

# BOIDS: High-Dimensional Bayesian Optimization via Incumbent-Guided Direction Lines and Subspace Embeddings

Lam Ngo<sup>1</sup>, Huong Ha<sup>1</sup>, Jeffrey Chan<sup>1</sup>, Hongyu Zhang<sup>2</sup>

<sup>1</sup>RMIT University, Australia

<sup>2</sup>Chongqing University, China

s3962378@student.rmit.edu.au, huong.ha@rmit.edu.au, jeffrey.chan@rmit.edu.au, hyzhang@cqu.edu.cn

## Abstract

When it comes to expensive black-box optimization problems, Bayesian Optimization (BO) is a well-known and powerful solution. Many real-world applications involve a large number of dimensions, hence scaling BO to high dimension is of much interest. However, state-of-the-art high-dimensional BO methods still suffer from the curse of dimensionality, highlighting the need for further improvements. In this work, we introduce BOIDS, a novel high-dimensional BO algorithm that guides optimization by a sequence of one-dimensional direction lines using a novel tailored line-based optimization procedure. To improve the efficiency, we also propose an adaptive selection technique to identify most optimal lines for each round of line-based optimization. Additionally, we incorporate a subspace embedding technique for better scaling to high-dimensional spaces. We further provide theoretical analysis of our proposed method to analyze its convergence property. Our extensive experimental results show that BOIDS outperforms state-of-the-art baselines on various synthetic and real-world benchmark problems.

**Code** — <https://github.com/LamNgo1/boids>

**Extended version** — <https://arxiv.org/abs/2412.12918>

## 1 Introduction

Bayesian Optimization (BO) (Shahriari et al. 2015; Garnett 2023) is a powerful optimization method, capable of solving expensive black-box objective functions. BO has applications in many fields, notably hyperparameter tuning for machine learning models (Turner et al. 2021; Wang et al. 2024), neural architecture search (Ru et al. 2021), reinforcement learning (Parker-Holder et al. 2022), engineering (Shields et al. 2021; Ament et al. 2023), robotics (Mayr et al. 2022).

BO is an iterative method that can optimize expensive, black-box objective functions in a sample-efficient manner. In each iteration, BO constructs a surrogate model from observed data to approximate the objective function. To propose the next data points for observation, BO optimizes an acquisition function derived from the surrogate model and an optimization policy. The acquisition function quantifies the potential information gain from new observations, guiding the search towards solutions that maximize this gain.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Common acquisition functions are computationally expensive to optimize and may even have tractable solutions under certain conditions. As a result, BO problems are transformed into less costly optimization tasks.

As recent applications involve optimizing problems with hundreds to thousands of dimensions, scaling BO to such high dimensions has become a prevalent research topic (Binois and Wycoff 2022; Wang et al. 2023). However, BO’s performance significantly degrades as the dimension increases, partly due to the over-exploration of the acquisition function (Eriksson et al. 2019; Ngo et al. 2024). Classical approaches for optimizing BO acquisition functions, such as gradient solvers or evolutionary algorithms, operate over the entire search space, which can suffer from the curse of dimensionality. In high dimensions, the exponentially large search space makes optimizing the acquisition function computationally expensive (if the entire space is considered) or less accurate (if approximation techniques are used). Among existing solutions, LineBO (Kirschner et al. 2019) introduces line-based optimization, which restricts the acquisition function optimization to a one-dimensional line (*guiding line*), enhancing efficiency. This approach from LineBO can be interpreted as moving the data points following predefined lines as guiding directions. Despite its strong theoretical foundation, LineBO is not empirically competitive against recent high-dimensional BO baselines (Letham et al. 2020; Bardou, Thiran, and Begin 2024).

In this work, we propose a novel high-dimensional BO method inspired by the line-based optimization approach, where the optimization process is guided by our proposed novel guiding lines learnt from observed data. Specifically, inspired by Particle Swarm Optimization (PSO) algorithm (Kennedy and Eberhart 1995), we propose to use incumbents - the best solutions found so far - to determine the directions of these guiding lines. We make use of two types of incumbents: (1) the global incumbent across all data points and (2) the personal incumbent of each data point throughout its history. By using these incumbents, we encourage exploitation in regions closer to the incumbents, mitigating the over-explorative behavior in high dimensions. We term our proposed guiding lines as *incumbent-guided* lines. We then introduce a new line-based optimization process tailored for these incumbent-guided lines, using a novel multi-objective acquisition function. Furthermore, to increase the sample

efficiency, we propose to select the best line in each iteration using a multi-armed bandit (MAB) technique. Finally, we also propose to incorporate a subspace embedding technique (Papenmeier, Nardi, and Poloczek 2022) to further improve the overall performance. We name our method *High-dimensional Bayesian Optimization via Incumbent-guided Direction Lines and Subspace Embeddings* (BOIDS). To provide deeper insight into the novel incumbent-guided lines, we provide theoretical analysis on the convergence property and derive simple regret bound for BOIDS. Our experimental results demonstrate that BOIDS outperforms state-of-the-art high-dimensional BO methods on a collection of synthetic and real-world benchmark problems. We summarize our contributions as follows.

- We propose BOIDS, a novel high-dimensional BO algorithm that incorporates incumbent-guided line directions into a tailored line-based optimization process, employs a multi-armed bandit technique to select the optimal directions, and uses subspace embedding technique to boost the overall performance.
- We provide theoretical analysis on the convergence property of our proposed method.
- We demonstrate that BOIDS empirically outperforms state-of-the-art baselines on a comprehensive set of synthetic and real-world benchmark problems.

## 2 Background

### 2.1 Bayesian Optimization

Without loss of generality, let us consider the minimization problem: given an *expensive black-box* objective function  $f : \mathcal{X} \rightarrow \mathbb{R}$  where  $\mathcal{X} \subset \mathbb{R}^d$  is the search space, the goal is to find the global optimum  $\mathbf{x}^*$  of the objective function  $f$ ,

$$\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (1)$$

using the least number of function evaluations.

BO iteratively proposes potential data point  $\mathbf{x}$  and obtain noisy function evaluation  $y = f(\mathbf{x}) + \varepsilon$ , where  $\varepsilon \sim \mathcal{N}(0, \sigma^2)$  represents observation noise. In each iteration, BO constructs a surrogate model from the observed dataset  $\mathcal{D} = \{\mathbf{x}^{(i)}, y^{(i)}\}_{i=1}^t$  up to iteration  $t$ , then constructs an acquisition function to decide next data point for observation.

The surrogate model for BO is a statistical model that approximates the objective function. The most common type of BO surrogate model is Gaussian Process (GP), which is characterized by a mean function for the prior belief and a kernel that captures the behavior of the objective function, such as smoothness (Williams and Rasmussen 2006). Other common types of BO surrogate model include TPE (Bergstra et al. 2011), Random Forest (Hutter, Hoos, and Leyton-Brown 2011), neural networks (Müller et al. 2023).

The acquisition function  $\alpha : \mathcal{X} \rightarrow \mathbb{R}$  quantifies the potential information gain at individual point  $\mathbf{x}$  across the search space. By maximizing the acquisition function  $\mathbf{x}^{(t+1)} \in \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha(\mathbf{x})$ , which maximizes the information gain, the next data point for observation is selected. Common choices of acquisition functions include Expected Improvement (Mockus, Tiesis, and Zilinskas 1978), Upper Confidence Bound (Srinivas et al. 2009), Thompson Sampling

(Thompson 1933). Despite their different motivations and properties, acquisition functions must strike a balance between exploitation - sampling in regions where the function values are expected to be optimal, and exploration - sampling in regions where there is significant uncertainty about the objective function.

### 2.2 Line-based Bayesian Optimization

LineBO (Kirschner et al. 2019) addresses high-dimensional BO problems by performing BO on a series of one-dimensional lines (*guiding lines*). In each iteration  $t$ , LineBO defines a guiding line  $\mathcal{L} = \mathcal{L}(\hat{\mathbf{x}}, \mathbf{v}) \in \mathcal{X}$  that passes through the current observed data point  $\hat{\mathbf{x}} = \mathbf{x}^{(t)}$  and follows a uniformly random direction  $\mathbf{v}$ , i.e.,  $\mathbf{v}$  is uniformly sampled from a  $d$ -dimensional unit sphere. Then LineBO finds the next data point for evaluation by maximizing the acquisition function over the guiding line:  $\mathbf{x}^{(t+1)} = \arg \max_{\mathbf{x} \in \mathcal{L}} \alpha(\mathbf{x})$ . This line-based optimization helps optimize the acquisition function more efficiently, especially in high dimensions, where the search space  $\mathcal{X}$  is exponentially large. Despite having strong theoretical property, the empirical performance of LineBO is not competitive with recent state-of-the-art high-dimensional BO methods. We hypothesize that this is due to (1) the uniformly random construction of  $\mathbf{v}$  and (2) the restriction of acquisition optimization on 1D lines which significantly reduces the number of potential data points. Inspired by theoretical and empirical foundations of LineBO, in this paper, we propose a novel algorithm for line-based optimization that guides optimization by direction  $\mathbf{v}$  learnt from observed data.

### 2.3 Subspace Embedding

Subspace embedding methods address the high-dimensional optimization problems by leveraging latent subspaces for optimization. In this work, we focus on subspace embedding via linear random transformation (Wang et al. 2016; Nayebi, Munteanu, and Poloczek 2019; Letham et al. 2020). The main idea is to assume the existence of a low-dimensional subspace (active subspace)  $\mathcal{Z}$  with dimension  $d_e \leq d$ , a function  $g : \mathcal{Z} \rightarrow \mathbb{R}$  and a projection matrix  $\mathbf{T} : \mathcal{X} \rightarrow \mathcal{Z}$  such that  $\forall \mathbf{x} \in \mathcal{X}, g(\mathbf{T}\mathbf{x}) = f(\mathbf{x})$ . Then subspace embedding reduces the high-dimensional optimization problem to a lower-dimension subspace, effectively mitigate the curse of dimensionality issue. As the effective dimension  $d_e$  is unknown in practice, subspace embedding methods instead operate in a low-dimensional target subspace  $\mathcal{A}$  with a pre-defined target dimension  $d_{\mathcal{A}}$ . In subspace embedding BO methods, each iteration proposes a data point  $\mathbf{x}_{\mathcal{A}} \in \mathcal{A}$ , then uses a projection matrix  $\mathbf{S} : \mathcal{A} \rightarrow \mathcal{X}$  for function evaluation. The projection matrix is defined as a sparse matrix  $\mathbf{S} \in \{0, \pm 1\}^{d_{\mathcal{A}} \times d}$  (Nayebi, Munteanu, and Poloczek 2019), therefore, the function evaluation is  $f(\mathbf{S}^T \mathbf{x}_{\mathcal{A}})$ .

## 3 Related Work

Various prior research have addressed the high-dimensional BO problems and can be categorized as follows.

**Effective Surrogate Models.** This approach replaces the common BO surrogate model, Gaussian Process (Williams and Rasmussen 2006), with other models that scale better to high-dimensional spaces, including Tree-Parzen Estimator (Bergstra et al. 2011), Random Forest (Hutter, Hoos, and Leyton-Brown 2011) and neural networks (Snoek et al. 2015; Springenberg et al. 2016). Recently, Prior-data Fitted Networks (PFNs) (Müller et al. 2021) was developed to mimic GP and Bayesian neural network to approximate the posterior predictive distribution for BO (Müller et al. 2023).

**Subspace Embedding.** This approach transforms the high-dimensional optimization problem into a low-dimensional problem on a target subspace. LineBO (Kirschner et al. 2019) operates on one-dimensional target subspaces. HESBO (Nayebi, Munteanu, and Poloczek 2019) and ALEBO (Letham et al. 2020) can handle target subspaces with higher dimensions, yet the target subspaces’ dimensions need to be predefined as a hyperparameter, which may lead to suboptimal performance. To overcome this, BAXUS (Papenmeier, Nardi, and Poloczek 2022) adaptively increases the target subspace’s dimension, which has been shown to be more effective.

**Variable Selection.** This approach learns a subset of variables to perform low-dimensional BO. DropOut (Li et al. 2018) is a classical approach that randomly selects the variables. SAASBO (Eriksson and Jankowiak 2021) learns the most relevant low-dimensional subspace using Sparse Axis-Aligned Subspace Prior for GP modelling.

**Search Space Partitioning.** This approach partitions the search spaces to smaller local regions, to enhance the surrogate modelling and encourage exploitation of potential regions that might contain the global optimum. TuRBO (Eriksson et al. 2019) defines the local regions as hyper-rectangles, updating their sizes based on the success or failure of the optimization process. As TuRBO’s local regions are heuristic, other works attempt to learn the local regions from data. LAMCTS (Wang, Fonseca, and Tian 2020) learns potential local regions with non-linear boundary, improving the performance at the cost of being more computational expensive. CMA-BO (Ngo et al. 2024) leverages the Covariance Matrix Adaptation strategy to learn hyper-ellipsoid local regions, which are updated by inexpensive computation.

**Objective Function Decomposition.** This approach assumes an additive structure of the objective functions and constructs multiple low-dimensional GPs to approximate the high-dimensional objective functions (Kandasamy, Schneider, and Póczos 2015; Hoang et al. 2018; Han, Arora, and Scarlett 2021). Recently, RDUCB (Ziomek and Ammar 2023) suggests that random decomposition without learning from data can result in competitive performance. DumBO (Bardou, Thiran, and Begin 2024) relaxes the assumption on the maximum dimension of each decomposition, which was required to be fixed and low in previous works.

**Evolutionary Algorithms (EA).** EA is a family of black-box optimization algorithms known for their empirical success in high-dimensional optimization problems. CMA-ES

(Hansen and Ostermeier 2001) is a well-known powerful black-box optimization method. Particle Swarm Optimization (PSO) (Kennedy and Eberhart 1995) makes use of the incumbents to guide the swarm (set of data points) for searching the optima, which serves as the inspiration for our proposed method. Empirically, EA methods are generally less sample-efficient than BO algorithms.

## 4 BOIDS: High-Dimensional Bayesian Optimization via Incumbent-guided Direction Lines and Subspace Embeddings

In this section, we present our proposed method BOIDS. As illustrated in Fig. 1, BOIDS maintains a set of  $m$  incumbent-guided lines simultaneously (Sec. 4.1). Then, the best line is selected using a multi-armed bandit strategy (Sec. 4.2). Subsequently, we sample the next data point using our novel line-based optimization and multi-objective acquisition functions (Sec. 4.3). Finally, our subspace embedding technique improves the efficiency of BOIDS (Sec. 4.4).

### 4.1 Incumbent-guided Direction Lines

Our method, BOIDS, is a line-based optimization approach in which we propose to compute the guiding lines using the *incumbents*, i.e., the best data points found so far. BOIDS inspires the idea from the PSO algorithm (Kennedy and Eberhart 1995), particularly the velocity formula, to define the incumbent-guided lines. At iteration  $t$ , let us denote  $\{\mathbf{x}_i^{(t)}\}_{i=1}^m$  as the set of  $m$  data points (aka the  $m$  particles in PSO). Let us denote  $\mathbf{p}_i^{(t)}$  as the personal incumbent (aka PSO personal best) of the  $i$ -th data point  $\mathbf{x}_i^{(t)}$ , which is its best historical location (in terms of function values) up to iteration  $t$ . Let us denote  $\mathbf{g}^{(t)}$  as the global incumbent (aka PSO global best), which is the best data point in observed dataset  $\mathcal{D}$  up to  $t$ . See the Appendix for more details on the incumbents. The incumbent-guided direction  $\mathbf{v}_i^{(t+1)}$  for the  $i$ -th data point  $\mathbf{x}_i^{(t)}$  in the next iteration is then defined as,

$$\mathbf{v}_i^{(t+1)} = w\bar{\mathbf{x}}_i^{(t)} + \mathbf{r}_1 c_1 \bar{\mathbf{p}}_i^{(t)} + \mathbf{r}_2 c_2 \bar{\mathbf{g}}^{(t)}, \quad (2)$$

where  $\bar{\mathbf{x}}_i^{(t)}$  is the displacement vector of the  $i$ -th point  $\mathbf{x}_i^{(t)}$  relative to its last update,  $\bar{\mathbf{p}}_i^{(t)} = \mathbf{p}_i^{(t)} - \mathbf{x}_i^{(t)}$  is the personal best direction of the  $i$ -th point  $\mathbf{x}_i^{(t)}$ , and  $\bar{\mathbf{g}}^{(t)} = \mathbf{g}^{(t)} - \mathbf{x}_i^{(t)}$  is the global best direction. The uniformly random vectors  $\mathbf{r}_1$  and  $\mathbf{r}_2$  are sampled from  $\mathcal{U}([0, 1]^d)$ , and  $w$ ,  $c_1$ ,  $c_2$  are coefficients that control exploration and exploitation. The set of incumbent-guided lines for the  $m$  data points are then defined as  $\mathbb{L} = \{\mathcal{L}(\hat{\mathbf{x}}_i, \mathbf{v}_i)\}_{i=1}^m$ , where each line  $\mathcal{L}_i = \mathcal{L}(\hat{\mathbf{x}}_i, \mathbf{v}_i)$  passes through a point  $\hat{\mathbf{x}}_i = \mathbf{x}_i^{(t)}$  and follows directions  $\mathbf{v}_i$  as in Eq. (2). These incumbent-guided lines encourage exploitation in promising regions by focusing the search on incumbents’ directions, while preserving the explorative behavior due to inherent randomness. As shown by the performance of PSO (Kennedy and Eberhart 1995), directions  $\mathbf{v}_i$  can point towards areas containing the global optimum, effectively guiding the line-based optimization approach towards it. Note the settings for incumbent-guided lines including  $m$ ,  $w$ ,  $c_1$ ,  $c_2$  can follow PSO (See Sec. 6.1).

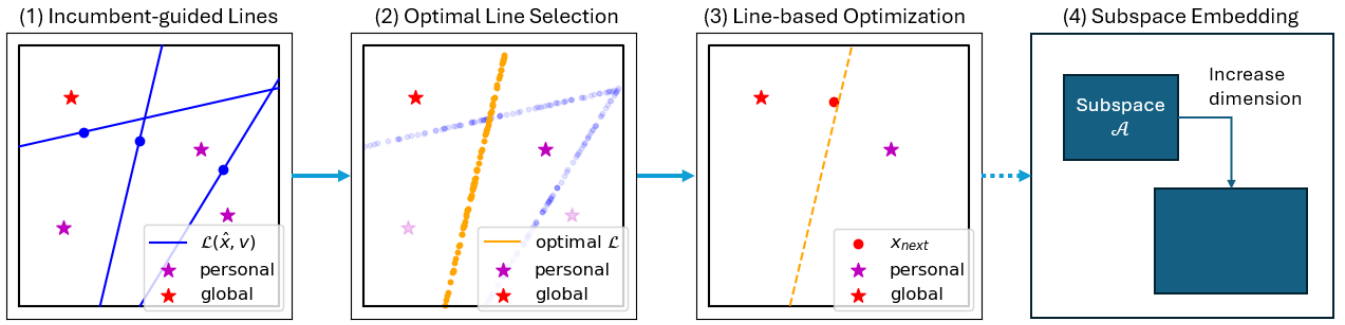


Figure 1: Illustration of our proposed BOIDS algorithm. (1) The incumbent-guided lines are constructed from the personal and global incumbents. (2) The optimal lines are selected based on Thompson Sampling MAB strategy. (3) Line-based Optimization is performed following the optimal line. (4) Subspace embedding technique is incorporated to enhance the performance.

## 4.2 Adaptive Line Selection

In this section, we describe our proposed technique that adaptively selects the most optimal direction from the set of  $m$  incumbent-guided line directions  $\mathbb{L} = \{\mathcal{L}_i\}_{i=1}^m$  formulated in Sec. 4.1. We formulate this problem as a multi-armed bandit (MAB) problem, in which each of  $m$  lines  $\mathcal{L}_i$  represents an arm to be pulled, and the goal is to choose the most optimal arm (i.e., the arm with the highest reward/performance). To solve this MAB problem, we employ Thompson Sampling (TS) (Thompson 1933), a well-known algorithm for decision-making tasks (Agrawal and Goyal 2012; Eriksson et al. 2019). We leverage the GP surrogate model, trained from all observed data so far, as the probability distribution to compute the reward of each arm. To compute the reward, we first sample a random realization  $g$  from the posterior of the GP surrogate model. Then we sample a pool of candidate data points across all  $m$  lines and compute the TS values based on the random realization  $g$ . The reward  $r_i$  of line  $\mathcal{L}_i$  is defined as the maximum TS values among all data points on  $\mathcal{L}_i$ , i.e.,  $r_i = \max_{\mathbf{x} \in \mathcal{L}_i} g(\mathbf{x})$ . Finally, we maximize the reward among  $m$  lines to identify the optimal line chosen for line-based optimization. Formally, the line  $\mathcal{L}_{i^*}$  is selected by maximizing the following equation,

$$i^* = \operatorname{argmax}_{i=1, \dots, m} \max_{\mathbf{x} \in \mathcal{L}_i} g(\mathbf{x}) \quad \text{where } g \sim \mathcal{GP}(\mathcal{D}), \quad (3)$$

and  $\mathcal{GP}(\mathcal{D})$  denotes the GP posterior trained from all the observed data  $\mathcal{D}$  so far. Note that even though we use TS for line direction selection, the line-based optimization process can employ different acquisition functions. The pseudo code for this procedure is shown in the Appendix.

## 4.3 Incumbent-guided Line-based Optimization

In this section, we present the line-based optimization process developed for the proposed incumbent-guided lines. Let the optimally chosen incumbent-guided direction line from Sec. 4.2 be denoted as  $\mathcal{L}_{i^*} = \mathcal{L}(\hat{\mathbf{x}}_{i^*}^{(t)}, \mathbf{v}_{i^*})$ . The main idea is to suggest the next data point for observation by guiding  $\hat{\mathbf{x}}_{i^*}^{(t)}$  following  $\mathcal{L}_{i^*}$ , while being quantified by a common acquisition function  $\alpha(\cdot)$ , such as EI, TS. Instead of optimizing  $\alpha(\cdot)$  solely over the guiding line  $\mathcal{L}_{i^*}$  (as done in LineBO),

we optimize  $\alpha(\cdot)$  in the entire search space  $\mathcal{X}$  while imposing additional constraints that incorporate the incumbent information. Specifically, we define the Euclidean distance constraints to the personal incumbent of  $\hat{\mathbf{x}}_{i^*}^{(t)}$  and the current global incumbents:  $L_{\mathbf{p}}(\mathbf{x}) = \|\mathbf{x} - \mathbf{p}_{i^*}^{(t)}\|$  and  $L_{\mathbf{g}}(\mathbf{x}) = \|\mathbf{x} - \mathbf{g}^{(t)}\|$ , respectively. These constraints encourage exploitation in the direction of incumbents, which can guide the search towards the global optimum. We then formulate the acquisition optimization problem as a multi-objective (MO) optimization problem where the Euclidean distances are treated as additional objectives alongside the acquisition function. The MO acquisition function for BOIDS is then,

$$\mathcal{P} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} (f_{\alpha}(\mathbf{x}), f_{\mathbf{p}}(\mathbf{x}), f_{\mathbf{g}}(\mathbf{x})), \quad (4)$$

where  $f_{\alpha}(\cdot) = \alpha(\cdot)$ ,  $f_{\mathbf{p}}(\cdot) = -L_{\mathbf{p}}(\cdot)$  and  $f_{\mathbf{g}}(\cdot) = -L_{\mathbf{g}}(\cdot)$ . In practice, to direct the acquisition optimization to follow the optimal incumbent-guided line  $\mathcal{L}_{i^*}$ , we initialize Eq. (4) with data points randomly sampled on  $\mathcal{L}_{i^*}$ , then optimize Eq. (4) using off-the-shelf MO solvers, such as evolutionary algorithms. Solving Eq. (4) yields a Pareto set of  $n_{\mathcal{P}}$  solutions  $\mathcal{P} = \{\mathbf{x}_i\}_{i=1}^{n_{\mathcal{P}}}$  and a Pareto front  $\mathcal{F} = \{f_{\alpha}(\mathbf{x}_i), f_{\mathbf{p}}(\mathbf{x}_i), f_{\mathbf{g}}(\mathbf{x}_i)\}_{i=1}^{n_{\mathcal{P}}}$ . Finally we select the next BO data points for observation by maximizing  $\mathcal{F}$  w.r.t.  $f_{\alpha}(\cdot)$ ,

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in \mathcal{P}} f_{\alpha}(\mathbf{x}). \quad (5)$$

Eq. (5) selects the solution with the best acquisition objective on the Pareto front, ensuring that the selected data points are optimal (based on acquisition function values), while also being non-dominated by other solutions in terms of the Euclidean distances to the incumbents (based on MO optimization). See the Appendix for the pseudo code.

## 4.4 Subspace Embedding

To further enhance BOIDS's capability in handling high-dimensional optimization problems, we employ a linear random subspace embedding strategy. Specifically, we leverage the adaptive expanding subspace embedding, known as the BAXUS embedding (Papenmeier, Nardi, and Poloczek 2022), and execute BOIDS through a series of low-dimensional subspace with varying dimensionalities. Compared to other random linear embedding rules, such as

HESBO (Nayebi, Munteanu, and Poloczek 2019), the BAxUS embedding produces the highest worst-case probability to contain the global optimum, given the same conditions (see further discussion in Sec. 5.2). In detail, given a low-dimensional subspace  $\mathcal{A}$  with dimension  $d_{\mathcal{A}}$ , BOIDS proposes data points  $\mathbf{x}_{\mathcal{A}} \in \mathcal{A}$ , which are projected to the original input space  $\mathcal{X}$  via a projection matrix  $\mathbf{S}$ . The optimization on  $\mathcal{A}$  is performed within a budget  $T_{d_{\mathcal{A}}}$ , after which the subspace  $\mathcal{A}$  is expanded to a subspace with higher dimensionality. During the expansion of  $\mathcal{A}$ , previous observations and the set of  $m$  incumbent-guided lines are preserved via embedding matrix  $\mathbf{S}$ . Note that we only employ the BAxUS embedding rule, not other components of the BAxUS algorithm such as trust regions. Specifically, we adopt the rule for increasing the subspace dimensionality per expansion and the number of iterations budget  $T_{d_{\mathcal{A}}}$  for each subspace  $\mathcal{A}$ . The subspace expansion process is conducted repeatedly until  $d_{\mathcal{A}}$  reaches the input dimension  $d$ , after which BOIDS operates in the input space  $\mathcal{X}$ .

#### 4.5 The Proposed BOIDS Algorithm

Overall, the BOIDS algorithm operates as follows (pseudo-code in Alg. 1 and illustration in Fig. 1). At the beginning, based on a predefined initial subspace dimension  $d_{\mathcal{A}}$ , we initialize the embedding matrix  $\mathbf{S}$ , compute the budget  $T_{d_{\mathcal{A}}}$ , randomly initialize observed dataset  $\mathcal{D}_0 \in \mathcal{A}$  and randomly choose  $m$  data points in  $\mathcal{D}_0$  (lines 3-4). Then, in each iteration, after constructing a GP surrogate model  $\mathcal{GP}$ , we sequentially compute the set of  $m$  incumbent-guided directions  $\mathbb{L} = \{\mathcal{L}(\hat{\mathbf{x}}_i, \mathbf{v}_i)\}_{i=1}^m$ , select the most optimal line direction  $\mathcal{L}_{i^*}$  and perform line-based optimization to suggest the next data point for observation  $\mathbf{x}_{\mathcal{A}} \in \mathcal{A}$  (lines 7-10). Then, we project  $\mathbf{x}_{\mathcal{A}}$  to  $\mathcal{X}$  via projection matrix  $\mathbf{S}$  and compute objective function evaluation (lines 11-12). The process repeats until the budget  $T_{d_{\mathcal{A}}}$  for subspace  $\mathcal{A}$  is depleted. After that, if  $d_{\mathcal{A}} < d$ , we expand the subspace  $\mathcal{A}$  dimensionality, update the matrix  $\mathbf{S}$  and recompute the budget  $T_{d_{\mathcal{A}}}$  (line 16). Otherwise, we discard all previous observed data, and restart optimization in the input subspace (line 18). The algorithm repeats until the predefined budget  $T$  is depleted, then returns the best solution found  $\mathbf{x}^*$  as the final solution.

### 5 Theoretical Analysis

In this section, we provide theoretical analysis of the convergence property of BOIDS, including local and global convergence properties. We prove that our incumbent-guided line-based optimization has a sub-linear simple regret bound, enabling BOIDS to converge to (local) optimum. Moreover, as we incorporate a subspace embedding technique, we discuss the global convergence property given the chosen embedding rule.

#### 5.1 Local Convergence

The main idea is to show that BOIDS, leveraging the proposed incumbent-guided directions, can converge to a (local) optimum, supported by a sub-linear simple regret bound. The simple regret  $r_t$  is defined as the difference between the global optimum value  $f^*$  and the best function

---

#### Algorithm 1: The BOIDS Algorithm

---

```

1: Input: Function  $f(\cdot)$ , budget  $T$ , initial dimension  $d_{\mathcal{A}}$ ,
   number of data points  $m$ 
2: Output: The optimum  $\mathbf{x}^*$ 
3: Initialize subspace  $\mathcal{A}$ , matrix  $\mathbf{S}$ , budget  $T_{d_{\mathcal{A}}}$ 
4: Initialize  $\mathcal{D}_0 \in \mathcal{A}$  and randomly choose  $m$  data points
5: while budget  $T$  not depleted do
6:   while budget  $T_{d_{\mathcal{A}}}$  not depleted do
7:     Build GP surrogate model  $\mathcal{GP}_{\mathcal{A}}$  from  $\mathcal{D}$ 
8:     Compute set of lines  $\mathbb{L} = \{\mathcal{L}(\hat{\mathbf{x}}_i, \mathbf{v}_i)\}_{i=1}^m$ 
9:      $\mathcal{L}_{i^*} \leftarrow \text{LINE-SELECT}(\mathcal{GP}_{\mathcal{A}}, \mathbb{L}) \quad \triangleright \text{Sec. 4.2}$ 
10:     $\mathbf{x}_{\mathcal{A}} \leftarrow \text{LINE-OPT}(\mathcal{GP}_{\mathcal{A}}, \mathcal{L}_{i^*}) \quad \triangleright \text{Sec. 4.3}$ 
11:    Project  $\mathbf{x}_{\mathcal{X}} \leftarrow \mathbf{S}^T \mathbf{x}_{\mathcal{A}}$ 
12:    Compute function evaluation  $y = f(\mathbf{x}_{\mathcal{X}}) + \varepsilon$ 
13:    Update  $\mathcal{D} \leftarrow \{\mathbf{x}_{\mathcal{A}}, \mathbf{x}_{\mathcal{X}}, y\}$ 
14:   end while
15:   if  $d_{\mathcal{A}} < d$  then
16:     Increase  $d_{\mathcal{A}}$ , update  $\mathbf{S}$ , compute  $T_{d_{\mathcal{A}}}$ 
17:   else
18:     Discard previous dataset, set  $d_{\mathcal{A}} = d$ ,  $\mathbf{S} = \mathbb{I}^d$ 
19:   end if
20: end while
21: Return  $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}^{(i)} \in \mathcal{D}} \{y^{(i)}\}_{i=1}^N$ 

```

---

value  $y_i$  found so far until iteration  $t$ . For a minimization problem,  $r_t = \min_{i=1}^t y^{(i)} - f^*$ . We will now proceed to derive the bound for the simple regret. We first impose the following assumption on the objective function  $f$ .

**Assumption 1 (RKHS)** *The objective function  $f$  is a member of a reproducing kernel Hilbert space  $\mathcal{H}(k)$  with known kernel  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  and bounded norm  $\|f\|_{\mathcal{H}} \leq B$ .*

We leverage the standard analysis of BO derived in (Srinivas et al. 2009) for our regret analysis. In particular, the regret bound is often expressed in terms of the *maximum information gain*  $\gamma_T$ , whose bounds are known for common kernels such as Matérn and RBF (Srinivas et al. 2009; Seeger, Kakade, and Foster 2008). Similar to LineBO (Kirschner et al. 2019), we impose the following assumptions on  $\gamma_T$ .

**Assumption 2 (Maximum Information Gain)** *Let  $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$  be a one-dimensional kernel. Let  $\kappa = 2/(2\nu + 2)$ , where  $\nu$  is the smoothness factor of kernel  $k$ . The maximum information gain is then*

$$\gamma_T(k) \leq \mathcal{O}(T^\kappa \log T). \quad (6)$$

In BOIDS, as we use the incumbent-guided line as the search direction, we derive the following lemma.

**Lemma 1 (Incumbent-guided Search Direction)** *Let  $\mathbf{r}_1, \mathbf{r}_2 \in \mathbb{R}^d$  be two random vectors sampled from a uniform distribution  $\mathcal{U}([0, 1]^d)$ ,  $\mathbf{h}_1, \mathbf{h}_2 \in \mathbb{R}^d$  are arbitrary fixed vectors, and the element-wise multiplications  $\mathbf{v}_1 = \mathbf{r}_1 \circ \mathbf{h}_1$  and  $\mathbf{v}_2 = \mathbf{r}_2 \circ \mathbf{h}_2$ . Then for all vectors  $\mathbf{g} \in \mathbb{R}^d$*

$$\mathbb{E}[\langle \mathbf{g}, \mathbf{v}_1 + \mathbf{v}_2 \rangle^2] \geq C \|\mathbf{g}\|^2, \quad (7)$$

where constant  $C \in \mathbb{R}$  is a constant such that  $C \leq \frac{1}{4}(\|\mathbf{h}_1\|^2 \cos^2 \theta_1 + \|\mathbf{h}_2\|^2 \cos^2 \theta_2)$ , and  $\theta_1, \theta_2$  are the angles between  $\mathbf{g}$  and  $\mathbf{h}_1, \mathbf{h}_2$ , respectively.

See Appendix for the proof. In Lemma 1,  $\mathbf{h}_1$  and  $\mathbf{h}_2$  represents the two incumbent directions,  $\bar{\mathbf{p}}$  and  $\bar{\mathbf{g}}$  in Eq. (2). From Lemma 1, we derive the following simple regret bound.

**Proposition 1** *Let Assumptions 1 and 2 hold. After  $T$  iterations of BOIDS with random directions that satisfy Lemma 1, the expected simple regret over the random vectors  $\mathbf{r}_1$  and  $\mathbf{r}_2$  is bounded as follows,*

$$\mathbb{E}[r_T] \leq \mathcal{O}((d \log T/T)^{1/2-\kappa}). \quad (8)$$

See the Appendix for the proof. In the case of our implementation with Matern 5/2 kernel for the experiments, the simple regret is then bounded by  $\mathbb{E}[r_T] \leq \mathcal{O}((d \log T/T)^{3/14})$ . Compared to LineBO (Kirschner et al. 2019), our simple regret is tighter bounded in terms of the input dimension  $d$ , while having a similar rate in terms of  $T$ .

## 5.2 Global Convergence

As BOIDS is a line-based optimization method, it inherits the global convergence property from LineBO (Theorem 1 of Kirschner et al., 2019). However, because BOIDS also incorporates the linear random subspace embedding technique, the global convergence property depends on whether or not the subspace  $\mathcal{A}$  contains the global optimum  $\mathbf{x}_{\mathcal{X}}^*$ , i.e., there exists a point  $\mathbf{x}_{\mathcal{A}}^* \in \mathcal{A}$  such that  $\mathbf{S}^\top \mathbf{x}_{\mathcal{A}}^* = \mathbf{x}_{\mathcal{X}}^*$ , where  $\mathbf{S}$  is the embedding matrix. Previous works analyze this property via *the probability of the embedding to contain the global optimum* (Nayebi, Munteanu, and Poloczek 2019; Letham et al. 2020; Papenmeier, Nardi, and Poloczek 2022). As BOIDS employs BAXUS embedding, we use the following theorem (Theorem 1 of Papenmeier, Nardi, and Poloczek, 2022) to reason about this probability,

**Theorem 1** *Let  $\alpha = \lfloor d/d_{\mathcal{A}} \rfloor$ ,  $\beta = \lceil d/d_{\mathcal{A}} \rceil$ . Denote  $\binom{a}{b}$  as the binomial coefficients. The probability of BAXUS embedding to contain the global optimum is,*

$$p^* = \frac{\sum_{i=0}^{d_e} \binom{d_{\mathcal{A}}(1+\alpha)-d}{i} \binom{d-d_{\mathcal{A}}\beta}{d_e-i} \alpha^i \beta^{d_e-i}}{\binom{d}{d_e}}. \quad (9)$$

A useful result from Theorem 1 is that the worst case success probability becomes one when  $d_{\mathcal{A}} = d$ , guaranteeing that the embedding in BOIDS always contains the global optimum. This property, together with the simple regret bound in Proposition 1, concludes BOIDS convergence property.

## 6 Experiments

In this section, we extensively evaluate our proposed method BOIDS against a set of baselines on a set of benchmark problems, showing that BOIDS outperforms state-of-the-art baselines. We additionally conduct ablation study to understand each component of the proposed method.

### 6.1 Experimental Settings and Baselines

We evaluate our proposed algorithm **BOIDS** against a comprehensive set of baselines: **Standard BO**, **LineBO** (Kirschner et al. 2019), **BAXUS** (Papenmeier, Nardi, and Poloczek 2022), **HESBO** (Nayebi, Munteanu, and Poloczek 2019), **ALEBO** (Letham et al. 2020), **SAASBO** (Eriksson and Jankowiak 2021), **TuRBO** (Eriksson et al. 2019),

**CMA-BO** (Ngo et al. 2024), **RDUCB** (Ziomek and Ammar 2023), **CMA-ES** (Hansen and Ostermeier 2001) and **PSO** (Kennedy and Eberhart 1995).

For ALEBO and HESBO, we run each method with different target dimensions  $d_{\mathcal{A}} = 10$  and  $d_{\mathcal{A}} = 20$ . For LineBO, we compare to CoordinateLineBO version, as it was identified as the best variant proposed by Kirschner et al., 2019. We run all methods 10 times with different random seeds and report the mean and standard error. Most methods are run with 1000 iterations on all problems, except for SAASBO (100 iterations), ALEBO (500 iterations) and RDUCB (2 days running time) due to high computational time and memory required. The settings of BOIDS are as follows. For the incumbent-guided direction lines, we use a common configuration as with PSO,  $m = 20$ ,  $w = 0.729$ ,  $c_1 = c_2 = 2.05w$  (Shi and Eberhart 1998; Regis 2014; Li et al. 2020). For the surrogate model, we use a GP with Matérn 5/2 kernel function. For acquisition function, we use Thompson Sampling and employ NSGA-II algorithm (Deb et al. 2002) as the MO optimizer for Eq. (4). Detailed implementation of BOIDS and baselines are in the Appendix.

### 6.2 Benchmark Problems

We conduct experiments on 3 synthetic and 3 real-world benchmark problems. The dimensions of these problems range from 100 to 500. For synthetic problems, we use *Ackley-100D*, *Branin-500D* and *Hartmann-500D*, which are widely-used in BO research works (Eriksson et al. 2019; Song et al. 2022; Ngo et al. 2024). For real-world benchmark problems, we use *Mopta-124D*, *LassoDNA-180D* and *HalfCheetah-102D*. Mopta-124D is a vehicle design task aiming to minimize the vehicle’s weight, with code implementation from Eriksson and Jankowiak, 2021. LassoDNA-180D is a hyperparameter tuning task involving a microbiology DNA dataset, with code implementation from Šehić et al., 2022. HalfCheetah-102D is a reinforcement learning task aiming to maximize a cumulative reward, with code implementation from Song et al., 2022. Details on each benchmark problem are in the Appendix.

### 6.3 Comparison with Baselines

Fig. 2 shows the performance of all baselines across all six benchmark problems. On synthetic problems (Ackley-100D, Branin-500D and Hartmann-500D), BOIDS significantly outperforms all baselines. BAXUS converges to less optimal values. SAASBO, despite being a strong baseline, especially on Branin-500D, is prohibitively expensive and is only suitable for very limited budgets (up to 64 GB Memory for 100 iterations). For real-world benchmark problems, BOIDS consistently outperforms all baselines. On Mopta-124D, BOIDS shows a slower start compared to TuRBO due to its initial focus on low-dimensional subspaces, a behavior observed in subspace embedding techniques including BAXUS. However, BOIDS quickly finds more optimal solutions as iterations progress. On LassoDNA-180D, even though both RDUCB and CMA-BO find similarly optimal solutions as BOIDS, BOIDS converges much faster. On HalfCheetah-102D, ALEBO ( $d_{\mathcal{A}} = 10$ ) is competitive to

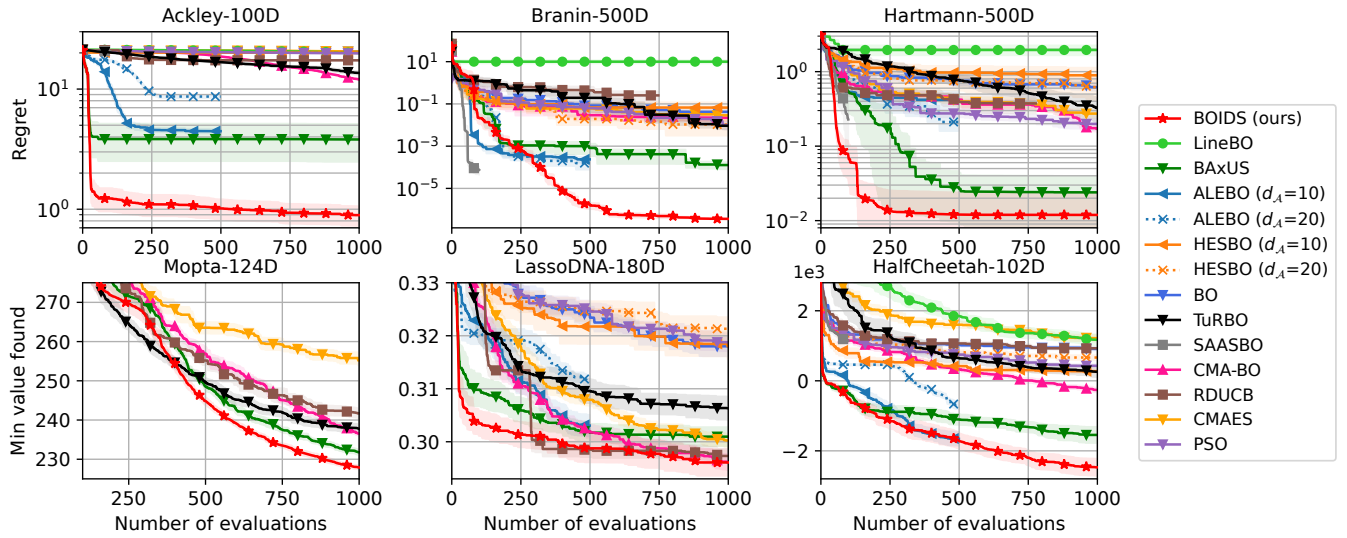


Figure 2: Comparison of our proposed method, BOIDS, against the state-of-the-art baselines on six minimization problems. Note that some methods (ALEBO, SAASBO and RDUCEB) only have limited iterations due to the prohibitively high computational cost and memory required. Overall, BOIDS outperforms all baselines significantly on most problems.

BOIDS but requires careful tuning of the target dimension, with ALEBO ( $d_A = 20$ ) showing significantly worse performance. Additionally, ALEBO is also prohibitively expensive to run (up to 40 hours for 500 iterations for each repeat). Overall, BOIDS demonstrates superior performance across these benchmark problems, indicating a robust and efficient solution for high-dimensional optimization problems.

#### 6.4 Ablation Study

We conduct an ablation study to investigate the effectiveness of each component in BOIDS by alternatively removing each of the four components (as described in Sec. 4.5) and comparing the four corresponding versions with the original BOIDS. First, in version “**W/o Guided Direction**”, we replace the incumbent-guided line component with uniformly random guiding lines  $\mathcal{L}(\hat{\mathbf{x}}, \mathbf{v})$ , where  $\mathbf{v} \sim \mathcal{U}([l-u, u-l]^d)$ . Second, in version “**W/o Line Select**”, we remove the MAB-based line selection component and instead select the lines randomly in each iteration,  $\mathcal{L}_i$  where  $i = \mathcal{U}([1 \dots m])$ . Third, in version “**W/o Line Opt**”, we omit the line-based optimization component, and instead optimize the acquisition function over the entire input space  $\mathcal{X}$ , i.e.,  $\mathbf{x}^{(t+1)} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \alpha(\mathbf{x})$ . Finally, in version “**W/o Embedding**”, we exclude the subspace embedding component and perform optimization directly in the input dimension  $d_A = d$  and  $\mathbf{S} = \mathbf{I}_d$ . All other settings in the four versions are kept similar to the original BOIDS.

We run these versions on two benchmark problems - LassoDNA-180D and HalfCheetah-102D - with 10 repeats, and compare them against the original method. We report the mean and standard error of the results in Fig. 3. Overall, the results show that the original BOIDS consistently shows the best performance. Among the components, the line selection component via MAB (Sec. 4.2) seems to be the least important, as removing it does not significantly degrade the perfor-

mance. Conversely, the incumbent-guided line component (Sec. 4.1) and line-based optimization component (Sec. 4.3) are crucial due to the significant performance degradation in the related versions. The subspace embedding component (Sec. 4.4) greatly affects the performance in HalfCheetah-102D, yet has less impact on LassoDNA-180D problem.




Figure 3: Impact of BOIDS components on performance when each is alternatively removed. The *incumbent-guided line* component (Sec. 4.1) and the tailored *line-based optimization* (Sec. 4.3) have the most significant influence.

## 7 Conclusion

In this paper, we address the high-dimensional optimization problems for expensive black-box objective functions. We propose BOIDS, a novel line-based BO algorithm using incumbent-guided direction lines. To further improve the efficiency, we select the most optimal line via a MAB approach. To further boost the performance, we employ a subspace embedding technique. We theoretically analyze BOIDS’s convergence property and empirically demonstrate that BOIDS outperforms state-of-the-art baselines on a variety of synthetic and real-world tasks.

## Acknowledgments

The first author (L.N.) would like to thank the School of Computing Technologies, RMIT University, Australia for providing computing resources for this project. Additionally, this project was undertaken with the assistance of computing resources from RACE (RMIT AWS Cloud Super-computing Hub).

## References

- Agrawal, S.; and Goyal, N. 2012. Analysis of Thompson Sampling for the Multi-armed Bandit Problem. In Mannor, S.; Srebro, N.; and Williamson, R. C., eds., *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, 39.1–39.26. Edinburgh, Scotland: PMLR.
- Ament, S.; Witte, A.; Garg, N.; and Kusuma, J. 2023. Sustainable Concrete via Bayesian Optimization. arXiv:2310.18288.
- Bardou, A.; Thiran, P.; and Begin, T. 2024. Relaxing the Additivity Constraints in Decentralized No-Regret High-Dimensional Bayesian Optimization. In *The Twelfth International Conference on Learning Representations*.
- Bergstra, J.; Bardenet, R.; Bengio, Y.; and Kégl, B. 2011. Algorithms for hyper-parameter optimization. *NeurIPS*, 24.
- Binois, M.; and Wycoff, N. 2022. A survey on high-dimensional Gaussian process modeling with application to Bayesian optimization. *ACM Transactions on Evolutionary Learning and Optimization*, 2(2): 1–26.
- Deb, K.; Pratap, A.; Agarwal, S.; and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation*, 6(2): 182–197.
- Eriksson, D.; and Jankowiak, M. 2021. High-dimensional Bayesian optimization with sparse axis-aligned subspaces. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 161.
- Eriksson, D.; Pearce, M.; Gardner, J.; Turner, R. D.; and Poloczek, M. 2019. Scalable Global Optimization via Local Bayesian Optimization. In *NeurIPS*.
- Garnett, R. 2023. *Bayesian Optimization*. Cambridge University Press.
- Han, E.; Arora, I.; and Scarlett, J. 2021. High-dimensional Bayesian optimization via tree-structured additive models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 7630–7638.
- Hansen, N.; and Ostermeier, A. 2001. Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation*, 9(2): 159–195.
- Hoang, T. N.; Hoang, Q. M.; Ouyang, R.; and Low, K. H. 2018. Decentralized high-dimensional Bayesian optimization with factor graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2011. Sequential Model-Based Optimization for General Algorithm Configuration. In *LION*, volume 6683 of *Lecture Notes in Computer Science*, 507–523. Springer.
- Kandasamy, K.; Schneider, J.; and Póczos, B. 2015. High dimensional Bayesian optimisation and bandits via additive models. In *International conference on machine learning*, 295–304. PMLR.
- Kennedy, J.; and Eberhart, R. 1995. Particle swarm optimization. In *Proceedings of ICNN'95-international conference on neural networks*, volume 4, 1942–1948. IEEE.
- Kirschner, J.; Mutny, M.; Hiller, N.; Ischebeck, R.; and Krause, A. 2019. Adaptive and safe Bayesian optimization in high dimensions via one-dimensional subspaces. In *International Conference on Machine Learning*, 3429–3438. PMLR.
- Letham, B.; Calandra, R.; Rai, A.; and Bakshy, E. 2020. Re-Examining Linear Embeddings for High-Dimensional Bayesian Optimization. In *NeurIPS*, volume 33.
- Li, C.; Gupta, S.; Rana, S.; Nguyen, V.; Venkatesh, S.; and Shilton, A. 2018. High dimensional Bayesian optimization using dropout. *arXiv preprint arXiv:1802.05400*.
- Li, F.; Shen, W.; Cai, X.; Gao, L.; and Wang, G. G. 2020. A fast surrogate-assisted particle swarm optimization algorithm for computationally expensive problems. *Applied Soft Computing*, 92: 106303.
- Mayr, M.; Ahmad, F.; Chatzilygeroudis, K.; Nardi, L.; and Krueger, V. 2022. Skill-based multi-objective reinforcement learning of industrial robot tasks with planning and knowledge integration. In *2022 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 1995–2002. IEEE.
- Mockus, J.; Tiesis, V.; and Zilinskas, A. 1978. The Application of Bayesian Methods for Seeking the Extremum. *Towards Global Optimization*, 2(117-129): 2.
- Müller, S.; Feurer, M.; Hollmann, N.; and Hutter, F. 2023. Pfn4bo: In-context learning for bayesian optimization. In *International Conference on Machine Learning*, 25444–25470. PMLR.
- Müller, S.; Hollmann, N.; Arango, S. P.; Grabocka, J.; and Hutter, F. 2021. Transformers Can Do Bayesian Inference. In *International Conference on Learning Representations*.
- Nayebi, A.; Munteanu, A.; and Poloczek, M. 2019. A Framework for Bayesian Optimization in Embedded Subspaces. In *ICML*, volume 97. PMLR.
- Ngo, L.; Ha, H.; Chan, J.; Nguyen, V.; and Zhang, H. 2024. High-dimensional Bayesian Optimization via Covariance Matrix Adaptation Strategy. *Transactions on Machine Learning Research*.
- Papenmeier, L.; Nardi, L.; and Poloczek, M. 2022. Increasing the scope as you learn: Adaptive bayesian optimization in nested subspaces. *NeurIPS*, 35.
- Parker-Holder, J.; Rajan, R.; Song, X.; Biedenkapp, A.; Miao, Y.; Eimer, T.; Zhang, B.; Nguyen, V.; Calandra, R.; Faust, A.; et al. 2022. Automated reinforcement learning (autorl): A survey and open problems. *Journal of Artificial Intelligence Research*, 74: 517–568.
- Regis, R. G. 2014. Particle swarm with radial basis function surrogates for expensive black-box optimization. *Journal of Computational Science*, 5(1): 12–23.

- Ru, B.; Wan, X.; Dong, X.; and Osborne, M. 2021. Interpretable Neural Architecture Search via Bayesian Optimisation with Weisfeiler-Lehman Kernels. In *International Conference on Learning Representations*.
- Seeger, M. W.; Kakade, S. M.; and Foster, D. P. 2008. Information consistency of nonparametric Gaussian process methods. *IEEE Transactions on Information Theory*, 54(5): 2376–2382.
- Šehić, K.; Gramfort, A.; Salmon, J.; and Nardi, L. 2022. Lassobench: A high-dimensional hyperparameter optimization benchmark suite for lasso. In *International Conference on Automated Machine Learning*, 2–1. PMLR.
- Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; and De Freitas, N. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175.
- Shi, Y.; and Eberhart, R. 1998. A modified particle swarm optimizer. In *1998 IEEE international conference on evolutionary computation proceedings. IEEE world congress on computational intelligence (Cat. No. 98TH8360)*, 69–73. IEEE.
- Shields, B. J.; Stevens, J.; Li, J.; Parasram, M.; Damani, F.; Alvarado, J. I. M.; Janey, J. M.; Adams, R. P.; and Doyle, A. G. 2021. Bayesian reaction optimization as a tool for chemical synthesis. *Nature*, 590(7844): 89–96.
- Snoek, J.; Rippel, O.; Swersky, K.; Kiros, R.; Satish, N.; Sundaram, N.; Patwary, M.; Prabhat, M.; and Adams, R. 2015. Scalable bayesian optimization using deep neural networks. In *International conference on machine learning*, 2171–2180. PMLR.
- Song, L.; Xue, K.; Huang, X.; and Qian, C. 2022. Monte Carlo Tree Search based Variable Selection for High Dimensional Bayesian Optimization. In *NeurIPS*.
- Springenberg, J. T.; Klein, A.; Falkner, S.; and Hutter, F. 2016. Bayesian optimization with robust Bayesian neural networks. *Advances in neural information processing systems*, 29.
- Srinivas, N.; Krause, A.; Kakade, S. M.; and Seeger, M. 2009. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*.
- Thompson, W. R. 1933. On the likelihood that on unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4): 285–294.
- Turner, R.; Eriksson, D.; McCourt, M.; Kiili, J.; Laaksonen, E.; Xu, Z.; and Guyon, I. 2021. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020. In *NeurIPS 2020 Competition and Demonstration Track*, 3–26. PMLR.
- Wang, L.; Fonseca, R.; and Tian, Y. 2020. Learning Search Space Partition for Black-box Optimization using Monte Carlo Tree Search. In *NeurIPS*, volume 33.
- Wang, X.; Jin, Y.; Schmitt, S.; and Olhofer, M. 2023. Recent advances in Bayesian optimization. *ACM Computing Surveys*, 55(13s): 1–36.
- Wang, Z.; Dahl, G. E.; Swersky, K.; Lee, C.; Nado, Z.; Gilmer, J.; Snoek, J.; and Ghahramani, Z. 2024. Pre-trained Gaussian Processes for Bayesian Optimization. *Journal of Machine Learning Research*, 25(212): 1–83.
- Wang, Z.; Hutter, F.; Zoghi, M.; Matheson, D.; and De Freitas, N. 2016. Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55: 361–387.
- Williams, C. K.; and Rasmussen, C. E. 2006. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.
- Ziomek, J. K.; and Ammar, H. B. 2023. Are random decompositions all we need in high dimensional Bayesian optimization? In *International Conference on Machine Learning*, 43347–43368. PMLR.