Received: 22 May 2022 | Revised: 5 July 2022 | Accepted: 31 July 2022

DOI: 10.1002/aic.17857

RESEARCH ARTICLE

Process Systems Engineering

AIChE JOURNAL

Constrained robust Bayesian optimization of expensive noisy black-box functions with guaranteed regret bounds

Akshay Kudva | Farshud Sorourifar | Joel A. Paulson

Department of Chemical and Biomolecular Engineering, The Ohio State University, Columbus, Ohio, USA

Correspondence
Joel A. Paulson, Department of Chemical and Biomolecular Engineering, The Ohio State University, Columbus, OH 43210, USA.
Email: paulson.82@osu.edu

Funding information
National Science Foundation, Grant/Award Number: 2029282; NSF Graduate Research Fellowship

Abstract

Many real-world design problems involve optimization of expensive black-box functions. Bayesian optimization (BO) is a promising approach for solving such challenging problems using probabilistic surrogate models to systematically trade-off between exploitation and exploration of the design space. Although BO is often applied to unconstrained problems, it has recently been extended to the constrained setting. Current constrained BO methods, however, cannot identify solutions that are robust to unavoidable uncertainties. In this article, we propose a robust constrained BO method, constrained adversarially robust Bayesian optimization (CARBO), that addresses this challenge by jointly modeling the effect of the design variables and uncertainties on the unknown functions. Using exact penalty functions, we establish a bound on the number of CARBO iterations required to find a near-global robust solution and provide a rigorous proof of convergence. The advantages of CARBO are demonstrated on two case studies including a non-convex benchmark problem and a realistic bubble column reactor design problem.

KEYWORDS
black-box optimization, derivative-free optimization under uncertainty, exact penalty functions, Gaussian processes, robust Bayesian optimization

1 | INTRODUCTION

Uncertainty is inevitably present in real-world problems due to noisy and incomplete datasets, unknown parameters, environmental disturbances (such as product demand and prices), and implementation errors. Therefore, so-called "optimal" solutions found by solving some underlying nominal optimization problem, which neglects these many sources of uncertainty can often be suboptimal, or even worse, infeasible.[1,2] Local and global sensitivity-based methods[3–5] are often used to study the impact of any critical uncertainties on specific designs; however, such methods focus on robustness analysis amongst a small, finite set of designs, meaning they are unable to systematically synthesize new designs with improved robustness directly. There has been a significant amount of work on so-called "optimization under uncertainty," as summarized in several review articles.[6–9] Two main classes of methods for optimization under uncertainty are stochastic and robust optimization. Stochastic optimization[10,11] models the uncertainty in the form of probability distributions, which must be either estimated from data or selected according to expert opinion/prior knowledge. An important challenge with this approach, however, is that the true distribution of the uncertainties is rarely known in practice, which limits its applicability. Robust optimization,[12–15] on the other hand, is an alternative strategy that adopts a "min–max" approach by defining the *robust optimal design* as the one with the best worst-case performance subject to worst-case constraint satisfaction. There has been a significant amount of work on the

development of robust optimization theory and algorithms for both convex and non-convex problems. In fact, it is worth noting that robust optimization can be interpreted as a special category of generalized semi-infinite programs (GSIP),[16] which are optimization problems involving a finite number of decision variables subject to a potentially infinite number of constraints in which the feasible set of parameters may depend on the decision variables. Whenever the feasible parameter set is independent of the decision variables, then the problem falls into the category of a standard semi-infinite program (SIP) for which we refer interested readers to Refs. [17-21] Within the framework of robust optimization, uncertainty sets described by convex inequality constraints can be efficiently handled, as discussed in, for example, Refs. [22–24]. Several methods have also been developed for SIPs defined in terms of general non-convex functions including cutting-plane algorithms,[25–27] interval methods,[28] and interior point techniques[29] to name a few.

Although solving (G)SIPs with both explicit and implicit function representations has been an active area of research for more than three decades, the methods summarized above all require the structure of the objective and constraint functions to be known and exploited within the algorithm. However, obtaining equation-oriented models whose structure can be exploited by these methods for each and every component of a complex system is frequently not possible. For example, in addition to having many types of challenging nonlinearities, engineering design problems often require the use of *expensive* computer simulations or experiments, implying that the functional relationship between the design variables and uncertain parameters and the objective and constraints are unknown (often referred to as "simulation-based" or "black-box" models). Some examples include the choice of laboratory experiments in material and drug design,[30] calibration of expensive simulators to experimental data,[5] airfoil shape design,[31,32] hyperparameter tuning in machine learning algorithms,[33] and automated design of advanced multivariable control structures[34–36] to name a few. Due to lack of known structure in these cases, one must resort to so-called derivative-free optimization (DFO) methods,[37] which only require the black-box functions to be queried at specific input values such that these methods can be very generally applied.

One of the first methods developed for robust optimization of unconstrained simulation-based problems was presented in Ref. [1], which was extended to constrained problems in Ref. [2]. These methods are analogous to local search techniques (e.g., gradient descent), as they iteratively take steps along descent directions that preserve robust feasibility. They also require one to determine the set of possible worst-case uncertainty values at each iteration, which requires one to repeatedly solve the lower-level maximization problem. In Refs. [1,2] this lower-level problem is practically solved using multiple gradient ascents from different starting designs to improve efficiency. Therefore, under this assumption,[1,2] are not DFO methods since one must be able to extract gradient information from the simulator, which limits their applicability. A natural extension of these methods would be to replace the gradient ascent method with some DFO method, which can be broadly categorized as either stochastic or deterministic. Stochastic DFO methods, such as genetic algorithm[38] or particle swarm optimization,[39] utilize random numbers within the search process and are known to require a large number of function evaluations, which prevents them from being applied to expensive black-box functions. Deterministic DFO methods, on the other hand, are often motivated by optimization of expensive functions and, in particular, model-based DFO methods, which construct a surrogate model of the unknown functions to guide the search process, have become very popular in recent years. Model-based DFO methods mainly differ by their choice of *scale* (e.g., local versus global approximation) and *type* of function approximation technique (e.g., polynomial, neural network, or radial basis function models). The Bayesian optimization (BO) framework[40–43] is a *global* model-based DFO method that takes advantage of *probabilistic* data-driven surrogate models, and has been found to empirically perform well on a variety of problems in many different disciplines. BO methods typically rely on Gaussian process (GP) models since they are non-parametric (i.e., can represent virtually any function given enough data) and directly quantify uncertainty in future predictions in the form of probability distributions.[44] By combining these uncertain predictions from the GP model with an expected utility (or acquisition function), BO can systematically address the exploration–exploitation tradeoff in a way that boosts data efficiency relative to many purely deterministic surrogate-based DFO methods (see, e.g., Refs. [45–48]).

Even though BO has been successfully applied to many different nominal optimization problems, it has proved difficult to extend BO to the robust optimization setting due to the min–max problem structure, which involves two competing optimization stages. For example, the fact that one would have to execute multiple BO runs at each iteration of Refs. [1,2] would likely prevent this hybrid approach from working on expensive functions. As opposed to using established BO methods for lower-level maximization sub-problems, there has also been work on extending BO to work directly on robust min–max optimization problems. The MiMaReK algorithm[49] is one of the earliest attempts to develop a robust BO strategy, which uses a two-level expected improvement acquisition function. The main disadvantage of MiMaReK is that a set of uncertainty samples are sequentially constructed such that it requires a growing number of expensive function evaluations at each iteration. The probabilistically robust BO (PRBO) approach[50] is an alternative to MiMaReK that selects the minimum number of uncertainty samples required to a priori provide a probabilistic robustness certificate at each iteration. However, both MiMaReK and PRBO may require a large number of function evaluations at every iteration, which still limits their applicability. The adversarially robust BO (ARBO) method[51,52] attempts to overcome this limitation by *simultaneously* modeling the effect of the design variables and uncertainties on the objective function. ARBO only requires a single expensive function evaluation at every iteration, which can potentially result in a drastic reduction in the number of evaluations needed to achieve convergence. Furthermore, ARBO can provide a bound on the number of iterations needed to achieve convergence by exploiting recent theory developed for the lower and upper confidence bounds of the GP model. Although powerful, ARBO does not directly handle

expensive-to-evaluate worst-case constraints, which are important in a wide-variety of safety-critical applications in engineering and beyond.

In this article, we propose a novel constrained extension of ARBO, referred to as constrained adversarially robust Bayesian optimization (CARBO), that is well-suited to constrained robust min-max optimization problems whose objective and/or constraint functions are defined in terms of noisy expensive black-box functions. Constraints are rigorously incorporated in CARBO by extending our recent work on the construction of convergent constrained BO algorithms using exact penalty functions in Ref. [53] to the min–max optimization setting. CARBO executes two steps in alternating fashion at each iteration: (i) select the design value that minimizes a lower confidence bound on a penalized version of the original robust problem (which we show can be interpreted as a probabilistic relaxation of the original problem that contains the true unknown feasible region with high probability) and (ii) select the uncertainty values that maximize the upper confidence bound of the unknown objective and constraint functions at the fixed design found in the previous step. In addition to providing detailed descriptions of the practical implementation aspects of CARBO, another key contribution of this article is a comprehensive theoretical analysis of CARBO's convergence properties. To prove convergence to a robust global optimum, we introduce the notion of *penalty-based robust-regret*, which measures the difference in the quality of recommended point by CARBO and the unknown robust global solution. By establishing that the cumulative penalty-based robust-regret is a sublinear function of the number of CARBO iterations under certain relatively mild assumptions, we can prove that there must exist a finite iteration such that the recommended point is arbitrarily close to the constrained robust global optimum. Since this bound decays to zero with an increasing number of CARBO iterations, we can immediately infer that CARBO is an asymptotically consistent and convergent algorithm. To the best of our knowledge, CARBO is the first method with guaranteed convergence properties that is applicable to constrained robust optimization problems involving noisy expensive black-box functions. In fact, the only other method we are aware of that is applicable to this class of problems was proposed in Ref. [54]. An important challenge with this method is that there are several internal parameters whose selection is not discussed, making it difficult to implement and compare against. Furthermore, Ref. [54] only selects uncertainty sample points based on the predicted worst-case constraint value, which may lead to a loss in worst-case performance, as demonstrated by our theoretical results.

The remainder of this article is organized as follows. In Section 2, we formulate the constrained robust optimization problem of interest and summarize all relevant assumptions on the objective and constraint functions. In Section 3, the GP modeling approach and the proposed CARBO algorithm are presented. The key theoretical convergence results for CARBO are presented in Section 4, while several practical implementation details are summarized in Section 5. The main advantages of CARBO are demonstrated on two case studies in Section 6, which includes a non-convex benchmark problem and a realistic engineering design problem defined in terms of an expensive bubble column reactor simulator for industrial waste gas recovery. Lastly, we conclude the article and discuss some interesting directions for future work in Section 7.

2 | PROBLEM FORMULATION

In this work, we are interested in solving the following constrained robust black-box optimization problem

$$\min_{\theta \in \Theta} \max_{w \in \mathcal{W}(\theta)} f(\theta, w), \tag{1a}$$

$$\text{s.t. } g_i(\theta, w) \le 0, \ \forall w \in \mathcal{W}(\theta), \ \forall i \in \{1, \ldots, m\}, \tag{1b}$$

where $\theta \in \mathrm{R}^{p_1}$ denotes the set of decision (or design) variables that must reside in a set $\Theta \subset \mathrm{R}^{p_1}$, $w \in \mathrm{R}^{p_2}$ denotes the set of uncertainties (or external disturbances) that are restricted to a set $\mathcal{W}(\theta) \subset \mathrm{R}^{p_2}$ that can depend on θ, and $f : \mathrm{R}^{p_1} \times \mathrm{R}^{p_2} \to \mathrm{R}$ and $g_i : \mathrm{R}^{p_1} \times \mathrm{R}^{p_2} \to \mathrm{R}$ are *unknown* black-box objective and constraint functions, respectively, for all $i = 1, \ldots, m$. The objective function (1a) can be interpreted as a sequential two-player game, with $\min_\theta \max_w$ reflecting the assumed order of the game, meaning the design θ must be selected first while the adversary w is able to adapt to any given choice of θ. The constraints (1b), on the other hand, may represent critical safety and/or performance requirements and must be satisfied for any feasible uncertainty realization. Such constraints are often referred to as "worst-case" constraints, as (1b) can be equivalently stated as $\max_{w \in \mathcal{W}(\theta)} g_i(\theta, w) \le 0$ for all $i = 1, \ldots, m$. It is important to note that, in general, different values of w produce the worst-case objective and constraint functions, which will be an important factor in the derivation of our algorithm in Sections 3 and 4. We are interested in finding globally optimal solutions to problem (1a and 1b), which falls under the class of generalized semi-infinite programming (GSIP) problems. A comprehensive review on GSIPs is provided in Ref. [16] while further details on state-of-the-art numerical methods for GSIPs can be found in, for example, Refs. [9,18,21]. The proper selection of an algorithm for solving (2) depends on the underlying characteristics of the functions $\mathcal{F} = \{f, g_1, \ldots, g_m\}$ and the constraint sets Θ and $\mathcal{W}(\theta)$. Currently, two main classes of numerical methods exist for GSIPs: discretization- and local reduction-based methods. A key implicit assumption in these methods is that the functions $\mathcal{F}$ are known (and often differentiable), which is not the case in this work, as summarized in the following assumption:

Assumption 1.

1. *The feasible sets Θ and $\mathcal{W}(\theta)$ are known and compact for all $\theta \in \Theta$.*
2. *The worst-case uncertainty values $w_v^\star(\theta) \in \mathrm{argmax}_{\mathcal{W}(\theta)} v(\theta, w)$ for all functions $v \in \mathcal{F}$ are unknown, so they cannot be determined/estimated from prior knowledge.*

3. *The objective $f(\theta, w)$ and constraint functions $\{g_i(\theta, w)\}_{i=1}^m$ are fully black-box in nature, meaning no closed-form expression or special structure (e.g., convexity) is known.*
4. *The set of functions $\mathcal{F}$ are expensive to evaluate and only provide zeroth order information, that is, we cannot obtain information about the derivatives of the function output with respect to the inputs.*
5. *The observations of $v(\theta, w)$ for all $v \in \mathcal{F}$ are corrupted by noise. That is, $y_v = v(\theta, w) + \varepsilon_v$ where $\varepsilon_v \sim \mathcal{N}(0, \sigma_{\varepsilon,v}^2)$.*
6. *The set of functions in $\mathcal{F}$ are independent such that they must be queried separately.*

Characteristics 1–3 imply that (2) can be used to represent a wide-variety of problems since we do not require any prior structural knowledge of the objective or constraint functions. Characteristic 4 is a key assumption in this work since it prevents the application of any currently available robust optimization methods that either exploit derivative information and/or require a large number of function evaluations (which implicitly assume cheap objective and constraint functions are readily available). In addition, characteristic 5 further increases the complexity of the problem since we do not assume access to highly accurate function evaluations, which may not be available in practice. Even though we assume the effective noise term is additive and normally distributed, our proposed algorithm can learn the parameters of this distribution as data is collected, which is discussed in more detail in Section 5. Lastly, characteristic 6 is a minor assumption that we make to study the limited information case and can be easily relaxed (see Remark 2). It is worth noting that, if Assumption 1 was replaced with standard assumptions from the GSIP literature, then we could directly apply established numerical methods recently summarized in Ref. [9]. However, to the best of our knowledge, no such method exists for functions satisfying Assumption 1, which is the main contribution of this work.

The most direct way to solve (1a and 1b) is to define an equivalent constrained black-box optimization problem

$$\min_{\theta \in \Theta} F(\theta), \text{ s.t. } G_i(\theta) \le 0, \ \forall i \in \{1, \ldots, m\}, \tag{2}$$

where $F(\theta) = \max_{\mathcal{W}(\theta)} f(\theta, w)$ and $G_i(\theta) = \max_{\mathcal{W}(\theta)} g_i(\theta, w)$. However, this requires one to solve a set of lower-level optimization problems at every iteration of the upper-level minimization algorithm. Such an approach is not applicable to expensive functions since one may spend the majority of the available evaluation budget on a single (potentially poor) choice of θ. Since we are considering a very challenging class of problems, as highlighted by Assumption 1, we make two additional assumptions to ensure that we can make progress on solving these problems:

Assumption 2. *The unknown functions $v \in \mathcal{F}$ are sufficiently smooth enough so that they can be modeled as GPs, which are formally defined in Section 3.*

Assumption 3. *At least one global solution $\theta^\star$ to (2) exists and the Mangasarian–Fromovitz Constraint Qualification (MFCQ) for the upper-level problem (3) holds at every global solution $\theta^\star$.*

Assumption 2 is a standard one considered in the Bayesian paradigm. Assumption 3, which is based on results established in Ref. [55] is used to provide some theoretical convergence properties of the proposed algorithm and is fairly weak since the MFCQ conditions only need to hold at global minimizers. Note that similar constraint qualification assumptions are routinely made in the GSIP literature to establish convergence properties, as discussed in, for example, Ref. [56]. Instead of approaching (1a and 1b) by applying established black-box optimization methods to the nested formulation (2), we look to reformulate (1a and 1b) as a so-called *bandit feedback* problem. In particular, we would like to develop a sequential learning (or bandit) algorithm that selects a batch of samples $\{(\theta_{v,t}, w_{v,t}) \in \Theta \times \mathcal{W}(\theta_{v,t})\}_{v \in \mathcal{F}}$ at which the objective and constraint functions should be queried at each iteration $t = 1, 2, \ldots$ (where "iteration" refers to the total evaluation period of the objective and constraints) and subsequently recommends a best point θ_t^r that ideally converges to a global solution $\theta_t^r \to \theta^\star$ with high probability. The authors of this work recently proposed an approach based on Ref. [51], termed adversarially robust Bayesian optimization (ARBO),[52] that provides such a convergence guarantee in the absence of constraints. The theoretical analysis of ARBO is based on the *instantaneous robust-regret* at iteration t, which is defined as follows

$$r_t^w = \max_{w \in \mathcal{W}(\theta_t)} f(\theta_t, w) - \max_{w \in \mathcal{W}(\theta^\star)} f(\theta^\star, w), \tag{3}$$

where $\theta^\star$ denotes any global solution to (1a and 1b). An algorithm that minimizes the *cumulative robust-regret* $R_T^w = \sum_{t=1}^T r_t^w$ defined over T iterations will ensure that we learn $\theta^\star$ as quickly as possible. It is not possible to minimize R_T^w directly since it is defined in terms of the unknown optimal solution $\theta^\star$. Instead, ARBO ensures that the *no robust-regret* property is satisfied, that is, $\lim_{T \to \infty} R_T^w / T = 0$. In the absence of constraints, the no robust-regret property ensures convergence, which we can see by analyzing the *simple robust-regret*

$$S_T^w = \min_{t \in \{1, \ldots, T\}} r_t^w = \min_{t \in \{1, \ldots, T\}} \left(\max_{w \in \mathcal{W}(\theta_t)} f(\theta_t, w) - \max_{w \in \mathcal{W}(\theta^\star)} f(\theta^\star, w) \right). \tag{4}$$

Since $r_t^w \ge 0$ must be non-negative if $\theta_t \in \Theta$ (trivial to guarantee since Θ is known) and no black-box constraints are present ($m = 0$), and the minimum of a sequence must be less than the average, that is, $0 \le S_T^w \le \frac{1}{T} R_T^w$, the no robust-regret property directly implies that $S_T^w \to 0$ can only happen when one of the sampled points belongs to the set of global solutions. This analysis breaks down in the presence of black-box constraints ($m > 0$) since the lower bound on S_T^w may not hold. In particular, we may have $r_t^w < 0$ for certain choices of $\theta_t \in \Theta$ that result in violation of the black-box constraints. Therefore, the goal of this article is to develop an extension of ARBO that is

applicable to constrained problems of the form (1a and 1b). We give a detailed description of the proposed algorithm in the next section. We then propose a new definition of robust-regret for constrained problems and demonstrate how this can be used to establish robust-regret bounds that lead to rigorous convergence guarantees in Section 4.

3 | CONSTRAINED ADVERSARIALLY ROBUST BAYESIAN OPTIMIZATION

In this section, we first summarize the GP regression methodology, which is used to construct probabilistic surrogate models for the unknown objective and constraint functions. We then present the proposed CARBO algorithm, followed by a discussion on how to reduce the number of samples needed at each CARBO iteration using constraint aggregation methods.

3.1 | Gaussian process regression

Let $\mathbf{x} = \left[\theta^T, w^T\right]^T$ denote the concatenated vector of design variables and uncertainties that are restricted to the space $\mathcal{X} = \Theta \times \mathcal{W}(\Theta) \subset \mathrm{R}^p$ where $p = p_1 + p_2$. Since we treat all of the black-box objective and constraint functions in a similar fashion, we focus the presentation in this section on a scalar function $v : \mathrm{R}^p \to \mathrm{R}$ for any $v \in \mathcal{F}$. For simplicity of notation, we also interchangeably denote $v(\theta, w)$ as $v(\mathbf{x})$ throughout the article for all $v \in \mathcal{F}$. In accordance with Assumption 2, we cannot make any parametric assumption about v due to lack of knowledge about its underlying structure and must instead model v as a sample of a GP, which is a commonly used class of models in non-parametric Bayesian inference. As discussed in detail in Ref. [44] GPs can be interpreted as an infinite collection of random variables for which any finite subset has a joint Gaussian distribution, meaning they generalize the notion of multivariate distributions over vector spaces to probability distributions over the space of functions. It should be noted that, in the absence of Assumption 2, any of the functions $v \in \mathcal{F}$ could be arbitrarily discontinuous at every input $\mathbf{x} \in \mathcal{X}$ in the worst-case. Therefore, by leveraging GP models, we are in essence encoding that some level of smoothness implicitly holds in these functions without making rigid parametric assumptions.

We let $v(\mathbf{x}) \sim \mathcal{GP}(\mu_v(\mathbf{x}), k_v(\mathbf{x}, \mathbf{x}'))$ denote a GP that is uniquely specified by its mean function $\mu_v(\mathbf{x})$ and covariance (or kernel) function $k_v(\mathbf{x}, \mathbf{x}')$, which are defined as follows

$$\mu_v(\mathbf{x}) = E_v\{v(\mathbf{x})\}, \tag{5a}$$

$$k_v(\mathbf{x}, \mathbf{x}') = E_v\{(v(\mathbf{x}) - \mu_v(\mathbf{x}))(v(\mathbf{x}') - \mu_v(\mathbf{x}'))\}. \tag{5b}$$

The properties of the fitted functions v are derived from the choice of k_v. The proposed CARBO algorithm is applicable to GPs with *any* covariance function including stationary and non-stationary kernels. The theoretical results established in Section 4, however, focus on covariance functions from the Matern class,[44] defined as follows

$$k_v(\mathbf{x}, \mathbf{x}'; \nu, \Psi) = \zeta^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} r(\mathbf{x}, \mathbf{x}')\right) B_\nu \left(\sqrt{2\nu} r(\mathbf{x}, \mathbf{x}')\right), \tag{6}$$

where $r(\mathbf{x}, \mathbf{x}') = \sqrt{(\mathbf{x} - \mathbf{x}') L^{-2} (\mathbf{x} - \mathbf{x}')}$ is the scaled Euclidean distance, $L = \mathrm{diag}(l_1, \ldots, l_p)$ is a diagonal scaling matrix composed of length-scale parameters $l_1, \ldots, l_p > 0$, ν is a parameter that dictates smoothness (i.e., the corresponding function is $\lceil \nu/2 - 1 \rceil$ times differentiable), ζ^2 is a scaling factor for the output variance, Γ and B_ν are the Gamma and modified Bessel functions, respectively, and $\Psi = \{l_1, \ldots, l_p, \zeta\}$ are the hyperparameters of the kernel for a fixed value of ν. Training a GP model corresponds to calibrating $\{\Psi, \sigma_{\varepsilon,v}\}$ to the available data. For now, we assume the kernel hyperparameters are known and discuss the training procedure further in Section 5. It should be noted that learning the length-scale parameters $l_1, \ldots, l_p$ is directly related to feature selection since the ith dimension $[\mathbf{x}]_i$ is unimportant (i.e., does not contribute to the output predictions) whenever $l_i \to \infty$.

In addition to being non-parametric models, GPs have simple analytic expressions for the mean and covariance of the posterior distribution when conditioned on (possibly noisy) function observations. Let us assume that we have access to t noisy observations at known input points, which are represented by the following matrices

$$\mathbf{X}_{v,t} = [\mathbf{x}_{v,1}, \mathbf{x}_{v,2}, \ldots, \mathbf{x}_{v,t}]^T \in \mathrm{R}^{t \times p}, \tag{7a}$$

$$\mathbf{y}_{v,t} = \left[y_{v,1}, y_{v,2}, \ldots, y_{v,t}\right]^T \in \mathrm{R}^{t \times 1}, \tag{7b}$$

where $y_{v,i} = v(\mathbf{x}_{v,i}) + \varepsilon_{v,i}$ and $\varepsilon_{v,i} \sim \mathcal{N}(0, \sigma^2_{\varepsilon,v})$ is i.i.d. Gaussian noise for all $i = 1, \ldots, t$. As shown in section 2 of Ref. [44], the posterior distribution $v(\mathbf{x}) \mid \mathbf{X}_t, \mathbf{y}_{v,t} \sim \mathcal{GP}\left(\mu_{v,t}(\mathbf{x}), k_{v,t}(\mathbf{x}, \mathbf{x}')\right)$ of the function $v(\mathbf{x})$ remains a GP with the following expressions for the posterior mean $\mu_{v,t}(\mathbf{x})$, covariance $k_{v,t}(\mathbf{x}, \mathbf{x}')$, and variance $\sigma^2_{v,t}(\mathbf{x})$

$$\mu_{v,t}(\mathbf{x}) = \mu_v(\mathbf{x}) + \mathbf{k}^T_{v,t}(\mathbf{x}) \left(\mathbf{K}_{v,t} + \sigma^2_{\varepsilon,v} \mathbf{I}_t\right)^{-1} \left(\mathbf{y}_{v,t} - \mu_v(\mathbf{x})\right), \tag{8a}$$

$$k_{v,t}(\mathbf{x}, \mathbf{x}') = k_v\left(\mathbf{x}, \mathbf{x}'\right) - \mathbf{k}^T_{v,t}(\mathbf{x}) \left(\mathbf{K}_{v,t} + \sigma^2_{\varepsilon,v} \mathbf{I}_t\right)^{-1} \mathbf{k}_{v,t}(\mathbf{x}'), \tag{8b}$$

$$\sigma^2_{v,t}(\mathbf{x}) = k_{v,t}(\mathbf{x}, \mathbf{x}), \tag{8c}$$

where $\mathbf{k}_{v,t}(\mathbf{x}) = k_v(\mathbf{X}_t, \mathbf{x}) \in \mathrm{R}^{t \times 1}$ is the vector of covariance values between the test point $\mathbf{x}$ and the observed data points $\mathbf{X}_t$, $\mathbf{K}_{v,t} = k_v(\mathbf{X}_t, \mathbf{X}_t) \in \mathrm{R}^{t \times t}$ is the covariance matrix between the observed data points, and $\mathbf{I}_t$ is the $t \times t$ identity matrix. A key advantage of the posterior GP expressions shown in (8a)–(8c) is that they quantify the degree of uncertainty in the predictions such that they can straightforwardly be used to generate *confidence bounds* for $v(\theta, w)$ at any desired test input $\mathbf{x} = (\theta, w) \in \Theta \times \mathcal{W}(\theta)$. In particular, given an exploration parameter β_{t+1}, the upper and lower confidence bounds on v are defined as follows

$$\mathrm{UCB}_{v,t}(\theta,w)=\mu_{v,t}(\theta,w)+\beta_{t+1}^{1/2}\sigma_{v,t}(\theta,w), \tag{9a}$$

$$\mathrm{LCB}_{v,t}(\theta,w)=\mu_{v,t}(\theta,w)-\beta_{t+1}^{1/2}\sigma_{v,t}(\theta,w), \tag{9b}$$

which can be readily computed from the GP posterior expressions in (8a)–(8c). It turns out that both the upper and lower confidence bounds provide important information that can be exploited by CARBO, as discussed next.

3.2 | Overview of the CARBO algorithm

Given the posterior GP surrogate models for all unknown functions $v\in\mathcal{F}$ from (8a)–(8c), we can come up with several different strategies for sequentially selecting sample points $\{\mathbf{x}_{v,t}\}_{v\in\mathcal{F}}$ given previous data $\{\mathbf{X}_{v,t-1},\mathbf{y}_{v,t-1}\}_{v\in\mathcal{F}}$. For example, we could attempt to spend our evaluation budget on learning globally accurate representations of the functions $v\in\mathcal{F}$ as quickly as possible. We can interpret this goal as a Bayesian experimental design (ED) problem,[57] which suggests that we select samples that maximize the predicted variance of each unknown function, that is,

$$\mathbf{x}_{v,t}\in\underset{\mathbf{x}\in\mathcal{X}}{\operatorname{argmax}}\quad \sigma_{v,t-1}^2(\mathbf{x}),\ \ \forall v\in\mathcal{F}. \tag{10}$$

This approach is good at *exploring* the functions $v\in\mathcal{F}$; however, even in the unconstrained case, it is not particularly well-suited to the optimization task since we only care about identifying values of $\mathbf{x}=\{\theta,w\}$ that minimize the worst-case value of $f(\theta,w)$. This issue is only exacerbated in the constrained case considered here since we have more functions that must be explored, which implies an even larger number of expensive function evaluations are needed to globally model every $v\in\mathcal{F}$. A natural alternative to the explorative Bayesian ED approach is to select a single $\mathbf{x}_t$ by solving the following min-max optimization problem

$$\min_{\theta\in\Theta}\max_{w\in\mathcal{W}(\theta)}\ \mu_{f,t-1}(\theta,w)\ \text{ s.t. }\ \mu_{g_i,t-1}(\theta,w)\le 0,\ \forall w\in\mathcal{W}(\theta),\ \forall i\in\{1,\ldots,m\}, \tag{11}$$

ALGORITHM 1 CARBO: A sequential learning algorithm for constrained robust black-box optimization problems

Input: The set of the design variables Θ and uncertainties $\mathcal{W}(\theta)$ for all $\theta\in\Theta$; GP priors $\{\mu_v,k_v\}_{v\in\mathcal{F}}$; exploration parameters $\{\beta_t\}_{t\ge 1}$; constraint penalty weight factor ρ; and total number of iterations T.

1: **for** $t=1$ to T **do**

2: Solve the following min–max optimization problem for θ_t

$$\theta_t\in\underset{\theta\in\Theta}{\operatorname{argmin}}\left(\max_{w\in\mathcal{W}(\theta)}\mathrm{LCB}_{f,t-1}(\theta,w)+\rho\sum_{i=1}^{m}\left(\max_{w\in\mathcal{W}(\theta)}\mathrm{LCB}_{g_i,t-1}(\theta,w)\right)^{+}\right), \tag{12}$$

where $h^{+}(x)=(h(x))^{+}=\max\{h(x),0\}$ denotes the positive part of a function $h(x)$.

3: Set $\theta_{v,t}=\theta_t$ for all $v\in\mathcal{F}$.

4: **for** every unknown function v in the set $\mathcal{F}$ **do**

5: Solve the following maximization problem for $w_{v,t}$

$$w_{v,t}\in\underset{w\in\mathcal{W}(\theta_t)}{\operatorname{argmax}}\ \mathrm{UCB}_{v,t-1}(\theta_{v,t},w). \tag{13}$$

6: Evaluate the unknown function v at $\{\theta_{v,t},w_{v,t}\}$, i.e., collect data $y_{v,t}=v(\theta_{v,t},w_{v,t})+\varepsilon_{v,t}$.

7: Update the GP posterior mean $\mu_{v,t}$, variance $\sigma_{v,t}^2$, and confidence bounds $\mathrm{LCB}_{v,t}$ and $\mathrm{UCB}_{v,t}$ according to (8a)–(8c) and (9a and 9b) with the newest available data.

8: **end for**

9: **end for**

10: Return the recommended design point $\theta_T^r=\theta_{t^\star}$ that corresponds to the sampled point with the smallest penalized upper confidence bound, defined as follows

$$t^\star\in\underset{t\in\{1,\ldots,T\}}{\operatorname{argmin}}\left(\max_{w\in\mathcal{W}(\theta_t)}\mathrm{UCB}_{f,t-1}(\theta_t,w)+\rho\sum_{i=1}^{m}\left(\max_{w\in\mathcal{W}(\theta_t)}\mathrm{UCB}_{g_i,t-1}(\theta_t,w)\right)^{+}\right). \tag{14}$$

which is derived by substituting the mean representation of the posterior objective and constraint functions into (1a and 1b). This type of mean-based search rule is known to be greedy in the sense that it can easily get stuck in shallow local solutions. Another major challenge with such an approach is the inaccuracies in the mean constraint functions can cutoff exploration in a large portion of the feasible region, especially in the early iterations when limited data is available. In other words, mean-based search strategies tends to *over-exploit* the currently available data. This over-exploitation is not a big issue when data is abundant and/or cheap to collect, but is a major limitation in the context of expensive function evaluations considered in Assumption 1. Furthermore, in the robust optimization case, (11) does not incentivize improved learning of the feasible region because the uncertainty values are selected based on the predicted worst-case objective function.

Motivated by the GP-LCB algorithm[58] that has been shown to work well in the nominal unconstrained case, we look to develop a strategy that takes advantage of the *confidence bounds* defined in (9a and 9b) to address the exploration-exploitation tradeoff. For sufficiently large choices of β_{t+1}, these confidence bounds can be used to guarantee that the true unknown function is contained within these bounds with high probability, which is discussed in Section 4. The proposed CARBO algorithm, which relies on the confidence bounds (9a and 9b), is summarized below in Algorithm 1.

The suggested θ_t is the one that leads to the minimum *robust penalized* lower confidence bound, as shown in (12). Note that we specifically penalize the constraints using a non-smooth penalty function, which will be needed to establish the theoretical convergence results in the next section. To better interpret (12), let us recast it as the following equivalent constrained min–max problem

$$\min_{\theta\in\Theta,\,\varepsilon_i\geq 0}\ \max_{w\in\mathcal{W}(\theta)}\ \mu_{f,t-1}(\theta,w)-\beta_t^{1/2}\sigma_{f,t-1}(\theta,w)+\rho\sum_{i=1}^{m}\varepsilon_i, \tag{15a}$$

$$\text{s.t.}\ \max_{w\in\mathcal{W}(\theta)}\left(\mu_{g_i,t-1}(\theta,w)-\beta_t^{1/2}\sigma_{g_i,t-1}(\theta,w)\right)\leq\varepsilon_i,\qquad \forall i\in\{1,\ldots,m\}, \tag{15b}$$

where ε_i denote slack variables that represent the degree of robust constraint violation for the ith constraint. We delay a discussion on how to practically solve (15a) until Section 5 and, for now, assume that we can identify a global solution. From (15a), we see that Algorithm 1 simplifies to the standard GP-LCB algorithm whenever there is no uncertainty (i.e., $\mathcal{W}(\theta)=\{\overline{w}\}$ is a singleton containing only a nominal value) and there are no constraints ($m=0$). However, the inclusion of uncertainty and constraints do fundamentally change the behavior of the method. We see that the constraints (15b) can be interpreted as a "relaxation" of $\{\theta : G_i(\theta)\leq 0, \forall i=1,\ldots,m\}$ that should contain the true feasible domain with high probability (as long as $\{\beta_t\}_{t\geq 1}$ are appropriately chosen). To handle uncertainty, once we have selected θ_t, we also need to select feasible a set of uncertainty values $\{w_{v,t}\}_{v\in\mathcal{F}}$. This is accomplished by solving (13), which corresponds to selecting the uncertainty value that maximizes the *upper* confidence bound for each unknown function. CARBO is thus based on two distinct principles that alternate within each step of the algorithm: (i) make optimistic selections under uncertainty for the design θ_t that is shared for all $v\in\mathcal{F}$ and (ii) make pessimistic selections under uncertainty for the anticipated worst-case values of $\{w_{v,t}\}_{v\in\mathcal{F}}$ that can differ between the unknown objective and constraint functions. The first concept is similar to traditional BO algorithms, while the second is unique to constrained robust BO algorithms that must mitigate the effect of a potential adversary simultaneously on multiple functions. We also see that CARBO reduces to the ARBO algorithm from Ref. [52] whenever $m=0$, which only requires one evaluation at each iteration because only the objective function is considered to be unknown.

Once the outer for loop is completed in Algorithm 1, a final recommended point is selected out the sequence $\{\theta_1,\ldots,\theta_T\}$ using the rule shown in (14). It should be noted that many different recommendation procedures are possible (e.g., simply returning the final point $\theta_T^r=\theta_T$); however, we select one that minimizes a pessimistic bound on a penalized variation of the robust-regret from (3), which is based on the theoretical results established in the next section. In addition, the key parameters of the algorithm, mainly $\{\beta_t\}_{t\geq 1}$ and ρ, have been left unspecified. We also analyze their impact on the convergence properties of CARBO next.

Remark 1. In CARBO, $m+1$ function evaluations are needed at each iteration, which scales with the number of constraints considered in the original problem (1a and 1b). Thus, it can be valuable to aggregate constraints to limit these evaluations. The simplest constraint aggregation approach is the maximum constraint method, which replaces (1b) with a single constraint

$$\max\{g_1(\theta,w),\ldots,g_m(\theta,w)\}\leq 0,\ \ \forall w\in\mathcal{W}(\theta). \tag{16}$$

An important challenge with this approach is that the resulting constraint may not satisfy the smoothness requirement of Assumption 2. A better alternative is to rely on the Kreisselmeier–Steinhauser (KS) function to aggregate the constraints into the following smooth composite function

$$KS(g_1(\theta,w),\ldots,g_m(\theta,w))=\frac{1}{\lambda}\ln\left[\sum_{i=1}^{m}e^{\lambda g_i(\theta,w)}\right]\leq 0,\ \ \forall w\in\mathcal{W}(\theta), \tag{17}$$

where $\lambda>0$ is a weight factor that ensures the KS function converges to the max operator as $\lambda\to\infty$. Not only does the KS function provide the desired smoothness properties, but also it can be shown to be an overestimate of the original constraints $g_i(\theta,w)\leq 0$,[59] so that it can still be used to establish robust feasibility guarantees in our context. Therefore, by replacing (1b) with (17), CARBO only requires two samples at each iteration, which is completely independent of m. It is worth noting that the computational cost of each CARBO iteration is also reduced when applying KS aggregation since only two GP models are included in the sub-problems (12) and (13). The main downside to this KS simplification is that

less information is provided to these sub-problem such that constraint predictions may not be as accurately represented at each iteration.

Remark 2. In the description of CARBO provided in Algorithm 1, we have focused on the case that the functions $v \in \mathcal{F}$ must be independently evaluated, as noted in characteristic 6 of Assumption 1. There are certain cases where the evaluation of the objective and constraints is coupled such that, when we query one of the functions, we can cheaply obtain data for some or all of the other functions of interest. One such example is when the functions $\{f, g_1, ..., g_m\}$ are defined in terms of a vector of outputs from a complex simulator that must be evaluated simultaneously. Since the case considered here is a special case of this more general setting, all of the results proved in this article still hold. Furthermore, we can easily incorporate this information into the GP models by augmenting the datasets (7a and 7b) with any additional measurements, which can be expected to practically improve the rate of convergence.

4 | THEORETICAL CONVERGENCE GUARANTEES FOR CARBO

The focus of this section is on theoretical analysis of the proposed CARBO algorithm discussed in Section 3. First, we provide an overview of exact penalty functions for constrained nonlinear programs of the form (2), which is needed to formulate our modified definition of regret. Then, we use established results on GP confidence bounds to determine an upper bound on our proposed regret quantity. Finally, we show how this newly derived upper regret-like bound is sufficient to establish convergence of CARBO to the constrained global minimax solution.

4.1 | Proposed definition of exact penalty-based robust-regret

As discussed previously, the original robust problem (1a and 1b) can be represented by (2) where $F(\theta)$ and $G_i(\theta)$ are defined in terms of the lower-level optimization problems. Let us consider a related optimization problem of the form

$$\min_{\theta \in \Theta} P(\theta;\rho) = F(\theta) + \rho \sum_{i=1}^{m} G_i^+(\theta), \tag{18}$$

where $\rho \geq 0$ is a non-negative constant and $P(\theta;\rho)$ denotes the non-differentiable penalty function. The key "exactness" property for $P(\cdot)$ that is relevant to this work is summarized in the following theorem.

Theorem 1. (Exact Penalty Function[55]) Let Assumption 3 hold. Then, there exists a finite threshold value $\bar{\rho}$ such that, for any $\rho \in [\bar{\rho}, \infty)$, every global solution to (2) is a global solution of (18) and vice versa.

It is important to note that exact penalty property summarized in Theorem 1 does not hold for all choices of penalty functions. As an example, it is common to formulate the Lagrangian function (by linearly penalizing constraints) and then apply duality theory to transform the original problem into a dual problem that yields a lower bound on the optimal objective function of the original (or primal) problem according to the weak duality theorem. Unless certain strong convexity conditions hold, however, there will exist a *duality gap* between the primal and dual solutions, meaning we cannot directly use the dual problem to solve the original primal problem. As discussed in Ref. [60] we can interpret the solution to (18) as an extended dual function that ensures zero duality gap (i.e., strong duality holds), which is a critical property for potentially highly non-convex problems because it provides us with a path forward to solving the original problem (2) (equivalent to (1a and 1b)) by solving simpler problems of the form (18).

Let $\mathbf{G}(\theta) = [G_1(\theta), ..., G_m(\theta)]^T$ denote a concatenated vector of the upper-level constraints. Using the exact penalty function concept in (18), we define the *instantaneous penalty-based robust-regret* at any iteration t as follows

$$\begin{aligned} r^w_{EP,t}(\rho) &= F(\theta_t) + \rho \,\|\, \mathbf{G}^+(\theta_t)\|_1 - F(\theta^\star), \\ &= \max_{w \in \mathcal{W}(\theta_t)} f(\theta_t, w) + \rho \sum_{i=1}^{m} \left(\max_{w \in \mathcal{W}(\theta_t)} g_i(\theta_t, w) \right)^+ - \max_{w \in \mathcal{W}(\theta^\star)} f(\theta^\star, w), \\ &= r^w_t + \rho \sum_{i=1}^{m} \left(\max_{w \in \mathcal{W}(\theta_t)} g_i(\theta_t, w) \right)^+. \end{aligned} \tag{19}$$

Note that the term $\rho \,\|\, \mathbf{G}^+(\theta^\star)\|_1 = 0$ by Assumption 3, that is, at least one global solution exists to the original constrained problem (1a and 1b). The second and third lines in (19) show that there is a clear relationship between the original robust-regret r^w_t and the penalized version $r^w_{EP,t}(\rho)$, with the additional term having important implications on convergence, as shown in detail in the subsequent sections.

Remark 3. The exact penalty function proposed in (18) can be categorized as an L_1 penalty function and is one of most recommended ones for (G)SIPs.[56] Other exact penalty functions have been proposed for this class of problems including modified L_1 penalty functions[61] and L_∞ penalty functions.[62] In principle, these other penalty functions can be used within the proposed CARBO algorithm, though it is not clear what their effect would be on the underlying numerics or rate of convergence. We believe better understanding the impact on the choice of penalty function used by CARBO would be an interesting direction for future work.

4.2 | Upper bound on cumulative penalty-based robust-regret

CARBO, presented in Algorithm 1, requires one to select the exploration parameters $\{\beta_t\}_{t \geq 1}$ that directly specify the width of the

confidence intervals as seen in (9a and 9b). To see the importance of these values, note that (15a and 15b) reduces to the mean-based search method (11) when $\beta_t = 0$ (and $\varepsilon_i = 0$), which results in variance information being completely ignored. We first extend a result from Ref. [58] for one GP to multiple GPs. For simplicity, we summarize this result for the case of a finite set $\mathcal{X} = \Theta \times \mathcal{W}(\Theta)$ first and discuss the extension to compact sets in Section 5.

Lemma 1. *(GP Confidence Bounds) Let $v(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ be a sample of a GP for which noisy observations of the form (7a and 7b) are available for all $v \in \mathcal{F}$. In addition, let $\beta_t = 2\log\left((m+1)|\mathcal{X}|t^2\pi^2/(6\alpha)\right)$ for a specified failure probability $\alpha \in (0,1)$ and discrete set $|\mathcal{X}| < \infty$. Then, the following simultaneous bounds on the objective and constraint functions*

$$v(\mathbf{x}) \in [\mathrm{LCB}_{t-1}(\mathbf{x}), \mathrm{UCB}_{t-1}(\mathbf{x})], \quad \forall \mathbf{x} \in \mathcal{X}, \forall t \geq 1, \forall v \in \mathcal{F}, \tag{20}$$

hold with probability $1-\alpha$ (over the set of GP posteriors at every iteration t).

Proof. The result [58, lemma 5.1] proves (20) holds for a single function v, that is, $|\mathcal{F}| = 1$. To extend this result, we replace α with a smaller value $\alpha/|\mathcal{F}|$. We then apply the union bound (aka Boole's inequality) as follows

$$\mathrm{P}\left\{\bigcup_{v \in \mathcal{F}} v(\mathbf{x}) \notin [\mathrm{LCB}_{t-1}(\mathbf{x})\mathrm{UCB}_{t-1}(\mathbf{x})], \forall \mathbf{x} \in \mathcal{X} \forall t \geq 1\right\} \leq \sum_{v \in \mathcal{F}} \mathrm{P}\{v(\mathbf{x}) \notin [\mathrm{LCB}_{t-1}(\mathbf{x})\mathrm{UCB}_{t-1}(\mathbf{x})], \forall \mathbf{x} \in \mathcal{X} \forall t \geq 1\}.$$

The right-hand side above can be reduced to $\sum_{v \in \mathcal{F}} \frac{\alpha}{|\mathcal{F}|} = \alpha$ such that the stated claim holds since $|\mathcal{F}| = m+1$.

We also need the following bound that can be derived from the properties of the positive part operator.

Lemma 2. *(Positive Part Difference) The inequality $a^+ - b^+ \leq (a-b)^+$ holds for any real constants $a, b \in \mathrm{R}$.*

Proof. There are four cases to consider: (i) $a, b < 0$, (ii) $a < 0$ and $b \geq 0$, (iii) $a \geq 0$, $b < 0$, and (iv) $a, b \geq 0$. For case (i), the left-hand side is 0 while the right-hand side must be ≥ 0 by definition. For case (ii), the left-hand side is $-b \leq 0$ while the right-hand side is 0. For case (iii), the left-hand side is a while the right-hand side is $a + |b| > a$. For case (iv), the left-hand side is $a - b$ while the right-hand side must always be $\geq a - b$.

Let us now convert the GP confidence bounds in Lemma 1 to a bound on our proposed regret measure $r^w_{EP,t}$ in (19) that depends only on constants and the GP posterior variance.

Lemma 3. *(Instantaneous Penalty-based Robust-Regret Upper Bound) Let $t \geq 1$ be fixed, the query points $\{\theta_{v,t}, w_{v,t}\}_{v \in \mathcal{F}}$ be selected according to CARBO (Algorithm 1), and Assumptions 2 and 3 hold. Then, if $\mathrm{LCB}_{v,t-1}(\mathbf{x}) \leq v(\mathbf{x}) \leq \mathrm{UCB}_{v,t-1}(\mathbf{x})$ holds for all $\mathbf{x} \in \mathcal{X}$ and $v \in \mathcal{F}$, the instantaneous penalty-based robust-regret in (19) satisfies the following bound*

$$r^w_{EP,t}(\rho) \leq 2\beta_t^{1/2} \sum_{v \in \mathcal{F}} \rho_v \sigma_{v,t-1}(\theta_{v,t}, w_{v,t}), \tag{21}$$

where $\rho_v = 1$ for $v = f$ and $\rho_v = \rho$ for all $v \in \mathcal{F} \setminus f$.

Proof. First, we need to establish a pessimistic estimate of $r^w_{EP,t}(\rho)$, which we define as follows

$$\bar{r}^w_{EP,t}(\rho) = \max_{w \in \mathcal{W}(\theta_t)} \mathrm{UCB}_{f,t-1}(\theta_t, w) + \rho \sum_{i=1}^{m} \left(\max_{w \in \mathcal{W}(\theta_t)} \mathrm{UCB}_{g_i,t-1}(\theta_t, w)\right)^+ - \max_{w \in \mathcal{W}(\theta^\star)} f(\theta^\star, w).$$

We clearly see that $r^w_{EP,t}(\rho) \leq \bar{r}^w_{EP,t}(\rho)$ given the upper bound assumption on all $v \in \mathcal{F}$. Now let us use our previous definitions and results to establish the following sequence of inequalities
where first line follows from the definition of $w_{v,t}$ in (13); the second line follows from $\theta^\star \in \Theta$ being feasible and satisfying worst-case constraints by Assumption 3; the third line follows from the lower bound

$$\begin{aligned}
\bar{r}^w_{EP,t}(\rho) &= \mathrm{UCB}_{f,t-1}(\theta_t, w_{f,t}) + \rho \sum_{i=1}^{m} \left(\mathrm{UCB}_{g_i,t-1}(\theta_t, w_{g_i,t})\right)^+ - \max_{w \in \mathcal{W}(\theta^\star)} f(\theta^\star, w), \\
&= \mathrm{UCB}_{f,t-1}(\theta_t, w_{f,t}) + \rho \sum_{i=1}^{m} \left(\mathrm{UCB}_{g_i,t-1}(\theta_t, w_{g_i,t})\right)^+ - \min_{\theta \in \Theta} \left[\max_{w \in \mathcal{W}(\theta)} f(\theta, w) + \rho \sum_{i=1}^{m} \left(\max_{w \in \mathcal{W}(\theta)} g_i(\theta, w)\right)^+\right], \\
&\leq \mathrm{UCB}_{f,t-1}(\theta_t, w_{f,t}) + \rho \sum_{i=1}^{m} \left(\mathrm{UCB}_{g_i,t-1}(\theta_t, w_{g_i,t})\right)^+ - \min_{\theta \in \Theta} \left[\max_{w \in \mathcal{W}(\theta)} \mathrm{LCB}_{f,t-1}(\theta, w) + \rho \sum_{i=1}^{m} \left(\max_{w \in \mathcal{W}(\theta)} \mathrm{LCB}_{g_i,t-1}(\theta, w)\right)^+\right], \\
&= \mathrm{UCB}_{f,t-1}(\theta_t, w_{f,t}) + \rho \sum_{i=1}^{m} \left(\mathrm{UCB}_{g_i,t-1}(\theta_t, w_{g_i,t})\right)^+ - \max_{w \in \mathcal{W}(\theta_t)} \mathrm{LCB}_{f,t-1}(\theta_t, w) - \rho \sum_{i=1}^{m} \left(\max_{w \in \mathcal{W}(\theta_t)} \mathrm{LCB}_{g_i,t-1}(\theta_t, w)\right)^+, \\
&\leq \mathrm{UCB}_{f,t-1}(\theta_t, w_{f,t}) - \mathrm{LCB}_{f,t-1}(\theta_t, w_{f,t}) + \rho \sum_{i=1}^{m} \left[\left(\mathrm{UCB}_{g_i,t-1}(\theta_t, w_{g_i,t})\right)^+ - \left(\mathrm{LCB}_{g_i,t-1}(\theta_t, w_{g_i,t})\right)^+\right],
\end{aligned}$$

on all $v \in \mathcal{F}$; the fourth line follows from the definition of θ_t in (12); and the fifth line follows from the fact that $\max_{w \in \mathcal{W}(\theta_t)} \mathrm{LCB}_{v,t-1}(\theta_t, w) \geq \mathrm{LCB}_{v,t-1}(\theta_t, w_{v,t})$ for any feasible choice of $w_{v,t} \in \mathcal{W}(\theta_t)$.

Now that we have bounded a pessimistic estimate of $r^w_{EP,t}(\rho)$ in terms of the upper and lower confidence bounds on the unknown functions $v \in \mathcal{F}$, we can use the definitions in (9a and 9b) to further simplify this expression:

$$\begin{aligned} \bar{r}^w_{EP,t}(\rho) &\leq \mathrm{UCB}_{f,t-1}(\theta_t, w_{f,t}) - \mathrm{LCB}_{f,t-1}(\theta_t, w_{f,t}) + \rho \sum_{i=1}^{m} \big(\mathrm{UCB}_{g_i,t-1}(\theta_t, w_{g_i,t}) \\ &\quad - \mathrm{LCB}_{g_i,t-1}(\theta_t, w_{g_i,t})\big)^+, \\ &= 2\beta_t^{1/2}\sigma_{f,t-1}(\theta_t, w_{f,t}) + \rho \sum_{i=1}^{m} 2\beta_t^{1/2}\sigma_{g_i,t-1}(\theta_t, w_{g_i,t}), \end{aligned}$$

where the first line follows from Lemma 2 and the second line follows from (9a and 9b). The stated result follows by noticing $\theta_{v,t} = \theta_t$ for all $v \in \mathcal{F}$ (Line 3 of Algorithm 1) and simple rearrangement of the last inequality.

Our final step is to convert the bound on instantaneous penalty-based robust-regret into a cumulative one, so that we are able to analyze the behavior of CARBO for more than just a single iteration. Before presenting this main result, we define the *maximum information gain* (MIG), which is a fundamental quantity in Bayesian ED that provides a measure of informativeness of any finite set of sampling points $\mathcal{A} \subset \mathcal{X}$. We will state our cumulative penalty-based robust-regret bound directly in terms of MIG.

Definition 1. (Maximum Information Gain) Let $\mathcal{A} \subset \mathcal{X}$ denote any subset of sampling points from $\mathcal{X}$. The maximum information gain (MIG) for any unknown function $v \in \mathcal{F}$ under the t noisy measurements from (7a and 7b) is defined as

$$\gamma_{v,t} = \max_{\mathcal{A} \subset \mathcal{X}: |\mathcal{A}| = t} \frac{1}{2} \log\det\left(\mathbf{I}_t + \sigma^{-2}\mathbf{K}_{v,\mathcal{A}}\right), \tag{22}$$

where $\mathbf{K}_{v,\mathcal{A}} = [k_v(\mathbf{x}, \mathbf{x}')]_{\mathbf{x},\mathbf{x}' \in \mathcal{A}}$ is the positive definite covariance matrix between the points $\mathcal{A}$. Note the term inside of the max in (22) is often referred to as the Shannon Mutual Information between v and the observations at points $\mathbf{x} \in \mathcal{A}$.

Theorem 2. (Cumulative Penalty-based Robust-Regret Upper Bound) Let $\alpha \in (0,1)$ denote the failure probability, $\beta_t = 2\log\left((m+1)|\mathcal{X}|t^2\pi^2/(6\alpha)\right)$, and Assumption 2 hold. Running CARBO (*Algorithm 1*), we obtain the following bound on the cumulative penalty-based robust-regret for all $\rho \geq 0$

$$\mathrm{P}\left\{R^w_{EP,T}(\rho) \leq \sqrt{T\beta_T\Psi_T(\rho)},\ \forall T \geq 1\right\} \geq 1 - \alpha, \tag{23}$$

where $R^w_{EP,T}(\rho) = \sum_{t=1}^{T} r^w_{EP,t}(\rho)$ denotes the cumulative penalty-based robust-regret and

$$\Psi_T(\rho) = \sum_{v \in \mathcal{F}} \sum_{q \in \mathcal{F}} \rho_v \rho_q \sqrt{\widetilde{\gamma}_{v,T}\widetilde{\gamma}_{q,T}}, \tag{24}$$

is an aggregated MIG factor with $\widetilde{\gamma}_{v,T} = \left(8\log\left(1 + \sigma_{\varepsilon,v}^{-2}\right)\right)\gamma_{v,T}$ denoting the scaled MIG for $v \in \mathcal{F}$ after T CARBO iterations.

Proof. Combining Lemmas 1 and 3, we see that the following event

$$\left\{\left(r^w_{EP,t}(\rho)\right)^2 \leq 4\beta_t \sum_{v \in \mathcal{F}} \sum_{q \in \mathcal{F}} \rho_v \rho_q \sigma_{v,t-1}(\theta_{v,t}, w_{v,t})\sigma_{q,t-1}(\theta_{q,t}, w_{q,t}),\ \forall t \geq 1\right\},$$

holds with probability $\geq 1 - \alpha$. Since β_t is non-decreasing with t, we know that the sum over T steps must satisfy

$$\begin{aligned} \sum_{t=1}^{T} \left(r^w_{EP,t}(\rho)\right)^2 &\leq 4\beta_T \sum_{v \in \mathcal{F}} \sum_{q \in \mathcal{F}} \sum_{t=1}^{T} \rho_v \rho_q \sigma_{v,t-1}(\theta_{v,t}, w_{v,t})\sigma_{q,t-1}(\theta_{q,t}, w_{q,t}), \\ &\leq 4\beta_T \sum_{v \in \mathcal{F}} \sum_{q \in \mathcal{F}} \rho_v \rho_q \sqrt{\sum_{t=1}^{T} \sigma^2_{v,t-1}(\theta_{v,t}, w_{v,t})} \sqrt{\sum_{t=1}^{T} \sigma^2_{q,t-1}(\theta_{q,t}, w_{q,t})}, \end{aligned}$$

where the second line follows from the Cauchy–Schwarz inequality. To further simplify the innermost terms, we need to rely on the following established result from [Ref. [58], lemma 5.4]

$$\sum_{t=1}^{T} \sigma^2_{v,t-1}(\theta_{v,t}, w_{v,t}) \leq \frac{2}{\log\left(1 + \sigma_{\varepsilon,v}^{-2}\right)}\gamma_{v,T} = \frac{\widetilde{\gamma}_{v,T}}{4}, \quad \forall v \in \mathcal{F},$$

which upper bounds the sum of posterior variances of sampled points in terms of the scaled MIG. Substituting this result into the inequality above results in

$$\sum_{t=1}^{T} \left(r^w_{EP,t}(\rho)\right)^2 \leq \beta_T \sum_{v \in \mathcal{F}} \sum_{q \in \mathcal{F}} \rho_v \rho_q \sqrt{\widetilde{\gamma}_{v,T}\widetilde{\gamma}_{q,T}} = \beta_T \Psi_T(\rho),$$

where we have substituted the definition of $\Psi_T(\rho)$ provided in (31). Finally, the stated result follows from the fact that $R^w_{EP,T}(\rho) \leq \sqrt{T\sum_{t=1}^{T}\left(r^w_{EP,t}(\rho)\right)^2}$ by the Cauchy–Schwarz inequality.

To explicitly determine the growth rate of $R^w_{EP,T}(\rho)$ with respect to the total number of iterations T, we need to establish bounds on $\Psi_T(\rho)$ defined in (24). This can be done by exploiting bounds on the MIG $\gamma_{v,T}$ for all $v \in \mathcal{F}$, which have been previously identified for common kernel choices. For example, the squared exponential (SE) kernel, which can be derived from (6) when $\nu \to \infty$, results in a MIG that satisfies $\gamma_{v,T} = O\left((\log T)^{p+1}\right)$. Substituting this expression into (24), we see that $\Psi_T(\rho) = O\left((\log T)^{p+1}\right)$ since all terms in the summation have the same order and $\{\rho_v\}_{v \in \mathcal{F}}$ are constants. Finally, substituting this order of magnitude result into (23), we can find that

$$R^{w}_{EP,T}(\rho) = O^{*}\left(\sqrt{T}(\log(T))^{\frac{p+1}{2}}\right), \tag{25}$$

where O^{*} is a variant of the traditional order of magnitude O notation that hides dimension-independent log factors. Similar results can be obtained for other kernel choices, implying $R^{w}_{EP,T}(\rho)$ grows *sublinearly* with respect to T for any $\rho \geq 0$ with high probability for a sufficiently small choice of α. This sublinear property of $R^{w}_{EP,T}(\rho)$ is crucial for establishing global convergence of CARBO, which we show next.

4.3 | Establishing convergence to a constrained robust global minimum

To establish convergence of CARBO to θ^{*}, we need to analyze the effect of the pessimistic recommendation procedure proposed in (14). We now summarize our main result, which shows under what conditions $\theta^{r}_{T}(\rho)$ converges to a global solution of the original problem (1a and 1b).

Theorem 3. (Convergence of CARBO) Let Assumptions 1–3 hold and $\overline{\rho}$ be the threshold value such that $P(\theta;\rho)$ in (25) is an exact penalty function for any $\rho \geq \overline{\rho}$, which must exist according to Theorem 1. Furthermore, let $\{\theta_1, \ldots, \theta_T\}$ be the sequence of points generated by CARBO (*Algorithm 1*) where $\{\beta_t\}_{t\geq 1}$ is chosen to satisfy the conditions of Theorem 2. Then, the recommended point $\theta^{r}_{T}(\rho)$ defined in *(14)* must satisfy the following

$$\mathrm{P}\left\{\lim_{T\to\infty} \theta^{r}_{T}(\rho) = \theta^{*}\right\} \geq 1-\alpha, \quad \forall \rho \in [\overline{\rho}, \infty), \tag{26}$$

where "=" implies in the set of global solutions.

Proof. Let us first define the simple penalty-based robust-regret as $S^{w}_{EP,T}(\rho) = \min_{t\in\{1,\ldots,T\}} r^{w}_{EP,t}(\rho)$, which is similar to (4) but also incorporates the constraint penalty term. Since $r^{w}_{EP,t}(\rho) \leq \overline{r}^{w}_{EP,t}(\rho)$ as shown in the proof of Lemma 3, we can immediately infer that $S^{w}_{EP,T}(\rho) \leq \min_{t\in\{1,\ldots,T\}} \overline{r}^{w}_{EP,t}(\rho) = \overline{r}^{w}_{EP,t^{*}}(\rho)$ where the index t^{*} is defined according the recommendation procedure in (14) that is equivalent to $t^{*} \in \mathit{argmin}_{t\in\{1,\ldots,T\}} \overline{r}^{w}_{EP,t}(\rho)$. A direct consequence of the exact penalty function property is that $S^{w}_{EP,T}(\rho) \geq 0$ for all $T \geq 1$ since $F(\theta) + \rho \,\|\mathbf{G}^{+}(\theta)\|_1 \geq F(\theta^{*})$ for all $\theta \in \Theta$ as long as $\rho \geq \overline{\rho}$. In addition, we note that the cumulative penalty-based robust-regret bound in (23) also holds for the pessimistic version $\overline{R}^{w}_{EP,T}(\rho) = \sum_{t=1}^{T} \overline{r}^{w}_{EP,t}(\rho)$, as the proofs rely on the bound established for $\overline{r}^{w}_{EP,t}(\rho)$. Since the minimum of a sequence of points must be less than or equal to the average of those points, we can establish the following lower and upper bounds on $S^{w}_{EP,T}(\rho)$:

$$0 \leq S^{w}_{EP,T}(\rho) \leq \overline{r}^{w}_{EP,t^{*}}(\rho) \leq \frac{1}{T}\overline{R}^{w}_{EP,T}(\rho), \quad \forall \rho \in [\overline{\rho}, \infty).$$

Because $\overline{R}^{w}_{EP,T}(\rho)$ is sublinear in T, as illustrated in, for example, (25), we must have that $\overline{R}^{w}_{EP,T}(\rho)/T \to 0$ as $T \to \infty$ with probability $\geq 1-\alpha$. Since both the lower and upper bound on $S^{w}_{EP,T}(\rho)$ converge to 0, we must have $S^{w}_{EP,T}(\rho) \to 0$ when the above inequalities holds, which only occurs when $\theta^{r}_{T}(\rho) \to \theta^{*}$. This immediately implies (26) must hold.

Theorem 3 provides a set of conditions that must be satisfied for CARBO to converge. There are two main practical implications of this result. First, the exploration constants $\{\beta_t\}_{t\geq 1}$ and the penalty weight ρ have important impacts on CARBO's convergence. In particular, if we would like a high probability of convergence, then we must correspondingly select larger values for the exploration constants, which may result in worse short-term performance. Furthermore, if we select ρ to be too small, then we are biased toward sampling and recommending infeasible points, implying conservative estimates for ρ should be preferred in practice. The second key insight is that the recommendation procedure plays an important role in the convergence result. We are unable to use the original regret definition since we cannot find the index t^{*} that minimizes the sequence $\left\{r^{w}_{EP,1}(\rho), \ldots, r^{w}_{EP,T}(\rho)\right\}$ since robust-regret is defined in terms of a max operator applied to the unknown functions. By using the pessimistic version of penalty-based robust-regret, we have enough information to identify t^{*} that minimizes the sequence $\left\{\overline{r}^{w}_{EP,1}(\rho), \ldots, \overline{r}^{w}_{EP,T}(\rho)\right\}$ since these terms are fully defined in terms of the known GP upper confidence bounds. In the next section, we discuss several important practical implementation issues including how to numerically implement Algorithm 1 and how to verify the assumptions required to prove Theorems 2 and 3.

Remark 4. It is interesting to note that the recommendation procedure summarized in (14) uses confidence bounds computed from different sets of data. A potentially more robust approach is to replace $\mathrm{UCB}_{v,t-1}$ with a common bound $\mathrm{UCB}_{v,T}$ for all $v \in \mathcal{F}$ so that the selection is based on the most recently available dataset. One can show that the Theorem 3 will still hold for this rule, as long as one intersects the set of confidence bounds at each iteration $\{t-1, \ldots, T\}$ so that they are monotonically shrinking.

5 | PRACTICAL CONSIDERATIONS AND IMPLEMENTATION DETAILS

In Section 4, we analyzed the theoretical convergence of CARBO under certain assumptions. Although these results provide various insights and serve as a useful guide for understanding performance,

our practical implementation differs in certain aspects that are elaborated on in detail in this section.

5.1 | Choice of exploration constant β_t

Lemmas 1 and 3 and Theorems 2 and 3 are presented for discrete spaces $\mathcal{X}$; however, many real-world problems are modeled with continuous spaces. In Ref. [58] a discretization technique was developed to generalize the bounds on the GP-LCB algorithm to compact and convex sets by enlarging the exploration constant β_t. This technique can be applied when the covariance functions $k_v(\mathbf{x}, \mathbf{x}')$ ensure the following high probability bounds on the derivatives of the propagated GP sample paths for f hold for some constants $a, b > 0$

$$\mathrm{P}\left\{\sup_{\mathbf{x}\in\mathcal{X}} \frac{\partial f}{\partial [\mathbf{x}]_i} > L\right\} \leq a e^{-(L/b)^2}, \quad \forall i = 1, \ldots, p. \tag{27}$$

Whenever this condition holds, we can generalize Theorems 2 and 3 to compact and convex sets of the form $\mathcal{X} \subset [0, r]^p$ by replacing β_t with the following larger value

$$\beta_t = 2\log\left(\frac{2(m+1)t^2\pi^2}{3\alpha}\right) + 2p\log\left(t^2 pbr\sqrt{\log(4pa/\alpha)}\right). \tag{28}$$

Note that it has been shown that these values of β_t are often conservative in practice such that a more aggressive strategy is usually preferred. We set $\beta_t = 4$ in our numerical experiments in Section 6 for simplicity, though alternative methods could easily be used.

5.2 | Data-dependent GP prior

Our previous results in Sections 3 and 4 assume that the hyperparameters $\{\Psi, \sigma_{\varepsilon,v}\}$ of the GP prior for all $v \in \mathcal{F}$ are known. However, this is often not true in practice, so that we must rely on some hyperparameter estimation scheme to train the GP models. In this work, we focus on the maximum likelihood estimation (MLE) framework, which determines $\{\Psi_t^\star, \sigma_{\varepsilon,v,t}^\star\}$ at every iteration t by maximizing the log-likelihood function

$$\{\Psi_t^\star, \sigma_{\varepsilon,v,t}^\star\} \in \underset{\Psi, \sigma_{\varepsilon,v}^2}{\operatorname{argmin}} \mathcal{L}_{v,t}(\Psi, \sigma_{\varepsilon,v}) = \log\left(p\left(\mathbf{y}_{v,t}|\mathbf{X}_{v,t}, \Psi, \sigma_{\varepsilon,v}\right)\right). \tag{29}$$

According to the GP prior, the measured data vector $\mathbf{y}_{v,t}$ must follow a multivariate Gaussian distribution, that is,

$$\mathbf{y}_{v,t} \sim \mathcal{N}(0, \Sigma_{v,t}(\Psi, \sigma_{\varepsilon,v})), \quad [\Sigma_{v,t}(\Psi, \sigma_{\varepsilon,v})]_{ij} = k_v\left(\mathbf{x}_i, \mathbf{x}_j|\Psi\right) + \sigma_{\varepsilon,v}^2\delta_{ij}, \quad \forall i,j \in \{1, \ldots, t\}, \tag{30}$$

where δ_{ij} denotes the Dirac delta function that is 1 when $i = j$ and 0 otherwise. We can then derive the following analytic expression for the log-likelihood function

$$\mathcal{L}_{v,t}(\Psi, \sigma_{\varepsilon,v}) = -\mathbf{y}_{v,t}^T\Sigma_{v,t}^{-1}(\Psi, \sigma_{\varepsilon,v})\mathbf{y}_{v,t} - \frac{1}{2}\log\left(\det(\Sigma_{v,t}(\Psi, \sigma_{\varepsilon,v}))\right) - \frac{p}{2}\log(2\pi). \tag{31}$$

Since (31) is a differentiable function, the GP training optimization problem (29) is a nonlinear program that can be efficiently solved to local optimality using state-of-the-art methods such as IPOPT.[63] To avoid getting stuck in shallow local minima, we first use the global solver DIRECT[64] to find a good initial condition to provide to IPOPT (i.e., the best solution found with DIRECT under its default settings is used to initialize IPOPT).

Note that re-training the set of GP models for all $v \in \mathcal{F}$ at each iteration of CARBO can be somewhat expensive depending on the size of the optimization problem (29), the number of data points, and the number of constraints. A simple way to reduce this cost is to only periodically update the hyperparameters (e.g., every 5 iterations). Although this may slightly reduce performance, it is a strategy that has been successfully employed in the Bayesian optimization literature on a variety of different problems. Another advantage of (periodically) updating the GP prior hyperparameters as new data is obtained is that one can monitor the estimated values to see if they converge as the number of iterations increase, which will occur whenever Assumption 2 is satisfied.

5.3 | Verification of upper-level MFCQ conditions

The proofs of Theorems 2 and 3 invoked the constraint qualification condition in Assumption 3, which we cannot easily verify holds a priori since the functions $v \in \mathcal{F}$ are fully black-box in nature. A reasonable alternative is to instead check if Assumption 3 holds after the CARBO algorithm is terminated. Under Assumption 2, the posterior mean function in (8a)–(8c) should provide a reasonably accurate representation of the true unknown functions, meaning we can verify Assumption 3 for the learned GP models for every $v \in \mathcal{F}$. To run this verification procedure, one would need to identify the set of global solutions to the mean-based approximation to (2) (using, e.g., a combination of a multistart procedure with some established numerical method for GSIPs) and check if the MFCQ conditions hold at each of these points. Since the posterior GP mean functions are differentiable and cheap to evaluate, this procedure should require significantly less computational cost than running the complete CARBO algorithm.

5.4 | Choice of penalty weight factor ρ

As shown in Theorem 3, we can achieve convergence for any value of $\rho \geq \bar{\rho}$. Since the functions $v \in \mathcal{F}$ are unknown, we cannot exactly determine $\bar{\rho}$ before running CARBO and must instead rely on conservative estimates. Assuming the functions $v \in \mathcal{F}$ have been reasonably well-scaled, the easiest approach is to select a large value for ρ (e.g., a value on the order to $10^3 - 10^5$). We set $\rho = 10^3$ in our simulation results. It

should be noted that large ρ values may make the min–max optimization problem (2) harder to solve numerically. An alternative approach would be to take advantage of the strong duality result established for (18), that is, $\max_{\rho \geq 0} \min_{\theta \in \Theta} P(\theta;\rho) = F(\theta^\star)$ by replacing $P(\theta;\rho)$ with its upper confidence bound $\mathrm{UCB}_{P,t-1}(\theta;\rho)$ that can be expressed in terms of (9a and 9b). One can practically solve this trilevel optimization problem by solving the inner min–max optimization problem for an increasing set of ρ values. A conservative estimate of $\bar{\rho}$ then corresponds to the ρ such that $\min_{\theta \in \Theta} \mathrm{UCB}_{P,t-1}(\theta;\rho)$ converges to a nearly constant value. Similarly to the hyperparameter training procedure described in Section 5.2, one can update the estimate of $\bar{\rho}$ periodically to reduce the computational cost of repeatedly solving the min–max optimization problems.

5.5 | Initialization procedure

Since the MLE-based training procedure described in the previous section can be unreliable when we have a very small amount of data, we follow the recommendations from Ref. [65] by initializing CARBO with N_{init} points selected uniformly at random. This procedure is meant to overcome limitations in the first few iterations whenever we have an uninformative prior for each v, which is often the case in practice. This procedure can easily be modified or replaced when additional information is available.

5.6 | Constrained min–max optimization of the CARBO acquisition function

The theoretical analysis in Section 4 assumes we can exactly optimize the acquisition functions defined in terms of the lower and upper confidence bounds in (12) and (13), respectively. The maximization problem (13) resembles a standard sub-problem that arises in BO (or GP-LCB to be specific) such that the same procedures can be exploited to practically solve this finite nonlinear program (NLP). In particular, we rely on a combination of derivative-free global search with a local gradient-based solver (IPOPT) to refine the solution. The min–max optimization problem (12), which can be equivalently formulated as the constrained problem (15a and 15b), is a much more challenging problem to solve. We could treat (15a and 15b) as a robust black-box optimization problem with constraints and apply the method developed in Ref. [2] since the GP mean and variance evaluations are cheap compared with the original unknown functions. However, this method will produce a limited convergence rate in practice since it does not exploit available derivative information from the posterior GP mean and variance equations. Instead, we recognize that (15a and 15b) can be formulated as a generalized semi-infinite program (GSIP),[66] that is, a mathematical program with a finite number of decision variables subject to an infinite number of constraints the index set of which is dependent on the decision variables, as follows

$$\begin{aligned} &\min_{\theta \in \Theta, \varepsilon_i \geq 0, z} z, \\ &\text{s.t. } \mu_{f,t-1}(\theta, w) - \beta_t^{1/2} \sigma_{f,t-1}(\theta, w) + \rho \sum_{i=1}^{m} \varepsilon_i - z \leq 0, \quad \forall w \in \mathcal{W}(\theta), \\ &\mu_{g_i,t-1}(\theta, w) - \beta_t^{1/2} \sigma_{g_i,t-1}(\theta, w) - \varepsilon_i \leq 0, \quad \forall i \in \{1, \ldots, m\}, \quad \forall w \in \mathcal{W}(\theta). \end{aligned} \tag{32}$$

Whenever $\mathcal{W} = \mathcal{W}(\theta)$ is independent of θ, which is a common case in practice, the GSIP reduces to a standard semi-infinite program (SIP) for which efficient algorithms have been developed. A particularly simple and efficient approach for SIPs with non-convex lower-level problems, such as (32), is described in Ref. [26] which we use in our numerical experiments in Section 6. This approach obtains an upper bound by restricting the right-hand side of the constraints above, that is, replacing ≤ 0 with a negative value that changes in the algorithm. To be able to evaluate this upper bound using a finite NLP, the infinite set $\mathcal{W}$ is replaced by a finite set $\mathcal{W}^{UBD} \subset \mathcal{W}$. The set $\mathcal{W}^{UBD}$ is sequentially populated in a way that guarantees convergence under relatively mild assumptions. A converging lower bound is also obtained in a similar fashion using successively tighter discretization following the principle of Ref. [25]. Therefore, in addition to exploiting derivative information, we can get a certificate of optimality for any user specified tolerance value, as long as we use global NLP solvers to solve some of the sub-problems.

5.7 | Multipoint recommendation procedures

There is a natural tradeoff between performance and constraint satisfaction in many real-world engineering problems, which can be complicated even further when dealing with the effect of uncertainties. Therefore, as commonly done in the multiobjective optimization literature,[67] it can be useful to recommend a set of viable solutions instead of a single value as done in (14). To this end, as opposed to just taking the minimum of the sequence, one may rank the results according to the metric defined by the objective function in (14). In addition to looking at this aggregated metric, however, one can also look at the individual pessimistic estimates of the worst-case objective and constraint values. We have found that CARBO often generates natural clusters of points, with some clusters clearly producing poor worst-case objective and constraint values that can be eliminated from further consideration. Out of the remaining points, there may be some that provide a similar value to the minimum of the aggregated penalty-based metric by improving the worst-case objective at the cost of small amount of (predicted) worst-case constraint violation. Since the metric in (14) is merely an estimate, we recommend allocating some additional budget to improve upon the worst-case estimates for the objective and any nearly active constraint by running a standard BO procedure for these functions at "good" fixed θ_t values. The definition of "good" and the allowable evaluation budget will be problem-dependent; however, a good rule-of-thumb is to prefer points whose worst-case upper confidence bound for the constraints is not too close to 0 whenever robust constraint satisfaction is essential.

5.8 | Modified search method to locate robustly feasible points

Whenever robust constraint satisfaction is a critical requirement and we have a limited evaluation budget, it may be preferable to initially ignore the objective function by setting $f(\theta, w) \approx 0$. Although CARBO is still applicable in this case, it turns out that the min–max optimization problem (12) does not preferentially select between different values of θ whose lower confidence bound robustly satisfies constraints. Therefore, it can be advantageous to replace (12) with the following alternative in such cases:

$$\theta_t \in \underset{\theta \in \Theta}{\operatorname{argmax}} \prod_{i=1}^{m} \mathrm{P}\left\{ \max_{w \in \mathcal{W}(\theta)} g_i(\theta, w) \le 0 \right\}, \tag{33}$$

which directly maximizes the probability of robust constraint satisfaction. Note that the evaluation of the probability operator in (33) is very difficult since the maximum of a GP is no longer a GP. Thus, in practice, it is likely preferred to reformulate the probability evaluation in terms of an expectation over a standard normal $Z_i \sim \mathcal{N}(0,1)$ using the reparametrization trick[68] as follows

$$\mathrm{P}\left\{ \max_{w \in \mathcal{W}(\theta)} g_i(\theta, w) \le 0 \right\} = \mathbb{E}_{Z_i}\left\{ \mathbf{1}_{\infty 0]}\left(\max_{w \in \mathcal{W}(\theta)} \left(\mu_{g_i, t-1}(\theta, w) + \sigma_{g_i, t-1}(\theta, w) Z_i \right) \right) \right\}, \tag{34}$$

where $\mathbb{E}_{Z_i}\{\cdot\}$ denotes the expectation with respect to random variable Z_i. One can then directly estimate this expected value for any $\theta \in \Theta$ using Monte Carlo (MC) integration. A hybrid method that initially uses (33) to find at least one robustly feasible θ_t value (with high probability) and then switches to (12) for improved objective performance would be an interesting approach to study further in future work.

6 | CASE STUDIES

In this section, we demonstrate the performance of CARBO on two problems. The first case study is a benchmark problem for robust optimization with constraints. Since we know the global solution to this problem, we can straightforwardly compute the simple penalty-based robust-regret measure at every iteration in order to test the convergence claim made in Theorem 3. For our second case study, we consider a challenging robust design problem for a bubble column reactor that converts industrial waste to valuable liquid fuels using cellular fermentation. Since this bubble column problem is defined in terms of an expensive high-fidelity simulator, we do not have exact knowledge of the true solution. Therefore, instead of analyzing penalty-based robust-regret, we demonstrate CARBO's ability to identify a set of robust design parameters that result in good worst-case performance and constraint satisfaction.

6.1 | Benchmark problem

We first consider a problem from Ref. [2] that is defined in terms of a polynomial objective function with polynomial constraint functions of the form

$$\begin{aligned} f_{\text{poly}}(\theta) = {} & 2\theta_1^6 - 12.2\theta_1^5 + 21.2\theta_1^4 - 6.4\theta_1^3 - 4.7\theta_1^2 + 6.2\theta_1 \\ & + \theta_2^6 - 11\theta_2^5 + 43.3\theta_2^4 - 74.8\theta_2^3 + 56.9\theta_2^2 - 10\theta_2 \\ & - 4.1\theta_1\theta_2 - 0.1\theta_2^2\theta_1^2 + 0.4\theta_2^2\theta_1 + 0.4\theta_1^2\theta_2, \end{aligned} \tag{35a}$$

$$g_{\text{poly},1}(\theta) = (\theta_1 - 1.5)^4 + (\theta_2 - 1.5)^4 - 10.125, \tag{35b}$$

$$g_{\text{poly},2}(\theta) = -(2.5 - \theta_1)^3 - (\theta_2 + 1.5)^3 + 15.75. \tag{35c}$$

The feasible set for the design variables is $\theta = [\theta_1, \theta_2]^T \in \Theta = [-1, 4]^2 \subset \mathrm{R}^2$. To define a problem of the form (1a and 1b), we incorporate the effective of implementation errors (additive uncertainty) as follows

$$f(\theta, w) = f_{\text{poly}}(\theta + w), \tag{36a}$$

$$g_1(\theta, w) = g_{\text{poly},1}(\theta + w), \tag{36b}$$

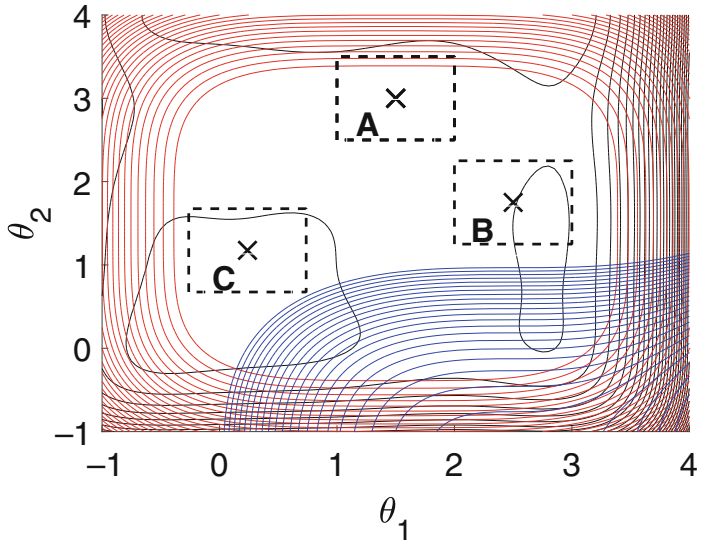

FIGURE 1 Contour plot of the nominal objective f_{poly} (black) and constraint $g_{\text{poly},1}$ (blue) and $g_{\text{poly},2}$ (red) functions for the benchmark problem. The points A, B, and C indicate three different design choices with the corresponding boxes showing the range of uncertainty. Point A is infeasible, Point B is feasible but results in large worst-case objective values, and Point C is the robust optimal solution.

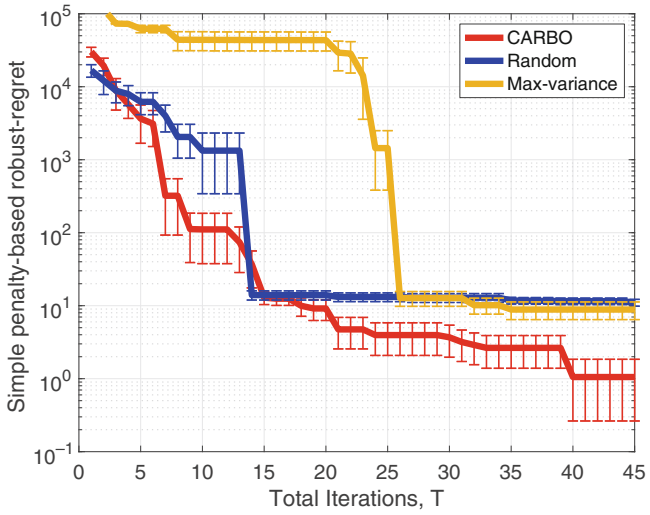

FIGURE 2 Simple penalty-based robust-regret $S^{w}_{EP,T}(\rho)$ with $\rho = 10^3$ for CARBO (red), random search (blue), and max-variance (yellow) on the benchmark problem. The runs are repeated five times from different random initial conditions and the estimated average $S^{w}_{EP,T}(\rho)$ with corresponding confidence intervals are shown as error bars.

$$g_2(\theta, w) = g_{\text{poly},2}(\theta + w), \tag{36c}$$

where the set of possible uncertainty values is given by $\mathcal{W} = [-0.5, 0.5]^2 \subset \mathrm{R}^2$. A contour plot of the nominal objective and constraint functions in (35a)–(35c) is shown in Figure 1. We have labeled three points A, B, and C to illustrate how uncertainty plays an important role in this problem. In particular, point A is feasible for the nominal set of constraints but becomes infeasible once uncertainty is considered in (36a)–(36c). Point B, on the other hand, remains feasible under perturbations but a relatively large worst-case objective value (that is close to a local minima for this problem). Point C is the global robust solution to this problem, as it clearly results in robust constraint satisfaction as well as a low worst-case objective value (since the function is relatively flat in this region).

Our goal is to identify the constrained robust optimal solution (Point C) in as few iterations as possible using only zeroth-order information from (36a)–(36c). In accordance with Theorem 3, we use the simple penalty-based robust-regret $S^{w}_{EP,T}(\rho) = \min_{t \in \{1,\ldots,T\}} r^{w}_{EP,t}(\rho)$ as our metric. As discussed in Section 5.5, we initially select $N_{\text{init}} = 5$ points uniformly at random before running CARBO (Algorithm 1). Since $S^{w}_{EP,T}(\rho)$ is a function of the randomly selected initial points, $S^{w}_{EP,T}(\rho)$ is a random variable. We thus repeat CARBO $N_{\text{repeat}} = 5$ times to obtain a sample average estimate for the expected simple penalty-based robust-regret

$$\mathbb{E}\{S^{w}_{EP,T}(\rho)\} \approx \frac{1}{N_{\text{repeat}}} \sum_{i=1}^{N_{\text{repeat}}} S^{w,(i)}_{EP,T}(\rho), \tag{37}$$

where $S^{w,(i)}_{EP,T}(\rho)$ denotes the simple penalty-based robust-regret for the ith CARBO run. Because this estimate is constructed from a finite number of samples, we also report confidence intervals estimated as 1.96 times the sample-based standard deviation divided by the number of replicates (based on the standard error formula).

As noted previously, CARBO is one of the first algorithms developed for the challenging class of problems considered in this work such that we have limited options to compare against. Thus, we select two baseline algorithms: (i) random search and (ii) the max-variance method, which are commonly used in the black-box optimization literature. Random search selects the sampled point (θ_t, w_t) by drawing samples uniformly at random from the set $\Theta \times \mathcal{W}$. The max-variance methods selects the sampled point according to (10). The simple penalty-based robust-regret plots for CARBO, random search, and max-variance are shown in Figure 2. We clearly see that CARBO converges at a much faster rate than random search and max-variance. We also see that, on average, the value of $S^{w}_{EP,T}(\rho)$ drops more than an order of magnitude within the first 10 iterations and drops another order of magnitude within the next 10 iterations, which highlights the quick progress toward the global robust solution. We also plot the set of recommended points using (14) at three different set of total iterations $T = 5, 33, 39$ for CARBO, random search, and max-variance in Figure 3 for additional insights. We see that the recommended points are relatively spread out in the early iterations (left, $T = 5$), but get progressively closer to the true robust global solution (black star) in the later iterations (right, $T = 33$ and bottom, $T = 39$). Random search and max-variance, on the other hand, consistently lead to recommended points that are far away from the true solution since neither method can navigate the exploration–exploitation tradeoff. This is a key limitation, especially when a significant portion of Θ leads to violation of the robust constraints.

6.2 | Design of a bubble column fermentation reactor for waste gas recovery

Gas fermentation has emerged a promising route for converting industrial waste gases and synthesis gas into renewable liquid fuels and chemicals using specially developed bacteria. *Clostridium autoethaogenum*, which is an acetogenic anaerobic bacterium, has the potential to be effective at fermenting carbon monoxide into ethanol and acetate through the Wood–Ljungdahi pathway; however, the yield of ethanol in the wild-type strain is known to be low for ethanol compared with acetate.[69] An improved strain of *C. autoethaogenum* that provides an increased biofuel yield has been developed by researchers at LanzaTech.[70] However, the majority of studies on this system have been done at the bench-scale using continuously stirred tank reactors (CSTRs), which are not feasible for large-scale production. Bubble column technology, on the other hand, are able to provide good heat and mass transfer efficiencies at low operating cost due to mixing from established gas sparging configurations. Bubble columns are more challenging to optimize due to spatial variation throughout the column, which can lead to significantly different growth environments as a function of column position. In Ref. [71] a detailed spatiotemporal metabolic model was developed that

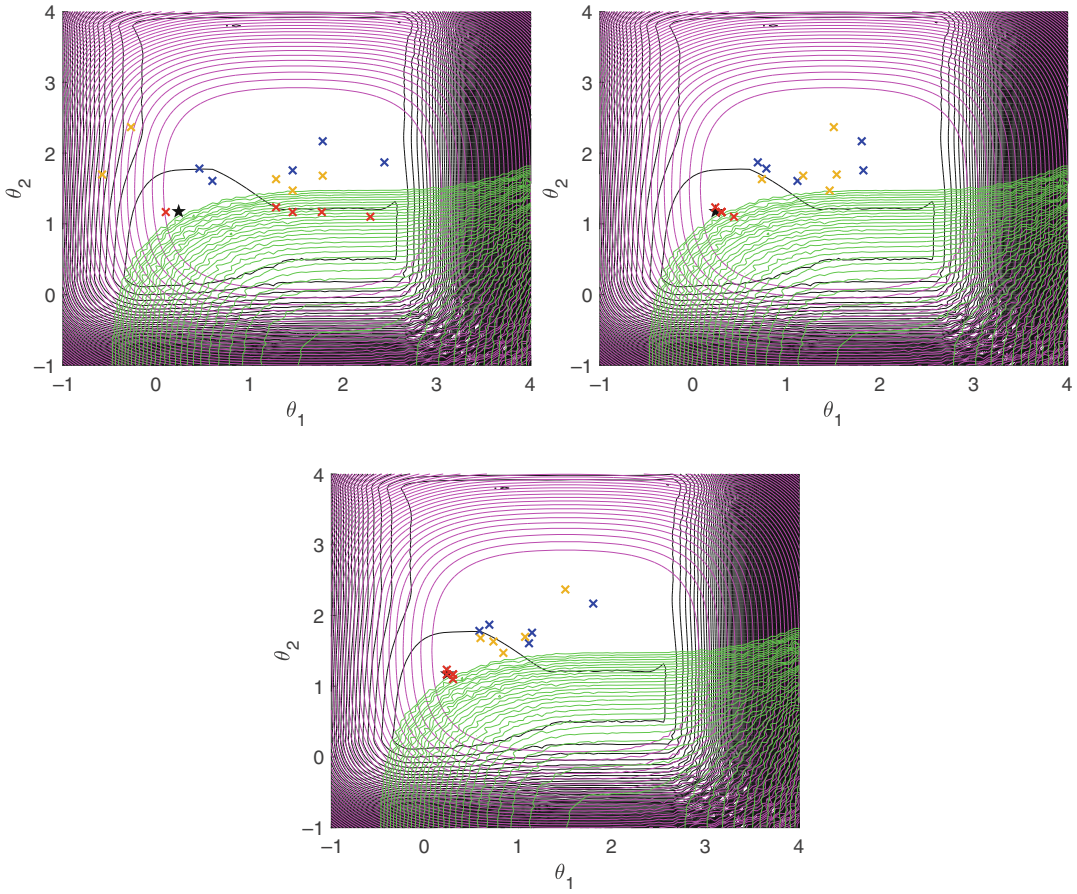

FIGURE 3 Recommended points based on the sequence of $\{\theta_t\}_{t=1}^{T}$ values generated by CARBO (red), random search (blue), and max-variance (yellow) for the 5 runs used to generate Figure 2 for $T = 5$ (left), $T = 33$ (right), and $T = 39$ (bottom). The black star denotes the true unknown constrained robust global minimum. The black, green, and magenta contour lines correspond to the worst-case objective, constraint 1, and constraint 2 functions, respectively.

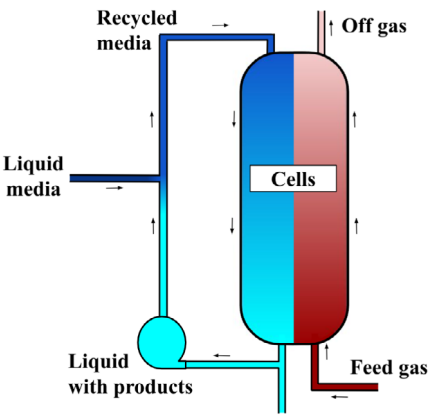

FIGURE 4 Schematic of the bubble column reactor for fermenting carbon monoxide into ethanol and acetate considered in Ref. [71].

involves two key components: (i) a set multiphase convection–dispersion equations that govern transport of carbon monoxide, secreted byproducts, and biomass and (ii) a genome-scale reconstruction of *C. autoethaogenum* metabolism using the flux balance analysis

TABLE 1 Minimum and maximum values for the considered design variables and uncertain variables in the bubble column reactor case study

Variable	Minimum	Maximum	Units
u_g	9.84	14.76	m/h
u_l	−190.2	−126.9	m/h
T_R	307.15	313.15	Kelvin
$k_L a_{CO}$	518	528	1/h

method. A schematic of the overall process is shown below in Figure 4. The focus of Ref. [71] was on the validation of this spatio-temporal model using experimental data. To the best of our knowledge, no systematic optimization studies have been performed on this model due to its computational complexity and the presence of uncertainty, making it a great candidate problem for CARBO, which is able to address both of these challenges.

We developed our optimization model based on recommendations from Ref. [71]; interested readers are referred to this publication for a complete description of the model. We focus on two decision variables, mainly the superficial gas velocity u_g and the liquid phase

TABLE 2 Top 10 design values sampled by CARBO for the bubble column reactor case study

Iteration #	Worst-case objective $\times 10^{-2}$	Worst-case constraint 1 $\times 10^{-3}$	Worst-case constraint 2 $\times 10^{-3}$	Design 1 (θ_1)	Design 2 (θ_2)
61	7.805	−2.316	−9.892	0.936	1
36	8.062	−3.946	−9.956	0.901	1
90	8.514	−4.911	−10.060	0.841	1
73	8.776	−6.994	−10.141	0.806	1
78	8.896	−7.425	−10.165	0.790	1
77	9.481	−9.131	−10.277	0.714	1
9	9.543	−6.307	−5.075	0.714	0
75	9.617	−9.557	−10.299	0.697	1
46	9.913	−7.708	−5.059	0.667	0
84	9.923	−10.106	−10.348	0.658	1

Note: The worst-case objective and constraint values correspond to the individual terms in (39).

velocity u_l, and two key uncertain parameters, mainly the reactor temperature T_R and the gas–liquid mass transfer coefficient for carbon monoxide $k_L a_{CO}$. The considered ranges and units for these variables are summarized in Table 1. We define θ and w in terms of scaled versions of these variables, that is,

$$\theta_1 = \frac{u_g - u_{g,\min}}{u_{g,\max} - u_{g,\min}}, \quad \theta_2 = \frac{u_l - u_{l,\min}}{u_{l,\max} - u_{l,\min}},$$
$$w_1 = \frac{T_R - T_{R,\min}}{T_{R,\max} - T_{R,\min}}, \quad w_2 = \frac{k_L a_{CO} - k_L a_{CO,\min}}{k_L a_{CO,\max} - k_L a_{CO,\min}}.$$

The objective is to maximize the production of ethanol while ensuring constraints on the selectivity (of ethanol vs. acetate production) and the time-to-steady-state (TSS) are satisfied. Consequently, we can state this problem in the form of the original problem (1a and 1b) as follows

$$\begin{aligned} \min_{\theta \in \Theta} \max_{w \in \mathcal{W}} \; & f(\theta, w) := \frac{-C_E(\theta, w) + 15\,\text{g/L}}{15\,\text{g/L}}, \\ \text{s.t. } & g_1(\theta, w) := \frac{\text{TSS}(\theta, w) - 825\,\text{h}}{825\,\text{h}} \le 0, \quad \forall w \in \mathcal{W}, \\ & g_2(\theta, w) := \frac{1.515\,\text{g ethanol/g acetate} - S_{E/A}(\theta, w)}{1.515\,\text{g ethanol/g acetate}} \le 0, \quad \forall w \in \mathcal{W}, \end{aligned} \tag{38}$$

where C_E, TSS, and $S_{E/A}$ denote the ethanol concentration, time-to-steady-state, and steady-state selectivity, respectively, which can be computed by evaluating the aforementioned model at any $(\theta, w) \in \Theta \times \mathcal{W}$. Note that we have appropriately scaled the objective and constraint functions to simplify the presentation of the results. We used the high-fidelity simulator code available online (http://www.ecs.umass.edu/che/henson_group/downloads.html) to perform the evaluations needed by CARBO.

Due to the complexity of the simulation-based model considered here, we do not have a priori knowledge of the exact robust global solution such that we cannot verify the theoretical convergence results (as we were able to do for the benchmark problem). We instead only conduct a single CARBO run, assuming a fixed number of iterations can be run, which is likely how CARBO would be applied to a real-world problem. In line with Section 5.5, we select the first $N_{\text{init}} = 10$ samples uniformly at random and the remaining evaluations using CARBO with $T = 90$. As discussed in Section 5.7, we use a multi-point recommendation procedure that ranks points θ_t based on the following metric:

$$\left(\max_{w \in \mathcal{W}(\theta_t)} \text{UCB}_{f,T}(\theta_t, w)\right) + \rho\left(\max_{w \in \mathcal{W}(\theta_t)} \text{UCB}_{g_1,T}(\theta_t, w)\right)^+ + \rho\left(\max_{w \in \mathcal{W}(\theta_t)} \text{UCB}_{g_2,T}(\theta_t, w)\right)^+. \tag{39}$$

The top 10 design values, ranked according to (39) from smallest to highest, are shown in Table 2. We see that the pessimistic estimate of the worst-case constraint violation is less than 0 for all 10 of these designs, such that the second two terms in (39) are 0. The first three designs indicate a tradeoff between worst-case performance and constraint satisfaction within the neighborhood of $0.85 \le \theta_1 \le 0.95$ and $\theta_2 = 1$. To validate this prediction, we perform additional testing on these three design values. In particular, we apply Monte Carlo (MC) sampling by running the high-fidelity simulator for 200 randomly sampled $w \in \mathcal{W}$ values for each of these fixed designs. Furthermore, to highlight the importance of considering uncertainty during the optimization process, we also apply the same MC sampling procedure to the nominal design estimated by solving the following optimization problem

$$\theta_{\text{nominal}} \in \operatorname*{argmin}_{\theta \in \Theta} f(\theta, w_{\text{nominal}}) \text{ s.t. } g_1(\theta, w_{\text{nominal}}) \le 0, \; g_2(\theta, w_{\text{nominal}}) \le 0, \tag{40}$$

where $w_{\text{nominal}} = [0.5, 0.5]^T$ is the nominal uncertainty value. The results of the MC sampling validation procedure are shown in Figure 5. From these results, we clearly see that the top three ranked designs perform much better than the nominal solution with respect to constraint satisfaction. As expected, this improved constraint satisfaction capability comes at the cost of performance (i.e., steady-state

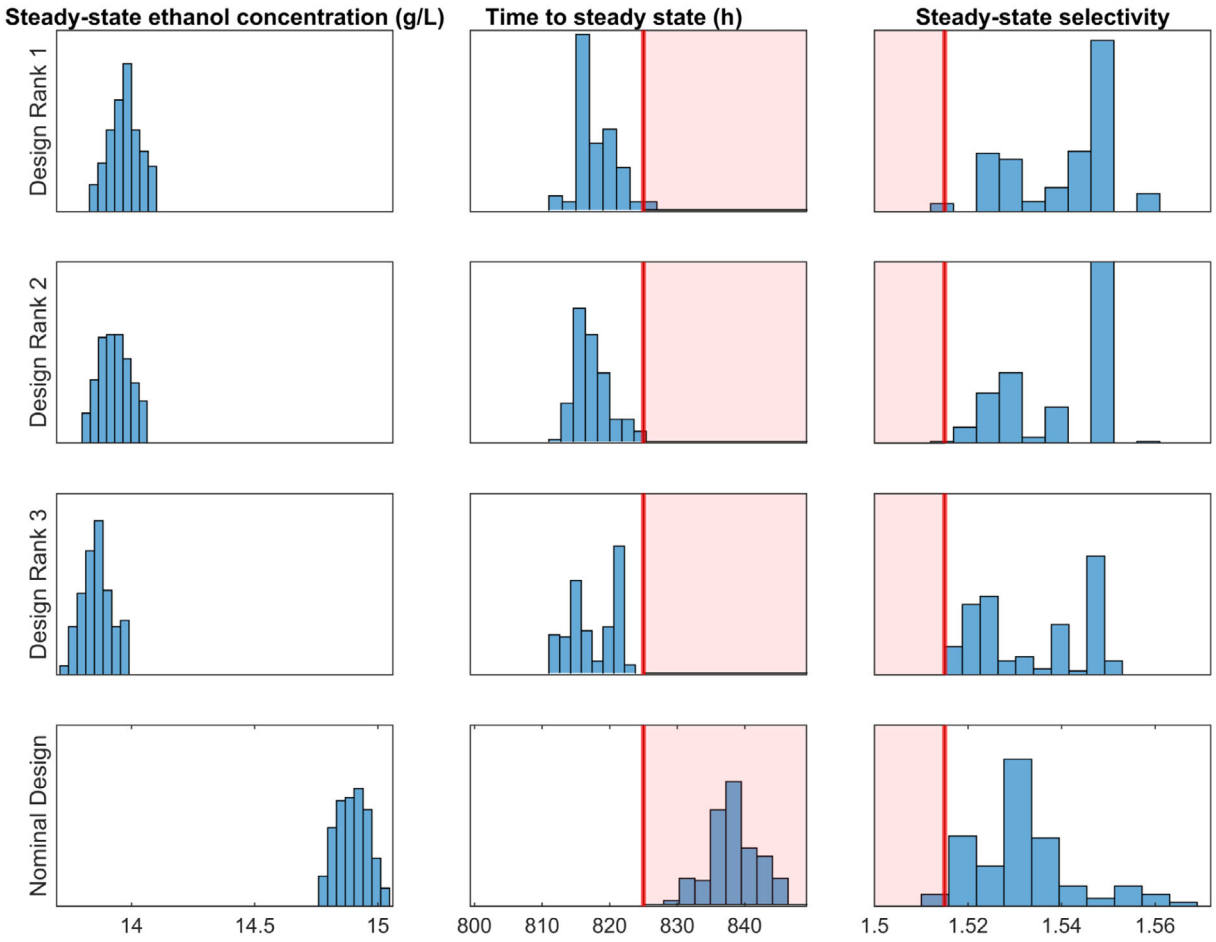

FIGURE 5 Histograms of the worst-case objective and constraint values obtained from 200 Monte Carlo samples of the high-fidelity simulator for the top three ranked designs from Table 2 (first three rows) and the nominal design from (40) (fourth row). The shaded red region in the middle and right columns indicates the infeasible constraint regions.

ethanol concentration), though average performance only drops around 5%. In addition, we see that constraint violation probability drops from around 2% to 0.5% to 0% for design ranks 1, 2, and 3, respectively, which exactly matches the predictions made in Table 2.

7 | CONCLUSIONS AND FUTURE WORK

This article presents a new algorithm, referred to as CARBO, for solving robust optimization problems with constraints in a derivative-free fashion wherein the underlying functions are defined in terms of noisy and expensive high-fidelity simulations and/or experiments. CARBO uses probabilistic surrogate models to jointly describe the effect of the design variables and uncertainties on the black-box objective and constraint functions. In particular, GP models are used because they allow for simple analytic expressions for the posterior mean and variance (given available data in the form of function evaluations at known points) that can be used to construct upper and lower confidence bounds for the unknown objective and constraint functions. At each iteration, CARBO follows two steps to decide the next batch of sample points: (i) select the design value that minimizes the lower confidence bound for a penalized version of the original robust problem and (ii) select the set of uncertainty values that maximize the upper confidence bound of the unknown objective and constraint functions at the fixed design value from the previous step. These alternating optimistic and pessimistic steps are repeated at every iteration until some maximum allowed evaluation budget is exhausted. As a final step, CARBO employs a novel recommendation procedure to recommend a point out of the set of sampled points that is likely the closest to the true global solution of the original robust problem. Using the concept of *exact penalty functions*, we theoretically prove that CARBO converges to the global solution with high probability under certain assumptions by establishing rigorous bounds on a penalized version of cumulative robust-regret (which is a commonly used metric for how close a set of sampled points are to the true solution). We also discuss several important practical implementation details and extensions of CARBO. To demonstrate its effectiveness, we apply CARBO to two simulation case studies including a non-convex benchmark problem and a realistic engineering design problem related to industrial waste gas recovery using microbial fermentation. Our results highlight that CARBO can consistently identify near-global robust solutions that

ensure worst-case constraint satisfaction using on the order of tens to hundreds of function evaluations, which is considerably less than alternative methods. Our future work will mainly focus on incorporating more complex non-Gaussian noise models, extending CARBO to work for multiobjective problems, and studying the effect of the confidence bound widths and penalty weight factor on the rate of convergence.

AUTHOR CONTRIBUTIONS

Akshay Kudva: Conceptualization (supporting); formal analysis (supporting); methodology (supporting); software (lead); visualization (lead); writing – original draft (lead); writing – review and editing (supporting). **Farshud Sorourifar:** Conceptualization (supporting); formal analysis (supporting); methodology (supporting); software (supporting); visualization (supporting); writing – original draft (supporting); writing – review and editing (supporting). **Joel Paulson:** Conceptualization (lead); formal analysis (lead); methodology (lead); resources (lead); software (supporting); visualization (supporting); writing – original draft (supporting); writing – review and editing (lead).

FUNDING INFORMATION

JAP acknowledges funding support from the National Science Foundation (Grant No. 2029282). FS acknowledges funding support from the NSF Graduate Research Fellowship.

DATA AVAILABILITY STATEMENT

The codes that support the findings of this study are openly available on GitHub at https://github.com/PaulsonLab.

ORCID

Joel A. Paulson https://orcid.org/0000-0002-1518-7985

REFERENCES

1. Bertsimas D, Nohadani O, Teo KM. Robust optimization for unconstrained simulation-based problems. *Oper Res*. 2010;58(1):161-178. doi:10.1287/opre.1090.0715
2. Bertsimas D, Nohadani O, Teo KM. Nonconvex robust optimization for problems with constraints. *INFORMS J Comput*. 2010;22(1):44-58. doi:10.1287/ijoc.1090.0319
3. Castillo E, Mnguez R, Castillo C. Sensitivity analysis in optimization and reliability problems. *Reliab Eng Syst Saf*. 2008;93(12):1788-1800.
4. Sudret B. Global sensitivity analysis using polynomial chaos expansions. *Reliab Eng Syst Saf*. 2008;93(7):964-979.
5. Paulson JA, Martin-Casas M, Mesbah A. Fast uncertainty quantification for dynamic flux balance analysis using non-smooth polynomial chaos expansions. *PLoS Comput Biol*. 2019;15(8):e1007308.
6. Sahinidis NV. Optimization under uncertainty: state-of-the-art and opportunities. *Comput Chem Eng*. 2004;28(6-7):971-983.
7. Ning C, You F. Optimization under uncertainty in the era of big data and deep learning: when machine learning meets mathematical programming. *Comput Chem Eng*. 2019;125:434-448.
8. Keith AJ, Ahner DK. A survey of decision making and optimization under uncertainty. *Ann Oper Res*. 2021;300(2):319-353.
9. Djelassi H, Mitsos A, Stein O. Recent advances in nonconvex semi-infinite programming: applications and algorithms. *EURO J Comput Optim*. 2021;9:100006.
10. Birge JR, Louveaux F. *Introduction to Stochastic Programming*. Springer Science & Business Media; 2011.
11. Shapiro A, Dentcheva D, Ruszczynski A. Lectures on stochastic programming: modeling and theory. *SIAM*. 2021. https://epubs.siam.org/doi/book/10.1137/1.9781611973433
12. Halemane KP, Grossmann IE. Optimal process design under uncertainty. *AIChE J*. 1983;29(3):425-433.
13. Ben-Tal A, Nemirovski A. Robust optimization–methodology and applications. *Math Program*. 2002;92(3):453-480.
14. Beyer HG, Sendhoff B. Robust optimization–a comprehensive survey. *Comput Methods Appl Mech Eng*. 2007;196(33–34):3190-3218.
15. Gabrel V, Murat C, Thiele A. Recent advances in robust optimization: An overview. *Eur J Oper Res*. 2014;235(3):471-483.
16. Vázquez FG, Rückmann J, Stein O, Still G. Generalized semi-infinite programming: a tutorial. *J Comput Appl Math*. 2008;217(2):394-419.
17. Hettich R, Kortanek KO. Semi-infinite programming: theory, methods, and applications. *SIAM Rev*. 1993;35(3):380-429.
18. Reemtsen R, Görner S. Numerical methods for semi-infinite programming: a survey. In: Reemtse R, ed. *Semi-Infinite Programming*. Springer; 1998:195-275.
19. Goberna MÁ, López MA. *Semi-Infinite Programming: Recent Advances*. Kluwer Academic Publishers; 2001.
20. Bhattacharjee B, Lemonidis P, Green WH Jr, Barton PI. Global solution of semi-infinite programs. *Math Program*. 2005;103(2):283-307.
21. Stein O. How to solve a semi-infinite optimization problem. *Eur J Oper Res*. 2012;223(2):312-320.
22. Žaković S, Pantelides C, Rustem B. An interior point algorithm for computing saddle points of constrained continuous minimax. *Ann Oper Res*. 2000;99(1):59-77.
23. Bertsimas D, Sim M. Tractable approximations to robust conic optimization problems. *Math Program*. 2006;107(1):5-36.
24. Ben-Tal A, Nemirovski A. Selected topics in robust convex optimization. *Math Program*. 2008;112(1):125-158.
25. Blankenship JW, Falk JE. Infinitely constrained optimization problems. *J Optim Theory Appl*. 1976;19(2):261-281.
26. Mitsos A. Global optimization of semi-infinite programs via restriction of the right-hand side. *Optimization*. 2011;60(10-11):1291-1308. doi:10.1080/02331934.2010.527970
27. Stuber MD, Barton PI. Semi-infinite optimization with implicit functions. *Ind Eng Chem Res*. 2015;54(1):307-317.
28. Zuhe S, Neumaier A, Eiermann M. Solving minimax problems by interval methods. *BIT Numer Math*. 1990;30(4):742-751.
29. Stein O, Still G. Solving semi-infinite optimization problems with interior point techniques. *SIAM J Control Optim*. 2003;42(3):769-788.
30. Schweidtmann AM, Clayton AD, Holmes N, Bradford E, Bourne RA, Lapkin AA. Machine learning meets continuous flow chemistry: automated optimization towards the Pareto front of multiple objectives. *Chem Eng J*. 2018;352:277-282.
31. Meliani M, Bartoli N, Lefebvre T, Bouhlel M, Martins J, Morlier J. *Multi-Fidelity Efficient Global Optimization: Methodology and Application to Airfoil Shape Design*. AIAA Aviation Forum; 2019:3236.
32. Liu Z, An W, Qu X, Liu X, Lyu H. Portfolio-based Bayesian optimization for airfoil design. *AIAA J*. 2021;59(6):1975-1989. doi:10.2514/1.J059812
33. Wu J, Toscano-Palmerin S, Frazier PI, Wilson AG. Practical multi-fidelity Bayesian optimization for Hyperparameter tuning. In Proceedings of the 35th Uncertainty in Artificial Intelligence Conference, Vol. 115 of Proceedings of Machine Learning Research PMLR; 2020: 788–798.
34. Paulson JA, Mesbah A. Data-driven scenario optimization for automated controller tuning with probabilistic performance guarantees. *IEEE Control Syst Lett*. 2020;5(4):1477-1482.
35. Sorourifar F, Choksi N, Paulson JA. Computationally efficient integrated design and predictive control of flexible energy systems using multi-fidelity simulation-based Bayesian optimization. *Optim Control*

Appl Methods. 2021. https://onlinelibrary.wiley.com/doi/full/10.1002/oca.2817?casa_token=wWcYuPcrEFAAAAAA%3ANh6oNT3ID2vLg63ohSrKf4xXKrlrZeKPwhW1_xtGRTfge-OkCjrsp0IKYNqd_wgZAIHjUDV5ZM1OsWbd

36. Lu Q, González LD, Kumar R, Zavala VM. Bayesian optimization with reference models: a case study in MPC for HVAC central plants. *Comput Chem Eng*. 2021;154:107491.
37. Rios LM, Sahinidis NV. Derivative-free optimization: a review of algorithms and comparison of software implementations. *J Glob Optim*. 2013;56(3):1247-1293.
38. Mukhopadhyay DM, Balitanas MO, Farkhod A, Jeon SH, Bhattacharyya D. Genetic algorithm: a tutorial review. *Int J Grid Distrib Comput*. 2009;2(3):25-32.
39. Eberhart R, Kennedy J. Particle swarm optimization. In Proceedings of the IEEE International Conference on Neural Networks, Vol. 4, Citeseer; 1995. p. 1942–1948.
40. Shahriari B, Swersky K, Wang Z, Adams RP, De Freitas N. Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE*. 2015;104(1):148-175.
41. Frazier PI. A Tutorial on Bayesian optimization. *arXiv*. 2018. https://arxiv.org/pdf/1807.02811.pdf
42. Balandat M, Karrer B, Jiang D, et al. BoTorch: a framework for efficient Monte-Carlo Bayesian optimization. *Adv Neural Inf Process Syst*. 2020;33:21524-21538.
43. Paulson JA, Lu C. COBALT: constrained Bayesian optimization of computationally expensive grey-box models exploiting derivative information. *Comput Chem Eng*. 2022;160:107700.
44. Rasmussen CE, Williams CKI. *Gaussian Processes for Machine Learning*. MIT Press; 2006.
45. Eason JP, Biegler LT. Advanced trust region optimization strategies for glass box/black box models. *AIChE J*. 2018;64(11):3934-3943.
46. Wilson ZT, Sahinidis NV. The ALAMO approach to machine learning. *Comput Chem Eng*. 2017;106:785-795.
47. Boukouvala F, Floudas CA. ARGONAUT: algorithms for global optimization of constrained grey-box computational problems. *Optim Lett*. 2017;11(5):895-913.
48. Kim SH, Boukouvala F. Surrogate-based optimization for mixed-integer nonlinear problems. *Comput Chem Eng*. 2020;140:106847.
49. Marzat J, Walter E, Piet-Lahanier H. A new expected-improvement algorithm for continuous minimax optimization. *J Glob Optim*. 2016;64(4):785-802.
50. Paulson JA, Shao K, Mesbah A. Probabilistically robust Bayesian optimization for data-driven design of arbitrary controllers with Gaussian process emulators. In Proceedings of the IEEE Conference on Decision and Control; 2021:3633–3639.
51. Bogunovic I, Scarlett J, Jegelka S, Cevher V. Adversarially robust optimization with Gaussian processes. *Adv Neural Inf Process Syst*. 2018;31. https://proceedings.neurips.cc/paper/2018/hash/60243f9b1ac2dba11ff8131c8f4431e0-Abstract.html
52. Paulson JA, Makrygiorgos G, Mesbah A. Adversarially robust Bayesian optimization for efficient auto-tuning of generic control structures under uncertainty. *AIChE J*. 2021;68(6):e17591.
53. Lu C, Paulson JA. No-regret Bayesian optimization with unknown equality and inequality constraints using exact penalty functions. In IFAC Symposium on Dynamics and Control of Process Systems, Including Biosystems; 2022.
54. ur Rehman S, Langelaar M. Expected improvement-based infill sampling for global robust optimization of constrained problems. *Optim Eng*. 2017;18(3):723-753.
55. Di Pillo G, Grippo L. Exact penalty functions in constrained optimization. *SIAM J Control Optim*. 1989;27(6):1333-1360.
56. Coope ID, Price CJ. Exact penalty function methods for nonlinear semi-infinite programming. In: Reemtse R, ed. *Semi-Infinite Programming*. Springer; 1998:137-157.
57. Chaloner K, Verdinelli I. *Bayesian Experimental Design: A Review*. Statistical Science; 1995:273-304.
58. Srinivas N, Krause A, Kakade SM, Seeger M. Gaussian process optimization in the bandit setting: no regret and experimental design. In International Conference on Machine Learning; 2015:2171–2180.
59. Martins JRRA, Poon NMK. On structural optimization using constraint aggregation. In Sixth World Congress on Structural and Multidisciplinary Optimization, Rio de Janeiro, Brasil Citeseer; 2005.
60. Chen Y, Chen M. Extended duality for nonlinear programming. *Comput Optim Appl*. 2010;47(1):33-59.
61. Conn AR, Gould NIM. An exact penalty function for semi-infinite programming. *Math Program*. 1987;37(1):19-40.
62. Price CJ, Coope ID. Numerical experiments in semi-infinite programming. *Comput Optim Appl*. 1996;6(2):169-189.
63. Biegler LT, Zavala VM. Large-scale nonlinear programming using IPOPT: An integrating framework for enterprise-wide dynamic optimization. *Comput Chem Eng*. 2009;33(3):575-582.
64. Jones DR, Perttunen CD, Stuckman BE. Lipschitzian optimization without the Lipschitz constant. *J Optim Theory Appl*. 1993;79(1):157-181.
65. Bull AD. Convergence rates of efficient global optimization algorithms. *J Mach Learn Res*. 2011;12(88):2879-2904.
66. Still G. Generalized semi-infinite programming: theory and methods. *Eur J Oper Res*. 1999;119(2):301-313.
67. Marler RT, Arora JS. Survey of multi-objective optimization methods for engineering. *Struct Multidiscipl Optim*. 2004;26(6):369-395.
68. Wilson JT, Moriconi R, Hutter F, Deisenroth MP. The reparameterization trick for acquisition functions. *arXiv*. 2017. https://arxiv.org/pdf/1712.00424.pdf
69. Marcellin E, Behrendorff JB, Nagaraju S, et al. Low carbon fuels and commodity chemicals from waste gases–systematic approach to understand energy metabolism in a model acetogen. *Green Chem*. 2016;18(10):3020-3028.
70. Heijstra BD, Kern E, Koepke M, Segovia S, Liew FM, Novel bacteria and methods of use thereof. Google patents. US patent 13/812029. 2013.
71. Chen J, Daniell J, Griffin D, Li X, Henson MA. Experimental testing of a spatiotemporal metabolic model for carbon monoxide fermentation with clostridium autoethanogenum. *Biochem Eng J*. 2018;129:64-73. doi:10.1016/j.bej.2017.10.018

How to cite this article: Kudva A, Sorourifar F, Paulson JA. Constrained robust Bayesian optimization of expensive noisy black-box functions with guaranteed regret bounds. *AIChE J*. 2022;68(12):e17857. doi:10.1002/aic.17857