

Trust Region Constrained Bayesian Optimization with Penalized Constraint Handling

Raju Chowdhury^a, Tanmay Sen^b, Prajamitra Bhuyan^c, Biswabrata Pradhan^d

^{abd} Statistical Quality Control and Operations Research Unit,
Indian Statistical Institute, Kolkata, India.

^c Operations Management Group,
Indian Institute of Management Calcutta, Kolkata, India.

Abstract

Constrained optimization in high-dimensional black-box settings is difficult due to expensive evaluations, the lack of gradient information, and complex feasibility regions. In this work, we propose a Bayesian optimization method that combines a penalty formulation, a surrogate model, and a trust region strategy. The constrained problem is converted to an unconstrained form by penalizing constraint violations, which provides a unified modeling framework. A trust region restricts the search to a local region around the current best solution, which improves stability and efficiency in high dimensions. Within this region, we use the Expected Improvement acquisition function to select evaluation points by balancing improvement and uncertainty.

The proposed Trust Region method integrates penalty-based constraint handling with local surrogate modeling. This combination enables efficient exploration of feasible regions while maintaining sample efficiency. We compare the proposed method with state-of-the-art methods on synthetic and real-world high-dimensional constrained optimization problems. The results show that the method identifies high-quality feasible solutions with fewer evaluations and maintains stable performance across different settings.

Keywords: black-box function, Expected Improvement, Gaussian Process, Penalty function, Hybrid optimization.

1 Introduction

In many real-world settings, optimization problems arise without an explicit mathematical form. These problems are treated as black-box optimization problems, where the objective and constraint functions can only be evaluated at selected input points. Such problems appear in engineering design, model tuning, and scientific simulations. The difficulty becomes more pronounced in high-dimensional spaces due to the curse of dimensionality (Powell, 2019). As the number of variables grows, the search space becomes large and harder to explore with limited evaluations. It becomes difficult to efficiently identify promising regions. The presence of constraints further reduces the feasible region, which is often small and irregular, making the search for feasible and optimal solutions more challenging.

Bayesian optimization (BO) (Shahriari et al., 2016; Gramacy, 2020; Garnett, 2023) provides a principled approach for solving black-box optimization problems when function evaluations are expensive. It builds a probabilistic surrogate model of the unknown objective, often using Gaussian Processes, which gives both a prediction of the function value and a measure of uncertainty in unexplored regions. An acquisition function then selects new evaluation points by balancing improvement in the objective with uncertainty in the predictions. The method proceeds sequentially, with each new observation updating the surrogate model and improving future decisions. This approach is sample efficient and is well-suited for high-dimensional, and constrained settings where evaluations are costly, and the feasible region may be limited.

However, in high-dimensional settings (Greenhill et al., 2020; Papenmeier et al., 2026), standard BO often performs poorly because surrogate models become less accurate over the full domain, and optimizing the acquisition function becomes difficult. This leads to inefficient exploration and unstable search behavior. As the dimension increases, the search space grows rapidly, making it harder to identify promising regions with limited evaluations. The presence of constraints further increases the difficulty (Amini et al., 2025), as the feasible region may be small, irregular, and hard to locate. Together, these challenges make it difficult to maintain reliable performance and efficiently identify feasible and optimal solutions.

The classical optimization literature has long studied methods for handling constraints, including penalty functions, augmented Lagrangian methods, and barrier functions (No-

cedal and Wright, 2006). These approaches transform constrained problems into unconstrained ones or decompose them into simpler subproblems that can be solved iteratively. Building on these ideas, Gramacy et al. (2016) combined classical optimization with BO and proposed an Augmented Lagrangian-based BO framework, where the acquisition function is defined on the predictive mean surface. Inspired by this approach, Ariaifar et al. (2019) introduced a framework, ADMMBO, that uses the big- M penalty and the Alternating Directions Method of Multipliers together with the Expected Improvement acquisition function. In a similar direction, Pourmohamad and Lee (2022) developed a log-barrier based BO framework and constructed the acquisition function using the predictive mean. Recently, Upadhye and Chowdhury (2024) proposed a big- M penalty-based approach, known as M -LCBBO, in which a constrained version of the Lower Confidence Bound acquisition function is derived.

High-dimensional constrained optimization is challenging due to the large search space and difficulty in identifying feasible regions. Global surrogate models often become unreliable, which motivates local modeling strategies. Trust region based methods restrict the search to local regions and adapt their size based on progress. TuRBO (Eriksson et al., 2019) demonstrates improved performance in high dimensions, while SCBO (Eriksson and Poloczek, 2021) extends this idea to constrained settings, defines a hyperrectangular trust region. FuRBO (Ascia et al., 2025) further modifies the trust region using a hyperspherical domain to improve exploration. These approaches demonstrate that local trust-region strategies are effective for high-dimensional constrained optimization.

We propose a trust region based BO framework, referred to as TR-MEI, to address high-dimensional constrained optimization problem. The framework integrates big- M penalty approach with surrogate modeling and a trust region strategy. In this setting, we use a modified Expected Improvement acquisition function defined over the penalized objective, balancing exploration and exploitation within the local trust region.

The article is organized as follows. Section 2 presents the relevant background and existing trust region based BO methods. Section 3 describes the proposed methodology along with the trust region strategy. Section 4 presents the empirical comparison with existing methods on synthetic and real-world problems. Section 5 concludes the article with a discussion and directions for future research.

2 Background

In this Section, we briefly describe Bayesian Optimization and its main components, including surrogate modeling and Expected Improvement acquisition function for selecting evaluation points. We also outline a trust region based Bayesian optimization approach for high-dimensional settings.

2.1 Bayesian Optimization

Bayesian Optimization (BO) (Shahriari et al., 2016; Frazier, 2018) is an effective method for optimizing black box functions. It is a derivative-free and sequential approach that aims to find the optimum with a small number of evaluations. BO has two main components. The first is a surrogate model that approximates the original function. The second is an acquisition function that guides the selection of the next sample point. A widely used surrogate is the Gaussian Process (GP) (Rasmussen and Williams, 2006), defined by its mean and covariance or kernel function. Among acquisition functions, Expected Improvement (EI) (Garnett, 2023) is well structured, easy to compute, and widely used. Below, we give a brief mathematical description of GP surrogate modelling and EI.

Surrogate Modelling:

Let $\mathcal{D}_n = \{x_i, f(x_i)\}_{i=1}^n$ denote the observed input output pairs. We model these data using a GP $Y(\cdot)$ with mean function $m(\cdot)$ and covariance function $k(\cdot, \cdot)$. The model is trained on $m(\cdot)$ and $k(\cdot, \cdot)$, estimated using $\mathcal{D}_n$. Next, for a new input x , the trained model $Y(\cdot)$ gives a posterior prediction that follows a normal distribution with mean $\mu(x)$ and variance $\sigma^2(x)$ (Rasmussen and Williams, 2006), that is

$$\begin{aligned} Y(x) \mid \mathcal{D}_n &\sim \mathcal{N}(\mu(x), \sigma^2(x)) \\ \mu(x) &= k(x, X) K^{-1} \mathbf{y} \\ \sigma^2(x) &= k(x, x) - k(x, X) K^{-1} k(X, x) \end{aligned}$$

where $X = [x_1, \dots, x_n]$, $\mathbf{y} = [f(x_1), \dots, f(x_n)]^\top$, and K is the covariance matrix with entries $K_{ij} = k(x_i, x_j)$. Thus, these posterior predictive quantities, $\mu(x)$ and $\sigma^2(x)$, are used in the EI acquisition function.

Expected Improvement (EI):

The idea of EI is based on the improvement statistic defined as

$$\mathcal{I}(x) = \max(0, f^* - Y(x)),$$

which measures the improvement over the current best value $f^* = \min(f(x))$. Based on this, Jones et al. (1998) defines the EI acquisition function as the expectation of $\mathcal{I}(x)$. The expression of EI is

$$\text{EI}(x) = (f^* - \mu(x))\Phi\left(\frac{f^* - \mu(x)}{\sigma(x)}\right) + \sigma(x)\phi\left(\frac{f^* - \mu(x)}{\sigma(x)}\right), \quad (1)$$

where $\mu(x)$ and $\sigma(x)$ are the posterior predictive mean and standard deviation of $Y(x)$ at some point x , and $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cumulative distribution function and probability density function of the standard normal distribution. EI balances the trade-off between exploration and exploitation during the search process. The first term emphasizes exploitation by favoring points with high predicted values. The second term emphasizes exploration by favoring points with high uncertainty. The next evaluation point is obtained by maximizing EI, that is, $x^* = \operatorname{argmax} \text{EI}(x)$, which is then used for the next function evaluation and model update.

2.2 Trust Region CBO Approaches

Solving a constrained optimization problem in high-dimensional space using the BO algorithm is challenging due to the large search space and the influence of constraints. To handle high dimensionality, the trust region method restricts the search to a local region around the current best solution. This region adapts during the optimization process; it expands when improvements are observed and shrinks when progress slows. This local search improves efficiency and stability in high-dimensional settings. Here, we discuss two recently proposed trust-region based BO approaches.

2.2.1 Scalable Constrained Bayesian Optimization (SCBO)

A framework based on the trust-region algorithm (Yuan, 2015), Eriksson et al. (2019) proposed the trust region BO (TuRBO) algorithm for solving the unconstrained black-box optimization problem in a high-dimensional setting. Later, Eriksson and Poloczek (2021)

extends TuRBO to Scalable Constrained Bayesian Optimization (SCBO) to handle black-box high-dimensional optimization problems with constraints. SCBO initialized a trust region, that is, a hyperrectangle, around the best feasible solution found so far. If no feasible point is available, it is centered at the point with the smallest constraint violation.

At each iteration, SCBO generates a candidate point within the current trust region using Thompson sampling (TS) (Thompson, 1933) for the next function evaluation. TS balance improvement from the current best value and the feasibility of constraints. Also, the size of the trust region is adjusted during the optimization process. It expands when successful improvements are observed and shrinks when the search fails to make progress. This adaptive strategy helps maintain efficient, stable performance in high-dimensional, constrained settings.

2.2.2 Feasibility-Driven Trust Region Bayesian Optimization (FuRBO)

The FuRBO framework (Ascia et al., 2025) builds on the TuRBO algorithm (Eriksson et al., 2019), but differs in the way the trust region is defined. Similar to SCBO, the trust region is centered at the best feasible point, or at the point with the smallest constraint violation when no feasible solution is available. However, FuRBO defines the trust region as a hypersphere instead of a hyperrectangle.

This spherical region provides an isotropic search space around the center, which avoids bias toward coordinate directions. Candidate points are sampled within this hypersphere, which leads to a more uniform exploration of the local region. The size of the trust region is adapted during the optimization process. It expands when improvements are observed and shrinks when progress slows down. This design improves the robustness of the search in high-dimensional constrained settings.

3 Methodology

This Section discusses the methodology that combines the big- M penalty approach for constraint handling with the trust region to address the high dimensionality of the optimization problem.

3.1 Formulation

The general constrained optimization problem is defined as

$$\begin{aligned} & \min_{\mathbf{x} \in \mathcal{B}} f(\mathbf{x}) \\ & \text{subject to, } g_j(\mathbf{x}) \leq 0, \quad j = 1, 2, \dots, J. \end{aligned} \tag{2}$$

where $\mathcal{B}$ is a high-dimensional search space, typically a bounded domain in $\mathbb{R}^D$. In this formulation, $f(\mathbf{x})$ denotes the objective function, and $g_j(\mathbf{x})$ represents the set of constraint functions.

All functions in (2) are treated as black-box. Their analytical forms are unknown, and gradient information is not available. Only point-wise evaluations of $f(\mathbf{x})$ and $g_j(\mathbf{x})$ can be obtained, which are often expensive to compute. The goal is to identify a feasible solution that satisfies all constraints while minimizing the objective function using a limited number of evaluations.

In recent years, converting a constrained problem into an unconstrained one has become common in the BO literature (Gramacy et al., 2016; Ariafar et al., 2019; Upadhye and Chowdhury, 2024; Zhao and Xu, 2024; Chowdhury and Upadhye, 2025). These approaches rely on penalty-based methods from classical optimization (Nocedal and Wright, 2006), where constraint violations are incorporated into the objective function.

In this article, we adopt the big- M penalty approach, which imposes a penalty when constraints are violated. The constrained problem is reformulated as the following unconstrained problem (Ariafar et al., 2019; Upadhye and Chowdhury, 2024; Chowdhury and Upadhye, 2025),

$$\min_{\mathbf{x} \in \mathcal{B}} F(\mathbf{x}) = f(\mathbf{x}) + M \sum_{j=1}^J \mathbb{I}(g_j(\mathbf{x}) > 0), \tag{3}$$

where $\mathbb{I}(\cdot)$ is the indicator function that returns 1 if the constraint is violated and 0 otherwise, and M is a large positive constant that penalizes infeasible solutions.

This formulation assigns an additional cost to any point that violates one or more constraints. As a result, feasible solutions are always preferred over infeasible ones when M is sufficiently large. The transformed objective $F(\mathbf{x})$ can then be optimized using standard BO methods without explicitly handling constraints.

Now, we model the unconstrained objective in (3) using independent GPs for the objective and each constraint. Let $Y_f(\mathbf{x})$ denote the GP model for the objective and $Y_{g_j}(\mathbf{x})$

denote the GP model for the j -th constraint. The resulting surrogate for the penalized objective is given by

$$Y_F(\mathbf{x}) = Y_f(\mathbf{x}) + M \sum_{j=1}^J \mathbb{I}(Y_{g_j}(\mathbf{x}) > 0). \quad (4)$$

This formulation combines the predictive distributions of the objective and constraints. Since each component is modeled independently, we can derive the predictive mean and variance of $Y_F(\mathbf{x})$ at a new point x as

$$\mathbb{E}[Y_F(x)] = \mu_F(x) = \mu_f(x) + M \sum_{j=1}^J p(Y_{g_j}(x) > 0), \quad (5)$$

$$\mathbb{V}[Y_F(x)] = \sigma_F^2(x) = \sigma_f^2(x) + M^2 \sum_{j=1}^J p(Y_{g_j}(x) > 0) \left(1 - p(Y_{g_j}(x) > 0)\right), \quad (6)$$

where $p(Y_{g_j}(x) > 0) = 1 - \Phi\left(-\frac{\mu_{g_j}(x)}{\sigma_{g_j}(x)}\right)$ denotes the probability of constraint violation.

The mean $\mu_F(x)$ increases with the probability of infeasibility, which penalizes regions that are likely to violate constraints. The variance $\sigma_F^2(x)$ reflects both the uncertainty in the objective and the uncertainty in constraint satisfaction.

Using these quantities, we define the penalized EI acquisition function (Chowdhury and Upadhye, 2025) as

$$MEI(x) = (F^* - \mu_F(x)) \Phi\left(\frac{F^* - \mu_F(x)}{\sigma_F(x)}\right) + \sigma_F(x) \phi\left(\frac{F^* - \mu_F(x)}{\sigma_F(x)}\right), \quad (7)$$

where $F^* = \min(F(x))$. This acquisition function guides the search toward regions with low objective values while accounting for the likelihood of constraint satisfaction.

Next, we integrate the MEI acquisition function with a trust region strategy to obtain Trust Region MEI (TR- MEI). Instead of optimizing MEI over the entire search space $\mathcal{B}$, the optimization is restricted to a local trust region centered at the current best solution.

3.2 Trust-Region MEI

In a high-dimensional constrained optimization setting, the trust region approach is effective for locating feasible and optimal solutions. In the literature, SCBO and FuRBO are two trust region based methods that differ in how the next sample point is selected during the BO procedure. As discussed in Section 2.2, SCBO defines the trust region over a hyperrectangle domain, while FuRBO defines it over a hypersphere domain.

In this article, we embrace the hyperrectangle domain setting and use coordinate-wise perturbation to generate candidate points within the trust region. We also employ the EI acquisition function within this local region.

At each iteration, a trust region is defined around the best observed point. Candidate points are generated within this region and evaluated using the *MEI* criterion. The size of the trust region is updated during the optimization process. It expands when improvements are observed and shrinks when no progress is made. This adaptive mechanism focuses the search in promising regions while maintaining stability.

By combining the big- M penalty-based formulation with a trust region strategy, TR-*MEI* improves efficiency in high-dimensional constrained problems. It reduces the effect of the large search space and guides the optimization toward feasible and optimal regions.

Further, we compare the proposed TR-*MEI* method with SCBO and FuRBO to highlight the differences:

- TR-*MEI* solves an unconstrained version of the constrained problem, while SCBO and FuRBO handle the constrained problem explicitly.
- TR-*MEI* and SCBO define the hyperrectangular trust region, while FuRBO uses hyperspherical.
- TR-*MEI* uses the standard EI acquisition function, while SCBO and FuRBO rely on Thompson sampling strategies.

The advantages of the proposed methodology are as follows. First, the TR-*MEI* acquisition function admits an explicit mathematical expression. This allows direct computation and avoids the need for approximation or sampling based estimates. As a result, the method remains simple to implement and computationally efficient.

Second, the method can start from infeasible initial points. The big- M penalty formulation assigns higher objective values to constraint violations, which guides the search toward feasible regions even when no feasible solution is available at the beginning. This makes the approach robust in practical scenarios where feasible initialization is difficult.

Finally, the integration of the trust region strategy enables the method to handle high-dimensional constrained optimization problems effectively. By restricting the search to a local region and adapting its size based on observed progress, the method improves

stability and reduces the effect of the large search space while maintaining convergence toward feasible and optimal solutions.

4 Experiments

This section presents the comparison results of the proposed algorithm on a range of synthetic and real-world high-dimensional constrained optimization problems. The goal is to evaluate the method’s performance in terms of solution quality, constraint satisfaction, and computational efficiency.

Algorithm 1 Trust Region M -EI

```

1: Evaluate initial design  $X$  and train surrogate models  $(f, g_j)$ 
2:  $x_{\text{best}} \leftarrow \arg \min_{x \in \mathcal{B}} S(X; f, g_j)$ 
3: while Optimization budget not exhausted do
4:    $TR \leftarrow \text{define\_TR}$ 
5:    $x_{\text{next}} \leftarrow \text{TR-MEI}((f, g_j), TR)$ 
6:    $f_{\text{next}} \leftarrow f(x_{\text{next}})$ 
7:    $g_j^{\text{next}} \leftarrow g_j(x_{\text{next}})$ 
8:    $S \leftarrow S \cup \{(x_{\text{next}}, f_{\text{next}}, g_j^{\text{next}})\}$ 
9:   Fit surrogate models over  $S$ 
10:  Update  $n_s$  and  $n_f$ 
11:  if  $n_s = \tau_s$  or  $n_f = \tau_f$  then
12:     $L \leftarrow \text{adjust}(L)$ 
13:  end if
14:   $x_{\text{best}} \leftarrow \arg \min_{x \in \mathcal{X}} S(X; f, g_j)$ 
15: end while
16: return  $x_{\text{best}}$ 

```

4.1 Experimental setup

We evaluate the proposed TR-MEI algorithm against SCBO and FuRBO. The comparison focuses on their performance in high-dimensional constrained optimization problems, where both feasibility and sample efficiency are important.

For evaluation, we consider three representative test functions, namely Ackley, Levy, and Rastrigin. These functions capture different landscape characteristics such as separable, ill-conditioned, and multimodal behavior. Each function is defined in 20 dimensions with moderate constraint difficulty. The same experimental setting is used for all methods to ensure a fair comparison. Each algorithm begins with an initial design of $2D$ points and is allowed a total evaluation budget of 50 function evaluations. All experiments are repeated multiple times with different initial designs and random seeds to account for variability in performance.

For a fair comparison across all frameworks, the number of initial random points is kept the same. We set $n = 2 \times d$, where d is the dimension of the problem. We use a zero mean function and a Matérn 5/2 kernel for fitting the GP surrogate model. The results are averaged over 30 runs and reported with one standard error to reflect variability.

All implementations of TR-MEI, FuRBO, and SCBO are carried out in Python using the GPyTorch (Gardner et al., 2018) and BoTorch (Balandat et al., 2020) libraries. A common implementation framework is used for all methods to ensure consistency in modeling and evaluation. This setup allows a fair comparison across methods while maintaining computational efficiency. Further Algorithm 1 provides the flow-chart of the TR-MEI methodology.

Feasible solutions are always preferred over infeasible ones. Any infeasible point is assigned the worst objective value observed for the corresponding problem across all methods. This ensures that the evaluation properly reflects both optimality and constraint satisfaction.

4.2 Results

In this subsection, we compare TR-MEI, FuRBO, and SCBO on the Ackley, Levy, and Rastrigin synthetic problems in a 20 dimensional constrained setting. These benchmark functions are widely used to evaluate optimization algorithms due to their complex landscapes. The Ackley function is characterized by a multimodal surface with many local minima, while the Levy function presents a challenging structure with strong nonlinearity. The mathematical formulations of the problems are given in (A1), (A2), (A3).

In Figures 1, 2, and 3, we visualize the performance of each method across different




Figure 1: Ackley

benchmark problems. For the Ackley problem, the results show the ability of each method to handle a multimodal landscape and identify feasible solutions. In Figure 1, TR-MEI converges to a feasible optimal solution faster than FuRBO and SCBO, which continue to improve over iterations. Table 1 shows that the average performance of TR-MEI is better than the other methods, which indicates consistent performance.

Table 1: Best valid optimal solution with mean $\pm$ standard deviation.

Problems	TR-MEI	FuRBO	SCBO
Ackley	0.0085 ± 0.0102	4.4120 ± 0.3733	7.3662 ± 1.0541
Levy	0.0071 ± 0.0002	4.8199 ± 0.5856	5.4482 ± 0.3754
Rastrigin	47.0507 ± 1.1983	173.7190 ± 40.0435	292.2801 ± 36.344

For the L  vy problem, the comparison shows how the methods handle strong nonlinearity and complex constraint structures. Figure 2 shows that TR-MEI reaches a feasible




Figure 2: Levy

optimal solution around 15 function evaluations, while FuRBO and SCBO approach the optimum only after about 50 evaluations. Table 1 shows the consistency of the solution in terms of average performance.

For the Rastrigin problem, the results show the performance of each method in a setting with many regularly distributed local minima, which makes global optimization difficult. Figure 3 shows that TR-MEI achieves better performance compared to FuRBO and SCBO. This observation is also supported by Table 1.




Figure 3: Rastrigin

4.3 Application to portfolio optimization

To be written.

5 Discussion

An optimization problem in a high-dimensional search space under constraints is challenging for identifying a feasible global optimum. In this article, we propose a hybrid BO methodology that combines the classical big M penalty, GP surrogate modeling, and a trust region algorithm. The proposed TR-MEI method uses the EI acquisition function, which balances exploration and exploitation within a local trust region.

The integration of EI within the trust region enables the method to select informative points based on the predicted mean and uncertainty. As a result, the proposed TR-MEI approach provides an effective balance between feasibility and optimality while maintaining sample efficiency in high-dimensional constrained optimization problems.

Finally, the comparison of TR-MEI with SCBO and FuRBO on both synthetic and real-world problems demonstrates its efficiency and effectiveness. The results show that TR-MEI is able to identify high-quality feasible solutions with fewer function evaluations. It maintains stable performance across different problem settings and scales well with increasing dimensionality. These observations highlight the advantage of combining the penalty-based formulation with a trust region strategy in constrained BO literature.

Appendix A:

The mathematical descriptions of the synthetic constrained optimization problems are given below:

Constrained Ackley Problem

Let $x = (x_1, \dots, x_d) \in \mathbb{R}^d$. The Ackley function is defined as

$$f(x) = -a \exp \left(-b \sqrt{\frac{1}{d} \sum_{i=1}^d x_i^2} \right) - \exp \left(\frac{1}{d} \sum_{i=1}^d \cos(cx_i) \right) + a + e,$$

where the standard parameter values are

$$a = 20, \quad b = 0.2, \quad c = 2\pi.$$

The constrained optimization problem is

$$\begin{aligned}
& \min_{x \in \mathbb{R}^d} f(x) \\
& \text{s.t.} \quad \sum_{i=1}^d x_i \leq 0, \\
& \quad \quad \|x\|_2 \leq 5.
\end{aligned} \tag{A1}$$

Constrained L  vy Problem

Let $x = (x_1, \dots, x_d) \in \mathbb{R}^d$. Define

$$w_i = 1 + \frac{x_i - 1}{4}, \quad i = 1, \dots, d.$$

The optimization problem is

$$\begin{aligned}
& \min_{x \in \mathbb{R}^d} \sin^2(\pi w_1) + \sum_{i=1}^{d-1} (w_i - 1)^2 [1 + 10 \sin^2(\pi w_i + 1)] \\
& \quad \quad + (w_d - 1)^2 [1 + \sin^2(2\pi w_d)] \\
& \text{s.t.} \quad \sum_{i=1}^d x_i \leq 0, \\
& \quad \quad \|x\|_2 \leq 5.
\end{aligned} \tag{A2}$$

Constrained Rastrigin Problem

Let $x = (x_1, \dots, x_d) \in \mathbb{R}^d$. The Rastrigin function is defined as

$$f(x) = Ad + \sum_{i=1}^d [x_i^2 - A \cos(2\pi x_i)],$$

where the standard parameter value is

$$A = 10.$$

The constrained optimization problem is

$$\begin{aligned}
& \min_{x \in \mathbb{R}^d} f(x) \\
& \text{s.t.} \quad \sum_{i=1}^d x_i \leq 0, \\
& \quad \quad \|x\|_2 \leq 5.
\end{aligned} \tag{A3}$$

Declaration of Interest

The authors declare no conflicts of interest. All authors contributed equally.

References

- Amini, S., Vannieuwenhuyse, I., and Morales-Hernández, A. (2025). Constrained Bayesian Optimization: A Review. *IEEE Access*, 13:1581–1593.
- Ariaifar, S., Coll-Font, J., Brooks, D., and Dy, J. (2019). ADMMBO: Bayesian optimization with unknown constraints using ADMM. *Journal of Machine Learning Research*, 20(123):1–26.
- Ascia, P., Raponi, E., Bäck, T., and Duddeck, F. (2025). Feasibility-Driven Trust Region Bayesian Optimization. In Akoglu, L., Doerr, C., van Rijn, J. N., Garnett, R., and Gardner, J. R., editors, *Proceedings of the Fourth International Conference on Automated Machine Learning*, volume 293 of *Proceedings of Machine Learning Research*, pages 2/1–28. PMLR.
- Balandat, M., Karrer, B., Jiang, D. R., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. (2020). BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*.
- Chowdhury, R. and Upadhye, N. S. (2025). Addressing mixed constraints: an improved framework for black-box optimization. *Journal of Global Optimization*, 92:889–908.
- Eriksson, D., Pearce, M., Gardner, J., Turner, R. D., and Poloczek, M. (2019). Scalable Global Optimization via Local Bayesian Optimization. *Advances in neural information processing systems*, 32.
- Eriksson, D. and Poloczek, M. (2021). Scalable Constrained Bayesian Optimization. In *International conference on artificial intelligence and statistics*, pages 730–738. PMLR.
- Frazier, P. I. (2018). Bayesian Optimization. In *Recent advances in optimization and modeling of contemporary problems*, pages 255–278. Informs.
- Gardner, J. R., Pleiss, G., Bindel, D., Weinberger, K. Q., and Wilson, A. G. (2018). Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*.
- Garnett, R. (2023). *Bayesian Optimization*. Cambridge University Press.
- Gramacy, R. B. (2020). *Surrogates*. Chapman and Hall/CRC.
- Gramacy, R. B., Gray, G. A., Le Digabel, S., Lee, H. K., Ranjan, P., Wells, G., and Wild, S. M. (2016). Modeling an augmented Lagrangian for blackbox constrained optimization. *Technometrics*, 58(1):1–11.
- Greenhill, S., Rana, S., Gupta, S., Vellanki, P., and Venkatesh, S. (2020). Bayesian Optimization for Adaptive Experimental Design: A Review. *IEEE Access*, 8:13937–13948.

- Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer, 2 edition.
- Papenmeier, L., Poloczek, M., and Nardi, L. (2026). Understanding high-dimensional Bayesian optimization. In *Proceedings of the 42nd International Conference on Machine Learning*, ICML’25. JMLR.org.
- Pourmohamad, T. and Lee, H. K. (2022). Bayesian optimization via barrier functions. *Journal of Computational and Graphical Statistics*, 31(1):74–83.
- Powell, W. B. (2019). A unified framework for stochastic optimization. *European journal of operational research*, 275(3):795–821.
- Rasmussen, C. E. and Williams, C. K. (2006). *Gaussian Processes for Machine Learning*, volume 2. MIT press Cambridge, MA.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. (2016). Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1):148–175.
- Thompson, W. R. (1933). On the Likelihood That One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3/4):285–294.
- Upadhye, N. S. and Chowdhury, R. (2024). Constrained Bayesian optimization with lower confidence bound. *Technometrics*, 66(4):561–574.
- Yuan, Y.-X. (2015). Recent advances in trust region algorithms. *Math. Program.*, 151(1):249–281.
- Zhao, J. and Xu, J. (2024). Bayesian Optimization via Exact Penalty. *Technometrics*.