a new variant of constrained BO methods for control parameter tuning that improves global convergence rates within a prescribed amount of constraint violation with guaranteed high probability under a time-varying contextual setting;
2. We propose a simple and tractable constrained auxiliary acquisition function optimization problem for

trading off performance improvement and constraint violation;

3. To incorporate the environmental conditions that impact the objective and constraints, we augment the input space by exogenous variables and propose a new acquisition function by extending the commonly used Constrained Expected Improvement (CEI) [14, 13] to the contextual setting; and,
4. We validate our algorithms on a set-point optimization problem using a high-fidelity VCS that has been calibrated on an industrial HVAC system, with ambient temperature and ambient humidity as two exogenous variables. Simulation results with real-world weather signals as context variables demonstrate that our method efficiently minimizes the power consumption while simultaneously managing constraint violations within a tolerable level.

We now state the organization of our paper. In Sec. 2, we present the statement of our problem and our proposed solution concept. Then in Sec. 3, we present our VACBO algorithm. After that, we give the application result of a case study on vapor compression system in Sec. 4. Finally, we conclude the paper in Sec. 5.

2. Preliminaries

2.1. Problem Statement

We consider closed-loop systems of the form

$$\xi_+ = F(\xi, \theta, z), \quad (1)$$

where $\xi, \xi_+ \in \mathbb{R}^{n_\xi}$ denote the system state and the successor state (respectively), $\theta \in \Theta \subset \mathbb{R}^{n_\theta}$ the control parameters (e.g., set-points) to be tuned, $z \in \mathcal{Z} \subset \mathbb{R}^{n_z}$ the exogenous variables (e.g., ambient temperature) that affect the dynamics, and $F(\cdot, \cdot, \cdot)$ the closed-loop dynamics with initial condition ξ_0 . We assume that the closed-loop system (1) is designed such that it is exponentially stable to a control parameter and context dependent equilibrium state $\xi^\infty(\theta, z)$ for every $(\theta, z) \in \Theta \times \mathcal{Z}$. We further assume that $\xi^\infty(\cdot, \cdot)$ is a continuous map on $\Theta \times \mathcal{Z}$. We also assume that ξ converges to the steady state $\xi^\infty(\theta, z)$ in a faster time scale than the variation of the control parameter θ and the context variable z . To further clarify the unmodelled cost/constraints and their relationships to the unmodelled dynamics, we introduce the instantaneous cost function $L(\xi, \theta, z)$ and the instantaneous constraint functions $G_i(\xi, \theta, z), i \in [N]$. Since ξ^∞ is an unmodelled function of (θ, z) in steady state, we can then introduce the steady-state cost and constraint functions as functions of (θ, z) ,

$$\ell(\theta, z) = L(\xi^\infty(\theta, z), \theta, z), \quad (2)$$

$$g_i(\theta, z) = G_i(\xi^\infty(\theta, z), \theta, z), i \in [N]. \quad (3)$$

Depending on the specific problems, it is possible that L and G_i are well-modelled and have closed-form expressions.

²Code is available at <https://github.com/PREDICT-EPFL/VACBO>.

However, since ξ^∞ is typically unmodelled, ℓ and g_i , as functions of (θ, z) , are treated as unmodelled black-box functions in this paper.

This is not unusual: while ℓ may be well-defined in terms of system outputs, it is often the case that the map from control parameters and exogenous variables to cost remains unmodeled; in fact, ℓ may not even admit a closed-form representation on Θ and $\mathcal{Z}$, c.f. [9, 7].

We also define N unmodeled constraints on the system outputs that require caution during tuning. The i -th such constraint is given by $g_i(\theta, z) : \mathbb{R}^{n_\theta+n_z} \rightarrow \mathbb{R}, i \in [N]$, where the notation $[N] \stackrel{\text{def}}{=} \{i \in \mathbb{N}, 1 \leq i \leq N\}$; we assume each $g_i(\cdot, \cdot)$ is continuous on $\Theta \times \mathcal{Z}$. We assume that the cost function $\ell(\theta, z)$ and every constraint $g_i(\theta, z), i \in [N]$ can be ascertained, either by measurement or estimation, during the hardware/simulation experiment. We introduce the brief notation $g(\cdot, \cdot) \leq 0 \triangleq g_i(\cdot, \cdot) \leq 0, i \in [N]$ and assume that an initial feasible set of solutions is available at design time.

Assumption 1. *The designer has access to a non-empty safe set $S_0 \subset \Theta \times \mathcal{Z}$ such that for any $(\theta, z) \in S_0$, all constraints are satisfied; that is, $g(\theta, z) \leq 0$ for every $(\theta, z) \in S_0$.*

While such an initial set S_0 can be derived using domain expertise, it is likely that S_0 contains only a few feasible (θ, z) , and at worst could even be a singleton set. Assumption 1 is a common assumption in the literature of safe optimization (e.g., [38]). Without this assumption, it is possible that in the initial steps, we may already sample some points with large violations, which makes the problem unlikely to be tractable at all. Furthermore, we notice that in many applications (e.g., vapor compression system set-points tuning), an initial set of feasible (maybe suboptimal) set-points are known based on domain knowledge. Therefore, Assumption 1 is a necessary but not restrictive assumption for our problem.

We cast the control parameter tuning problem as a black-box constrained optimization problem, formally described by

$$\min_{\theta \in \Theta} \ell(\theta, z), \quad (4a)$$

$$\text{subject to: } g_i(\theta, z) \leq 0, \quad \forall i \in [N]. \quad (4b)$$

The exogenous variable z represents some quantities reflecting the environmental conditions, which we can measure but not directly control at each step. We give an example in the following.

Example. *Ambient temperature and ambient humidity are two important exogenous variables that impact the operation of the vapor compression system. Both of them can be measured but not directly controlled during the operations of the vapor compression systems.*

Furthermore, z can be time-varying with its own dynamics. We use z_t to denote the value of z at time step t . Our **objective** is to solve the constrained optimization

problem (4) with limited constraint violations during the optimization process. Since the constraints are assumed to be unmodeled and a limited set of feasible solutions is known at design time, we do not expect a guarantee of zero constraint violation. The tolerable amount and duration of constraint violations are problem-dependent. In some applications, such as vapor compression systems, small constraint violations over a short-term are acceptable, while large constraint violations are strongly discouraged. In such cases, instead of being overly cautious and ending up with suboptimal solutions, we allow small constraint violations *as long as the resulting knowledge gathered by evaluating an infeasible (in terms of constraint violation) θ accelerates the optimization process or helps avoid local minima.*

Remark 1. *Our formulation (4) can also optimize batch processes over finite-time horizons, say T_h . This would involve defining the objective and constraints over a batch trajectory with stage loss $\ell(\theta, z) := \frac{1}{T_h} \int_0^{T_h} l(\tau, \theta, z) d\tau$.*

2.2. Proposed Solution

We propose a modified Bayesian optimization framework to solve the problem (4) that is violation-aware: the algorithm automatically updates the degree of risk-taking in the current iteration based on the severity of constraint violations in prior iterations. Concretely, for an infeasible θ , the constraint violation cost is given by

$$\bar{c}_i(\theta, z) \triangleq c_i([g_i(\theta, z)]^+), \quad i \in [N] \quad (5)$$

where $[g_i(\cdot, \cdot)]^+ := \max\{g_i(\cdot, \cdot), 0\}$ and $c_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. Note that g_i corresponds to physically meaningful system outputs that we can measure, e.g., temperature. This violation cost function c_i is user-defined as a means to explicitly weigh the severity of ‘small’ versus ‘large’ constraint violations. While the function c_i is at the discretion of the designer, it needs to satisfy the following mild assumptions in order to achieve desirable theoretical properties; see §3.

Assumption 2. *The violation cost function c_i satisfies:*

$$(A1) \quad c_i(0) = 0,$$

$$(A2) \quad c_i(s_1) \geq c_i(s_2), \text{ if } s_1 > s_2 \geq 0,$$

$$(A3) \quad c_i \text{ is left continuous on } \mathbb{R}_{\geq 0}.$$

Assumption 2 captures some intuitive properties required of the violation cost function. According to (A1), there is no cost associated with no violation. From (A2), we ensure that the violation cost is monotonically non-decreasing with increased violations. Finally (A3) ensures that this monotonic increase is smooth and does not exhibit discontinuous jumps from the left. Although c_i is typically a smooth function (e.g., linear function), the composite cost function $\bar{c}_i(\theta, z)$ may be non-smooth because the function $[\cdot]^+$ is non-differentiable at the point 0.

To adapt the degree of risk-taking based on prior data obtained, we define a violation budget over a horizon of $T \in \mathbb{N}$ optimization iterations. Our goal is to sequentially

search over T iterations $\{\theta_t\}_{t=1}^T$ while using a prescribed budget of constraint violations in order to obtain a feasible and optimal set of parameters

$$\begin{aligned} & \min_{\theta_t \in \Theta} \sum_{t=1}^T \ell(\theta_t, z_t) \\ \text{subject to: } & \sum_{t=1}^T \bar{c}_i(\theta_t, z_t) \leq B_i, \quad i \in [N] \end{aligned} \quad (6)$$

where B_i denotes a budget allowed for the i -th violation cost. Note that this formulation is a generalization of well-known constrained/safe Bayesian optimization formulations proposed in the literature. If we set all $B_i \equiv 0$, then our formulation is closely related to safe BO [38, 37]. Alternatively, setting $B_i \equiv \infty$ reduces our problem to constrained BO agnostic to violation cost [13, 14].

3. Violation-Aware Contextual Bayesian Optimization

3.1. Bayesian Optimization Preliminaries

For Bayesian optimization, one models $\ell(x)$ and $g(x)$ as functions sampled from independent Gaussian processes. In our case, the input x to the Gaussian process consists of control parameters θ and the exogenous variables z . At iteration t , conditioned on previous input and function evaluation data $\mathcal{D} := \{(\theta, z)_{1:t}, \ell((\theta, z)_{1:t})\}$, the posterior mean and standard deviation of ℓ is given by

$$\mu_{\ell}(x|\mathcal{D}) = k_{\ell}^{\top}(x, x_{\mathcal{D}})K_{\ell}^{-1}\Delta y_{\ell} + \mu_{\ell,0}(x) \quad (7)$$

and

$$\sigma_{\ell}^2(x|\mathcal{D}) = k_{\ell}(x, x) - k_{\ell}^{\top}(x, x_{\mathcal{D}})K_{\ell}^{-1}k_{\ell}(x_{\mathcal{D}}, x), \quad (8)$$

where $x_{\mathcal{D}} = (\theta, z)_{1:t}$ is the set of control parameters and exogenous variables with which previous experiments/simulations have been performed. Here,

$$\begin{aligned} k_{\ell}(x, x_{\mathcal{D}}) &\triangleq [k_{\ell}(x, x_i)]_{x_i \in x_{\mathcal{D}}}, \\ k_{\ell}(x_{\mathcal{D}}, x) &\triangleq [k_{\ell}(x_i, x)]_{x_i \in x_{\mathcal{D}}}, \\ K_{\ell} &\triangleq (k_{\ell}(x_i, x_j))_{x_i, x_j \in x_{\mathcal{D}}}, \\ \Delta y_{\ell} &\triangleq [\ell(x_i) - \mu_{\ell,0}(x_i)]_{x_i \in x_{\mathcal{D}}}, \end{aligned}$$

and $k_{\ell}(\cdot, \cdot)$ is a user-defined kernel function and $\mu_{\ell,0}$ is the prior mean function, both associated with ℓ ; see [12] for more details on kernel and prior mean selection. The above quantities are all column vectors, except K_{ℓ} , which is a positive-definite matrix. Here, note that $K_{\ell} \in \mathbb{R}^{|D| \times |D|}$ and thus the inversion of K_{ℓ} requires $\mathcal{O}(|D|^3)$ time complexity. In this paper, we mainly deal with the scenario where $|D|$ is small to medium and thus the computation complexity is tolerable. This is common for Bayesian optimization since it finds the optimal solution within a small number of samples. For the regime of large $|D|$, there are many works on computationally efficient Gaussian process inference [35, 25].

For the constraint functions g , similar expressions for the posterior mean $\mu_{g_i}(x|\mathcal{D})$ and standard deviation $\sigma_{g_i}(x|\mathcal{D})$ can be obtained.

The kernelized functions above provide tractable approximations of the cost of the closed-loop system, along with the constraint functions, both of which were hitherto unmodeled/unknown. Classical BO methods use the statistical information embedded within these approximations to intelligently explore the search space Θ via acquisition functions. A specific instance of an acquisition function commonly used in constrained BO is the constrained expected improved (CEI) function [13]. It is defined as the expectation of the multiplication of improvement as compared to the incumbent best objective sampled so far and the feasibility indicator. However, in the contextual case, the objective may heavily rely on context. The incumbent best objective value under a favorable context may mislead the parameter search under the adversarial context.

Therefore, instead of using the incumbent best objective, we propose a two-step approach. In the first step, we use the minimum value of the posterior mean of ℓ to construct a proxy of the best objective sampled so far under a different context as in (9).

$$\hat{\ell}_t^{\min}(z) = \min_{\theta \in \Theta} \mu_{\ell}((\theta, z)|\mathcal{D}) \quad (9)$$

In the second step, we use the minimum value of posterior mean $\hat{\ell}_t^{\min}$ to construct a new acquisition function, the Constrained Proxy Expected Improvement (CPEI), given by

$$\text{CPEI}((\theta, z)|\mathcal{D}) = \mathbb{E} \left(\prod_{i \in [N]} \mathbf{1}_{g_i(\theta, z) \leq 0} I(\theta, z)|\mathcal{D} \right), \quad (10)$$

where $\mathbf{1}$ denotes the indicator function, $\mathbb{E}$ denotes the expectation operator, and $I(\theta, z) = \max\{0, \hat{\ell}_t^{\min}(z) - \ell(\theta, z)\}$ is the improvement of (θ, z) over the proxy $\hat{\ell}_t^{\min}(z)$ for the best incumbent solution over t iterations.

Since we use independent Gaussian processes to model $g_i(\theta, z)$, $\forall i \in [N]$ and $\ell(\theta, z)$, we deduce

$$\text{CPEI}(\theta, z|\mathcal{D}) = \prod_{i \in [N]} \mathbb{P}(g_i(\theta, z) \leq 0|\mathcal{D}) \mathbb{E}(I(\theta, z)|\mathcal{D}). \quad (11)$$

For the unsampled point (θ, z) , the corresponding output $g_i(\theta, z)$ and $\ell(\theta, z)$ are modeled as independent random variables, due to the assumption that these Gaussian processes are independent. This is standard assumption in the literature of Bayesian optimization with black-box unknown constraints [13, 14]. Further exploiting the correlation between $g_i(\theta, z)$, $\forall i \in [N]$ and $\ell(\theta, z)$ to speed up the learning of g_i and ℓ is left as future work.

We have $\mathbb{P}(g_i(\theta, z) \leq 0|\mathcal{D}) = \Phi \left(\frac{-\mu_{g_i}(\theta, z|\mathcal{D})}{\sigma_{g_i}(\theta, z|\mathcal{D})} \right)$, and the closed-form expression of expected improvement [19],

$$\mathbb{E}(I(\theta, z)|\mathcal{D}) = \Delta I(\theta, z|\mathcal{D})\Phi(w) + \sigma_{\ell}(\theta, z|\mathcal{D})\phi(w), \quad (12)$$

where $\Delta l(\theta, z|D) = \hat{\ell}_t^{\min}(z) - \mu_{\ell}(\theta, z|D)$, $w = \frac{\hat{\ell}_t^{\min}(z) - \mu_{\ell}(\theta, z|D)}{\sigma_{\ell}(\theta, z|D)}$, $\Phi(\cdot)$ and $\phi(\cdot)$ are the standard normal cumulative distribution and probability density functions, respectively.

3.2. VACBO Algorithm

Our VACBO algorithm proposes an auxiliary optimization problem that leverages the constrained proxy expected improvement acquisition function to guide the search of feasible points with potentially lower objective to evaluate while ensuring (with high probability) that the violation cost will remain within a prescribed budget.

Given the total violation cost budget, a question is how to allocate the budget across different samples. Intuitively, it may be beneficial to dynamically adjust the violation cost budget allocated to a single step. For example, if we find that we incur no violation cost for several steps, it is possible that we are overly cautious in those steps and may get stuck in a local minimum. So we can then increase the violation cost budget allocated for the next step. It is also possible that we incur significant violation in some step due to the over-confidence in the constraint function prediction by Gaussian process regression. In this case, we need to decrease the violation cost budget allocated to one single step. To capture this intuition in our algorithm, we design a violation cost budget allocation scheme to dynamically adjust the violation cost allocated to one single step. We use $B_{i,t}$ to denote the violation cost budget allocated to the step t for the i -th constraint. Our violation cost budget allocation scheme is given as,

$$B_{i,t} \triangleq \min \left\{ \max \left\{ B_i S_{i,t} - \sum_{\tau=1}^{t-1} \bar{c}_i(\theta_{\tau}, z_{\tau}), 0 \right\}, B_i^{\max} \right\}, \quad (13)$$

where $S_{i,t}$ is a non-negative and non-decreasing sequence that satisfies $S_{i,T} = 1$ and $B_i^{\max}$ is the maximum violation cost tolerable for the i -th constraint, which is a user-provided parameter that captures the user's maximum tolerance for constraint violation in one single step.

At this iteration, after observing the exogenous variables, we solve the following *auxiliary* problem

$$\theta_t^* := \arg \max_{\theta \in \Theta} \text{CPEI}(\theta, z|D), \quad (14a)$$

$$\text{subject to: } \prod_{i \in [N]} \mathbb{P}(\bar{c}_i(\theta, z) \leq B_{i,t}) \geq 1 - \epsilon_t, \quad (14b)$$

to compute the next control parameter candidate θ_t^* , where $0 < \epsilon_t \ll 1$ determines the probability of large constraint violation. Note that (14a) involves maximizing the constrained expected improvement type objective, which is common to cBO algorithms; c.f. [13]. Our modification using the budget, as written in (14b), enforces that the next sampled point will not use up more than the violation cost budget $B_{i,t}$ for all constraints with a probability of at least $1 - \epsilon_t$, conditioned on the data seen so far. This modification

allows us to trade a prescribed level of violation risk for more aggressive exploration, leading to faster convergence. Although we use the constrained proxy expected improvement acquisition function here, we can easily generalize to other acquisition functions by simply replacing CPEI in the objective (14a).

We now discuss how to efficiently solve the auxiliary problem (14). Recall from Assumption 2 that c_i is non-decreasing on $\mathbb{R}_{\geq 0}$ for every $i \in [N]$. Therefore, we can define an inverse violation function $c_i^{-1}(s) = \sup\{r \in \mathbb{R}_{\geq 0} \mid c_i(r) \leq s\}$ for any $s \in \mathbb{R}_{\geq 0}$. Therefore, we can write $\mathbb{P}(\bar{c}_i(\theta, z) \leq B_{i,t}|D) = \mathbb{P}([g_i(\theta, z)]^+ \leq c_i^{-1}(B_{i,t})|D) = \mathbb{P}(g_i(\theta, z) \leq c_i^{-1}(B_{i,t})|D)$. Since $g_i(\theta, z)$ follows a Gaussian distribution with mean $\mu_{g_i}(\theta, z|D)$ and variance $\sigma_{g_i}^2(\theta, z|D)$, we get

$$\mathbb{P}(\bar{c}_i(\theta, z) \leq B_{i,t}|D) = \Phi \left(\frac{c_i^{-1}(B_{i,t}) - \mu_{g_i}(\theta|D)}{\sigma_{g_i}(\theta|D)} \right). \quad (15)$$

When the number of control parameters n_{θ} is small (e.g., < 6), we can place a grid on Θ and evaluate the cost and constraints of (14) at all the grid nodes. The maximum feasible solution can then be used as the solution to the auxiliary problem. When the number of control parameters is large (e.g., $n_{\theta} > 6$), we can use gradient-based methods with multiple starting points to solve problem (14), since evaluating the learned GPs approximating ℓ and g require very little computational time or effort when T is not large (empirically, < 2000). Indeed, the main contribution of our paper is designing a new violation-aware Bayesian optimization algorithm by formulating a constrained auxiliary problem (Problem (14)) to acquire the next sample point. How to solve the auxiliary problem is not the main focus of this paper. Indeed, since the auxiliary constrained optimization problem (14) is a nonlinear programming problem with explicit closed-form expressions, there are many existing mature non-linear programming solvers that can reliably solve the auxiliary non-linear programming problem (such as Ipopt [41]). From the discussion in the last paragraph in Sec. 3.1 (Eq. (11) and Eq. (12)), it can be seen that the objective function of CPEI has a closed-form expression in the variables of posterior means and posterior standard deviations. From Eq. (15), we can also observe that the constraints in Eq. (14) also have closed-form expressions in the posterior means and the posterior standard deviations. Although the composite function $\bar{c}_i(\theta, z)$ may be non-smooth, Eq. (15) indicates that the probability $\mathbb{P}(\bar{c}_i(\theta, z) \leq B_{i,t}|D)$ is differentiable with respect to the posterior mean μ_{g_i} and the standard deviation σ_{g_i} , which are differentiable functions of θ . Furthermore, from Eq. (7) and Eq. (8), we can see that the posterior means and the posterior standard deviations also have closed-form expressions in the input variable $x = (\theta, z)$. Therefore, we can apply the chain rules to derive the explicit gradient of the objective/constraints with respect to the control parameters θ using the closed-form gradients of the objective/constraints with respect to

the posterior means/standard deviations and the closed-form gradients of the posterior means/standard deviations with respect to the control parameters θ . Since both the objectives and the constraints in Eq. (14) may potentially be non-convex, theoretical guarantees on convergence to the global optimum can not be derived. However, we start from random initial points, which can improve the chance of finding the global optimum. Indeed, when the number of the random initial points is large enough, we can identify the global optimum with high probability. This technique is commonly adopted in the Bayesian optimization literature. We provide pseudocode for implementation in Algorithm 1.

The following proposition provides a probabilistic guarantee of keeping the violation cost below the given budget. It highlights the “violation-awareness” property exhibited by VACBO.

Proposition 1. Fix $\delta \in (0, 1)$ and $T \in \mathbb{N}$. If $\epsilon_t, t \in [T]$ are chosen such that $\delta = 1 - \prod_{t=1}^T (1 - \epsilon_t)$, then the VACBO algorithm satisfies the probability that

$$\left\{ \max_{i \in [N]} \bar{c}_i(\theta_t, z_t) \leq B_i^{\max} \text{ and } \sum_{i=1}^T \bar{c}_i(\theta_t, z_t) \leq B_i, \forall i \in [N] \right\}$$

is at least $1 - \delta$.

Proof. Let

$$\mathcal{E}_t^1 := \left\{ \max_{\tau \in [t]} \bar{c}_i(\theta_\tau, z_\tau) \leq B_i^{\max}, \forall i \in [N] \right\}$$

$$\mathcal{E}_t^2 := \left\{ \sum_{\tau=1}^t \bar{c}_i(\theta_\tau, z_\tau) \leq B_i S_t^i, \forall i \in [N] \right\}$$

$$\mathcal{E}_t := \mathcal{E}_t^1 \cap \mathcal{E}_t^2,$$

where $t \in [T]$. Furthermore, we let

$$\Delta \mathcal{E}_t^1 := \{ \bar{c}_i(\theta_t, z_t) \leq B_i^{\max}, \forall i \in [N] \}$$

Notice that $\mathcal{E}_{t+1}^1 = \mathcal{E}_t^1 \cap \Delta \mathcal{E}_{t+1}^1$. We have

$$\begin{aligned} \mathbb{P}(\mathcal{E}_T) &\geq \mathbb{P}(\mathcal{E}_{T-1}) \mathbb{P}(\mathcal{E}_T | \mathcal{E}_{T-1}) \\ &= \mathbb{P}(\mathcal{E}_{T-1}) \mathbb{P}(\mathcal{E}_{T-1}^1 \cap \Delta \mathcal{E}_T^1 \cap \mathcal{E}_T^2 | \mathcal{E}_{T-1}) \\ &= \mathbb{P}(\mathcal{E}_{T-1}) \mathbb{P}(\Delta \mathcal{E}_T^1 \cap \mathcal{E}_T^2 | \mathcal{E}_{T-1}) \\ &= \mathbb{P}(\mathcal{E}_{T-1}) \mathbb{P}(\bar{c}_i(\theta_T, z_T) \leq B_{i,T}, \forall i \in [N] | \mathcal{E}_{T-1}) \\ &\geq \mathbb{P}(\mathcal{E}_{T-1})(1 - \epsilon_T), \end{aligned} \quad (16)$$

where the last equality follows by that

$$B_{i,T} \triangleq \min \left\{ B_i S_{i,T} - \sum_{\tau=1}^{T-1} \bar{c}_i(\theta_\tau, z_\tau), B_i^{\max} \right\}$$

conditioned on $\mathcal{E}_{T-1}$. By recursion, we have

$$\mathbb{P}(\mathcal{E}_T) \geq \mathbb{P}(\mathcal{E}_1) \prod_{t=2}^T (1 - \epsilon_t) \geq \prod_{t=1}^T (1 - \epsilon_t) = 1 - \delta,$$

which concludes the proof since $S_T^i = 1$. $\square$

Algorithm 1 Violation-Aware Contextual Bayesian Optimization

- 1: **Require:** VACBO horizons T , violation total budget $B_i, \forall i \in [N]$, maximum allowed violation cost for a single step $B_i^{\max}, \forall i \in [N]$, and an initial safe set of points $\mathcal{X}_0$
- 2: Evaluate $\ell(\theta, z), g_i(\theta, z), \forall i \in [N]$ for $(\theta, z) \in \mathcal{X}_0$ by performing experiments or simulation, or using historical data
- 3: Initialize dataset

$$D = \{(\theta_t, z_t), \ell(\theta_t, z_t), g(\theta_t, z_t) \mid (\theta_t, z_t) \in \mathcal{X}_0\}$$

- 4: **for** $t \in [T]$ **do**
- 5: Observe the exogenous variables z_t for step t
- 6: $B_{i,t} \triangleq \min \left\{ \max \left\{ B_i S_{i,t} - \sum_{\tau=1}^{t-1} \bar{c}_i(\theta_\tau, z_\tau), 0 \right\}, B_i^{\max} \right\}$
- 7: $\Theta_t = \left\{ \prod_{i \in [N]} \mathbb{P}(\bar{c}_i(\theta, z_t) \leq B_{i,t} | D) \geq 1 - \epsilon_t \mid \theta \in \Theta \right\}$
- 8: $\theta_t^* = \arg \max_{\theta \in \Theta_t} \text{CPEI}(\theta, z_t | D) \quad \triangleright \text{Solving (14)}$
- 9: $\ell(\theta_t^*, z_t), g(\theta_t^*, z_t) \leftarrow$ perform experiment with θ_t^* under the context z_t
- 10: Update dataset,

$$D \leftarrow D \cup \{(\theta_t^*, z_t), \ell(\theta_t^*, z_t), g(\theta_t^*, z_t)\}$$

- 11: Update Gaussian process posterior
-

4. Case Study: Constrained VCS Optimization

In this section, we apply the violation-aware contextual Bayesian optimization framework to safely tune the set-points of a vapor compression system (VCS). As shown in Fig. 2, a VCS typically consists of a compressor, a condenser, an expansion valve, and an evaporator. While physics-based models of these systems can be formulated as large sets of nonlinear differential algebraic equations to predict electrical power consumption, there are a variety of challenges in developing and calibrating these models. This motivates interest in directly using measurements of the power under different operating conditions to search for optimal set-points to the VCS actuators using data-driven, black-box optimization methods such as BO, to minimize the power consumption.

During the tuning process, one constraint, considered here, is on the temperature of the refrigerant leaving the compressor, also referred to as the ‘discharge temperature’ (see Fig. 2). The discharge temperature must be managed because compressors are designed to operate within specific temperature ranges; excessively high temperatures can result in the breakdown of refrigerant oils and increase wear and tear. In addition, high temperatures are often correlated with high pressures, which can cause metal fatigue and compromise the integrity of the pressurized refrigerant pipes in extreme cases. While managing the constraints mentioned above in the long run is critical, we also observe that small violations over short periods of time have limited




Figure 2: Schematic diagram of vapor compression system with our proposed VACBO controlling the EEV (electronic expansion valve) and the two fans' speed. For simplicity, we do not show other measurements and controls.

harmful effects. Indeed, it may be beneficial to take the risk of short-period limited violation to accelerate the tuning process.

Meanwhile, another challenge of tuning comes from the time-varying ambient conditions. Two major ambient conditions that have significant impact on the performance of vapor compression systems are the ambient temperature and the ambient humidity. The feasibility and optimality of VCS setpoints may be strongly correlated with these ambient conditions. For example, one set of fixed setpoints may lead to a feasible discharge temperature with low ambient temperature, but result in discharge temperature violation when the ambient temperature becomes high. It is necessary to adapt the set-points to the ambient conditions.

We cast the VCS optimization problem in the same form as (4), with $\ell((\theta, z))$ denoting the steady-state power of the VCS with set points θ and exogenous variables z . The constraint $g \leq 0$ is given by $T_d((\theta, z)) - \hat{T}_d \leq 0$, where $T_d((\theta, z))$ is the steady-state discharge temperature with set points θ and exogenous variables z , and $\hat{T}_d$ is a safe upper bound. We close a feedback loop from compressor frequency to room temperature, leaving the set of 3 tunable set points θ as the expansion valve position and the indoor/outdoor fan speeds. We set the exogenous variables to be ambient temperature and the ambient humidity, which are the two major ambient factors influencing the performance of vapor compression systems. The effects of these set points and exogenous variables on power and discharge temperature are not easy to model, and no simple closed-form representation exists. In practice, we assign a setpoint θ under contexts z , wait for an adequate amount of time until the power signal

resides within a 95% settling zone, and use that power value as $\ell((\theta, z))$ and the corresponding discharge temperature as $T_d((\theta, z))$.

Implementation Details

We use a high-fidelity model of the dynamics of a prototype VCS³ written in the Modelica language [27] to collect data and optimize the set-points on-the-fly. A complete model description is available in [9]. The model was first developed in the Dymola [10] environment, and then exported as a functional mockup unit (FMU) [28]. Its current version comprises 12,114 differential equations. We sample the system state each second. To leave enough time for the system state to converge to the steady state, we update the set-points every 180s. Bayesian optimization is implemented in GPY [15].

We define our set-point search space $\Theta := [200, 350] \times [300, 450] \times [500, 850]$, in expansion valve counts, indoor fan rpm (revolutions per minute), and outdoor fan rpm, respectively. We aim to keep the discharge temperature below $\hat{T}_d = 333$ K; these constraints are set according to domain knowledge [7]. We initialize the simulator at an expansion valve position of 340 counts, an indoor fan speed of 440 rpm, and an outdoor fan speed of 840 rpm, which is known to be a feasible set-point based on experience.

Constraint violations are penalized with the function $c_i(s) = s^2$, $s \in \mathbb{R}^+$. The quadratic nature of the violation cost implies that minor violations are not as heavily

³Note that while the behavior of this model has been validated against a real VCS, the numerical values and/or performance presented in this work are not representative of any product.

penalized as larger ones. The reason for this is that small violations over a small period of time are unlikely to prove deleterious to the long-term health of the VCS, whereas large violations could have more significant effects, even over short periods of time; for instance, damage to motor winding insulation or exceeding mechanical limits on the pressure vessel of the compressor. These constraints have been incorporated into the selection of the thresholds $\hat{T}_d$. Of course, the threshold values and a parameterized violation cost could be considered to be hyperparameters, and could be optimized via further experimentation. We choose the RBF kernel for our problem, which is commonly used in Bayesian optimization [12], and compare our method to two other state-of-the-art BO methods, namely, safe BO [6, 11] and generic constrained BO (cBO) [14]. To ensure a fair comparison, we also extend the other two state-of-the-art BO methods to the contextual setting. More specifically, we augment the input space of Gaussian processes with exogenous variables and inherit the algorithms of both safe BO and generic cBO. Based on our domain experience and prior knowledge, we set the kernel variance to be 15.0 (2.0, respectively), the kernel lengthscales in control parameters to be [50, 60, 70] ([20, 24, 28], respectively) for the objective (constraint) and the kernel lengthscales in exogenous variables to be [1.0, 0.06] ([1.0, 0.06], respectively) for the objective (constraint).

We showcase the effect of varying $B_i, i = 1$ by selecting $B_1 = 0, 10$, and 20 , and set $S_{1,t} = a_i + b_i \frac{t}{T}$, where $a_i + b_i = 1$ and $a_i \geq 0$ and $b_i \geq 0$. This choice allows the algorithm to use an increasing fraction of the total budget minus the violation cost incurred in the previous steps when approaching the sampling limit T . Intuitively, such a choice of budget sequence gradually increases the violation risk level without sampling too aggressively in the initial several steps and allows us to reduce the budget based on the violation cost in the previous steps. Proposition 1 gives a conservative way of choosing ϵ_t . In this case study, setting ϵ_t to a small constant 0.01 works well. We use “VACBO B ” to indicate violation aware contextual BO with budget $B_1 = B$.

Generation of Context Sequences. We consider both artificial recurring contexts and real-world contexts. In Sec. 4.1, we use artificially generated recurring contexts to showcase the effectiveness of our algorithm in optimizing the power consumption while managing the discharge temperature violation well. In Sec. 4.2, we apply our algorithm to more realistic setting with real-world historic context sequences measured in Zurich, Switzerland.

4.1. Artificial Recurring Contexts

To showcase the performance improvement brought by our method, we first apply an artificial sequence of recurring contexts. We periodically choose context z_t from a predefined context list $(z_i^p)_{i \in [n]}$ with additive Gaussian noise. Fig. 3 shows the change of the two exogenous variables with respect to time in our experiment.




Figure 3: Recurring contexts used in our experiment.

Results and Discussion

Fig. 4 illustrates the experimental results of violation-aware contextual Bayesian optimization compared to a fixed solution, which fixes the tuning parameters to default feasible values. The fixed solution can maintain the feasibility of discharge temperature constraint all the time. However, without any set-point exploration and optimization, the fixed solution maintains a relatively high operating power all the time. In contrast, our VACBO algorithm strategically explores the parameter space and optimizes the power function, with only small short-term tolerable violation. With more and more samples collected, the VACBO algorithm achieves lower and lower power, significantly reducing the energy consumption as compared to the fixed-parameter solution.

We then compare our VACBO algorithm to two state-of-the-art Bayesian optimization methods with constraints, safe Bayesian optimization [38] and generic constrained Bayesian optimization [14, 13]. Fig. 5 gives the comparison of average power and the maximum discharge temperature violation. Our VACBO method reduces the average power by about 12.2% as compared to fixing the setpoints, while managing the discharge temperature well as shown in Fig. 4. In comparison, generic constrained BO incurs discharge temperature that violates the constraint by as large as more than 150K, which is very dangerous for the system. As compared to the fixed solution, safe BO shows a very minor improvement in terms of the average power.




Figure 4: The evolution of power and discharge temperature with respect to time under recurring contexts.

Table 1

Average solving time of the auxiliary problem per step for different algorithms. All the results are obtained by running over a computer with one Intel64 Family 6 Model 142 Stepping 12 GenuineIntel 1803 Mhz processor and 16.0 GB RAM.

Algorithms	CEI	Safe BO	VACBO
Solving time/seconds	97.3	28.7	11.5

Computational time. Tab. 1 shows the solving time per step for different algorithms. It can be observed that our algorithm uses the least amount of time to solve the constrained auxiliary problem in each step. More importantly, the amount of time VACBO uses in each step is only 11.5 seconds, which is one order of magnitude less than the time between two setpoints (i.e., 180 seconds). This implies that our algorithm is computationally feasible for this setpoint tuning problem.




Figure 5: Comparison of average power and maximum constraint violation.

4.2. Real-world Contexts

To further demonstrate the effectiveness of our method under real-world contexts, we use the ambient temperature and the ambient humidity in Zurich, Switzerland as the contexts.




Figure 6: The ambient temperature and the ambient humidity in Zurich, Switzerland, starting at 1:30 PM on 2nd, July, 2019.

Fig. 6 shows the context data used for the real-world contexts. Fig. 7 shows the evolution of power and discharge temperature with respect to time under real-world contexts for a set of fixed set-points, the generic constrained Bayesian optimization method (CEI [14, 13]), and VACBO.

Similar to the results under recurring contexts, without exploration and optimization of the set-point parameter space, the fixed solution keeps operating with high power. We further show the average power and maximum discharge temperature violation in Fig. 8. Again, VACBO achieves significant power reduction compared to fixed set-point solution and other state-of-the-art method. Furthermore, our method incurs small and tolerable constraint violation ($< 0.5K$), in sharp contrast to the significant violation ($\geq 2.0K$) of generic contextual constrained BO (CEI method).




Figure 7: The evolution of power and discharge temperature with respect to time under real-world contexts.

5. Conclusions

In this paper, we design a sample-efficient and violation-aware contextual Bayesian optimization (VACBO) algorithm to solve the closed-loop control performance optimization problem with unmodeled constraints and time-varying exogenous factors, by leveraging the fact that small violations over a short period only incur limited costs during the optimization process in many applications such as for




Figure 8: The average power and maximum discharge temperature violation for different methods.

vapor-compression cycles. We strategically trade the budgeted violations for faster convergence by solving a tractable auxiliary problem with probabilistic budget constraints at each step. Our experiments on a VCS show that, as compared to existing safe BO and generic constrained BO, our method simultaneously exhibits improved optimization performance and manages the violation cost well.

Acknowledgements

This research was supported by the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40_180545, and in part by the Swiss Data Science Center, grant agreement C20-13.

Declaration of competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Wenjie Xu: <Credit authorship details>. **Colin N Jones:** . **Bratislav Svetozarevic:** . **Christopher R. Laughman:** . **Ankush Chakrabarty:** .

References

- [1] Baheri, A., Vermillion, C., 2017. Altitude optimization of airborne wind energy systems: A Bayesian optimization approach, in: 2017 Amer. Control Conf. (ACC), IEEE. pp. 1365–1370.
- [2] Baheri, A., Vermillion, C., 2020. Waypoint optimization using Bayesian optimization: A case study in airborne wind energy systems, in: 2020 Amer. Control Conf. (ACC), IEEE. pp. 5102–5017.
- [3] Bansal, S., Calandra, R., Xiao, T., Levine, S., Tomlin, C.J., 2017. Goal-driven dynamics learning via Bayesian optimization, in: 2017 IEEE 56th Annu. Conf. on Decis. and Control (CDC), IEEE. pp. 5168–5173.
- [4] Beckers, T., Bansal, S., Tomlin, C.J., Hirche, S., 2019. Closed-loop model selection for kernel-based models using Bayesian optimization, in: 2019 IEEE 58th Conf. on Decis. and Control (CDC), IEEE. pp. 828–834.

- [5] Ben-Tal, A., El Ghaoui, L., Nemirovski, A., 2009. Robust optimization. volume 28. Princeton university press.
- [6] Berkenkamp, F., Schoellig, A.P., Krause, A., 2016. Safe controller optimization for quadrotors with Gaussian processes, in: 2016 IEEE Int. Conf. on Robot. and Automat. (ICRA), IEEE. pp. 491–496.
- [7] Burns, D.J., Laughman, C.R., Guay, M., 2018. Proportional-integral extremum seeking for vapor compression systems. IEEE Trans. Control Syst. Technol. 28, 403–412.
- [8] Chakrabarty, A., Benosman, M., 2021. Safe learning-based observers for unknown nonlinear systems using Bayesian optimization. Automatica 133, 109860.
- [9] Chakrabarty, A., Danielson, C., Bortoff, S.A., Laughman, C.R., 2021. Accelerating self-optimization control of refrigerant cycles with Bayesian optimization and adaptive moment estimation. Appl. Therm. Eng. 197, 117335. doi:<https://doi.org/10.1016/j.applthermaleng.2021.117335>.
- [10] Dassault Systemes, 2019. Dymola 2020.
- [11] Fiducioso, M., Curi, S., Schumacher, B., Gwerder, M., Krause, A., 2019. Safe contextual Bayesian optimization for sustainable room temperature PID control tuning, in: Proc. 28th Int. Joint Conference on Artificial Intelligence (IJCAI), pp. 5850–5856.
- [12] Frazier, P.I., 2018. A tutorial on Bayesian optimization. arXiv preprint [arXiv:1807.02811](https://arxiv.org/abs/1807.02811).
- [13] Gardner, J.R., Kusner, M.J., Xu, Z.E., Weinberger, K.Q., Cunningham, J.P., 2014. Bayesian optimization with inequality constraints., in: ICML, pp. 937–945.
- [14] Gelbart, M.A., Snoek, J., Adams, R.P., 2014. Bayesian optimization with unknown constraints, in: Proc. of the 30th Conf. on Uncertainty in Artif. Intell., AUAI Press, Arlington, Virginia, USA. p. 250–259.
- [15] GPy, since 2012. GPy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPy>.
- [16] Grossmann, I.E., Halemane, K.P., Swaney, R.E., 1983. Optimization strategies for flexible chemical processes. Computers & Chemical Engineering 7, 439–462.
- [17] Halemane, K.P., Grossmann, I.E., 1983. Optimal process design under uncertainty. AIChE Journal 29, 425–433.
- [18] Jeong, S., Murayama, M., Yamamoto, K., 2005. Efficient optimization design method using kriging model. Journal of aircraft 42, 413–420.
- [19] Jones, D.R., Schonlau, M., Welch, W.J., 1998. Efficient global optimization of expensive black-box functions. J. Global Optim. 13, 455–492.
- [20] Khosravi, M., Eichler, A., Schmid, N., Smith, R.S., Heer, P., 2019. Controller tuning by Bayesian optimization an application to a heat pump, in: 2019 18th Eur. Control Conf. (ECC), IEEE. pp. 1467–1472.
- [21] König, C., Khosravi, M., Maier, M., Smith, R.S., Rupenyan, A., Lygeros, J., 2020. Safety-aware cascade controller tuning using constrained Bayesian optimization. arXiv preprint [arXiv:2010.15211](https://arxiv.org/abs/2010.15211).
- [22] Krause, A., Ong, C., 2011. Contextual Gaussian process bandit optimization. Advances in Neural Information Processing Systems 24.
- [23] Lam, R., Willcox, K., Wolpert, D.H., 2016. Bayesian optimization with a finite budget: An approximate dynamic programming approach. Advances in Neural Inf. Process. Syst. 29, 883–891.
- [24] Lee, E.H., Perrone, V., Archambeau, C., Seeger, M., 2020. Cost-aware Bayesian optimization. arXiv preprint [arXiv:2003.10870](https://arxiv.org/abs/2003.10870).
- [25] Liu, H., Ong, Y.S., Shen, X., Cai, J., 2020. When gaussian process meets big data: A review of scalable gps. IEEE transactions on neural networks and learning systems 31, 4405–4423.
- [26] Marco, A., von Rohr, A., Baumann, D., Hernández-Lobato, J.M., Trimpe, S., 2020. Excursion search for constrained Bayesian optimization under a limited budget of failures. arXiv preprint [arXiv:2005.07443](https://arxiv.org/abs/2005.07443).
- [27] Modelica Association, 2017. Modelica specification, Version 3.4 URL: www.modelica.org.
- [28] Modelica Association, 2019. Functional Mockup Interface for Model Exchange and Co-Simulation, Version 2.0.1 URL: www.fmi-standard.org.
- [29] Pal, A., Zhu, L., Wang, Y., Zhu, G.G., 2020. Multi-objective stochastic Bayesian optimization for iterative engine calibration, in: Proc. of the Amer. Control Conf., IEEE. pp. 4893–4898.
- [30] Park, J., 2020. Contextual Bayesian optimization with trust region (CBOTR) and its application to cooperative wind farm control in region 2. Sustainable Energy Technologies and Assessments 38, 100679.
- [31] Paulson, J.A., Mesbah, A., 2020. Data-driven scenario optimization for automated controller tuning with probabilistic performance guarantees. IEEE Contr. Syst. Lett. 5, 1477–1482.
- [32] Piga, D., Forgone, M., Formentin, S., Bemporad, A., 2019. Performance-oriented model learning for data-driven MPC design. IEEE Contr. Syst. Lett. 3, 577–582.
- [33] Shang, C., You, F., 2019. A data-driven robust optimization approach to scenario-based stochastic model predictive control. Journal of Process Control 75, 24–39.
- [34] da Silva, W.G., Acarnley, P.P., Finch, J.W., 2000. Application of genetic algorithms to the online tuning of electric drive speed controllers. IEEE Trans. Ind. Electron. 47, 217–219.
- [35] Snelson, E., Ghahramani, Z., 2005. Sparse gaussian processes using pseudo-inputs. Advances in neural information processing systems 18.
- [36] Snoek, J., Larochelle, H., Adams, R.P., 2012. Practical Bayesian optimization of machine learning algorithms. Advances in Neural Inf. Process. Syst. 25.
- [37] Sui, Y., Burdick, J., Yue, Y., et al., 2018. Stage-wise safe Bayesian optimization with Gaussian processes, in: Proc. of the Int. Conf. on Mach. Learn., pp. 4781–4789.
- [38] Sui, Y., Gotovos, A., Burdick, J., Krause, A., 2015. Safe exploration for optimization with Gaussian processes, in: Proc. of the Int. Conf. on Mach. Learn., pp. 997–1005.
- [39] Telen, D., Vallerio, M., Cabianca, L., Houska, B., Van Impe, J., Logist, F., 2015. Approximate robust optimization of nonlinear systems under parametric uncertainty and process noise. Journal of Process Control 33, 140–154.
- [40] Turchetta, M., Berkenkamp, F., Krause, A., 2020. Safe exploration for interactive machine learning. Advances in Neural Inf. Process. Syst. 32, 2868–2878.
- [41] Wächter, A., Biegler, L.T., 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Mathematical programming 106, 25–57.
- [42] Xu, W., Jiang, Y., Jones, C.N., 2022a. Constrained efficient global optimization of expensive black-box functions. arXiv preprint [arXiv:2211.00162](https://arxiv.org/abs/2211.00162).
- [43] Xu, W., Jiang, Y., Maddalena, E.T., Jones, C.N., 2022b. Lower bounds on the worst-case complexity of efficient global optimization. arXiv preprint [arXiv:2209.09655](https://arxiv.org/abs/2209.09655).
- [44] Xu, W., Jiang, Y., Svetozarevic, B., Jones, C.N., 2022c. CONFIG: Constrained efficient global optimization for closed-loop control system optimization with unmodeled constraints. arXiv preprint [arXiv:2211.11822](https://arxiv.org/abs/2211.11822).
- [45] Xu, W., Jones, C.N., Svetozarevic, B., Laughman, C.R., Chakrabarty, A., 2022d. VABO: Violation-aware Bayesian optimization for closed-loop control performance optimization with unmodeled constraints, in: Proc. American Control Conference, pp. 5288–5293.