

LSR-BO: Local Search Region Constrained Bayesian Optimization for Performance Optimization of Vapor Compression Systems

Paulson, Joel A.; Sorouifar, Farshud; Laughman, Christopher R.; Chakrabarty, Ankush

TR2023-057 June 01, 2023

Abstract

Bayesian optimization (BO) has recently been demonstrated as a powerful tool for efficient derivative-free optimization of expensive black-box functions, such as those prevalent in performance optimization of complex energy systems. Classical BO algorithms ignore the relationship between consecutive optimizer candidates, resulting in jumps in the admissible search space which can lead to fail-safe mechanisms being triggered, or undesired transient dynamics that violate operational constraints. In this paper, we propose LSR-BO, a novel global optimization methodology that enforces local search region (LSR) constraints by design, which restricts how much the optimizer candidate can be changed at every iteration. We demonstrate that naively incorporating LSR constraints into BO causes the algorithm to get stuck in local sub-optimal solutions, and overcome this challenge through the development a novel exploration strategy that can gracefully navigate the trade-off between short-term “local”, and long-term “global”, performance improvement. Furthermore, we provide theoretical guarantees on the convergence of LSR-BO. Finally, we verify the effectiveness of our proposed LSR-BO method on an illustrative benchmark and a real-world energy minimization problem for a commercial vapor compression system.

American Control Conference (ACC) 2023

© 2023 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

LSR-BO: Local Search Region Constrained Bayesian Optimization for Performance Optimization of Vapor Compression Systems

Joel A. Paulson, Farshud Sorouifar, Christopher R. Laughman, and Ankush Chakrabarty

Abstract—Bayesian optimization (BO) has recently been demonstrated as a powerful tool for efficient derivative-free optimization of expensive black-box functions, such as those prevalent in performance optimization of complex energy systems. Classical BO algorithms ignore the relationship between consecutive optimizer candidates, resulting in jumps in the admissible search space which can lead to fail-safe mechanisms being triggered, or undesired transient dynamics that violate operational constraints. In this paper, we propose LSR-BO, a novel global optimization methodology that enforces local search region (LSR) constraints by design, which restricts how much the optimizer candidate can be changed at every iteration. We demonstrate that naively incorporating LSR constraints into BO causes the algorithm to get stuck in local sub-optimal solutions, and overcome this challenge through the development a novel exploration strategy that can gracefully navigate the trade-off between short-term “local”, and long-term “global”, performance improvement. Furthermore, we provide theoretical guarantees on the convergence of LSR-BO. Finally, we verify the effectiveness of our proposed LSR-BO method on an illustrative benchmark and a real-world energy minimization problem for a commercial vapor compression system.

Index Terms—Probabilistic machine learning, Black-box optimization, Bayesian methods, Energy systems.

I. INTRODUCTION

Closed-loop systems can be further optimized after deployment by adapting certain design parameters (e.g., controller gains, reference inputs, or internal model parameters) given some useful performance metric defined in terms of measured operational data. Manually tuning these control parameters can be highly inefficient and/or requires considerable task-specific expertise that cannot easily be generalized to new types of systems. This has motivated recent work on the development of *auto-tuning* algorithms that are capable of automatically adjusting the control parameters to achieve optimal performance as a way to save time, manual effort, and cost of experimentation [1]–[4].

Since the map between the control parameters and closed-loop performance function is often unmodeled (unknown) and may be highly nonlinear, it is common to treat this map as a black-box function with fully unknown structure. Therefore, we can treat auto-tuning as a black-box optimization problem, which can be tackled using any derivative-

free optimization (DFO) method; c.f. [5] for a detailed overview of DFO algorithms. However, since closed-loop experimentation or high-fidelity software simulations (like so-called digital twins) are needed to accurately represent true system behavior, the evaluation of which is expensive, auto-tuning algorithms must be designed to require as few experiments/evaluations as possible. Bayesian optimization (BO) is a *sample-efficient* DFO method that uses a probabilistic machine learning model to intelligently search over feasible parameter spaces [6], [7]. Due to its sample efficiency, BO has received significant attention in the context of auto-tuning including real-world applications in wind energy systems [8], engines [9], and space cooling [10]. In addition, BO has been shown to generalize well across a wide-variety of complex control problems such as controller parameter/reference setpoint tuning [11]–[13], cascade controller tuning [14], and MPC tuning [15]–[18].

An important challenge in the online deployment of such BO-based auto-tuning strategies is that they explore the search space without considering the distance between consecutive optimization candidates. BO thus has the tendency to suggest candidates at consecutive iterations that are far apart in the search space. While this is not a problem for static systems, these types of large “jumps” in control parameter values can induce aggressive dynamic behavior that leads to violation of safety constraints or triggers fail-safe mechanisms that limit performance. For example, energy consumption in HVAC systems is correlated with the electronic expansion valve (EEV) position setpoint. Aggressively exploring the search space is tantamount to aggressively opening and closing the EEV, which causes excessive wear and tear of the mechanical components of the valve and can produce oscillations in the refrigerant flow dynamics that result in compromised heating/cooling performance and can destabilize feedback loops that regulate the behavior of other system actuators (e.g., compressor speed).

In this paper, we propose a local search region (LSR) constrained Bayesian optimization method that ensures candidates suggested by classical BO are restricted to a predefined ball centered at the previous candidate. This limits the rate of change of the control parameter candidates at each iteration using so-called ‘LSR constraints’. Although one can naively enforce these LSR constraints at each BO step, we show that the resulting strategy often gets stuck in local sub-optimal solutions. To counteract this, we propose the novel ‘LSR-BO’ algorithm, that trades off short-term “local” improvement (i.e., take a nearby safe step when it has a good chance of improving upon our current best point) and long-term

J.A. Paulson and F. Sorouifar are with the Department of Chemical and Biomolecular Engineering, The Ohio State University, Columbus, OH 43210, USA. {sorouifar.1, paulson.82}@osu.edu

C. Laughman and A. Chakrabarty are with Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, USA. laughman@merl.com, achakrabarty@ieee.org

J.A. Paulson and F. Sorouifar were partially supported by the National Science Foundation Graduate Research Fellowship and Grant 2237616.

“global” improvement (i.e., take multiple safe steps towards a new region of the parameter space when we have exhausted local improvement). We also prove convergence of LSR-BO and describe sufficient conditions to ensure convergence to a global optimum.

To the best of the authors’ knowledge, our recent work [19] is the first and only paper to explicitly consider LSR constraints in the context of BO. The focus of [19] was to characterize the optimal BO sampling policy using dynamic programming (DP) and provide a tractable approximation. However, this type of strategy is significantly more computationally expensive than traditional BO and does not provide any convergence guarantees. As such, the main result of this paper is to develop a simple and convergent LSR-BO algorithm that can be straightforwardly implemented on real-world control applications.

In summary, our **major contributions** include:

- C1) development of a novel BO method, LSR-BO, that guarantees satisfaction of local “safety” constraints, while balancing local and global performance improvement;
- C2) a proof of convergence of the proposed LSR-BO method through a simple-to-enforce rolling terminal constraint that ensures dense sampling in the search space; and,
- C3) demonstration of LSR-BO’s ability to ensure safe energy optimization of industrial heat pump systems using high-fidelity simulations.

The rest of the paper is organized as follows. In Section II, we describe the importance of the LSR-constrained BO problem, and provide our proposed solution in Section III. We provide guarantees on LSR-BO convergence in Section IV and demonstrate that the method outperforms other constraint-enforcing BO algorithms on a benchmark and real-world energy application in Section V. We present concluding remarks in Section VI.

II. PROBLEM STATEMENT

We consider stable closed-loop systems of the form

$$x_+ = f(x, \theta) \quad (1)$$

where $x \in \mathbb{R}^{n_x}$ and $x_+ \in \mathbb{R}^{n_x}$ denote the current and updated state of the system, respectively, and $\theta \in \Theta \subset \mathbb{R}^d$ denotes a set of control-relevant parameters to be assigned. The parameters θ can represent any tunable variable in the control policy, which includes parameter choices (e.g., controller gains), references (e.g., setpoints), and structural choices (e.g., turning on or off a component of the policy).

Assumption 1: The admissible set of control parameters Θ is known at design time. \triangleleft

To evaluate the performance of the closed-loop system (1), we assume the existence of a steady-state performance function $J : \Theta \rightarrow \mathbb{R}$ such that $J(\theta)$ effectively quantifies system performance for any fixed $\theta \in \Theta$, once the closed-loop system achieves steady-state. The closed-loop performance optimization task then requires solving for

$$\theta^* \in \operatorname{argmax}_{\theta \in \Theta} J(\theta), \quad (2)$$

where θ^* denotes the globally optimal control policy parameters. The characteristics of the function J and domain Θ restrict the types of algorithms that can be used to solve (2). In many real-world applications, the structure of J is very often unmodeled and could potentially be highly non-convex in nature. This is especially true when attempting to solve (2) using closed-loop data collected from an online experimental system wherein the relationship between θ and $J(\theta)$ contains the dynamics (1), and therefore, may be complicated.

A more general representation of (2) is to treat J as a *black-box* function, implying we can only “learn” about it by evaluating $J(\theta)$ at specific θ values. The following assumption, which is trivially satisfied by the energy systems of interest in this work, greatly limits the amount of data that can be collected.

Assumption 2: The performance function $J(\theta)$ can be evaluated at any selected $\theta \in \Theta$ but each evaluation requires significant *expenditure*. \triangleleft

Expenditure can manifest itself in many forms. For example, $J(\theta)$ might take a long time to simulate, specialized sensors/resources for experimental, or significant manual labor to implement and evaluate. Without additional assumptions, we would need to apply a brute-force search to solve (2) since $J(\theta)$ could be highly discontinuous at every point $\theta \in \Theta$ (in the worst-case). Fortunately, most real-world problems demonstrate some degree of continuity, which is a useful property for constructing a more efficient search algorithm. Therefore, we make one additional assumption about the performance function:

Assumption 3: $J(\theta)$ is bounded, continuous, and at least once differentiable at every point $\theta \in \Theta$. \triangleleft

Assumption 3 is not strong, as it requires that the closed-loop system (1) must be globally asymptotically stable to a parameter-dependent equilibrium state $x^\infty(\theta)$, and that the map $x^\infty(\cdot)$ is continuous on Θ . This property can be (at least locally) ensured by proper design of the control policy. Continuity of $x^\infty(\theta)$ is sufficient to guarantee continuity of $J(\theta)$, as long as the closed-loop performance is defined by a continuous transformation of the equilibrium states.

Recent advances in Bayesian optimization (BO) have shown that one can solve (2) under Assumptions 1–3 by iteratively constructing a probabilistic surrogate representation of $J(\theta)$ and exploiting the predicted uncertainty estimates to efficiently search over Θ . Although BO represents a powerful framework for solving (2), it has the tendency to produce consecutive samples θ_k and θ_{k+1} that are a large distance apart to ensure efficient exploration of Θ . However, in many practical applications, one cannot search across arbitrarily large distances of Θ within one iteration of BO. In the context of closed-loop systems (1), large changes in the control parameter candidates can induce aggressive dynamics causing safety constraint violation and/or the triggering of fail-safe mechanism that limit performance quality.

We propose to avoid these undesired behaviors by restricting the next candidate sample θ_{k+1} to be within a domain-informed safe neighborhood of the current candidate θ_k .

Concretely, the constraint is given by

$$\theta_{k+1} \in \mathcal{B}_\delta(\theta_k) \quad \forall k \in \{0, \dots, N-1\}, \quad (3a)$$

where

$$\mathcal{B}_\delta(\theta) = \{\theta_+ \in \Theta : \|\theta_+ - \theta\|_p \leq \delta\}, \quad (3b)$$

denotes a p -norm ball of radius $\delta > 0$ centered at θ . We refer to the proposed constraints (3) as local search region (LSR) constraints. Traditional BO methods are not equipped to handle such restrictions on their exploration capabilities due to their greedy ‘‘one-step’’ nature. In addition, we would like to have some assurances on its convergence properties to avoid the designed sequence $\{\theta_1, \theta_2, \dots\}$ from getting stuck in a local maximum, which is discussed in Section IV.

III. BAYESIAN OPTIMIZATION WITH LOCAL SEARCH REGION CONSTRAINTS (LSR-BO)

A. Bayesian Optimization Preliminaries

In Bayesian optimization, one models the black-box (unknown) performance function $J(\cdot)$ as a realization of a Gaussian process (GP) [20]. To define a GP model, we must define a prior distribution π , which is uniquely specified by a prior mean function $\mu(\theta) = \mathbb{E}_\pi [J(\theta)]$ and prior kernel covariance function $\kappa(\theta, \theta') = \mathbb{E}_\pi [(J(\theta) - \mu(\theta))(J(\theta') - \mu(\theta'))]$. Given the prior π along with a current set of observations $\mathcal{D}_k = \{(\theta_i, J(\theta_i))\}_{i=1}^k$, we may derive the posterior distribution for $J(\theta)$ at any future test point as follows

$$J(\theta) | \mathcal{D}_k \sim \mathcal{N}(\mu_k(\theta), \sigma_k^2(\theta)), \quad (4)$$

where

$$\begin{aligned} \mu_k(\theta) &= \mu(\theta) + \mathbf{k}_k^\top(\theta) \mathbf{K}_k^{-1} (\mathbf{J}_k - \boldsymbol{\mu}_k), \\ \sigma_k^2(\theta) &= \kappa(\theta, \theta) - \mathbf{k}_k^\top(\theta) \mathbf{K}_k^{-1} \mathbf{k}_k(\theta), \end{aligned}$$

and

$$\begin{aligned} \mathbf{k}_k(\theta) &= [\kappa(\theta_1, \theta), \dots, \kappa(\theta_k, \theta)]^\top && \in \mathbb{R}^{k \times 1}, \\ \mathbf{J}_k &= [J(\theta_1), \dots, J(\theta_k)]^\top && \in \mathbb{R}^{k \times 1}, \\ \boldsymbol{\mu}_k &= [\mu(\theta_1), \dots, \mu(\theta_k)]^\top && \in \mathbb{R}^{k \times 1}, \\ \mathbf{K}_k &= [\kappa(\theta_i, \theta_j)]_{(i,j) \in \{1, \dots, k\} \times \{1, \dots, k\}} && \in \mathbb{R}^{k \times k}. \end{aligned}$$

The choice of the prior π determines the properties of the fitted functions and plays an important role in the accuracy of the GP; interested readers are referred to [20] for details on how to select μ and κ . We will focus on covariance functions belonging to the Matern class that have a parameter ν that controls the degree of smoothness of the performance function (i.e., J has is at least $\lceil \nu - 1 \rceil$ times differentiable).

Traditional BO methods take advantage of the statistical information embedded in the GP approximations to intelligently explore the search space Θ by defining a corresponding acquisition function. Ideally, the value of the acquisition function at every $\theta \in \Theta$ should provide a good measure of the potential benefit of querying J at that point in the future. Several acquisition functions have been proposed in

the literature, with one of the most popular being the so-called expected improvement (EI) function, given by

$$\text{EI}_k(\theta) = \mathbb{E}_\pi [(J(\theta) - J_k^*)^+ | \mathcal{D}_k],$$

where $a^+ = \max(a, 0)$ and $J_k^* = \max_{(\theta, J(\theta)) \in \mathcal{D}_k} J(\theta)$ is the incumbent solution that corresponds to the best objective value observed so far. The following closed-form expression for EI has been derived for GP approximators [21]:

$$\text{EI}_k(\theta) = \rho(\mu_k(\theta) - J_k^*, \sigma_k(\theta)),$$

where

$$\rho(y, s) = \begin{cases} y\Phi(y/s) + s\phi(y/s), & s > 0, \\ \max(y, 0), & s = 0, \end{cases} \quad (5)$$

and Φ and ϕ denote the cumulative density function and probability density function of a zero-mean unit-variance normal distribution, respectively. Since EI_k is inexpensive, it can be maximized using standard techniques and traditional BO proceeds by selecting $\theta_{k+1} \in \arg\max_{\theta \in \Theta} \text{EI}_k(\theta)$.

B. Proposed LSR-BO Algorithm

Herein, we describe the LSR-BO algorithm, which hinges on modifying the EI acquisition function to automatically trade-off local and global information predicted by the GP.

We compute a local optimizer for the EI acquisition over the local constraint region (3) to determine the nearest point $\theta_{k+1}^{\text{local}}$ that has the best chance of improving upon the incumbent solution. That is,

$$\theta_{k+1}^{\text{local}} \in \arg\max_{\theta \in \mathcal{B}_\delta(\theta_k)} \text{EI}_k(\theta). \quad (6)$$

By construction, $\theta_{k+1}^{\text{local}}$ must satisfy LSR constraints; however, it is possible that no nearby point is likely to improve upon the incumbent. This will be the case whenever the largest EI value in the local region is small, indicating this region has been sufficiently well-explored and, thus, $\theta_{k+1}^{\text{local}}$ is not a good sample choice. In such situations, an alternative strategy is to globally optimize EI to find the best possible point $\theta_{k+1}^{\text{global}}$ in the admissible search space Θ , given by

$$\theta_{k+1}^{\text{global}} \in \arg\max_{\theta \in \Theta} \text{EI}_k(\theta). \quad (7)$$

Since $\theta_{k+1}^{\text{global}}$ may not satisfy (3), we cannot directly evaluate $J(\theta_{k+1}^{\text{global}})$. The most straightforward approach to guarantee constraint satisfaction is to project $\theta_{k+1}^{\text{global}}$ onto $\mathcal{B}_\delta(\theta_k)$, which can be interpreted as taking the largest possible feasible step in the direction of the globally optimal sample. Using such an approach at every iteration would result in us missing out on potential nearby points that could lead to large possible improvements. Thus, the key idea in LSR-BO is to tradeoff between the ‘‘local’’ and ‘‘global’’ steps, respectively, defined by (6) and (7), which is achieved by

$$\theta_{k+1} = \begin{cases} \theta_{k+1}^{\text{local}}, & \text{EI}_k(\theta_{k+1}^{\text{local}}) \geq \gamma, \\ \text{Proj}_{\mathcal{B}_\delta(\theta_k)}(\theta_{k+1}^{\text{global}}), & \text{otherwise,} \end{cases} \quad (8)$$

where γ is a user-defined tuning parameter that controls the degree of local search. Based on the properties of the EI

function, (8) reduces to greedy local and a global search as $\gamma \rightarrow 0$ and $\gamma \rightarrow \infty$, respectively. We found that a reasonably small value of $\gamma = 0.1$ provided good results in all of our case studies (Section V); we plan on systematically studying the effect of γ on performance in future work.

Since the GP model (4) is likely imprecise with only a few data points, we assume that the initial points are chosen in a random quasi-uniform fashion. This initial dataset, denoted by \mathcal{D}_0 , is assumed to be generated offline in the sense that the experiments can be reset and, therefore, the points do not have to satisfy (3).

We can now formally define the sequential LSR-BO sample selection strategy.

Definition 1: The LSR-BO(π) strategy selects a sequence of sample points $\{\theta_1, \theta_2, \dots\}$ for a given prior π by selecting θ_{k+1} according to (8) for all $k \geq 0$. \triangleleft

This approach is extremely simple to implement, with computational and memory requirements on the same order as classical BO methods. Another important advantage of this approach is we can guarantee the incumbent converges to the global solution, i.e., $J_k^* \rightarrow J(\theta^*)$ as $k \rightarrow \infty$, as shown in the following section. Equivalently, the simple regret $J(\theta^*) - J_k^*$ decays to zero asymptotically.

IV. CONVERGENCE PROPERTIES OF LSR-BO

The convergence properties of EI (in the absence of LSR constraints) have been analyzed in [22], which requires a non-degeneracy assumption on the covariance function κ .

Definition 2: A GP model (4) with continuous, positive-definite covariance kernel κ has the no-empty-ball (NEB) property if the following holds:

$$\inf_{\substack{k \in \mathbb{N} \\ \theta_1, \dots, \theta_k \in \Theta \\ \|\theta_i - \theta_0\| \geq \epsilon, \forall i}} \sigma_k^2(\theta_0) > 0$$

for any sequence $\{\theta_k\}_{k \geq 1}$, any $\epsilon > 0$, and any $\theta_0 \in \Theta$. \triangleleft

The NEB property directly implies $\sigma_k^2(\theta_0) \rightarrow 0$ when θ_0 is an adherent point of the set $\{\theta_k\}_{k \geq 1}$, which is equivalent to saying that, if the prediction error covariance $\sigma_k^2(\theta_0)$ goes to zero, then there can be no ‘‘empty ball’’ centered at θ_0 . As shown in [22, Proposition 10], kernel functions from the Matérn class have the NEB property and such kernels can be used to reconstruct performance functions that satisfy Assumption 3. The following result can then be established under the NEB property.

Lemma 1 (Global EI Convergence [22]): Let κ be a kernel that exhibits the NEB property. For any initial set $\theta_1, \dots, \theta_{k_0} \in \Theta$, the sequence $\{\theta_k\}_{k \geq 1}$ generated by $\theta_{k+1} = \theta_{k+1}^{\text{global}}$ for all $k > k_0$, where $\theta_{k+1}^{\text{global}}$ is given by (7), is almost surely (a.s.) dense in Θ .

Note that convergence of the ‘global’ EI strategy is an immediate consequence of the kernel allowing dense sampling of Θ . That is, $\max_{i \in \{1, \dots, k\}} J(\theta_i) \rightarrow \max_{\theta \in \Theta} J(\theta)$ as $k \rightarrow \infty$ (almost surely). This result immediately carries over to the LSR-BO(π) at the limit $\delta \rightarrow \infty$. In particular, as the radius of the LSR constraints becomes arbitrarily large, (8) naturally reduces to $\theta_{k+1} = \theta_{k+1}^{\text{global}}$ since we are able to jump

across Θ arbitrarily far. However, as we show next, the same concept further generalizes to all positive δ values under a certain assumption that can be easily satisfied through a slight modification to (8).

Theorem 1 (LSR-BO Convergence): Let κ exhibit the NEB property. Suppose there exist a sequence of integers $\{m_k\}_{k \geq 1}$ and a finite integer $M < \infty$ such that the sequence generated by LSR-BO(π) satisfies $\theta_{m_{k+1}} = \theta_{m_{k+1}}^{\text{global}}$ with $|m_{k+1} - m_k| \leq M$ for all $k \geq 0$. Then, the LSR-BO(π) strategy generates a dense set of sample points in Θ (a.s.). Consequently, $J_k^* \rightarrow J(\theta^*)$ as $k \rightarrow \infty$ (a.s.).

Proof: By assumption, we can always take a subsequence $\{\theta_{m_0}^{\text{global}}, \dots, \theta_{m_k}^{\text{global}}\}$ from the sequence generated by LSR-BO(π) for any $k \geq 1$. These global samples can be at most M steps apart in this sequence, meaning at most M samples are added into the GP model between k and $k+1$. By treating all of these additional (non-global) sample points as ‘‘initial’’ points, we can then directly apply Lemma 1 to infer that the LSR-BO(π) sequence provides dense sampling in Θ a.s. as $k \rightarrow \infty$ (or equivalently $m_k \rightarrow \infty$). \square

The assumption made about the existence of $\{m_k\}$ deserves some further elaboration. This sequence assumption allows us to exploit the properties of the global EI policy in a straightforward manner; however, we have not shown that LSR-BO(π) always generates such a sequence. Due to the potential to switch back and forth between local and global steps in (8), we do not attempt to formally prove that this sequence exists. Instead, we propose a simple modification such that the sequence assumption is satisfied by construction. To this end, we employ a ‘‘terminal constraint’’ that enforces $\theta_{k+M} = \theta_{k+1}^{\text{global}}$ for all $k \in \{0, M, 2M, \dots\}$, while still satisfying LSR constraints. This can be enforced by doing one more projection step after (8) as follows

$$\theta_{k+1} \leftarrow \text{Proj}_{\mathcal{T}(n(M,k), \theta_k, \theta_{k+1}^{\text{global}})}(\theta_{k+1}), \quad (9)$$

where $\mathcal{T}(j, \theta_0, \theta_f)$ denotes the set of all next samples that satisfy LSR constraints and the terminal condition given j steps remaining, initial sample θ_0 , and final sample θ_f , and $n(M, k) = M - 1 - \text{mod}(k - 1, M)$ is the number of steps remaining until step M is reached in current interval k . The choice of M would need to be at least as large as the minimum number of steps that it takes to move between any two points in Θ , which is finite for any $\delta > 0$. Although this modified version of LSR-BO(π) will converge independently of M , the rate of convergence is expected to depend on M , which we plan to study more in future work.

Remark 1: Note that the terminal constraint is a *sufficient condition* for convergence; we found the assumptions in Theorem 1 were automatically satisfied by the original LSR-BO(π) algorithm with $\gamma > 0$ in practice.

V. CASE STUDIES

A. Example 1: Illustrative 2-D Benchmark Problem

We first test our proposed LSR-BO method on a modified version of the Branin function from [19], which is a standard benchmark problem for global optimization algorithms due

to its non-convex nature with multiple local optima and one global optimum.

The compact domain is given by $\Theta = [-5, 10] \times [0, 15]$, while we also incorporate the following LSR constraints to mimic safe exploration in this problem

$$|\theta_{1,k+1} - \theta_{1,k}| \leq 0.5\delta, \quad |\theta_{2,k+1} - \theta_{2,k}| \leq 1.5\delta, \quad (10)$$

that must hold for all $k \geq 0$. These constraints force the next sample to be within a hyperrectangle of the current sample, which can be straightforwardly mapped to a hypercube of the form (3) with $p = \infty$.

Our proposed LSR-BO algorithm requires the specification of the prior GP functions as well as a numerical procedure to optimize the EI acquisition functions. For the prior, we selected $\mu(\theta) = 0$ and a Matern-2.5 function for κ . The length- and output-scale parameters were estimated at each iteration using the MLE approach summarized in [20, Section 2]. We use a multi-start version of the L-BFGS-B algorithm [23] to optimize the EI functions in (6) and (7) wherein the best 10 points from an initial set of 1000 samples (ranked by their EI values) are used as the starting points.

We use so-called simple regret to measure the performance of the generated sequence of performance function evaluations $\mathcal{D}_{k+1} = \mathcal{D}_k \cup \{(\theta_{k+1}, J(\theta_{k+1}))\}$, i.e., $S_k(\mathcal{D}_0) = J(\theta^*) - J_k^*$ for any iteration $k \geq 0$, which depends on the initial dataset \mathcal{D}_0 . To illustrate that LSR-BO is capable of performing well across a wide array of starting points, we randomly generate 50 different \mathcal{D}_0 composed of 10 samples generated in quasi-uniform random fashion inside of Θ . The starting value θ_0 is always selected to the point in \mathcal{D}_0 that maximizes the performance function. We use these replicate runs to estimate statistical properties of the sequences generated by the LSR-BO algorithm. To highlight the advantages of the LSR-BO, we compare its performance to the following three baseline algorithms that have all been constructed to satisfy the LSR constraints by design.

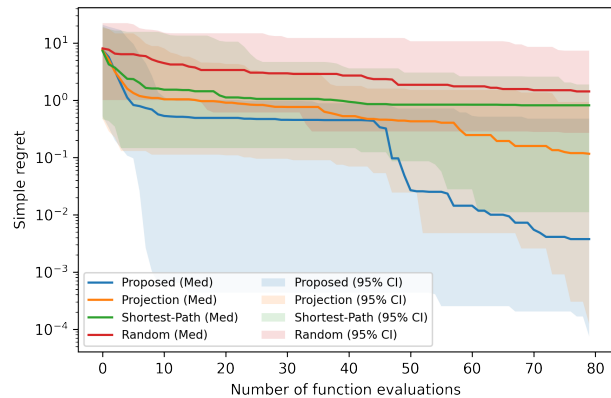
Shortest-Path: The global EI problem is solved to generate the next desired sample. If this point does not satisfy LSR constraints, we sample the sequence of points along the shortest path between the current and desired global point.

Projection: We sample by projecting the global EI point onto the LSR constraints, i.e., $\theta_{k+1} = \text{Proj}_{\mathcal{B}_\delta(\theta_k)}(\theta_{k+1}^{\text{global}})$. We can think of this as a special case of the LSR-BO sampling method (8), when $\gamma \rightarrow \infty$.


Random: The procedure is the same as the shortest-path algorithm, except that the desired point is selected uniformly at random. This approach is thus a simple modification to random search to ensure satisfaction of LSR constraints.

The statistical results for the simple regret $S_k(\mathcal{D}_0)$ for all methods for up to $k = 80$ iterations is shown in Fig. 1 for LSR constraint radius values $\delta = 0.5, 1.5$ to illustrate the effect of the radius of the LSR on the performance of the LSR-BO algorithm. We see that LSR-BO outperforms all considered methods in both cases; the regret decay is consistently faster and the final solution better than the competitor algorithms. In addition, the improvements become even more pronounced as the size of the LSR constraint

region is reduced. It is interesting to observe that the LSR-BO algorithm initially converges quickly to a local optimal value (within the first 10–20 iterations) and then expands out from the local region to achieve the global optimal value.



(a) $\delta = 0.5$



(b) $\delta = 1.5$

Fig. 1: Comparison of simple regret statistics (median and 95% confidence interval) versus number of iterations for the proposed (LSR-BO), shortest-path, projection, and random search algorithms for two LSR constraint radius values δ .

B. Example 2: Industrial Heat Pump Energy Optimization

We also study a real-world problem of tuning setpoints of a vapor-compression heat pump to minimize the operating power consumption. A block diagram of the heat pump of interest is shown in Fig. 2A, which consists of a compressor, a condenser, an expansion valve, and an evaporator that exchanges heat between an indoor occupied setting and the ambient environment. We can cast this problem in the form of (2) by defining the performance function $J(\theta) = -P_\infty(\theta)$ as the negative of the steady-state power consumption P_∞ that is a function of the heat pump setpoints θ . In the system of interest, we first close a feedback loop from compressor frequency to room temperature, which leaves us with three tunable setpoints θ representing the electrical expansion valve (EEV) position, the indoor fan speed (IFS), and the outdoor fan speed (OFS). By assigning fixed setpoint values θ , we wait for an adequate amount of time until the

power signal resides within a 95% settling zone and use that to represent $J(\theta)$. The setpoints must lie in the following known ranges: EEV $\in [200, 300]$ counts, IFS $\in [200, 500]$ rpm, and OFS $\in [500, 1000]$ rpm.

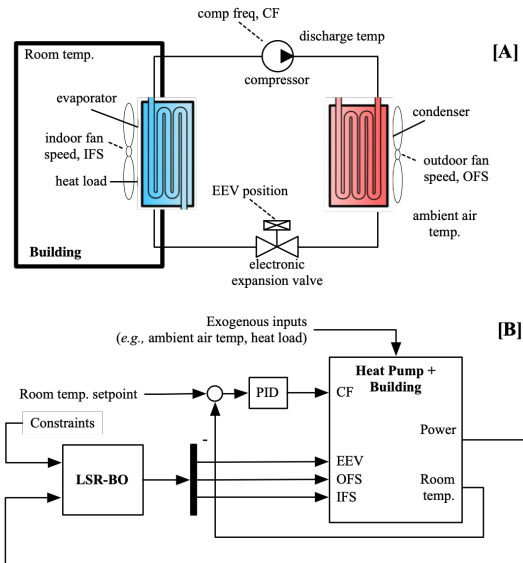


Fig. 2: (A) Schematic of industrial heat pump system. (B) Illustration of the proposed LSR-BO method for minimizing power consumption through manipulation of setpoints.

Reckless alterations of the setpoints can result in wear and tear of the actuators, and lead to undesired transients in the refrigerant flow and thermal dynamics. This prevents the direct use of traditional BO methods that demonstrate significant jumps in θ values between consecutive iterations. We avoid these harmful operational modes by incorporating LSR constraints (3) that limit the change in EEV position to ± 5 counts, IFS to ± 10 rpm, and OFS to ± 25 rpm. An illustrative schematic of the proposed data-driven optimization approach for the heat pump system is shown in Fig. 2B.

Implementation Details: We use a high-fidelity dynamic model of a prototype vapor-compression system (VCS)¹ written in the Modelica language [24] to collect power consumption data and optimize the set-points on-the-fly. A complete description of the model is available in [10]. The model was first developed in the Dymola [25] environment and then exported as a functional mockup unit (FMU) [26]. Its current version is comprised of a total of 12,114 differential equations. We use the same GP model and optimization settings as discussed in Example 1.

Results and Discussion: We again compare the LSR-BO method to the shortest-path and projection methods defined previously. Since the true global optimum is unknown in this problem, we directly compare the power consumption values produced over 40 allowed high-fidelity function evaluations. Similar to Example 1, we initially populate \mathcal{D}_0 with high-fidelity evaluations at 10 domain-informed samples in Θ ,

¹Note that while the behavior of this model have been validated against a real VCS, the numerical values and/or performance presented in this work is not representative of any product.

and select θ_0 as the point within this set that produces the largest $J(\theta)$ (minimal power consumption). We again repeat all algorithms 50 times to estimate statistical properties of the generated sample trajectories. The resulting minimal power consumption profiles, i.e., $\min_{(\theta), P_{\infty}(\theta) \in \mathcal{D}_k} P_{\infty}(\theta)$ for the LSR-BO, shortest-path, and projection algorithms are shown in Fig. 3. We see that that the proposed method achieves tighter confidence intervals, implying LSR-BO can reliably find setpoint values that lead to lower power consumption in a shorter amount of time. The complete dynamic simulation profiles for the VCS for the median LSR-BO run are also shown in Fig. 4. Note that the grey shading in the figure denotes the ‘offline’ experiments on the heat pump system: they are actually done online, which is why they appear in the trace, but with manual safety configurations running to avoid deleterious behavior. This is why there are LSR constraint violations in the grey shaded areas, and after that, the LSR constraints are consistently satisfied. Furthermore, we see a clear reduction in the power consumption over time, without generating any intermediate harmful transient states. We also note that the sensitivity of the power varies among the actuators; because the power is more sensitive to the EEV position, the EEV tends to converge to a steady-state value as the power is minimized. In comparison, the lower sensitivity of the power to the fan speeds, and in particular OFS, causes these actuators to continue to vary after power has converged to a minimum value because these variations do not translate to a significant change in the power consumption.

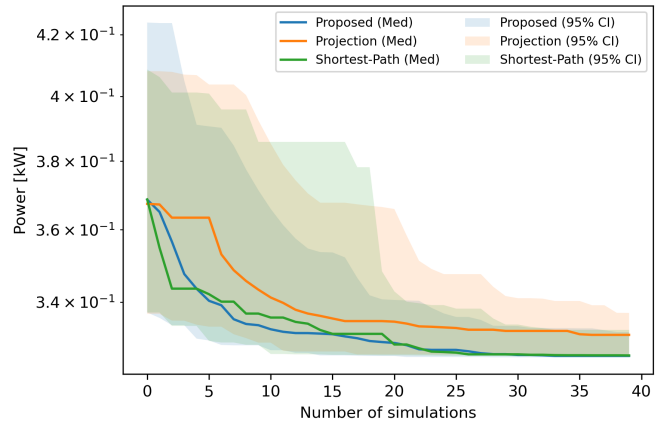


Fig. 3: Comparison of minimum power sequence statistics (median and 95% confidence interval) versus number of iterations for proposed (LSR-BO) method, shortest-path, and projection algorithms on real-world VCS problem.

VI. CONCLUSIONS

This paper presents a new, simple-to-implement variation of Bayesian optimization (BO) suitable for global optimization of expensive-to-evaluate objective functions with local search rate (LSR) constraints imposed on the sequence of candidate points. The newly added LSR constraints must be considered when there is a significant cost to changing

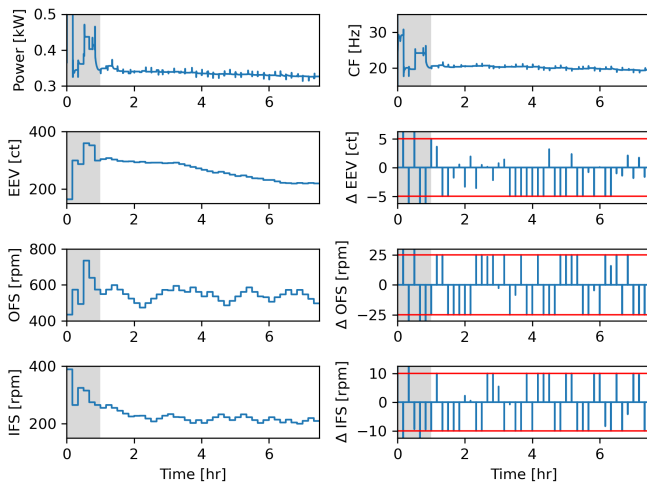


Fig. 4: Heat pump power, compressor frequency (CF), EEV position, OFS, and IFS values over time for median LSR-BO run. Δ EEV, Δ OFS, and Δ IFS represent the change in setpoints at each iteration with the red lines denoting the LSR constraints. The grey shaded regions denote the offline experiments used to initialize the GP model.

the design variables at every optimization iteration. We are particularly interested in closed-loop performance optimization of energy systems, which require limited alterations to control design parameters to avoid pushing the system into unsafe operating modes. The proposed method, LSR-BO, restricts how much the candidate point can be changed at each iteration while also avoiding getting stuck in sub-optimal local solutions. In addition to having computational and memory requirements on the same order as classical BO, we prove that LSR-BO inherits its established convergence properties. We compare the efficiency of LSR-BO to three competing alternative methods on two challenging case studies including one derived from a commercial vapor-compression system. Further, we analyze the statistical properties of the performance evaluation sequences generated by all methods under different random initializations and show that LSR-BO consistently generates better quality solutions including up to multiple orders of magnitude better median and 95% confidence interval performance values. Future work will include the development of multi-objective and constrained extensions of LSR-BO as well as theoretical guarantees on the expected and worst-case rate of convergence.

REFERENCES

- [1] J. L. Garriga and M. Soroush, "Model predictive control tuning methods: A review," *Industrial & Engineering Chemistry Research*, vol. 49, no. 8, pp. 3505–3515, 2010.
- [2] J. A. Paulson and A. Mesbah, "Data-driven scenario optimization for automated controller tuning with probabilistic performance guarantees," *IEEE Control Systems Letters*, vol. 5, no. 4, pp. 1477–1482, 2020.
- [3] J. A. Paulson, K. Shao, and A. Mesbah, "Probabilistically robust Bayesian optimization for data-driven design of arbitrary controllers with Gaussian process emulators," in *Proceedings of the Conference on Decision and Control*, pp. 3633–3639, IEEE, 2021.

- [4] 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 *Proc. of the American Control Conference*, pp. 5288–5293, IEEE, 2022.
- [5] L. M. Rios and N. V. Sahinidis, "Derivative-free optimization: a review of algorithms and comparison of software implementations," *Journal of Global Optimization*, vol. 56, no. 3, pp. 1247–1293, 2013.
- [6] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2015.
- [7] P. I. Frazier, "A tutorial on Bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.
- [8] A. Baheri and C. Vermillion, "Waypoint optimization using Bayesian optimization: A case study in airborne wind energy systems," in *Proceedings of the American Control Conference*, pp. 5102–5017, IEEE, 2020.
- [9] A. Pal, L. Zhu, Y. Wang, and G. G. Zhu, "Multi-objective stochastic Bayesian optimization for iterative engine calibration," in *Proceedings of the American Control Conference*, pp. 4893–4898, IEEE, 2020.
- [10] A. Chakrabarty, C. Danielson, S. A. Bortoff, and C. R. Laughman, "Accelerating self-optimization control of refrigerant cycles with Bayesian optimization and adaptive moment estimation," *Applied Thermal Engineering*, vol. 197, p. 117335, 2021.
- [11] R. R. Duivendoorn, F. Berkenkamp, N. Carion, A. Krause, and A. P. Schoellig, "Constrained Bayesian optimization with particle swarms for safe adaptive controller tuning," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 11800–11807, 2017.
- [12] M. Khosravi, A. Eichler, N. Schmid, R. S. Smith, and P. Heer, "Controller tuning by Bayesian optimization an application to a heat pump," in *Proceedings of the European Control Conference*, pp. 1467–1472, IEEE, 2019.
- [13] A. Lederer, A. Capone, and S. Hirche, "Parameter optimization for learning-based control of control-affine systems," in *Learning for Dynamics and Control*, pp. 465–475, PMLR, 2020.
- [14] M. Khosravi, V. N. Behrunani, P. Myszkowski, R. S. Smith, A. Rupenyan, and J. Lygeros, "Performance-driven cascade controller tuning with Bayesian optimization," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 1, pp. 1032–1042, 2021.
- [15] D. Piga, M. Forgiione, S. Formentin, and A. Bemporad, "Performance-oriented model learning for data-driven MPC design," *IEEE Control Systems Letters*, vol. 3, no. 3, pp. 577–582, 2019.
- [16] F. Soroufifar, G. Makrygiorgos, A. Mesbah, and J. A. Paulson, "A data-driven automatic tuning method for MPC under uncertainty using constrained Bayesian optimization," *IFAC-PapersOnLine*, vol. 54, no. 3, pp. 243–250, 2021.
- [17] Q. Lu, L. D. González, R. Kumar, and V. M. Zavala, "Bayesian optimization with reference models: A case study in MPC for HVAC central plants," *Computers & Chemical Engineering*, vol. 154, p. 107491, 2021.
- [18] J. A. Paulson, G. Makrygiorgos, and A. Mesbah, "Adversarially robust Bayesian optimization for efficient auto-tuning of generic control structures under uncertainty," *AIChE Journal*, vol. 68, no. 6, p. e17591, 2022.
- [19] J. A. Paulson, F. Soroufifar, and A. Chakrabarty, "Efficient multi-step lookahead Bayesian optimization with local search constraints," in *Proceedings of the Conference on Decision and Control*, IEEE, 2022 (accepted).
- [20] C. K. Williams and C. E. Rasmussen, *Gaussian Processes For Machine Learning*, vol. 2. MIT press Cambridge, MA, 2006.
- [21] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global Optimization*, vol. 13, no. 4, pp. 455–492, 1998.
- [22] E. Vazquez and J. Bect, "Convergence properties of the expected improvement algorithm with fixed mean and covariance functions," *Journal of Statistical Planning and Inference*, vol. 140, no. 11, pp. 3088–3095, 2010.
- [23] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Mathematical Programming*, vol. 45, no. 1, pp. 503–528, 1989.
- [24] Modelica Association, "Modelica specification, Version 3.4," 2017.
- [25] Dassault Systemes, "Dymola 2020," 2019.
- [26] Modelica Association, "Functional Mockup Interface for Model Exchange and Co-Simulation, Version 2.0.1," 2019.