CONFIG: Constrained Efficient Global Optimization for Closed-Loop Control System Optimization with Unmodeled Constraints

Wenjie Xu [*,**] **Yuning Jiang** [*] **Bratislav Svetozarevic** [**] **Colin N. Jones** [*]

[*] *Automatic Control Lab, EPFL, Switzerland (e-mail: wenjie.xu, yuning.jiang colin.jones@epfl.ch)*
[**] *Swiss Federal Laboratories for Materials Science and Technology (EMPA), Switzerland (e-mail: bratislav.svetozarevic@empa.ch)*

Abstract: In this paper, the CONFIG algorithm, a simple and provably efficient constrained global optimization algorithm, is applied to optimize the closed-loop control performance of an unknown system with unmodeled constraints. Existing Gaussian process based closed-loop optimization methods, either can only guarantee local convergence (e.g., SafeOPT), or have no known optimality guarantee (e.g., constrained expected improvement) at all, whereas the recently introduced CONFIG algorithm has been proven to enjoy a theoretical global optimality guarantee. In this study, we demonstrate the effectiveness of CONFIG algorithm in the applications. The algorithm is first applied to an artificial numerical benchmark problem to corroborate its effectiveness. It is then applied to a classical constrained steady-state optimization problem of a continuous stirred-tank reactor. Simulation results show that our CONFIG algorithm can achieve performance competitive with the popular CEI (Constrained Expected Improvement) algorithm, which has no known optimality guarantee. As such, the CONFIG algorithm offers a new tool, with both a provable global optimality guarantee and competitive empirical performance, to optimize the closed-loop control performance for a system with soft unmodeled constraints. Last, but not least, the open-source code is available as a python package to facilitate future applications.

Keywords: Constrained Efficient Global Optimization, Closed-loop Performance, Unmodeled Constraints, Bayesian Optimization, Controller Tuning.

1. INTRODUCTION

The performance of closed-loop systems can typically be optimized by tuning control parameters (e.g., gains or set-points) under the guidance of the operational data. Manual tuning of these parameters often takes significant human effort and domain-specific engineering expertise, which may make controller tuning economically infeasible, since the personnel cost may outweigh the benefit from optimized performance. Therefore, algorithms that can automatically adjust these control parameters, without human effort, to optimize closed-loop performance are of great practical interest.

Closed-loop performance of a control system can typically be defined as functions that take inputs and outputs measured from the closed-loop system as the arguments. For example, in building thermal control, discomfort can be defined as the integration of temperature deviation from a comfort range. While how performance is determined by inputs and outputs is typically clearly understood, the map from control parameters, which are actually tuned, to the closed-loop performance is largely unknown and behaves like a black-box. It is either because an identified system model is unavailable during parameter tuning or the models are so complicated that it is almost impossible to directly construct the explicit mapping from controller parameters to closed-loop performance.

Therefore, data-driven tuning methods, where the control parameters-to-performance map is considered as a black-box function that can be learned with closed-loop operational data, are receiving more and more research attention. To get closed-loop operation data, we can perform experiments or simulate a model. However, both experimentation and high-fidelity software simulations are expensive in terms of hardware use or very long computation time. Therefore, it is desired that the tuning algorithms are able to identify a near-optimal set of control parameters with as few experiments/simulations as possible.

Due to the above reasons, efficient global optimization [1] (Jones et al., 1998) has been widely applied to the closed-loop controller tuning problem, thanks to its black-box modelling capability and superior empirical sample

[1] Also referred to as Gaussian process optimization (e.g., in (Srinivas et al., 2012)) or Bayesian optimization (e.g., in (Frazier, 2018)) in machine learning literatures or kriging method (e.g., in (Jeong et al., 2005)) in some engineering literature.

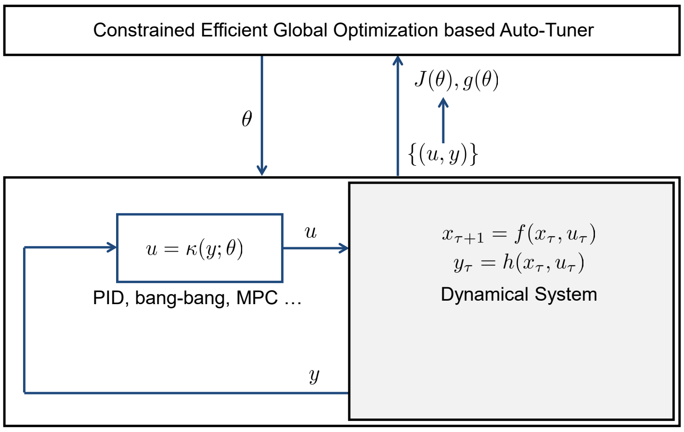

Fig. 1. Constrained efficient global optimization based auto-tuner optimizes the performance of a dynamical system controlled by a controller parameterized by θ. By conducting software simulation or physical experiments, we can collect input-ouput trajectories $\{(u, y)\}$ and derive the values of the objective $J(\theta)$ and the constraints $g(\theta)$. On top of the closed-loop dynamical system, we apply our constrained efficient global optimization method to guide the search of optimal feasible parameters θ.

efficiency (Xu et al., 2022b). Efficient global optimization is a type of sample-efficient derivative-free global optimization method (Jones et al., 1998; Frazier, 2018) that utilizes Gaussian process regression as a surrogate to adaptively search through parameter spaces. In recent work, efficient global optimization has demonstrated potential in controller gain tuning (Lederer et al., 2020; Duivenvoorden et al., 2017; Khosravi et al., 2019; König et al., 2020), MPC tuning (Bansal et al., 2017; Piga et al., 2019; Paulson and Mesbah, 2020) and in various other real-world control applications, such as wind energy systems (Baheri and Vermillion, 2017, 2020), engines (Pal et al., 2020) and space cooling (Chakrabarty et al., 2021).

Meanwhile, in a large variety of applications, another challenge for closed-loop controller tuning is the existence of unmodelled black-box constraints (Xu et al., 2022c). For example, when tuning the temperature controller parameters of a building, we need to minimize the energy consumption while keeping the comfort above the occupant-desired level. Both the energy and comfort are unknown black-box functions of the controller parameters. For these applications, we need to extend efficient global optimization based tuning methods to the constrained case.

Fig. 1 demonstrates the overall structure of the constrained efficient global optimization based controller tuning method, which optimizes the performance of a dynamical system controlled by a controller parameterized by θ. By conducting software simulation or physical experiments, we can collect input-output trajectories $\{(u, y)\}$ and derive the values of the objective $J(\theta)$ and the constraints $g(\theta)$. On top of the closed-loop system, we apply our constrained efficient global optimization method to guide the search of optimal feasible parameters θ.

To handle these unknown black-box constraints, a variety of efficient global optimization methods with constraints have been developed. We can roughly classify them into different groups based on whether constraint violations are allowed during the optimization process. For the setting where no constraint violations are tolerated, a group of safe Bayesian optimization methods have been developed (Sui et al., 2015, 2018; Turchetta et al., 2020). However, due to hard safety-critical constraints, these algorithms may get stuck at a local minimum. For another setting where constraint violations are allowed and incur no harm, there exist generic constrained Bayesian optimization methods. For example, a group of popular methods (Gardner et al., 2014; Gelbart et al., 2014) encode the constraint information into the acquisition function (e.g., constrained Expected Improvement). However, there are no known theoretical guarantees on the optimality, constraint violations or convergence rate for this set of methods.

More recently, a third group of approaches is receiving more attention (e.g., (Xu et al., 2022c; Zhou and Ji, 2022)), where constraint violations are allowed, but need to be managed well. In this setting, two more recent works adopt a penalty-function approach (Lu and Paulson, 2022) and a primal-dual approach (Zhou and Ji, 2022) to solve the constrained efficient global optimization problem. The common idea of the two works is the addition of a penalizing term of the constraint violation in the objective and transforming the constrained problem into an unconstrained one. However, both these methods require one to choose values for some critical parameters (e.g., penalty coefficient (Lu and Paulson, 2022) and dual update step size (Zhou and Ji, 2022)). The performance of these methods may be impacted heavily by the chosen parameters, but no clear guideline on how to select those critical parameters are known for implementation. Furthermore, the work of (Lu and Paulson, 2022) only analyzes the penalty-based regret, which is defined as the suboptimality plus penalized constraint violations, but does not derive separate bounds for cumulative regret (suboptimality) and violations. In practice, improper choice of penalty parameter may lead to the convergence to suboptimal or infeasible solution. Additionally, the cumulative violations considered in (Zhou and Ji, 2022) are the violations of the cumulative constraint values, not the real total violations. Such a weak form of cumulative violation bounds can not rule out the case of constraint oscillation, where points with severe violations and small constraint values are alternatingly sampled, keeping a low cumulative constraint value. Besides these works with theoretical guarantees, (Priem et al., 2020) also utilizes upper trust bound to approximate the feasible set and maximize some acquisition function inside this approximate feasible set. However, (Priem et al., 2020) neither gives systematic design of upper trust bound and acquisition function nor provides any theoretical feasibility and optimality guarantees.

In this paper, we apply the recently introduced CONFIG (**CON**strained ef**FI**cient **G**lobal optimization) algorithm (Xu et al., 2022a), a simple and provably efficient global optimization algorithm, to optimize the closed-loop control performance of an unknown system with unmodeled constraints. In sharp contrast to existing popoular Gaussian process based closed-loop optimization methods, which either can only guarantee local convergence (e.g., SafeOPT), or have no known optimality guarantee (e.g., constrained expected improvement), CONFIG algorithm has been proven to enjoy a theoretical global optimality guarantee (Xu et al., 2022a). Tab. 1 summarizes the comparison of this method to existing state-of-the-art constrained efficient global optimization methods.

In this study, the algorithm is first applied to artificial numerical problems to indicate its effectiveness. It is then applied to a classical constrained steady-state optimization problem of a continuous stirred-tank reactor. Simulation results show that CONFIG algorithm can achieve a performance competitive with the popular CEI (Constrained Expected Improvement) algorithm, which has no known optimality guarantee. Our algorithm offers a new tool, with a provable global optimality guarantee, to optimize the closed-loop control performance for a system with non-safety critical unmodeled constraints. Last but not least, the open-source code is available as a python package to facilitate future applications.

The **contributions** of the paper include:

(1) A recent variant of the constrained efficient global optimization method, CONFIG, which achieves provable global optimality, is proposed to do closed-loop control system optimization with unmodeled constraints.
(2) The effectiveness and superior performance of the algorithm on artificial problem instances is demonstrated.
(3) The method is applied to a constrained steady-state optimization problem of a continuous stirred-tank reactor. Simulation results show that this new method achieves performance competitive with the popular state-of-the-art methods that operate without known global optimality guarantees.
(4) An open source python implementation of CONFIG is provided. [2]

2. PROBLEM STATEMENT

We consider controlled closed-loop systems of the form

$$x^+ = F(x, \theta), \tag{1}$$

where $x, x^+ \in \mathbb{R}^{n_x}$ denote the system state and its update respectively, $\theta \in \Theta \subset \mathbb{R}^{n_\theta}$ are the control parameters (e.g., set-points) to be tuned, and $F(\cdot,\cdot)$ the closed-loop dynamics with some initial state x_0.

To determine the system performance, we define a continuous cost function $J(\theta) : \mathbb{R}^{n_\theta} \to \mathbb{R}$ to be minimized, which is an unknown/unmodeled black-box function of the parameters θ. We also define N unmodeled constraints on the system outputs that require management during tuning. The i-th such constraint is $g_i(\theta) : \mathbb{R}^{n_\theta} \to \mathbb{R}, i \in [N]$, where the notation $[N] \triangleq \{i \in \mathbb{N}, 1 \leq i \leq N\}$.

Remark 1. Our formulation can capture two typical application scenarios.

- **Steady state optimization.** In steady state optimization (e.g., in (Xu et al., 2022c)), we assume the state converges to a steady state $x^\infty(\theta)$, which is a function of θ. In this case, the cost function $J(\theta)$ is given as a cost function of the steady state
$$J(\theta) \triangleq \ell(x^\infty(\theta)).$$
and the constraints are also defined accordingly.
- **Batch performance optimization.** In batch performance optimization, we optimize batch processes over finite time-horizons, say T_h. Specifically, we define the objective over a batch trajectory as the integration of a stage cost,
$$J(\theta) \triangleq \frac{1}{T_h}\int_0^{T_h} l(x(\tau,\theta))\,\mathrm{d}\tau.$$
Constraints over the period can also be defined similarly. For example, in room temperature controller tuning (Fiducioso et al., 2019), the aim is to minimize energy consumption, which is the integration of power, subject to constraints including integration of temperature reference tracking error.

We formulate the control parameter tuning problem as a black-box constrained optimization problem in Eq. (2).

$$\min_{\theta \in \Theta} \quad J(\theta) \tag{2a}$$

$$\text{subject to:} \quad g_i(\theta) \leq 0, \quad \forall i \in [N]. \tag{2b}$$

Our **objective** is to solve the constrained optimization problem (2) while managing constraint violations during the optimization process.

3. METHODOLOGY

In this section, we introduce the preliminaries of using Gaussian process to model the black box functions $J(\theta)$ and $g_i(\theta), i \in [N]$, before introducing the algorithm.

3.1 Primer on Gaussian process regression

As in existing efficient global optimization based controller tuning works (e.g., (Xu et al., 2022c)), we use Gaussian process surrogate models to learn the unknown functions $J(\theta)$ and $g_i(\theta)$. As in (Chowdhury and Gopalan, 2017), we artificially introduce a Gaussian process $\mathcal{GP}(0, k_i(\cdot,\cdot)), i \in \{0\}\cup[N]$ for the surrogate modelling of the unknown black-box functions J and g_i. We also adopt an i.i.d Gaussian zero-mean noise model with noise variance λ. We use $y_{0,t}$ to denote the measurement of the tuning objective J at step t with the sample θ_t, which is corrupted by independently distributed σ_0 sub-Gaussian noise. Similarly, we use $y_{i,t}$ to denote the noisy measurement of the tuning constraint g_i at step t with the sample θ_t. We use Θ_t to denote the sample sequence $(\theta_1, \theta_2, \cdots, \theta_t)$. We introduce the following functions of θ, θ',

$$\mu_{0,t}(\theta) = k_0(\theta_{1:t}, \theta)^\top \left(K_{0,t} + \lambda I\right)^{-1} y_{0,1:t}, \tag{3a}$$

$$k_{0,t}(\theta,\theta') = k_0(\theta,\theta') - k_0(\theta_{1:t},\theta)^\top \left(K_{0,t}+\lambda I\right)^{-1} k_0(\theta_{1:t},\theta'), \tag{3b}$$

$$\sigma^2_{0,t}(\theta) = k_{0,t}(\theta,\theta), \tag{3c}$$

where $k_0(\theta_{1:t},\theta) = [k_0(\theta_1,\theta), k_0(\theta_2,\theta), \cdots, k_0(\theta_t,\theta)]^\top$, $K_{0,t} = (k_0(\theta,\theta'))_{\theta,\theta'\in\Theta_t}$ and $y_{0,1:t} = [y_{0,1}, y_{0,2}, \cdots, y_{0,t}]^\top$. Similarly, we can get $\mu_{i,t}(\cdot), k_{i,t}(\cdot,\cdot), \sigma_{i,t}(\cdot), \forall i \in [N]$ for the constraints. We also introduce the maximum information gain for the objective f as in (Srinivas et al., 2012),

$$\gamma_{0,t} := \max_{A\subset\Theta;|A|=t} \frac{1}{2}\log\left|I + \lambda^{-1}K_{0,A}\right|, \tag{4}$$

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

[2] Code is available at: https://github.com/JackieXuw/CONFIG

Table 1. The comparison with existing constrained efficient global optimization methods.

Works	Cumulative Regret Bound	Cumulative Violation Bound	Optimality Guarantee	Infeasibility Declaration
SafeOPT (Sui et al., 2015), etc.	✗	No violation	Local convergence	✗
Constrained EI (Gelbart et al., 2014), etc.	✗	✗	✗	✗
Primal-Dual method (Zhou and Ji, 2022), etc.	✓	weak-form	✗	✗
Penalty function method (Lu and Paulson, 2022), etc.	Penalty-based Regret Bound		Global convergence with penalty-dependent rate	✗
CONFIG (Ours) (Xu et al., 2022a)	✓	✓	Global convergence with rate	✓

3.2 *CONFIG algorithm for closed-loop controller tuning*

The CONFIG (**CON**strained ef**FI**cient **G**lobal optimization) based tuning algorithm is shown in Alg. 1, where $\beta_{i,t}^{1/2}$ is a weighting parameter balancing exploration and exploitation. Alg. 1 adopts the principle of *optimism in the face of uncertainty* for both the objective and constraints. Specifically, we solve a constrained auxiliary problem with the original objective and constraints replaced by their lower confidence bound surrogates, which are posterior mean minus some weighting factor times the posterior standard deviation. Before we solve the auxiliary problem, we check its feasibility. If the auxiliary problem is infeasible, we declare infeasibility for the original controller tuning problem. (Xu et al., 2022a) gives a way to select $\beta_{i,t}^{1/2}$ so as to guarantee the global optimality for our algorithm.

Algorithm 1 CONFIG for Closed-loop Controller Tuning

for $t \in [T]$ **do**
  if $\max_{i\in[N]} \min_{\theta\in\Theta}(\mu_{i,t}(\theta)-\beta_{i,t}^{1/2}\sigma_{i,t}(\theta)) > 0$ **then**
    Declare infeasibility of (2).
  Update controller parameters with

$$\theta_t \in \arg\min_{\theta\in\Theta} \quad \mu_{0,t}(\theta) - \beta_{0,t}^{1/2}\sigma_{0,t}(\theta)$$

  subject to $\mu_{i,t}(\theta) - \beta_{i,t}^{1/2}\sigma_{i,t}(\theta) \leq 0, \forall i \in [N]$.
  Run closed-loop system to get measurements of J and $g_i, i \in [N]$.
  Update Gaussian process posterior mean and covariance with the new evaluations added.

We now give the global optimality guarantee of the Alg. 1.

Theorem 1. ((Xu et al., 2022a), Theorem 4.2). Let the same assumptions of Theorem 4.2 in (Xu et al., 2022a) hold. Then we have, with probability at least $1-\delta, \forall\delta \in (0,1)$, there exists $\tilde{\theta}_T \in \{\theta_1, \theta_2, \cdots, \theta_T\}$, such that,

$$J(\tilde{\theta}_T) - J^* \leq \mathcal{O}\left(\frac{\sum_{i=0}^N \gamma_{i,T}}{\sqrt{T}}\right), \tag{5a}$$

$$[g_i(\tilde{\theta}_T)]^+ \leq \mathcal{O}\left(\frac{\sum_{i=0}^N \gamma_{i,T}}{\sqrt{T}}\right), \quad \forall i \in [N], \tag{5b}$$

where J^* is the optimal value of the problem (2).

In the case that the original problem is infeasible, (Xu et al., 2022a, Theorem 5.1) also gives a guarantee that infeasibility is declared in finite steps with high probability.

4. EXPERIMENTS

In this section, we will first demonstrate the effectiveness of our algorithm on an artificial numerical problem. We then apply it to a classical constrained steady-state optimization problem of a continuous stirred-tank reactor.

4.1 *Artificial numerical problem*

We consider the constrained global optimization problem,

$$\min_{(\theta^1,\theta^2)\in[-10,10]^2} \quad \cos(2\theta^1)\cos(\theta^2) + \sin(\theta^1), \tag{6a}$$

$$\text{subject to} \quad \cos(\theta^1 + \theta^2) - g_{\text{thr}} \leq 0, \tag{6b}$$

where $g_{\text{thr}} \in (-1, 1)$. To measure the convergence performance of different algorithms, we introduce the concept of constrained regret as

$$\min_{t\in[T]} \quad [J(\theta_t) - J^*]^+ + \sum_{i=1}^N [g_i(\theta_t)]^+,$$

which is essentially the sum of the positive part of sub-optimality and the constraint violations. Fig. 2 shows the performance of different algorithms in terms of constrained regret, when running on the problem (6) with $g_{\text{thr}} = -0.6$

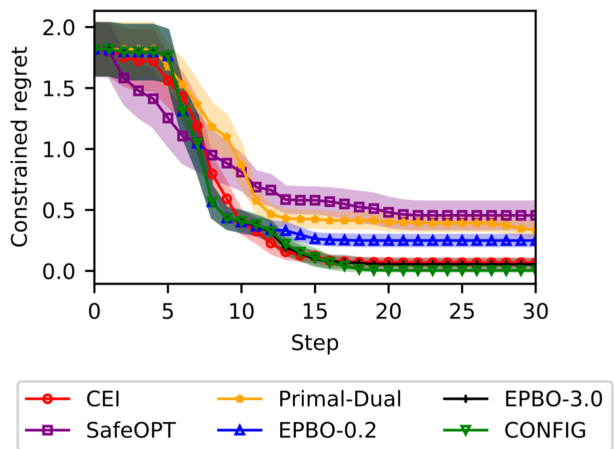

Fig. 2. Performance of different algorithms in terms of constrained regret with $g_{\text{thr}} = -0.6$ over 30 instances from randomly selected feasible starting points. Shaded area represents ± 0.3 standard deviation. EPBO-ρ represents EPBO (Lu and Paulson, 2022) method with the penalty coefficient ρ.

and 30 different feasible starting points. Our method can identify the global optimal solution within 20 steps in this particular example with all the 30 feasible starting points. In sharp contrast, SafeOPT and the primal-dual method can get stuck at a local minimum or infeasible solution, and suffer from a strictly positive average constrained regret. The popular CEI method is less efficient and can fail to give an almost optimal solution within 30 steps.

As for EPBO method (Lu and Paulson, 2022), the penalty can be tricky to tune since we do not have information on the gradient of the black-box functions. Randomly selecting a penalty may lead to a suboptimal or infeasible solution if it is not large enough. For example, as shown in Fig. 2, EPBO with penalty 0.2 converges to a strictly positive constrained regret, which means it can fail to find the global optimal feasible solution. On the other hand, too large penalty may also lead to numerical issues. Essentially, CONFIG method avoids the tuning effort of the penalty parameter and recovers the behavior of EPBO with sufficiently large penalty.

4.2 Steady state optimization: Williams-Otto problem

In this section, we consider the classical Williams-Otto benchmark problem (del Rio Chanona et al., 2021). In this problem, a continuous stirred-tank reactor (CSTR) is fed with two pure components A and B. The reactor operates at steady state under the temperature T_r. During the reaction, a byproduct G is also produced. We use X_A and X_G to denote the residual mass fractions of A and G at the reactor outlet, which need to be managed during the operation. See (del Rio Chanona et al., 2021, Sec. 4.1) for more details. Our goal is to tune the feedrate F_B of the component B and the reaction temperature T_r, so as to maximize the economic profit from the reaction. Our problem can then be formulated as,

$$\begin{aligned} \min_{F_\mathrm{B},T_\mathrm{r}} \quad & J(F_\mathrm{B},T_\mathrm{r}) \\ \text{subject to} \quad & \text{CSTR model (Mendoza et al., 2016)} \\ & g_1(F_\mathrm{B},T_\mathrm{r}) \stackrel{\text{def}}{=} X_\mathrm{A}(F_\mathrm{B},T_\mathrm{r}) - 0.12 \leq 0 \\ & g_2(F_\mathrm{B},T_\mathrm{r}) \stackrel{\text{def}}{=} X_\mathrm{G}(F_\mathrm{B},T_\mathrm{r}) - 0.08 \leq 0, \\ & F_\mathrm{B} \in [4,7], T_\mathrm{r} \in [70,100] \end{aligned} \tag{7}$$

where $J(F_\mathrm{B},T_\mathrm{r})$ is the minimization objective that is opposite to the economic profit, $g_1(F_\mathrm{B},T_\mathrm{r})$ and $g_2(F_\mathrm{B},T_\mathrm{r})$ are threshold constraints on the residual mass fractions at the reactor outlet. To measure the quality of a solution $\theta = (F_\mathrm{B},T_\mathrm{r})$, we use the normalized positive regret plus normalized violations as shown in

$$\frac{[J(\theta)-J^*]^+}{\sigma_J} + \frac{[g_1(\theta)]^+}{\sigma_{g_1}} + \frac{[g_2(\theta)]^+}{\sigma_{g_2}}, \tag{8}$$

where σ_J, σ_{g_1} and σ_{g_2} are standard deviations of J, g_1, g_2 over a set of sampled points.

Fig. 3 shows the average performance of different algorithms in log scale with 30 different random starting points. We observe that only the CEI method, EPBO-3.0 and CONFIG converge to below 1×10^{-10} within 20 steps, which means they effectively find the approximately constrained optimal solution within 20 steps. Since this example is a black-box, we do not know a feasible solution beforehand and thus, we do not show the result of

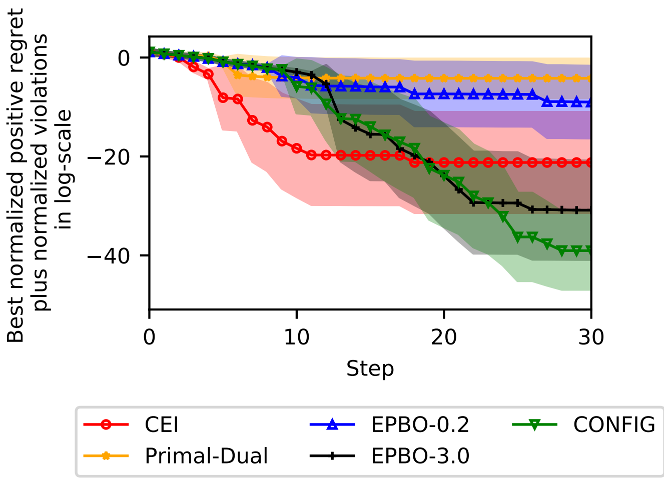

Fig. 3. Average performance of different algorithms in log scale from 30 different random starting points. The shaded area represents ±0.5 standard deviation.

SafeOPT, which requires an initial feasible solution. As for the EPBO method, when the penalty is set to be too small, it can converge to suboptimal or infeasible solution. Interestingly, with penalty parameter as 3.0, EPBO performs similarly to CONFIG in the first 20 steps, but achieves little improvement in the last 10 steps, while CONFIG further improves the solution.

5. COMPUTATIONAL TOOLBOX

To facilitate the future use of CONFIG in other applications, we build an open-source python package, config, which implements not only the CONFIG algorithm but also other popular constrained Bayesian optimization methods, to allow more flexibility in the choice of algorithms.

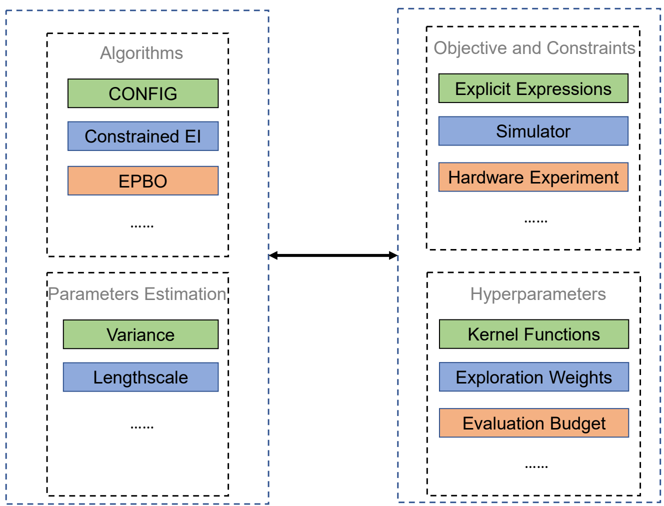

Fig. 4. An overview of the major components of our toolbox.

Fig. 4 shows an overview of the major components for our toolbox to work. Users need to specify the black-box optimization problems, where the objective and constraints can be explicit mathematical expressions, software simulator or hardware experiment. The users also need to specify the choices of some hyperparameters, including kernel functions and evaluation budget. If the user can not provide the parameters in the kernel functions based on domain knowledge, our toolbox also includes simple procedures to estimate the parameters in the kernel functions. The toolbox is extensible, and users can define their own acquisition policies easily using the interface provided.

6. CONCLUSION

In this paper, we have presented the application of the CONFIG algorithm, a simple and provably efficient constrained global optimization algorithm, to the optimization of the closed-loop control performance of an unknown system with unmodeled constraints. In sharp contrast to currently popular Gaussian process based closed-loop optimization methods, our CONFIG algorithm enjoys a theoretical global optimality guarantee. Simulation results on an artificial numerical problem and a classical constrained steady-state optimization problem of a continuous stirred-tank reactor both demonstrate that our CONFIG algorithm can achieve performance competitive with the popular CEI (Constrained Expected Improvement) algorithm, which has no known optimality guarantee. Our algorithm offers a new tool, with a provable global optimality guarantee, to optimize the closed-loop control performance for a system with non-safety critical unmodeled constraints. Last but not least, the open-source code is available as a python package to facilitate future applications.

ACKNOWLEDGEMENTS

This work was supported by the Swiss National Science Foundation under the NCCR Automation project (grant agreement 51NF40_180545), the RISK project (Risk Aware Data-Driven Demand Response, grant number 200021_175627), and by the Swiss Data Science Center (grant number C20-13).

REFERENCES

Baheri, A. and Vermillion, C. (2020). Waypoint optimization using Bayesian optimization: A case study in airborne wind energy systems. In *2020 Amer. Control Conf. (ACC)*, 5102–5017. IEEE.

Baheri, A. and Vermillion, C. (2017). Altitude optimization of airborne wind energy systems: A Bayesian optimization approach. In *2017 Amer. Control Conf. (ACC)*, 1365–1370. IEEE.

Bansal, S., Calandra, R., Xiao, T., Levine, S., and Tomlin, C.J. (2017). Goal-driven dynamics learning via Bayesian optimization. In *2017 IEEE 56th Annu. Conf. on Decis. and Control (CDC)*, 5168–5173. IEEE.

Chakrabarty, A., Danielson, C., Bortoff, S.A., and 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.

Chowdhury, S.R. and Gopalan, A. (2017). On kernelized multi-armed bandits. In *Int. Conf. on Mach. Learn.*, 844–853. PMLR.

del Rio Chanona, E.A., Petsagkourakis, P., Bradford, E., Graciano, J.A., and Chachuat, B. (2021). Real-time optimization meets Bayesian optimization and derivative-free optimization: A tale of modifier adaptation. *Comput. & Chem. Eng.*, 147, 107249.

Duivenvoorden, R.R., Berkenkamp, F., Carion, N., Krause, A., and Schoellig, A.P. (2017). Constrained Bayesian optimization with particle swarms for safe adaptive controller tuning. *IFAC-PapersOnLine*, 50(1), 11800–11807.

Fiducioso, M., Curi, S., Schumacher, B., Gwerder, M., and Krause, A. (2019). Safe contextual Bayesian optimization for sustainable room temperature PID control tuning. In *Proc. of the Twenty-Eighth Int. Joint Conf. on Artif. Intell.*, 5850–5856. IJCAI.

Frazier, P.I. (2018). A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*.

Gardner, J.R., Kusner, M.J., Xu, Z.E., Weinberger, K.Q., and Cunningham, J.P. (2014). Bayesian optimization with inequality constraints. In *Proc. of the Int. Conf. on Mach. Learn.*, volume 2014, 937–945.

Gelbart, M.A., Snoek, J., and Adams, R.P. (2014). Bayesian optimization with unknown constraints. In *Proc. of the 30th Conf. on Uncertainty in Artif. Intell.*, UAI'14, 250–259. AUAI Press, Arlington, Virginia, USA.

Jeong, S., Murayama, M., and Yamamoto, K. (2005). Efficient optimization design method using kriging model. *J. of Airc.*, 42(2), 413–420.

Jones, D.R., Schonlau, M., and Welch, W.J. (1998). Efficient global optimization of expensive black-box functions. *J. Global Optim.*, 13(4), 455–492.

Khosravi, M., Eichler, A., Schmid, N., Smith, R.S., and Heer, P. (2019). Controller tuning by Bayesian optimization an application to a heat pump. In *2019 18th Eur. Control Conf. (ECC)*, 1467–1472. IEEE.

König, C., Khosravi, M., Maier, M., Smith, R.S., Rupenyan, A., and Lygeros, J. (2020). Safety-aware cascade controller tuning using constrained Bayesian optimization. *arXiv preprint arXiv:2010.15211*.

Lederer, A., Capone, A., and Hirche, S. (2020). Parameter optimization for learning-based control of control-affine systems. In *Learn. for Dyn. and Control*, 465–475. PMLR.

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.

Mendoza, D.F., Graciano, J.E.A., dos Santos Liporace, F., and Le Roux, G.A.C. (2016). Assessing the reliability of different real-time optimization methodologies. *The Can. J. of Chem. Eng.*, 94(3), 485–497.

Pal, A., Zhu, L., Wang, Y., and Zhu, G.G. (2020). Multi-objective stochastic Bayesian optimization for iterative engine calibration. In *Proc. of the Amer. Control Conf.*, 4893–4898. IEEE.

Paulson, J.A. and Mesbah, A. (2020). Data-driven scenario optimization for automated controller tuning with probabilistic performance guarantees. *IEEE Contr. Syst. Lett.*, 5(4), 1477–1482.

Piga, D., Forgione, M., Formentin, S., and Bemporad, A. (2019). Performance-oriented model learning for data-driven mpc design. *IEEE Contr. Syst. Lett.*, 3(3), 577–

582.
Priem, R., Bartoli, N., Diouane, Y., and Sgueglia, A. (2020). Upper trust bound feasibility criterion for mixed constrained Bayesian optimization with application to aircraft design. *Aerosp. Sci. and Technol.*, 105, 105980.

Srinivas, N., Krause, A., Kakade, S.M., and Seeger, M.W. (2012). Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Trans. on Inf. Theory*, 58(5), 3250–3265.

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.*, 4781–4789.

Sui, Y., Gotovos, A., Burdick, J., and Krause, A. (2015). Safe exploration for optimization with Gaussian processes. In *Proc. of the Int. Conf. on Mach. Learn.*, 997–1005.

Turchetta, M., Berkenkamp, F., and Krause, A. (2020). Safe exploration for interactive machine learning. *Advances in Neural Inf. Process. Syst. 32*, 4, 2868–2878.

Xu, W., Jiang, Y., and Jones, C.N. (2022a). Constrained efficient global optimization of expensive black-box functions. *arXiv preprint arXiv:2211.00162.*

Xu, W., Jiang, Y., Maddalena, E.T., and Jones, C.N. (2022b). Lower bounds on the worst-case complexity of efficient global optimization. *arXiv preprint arXiv:2209.09655.*

Xu, W., Jones, C.N., Svetozarevic, B., Laughman, C.R., and Chakrabarty, A. (2022c). Vabo: Violation-aware bayesian optimization for closed-loop control performance optimization with unmodeled constraints. In *2022 American Control Conference (ACC)*, 5288–5293. IEEE.

Zhou, X. and Ji, B. (2022). On kernelized multi-armed bandits with constraints. *arXiv preprint arXiv:2203.15589.*