

Primal-Dual Contextual Bayesian Optimization for Control System Online Optimization with Time-Average Constraints

Wenjie Xu, Yuning Jiang, Bratislav Svetozarevic, Colin N. Jones

Abstract—This paper studies the problem of online performance optimization of constrained closed-loop control systems, where both the objective and the constraints are unknown black-box functions affected by exogenous time-varying contextual disturbances. A primal-dual contextual Bayesian optimization algorithm is proposed that achieves sublinear cumulative regret with respect to the dynamic optimal solution under certain regularity conditions. Furthermore, the algorithm achieves zero time-average constraint violation, ensuring that the average value of the constraint function satisfies the desired constraint. The method is applied to both sampled instances from Gaussian processes and a continuous stirred tank reactor parameter tuning problem; simulation results show that the method simultaneously provides close-to-optimal performance and maintains constraint feasibility on average. This contrasts current state-of-the-art methods, which either suffer from large cumulative regret or severe constraint violations for the case studies presented.

I. INTRODUCTION

Constrained closed-loop control systems are critical to a wide range of applications, and optimizing their online performance by tuning control parameters is a common challenge. This paper studies the time-varying black-box optimization problems that arise in optimizing system performance with constraints, where both the objective and the constraints are unknown. For instance, consider building control, which aims to minimize energy consumption while meeting occupant comfort requirements [1]. Solving the controller tuning problem in such scenarios is challenging for several reasons. Firstly, it is hard to model the mapping from control parameters to the performance metric of the closed-loop system. Secondly, unknown constraint violations need to be taken care of during the optimization process. Finally, evaluating the system’s performance is often expensive in practice.

To tackle these challenges, Bayesian optimization has shown promise as a sample-efficient, derivative-free black-box optimization method [2]. Bayesian optimization constructs a surrogate model using Gaussian process regression [3] and uses this model to guide the sampling of the black-box functions. Variants of Bayesian optimization methods have been proposed to handle constrained optimization problems. One common method is to maximize constrained expected improvement (CEI) [4], [5] to select the next

sample in each step. Another line of research developed safe Bayesian optimization (Safe BO) methods by restricting the sampling to only feasible points [6]–[9]. These methods have found wide applications in control system optimization. For example, CEI method is applied to trajectory optimization for path following control [10]. Safe BO is applied to tune the PI controller of a room heating system subject to a group of safe constraints [11]. However, the CEI method may not guarantee constraint feasibility and can suffer from severe constraint violations, while safe Bayesian optimization methods can be too cautious and lose performance due to the strict requirement of sampling feasible points.

In the context of controller tuning, one additional challenge that needs to be considered is that the unknown objective and constraints are time-varying [12]. To address this challenge, several time-varying variants of Bayesian optimization have been proposed in the literature [12]–[14]. In practice, these variations can be due to the change of ‘contextual variables’ observed by the decision-maker before selecting a new set of parameters [15]. Affected by the time-varying contextual disturbances, guaranteeing constraint satisfaction at every time instant can be challenging and may not be necessary [16]. But rather, it can be more of interest to satisfy the *time-average* constraint [17], especially when it represents some economic cost that accumulates over time. Examples include performance optimization of a machine subject to fatigue constraints (see, e.g., [18]) and data center cooling subject to the constraints on the number of delayed queries per unit of time [19].

Motivated by the aforementioned observations, we formulate the time-varying black-box optimization problem as a constrained contextual Bayesian optimization problem. To solve it, we extend the primal-dual Bayesian optimization framework [20] to the contextual setting [21]. In the proposed algorithm, the contextual variable that impacts the response surface of the black-box functions can be observed at each step before selecting a new set of parameters, and then the system performance is optimized over the current context. The detailed contribution is summarized as follows:

- We propose a **Primal-Dual Contextual Bayesian Optimization (PDCBO)** algorithm for time-varying constrained black-box optimization problems. In contrast to the regret with respect to a static optimal value in the non-contextual setting [20], we provide bounds on the contextual regret with respect to the time-varying optimal solutions. Furthermore, our algorithm can achieve *zero* time-averaged constraint violations, even under the *adversarial* time-varying contextual setting;

This work was supported by the Swiss National Science Foundation under the NCCR Automation (grant agreement 51NF40_180545).

W. Xu, Y. Jiang, and C. N. Jones are with the Automatic Control Laboratory, EPFL, Switzerland. W. Xu and B. Svetozarevic are with Swiss Federal Laboratories for Materials Science and Technology (Empa), Switzerland. {wenjie.xu, yuning.jiang, colin.jones}@epfl.ch, bratislav.svetozarevic@empa.ch

- The proposed method is deployed to both sampled instances from Gaussian processes and to a continuous stirred-tank reactor parameter tuning problem. Compared to other state-of-the-art methods, our method can simultaneously achieve the lowest cumulative objective while satisfying all the constraints on average.

II. PROBLEM FORMULATION

This work aims to sequentially optimize the control parameters $\theta_t \in \Theta \subset \mathbb{R}^{n_\theta}$ after observing some contextual variable $z_t \in \mathcal{Z} \subset \mathbb{R}^{n_z}$ in each step t , where Θ is the candidate set of the control parameters and \mathcal{Z} is the set of possible contextual variables. Our problem can be thus, formulated as follows:

$$\min_{\theta \in \Theta} f(\theta, z), \quad (1a)$$

$$\text{subject to: } g_i(\theta, z) \leq 0, \quad \forall i \in [N], \quad (1b)$$

where $[N]$ is defined as the set $\{1, 2, \dots, N\} \subset \mathbb{Z}$, $f : \mathbb{R}^{n_\theta} \times \mathbb{R}^{n_z} \rightarrow \mathbb{R}$ is the black-box objective function, and $g_i : \mathbb{R}^{n_\theta} \times \mathbb{R}^{n_z} \rightarrow \mathbb{R}$, $i \in [N]$ is a set of black-box constraints to be satisfied. We use $\theta^*(z)$ to denote the optimal solution of the problem (1), which depends on the current context variable z . We use g to denote the concatenation $(g_i)_{i=1}^N$ and $g(\theta, z)$ to denote the concatenation $(g_i(\theta, z))_{i=1}^N$.

We make some assumptions regarding the regularities of the problem.

Assumption 1 *Both \mathcal{Z} and Θ are compact.*

Assumption 1 is commonly seen in practice. For example, we can restrict the control parameter set Θ to a hyper box and often know an a priori bound on the range of contextual variables.

Assumption 2 *$f \in \mathcal{H}_0$, $g_i \in \mathcal{H}_i$, where $\mathcal{H}_i, i \in \{0\} \cup [N]$ is a reproducing kernel Hilbert space (RKHS) equipped with kernel function $k_i(\cdot, \cdot)$ (See [22]). Furthermore, $\|f\| \leq C_0$, $\|g_i\| \leq C_i, \forall i \in [N]$, where $\|\cdot\|$ is the norm induced by the inner product of the corresponding RKHS without further notice.*

Assumption 2 requires that the underlying functions are regular in the sense of having a bounded norm in an RKHS. In the existing literature, having a bounded norm in an RKHS is a commonly adopted assumption (e.g., [20], [23]). Intuitively, it means the black-box function has a ‘smoothness’ property at least to a certain level (See [22]).

Assumption 3 *We only have access to a noisy zero-order oracle, which means each round of query θ_t with contextual variable z_t returns the noisy function evaluations,*

$$y_{0,t} = f(x_t) + \nu_{0,t}, \quad (2a)$$

$$y_{i,t} = g_i(x_t) + \nu_{i,t}, \quad i \in [N], \quad (2b)$$

where x_t denotes the concatenation (θ_t, z_t) and $\nu_{i,t}, i \in \{0\} \cup [N]$ are independent and identically distributed σ -sub-Gaussian noise.

Assumption 3 is reasonable as, in many applications, the performance value of a control system can be obtained through experiment or simulation with noise. However, the gradient or other higher-order information can be difficult to acquire.

Throughout the rest of this paper, we use the notation $X_t := (x_1, x_2, \dots, x_t)$ to define the sequence of sampled points up to step t . Therefore, the historical evaluations are

$$\mathcal{D}_{1:t} := \{(x_\tau, y_{0,\tau}, \dots, y_{N,\tau})\}_{\tau=1}^t.$$

Assumption 4 *The kernel function is normalized, such that, $k_i(x, x) \leq 1, \forall x \in \mathcal{X}, i \in \{0\} \cup [N]$, where $\mathcal{X} = \mathcal{Z} \times \Theta$.*

Assumption 5 *There exists $\xi > 0$ independent of the context $z \in \mathcal{Z}$ and $\theta(z) \in \Theta$, which can be dependent on z , such that $g(\theta(z), z) \leq -\xi e$, where $e \in \mathbb{R}^N$ is the vector with all 1s and the inequality is interpreted elementwise.*

Assumption 5 is the so-called Uniform Slater Condition. Assump. 5 is quite mild since we only require that a small ξ exists rather than knowing a feasible solution as assumed in safe Bayesian optimization literature [6]. As an immediate implication of Assumption 5, we have the following result.

Proposition 1 *There exists a probability distribution $\bar{\pi}(z)$ over Θ , such that $\mathbb{E}_{\bar{\pi}(z)}(g(\theta, z)) \leq -\xi e$.*

III. PERFORMANCE METRIC

We compare the controller parameters selected by our algorithm with the optimal tracking solution of problem (1), given the context z taking the value z_t at step t .

We are interested in two metrics,

$$\mathcal{R}_T = \sum_{t=1}^T (f(\theta_t, z_t) - f(\theta^*(z_t), z_t)), \quad (3a)$$

$$\mathcal{V}_T = \left\| \left[\sum_{t=1}^T g(\theta_t, z_t) \right]^+ \right\|, \quad (3b)$$

which are the cumulative regret compared to the optimal tracking solutions and the cumulative constraint violations. The regret here is essentially the *contextual regret* [21], which is stronger than the regret considered in [20].

The form of constraint violation in (3b) is the violation of cumulative constraint value. \mathcal{V}_T/T gives the violation of the average constraint value, which is common in practice. For example, when g represents some economic cost (e.g., monetary expenses and energy consumption), it is usually of more interest to bound the average constraint value during a period rather than the whole constraint function sequence.

IV. GAUSSIAN PROCESS REGRESSION

As in the existing popular Bayesian optimization methods, we use Gaussian process surrogates to learn the unknown functions. As in [24], we artificially introduce a Gaussian process $\mathcal{GP}(0, k_0(\cdot, \cdot))$ for the surrogate modelling of the

unknown black-box function f . We also adopt an i.i.d Gaussian zero-mean noise model with noise variance $\lambda > 0$, which can be chosen by the algorithm.

Recall that x_t is defined as (θ_t, z_t) . We introduce the following functions of (x, x') ,

$$\mu_{0,t}(x) = k_0(x_{1:t}, x)^\top (K_{0,t} + \lambda I)^{-1} y_{0,1:t}, \quad (4a)$$

$$k_{0,t}(x, x') = k_0(x, x') - k_0(x_{1:t}, x)^\top (K_{0,t} + \lambda I)^{-1} k_0(x_{1:t}, x'), \quad (4b)$$

$$\sigma_{0,t}^2(x) = k_{0,t}(x, x) \quad (4c)$$

with $k_0(x_{1:t}, x) = [k_0(x_1, x), k_0(x_2, x), \dots, k_0(x_t, x)]^\top$, $K_{0,t} = (k_0(x, x'))_{x, x' \in X_t}$, $y_{0,1:t} = [y_{0,1}, y_{0,2}, \dots, y_{0,t}]^\top$. Similarly, we can get $\mu_{i,t}(\cdot), k_{i,t}(\cdot, \cdot), \sigma_{i,t}(\cdot), \forall i \in [N]$ for the constraints.

To characterize the complexity of the Gaussian processes and the RKHSs corresponding to the kernel functions, we further introduce the maximum information gain for the objective f as in [23],

$$\gamma_{0,t} := \max_{A \subset \Theta \times \mathcal{Z}; |A|=t} \frac{1}{2} \log |I + \lambda^{-1} K_{0,A}|, \quad (5)$$

where $K_{0,A} = (k_0(x, x'))_{x, x' \in A}$. Similarly, we introduce $\gamma_{i,t}, \forall i \in [N]$ for the constraints.

Remark 1 *Note that the Gaussian process model above is only used to derive posterior mean functions, covariance functions, and maximum information gain for algorithm design and theoretical analysis. It does not change our setup that f is a deterministic function and that the observation noise only needs to be sub-Gaussian.*

Based on the aforementioned preliminaries of Gaussian process regression, we then derive the lower confidence, and upper confidence bound functions.

Lemma 1 *Let Assumptions 1 and 2 hold. With probability at least $1 - \delta, \forall \delta \in (0, 1)$, the following holds for all $x \in \mathcal{X}$ and $t \geq 1$,*

$$g_i(x) \in [\underline{g}_{i,t}(x), \bar{g}_{i,t}(x)], \quad \forall i \in [N] \quad (6)$$

$$\text{and } f(x) \in [\underline{f}_t(x), \bar{f}_t(x)], \quad (7)$$

where for all $i \in [N]$,

$$\underline{f}_t(x) := \max\{\mu_{0,t-1}(x) - \beta_{0,t}^{1/2} \sigma_{0,t-1}(x), -C_0\}, \quad (8a)$$

$$\bar{f}_t(x) := \min\{\mu_{0,t-1}(x) + \beta_{0,t}^{1/2} \sigma_{0,t-1}(x), C_0\}, \quad (8b)$$

$$\underline{g}_{i,t}(x) := \max\{\mu_{i,t-1}(x) - \beta_{i,t}^{1/2} \sigma_{i,t-1}(x), -C_i\}, \quad (8c)$$

$$\bar{g}_{i,t}(x) := \min\{\mu_{i,t-1}(x) + \beta_{i,t}^{1/2} \sigma_{i,t-1}(x), C_i\}, \quad (8d)$$

$$\beta_{i,t}^{1/2} := C_i + \sigma \sqrt{2(\gamma_{i,t-1} + 1 + \ln((N+1)/\delta))}. \quad (8e)$$

Proof: By Corollary 2.6, [25], with probability at least $1 - \delta, \forall \delta \in (0, 1)$, for all $x \in \mathcal{X}$ and $t \geq 1$,

$$\mu_{i,t-1}(x) - \beta_{i,t}^{1/2} \sigma_{i,t-1}(x) \leq g_i(x) \leq \mu_{i,t-1}(x) + \beta_{i,t}^{1/2} \sigma_{i,t-1}(x).$$

Furthermore, $|g_i(x)| = |\langle g_i, k_i(x, \cdot) \rangle| \leq \|g_i\| \|k_i(x, \cdot)\| \leq C_i, \forall i \in [N]$. Therefore, $g_i(x) \in [\underline{g}_{i,t}(x), \bar{g}_{i,t}(x)]$. Similarly, $f(x) \in [\underline{f}_t(x), \bar{f}_t(x)]$. ■

Without further notice, all the following results are conditioned on the high probability event in Lem. 1 happening.

V. ALGORITHM AND THEORETICAL GUARANTEES

Our primal-dual algorithm is shown in Alg. 1, where η is a small scaling factor to be chosen, λ_t can be interpreted as the dual variable up to the scaling of η , $0 < \epsilon \leq \frac{\xi}{2}$ is a slackness parameter, and $[\cdot]^+$ is interpreted element-wise. Intuitively, the larger η is, the more emphasis is given to the constraints.

Algorithm 1 Primal-Dual Contextual Bayesian Optimization (PDCBO).

1: **for** $t \in [T]$ **do**

2: Observe contextual variables z_t .

3: **Primal update:**

$$\theta_t = \arg \min_{\theta \in \Theta} \{ \underline{f}_t(\theta, z_t) + \eta \lambda_t^T \underline{g}_t(\theta, z_t) \}. \quad (9)$$

4: **Dual update:**

$$\lambda_{t+1} = [\lambda_t + \underline{g}_t(\theta_t, z_t) + \epsilon e]^+. \quad (10)$$

5: Evaluate f and $g_i, i \in [N]$ at $x_t = (\theta_t, z_t)$ with noise.

6: Update $(\mu_{i,t}, \sigma_{i,t}), i \in \{0\} \cup [N]$ with the new data.

7: **end for**

The Alg. 1 is based on the well-known primal-dual optimization method. The Lagrangian of the original problem is $\mathcal{L}(\theta, z, \phi) = f(\theta, z) + \phi^T g(\theta, z)$ and the dual function is $\mathcal{D}(z, \phi) = \min_{\theta} \mathcal{L}(\theta, z, \phi)$. In the primal update step, we optimize the Lagrangian. In the dual update step, we apply the dual ascent method. In primal and dual updates, we replace the unknown black-box functions with their lower confidence bounds, following the principle of *optimism in the face of uncertainty*.

The primal update (9) requires solving a potentially non-convex auxiliary optimization problem. When the dimension of θ is small (e.g., < 6) we can use a pure grid search method to solve it. For larger dimensions, one can, for example, apply the gradient descent method from multiple different initial points. Line 5 can be done by querying a simulator or doing an experiment in practice.

Remark 1 *In Alg. 1, λ_t is not exactly the dual variable, but the dual variable scaled by $\frac{1}{\eta}$. Indeed, λ_t can be interpreted as virtual queue length [20]. The intuition of ϵ is to introduce some level of constant pessimistic drift to control the cumulative violation.*

We now give the theoretical guarantees in Thm. 1.

Theorem 1 *Let the Assumptions 1, 2, 3, 4 and 5 hold. We further assume the maximum information gain term satisfies*

$\lim_{T \rightarrow \infty} \frac{\|\gamma_T^g\|}{\sqrt{T}} = 0$, where $\gamma_T^g := (\gamma_{1,T}, \dots, \gamma_{N,T})$. We set $\eta = \frac{1}{\sqrt{T}}$, $\lambda_1 = \left(\frac{4C_0}{\eta\xi} + \frac{4\|C\|^2}{\xi}\right)e$ and

$$\epsilon = \frac{\sqrt{N \left(\frac{4C_0}{\eta\xi} + \frac{4\|C\|^2}{\xi}\right)^2 + \frac{4C_0}{\eta} + 4\|C\|^2 + 8\|\beta_T^g\|\sqrt{T}\|\gamma_T^g\|}}{T},$$

where $\beta_T^g := (\beta_{1,T}, \dots, \beta_{N,T})$ and $C := (C_1, \dots, C_N)$. Let T be large enough such that $\epsilon = \mathcal{O}(\|\gamma_T^g\|/\sqrt{T}) \leq \xi/2$. We have

$$\mathcal{R}_T = \mathcal{O}\left(\sum_{i=0}^N \gamma_{i,T} \sqrt{T}\right) \text{ and } \mathcal{V}_T = 0.$$

Proof Sketch Our proof involves three key ingredients. First, the cumulative cost paid to learn the black-box functions can be bounded by some term involving the maximum information gain (5). Second, the regret and the violations in the primal values can be bounded by the drift of dual variables. Third, by properly selecting the constant dual drift ϵ^1 , we can balance the regret and violation such that *time-average* constraints are strictly satisfied while the cumulative *contextual regret* is bounded. The detailed proofs can be found in the Appendix of our technical report [26].

Remark 2 *Thm. 1 highlights that for kernels with the maximum information gain terms that grow slower than \sqrt{T} , our algorithm achieves average contextual regret converging to zero while incurring no constraint violation on average. The slow growth assumption holds for most popular kernels including Squared Exponential and Matérn (under the assumption shown in Tab. I) kernels [23].*

By applying the bounds on maximum information gains from [23], [27] to Thm. 1, we derive the kernel-specific bounds as in Tab. I, where $d = n_\theta + n_z$, ν is the smoothness parameter of the Matérn kernel, and $\tilde{\mathcal{O}}(\cdot)$ hides some polylogarithmic term.

TABLE I
KERNEL-SPECIFIC REGRET BOUNDS.

Kernel	Linear	Squared Exponential	Matérn ($\frac{\nu}{d} > \frac{1}{2}$)
\mathcal{R}_T	$\tilde{\mathcal{O}}(d\sqrt{T})$	$\tilde{\mathcal{O}}(\sqrt{T})$	$\tilde{\mathcal{O}}\left(T^{\frac{2\nu+3d}{4\nu+2d}}\right)$

Remark 3 *In Thm. 1, the running steps T is known beforehand. To extend to the case where T is unknown, we can apply the doubling trick [28]. The idea is to start the algorithm with a small T and restart it with T doubled every time the running steps are exhausted.*

VI. EXPERIMENTS

We consider two sets of experiments. In the first set, we use the objective and constraint functions sampled from Gaussian processes. In the second set, we consider the classical Williams-Otto benchmark problem [29]. We compare our

¹In practice, we can set ϵ to be a small value as compared to the range of the constraint function or even just 0.

method to the state-of-the-art Bayesian optimization methods with constraints consideration, including SafeOPT [6], [30] and CEI (Constrained Expected Improvement) method [4], [5]. The experiments are implemented in python, based on the package GPpy [31]. The code for the experiments is available through <https://github.com/PREDICT-EPFL/PDCBO>.

Choice of Parameters. The performance of PDCBO is mainly impacted by $\beta_{i,t}^{1/2}$. In practice, $\beta_{i,t}^{1/2}$ can usually be set as a constant. Indeed, when the kernel choices and the kernel hyperparameters fit the black-box functions well, setting $\beta_{i,t}^{1/2} = 3$ typically works well. In our experiments, manually setting $\beta_{i,t}^{1/2} = 1.0$ works well. We also set $\lambda = 0.05^2$ as the noise variance for the Gaussian process modeling. We use the common squared exponential kernel. For the second example, we sample a few points randomly and maximize the likelihood function to get the hyperparameters of the kernel.

Computational Time. In our experiments, all the problems have low-dimensional inputs ($n_\theta \leq 3$). So we use pure grid search to solve the auxiliary problem for primal update, which is relatively cheap due to the explicit expressions of the lower confidence bound functions as compared to the simulation of the Williams-Otto benchmark problem.

Performance Metrics. To measure the quality of sample sequences, we use constrained contextual regret shown in Eq. 3a. To measure the violations, we introduce the cumulative constraint value $\sum_{\tau=1}^t g_i(\theta_\tau, z_\tau)$, $i \in [N]$, and the time-average constraint value $\frac{1}{t} \sum_{\tau=1}^t g_i(\theta_\tau, z_\tau)$, $i \in [N]$.

A. Sampled Instances from Gaussian Process

We consider the problem (1) with only one constraint, where $\theta \in \Theta = [-10, 10]$, $z \in \mathcal{Z} = [-10, 10]$, and both f and g are unknown black-box functions sampled from a Gaussian process. We use the squared exponential kernel,

$$k((\theta_1, z_1), (\theta_2, z_2)) = \sigma_{SE}^2 \exp \left\{ - \left(\frac{\theta_1 - \theta_2}{l_\theta} \right)^2 - \left(\frac{z_1 - z_2}{l_z} \right)^2 \right\}, \quad (11)$$

where $\sigma_{SE}^2 = 2.0$, and $l_\theta = l_z = 1.0$. We run different algorithms for 500 steps over 50 different sampled instances. The context sequence is randomly generated from a uniform distribution over the context set \mathcal{Z} . Fig. 1 shows the cumulative contextual regret and cumulative constraint value of different algorithms over 50 different trajectories. It can be observed that our PDCBO achieves sublinear cumulative contextual regret while keeping the cumulative constraint value below zero with high confidence. In contrast, the CEI method suffers from almost linear cumulative contextual regret and fails to keep average constraint feasibility with high probability. Although Safe BO can keep average constraint feasibility, its cumulative contextual regret is about 62% higher than our PDCBO due to its being overly cautious to keep feasibility in every single step.

B. Williams-Otto Classical Problem

In this section, we consider the problem of the classical Williams-Otto benchmark problem [29]. In this problem,




Fig. 1. The cumulative context regret and cumulative constraint value of different algorithms. The statistics are obtained over 50 different trajectories. The shaded area represents ± 0.5 standard deviation.

we feed a continuous stirred-tank reactor (CSTR) with two components A and B. The reactor then operates at the steady state under the reaction temperature T_r . During the chemical reaction, a byproduct G is produced. We use X_A and X_G to denote the steady-state residual mass fractions of A and G at the reactor outlet. Both X_A and X_G need to be managed. That is, we need to control the steady-state residual mass fractions of A and G to be lower than some thresholds. See [29, Sec. 4.1] for more details. Meanwhile, a key factor that impacts the cumulative profit of the chemical reaction process is the time-varying nature of the product price and the costs of the raw materials [32]. To incorporate these time-varying factors, we model them as contextual variables.

We want to adaptively tune the feed rate F_B of the component B and the reaction temperature T_r to maximize the cumulative economic profit from the reaction while managing X_A and X_G . Our problem can be formulated as,

$$\begin{aligned}
 & \min_{F_B, T_r} J(F_B, T_r, P) \\
 & \text{subject to} \quad \text{CSTR model [33]} \\
 & \quad g_1(F_B, T_r) := X_A(F_B, T_r) - 0.12 \leq 0 \quad (12) \\
 & \quad g_2(F_B, T_r) := X_G(F_B, T_r) - 0.08 \leq 0, \\
 & \quad F_B \in [4, 7], T_r \in [70, 100]
 \end{aligned}$$

where $J(F_B, T_r, P)$ is the minimization objective that is opposite to the net economic profit, $P \in \mathbb{R}^4$ is the price

vector of the product and the raw materials, $g_1(F_B, T_r)$ and $g_2(F_B, T_r)$ are threshold constraints on the residual mass fractions. To generate the context P^t at step t , we randomly sample P_i^t uniformly from the interval $[(1-\alpha)\bar{P}_i, (1+\alpha)\bar{P}_i]$, where \bar{P}_i is a predefined normal price and $\alpha = 0.2$.

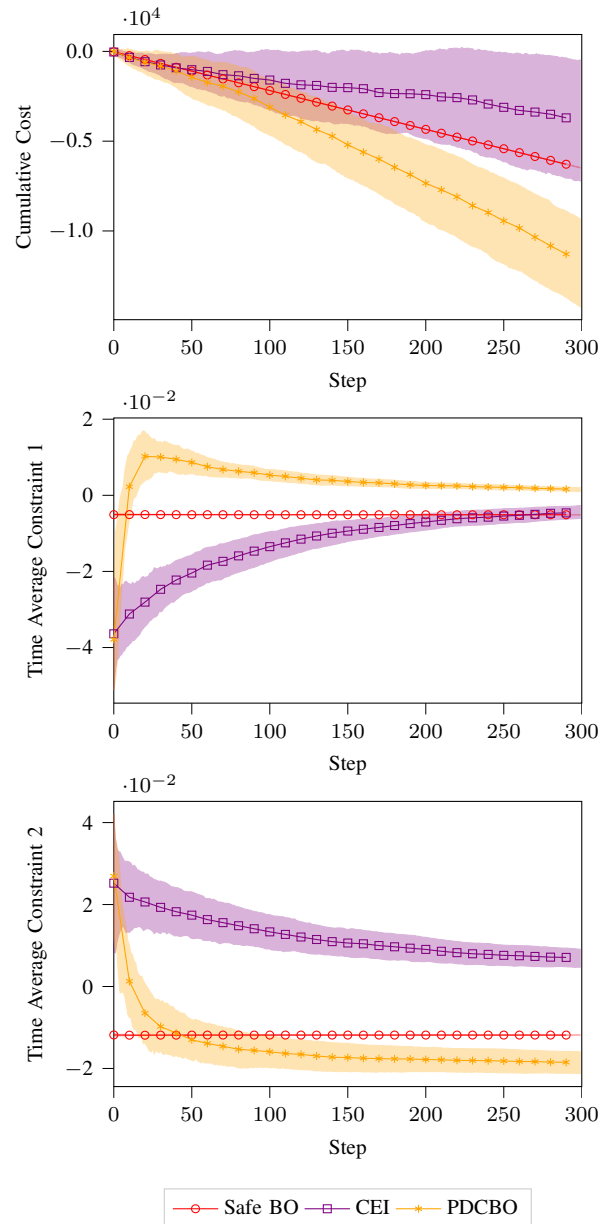


Fig. 2. Cumulative cost and average constraints of the chemical reaction over 50 different context trajectories. The shaded area represents ± 1.0 standard deviation.

Fig. 2 shows the cumulative cost and the running average of the constraints with 50 different context trajectories. We observe that our PDCBO achieves the lowest cumulative cost (the highest cumulative profit, equivalently) while satisfying the constraints on average. In contrast, Safe BO is overly cautious and can not achieve all the potential economic profit. Finally, the generic CEI method suffers from both low profit and severe constraint violation for the second constraint.

VII. CONCLUSION

This paper studied the problem of optimizing the closed-loop performance of constrained control systems, where both the objective and the constraints are unknown black-box functions affected by exogenous time-varying contextual disturbances. A primal-dual contextual Bayesian optimization algorithm was proposed, and it was shown to achieve sub-linear cumulative regret with respect to the dynamic optimal solution under certain regularity conditions. Furthermore, the algorithm achieves zero time-average constraint violation, ensuring that the average value of the constraint function satisfies the desired constraint. Our method is particularly useful when the constraint functions correspond to some economic costs that accumulate with time. The method is applied to both sampled instances from Gaussian processes and to a continuous stirred tank reactor feedrate and reaction temperature tuning problem; simulation results on both problems show that the method simultaneously provides close-to-optimal performance and maintains constraint feasibility on average. This contrasts current state-of-the-art methods, which either suffer from large cumulative regret or severe constraint violations for the case studies presented.

REFERENCES

- [1] W. Xu, C. N. Jones, B. Svetozarevic, C. R. Laughman, and A. Chakrabarty, "VABO: Violation-Aware Bayesian Optimization for closed-loop control performance optimization with unmodeled constraints," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 5288–5293.
- [2] W. Xu, Y. Jiang, E. T. Maddalena, and C. N. Jones, "Lower bounds on the worst-case complexity of efficient global optimization," *arXiv preprint arXiv:2209.09655*, 2022.
- [3] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [4] M. A. Gelbart, J. Snoek, and R. P. Adams, "Bayesian optimization with unknown constraints," in *Proc. of the 30th Conf. on Uncertainty in Artif. Intell.*, ser. UAI'14. Arlington, Virginia, USA: AUAI Press, 2014, p. 250–259.
- [5] J. R. Gardner, M. J. Kusner, Z. E. Xu, K. Q. Weinberger, and J. P. Cunningham, "Bayesian optimization with inequality constraints," in *Proc. of the Int. Conf. on Mach. Learn.*, vol. 2014, 2014, pp. 937–945.
- [6] Y. Sui, A. Gotovos, J. Burdick, and A. Krause, "Safe exploration for optimization with Gaussian processes," in *Proc. of the Int. Conf. on Mach. Learn.*, 2015, pp. 997–1005.
- [7] Y. Sui, J. Burdick, Y. Yue *et al.*, "Stage-wise safe Bayesian optimization with Gaussian processes," in *Proc. of the Int. Conf. on Mach. Learn.*, 2018, pp. 4781–4789.
- [8] D. Bergmann and K. Graichen, "Safe Bayesian optimization under unknown constraints," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 3592–3597.
- [9] D. Baumann, A. Marco, M. Turchetta, and S. Trimpe, "Gosafe: Globally optimal safe robot learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4452–4458.
- [10] A. Rupenyan, M. Khosravi, and J. Lygeros, "Performance-based trajectory optimization for path following control using Bayesian optimization," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 2116–2121.
- [11] M. Fiducioso, S. Curi, B. Schumacher, M. Gwerder, and A. Krause, "Safe contextual Bayesian optimization for sustainable room temperature PID control tuning," in *Proc. of the Twenty-Eighth Int. Joint Conf. on Artif. Intell.* IJCAI, 2019, pp. 5850–5856.
- [12] P. Brunzema, A. Von Rohr, and S. Trimpe, "On controller tuning with time-varying Bayesian optimization," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 4046–4052.
- [13] I. Bogunovic, J. Scarlett, and V. Cevher, "Time-varying Gaussian process bandit optimization," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 314–323.
- [14] Y. Deng, X. Zhou, B. Kim, A. Tewari, A. Gupta, and N. Shroff, "Weighted Gaussian process bandits for non-stationary environments," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp. 6909–6932.
- [15] W. Xu, C. N. Jones, B. Svetozarevic, C. R. Laughman, and A. Chakrabarty, "Violation-aware contextual Bayesian optimization for controller performance optimization with unmodeled constraints," *arXiv preprint arXiv:2301.12099*, 2023.
- [16] J. Fleming and M. Cannon, "Time-average constraints in stochastic Model Predictive Control," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 5648–5653.
- [17] M. Korda, R. Gondhalekar, F. Oldewurtel, and C. N. Jones, "Stochastic MPC framework for controlling the average constraint violation," *IEEE Transactions on Automatic Control*, vol. 59, no. 7, pp. 1706–1721, 2014.
- [18] M. Cannon, B. Kouvaritakis, and X. Wu, "Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty," *IEEE Transactions on Automatic Control*, vol. 54, no. 7, pp. 1626–1632, 2009.
- [19] J. Kim, M. Ruggiero, and D. Atienza, "Free cooling-aware dynamic power management for green datacenters," in *2012 International Conference on High Performance Computing & Simulation (HPCS)*. IEEE, 2012, pp. 140–146.
- [20] X. Zhou and B. Ji, "On kernelized multi-armed bandits with constraints," in *Advances in Neural Information Processing Systems*, 2022.
- [21] A. Krause and C. Ong, "Contextual Gaussian process bandit optimization," *Advances in Neural Information Processing Systems*, vol. 24, 2011.
- [22] B. Schölkopf, R. Herbrich, and A. J. Smola, "A generalized representer theorem," in *International conference on computational learning theory*. Springer, 2001, pp. 416–426.
- [23] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for Gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.
- [24] S. R. Chowdhury and A. Gopalan, "On kernelized multi-armed bandits," in *Int. Conf. on Mach. Learn.* PMLR, 2017, pp. 844–853.
- [25] W. Xu, Y. Jiang, B. Svetozarevic, and C. Jones, "Constrained efficient global optimization of expensive black-box functions," in *International Conference on Machine Learning*. PMLR, 2023, pp. 38485–38498.
- [26] W. Xu, Y. Jiang, B. Svetozarevic, and C. N. Jones, "Primal-dual contextual bayesian optimization for control system online optimization with time-average constraints," *arXiv preprint arXiv:2304.06104*, 2023.
- [27] S. Vakili, K. Khezeli, and V. Picheny, "On information gain and regret bounds in Gaussian process bandits," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 82–90.
- [28] L. Besson and E. Kaufmann, "What doubling tricks can and can't do for multi-armed bandits," *arXiv preprint arXiv:1803.06971*, 2018.
- [29] E. A. del Rio Chanona, P. Patsagkourakis, E. Bradford, J. A. Graciano, and B. Chachuat, "Real-time optimization meets Bayesian optimization and derivative-free optimization: A tale of modifier adaptation," *Comput. & Chem. Eng.*, vol. 147, p. 107249, 2021.
- [30] F. Berkenkamp, A. Krause, and A. P. Schoellig, "Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics," *Machine Learning*, pp. 1–35, 2021.
- [31] GPpy, "GPpy: A Gaussian process framework in python," <http://github.com/SheffieldML/GPpy>, since 2012.
- [32] R. Yadbantung and P. Bumroongsri, "Periodically time-varying economic model predictive control with applications to nonlinear continuous stirred tank reactors," *Computers & Chemical Engineering*, vol. 157, p. 107602, 2022.
- [33] D. F. Mendoza, J. E. A. Graciano, F. dos Santos Liporace, and G. A. C. Le Roux, "Assessing the reliability of different real-time optimization methodologies," *The Can. J. of Chem. Eng.*, vol. 94, no. 3, pp. 485–497, 2016.