

ECCBO: An Inherently Safe Bayesian Optimization with Embedded Constraint Control for Real-Time Optimization^{*}

Dinesh Krishnamoorthy

*Department of Mechanical Engineering, Eindhoven University of
Technology, 5600 MB, Eindhoven, The Netherlands.
(email: d.krishnamoorthy@tue.nl)*

Abstract:

This paper introduces a model-free real-time optimization (RTO) framework based on unconstrained Bayesian optimization with embedded constraint control. The main contribution lies in demonstrating how this approach simplifies the black-box optimization problem while ensuring “always-feasible” setpoints, addressing a critical challenge in real-time optimization with unknown cost and constraints. Noting that controlling the constraint does not require detailed process models, the key idea of this paper is to control the constraints to “some” setpoint using simple feedback controllers. Bayesian optimization then computes the optimum setpoint for the constraint controllers. By searching over the setpoints for the constraint controllers, as opposed to searching directly over the RTO degrees of freedom, this paper achieves an inherently safe model-free steady-state RTO. In particular, this paper shows that the proposed approach can achieve zero cumulative constraint violation without relying on assumptions about the Gaussian process model used in Bayesian optimization. The effectiveness of the proposed approach is demonstrated on a benchmark Williams-Otto reactor example.

Keywords: Real-time optimization, data-driven optimization, constraint control, Bayesian optimization

1. INTRODUCTION

Real-time optimization (RTO) plays a pivotal role in process operations due to its ability to adapt to changing operating conditions. It ensures that processes can efficiently respond to changing factors like fluctuating demand or supply, thereby optimizing production processes in real-time. More importantly, it helps maintain process safety and desired product qualities. The standard approach to real-time optimization involves developing detailed process models that are updated using measurements corresponding to steady-state process operation. The updated process models are then used within an optimization problem to compute the optimal setpoints (Chen and Joseph, 1987; Darby et al., 2011). However, developing and maintaining detailed process models can be a major hurdle that impedes practical implementation of RTO in industrial practice.

Model-free approach to online process optimization is an attractive alternative that circumvents the need for rigorous process models. There have been several attempts to develop model-free RTO schemes, most of which are based on online gradient estimation. The key underlying idea behind these approaches is to estimate the steady-state cost gradient directly from the process measurements, and drive it to zero using integral action, such that the necessary condition of optimality is asymptotically satisfied. See

^{*} Financial support from the NWO Veni Early Career Talent Scheme is gratefully acknowledged.

for example Garcia and Morari (1981); Golden and Ydstie (1989); Bamberger and Isermann (1978); François et al. (2005); Dochain et al. (2011). However, estimating the steady-state cost gradient often leads to prohibitively slow convergence, especially for process with long settling times (Srinivasan et al., 2011; Krishnamoorthy and Skogestad, 2022). Zeroth-order model-free optimization schemes such as Bayesian optimization is a promising alternative that avoids the need to estimate steady-state cost gradients.

Bayesian optimization is an effective method for optimizing unknown systems that finds the optimum by iteratively interacting with the system. BO uses a probabilistic surrogate model for the cost function, e.g., Gaussian process (GP), which is updated with each new observation. An acquisition function (also known as utility function) is induced from the GP surrogate model that guides the search for the optimum by systematically trading-off exploration and exploitation. Unlike gradient-based methods, Bayesian optimization calculates the next action independently of the current one, allowing for the possibility of significant jumps in the search space. While this notion of exploration speeds up the search for the global optimum, it can also lead to the exploration of infeasible setpoints that may violate safety-critical constraints at steady-state.

Constrained Bayesian optimization is still an active area of research, especially for engineering applications (Berkenkamp et al., 2021; Lu and Paulson, 2022; Krishnamoorthy and Doyle, 2022; Xu et al., 2022). The classical

approach to handling constraints is to scale the acquisition function with the probability of constraint feasibility computed from the GP surrogates for the constraints (Gelbert et al., 2012). Although such an approach would eventually converge to a feasible setpoint, it could potentially explore an infeasible setpoint. This is because the probability of constraint (in)feasibility at any point will be known with high probability only after observing the constraint at that point. As such, this approach does not have any guarantees on the cumulative violation of the constraints. Another approach, known as Safe Bayesian optimization, ensures no constraint violation with high probability (Berkenkamp et al., 2021; Krishnamoorthy and Doyle, 2022; Krishnamoorthy and Doyle III, 2023) under the assumption that the constraint GPs are well calibrated (i.e., the true function is contained within the confidence intervals with high probability). Satisfying such an assumption in practice can often result in an overly cautious algorithm, since the safe set may expand very slowly. A third class of approach uses penalty functions, where the constraint violation is within a prescribed budget, i.e. bounded cumulative constraint violation (Lu and Paulson, 2022; Xu et al., 2022).

To this end, all the existing approaches to constrained BO require probabilistic surrogate models not just for the cost but also for all the unknown constraints. Safety guarantees and /or cumulative constraint violation bounds only hold under the assumption that the GP models for the constraints are well calibrated. Simply put, calibrating the hyperparameters of the constraint GPs can impact the constraint handling capability. Furthermore, the different approaches referenced above also assume that the constraints are independent, and uses individual GP surrogate models for each unknown constraint.

In the context of real-time process optimization, the information flow can be vertically decomposed, where the upper layer uses the setpoints to the lower layer as degrees of freedom, and the lower layer implements the actions to achieve the setpoint. Based on the fact that controlling the constraint to a setpoint does not require detailed process models, the key idea of this paper is to exploit decentralized feedback controllers (such as PID) to control the constraints to “some” setpoint, and use Bayesian optimization to find the optimal setpoints to the lower level constraint controllers. If a constraint is optimally active, then the setpoint computed by Bayesian optimization will converge to the limiting value. Whereas, if a constraint is not optimally active, then the Bayesian optimization will converge to some feasible setpoint that optimizes the cost. By transforming the original optimization problem with embedded constraint controllers, BO now searches over the setpoint space, as opposed to searching directly over the RTO degrees of freedom. As such, the decision variables for the BO now only have simple box constraints, which can be trivially handled using standard Bayesian optimization, and the box constraints inherently ensure that the RTO layer will never compute an infeasible setpoint.

To this end, this paper presents an inherently safe Bayesian optimization with constraint control embedded (ECCBO) that achieves zero cumulative constraint violation, without relying on assumptions about the Gaussian process model used in Bayesian optimization. Since the focus of the paper is on the steady-state RTO layer, “safety” is defined w.r.t

the constraints at steady-state operation. The remainder of the paper is organized as follows: Section 2 recalls the preliminaries of RTO, constraint control, and Bayesian optimization. The proposed inherently safe Bayesian optimization algorithm with embedded constraint control (EC-CBO) is presented in Section 3. Section 4 demonstrates the proposed approach using a benchmark Williams-Otto reactor case study, before concluding the paper in Section 5.

2. PRELIMINARIES

Real-time optimization and constraint control: Consider the steady-state real-time optimization problem

$$\min_{\mathbf{u} \in \mathcal{U}} J(\mathbf{u}, \mathbf{d}) \quad (1a)$$

$$\text{s.t. } g_i(\mathbf{u}, \mathbf{d}) \geq 0, i \in \mathbb{I}_{1:n} \quad (1b)$$

where $\mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}^m$ denotes the set of decision variables for the RTO layer, $\mathbf{d} \in \mathbb{R}^l$ denotes the set of disturbances, $J : \mathcal{U} \times \mathbb{R}^l \rightarrow \mathbb{R}$ denotes the cost function, and $g_i : \mathcal{U} \times \mathbb{R}^l \rightarrow \mathbb{R}$ for $i = 1, \dots, n$ denotes the set of n constraints. The input constraint set $\mathcal{U} = \mathcal{U}_1 \times \dots \times \mathcal{U}_m$ is a simple set composed of upper and lower bounds. Typically, the degrees of freedom \mathbf{u} for the RTO layer are provided as setpoints to the regulatory control layer below.

For a process with n constraints, there can be utmost 2^n possible combinations of active constraints. The specific set of optimally active constraints, denoted as $\mathbf{g}_A \subseteq \mathbf{g} := [g_1, \dots, g_n]^T$, may change depending on the operating conditions. If a constraint is known to be optimally active, indicated by $g_i \in \mathbf{g}_A$ and $g_i^* = 0$, it can be controlled to a setpoint of its limiting value (which is 0 when expressed in the positive null form as done in (1)) using simple feedback controllers like PID, a technique known as active constraint control. However, for constraints that are not optimally active, denoted as $g_j \notin \mathbf{g}_A$, determining their non-trivial optimal values $g_j(\mathbf{u}^*, \mathbf{d}) = z_j^* > 0$ becomes challenging.

Typically, the RTO layer uses detailed process models to solve an optimization problem numerically to identify the set of active constraints and the optimal values $z_j^* > 0$ for the inactive constraints (assuming that the model used in the optimization solver captures the plant accurately). However, this can be challenging if detailed process models are not available, which is the case that is considered in this paper.

Bayesian Optimization: Bayesian optimization aims to optimize an unknown black-box function

$$\min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x})$$

where \mathcal{X} is a simple set. Bayesian optimization uses a probabilistic model, such as Gaussian processes as a surrogate for the unknown function $F(\mathbf{x})$.

$$F(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \quad (2)$$

where $\mu(\mathbf{x}) := \mathbb{E}(F(\mathbf{x}))$ denotes the mean and $k(\mathbf{x}, \mathbf{x}') := \mathbb{E}[(F(\mathbf{x}) - \mu(\mathbf{x}))(F(\mathbf{x}') - \mu(\mathbf{x}'))]$ denotes the covariance of the unknown function. Given noisy measurements y of the unknown function $F(\mathbf{x})$, the Gaussian process model is updated by conditioning on the observations $\{(\mathbf{x}_i, y_i)\}_i$ (Rasmussen and Williams, 2006). An acquisition function

$\alpha(\mathbf{x})$ is induced from the Gaussian process surrogate, which is then optimized to find the next action to take.

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \alpha_t(\mathbf{x}) \quad (3)$$

where $\alpha_t : \mathcal{X} \rightarrow \mathbb{R}$ is the acquisition function induced from the Gaussian process model conditioned on the past t measurements $\{(\mathbf{x}_i, y_i)\}_{i=1}^t$. There are several choices for acquisition functions such as expected improvement, lower confidence bound, Thompson sampling etc. See for example Frazier (2018) and the references therein for further details. The crucial idea here is that the acquisition function balances exploration and exploitation to find the optimum of the unknown function by sequentially interacting with the system.

Contextual Bayesian optimization: To account for variables that are not directly under our control, such as prices, feed rates etc., the Bayesian optimization problem can be extended to account for these external factors. Here the optimization problem is formulated as

$$\min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}, \mathbf{d})$$

where $\mathbf{d} \in \mathbb{R}^l$ denotes the external factors, also known as contexts. In this case, the GP model is used as a surrogate for $F(\mathbf{x}, \mathbf{d})$ by conditioning on the observations $\{([\mathbf{x}_i, \mathbf{d}_i]^\top, y_i)\}_i$. For a given contextual information \mathbf{d}_{t+1} at time step $t + 1$, the next action is computed by optimizing

$$\mathbf{x}_{t+1} = \arg \min_{\mathbf{x} \in \mathcal{X}} \alpha_t(\mathbf{x}, \mathbf{d}_{t+1}) \quad (4)$$

where $\alpha_t : \mathcal{X} \times \mathbb{R}^l \rightarrow \mathbb{R}$ is any suitable acquisition function derived from the GP surrogate conditioned on the past t observations. The literature on constrained Bayesian optimization that accounts for additional nonlinear constraints is not included, as these will not be utilized in this paper.

3. BAYESIAN OPTIMIZATION WITH EMBEDDED CONSTRAINT CONTROL (ECCBO)

Assumption 1. The cost and constraint functions $J(\cdot, \cdot)$ and $g_i(\cdot, \cdot)$ in (1) are unknown, but are measured in real-time.

The main objective is to find the optimum of the real-time optimization problem (1), under Assumption 1. In search for the optimum using Bayesian optimization, we want to ensure that no infeasible setpoints are explored.

In this paper, we propose to exploit constraint control to safely find the optimum of the black-box system. To do so, we transform the inequality constrained problem (1) to the following equality constrained optimization problem.

$$\min_{\mathbf{u}, \mathbf{z}} J(\mathbf{u}, \mathbf{d}) \quad (5a)$$

$$\text{s.t. } g_i(\mathbf{u}, \mathbf{d}) = z_i, i \in \mathbb{I}_{1:n} \quad (5b)$$

where $\mathbf{z} = [z_1, \dots, z_n]^\top$ denotes the optimum value of the constraints.

We use n degrees of freedom to tightly control the n constraints to their corresponding setpoints z_i using decentralized feedback control loops. The n degrees of freedom, and the n constraints are paired such that the steady-state relative gain array (RGA) of the resulting transfer matrix

is non-negative and close to identity matrix at crossover frequencies. In order to minimize large time delays it would be preferable to pair constraints that are physically close to the manipulated variable (pair-close rule). One must also try to avoid pairing important constraints with degree of freedom u_i with tight bounds \mathcal{U}_i (Krishnamoorthy and Skogestad, 2022).

We know that if a constraint g_i is optimally active, then $z_i = 0$, and if a constraint g_j is not optimally active, then we want $z_j > 0$. However, there are two challenges. Firstly, we do not know which of the constraints are optimally active, and which are not. Secondly, for the inactive constraints $g_j \notin \mathbf{g}_A$, we do not know the exact value of $z_j > 0$. This paper proposes to use unconstrained Bayesian optimization to find the optimal setpoints $z_i \in [0, \bar{z}_i]$ for $i = 1, \dots, n$ constraint controllers, and for the remaining unconstrained degrees of freedom u_{n+1}, \dots, u_m .

Assumption 2. (Perfect control). The constraint controllers are designed such that the desired setpoint $z_i \in [0, \bar{z}_i]$ is asymptotically achieved for all $i = 1, \dots, n$.

Since the constraints are tightly controlled to their corresponding setpoints z_i , we can equivalently rewrite (5) as

$$\min_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}, \mathbf{d}) = \min_{\mathbf{u} \in \mathcal{U}, \mathbf{z} \in \mathcal{Z}} J(\mathbf{u}, \mathbf{d}) \quad (6)$$

$$\text{s.t. } g_i(\mathbf{u}, \mathbf{d}) = z_i, i \in \mathbb{I}_{1:n}$$

where $\mathbf{x} := [z_1, \dots, z_n, u_{n+1}, \dots, u_m]^\top$. Let us define $\mathcal{Z}_i = [0, \bar{z}_i]$, where \bar{z}_i is some arbitrarily large upper bound for the constraint set points. This leads to the constraint set $\mathcal{X} := \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n \times \mathcal{U}_{n+1} \times \dots \times \mathcal{U}_m$.

We can see that (6) is an unconstrained black-box optimization problem, which can be solved using standard contextual Bayesian optimization. Here, we use a Gaussian process as a surrogate for the cost function $F(\mathbf{x}, \mathbf{d})$ as a function of \mathbf{x} and \mathbf{d} . The GP is conditioned on the cost observed with the embedded constraint controllers that control g_i to a setpoint of z_i using u_i as the manipulated variable.

Remark 1. (Steady-state wait-time). Note that the Gaussian process model used in the Bayesian optimization algorithm is updated based on steady-state cost measurement. As such, the BO-based RTO scheme requires a steady-state detection algorithm to ensure that GP models are conditioned on steady-state data.

Let us define the violation of the i^{th} constraint at steady-state as $v_{i,t} = -[g_i(\mathbf{u}_t, \mathbf{d}_t)]^-$ and the cumulative violation over T RTO iterations as

$$\mathcal{V}_T = \sum_{i=1}^n \sum_{t=1}^T v_{i,t}$$

Theorem 1. Under Assumption 2, the ECCBO framework achieves a cumulative violation of $\mathcal{V}_T = 0$ for any acquisition function.

Proof. The Bayesian optimization using any acquisition function by design searches over the setpoints $z_i \in [0, \bar{z}_i]$ for $i = 1, \dots, n$. By controlling the constraints to a setpoint $z_i \geq 0$ using the degrees of freedom u_i , the constraint violation at steady-state $v_{i,t} = 0$ for all $t \geq 1$, for all $i = 1, \dots, n$, which results in $\mathcal{V}_T = 0$.

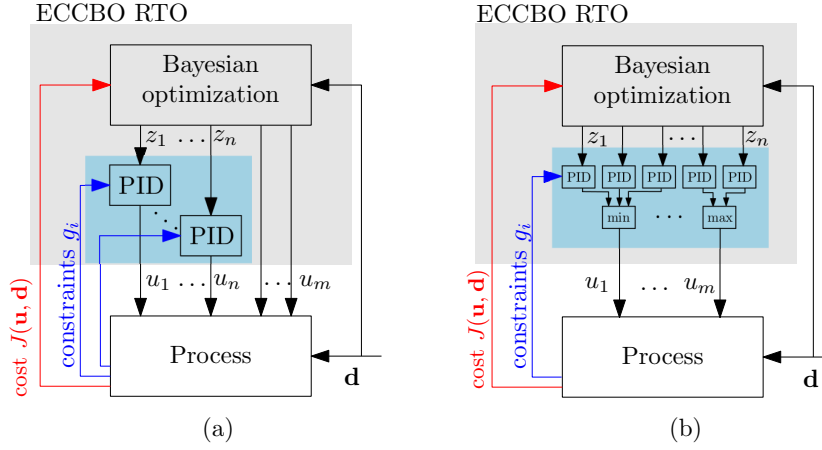


Fig. 1. Proposed real-time optimization scheme using Bayesian optimization highlighted in gray, with embedded constraint control (ECCBO) highlighted in blue. The process block shown here includes the lower level setpoint control layer and the plant. (a) More degrees of freedom than constraint $m > n$. (b) More constraints than degrees of freedom $m < n$. Note that the process block may further contain lower level regulatory controllers.

Denoting F^* as the minimum of (6) for a given disturbance \mathbf{d} , querying an action \mathbf{x}_t at time t may lead to a gap at steady-state, which we define as the steady-state regret $r_t = F(\mathbf{x}_t, \mathbf{d}) - F^*$. The cumulative regret over T RTO runs is then defined as $R_T = \sum_{t=0}^T r_t$.

Remark 2. (Regret). Since the Bayesian optimization in ECCBO reduces to a standard unconstrained Bayesian optimization, the regret bound for R_T follows from BO literature for specific acquisition functions under appropriate conditions (Chowdhury and Gopalan, 2017). Bounding the contextual regret is in general a challenging problem, as regret is evaluated in relation to the best action for each specific context. Consequently, the worst-case regret, e.g., when employing GP-LCB, to identify the optimal action \mathbf{x}^* can be determined after multiple instances of encountering a particular context. This is also reasonable in online process optimization of continuous processes, where the goal is optimize for the steady-state operating conditions, i.e. for a given \mathbf{d} . As such, the standard regret results from the BO literature will also hold for ECCBO.

Remark 3. (Overconstrained case $m < n$). If the number of degrees of freedom is less than the number of constraints, we can have utmost m constraints optimally active. In this case, max/min selector blocks can be used along with the n embedded constraint controllers. The choice of the selectors and the grouping of the constraint controllers for each selector block follows the selector design procedure as explained by Krishnamoorthy and Skogestad (2020). The setpoint to the n controllers will be determined by the unconstrained Bayesian optimization. That is, the decision variables for the transformed Bayesian optimization problem are $\mathbf{x} = [z_1, \dots, z_n]^T$. The is schematically shown in Fig. 1b.

The design procedure for the proposed model-free RTO scheme, known as ECCBO-RTO can be summarized as follows:


- Offline:
 - Pair each constraint, denoted as g_i , with an available degree of freedom using established pair-

ing rules from process control literature (Krishnamoorthy and Skogestad, 2022).

- Tightly tune the constraint controllers to minimize steady-state wait time, and at the same time ensuring clear time-scale separation from any lower level controllers.
- Build a Gaussian process to model the cost with respect to the contexts \mathbf{d} , the setpoints z_i to the constraint controllers $i = 1, \dots, n$, and any remaining unconstrained degrees of freedom u_{n+1}, \dots, u_m .
- Online:
 - If the process is operating at steady-state, update the Gaussian process model by conditioning on the past t observations of the cost measurement.
 - Induce an acquisition function $\alpha_t(\mathbf{x}, \mathbf{d})$ based on the updated Gaussian process model, where $\mathbf{x} := [z_1, \dots, z_n, u_{n+1}, \dots, u_m]^T$ contains the set of setpoints for the constraint controllers, as well as any remaining unconstrained degrees of freedom.
 - For a given context \mathbf{d}_{t+1} , optimize the acquisition function $\alpha(\mathbf{x}, \mathbf{d}_{t+1})$ to compute the actions \mathbf{x}_{t+1} , which is implemented on the process.

4. ILLUSTRATIVE EXAMPLE

This section demonstrates the proposed approach on a benchmark Williams-Otto reactor example, which converts raw materials A and B to useful products P and E, through a series of reactions



The feedrate F_A with pure A component is an external disturbance. The feedrate F_B and the reactor temperature T_r are the two degrees of freedom, that are used to optimize the process.

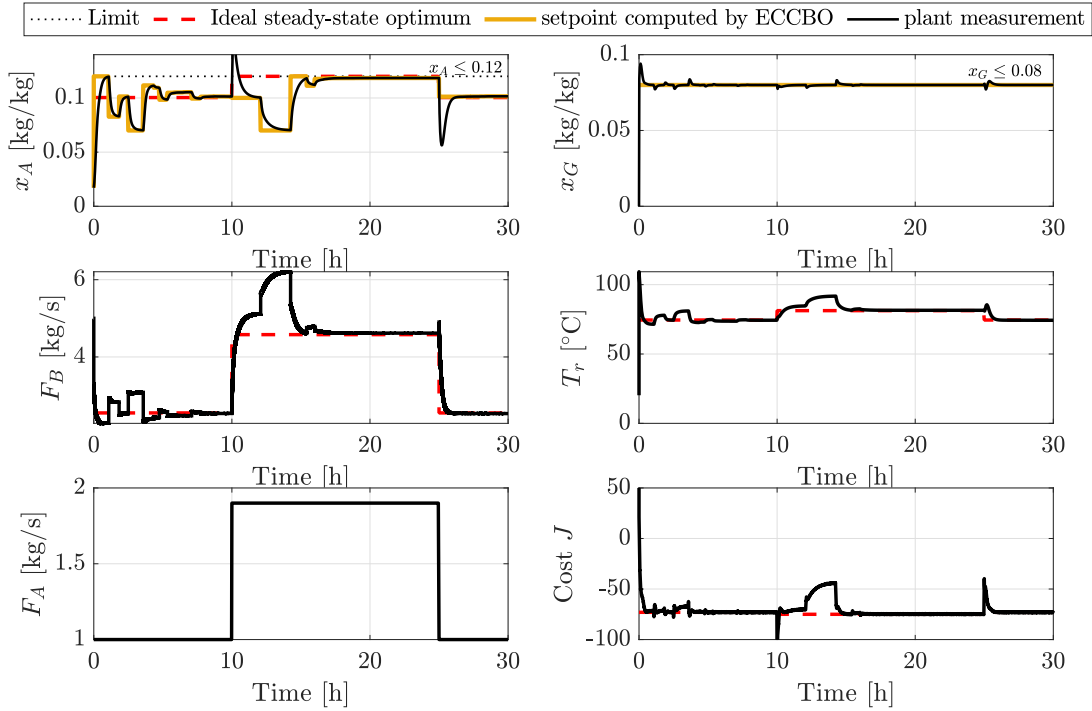


Fig. 2. Simulation results shown the performance of our proposed approach (solid black), compared with the true steady-state optimum (red dashed lines). The decision variable for the Bayesian optimization, namely the setpoint for the constraint controller is shown in yellow solid lines.

$$\begin{aligned}
 \max_{F_B, T_r} \quad & 1043.38x_P(F_A + F_B) + 20.92x_E(F_A + F_B) \\
 & - 79.23F_A - 118.34F_B \\
 \text{s.t.} \quad & x_G \leq 0.08 \\
 & x_A \leq 0.12
 \end{aligned} \quad (7)$$

To design an ECCBO-RTO for this problem, we employ two SISO constraint control loops.

- (1) Control the concentration of component G to a setpoint z_G using reactor temperature T_r .
- (2) Control the concentration of component A to a setpoint z_A using feed rate F_B .

Contextual Bayesian optimization is then used to find the optimal setpoints $z_G \in [0.07, 0.08]$ kg/kg and $z_A \in [0.07, 0.12]$ kg/kg, for the two control loops, respectively, for a given feed rate F_A . Simply put, the decision variables for the Bayesian optimization are z_G and z_A , and since the search space for this decision variables do not exceed 0.08 and 0.12, respectively, we can guarantee by design that the ECCBO-RTO layer will not compute any infeasible setpoints.

The plant measurements and the constraint controllers are sampled at a 1-second interval. For the two constraint control loops, the PI controllers are tuned using SIMC tuning rules. The GP surrogate for the cost used by the Bayesian optimization algorithm uses the RBF kernel along with a bias kernel, and the Bayesian optimization uses the lower confidence bound (GP-LCB) acquisition function¹. The Gaussian process is modeled using GPy

package (GPy, since 2012), and the acquisition function is optimized using L-BFGS algorithm from the SciPy library.

To initiate the Bayesian optimization routine, we employ a well-established steady-state detection algorithm, widely adopted in commercially available RTO tools. This algorithm is based on a statistical test that compares the total variance between consecutive signal points, ensuring the process operates at steady-state conditions as described by Cao and Rhinehart (1995); Câmara et al. (2016).

The simulation initiates with a feed rate of $F_A = 1$ kg/s, and changes to $F_A = 1.9$ kg/s at time $t = 10$ h, and back to $F_A = 1$ kg/s at time $t = 25$ h. The set of active constraint changes, as the feed rate F_A varies. The true steady-state optimum for the different conditions are shown in red dashed lines in Fig. 2. Initially, when the feed rate $F_A = 1$ kg/s, the constraint on x_A is not optimally active. Although this is unknown, the Bayesian optimization algorithm finds the optimum setpoint for x_A by sequentially interacting with the system. Since the search space for $z_A \in [0.07, 0.12]$, we can see that the setpoints are always below 0.12, thus always ensuring constraint feasibility. The setpoint z_A computed by the BO algorithm is shown in yellow in the top left subplot in Fig. 2. For feedrate $F_A = 1.9$ kg/s, the constraint on x_A will be optimally active. In this case, the Bayesian optimization algorithm converges to the limiting value of $z_A = 0.12$ kg/kg, without exploring any infeasible setpoints. Although at $t = 10$ h, the constraint is dynamically violated, the setpoint computed by ECCBO-RTO is feasible (i.e., no constraint violation at steady-state). Finally, when the feed rate reduces

¹ Traditional BO literature considers maximization problem, and hence uses the upper confidence bound (GP-UCB). The equivalent

acquisition function for a minimization problem is the lower confidence bound.

back to $F_A = 1\text{kg/s}$, the setpoint computed by the BO algorithm converges to the optimal setpoint in one step. This is because, the GP has previously seen the data for this “context”. Thanks to Gaussian processes conditioned on previously observed data, our algorithm converges to the optimum in one step. For this process, the purity constraint is very low and from engineering insight we know that this constraint will likely be optimally active. This can be incorporated in the GP prior, to avoid unnecessary explorations. This is reflected in the setpoint z_G .

In prior research, we assessed interior point-based safe Bayesian optimization and constrained Bayesian optimization on the same Williams-Otto reactor example. The findings illustrated in (Krishnamoorthy and Doyle III, 2023, Fig. 5) demonstrated that traditional constrained Bayesian optimization (Gelbert et al., 2012) frequently explored infeasible setpoints, operating far beyond safe limits until the next RTO update. This limitation makes it impractical for real-time optimization in systems with safety-critical constraints. In contrast, safe Bayesian optimization, depicted in (Krishnamoorthy and Doyle III, 2023, Fig. 4), avoided infeasible setpoints but required independent Gaussian process models for each constraint, which must be carefully calibrated. For the sake of brevity, these results are not reiterated in this paper, and readers are referred to Krishnamoorthy and Doyle III (2023).

5. CONCLUSION

In conclusion, this paper introduced a novel model-free Real-Time Optimization (RTO) scheme, ECCBO (Bayesian optimization with Embedded Constraint Control). By leveraging the simplicity of constraint control, and the power of Bayesian optimization, this approach addresses a critical challenge in model-free real time optimization of process systems with unknown constraints. The main advantage of ECCBO is that it offers a practical solution that ensures feasibility and safety using well established control tools within the process control industry. Specifically, we demonstrated that ECCBO achieves zero cumulative constraint violation without imposing assumptions on the GP model used in Bayesian optimization. Since ECCBO leverages the vertical decomposition with respect to the information flow in plant-wide control, it is important to note that this result is not generalizable to the broader constrained Bayesian optimization literature, where such lower level controllers are not implementable.

REFERENCES

- Bamberger, W. and Isermann, R. (1978). Adaptive on-line steady-state optimization of slow dynamic processes. *Automatica*, 14(3), 223–230.
- Berkenkamp, F., Krause, A., and Schoellig, A.P. (2021). Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *Machine Learning*, 1–35.
- Câmara, M.M., Quelhas, A.D., and Pinto, J.C. (2016). Performance evaluation of real industrial rto systems. *Processes*, 4(4), 44.
- Cao, S. and Rhinehart, R.R. (1995). An efficient method for on-line identification of steady state. *Journal of Process Control*, 5(6), 363–374.
- Chen, C.Y. and Joseph, B. (1987). On-line optimization using a two-phase approach: an application study. *Industrial & engineering chemistry research*, 26(9), 1924–1930.
- Chowdhury, S.R. and Gopalan, A. (2017). On kernelized multi-armed bandits. In *International Conference on Machine Learning*, 844–853. PMLR.
- Darby, M.L., Nikolaou, M., Jones, J., and Nicholson, D. (2011). RTO: An overview and assessment of current practice. *Journal of Process Control*, 21(6), 874–884.
- Dochain, D., Perrier, M., and Guay, M. (2011). Extremum seeking control and its application to process and reaction systems: A survey. *Mathematics and Computers in Simulation*, 82(3), 369–380.
- François, G., Srinivasan, B., and Bonvin, D. (2005). Use of measurements for enforcing the necessary conditions of optimality in the presence of constraints and uncertainty. *Journal of Process Control*, 15(6), 701–712.
- Frazier, P.I. (2018). A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*.
- Garcia, C.E. and Morari, M. (1981). Optimal operation of integrated processing systems. part i: Open-loop on-line optimizing control. *AIChE Journal*, 27(6), 960–968.
- Gelbert, G., Moeck, J.P., Paschereit, C.O., and King, R. (2012). Advanced algorithms for gradient estimation in one-and two-parameter extremum seeking controllers. *Journal of Process Control*, 22(4), 700–709.
- Golden, M.P. and Ydstie, B.E. (1989). Adaptive extremum control using approximate process models. *AIChE journal*, 35(7), 1157–1169.
- GPy (since 2012). GPy: A gaussian process framework in python. <http://github.com/SheffieldML/GPy>.
- Krishnamoorthy, D. and Doyle, F.J. (2022). Safe bayesian optimization using interior-point methods -applied to personalized insulin dose guidance. *IEEE Control Systems Letters*, 1–1. doi:10.1109/LCSYS.2022.3179330.
- Krishnamoorthy, D. and Doyle III, F.J. (2023). Model-free real-time optimization of process systems using safe bayesian optimization. *AIChE Journal*, 69(4), e17993.
- Krishnamoorthy, D. and Skogestad, S. (2020). Systematic design of active constraint switching using selectors. *Computers & Chemical Engineering*, 143, 107106.
- Krishnamoorthy, D. and Skogestad, S. (2022). Real-time optimization as a feedback control problem—a review. *Computers & Chemical Engineering*, 107723.
- Lu, C. and Paulson, J.A. (2022). No-regret bayesian optimization with unknown equality and inequality constraints using exact penalty functions. *IFAC-PapersOnLine*, 55(7), 895–902.
- Rasmussen, C.E. and Williams, C. (2006). *Gaussian processes for machine learning*. MIT Press.
- Srinivasan, B., François, G., and Bonvin, D. (2011). Comparison of gradient estimation methods for real-time optimization. In *21st European Symposium on Computer Aided Process Engineering-ESCAPE 21*, EPFL-CONF-155235, 607–611. Elsevier.
- Xu, W., Jones, C.N., Svetozarevic, B., Laughman, C.R., and Chakrabarty, A. (2022). Vabo: Violation-aware bayesian optimization for closed-loop control performance optimization with unmodeled constraints. In *2022 American Control Conference (ACC)*, 5288–5293. IEEE.