

# A Bayesian Optimization Framework for Finding Local Optima in Expensive Multimodal Functions

Yongsheng Mei<sup>a,\*</sup>, Tian Lan<sup>a</sup>, Mahdi Imani<sup>b</sup> and Suresh Subramaniam<sup>a</sup>

<sup>a</sup>The George Washington University

<sup>b</sup>Northeastern University

**Abstract.** Bayesian optimization (BO) is a popular global optimization scheme for sample-efficient optimization in domains with expensive function evaluations. The existing BO techniques are capable of finding a single global optimum solution. However, finding a set of global and local optimum solutions is crucial in a wide range of real-world problems, as implementing some of the optimal solutions might not be feasible due to various practical restrictions (e.g., resource limitation, physical constraints, etc.). In such domains, if multiple solutions are known, the implementation can be quickly switched to another solution, and the best possible system performance can still be obtained. This paper develops a multimodal BO framework to effectively find a set of local/global solutions for expensive-to-evaluate multimodal objective functions. We consider the standard BO setting with Gaussian process regression representing the objective function. We analytically derive the joint distribution of the objective function and its first-order derivatives. This joint distribution is used in the body of the BO acquisition functions to search for local optima during the optimization process. We introduce variants of the well-known BO acquisition functions to the multimodal setting and demonstrate the performance of the proposed framework in locating a set of local optimum solutions using multiple optimization problems.

## 1 Introduction

Bayesian optimization (BO) is a popular global optimization scheme for sample-efficient optimization in domains with expensive function evaluations [34, 40]. The BO iteratively builds a statistical model of the objective function according to all the past evaluations and sequentially selects the next evaluation by maximizing an acquisition function. BO has shown tremendous success in a wide range of domains with no analytical formulation of objective functions, including simulation optimizations [1], device tuning/calibration [7, 36], material/drug design [43, 14], and many more.

Despite several variants of BO in recent years [3, 11, 38], the focus of all methods has been on finding a single global optimum solution. However, many real-world problems can be considered multimodal optimization [8], where it is desired to find all or most global and local optimum solutions of multimodal functions. The rationale is that implementing some of the optimal solutions might not be feasible due to various practical restrictions (e.g., resource limitation, physical constraints, etc.); in such a scenario, if multiple solutions are

known, the implementation can be quickly switched to another solution and the best possible system performance can still be obtained.

Besides BO techniques, several deterministic and stochastic techniques have been developed for multi-model optimization. These include the gradient descent method, the quasi-Newton method [5, 22], and the Nelder-Mead's simplex method [4], which require the analytical form of the multimodal function and tend to be trapped into a local optimum. Evolutionary optimization is a class of techniques applicable to domains with no available analytical representations. Examples include variants of genetic algorithms [23, 24], clonal selection algorithms [9], and artificial immune networks [10]. However, evolutionary techniques' reliance on heuristics and excessive function evaluations prevent their reliable applications in domains with expensive-to-evaluate functions.

This paper develops a multimodal BO framework to effectively find a set of local/global solutions for multimodal objective functions. We consider the standard Gaussian process (GP) regression as the surrogate model for representing the objective function. Since characterizing local/global optima requires accessing both the values and first-order conditions of the objective function, we analytically derive the joint distribution of the objective function and its first-order gradients using the kernel function properties and derivatives, illustrated by Fig. 1. This joint distribution is used in the body of the BO acquisition functions to search for local optima during the optimization process. We introduce the variants of the well-known BO acquisition functions to the multimodal setting, including joint expected improvement and probability of improvement. The performance of the proposed framework in locating a set of local optimum solutions is demonstrated using multiple optimization problems, including evaluations using synthetic function, well-known multimodal benchmarks such as Griewank function and Shubert function, as well as through hyperparameter tuning problems for image classification on CIFAR-10 dataset [18]. Our proposed solution can effectively capture multiple local/global optima in these evaluations.

The main contributions of our work are as follows:

- We develop a new BO framework for finding local optima by analytically deriving the joint distribution of the objective function and its first-order gradients according to the kernel function of Gaussian process regression.
- We introduce new acquisition functions, such as joint expected improvement and probability of improvement, to search local optima during the optimization process of our framework.
- Experimental results on multimodal functions and real-world im-

\* Corresponding Author. Email: ysmei@gwu.edu.

age classification problem demonstrates the effectiveness of our framework in capturing multiple local and global optima.

## 2 Related Work

### 2.1 Bayesian Optimization and Acquisition Functions

Among many optimization frameworks [37, 32, 41], Bayesian optimization has emerged as a popular method for the sample-efficient optimization of expensive objective functions [35, 13]. It is capable of optimizing objective functions with no available closed-form expression, where only point-based (and possibly noisy) evaluation of the function is possible. The sampling-efficiency of BO schemes has made them suitable for domains with non-convex and costly function evaluations. Fundamentally, BO is a sequential model-based approach to solving the problem of finding a global optimum of an unknown objective function  $f(\cdot): \mathcal{X}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ , where  $\mathcal{X} \subset \mathbb{R}^d$  is a compact set. It performs a sequential search, and at each iteration  $k$ , selects a new location  $\mathbf{x}_{k+1}$  to evaluate  $f$  and observe its value, until making a final recommendation of the best estimate of the optimum  $\mathbf{x}^*$ .

The sequential selection is achieved through the acquisition function  $a: \mathcal{X} \rightarrow \mathbb{R}$ , defined over the posterior of GP model, where BO selects a sample in the search space with the highest acquisition value. Upon evaluating the objective function at the selected input, the surrogate GP model gets updated. Various BO policies have been developed depending on the acquisition function choices, including expected improvement (EI) [27], knowledge gradient (KG) [12], probability of improvement (PI) [19], upper-confidence bounds (UCB) [20], and entropy search (ES) [15]. Recent studies have also considered accounting for future improvements in solution quality [44, 17, 39] and BO methods for constrained problems [2, 30]. However, existing work often focuses on finding the global optimum rather than a set of local/global optima, which are also necessary for multimodal objective functions and will be explored in this paper.

### 2.2 Local Maxima in Gaussian Random Fields

A separate line of statistical applications concentrates on the tail distribution of the heights of local maxima in non-stationary Gaussian random fields, such as in peak detection problems. Such a tail distribution is defined as the probability that the height of the local maximum surpasses a given threshold at the point  $\mathbf{x}$ , conditioned on the case that the point is a local maximum of the Gaussian process model, given by the following equation:  $\Pr(f(\mathbf{x}) > \xi | f(\mathbf{x}) \text{ is one local maximum})$ , where  $\xi$  is the threshold. Existing work [6] has probed into this problem in the Gaussian process model. The general formulae are derived in [6] for non-stationary Gaussian fields and a subset of Euclidean space or Riemannian manifold of arbitrary dimension [21], which depends on local properties of the Gaussian process model rather than the global supremum of the field. Although the goal is to characterize certain local properties rather than finding a set of local/global optima, the contributions of the above work provide the key intuitions to our new BO framework design.

### 2.3 First-Order Derivative in Bayesian Optimization

The first-order derivative information has been exploited in existing works of BO, such as [25, 45]. Besides, [26] make gradient information available for hyperparameter tuning. Additionally, adjoint

methods provide gradients cheaply in the optimization of engineering systems [16, 31]. The gradients provide useful information about the objective function and can help the BO during the optimization process. For instance, [40] develop a derivative-enabled knowledge-gradient algorithm by incorporating derivative information into GP for BO, given by  $(f(\mathbf{x}_{k+1}), \nabla f(\mathbf{x}_{k+1}))^T | f(\mathbf{x}_{1:k}), \nabla f(\mathbf{x}_{1:k}) \sim \mathcal{N}((f(\mathbf{x}_{1:k}), \nabla f(\mathbf{x}_{1:k}))^T, \text{diag}(\sigma^2))$ , where  $\sigma^2$  is the variance. These methods assume the gradients of the objective function can be queried along with the objective function during the optimization process. GIBO[28] alternates between minimizing the variance of the estimate of the gradient and moving in the direction of the expected gradient. Later proposed MPD [29] extended and refined it by finding the maximum look-ahead gradient descent direction. However, we propose the BO framework to find local optima by computing and updating the joint distribution of the prediction with its first-order derivative regarding kernel functions. Besides, existing methods with gradients concentrate on using gradients as additional information to improve the traditional BO model targeting global optimum, while our algorithm aims to find as many local optima as possible. Despite the similarity, we do not directly access the objective's gradient.

## 3 Background

*Gaussian model:* Gaussian process (GP) model is the most commonly used model for standard BO and also adopted by our framework, providing the posterior distribution of the objective function according to all the past evaluations. In this part, we introduce several preliminaries of the GP model. Considering the objective function  $f(\cdot)$  and the GP model with  $k+1$  input samples  $\mathbf{x}$  of dimension  $n$ , the prior of the model is:

$$f(\mathbf{x}_{1:k+1}) \sim \mathcal{N}(\mu_{\mathbf{x}_{1:k+1}}, \Sigma_{\mathbf{x}_{1:k+1}, \mathbf{x}_{1:k+1}}),$$

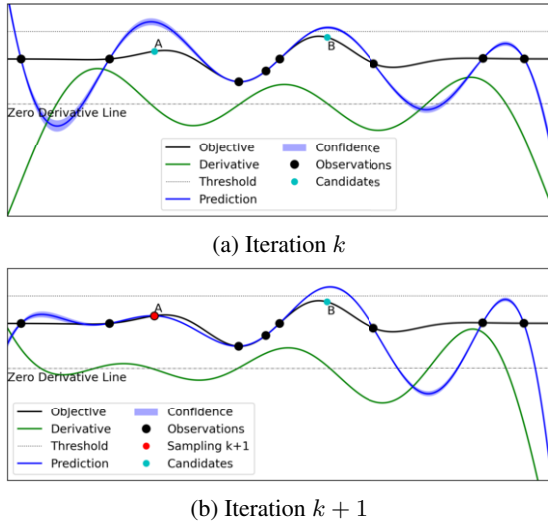
where we use  $\mu_{\mathbf{x}_{1:k+1}}$  to denote the mean of the prior and  $\Sigma_{\mathbf{x}_{1:k+1}, \mathbf{x}_{1:k+1}}$  to represent the initial covariance of  $k+1$  input samples. If we know the first  $k$  samples' values as observations  $f(\mathbf{x}_{1:k})$ , based on the prior, the posterior of the GP model representing the objective function at the next sampling point  $\mathbf{x}_{k+1}$  can be obtained as:

$$p_f \stackrel{\text{def}}{=} f(\mathbf{x}_{k+1}) | f(\mathbf{x}_{1:k}) \sim \mathcal{N}(\mu_{\mathbf{x}_{k+1}}, \Sigma_{\mathbf{x}_{k+1}, \mathbf{x}_{k+1}}), \quad (1)$$

where  $\mu_{\mathbf{x}_{k+1}}$  and  $\Sigma_{\mathbf{x}_{k+1}, \mathbf{x}_{k+1}}$  are the mean and variance respectively at this step. For simplicity purpose, we use  $p_f$  as the short notation of the posterior  $f(\mathbf{x}_{k+1}) | f(\mathbf{x}_{1:k})$ .

## 4 Finding Local Optima via Bayesian Optimization

We aim to develop a multimodal BO framework capable of efficiently computing optimal solutions. Our method will achieve optimization regarding a joint distribution containing the Gaussian posterior and its first-order derivative: the Gaussian posterior reflects the surrogate model as we observe the objectives in new places, and the first-order derivative of the posterior informs us about the latent location of local/global optima. Therefore, to this end, we compute the first-order derivative of the objective function to get the gradient and apply it to our new BO framework. Since the derivative operation is linear, we can combine the Gaussian posterior of the objective and its first-order derivative into a joint distribution and then compute the joint mean/variance for the acquisition functions. In this framework, we design new acquisition functions to fulfill the needs of finding local optima. Inspired by existing work [6], we first introduce a particular



**Figure 1:** Illustration of our proposed BO method for finding a set of local/global optima by considering both GP model prediction and its first-order derivatives.

threshold for the objective posterior to find candidates above/below the threshold value (depending on finding maxima/minima). Besides, as the optimum point's gradient is usually zero, we limit the posterior of the calculated derivative to a small interval around zero to further approach the optimum.

*An illustrative example.* Fig. 1 shows an illustrative example of our proposed method. In iteration  $k$ , we update the Gaussian posterior of the objective and its first-order derivative. Existing BO methods for finding a global maximum would select  $\mathbf{x}_{k+1}$  (and subsequent locations) to maximize the expected improvement or the probability of improvement, failing to recognize other local/global maxima in this multimodal function. For instance, probability of improvement can be described by the equation:  $\Pr(f(\mathbf{x}) > \xi)$ , where  $\xi$  is the given threshold.

To find other local/global optima, we derive the prediction's first-order derivative and let it satisfies:  $\Pr(f'(\mathbf{x}) < \epsilon)$ , where  $\epsilon$  denotes a small constant and hence indicates the place where the first-order condition approximates to zero.

On the other hand, if we only consider the first-order condition in choosing  $\mathbf{x}_{k+1}$ , the sequential search could easily get trapped in stationary/saddle points without improving the objective value toward local maxima. Therefore, our method leverages the joint distribution of objective function and its first-order derivative to find a set of local/global maxima by jointly computing and updating the probability:  $\Pr(f(\mathbf{x}) > \xi, f'(\mathbf{x}) < \epsilon)$ .

As shown in Fig. 1a, we have two latent local/global maxima, A and B, denoted by the cyan dots. Given the threshold  $\xi$  and first-order condition, the proposed algorithm will select local maximum A as the next sampling point rather than global maximum B at iteration  $k$ . Although our method can find the local maximum, it is not oblivious to other possible solutions, given by Fig. 1b, which suggests B as the latent sampling point for the next optimum solution. This illustrative example demonstrates our method's capability of finding all latent optimum solutions of multimodal functions by considering both GP model prediction and the first-order condition. Furthermore, analyzing such joint distribution would also be useful for deriving other special acquisition functions such as the expected improvement.

## 4.1 Multimodal Bayesian Optimization

We characterize a local/global optimum solution by leveraging the estimates of both the objective function and the first-order derivative of the function in the BO framework. We begin by considering the first-order derivative of the objective function  $f(\cdot)$ . Due to the linear property of the differentiation operation, we have the derivative of  $f(\cdot)$  also subject to the GP model [33]. Then, given input sample  $\mathbf{y}_{k+1}$  that can be a different point from  $\mathbf{x}_{k+1}$ , we derive another posterior:

$$p_{f'} \stackrel{\text{def}}{=} f'(\mathbf{y}_{k+1})|f(\mathbf{x}_{1:k}) \sim \mathcal{N}(\mu_{\mathbf{y}_{k+1}}, \Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}}), \quad (2)$$

where, for simplicity, we let  $p_{f'}$  be short for the posterior  $f'(\mathbf{y}_{k+1})|f(\mathbf{x}_{1:k})$ .

Since both  $\mathbf{x}_{k+1}$  and  $\mathbf{y}_{k+1}$  are related to the objective function  $f(\cdot)$ , posteriors  $f(\mathbf{x}_{k+1})|f(\mathbf{x}_{1:k})$  and  $f'(\mathbf{y}_{k+1})|f(\mathbf{x}_{1:k})$  will not be independent and identically distributed (i.i.d.). Besides, based on (1) and (2), and as both posteriors comply with the Gaussian distribution, we can further acquire a joint Gaussian distribution as follows:

$$p_{f, f'} \sim \mathcal{N} \left( \begin{bmatrix} \mu_{\mathbf{x}_{k+1}} \\ \mu_{\mathbf{y}_{k+1}} \end{bmatrix}, \begin{bmatrix} \Sigma_{\mathbf{x}_{k+1}, \mathbf{x}_{k+1}} & \Sigma_{\mathbf{x}_{k+1}, \mathbf{y}_{k+1}} \\ \Sigma_{\mathbf{y}_{k+1}, \mathbf{x}_{k+1}} & \Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}} \end{bmatrix} \right). \quad (3)$$

Before solving the mean and variance in joint distribution (3), we introduce the sufficient condition leveraging dominated convergence theorem in a lemma for interchanging derivative with expectation.

**Lemma 1** (Interchangeable condition). *Let  $Y \in \mathcal{Y}$  be a random variable.  $g: \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$  is a function such that  $g(t, Y)$  is integrable for all  $t$  and  $g$  is continuously differentiable with respect to  $t$ . Assume that there is a random variable  $Z$  such that  $|\frac{\partial}{\partial t} g(t, Y)| \leq Z$  almost surely for all  $t$  and  $\mathbb{E}(Z) < \infty$ . Then, we have:*

$$\frac{\partial}{\partial t} \mathbb{E}[g(t, Y)] = \mathbb{E} \left[ \frac{\partial}{\partial t} g(t, Y) \right].$$

*Proof.* See Appendix A. □

Next, we derive the mean and variance of the target joint distribution (3). Based on standard mean and variance results of Gaussian posterior model, we can obtain our mean of designed algorithm given by the following steps:

$$\begin{aligned} \mu_{k+1} &= \mathbb{E} \begin{bmatrix} f(\mathbf{x}_{k+1})|f(\mathbf{x}_{1:k}) \\ f'(\mathbf{y}_{k+1})|f(\mathbf{x}_{1:k}) \end{bmatrix} \\ &\stackrel{(a)}{=} \frac{d}{d\mathbf{y}_{k+1}} \mathbb{E} \begin{bmatrix} f(\mathbf{x}_{k+1})|f(\mathbf{x}_{1:k}) \\ f(\mathbf{y}_{k+1})|f(\mathbf{x}_{1:k}) \end{bmatrix} \\ &= \begin{bmatrix} k(\mathbf{x}_{k+1}, \mathbf{x}_{1:k}) \\ \frac{d}{d\mathbf{y}_{k+1}} k(\mathbf{y}_{k+1}, \mathbf{x}_{1:k}) \end{bmatrix} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} (f(\mathbf{x}_{1:k}) - \mu_0(\mathbf{x}_{1:k})), \end{aligned} \quad (4)$$

where (a) uses the interchangeable condition in Lemma 1, and  $\mu_0$  denotes the initial mean of the first  $k$  samples.

Furthermore, as the joint posterior is constructed by objective prediction and its first-order condition, the first covariance matrix  $\mathbf{K}_0$  is given by:

$$\mathbf{K}_0 = \begin{bmatrix} k(\mathbf{x}_{k+1}, \mathbf{x}_{k+1}) & \frac{d}{d\mathbf{y}_{k+1}} k(\mathbf{x}_{k+1}, \mathbf{y}_{k+1}) \\ \frac{d}{d\mathbf{y}_{k+1}} k(\mathbf{y}_{k+1}, \mathbf{x}_{k+1}) & \frac{d^2}{d\mathbf{y}_{k+1} d\mathbf{y}_{k+1}} k(\mathbf{y}_{k+1}, \mathbf{y}_{k+1}) \end{bmatrix}$$

Therefore, according to the variance of GP model, the variance of (3) is:

$$\sigma_{k+1}^2 = \mathbf{K}_0 - \begin{bmatrix} k(\mathbf{x}_{k+1}, \mathbf{x}_{1:k}) \\ \frac{d}{d\mathbf{y}_{k+1}} k(\mathbf{y}_{k+1}, \mathbf{x}_{1:k}) \end{bmatrix} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} \begin{bmatrix} k(\mathbf{x}_{k+1}, \mathbf{x}_{1:k}) \\ \frac{d}{d\mathbf{y}_{k+1}} k(\mathbf{y}_{k+1}, \mathbf{x}_{1:k}) \end{bmatrix}^T. \quad (5)$$

Therefore, we conclude the mean and variance of joint distribution (3) in a theorem by leveraging Lemma 1 as follows:

**Theorem 1** (Bayesian optimization for local optima). *Under the interchangeable condition, the mean of the joint Gaussian distribution posterior in (3) is:*

$$\begin{aligned} \mu_{k+1} &\stackrel{\text{def}}{=} \begin{bmatrix} \mu_{\mathbf{x}_{k+1}} \\ \mu_{\mathbf{y}_{k+1}} \end{bmatrix} \\ &= \mathbf{A} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} (f(\mathbf{x}_{1:k}) - \mu_0(\mathbf{x}_{1:k})), \end{aligned} \quad (6)$$

where  $\mu_0$  denotes the initial mean of the first  $k$  samples, and the auxiliary matrix  $\mathbf{A}$  is provided as follows:

$$\mathbf{A} = \begin{bmatrix} k(\mathbf{x}_{k+1}, \mathbf{x}_{1:k}) \\ \frac{d}{d\mathbf{y}_{k+1}} k(\mathbf{y}_{k+1}, \mathbf{x}_{1:k}) \end{bmatrix}.$$

Additionally, the solution of variance in the posterior is given by:

$$\begin{aligned} \sigma_{k+1}^2 &\stackrel{\text{def}}{=} \begin{bmatrix} \Sigma_{\mathbf{x}_{k+1}, \mathbf{x}_{k+1}} & \Sigma_{\mathbf{x}_{k+1}, \mathbf{y}_{k+1}} \\ \Sigma_{\mathbf{y}_{k+1}, \mathbf{x}_{k+1}} & \Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}} \end{bmatrix} \\ &= \mathbf{K}_0 - \mathbf{A} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} \mathbf{A}^T. \end{aligned} \quad (7)$$

In Theorem 1,  $k(\cdot)$  in (6) denotes the kernel function, such as the squared-exponential and Matern kernels, which define the influence of a solution on the performance and confidence estimations of untested nearby solutions. Apart from that, in variance (7),  $\mathbf{K}$  represents the covariance matrix, where  $[\mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$

Furthermore, we provide two examples using the squared-exponential kernel and polynomial kernel, respectively, to illustrate Theorem 1 as below.

**Example 1.1** (Square-exponential kernel). *As a popular kernel function widely used in many existing works, the squared-exponential kernel can strongly connect spatially adjacent sampling points as those points are more similar. The kernel is defined by the equation below:*

$$k(\mathbf{x}_i, \mathbf{x}_j) = \alpha \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|_2^2}{2l^2}\right), \quad (8)$$

where  $\alpha$  and  $l$  are scale factor and length scale parameter introduced by the kernel, respectively.

For simplification, we assume  $\mathbf{y}_{k+1} = \mathbf{x}_{k+1}$ . By leveraging the given square-exponential kernel (8) and computing its gradient, the mean in Theorem 1 becomes:

$$\mu_{k+1} = \bar{\mathbf{A}} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} f(\mathbf{x}_{1:k}), \quad (9)$$

where  $f(\mathbf{x}_{1:k})$  are known initial sample values and  $\bar{\mathbf{A}}$  is given by:

$$\bar{\mathbf{A}} = \alpha \begin{bmatrix} \exp\left(-\frac{\|\mathbf{x}_{k+1} - \mathbf{x}_{1:k}\|_2^2}{2l^2}\right) \\ -\frac{\mathbf{x}_{k+1} - \mathbf{x}_{1:k}}{l^2} \exp\left(-\frac{\|\mathbf{x}_{k+1} - \mathbf{x}_{1:k}\|_2^2}{2l^2}\right) \end{bmatrix}.$$

Likewise, we obtain the variance as:

$$\sigma_{k+1}^2 = \alpha \begin{bmatrix} 1 & \mathbf{0} \\ \mathbf{0} & \frac{1}{l^2} \mathbf{I}_n \end{bmatrix} - \bar{\mathbf{A}} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} \bar{\mathbf{A}}^T, \quad (10)$$

where  $\mathbf{I}_n$  represents the identity matrix of dimension  $n$ .

**Example 1.2** (Polynomial kernel). *In this example, we consider another frequently-used kernel, the Polynomial kernel, which represents the similarity of sampling points in the space over polynomials, allowing the learning of non-linear models. The kernel is defined as:*

$$k(\mathbf{x}_i, \mathbf{x}_j) = \bar{\alpha} (\mathbf{x}_i \cdot \mathbf{x}_j - c)^\delta, \quad (11)$$

where  $c \geq 0$  is a free parameter trading off the influence of higher-order versus lower-order terms in the polynomial,  $\bar{\alpha}$  is the scale factor, and  $\delta$  denotes the polynomial's dimension.

For simplification, let  $\mathbf{y}_{k+1} = \mathbf{x}_{k+1}$ . In the most typical situation, the kernel is a homogeneous polynomial kernel where  $c = 0$ , and the dimension  $\delta = 2$ . Then, we compute its gradient by leveraging the given polynomial kernel (11), and the mean in Theorem 1 becomes:

$$\mu_{k+1} = \bar{\mathbf{A}} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} f(\mathbf{x}_{1:k}), \quad (12)$$

where  $f(\mathbf{x}_{1:k})$  are known initial sample values and  $\bar{\mathbf{A}}$  is given by:

$$\bar{\mathbf{A}} = \bar{\alpha} \begin{bmatrix} (\mathbf{x}_{k+1} \cdot \mathbf{x}_{1:k})^2 \\ 2(\mathbf{x}_{k+1} \cdot \mathbf{x}_{1:k}) \cdot \mathbf{x}_{1:k} \end{bmatrix}.$$

Likewise, we obtain the variance as:

$$\begin{aligned} \sigma_{k+1}^2 &= \bar{\alpha} \begin{bmatrix} \|\mathbf{x}_{k+1}\|^4 & 2\|\mathbf{x}_{k+1}\|^2 \cdot \mathbf{x}_{k+1} \\ 2\|\mathbf{x}_{k+1}\|^2 \cdot \mathbf{x}_{k+1} & 2(\mathbf{x}_{k+1} \otimes \mathbf{x}_{k+1} + \|\mathbf{x}_{k+1}\|^2 \mathbf{I}_n) \end{bmatrix} \\ &\quad - \bar{\mathbf{A}} \mathbf{K}_{\mathbf{x}_{1:k}, \mathbf{x}_{1:k}}^{-1} \bar{\mathbf{A}}^T, \end{aligned} \quad (13)$$

where  $\otimes$  denotes the outer product and  $\mathbf{I}_n$  represents the identity matrix of dimension  $n$ .

## 4.2 Determining Local Optima

We manage to access the posterior of any  $\mathbf{x}_{k+1}$  via the designed GP model. The next step is determining the next test point to evaluate, which can be obtained through the acquisition function (AF). Due to the joint distribution in (3), to guarantee the discovery of a local optimum solution, we need 1) the value of the objective function at the test point to be larger than elsewhere nearby, and 2) the first-order derivative of the objective function at the same point to be close to zero. This step is another optimization problem regarding  $p_{f, f'}$ , but does not require evaluating the objective function  $f(\cdot)$ .

To start with, we use conditional distribution expansion to change the form of the joint probability distribution. The new joint posterior becomes:

$$p_{f, f'} = (p_f | p_{f'}) \cdot p_{f'}, \quad (14)$$

where  $p_f | p_{f'}$  obeys the Gaussian distribution with the mean of:

$$\bar{\mu}_{k+1} = \mu_{\mathbf{x}_{k+1}} + \Sigma_{\mathbf{x}_{k+1}, \mathbf{y}_{k+1}} \Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}}^{-1} (f'(\mathbf{y}_{k+1}) - \mu_{\mathbf{y}_{k+1}}), \quad (15)$$

and variance of:

$$\bar{\sigma}_{k+1}^2 = \Sigma_{\mathbf{x}_{k+1}, \mathbf{x}_{k+1}} - \Sigma_{\mathbf{x}_{k+1}, \mathbf{y}_{k+1}} \Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}}^{-1} \Sigma_{\mathbf{y}_{k+1}, \mathbf{x}_{k+1}}. \quad (16)$$

To satisfy the two aforementioned requirements, we let  $\mathbf{x}_{k+1}$  and  $\mathbf{y}_{k+1}$  be at the same position, where we have:

$$p_{f'} = f'(\mathbf{y}_{k+1}) | f(\mathbf{x}_{1:k}) = f'(\mathbf{x}_{k+1}) | f(\mathbf{x}_{1:k}).$$

According to (14), we further expand the joint distribution in (3) to fulfill the needs of the AF design. After that, referring to the original

AF definition, we newly define several AFs for determining the local optimum solution given in the following corollaries as examples.

Firstly, we provide the joint probability of improvement, which evaluates the objective function at the point most likely to improve upon the value of the current observations. Under such criterion, the point with the highest probability of improvement will be selected.

**Corollary 1** (Joint probability of improvement). *In our case, when  $\mathbf{x}_{k+1} = \mathbf{y}_{k+1}$ , the AF objective is:*

$$\begin{aligned} a_{PI}(\mathbf{x}_{k+1}) & \\ & \stackrel{(b)}{\approx} \int_{\xi}^{\infty} (p_f | p_{f'}) d\xi \cdot \int_{-\epsilon}^{\epsilon} p_{f'} d\epsilon \\ & = Q\left(\frac{\xi - \bar{\mu}_{k+1}}{\bar{\sigma}_{k+1}}\right) \left[ Q\left(\frac{-\epsilon - \mu_{\mathbf{y}_{k+1}}}{\Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}}}\right) - Q\left(\frac{\epsilon - \mu_{\mathbf{y}_{k+1}}}{\Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}}}\right) \right], \end{aligned} \quad (17)$$

where (b) uses expansion (14) and an approximation in calculating integral,  $\xi$  is the probability of improvement threshold,  $\epsilon$  is a small constant to restrict the first-order derivative, and  $Q(\cdot)$  denotes the Q-function with  $\bar{\mu}_{k+1}$  and  $\bar{\sigma}_{k+1}$  given in (15) and (16), respectively.

Note that we split the original integral function and let  $p_f | p_{f'} \geq \xi$  and  $|p_{f'}| \leq \epsilon$  to suit our purpose in determining local optimum solution. In (a) of (17), when calculating acquisition, due to small  $\epsilon$ , the standard double integral can be approximated by the product of two integrals. As  $p_f | p_{f'}$  and  $p_{f'}$  follow Gaussian distribution, using Q-function, which is the tail distribution function of the standard Gaussian distribution, simplifies the AF in the expression.

Next, we define a joint expected improvement that evaluates the objective function at the point that improves upon the value of the current observations in terms of expectation. Under this criterion, the point with the greatest expected improvement will be selected.

**Corollary 2** (Joint expected improvement). *In our setting, when  $\mathbf{x}_{k+1} = \mathbf{y}_{k+1}$ , the AF objective becomes:*

$$\begin{aligned} a_{EI}(\mathbf{x}_{k+1}) & \\ & \stackrel{(c)}{\approx} \int_{\xi}^{\infty} (f(\mathbf{x}_{k+1}) - \xi)(p_f | p_{f'}) d\xi \cdot \int_{-\epsilon}^{\epsilon} p_{f'} d\epsilon \\ & = \int_{\xi}^{\infty} (f(\mathbf{x}_{k+1}) - \xi)(p_f | p_{f'}) d\xi \\ & \cdot \left[ Q\left(\frac{-\epsilon - \mu_{\mathbf{y}_{k+1}}}{\Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}}}\right) - Q\left(\frac{\epsilon - \mu_{\mathbf{y}_{k+1}}}{\Sigma_{\mathbf{y}_{k+1}, \mathbf{y}_{k+1}}}\right) \right], \end{aligned} \quad (18)$$

where (c) adopts expansion (14) and an approximation in calculating integral. Other parameters are defined the same as those in Corollary 1.

Two AFs introduced in the corollaries will be used with our new BO framework for determining the local optima.

## 5 Experiment

In this section, we perform thorough experiments on the synthetic multimodal functions and real-world implementation to test our algorithm's effectiveness in finding local optima. Additionally, we report the ablations and scalability experiments of the designed algorithm on multimodal functions. The code has been made available at: [https://github.com/ysmei97/local\\_bo](https://github.com/ysmei97/local_bo).

### 5.1 Multimodal Function

We implement our framework on two common 2D functions (i.e., the Griewank function and the Shubert function) as benchmark objective functions for this experiment. The results are shown in Fig. 2 and Fig. 3. The red dots in the figures represent current sampling points, and the blue dots denote previously sampled points. To visualize the sampling order, we show the number of the current step above each dot. Each experiment is tested for 40 iterations.

#### 5.1.1 Benchmark 1: Griewank Function

The Griewank function has many widespread local minima, which are regularly distributed. We consider the case where input dimension  $m$  equals 2 and the range of input in each dimension satisfies  $x_i \in [-5, 5]$ . The function's graph shows four optima under such conditions in Figure 2.

We adopt the squared-exponential kernel with  $\alpha$  of 10 and  $l$  of 0.1, and select three random sampling points as priors, which are fixed in each test. Fig. 2b shows the final result of using joint PI with  $\xi$  of 1 and  $\epsilon$  of 0.1. We manage to reveal a total of two local optima: the first local optimum at step 4 (Fig. 2a) and the second one at step 40 (Fig. 2b). The acquisition values for potential sampling points that can be the local optima are higher, while for selected points and their neighbors are zero. Since all optima on the Griewank function are regularly distributed, the points on the grid have nearly similar acquisition values.

#### 5.1.2 Benchmark 2: Shubert Function

As another commonly-used function in optimization problem analysis, the Shubert function is a 2D multimodal function with multiple local and global optima. This function is often evaluated on the square  $x_1, x_2 \in [-10, 10]$ . To allow for easier viewing, we restrict the domain to  $x_1, x_2 \in [-2, 0]$ , presented in Figure 3.

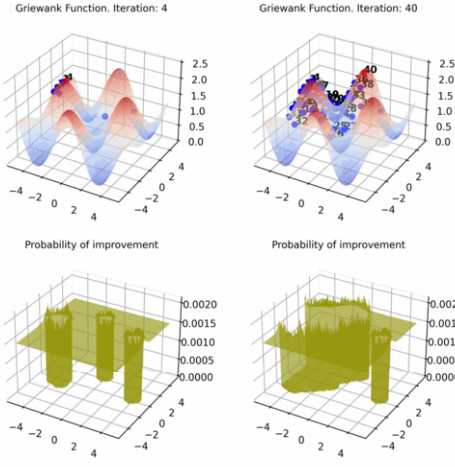
We use designed joint EI as AF with  $\xi$  of 0 and  $\epsilon$  of 0.1. As shown in Fig. 3, our framework can find the most local optima over time. We list several key steps where local optima are found in Fig. 3a, Fig. 3b, and Fig. 3c, which illustrates the effectiveness of our design.

### 5.2 Hyperparameter Tuning in Classification Task

Adopting the best hyperparameters, even the known ones, are not always feasible in many scenarios where we have additional hardware or runtime restrictions, especially for some expensive machine learning tasks. In these situations, we leverage the local optimal hyperparameters as a compromise. In this part, as a simulation of mentioned cases, we aim to verify our algorithm's capability for finding local optimum hyperparameter combinations on a real-world image classification problem.

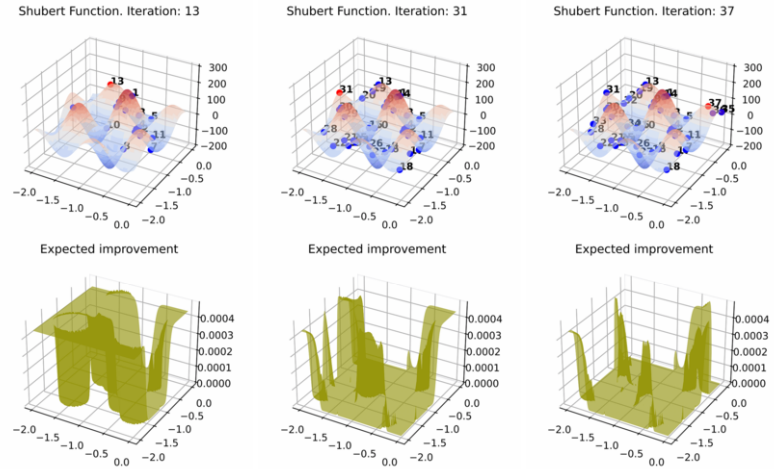
#### 5.2.1 Dataset

The dataset we use for the image classification task is CIFAR-10, a standard dataset consisting of 60000  $32 \times 32$  color images in 10 mutually exclusive classes with 50000 training images and 10000 test images. To fully utilize the dataset, we did necessary preprocessing, including reshaping and padding, with no loss of the information.



(a) Step 13: 1st optimum (b) Final result

Figure 2: Finding local optima on Griewank function via  $a_{PI}$ .



(a) Step 13: 1st optimum (b) Step 31: 2nd optimum (c) Step 37: 3rd optimum

Figure 3: Finding local optima on Shubert function via  $a_{EI}$ .

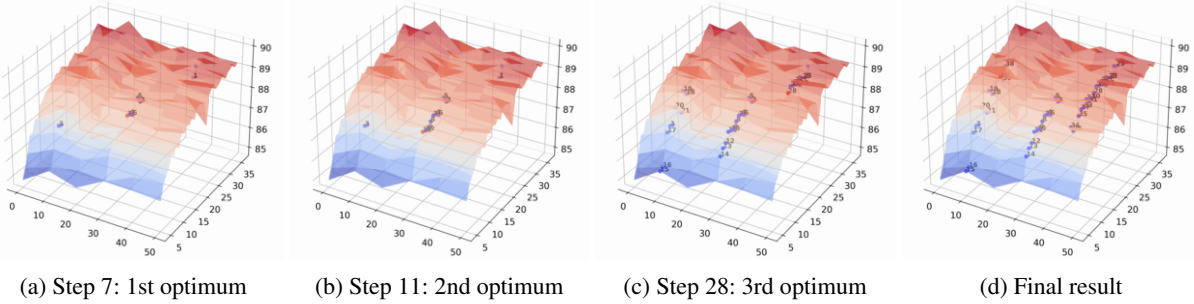


Figure 4: Searching for local optima with  $a_{EI}$  on CIFAR-10 image classification model of backbone ResNeXt.

### 5.2.2 Modeling and Experimentation

We build the model using ResNeXt [42] as the backbone. Among hyperparameters introduced by ResNeXt, we mainly focus on the *layer* of the network and the *width* of the subblock and leave the *cardinality* fixed. Based on this, we create an image classification model as the objective function, where the inputs are the layer and width and output is the validation accuracy. Since training online introduces latent nuance in searching local optima every step, we test our framework to determine local optimum hyperparameter combinations in an offline manner, i.e., we apply our algorithm to a trained model with all the samples ready.

In the experiment, we set cardinality to 4 and vary width from 5 to 35. Besides, we consider 6 layer candidates, which are 29, 38, 50, 68, 86, and 101. Using an interval of 10, we map the 6 layer candidates into a range from 0 to 50 as approximated representations. As shown in Fig. 4, the accuracy increases with width and layer projection values. However, such an increase is not monotonic, and we can spot many local optimum solutions. Since the mean of all results is around 87.0, we set the  $\xi$  as 87 and  $\epsilon$  as 0.01. We adopt squared-exponential kernel, with  $\alpha$  of 5 and  $l$  of 2.

Fig. 4 exemplifies some steps in which our method finds local optima. Compared to the width, the layer in this scenario impacts more on output accuracy. Therefore, the algorithm prioritizes to explore the local optima of different widths under a specific layer and then searches elsewhere for different layers. For instance, in Fig. 4b and

Fig. 4c, we spend multiple iterations searching for the optima with the layer of 68 before switching to explore that with the layer of 86.

### 5.3 Comparison with other baselines

In this section, we compare our method with several selected baselines algorithms: (1) MPD [29], which uses maximum look-ahead probability; and (2) GIBO [28], which performs local search by minimizing the trace of the posterior covariance of the gradient only; (3) vanilla BO with PI; (4) vanilla BO with EI. We use a 1D multimodal synthetic function of which the inputs range from 0 to 1 with 200 sampling instances to verify the methods' capability of finding the local and global maxima. Each run has a budget of 100 steps. We show the results in Table 1, where our framework can locate the objective local maxima faster than other baseline methods. Given a specific threshold, our framework will prioritize locating the nearby local maximum solutions. Besides, compared to the model with the joint probability of improvement, the one with joint expected improvement has better performance owing to its awareness of the potential improvement amount.

We also compute the average distance by adding the distance between new observation  $\mathbf{x}_i$  and ground truth  $\mathbf{x}^*$  of each step  $i$  and taking the average based on the number of already-taken steps  $n$ , given by  $\frac{\sum_i^n \|\mathbf{x}_i - \mathbf{x}^*\|}{n}$ . For every new observation, the distance is recognized as the nearest ground truth  $\mathbf{x}^*$  to  $\mathbf{x}_i$ . According to Table 1, our framework achieves lower regret than other baseline methods,

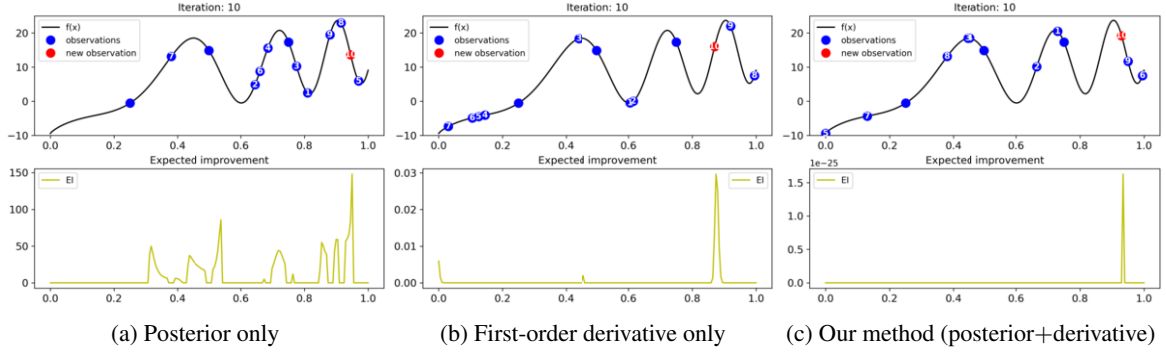


Figure 5: Ablations regarding the Gaussian posterior of the objective (standard BO) or its first-order derivative only.

Table 1: Comparison between proposed framework with other baselines in finding the local and global optimum solutions.

Baselines	1: Global maxima		2: 2nd large local maxima		3: 3rd large local maxima		Average Distance every 30 steps		
	Step	Average Distance	Step	Average Distance	Step	Average Distance	Step 30	Step 60	Step 90
Joint PI	78	0.062	36	0.040	12	0.034	0.043	0.068	0.059
Joint EI	80	0.061	22	0.055	6	0.022	0.055	0.065	0.057
MPD	20	0.092	37	0.109	47	0.089	0.128	0.080	0.081
GIBO	25	0.151	31	0.135	60	0.131	0.139	0.131	0.134
Vanilla PI	28	0.170	—	—	54	0.209	0.204	0.197	0.288
Vanilla EI	18	0.090	66	0.174	81	0.187	0.135	0.176	0.179

especially in deciding the local optimum solutions in experiments 2 and 3, demonstrating the effectiveness of finding local optima by adopting our approach.

#### 5.4 Scalability Experiment

To demonstrate our method’s scalability, we experiment with 3D Griewank functions for locating local/global optimal solutions. Due to the visualization limitation of high-dimensional functions, we post several steps of detected local maxima in Table 2, where we show that the distance to ground truth becomes smaller in the process for three maxima. We run the experiment for 300 steps in total, with a minimum sampling distance of 0.1. As seen from the table, our algorithm can be easily extended to cope with optimization problems in higher dimensions following the same procedure as processing lower-dimensional multimodal functions. Specifically, for the 3D Griewank function shown in Table 2, we successfully locate the solutions at steps 29, 169, and 245 with the distances to ground truth maxima as 0.024, 0, and 0.1, respectively.

#### 5.5 Ablation Experiment

In this experiment, we present the ablation results by removing the first-order derivative and posterior of the objective in AF (which is expected improvement) and run the test on a synthetic function. We present the results at the tenth iteration. As can be seen in Fig. 5a, once we remove the first-order derivative condition, the AF will be similar to the standard BO, and the algorithm will start looking for the points to maximize the acquisition value, failing to rapidly recognize other local/global maxima in this multimodal function. The two local maxima are not found during iterations, while our method successfully reveals them. In Fig. 5b, if we only consider the first-order derivative condition in choosing the next sampling point, the sequential search could easily get trapped in minima or stationary points without improving the objective value toward local maxima, such as points 1 and 2 or points 4, 5, and 6. In contrast, our algorithm

Table 2: Finding maxima on 3D Griewank function.

Step	Maxima	$x_1$	$x_2$	$x_3$	$f(x_1, x_2, x_3)$	Distance
27		-3.6	0.9	0.0	1.725	1.030
28		-3.0	-0.8	0.2	1.833	0.831
<b>29</b>	1	-3.0	-0.2	0.0	<b>1.982</b>	<b>0.224</b>
30		-3.4	-0.2	0.0	1.907	0.361
31		-2.9	0.2	-0.5	1.924	0.574
167		0.4	-4.2	0.0	1.912	0.447
168		-0.1	-4.4	0.2	1.993	0.224
<b>169</b>	2	0.0	-4.4	0.0	<b>2.004</b>	<b>0.000</b>
170		0.0	-4.3	-0.3	1.985	0.316
171		-0.3	-4.4	0.2	1.953	0.361
243		-0.5	4.2	-0.2	1.863	0.574
244		0.2	4.3	-0.1	1.978	0.245
<b>245</b>	3	0.0	4.4	-0.1	<b>2.003</b>	<b>0.100</b>
246		-0.3	4.4	0.5	1.934	0.583
247		-0.1	4.5	-0.3	1.984	0.332

can explicitly find the local maximum solutions by considering both conditions, given by Fig. 5c, where the maxima are points 1 and 4. The ablations indicate that our method can detect the global maximum and local maxima for given objective functions.

## 6 Conclusion

We propose a BO framework for finding a set of local/global optima in multimodal objective functions that are expensive to evaluate. Given a posterior modeled as a Gaussian Process, our method considers the joint distribution of the objective function and its first-order derivative. Evaluations on benchmark functions, image classification tasks, and hyperparameter tuning problems demonstrate the effectiveness of our solution in finding local/global optima.

## Acknowledgements

This research is based on work supported by the Army Research Office (ARO) under Grant W911NF2110299.

## References

- [1] Luigi Acerbi and Wei Ji Ma, 'Practical bayesian optimization for model fitting with bayesian adaptive direct search', *Advances in neural information processing systems*, **30**, (2017).
- [2] Setareh Ariafar, Jaume Coll-Font, Dana H Brooks, and Jennifer G Dy, 'Admbo: Bayesian optimization with unknown constraints using adm.', *J. Mach. Learn. Res.*, **20**(123), 1–26, (2019).
- [3] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy, 'Botorch: a framework for efficient monte-carlo bayesian optimization', *Advances in neural information processing systems*, **33**, 21524–21538, (2020).
- [4] Khurram Butt, Ramhuzaini A Rahman, Nariman Sepehri, and Shaahin Filizadeh, 'Globalized and bounded nelder-mead algorithm with deterministic restarts for tuning controller parameters: Method and application', *Optimal Control Applications and Methods*, **38**(6), 1042–1055, (2017).
- [5] Richard H Byrd, Samantha L Hansen, Jorge Nocedal, and Yoram Singer, 'A stochastic quasi-newton method for large-scale optimization', *SIAM Journal on Optimization*, **26**(2), 1008–1031, (2016).
- [6] Dan Cheng and Armin Schwartzman, 'Distribution of the height of local maxima of gaussian random fields', *Extremes*, **18**(2), 213–240, (2015).
- [7] Valentin Dalibard, Michael Schaarschmidt, and Eiko Yoneki, 'Boat: Building auto-tuners with structured bayesian optimization', in *Proceedings of the 26th International Conference on World Wide Web*, pp. 479–488, (2017).
- [8] Swagatam Das, Sayan Maity, Bo-Yang Qu, and Ponnuthurai Nagarathnam Suganthan, 'Real-parameter evolutionary multimodal optimization—a survey of the state-of-the-art', *Swarm and Evolutionary Computation*, **1**(2), 71–88, (2011).
- [9] L Nunes De Castro and Fernando J Von Zuben, 'The clonal selection algorithm with engineering applications', in *Proceedings of GECCO*, volume 2000, pp. 36–39, (2000).
- [10] Leandro Nunes De Castro and Jon Timmis, 'An artificial immune network for multimodal function optimization', in *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No. 02TH8600)*, volume 1, pp. 699–704. IEEE, (2002).
- [11] David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek, 'Scalable global optimization via local bayesian optimization', *Advances in neural information processing systems*, **32**, (2019).
- [12] Peter Frazier, Warren Powell, and Savas Dayanik, 'The knowledge-gradient policy for correlated normal beliefs', *INFORMS journal on Computing*, **21**(4), 599–613, (2009).
- [13] Peter I Frazier, 'A tutorial on bayesian optimization', *arXiv preprint arXiv:1807.02811*, (2018).
- [14] Ryan-Rhys Griffiths and José Miguel Hernández-Lobato, 'Constrained bayesian optimization for automatic chemical design using variational autoencoders', *Chemical science*, **11**(2), 577–586, (2020).
- [15] José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani, 'Predictive entropy search for efficient global optimization of black-box functions', *Advances in neural information processing systems*, **27**, (2014).
- [16] Antony Jameson, 'Re-engineering the design process through computation', *Journal of Aircraft*, **36**(1), 36–50, (1999).
- [17] Shali Jiang, Daniel Jiang, Maximilian Balandat, Brian Karrer, Jacob Gardner, and Roman Garnett, 'Efficient nonmyopic bayesian optimization via one-shot multi-step trees', *Advances in Neural Information Processing Systems*, **33**, 18039–18049, (2020).
- [18] Alex Krizhevsky, Geoffrey Hinton, et al., 'Learning multiple layers of features from tiny images'. Citeseer, (2009).
- [19] Harold J Kushner, 'A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise', (1964).
- [20] Tze Leung Lai, Herbert Robbins, et al., 'Asymptotically efficient adaptive allocation rules', *Advances in applied mathematics*, **6**(1), 4–22, (1985).
- [21] John M Lee, *Introduction to Riemannian manifolds*, volume 176, Springer, 2018.
- [22] Adrian S Lewis and Michael L Overton, 'Nonsmooth optimization via quasi-newton methods', *Mathematical Programming*, **141**(1), 135–163, (2013).
- [23] Jian-Ping Li, 'Truss topology optimization using an improved species-conserving genetic algorithm', *Engineering Optimization*, **47**(1), 107–128, (2015).
- [24] Yong Liang and Kwong-Sak Leung, 'Genetic algorithm with adaptive elitist-population strategies for multimodal function optimization', *Applied Soft Computing*, **11**(2), 2017–2034, (2011).
- [25] Daniel James Lizotte, Practical bayesian optimization, 2008.
- [26] Dougal Maclaurin, David Duvenaud, and Ryan Adams, 'Gradient-based hyperparameter optimization through reversible learning', in *International conference on machine learning*, pp. 2113–2122. PMLR, (2015).
- [27] Jonas Močkus, 'On bayesian methods for seeking the extremum', in *Optimization techniques IFIP technical conference*, pp. 400–404. Springer, (1975).
- [28] Sarah Müller, Alexander von Rohr, and Sebastian Trimpe, 'Local policy search with bayesian optimization', *Advances in Neural Information Processing Systems*, **34**, 20708–20720, (2021).
- [29] Quan Nguyen, Kaiwen Wu, Jacob Gardner, and Roman Garnett, 'Local bayesian optimization via maximizing probability of descent', *Advances in neural information processing systems*, **35**, 13190–13202, (2022).
- [30] Valerio Perrone, Iaroslav Shcherbatyi, Rodolphe Jenatton, Cedric Archambeau, and Matthias Seeger, 'Constrained bayesian optimization with max-value entropy search', *arXiv preprint arXiv:1910.07003*, (2019).
- [31] R-E Plessix, 'A review of the adjoint-state method for computing the gradient of a functional with geophysical applications', *Geophysical Journal International*, **167**(2), 495–503, (2006).
- [32] Parisa Rafiee, Mustafa Oktay, and Omur Ozel, 'Intermittent status updating with random update arrivals', in *2021 IEEE International Symposium on Information Theory (ISIT)*, pp. 3121–3126. IEEE, (2021).
- [33] Carl Edward Rasmussen and Christopher K. I. Williams, *Gaussian processes for machine learning*, Adaptive computation and machine learning, MIT Press, 2006.
- [34] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas, 'Taking the human out of the loop: A review of bayesian optimization', *Proceedings of the IEEE*, **104**(1), 148–175, (2015).
- [35] Jasper Snoek, Hugo Larochelle, and Ryan P Adams, 'Practical bayesian optimization of machine learning algorithms', *Advances in neural information processing systems*, **25**, (2012).
- [36] Rodrigo A Vargas Hernández, 'Bayesian optimization for calibrating and selecting hybrid-density functional models', *The Journal of Physical Chemistry A*, **124**(20), 4053–4061, (2020).
- [37] Su Wang, Yichen Ruan, Yuwei Tu, Satyavrat Wagle, Christopher G Brinton, and Carlee Joe-Wong, 'Network-aware optimization of distributed learning for fog computing', *IEEE/ACM Transactions on Networking*, **29**(5), 2019–2032, (2021).
- [38] Jia Wu, Xiu-Yun Chen, Hao Zhang, Li-Dong Xiong, Hang Lei, and Si-Hao Deng, 'Hyperparameter optimization for machine learning models based on bayesian optimization', *Journal of Electronic Science and Technology*, **17**(1), 26–40, (2019).
- [39] Jian Wu and Peter Frazier, 'Practical two-step lookahead bayesian optimization', *Advances in neural information processing systems*, **32**, (2019).
- [40] Jian Wu, Matthias Poloczek, Andrew G Wilson, and Peter Frazier, 'Bayesian optimization with gradients', *Advances in neural information processing systems*, **30**, (2017).
- [41] Xidong Wu, Zhengmian Hu, Jian Pei, and Heng Huang, 'Serverless federated auprc optimization for multi-party collaborative imbalanced data mining', in *SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, (2023).
- [42] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He, 'Aggregated residual transformations for deep neural networks', in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, (2017).
- [43] Yichi Zhang, Daniel W Apley, and Wei Chen, 'Bayesian optimization for materials design with mixed quantitative and qualitative variables', *Scientific reports*, **10**(1), 1–13, (2020).
- [44] Yunxiang Zhang, Xiangyu Zhang, and Peter I Frazier, 'Two-step lookahead bayesian optimization with inequality constraints', *arXiv preprint arXiv:2112.02833*, (2021).
- [45] A Žilinskas, Practical mathematical optimization: An introduction to basic optimization theory and classical and new gradient-based algorithms, 2006.