



Epsilon-Greedy Thompson Sampling to Bayesian Optimization

Bach Do

Civil and Environmental Engineering,
Uncertainty Quantification Laboratory,
University of Houston,
Houston, TX 77204
e-mail: bdo3@uh.edu

Taiwo Adebisi

Civil and Environmental Engineering,
Uncertainty Quantification Laboratory,
University of Houston,
Houston, TX 77204
e-mail: taadebi2@cougarnet.uh.edu

Ruda Zhang¹

Civil and Environmental Engineering,
Uncertainty Quantification Laboratory,
University of Houston,
Houston, TX 77204
e-mail: rudaz@uh.edu

Bayesian optimization (BO) has become a powerful tool for solving simulation-based engineering optimization problems thanks to its ability to integrate physical and mathematical understandings, consider uncertainty, and address the exploitation–exploration dilemma. Thompson sampling (TS) is a preferred solution for BO to handle the exploitation–exploration tradeoff. While it prioritizes exploration by generating and minimizing random sample paths from probabilistic models—a fundamental ingredient of BO—TS weakly manages exploitation by gathering information about the true objective function after it obtains new observations. In this work, we improve the exploitation of TS by incorporating the ϵ -greedy policy, a well-established selection strategy in reinforcement learning. We first delineate two extremes of TS, namely the generic TS and the sample-average TS. The former promotes exploration, while the latter favors exploitation. We then adopt the ϵ -greedy policy to randomly switch between these two extremes. Small and large values of ϵ govern exploitation and exploration, respectively. By minimizing two benchmark functions and solving an inverse problem of a steel cantilever beam, we empirically show that ϵ -greedy TS equipped with an appropriate ϵ is more robust than its two extremes, matching or outperforming the better of the generic TS and the sample-average TS.

[DOI: 10.1115/1.4066858]

Keywords: Thompson sampling, Bayesian optimization, ϵ -greedy policy, exploitation–exploration dilemma, cyclic constitutive law, computational foundations for engineering optimization, inverse methods for engineering applications, inverse methods for engineering applications, machine learning for engineering applications

1 Introduction

Consider the following minimization problem:

$$\begin{aligned} \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\ \text{subject to} \quad & \mathbf{x} \in \mathcal{X} \end{aligned} \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the vector of d input variables selected in a bounded, compact domain $\mathcal{X} \subset \mathbb{R}^d$, and $f(\mathbf{x}): \mathcal{X} \mapsto \mathbb{R}$ is a real-valued objective function. In science and engineering applications, the objective function $f(\mathbf{x})$ is often a black-box function evaluated via costly simulations, which unfortunately hinders the use of any mature, gradient-based numerical optimization algorithms.

Bayesian optimization (BO) [1–6] is a global sequential optimization technique well-suited for solving small- and moderate-dimensional optimization problems with costly or black-box objective functions. It finds application in diverse domains of science and engineering, including hyperparameter tuning of machine learning algorithms [1], identification of material parameters [7,8], material design [9,10], airfoil design [11], adaptive experimental design [12], and accelerator physics [13]. For a recent comprehensive review of its applications in engineering design, see Ref. [5]. At

its core, BO guides the optimization process using a probabilistic surrogate model \hat{f} of the objective function, usually a Gaussian process (GP), coupled with an optimization policy [4]. Given a dataset that has several observations of the input variables and the corresponding objective function values, a GP posterior built from this dataset often serves as the probabilistic model \hat{f} encapsulating our beliefs about the black-box objective function. The optimization policy specifying what we value in the dataset is defined through an acquisition function $\alpha(\mathbf{x})$. This acquisition function can be deterministic or stochastic, and is cost-effective to evaluate given \hat{f} , making it convenient for processing optimization. Different considerations that should be taken into account when formulating $\alpha(\mathbf{x})$ from \hat{f} include, for example, the value of the objective function [14], the information about the minimum location [15–17], and the information about the minimum value [18].

Various algorithms have been developed to address the classic exploitation–exploration dilemma in BO, where exploitation involves selecting new solutions anticipated to improve the objective function value immediately and exploration focuses on sampling new solutions to reduce uncertainty in predictions of the objective function for a long-term improvement in the solution. An effective BO algorithm is deemed to balance these opposing concerns [4]. Notable deterministic acquisition functions, such as expected improvement (EI) [19], weighted EI [20], GP upper confidence bound [21], knowledge gradient [22], and likelihood-weighting [23], can manage the exploitation–exploration tradeoff.

¹Corresponding author.

Manuscript received May 2, 2024; final manuscript received September 24, 2024; published online November 5, 2024. Assoc. Editor: Guang Lin.

However, they are considered myopic policies, and therefore tend to favor exploitation [14].

Thompson sampling (TS) is a preferred algorithm to address the exploitation–exploration dilemma in solving multi-armed bandit problems [24–27]. More specifically, TS selects arms (or actions) from a finite set of possible arms over a limited number of iterations. Each arm corresponds to a stochastic reward drawn from an unknown distribution. The goal is to craft a sequence of arms that maximizes the cumulative reward assuming that the rewards are independent of time and conditionally independent given the selected arms. When applied to GP-based BO, TS (or GP-TS [4]) generates a sequence of new solution points using the mechanism that involves random sampling from unknown posterior distributions of the global minimum location \mathbf{x}^* [28], which is due to the imperfect knowledge of f described by the probabilistic model \hat{f} . Rather than finding the posterior distributions directly, TS generates functions from \hat{f} and minimizes them for samples of \mathbf{x}^* . Thus, we can consider TS a BO method that minimizes stochastic acquisition functions generated from the model posteriors for selecting new solution points. It is also worth noting that several information-theoretic optimization policies of BO, such as predictive entropy [17] and max-value entropy [18], compute their acquisition functions based on a set of samples generated by TS.

While TS naturally manages the exploitation–exploration trade-off, its randomness can reduce the role of exploitation in optimization. When the GP posterior exhibits high uncertainty at the beginning of the optimization process, TS prioritizes exploration as it can diversify the selection of new solutions when limited information has been gained about the optimal solution \mathbf{x}^* . As the number of observations increases and the GP posterior becomes concentrated, the algorithm transitions to exploiting knowledge about the true objective function. This exploitation strategy, however, is inferior due to the randomness of TS, motivating the quest for an intervention to improve its exploitation.

In this work, we incorporate the ϵ -greedy policy into TS to improve its exploitation. This policy is a selection strategy of reinforcement learning to address the tension between exploitation and exploration [29]. Given $\epsilon \in (0, 1)$, the policy chooses an action by either maximizing an average reward function with probability $1 - \epsilon$ (i.e., exploitation or greedy) or selecting it randomly with probability ϵ (i.e., exploration). The selection strategy is pure exploitation when $\epsilon = 0$, or pure exploration when $\epsilon = 1$. Similarly, our proposed approach implements the generic TS (with probability ϵ) for exploration and a new fashion of TS called sample-average TS (with probability $1 - \epsilon$) for exploitation.

Several works have explored the ϵ -greedy policy in BO and multi-armed bandit problems. De Ath et al. [30] proposed two schemes for applying the policy to BO. The first scheme performs exploration (with probability ϵ) by randomly selecting a point on the Pareto frontier obtained by simultaneously minimizing the posterior mean and maximizing the posterior standard deviation. The second scheme performs exploration by randomly selecting a point in the input variable space. Both schemes implement exploitation (with probability $1 - \epsilon$) by minimizing the posterior mean function. Jin et al. [31] introduced the so-called ϵ -exploring TS (ϵ -TS) to multi-armed bandit problems. Given the posterior distributions of arms, exploration (with probability ϵ) sets the estimated reward function for each arm as a random sample drawn from the associated posterior distribution, while exploitation (with probability $1 - \epsilon$) sets the estimated reward function as the sample mean function. Yet the performance of ϵ -greedy TS for BO remains unexplored.

Our contributions are as follows: (1) ϵ -greedy TS to BO, which is a simple, effective method to improve the exploitation of TS; and (2) empirical evaluations demonstrating that ϵ -greedy TS with an appropriate ϵ is more robust than its two extremes (i.e., the generic TS and the sample-average TS), matching or outperforming the better of the two across various problems.

The rest of this paper progresses as follows. Section 2 provides a background essential for the development of ϵ -greedy TS. Section 3 describes the method in detail. Section 4 presents the empirical

performance of ϵ -greedy TS on minimizing two benchmark functions and finding constitutive parameters for a steel cantilever beam. Finally, Sec. 5 concludes this paper.

2 Background

This section provides an overview of GP modeling (Sec. 2.1), the implementation of the generic TS for BO (Sec. 2.2), and a simple method to generate sample paths from GP posteriors for use of TS (Sec. 2.3).

2.1 Overview of Gaussian Processes. Let $\mathcal{D} = \{(\mathbf{X}, \mathbf{y})\} = \{(\mathbf{x}^i, y^i)\}_{i=1}^N$ denote a training dataset, where \mathbf{x}^i are observations of a d -dimensional vector of input variables and y^i the corresponding observations of the objective function. We wish to build from \mathcal{D} an observation model $y(\mathbf{x}) = f(\mathbf{x}) + \epsilon_n: \mathbb{R}^d \mapsto \mathbb{R}$, where $\epsilon_n \sim \mathcal{N}(0, \sigma_n^2)$ is additive zero-mean Gaussian noise with variance σ_n^2 . This observation noise is assumed to be independent and identically distributed.

A GP model assumes that any finite subset of an infinite set of objective function values has a joint Gaussian distribution [32]. Oftentimes this assumption is encoded using the following GP prior:

$$f(\cdot) \sim \mathcal{GP}(0, \kappa(\cdot, \cdot | \boldsymbol{\theta})) \quad (2)$$

where $\kappa(\mathbf{x}, \mathbf{x}' | \boldsymbol{\theta}) = \text{cov}[f(\mathbf{x}), f(\mathbf{x}')]: \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ is a positive semi-definite covariance function parameterized by a vector $\boldsymbol{\theta}$ of hyperparameters. Common-used covariance functions include squared exponential (SE), Matérn 3/2, Matérn 5/2, and automatic relevance determination squared exponential (ARD SE) [4,32]. By conditioning on \mathcal{D} the model prior given in Eq. (2) and utilizing the observation noise assumption, it is straightforward to show that the vector of observations $[y(\mathbf{x}^1), \dots, y(\mathbf{x}^N)]^\top$ is distributed according to an N -variate Gaussian with zero mean and covariance matrix $\mathbf{C} = \mathbf{K} + \sigma_n^2 \mathbf{I}_N$, where the (i, j) th element of \mathbf{K} is $\kappa(\mathbf{x}^i, \mathbf{x}^j | \boldsymbol{\theta})$ and \mathbf{I}_N the N -by- N identity matrix. By further applying the conditional multivariate Gaussian [33], we obtain the posterior predictive distribution of the objective function value at an unseen input variable vector \mathbf{x}_* , i.e., $p(f(\mathbf{x}_*) | \mathbf{x}_*, \mathcal{D}) = \mathcal{N}(\mu_f(\mathbf{x}_*), \sigma_f^2(\mathbf{x}_*))$, which encodes information about the model we wish to build. The mean and variance of the predictive distribution read

$$\mu_f(\mathbf{x}_*) = \mathbf{K}_*^\top \mathbf{C}^{-1} \mathbf{Y} \quad (3)$$

$$\sigma_f^2(\mathbf{x}_*) = \kappa(\mathbf{x}_*, \mathbf{x}_* | \boldsymbol{\theta}) - \mathbf{K}_*^\top \mathbf{C}^{-1} \mathbf{K}_* \quad (4)$$

where

$$\mathbf{K}_* = [\kappa(\mathbf{x}_*, \mathbf{x}^1 | \boldsymbol{\theta}), \dots, \kappa(\mathbf{x}_*, \mathbf{x}^N | \boldsymbol{\theta})]^\top \quad (5)$$

Training GP models involves finding an optimal set of hyperparameters $\boldsymbol{\theta}$ that ensures robust predictive performance of trained GPs across unseen input variables. Exact training methods emulating all possible hyperparameter hypotheses are often impractical, leading to the prevalent use of approximate training methods. These approximate training methods fall into two main categories: deterministic approximations and Monte Carlo methods [34]. Deterministic approximations include the maximum likelihood method identifying a set of hyperparameters that maximizes the parameter likelihood $p(\mathcal{D} | \boldsymbol{\theta})$, and Laplace's method approximating the posterior $p(\boldsymbol{\theta} | \mathcal{D})$ as a multivariate Gaussian with a posterior mean equal to the maximum likelihood estimation and a posterior covariance estimated from Laplace approximation. Monte-Carlo methods, such as importance sampling, the Metropolis, Gibbs sampling, and slice sampling, generate approximate samples of $\boldsymbol{\theta}$ from the posterior $p(\boldsymbol{\theta} | \mathcal{D})$ given the likelihood $p(\mathcal{D} | \boldsymbol{\theta})$ and a prior $p(\boldsymbol{\theta})$. Several of the aforementioned approximate training methods have been incorporated in reliable GP toolboxes such as DACE [35], GPML [36], GPstuff [37], pyGPs [38], and GPflow [39].

Algorithm 1 Generic Thompson sampling for Bayesian optimization

```

1: Input: input variable domain  $\mathcal{X}$ , number of initial observations  $N$ ,
   threshold for number of BO iterations  $K$ , standard deviation of obser-
   vation noise  $\sigma_n$ 
2: Generate  $N$  initial samples of  $\mathbf{x}$ 
3: for  $\{i = 1 : N\}$ 
4:    $y^i \leftarrow f(\mathbf{x}^i) + \varepsilon_n^i$ 
5: end for
6:  $\mathcal{D}^0 \leftarrow \{\mathbf{x}^i, y^i\}_{i=1}^N$ 
7:  $\{\mathbf{x}_{\min}, y_{\min}\} \leftarrow \min\{y^i, i = 1, \dots, N\}$ 
8: for  $k = 1 : K$  do
9:   Build a GP posterior  $\hat{f}^k(\mathbf{x}) | \mathcal{D}^{k-1}$ 
10:  Generate a sample path  $g(\mathbf{x} | \mathcal{D}^{k-1})$  from  $\hat{f}^k(\mathbf{x}) | \mathcal{D}^{k-1}$ 
11:   $\mathbf{x}^k \leftarrow \arg \min_{\mathbf{x}} g(\mathbf{x} | \mathcal{D}^{k-1})$  s.t.  $\mathbf{x} \in \mathcal{X}; \mathbf{x} \notin \mathcal{D}^{k-1}$ 
12:   $y^k \leftarrow f(\mathbf{x}^k) + \varepsilon_n^k$ 
13:   $\mathcal{D}^k \leftarrow \mathcal{D}^{k-1} \cup \{\mathbf{x}^k, y^k\}$ 
14:   $\{\mathbf{x}_{\min}, y_{\min}\} \leftarrow \min\{y_{\min}, y^k\}$ 
15: end for
16: return  $\{\mathbf{x}_{\min}, y_{\min}\}$ 

```

2.2 Bayesian Optimization Via Thompson Sampling. As briefly described in Sec. 1, the generic TS generates a sequence of solution points \mathbf{x}^k ($k = 1, \dots, K$) by randomly sampling from an unknown posterior distribution $p(\mathbf{x}^* | \mathcal{D}^{k-1})$ of the global minimum \mathbf{x}^* , where K represents a finite budget on the number of BO iterations. Leveraging the fact that \mathbf{x}^* is fully determined by the objective function, the generic TS follows two simple steps in each iteration. It first generates a sample path from the GP posterior $\hat{f}^k(\mathbf{x}) | \mathcal{D}^{k-1}$ for which the detailed implementation is described in Sec. 2.3. It then minimizes the generated sample path to find a minimum location \mathbf{x}^* . By doing so, \mathbf{x}^* , assigned as the new solution \mathbf{x}^k , is a sample from $p(\mathbf{x}^* | \mathcal{D}^{k-1})$. Algorithm 1 summarizes the implementation of the generic TS.

Figure 1 shows sample paths generated from the GP posterior of univariate function $f(x) = x \sin x$ for $x \in [0, 20]$ and an approximate distribution derived from their minimum locations. Specifically, we build from ten observations a GP model with zero mean and SE covariance function. From this GP model, we generate 50 sample paths, and five of them are shown in Fig. 1(a). Minimizing the generated sample paths and utilizing a kernel density estimation on the obtained solutions result in an approximate distribution $p(\mathbf{x}^* | \mathcal{D})$, as shown in Fig. 1(b). Note that the generic TS in Algorithm 1 generates and minimizes only one sample path in each iteration and does not attempt to approximate the posterior distribution of the minimum location. The following section describes how we can randomly generate sample path $g(\mathbf{x} | \mathcal{D}^{k-1})$ in Line 10 of Algorithm 1.

Algorithm 2 Generation of sample paths using random features

```

1: Input: dataset  $\mathcal{D}^{k-1}$ , type of stationary covariance function  $\kappa(\cdot, \cdot | \boldsymbol{\theta})$ ,
   number of spectral points  $N_p$ , standard deviation of observation
   noise  $\sigma_n$ 
2: Build a GP model from  $\mathcal{D}^{k-1}$ 
3: Formulate  $p(\mathbf{s})$  from the posterior covariance function  $\kappa(\cdot, \cdot | \boldsymbol{\theta})$ , see
   Table 1
4: for  $j = 1 : N_p$  do
5:    $[\mathbf{W}]_j \sim p(\mathbf{s})$ 
6:    $[\mathbf{b}]_j \sim \mathcal{U}[0, 2\pi]$ 
7: end for
8: Formulate  $\boldsymbol{\phi}(\mathbf{x})$ , see Eq. (7)
9:  $\boldsymbol{\Phi} \leftarrow [\boldsymbol{\phi}(\mathbf{x}_1), \dots, \boldsymbol{\phi}(\mathbf{x}_N)]^\top$ 
10: Compute  $\boldsymbol{\mu}_\beta$  and  $\boldsymbol{\Sigma}_\beta$ , see Eq. (9)
11:  $\boldsymbol{\beta} \sim \mathcal{N}(\boldsymbol{\mu}_\beta, \boldsymbol{\Sigma}_\beta)$ 
12: return  $g(\mathbf{x} | \mathcal{D}^{k-1}) \leftarrow \boldsymbol{\beta}^\top \boldsymbol{\phi}(\mathbf{x})$ 

```

2.3 Sampling From Gaussian Process Posteriors. Given the GP posterior $\hat{f}^k(\mathbf{x}) | \mathcal{D}^{k-1}$ characterized by a stationary covariance




Fig. 1 Sample paths from the GP posterior for $f(x) = x \sin(x)$ and distribution of their minimum locations: (a) GP predictions and five sample paths drawn from the GP posterior and (b) approximate conditional distribution $p(x^* | \mathcal{D})$ obtained from minimum locations of 50 sample paths

function, we follow the spectral sampling approach to generate a sample path $g(\mathbf{x} | \mathcal{D}^{k-1})$, see, e.g., Refs. [41] and [17]. This approach approximates the GP prior in Eq. (2) using a Bayesian linear model of randomized basis functions to avoid the computational cost due to exhaustive sampling from marginal distributions of the objective function values at finite sets of input locations, which scales cubically in the number of input locations. Such an approximation is rooted in Bochner's theorem that guarantees the existence of a Fourier dual $S(\mathbf{s})$ ($\mathbf{s} \in \mathbb{R}^d$) of the stationary covariance function, which is called spectral density when a finite non-negative Borel measure is interpreted as a distribution [42]. The spectral density functions associated with several covariance functions of Matérn class and ARD SE are listed in Table 1. Once $S(\mathbf{s})$ is determined, we can represent $\kappa(\mathbf{x}, \mathbf{x}' | \boldsymbol{\theta})$ using the following randomized feature map [41]

$$\kappa(\mathbf{x}, \mathbf{x}' | \boldsymbol{\theta}) = \boldsymbol{\phi}(\mathbf{x})^\top \boldsymbol{\phi}(\mathbf{x}') \quad (6)$$

where the feature map $\boldsymbol{\phi}(\mathbf{x}) \in \mathbb{R}^{N_p}$ can be approximated by Rahimi and Recht [41] and Hernández-Lobato et al. [17]

$$\boldsymbol{\phi}(\mathbf{x}) = \sqrt{2\kappa(0 | \boldsymbol{\theta})/N_p} \cos(\mathbf{W}\mathbf{x} + \mathbf{b}) \quad (7)$$

Here $\mathbf{W} \in \mathbb{R}^{N_p \times d}$ and $\mathbf{b} \in \mathbb{R}^{N_p}$ stack N_p spectral points generated from the normalized spectral density $p(\mathbf{s}) = S(\mathbf{s})/\kappa(0 | \boldsymbol{\theta})$ and N_p points drawn from the uniform distribution $\mathcal{U}[0, 2\pi]$, respectively. $\kappa(0 | \boldsymbol{\theta})$ is well defined because $\kappa(\mathbf{x}, \mathbf{x}' | \boldsymbol{\theta})$ is a stationary covariance function satisfying $\kappa(\mathbf{x}, \mathbf{x}' | \boldsymbol{\theta}) = \kappa(\|\mathbf{x} - \mathbf{x}'\| | \boldsymbol{\theta})$.

Figure 2 shows approximate SE and Matérn 5/2 covariance functions computed from Eq. (6) for different numbers of spectral point samples. We see that the approximate covariance functions converge to the true ones when the number of spectral points increases.

Once the approximate feature map is formulated, the GP prior as a kernel machine can be approximated by the following Bayesian

Table 1 Spectral density functions for common-used stationary covariance functions [40]

Covariance function	$\kappa(\mathbf{x}, \mathbf{x}' \boldsymbol{\theta})$	Spectral density function $S(\mathbf{s})$
SE	$\sigma_f^2 \exp\left(-\frac{1}{2} \frac{r^2}{l^2}\right)^a$	$\sigma_f^2 (\sqrt{2\pi})^d \exp\left(-\frac{1}{2} l^2 \mathbf{s}^\top \mathbf{s}\right)$
ARD SE	$\sigma_f^2 \exp\left(-\frac{1}{2} \sum_{i=1}^d \frac{(x_i - x'_i)^2}{l_i^2}\right)^b$	$\sigma_f^2 (\sqrt{2\pi})^d \left(\prod_{i=1}^d l_i\right) \exp\left(-\frac{1}{2} \sum_{i=1}^d l_i^2 s_i^2\right)$
Matérn 3/2	$\sigma_f^2 \left(1 + \frac{\sqrt{3}r}{l}\right) \exp\left(-\frac{\sqrt{3}r}{l}\right)^c$	$\sigma_f^2 \frac{2^d \pi^{d/2} \Gamma\left(\frac{d+3}{2}\right) 3^{3/2}}{\frac{1}{2} \sqrt{\pi} l^3} \left(\frac{3}{l^2} + \mathbf{s}^\top \mathbf{s}\right)^{-\frac{d+3}{2}}$
Matérn 5/2	$\sigma_f^2 \left(1 + \frac{\sqrt{5}r}{l} + \frac{5r^2}{3l^2}\right) \exp\left(-\frac{\sqrt{5}r}{l}\right)^c$	$\sigma_f^2 \frac{2^d \pi^{d/2} \Gamma\left(\frac{d+5}{2}\right) 5^{5/2}}{\frac{3}{4} \sqrt{\pi} l^5} \left(\frac{5}{l^2} + \mathbf{s}^\top \mathbf{s}\right)^{-\frac{d+5}{2}}$

^a $r = \sqrt{(\mathbf{x} - \mathbf{x}')^\top (\mathbf{x} - \mathbf{x}')}$ and $\boldsymbol{\theta} = [\sigma_f, l]^\top$, where $l > 0$ and $\sigma_f > 0$ are length scale and marginal standard deviation, respectively. ^b $\boldsymbol{\theta} = [\sigma_f, l_1, \dots, l_d]^\top$, where $l_i > 0$ ($i = 1, \dots, d$) is length scale for the i th input variable ^c $r = \sqrt{(\mathbf{x} - \mathbf{x}')^\top (\mathbf{x} - \mathbf{x}')}$ and $\boldsymbol{\theta} = [\sigma_f, l]^\top$ ^d $\Gamma(\cdot) =$ Gamma function.




Fig. 2 Approximate covariance functions using random features from different numbers of spectral point samples: (a) SE and (b) Matérn 5/2

linear model:

$$f(\mathbf{x}) \approx \boldsymbol{\beta}^\top \boldsymbol{\phi}(\mathbf{x}) \quad (8)$$

where the prior of $\boldsymbol{\beta}$ is $\mathcal{N}(\mathbf{0}, \mathbf{I}_{N_p})$. By further conditioning Eq. (8) on the data, we obtain the following mean and covariance of the posterior of $\boldsymbol{\beta}$ [17]

$$\boldsymbol{\mu}_\beta = (\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma_n^2 \mathbf{I}_{N_p})^{-1} \boldsymbol{\Phi}^\top \mathbf{y} \quad (9a)$$

$$\boldsymbol{\Sigma}_\beta = (\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma_n^2 \mathbf{I}_{N_p})^{-1} \sigma_n^2 \quad (9b)$$

where $\boldsymbol{\Phi} = [\boldsymbol{\phi}(\mathbf{x}_1), \dots, \boldsymbol{\phi}(\mathbf{x}_{N_p})]^\top \in \mathbb{R}^{N_p \times N_p}$. To this end, we can approximate the GP posterior using a Bayesian linear model that

weights the randomized basis functions using samples from the posterior of $\boldsymbol{\beta}$.

The generation of $g(\mathbf{x} | \mathcal{D}^{k-1})$ using random features is detailed in Algorithm 2. Note that Line 10 of Algorithm 2 involves the Cholesky decomposition of matrix $\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \sigma_n^2 \mathbf{I}_{N_p}$ to facilitate the calculation of its inverse. Alternative posterior sampling techniques such as the decoupled method [43], which generates GP posteriors by updating the corresponding GP prior samples, can be adopted to increase the sample path accuracy.

Algorithm 3 ϵ -greedy Thompson sampling for Bayesian optimization

- 1: **Input:** input variable domain \mathcal{X} , number of initial observations N , threshold for number of BO iterations K , number of spectral points N_p , number of sample paths N_s , value of $\epsilon \in (0, 1)$, standard deviation of observation noise σ_n
- 2: Generate N samples of \mathbf{x}
- 3: **for** $i = 1 : N$ **do**
- 4: $y^i \leftarrow f(\mathbf{x}^i) + \epsilon_n^i$
- 5: **end for**
- 6: $\mathcal{D}^0 \leftarrow \{\mathbf{x}^i, y^i\}_{i=1}^N$
- 7: $\{\mathbf{x}_{\min}, y_{\min}\} \leftarrow \min\{y^i, i = 1, \dots, N\}$
- 8: **for** $k = 1 : K$ **do**
- 9: Build a GP posterior $\hat{f}^k(\mathbf{x}) | \mathcal{D}^{k-1}$
- 10: Generate $r \sim \mathcal{U}[0, 1]$
- 11: **if** $r \leq \epsilon$ **then**
- 12: Sample $g(\mathbf{x} | \mathcal{D}^{k-1})$
- 13: **else**
- 14: **for** $s = 1 : N_s$ **do**
- 15: Sample $h^s(\mathbf{x} | \mathcal{D}^{k-1})$
- 16: **end for**
- 17: $g(\mathbf{x} | \mathcal{D}^{k-1}) \leftarrow \frac{1}{N_s} \sum_{s=1}^{N_s} h^s(\mathbf{x} | \mathcal{D}^{k-1})$
- 18: **end if**
- 19: $\mathbf{x}^k \leftarrow \underset{\mathbf{x}}{\operatorname{argmin}} g(\mathbf{x} | \mathcal{D}^{k-1})$ s.t. $\mathbf{x} \in \mathcal{X}; \mathbf{x} \notin \mathcal{D}^{k-1}$
- 20: $y^k \leftarrow f(\mathbf{x}^k) + \epsilon_n^k$
- 21: $\mathcal{D}^k \leftarrow \mathcal{D}^{k-1} \cup \{\mathbf{x}^k, y^k\}$
- 22: $\{\mathbf{x}_{\min}, y_{\min}\} \leftarrow \min\{y_{\min}, y^k\}$
- 23: **end for**
- 24: **return** $\{\mathbf{x}_{\min}, y_{\min}\}$

3 ϵ -Greedy Thompson Sampling

In this section, we first introduce another version of TS called sample-average TS (Sec. 3.1). We then describe ϵ -greedy TS using the sample-average and generic TS in the context of ϵ -greedy policy (Sec. 3.2).

3.1 Sample-Average Thompson Sampling. To enhance the exploitation of TS, we implement another version of TS called sample-average TS (or averaging TS) [44]. Specifically, we call Algorithm 2 to generate in Line 10 of Algorithm 1 a total of N_s sample paths $h^s(\mathbf{x} | \mathcal{D}^{k-1})$ ($s = 1, \dots, N_s$). These sample paths are used to define the following average sample path:

$$g(\mathbf{x} | \mathcal{D}^{k-1}) = \frac{1}{N_s} \sum_{s=1}^{N_s} h^s(\mathbf{x} | \mathcal{D}^{k-1}) \quad (10)$$

Another approach to generate an average sample path is to use a decoupled representation per ‘‘Matheron’s rule’’ [45], which can be computed at the same cost of generating one sample path, regardless of the nominal number N_s of sample paths being averaged.

We then minimize the average sample path $g(\mathbf{x} | \mathcal{D}^{k-1})$ to find a new solution in each iteration of BO. This approach differs from that of Balandat et al. [44] which generates sample paths using randomized quasi Monte-Carlo techniques.

The sample-average TS and the generic TS can be considered two extremes of TS. If $N_s = \infty$, $g(\mathbf{x} | \mathcal{D}^{k-1})$ in Eq. (10) is indeed the GP posterior mean in Eq. (3) whose minimum promotes exploitation. If $N_s = 1$, the sample-average TS recovers the generic TS that favors exploration. Thus, there exists an optimal state between these two extremes that corresponds to an unknown value of N_s where TS balances exploitation and exploration. However, we do not attempt to find such an optimal state in this work. We instead set N_s at a sufficiently large value, say $N_s = 50$, to enforce the exploitation of TS.

3.2 ϵ -greedy Thompson Sampling. The two distinct extremes of TS motivate the use of ϵ -greedy policy to randomly switch between them. Specifically, we implement the generic TS with probability ϵ to explore the input variable space. We invoke the sample-average TS with probability $1 - \epsilon$ to guide the search toward exploitation. We do not use the posterior mean in Eq. (3) for exploitation because we observe that computing its derivatives is more expensive than computing the derivatives of the average sample path in Eq. (10), which adds computational cost if optimization is performed via a gradient-based algorithm. It is also straightforward to recover the generic TS from ϵ -greedy TS by simply setting $\epsilon = 1$ or $N_s = 1$. Algorithm 3 details the proposed ϵ -greedy TS.

The proposed algorithm addresses the exploitation–exploration dilemma by the following factors: (1) ϵ —a small value of ϵ promotes exploitation and (2) N_s —a sufficiently large value of N_s encourages exploitation.

4 Numerical Examples

We test the empirical performance of ϵ -greedy TS on two challenging minimization problems of the 2d Ackley and 6d Rosenbrock functions (Sec. 4.1), and on the identification problem of a steel cantilever beam subjected to cyclic loading (Sec. 4.2).

4.1 Benchmark Functions. Consider minimization problems of the 2d Ackley and 6d Rosenbrock functions [47]. The analytical expressions for these functions and their global minimums are given as follows:

2d Ackley function

$$f(\mathbf{x}) = -a \exp \left(-b \sqrt{\frac{1}{2} \sum_{i=1}^2 x_i^2} \right) - \exp \left(\frac{1}{2} \sum_{i=1}^2 \cos(cx_i) \right) + a + \exp(1) \quad (11)$$

where $a = 20$, $b = 0.2$, and $c = 2\pi$. The function is evaluated on $\mathcal{X} = [-10, 10]^2$ and has a global minimum at $\mathbf{x}^* = [0, 0]^T$ with $f^* = f(\mathbf{x}^*) = 0$.

6d Rosenbrock function

$$f(\mathbf{x}) = \sum_{i=1}^5 [100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2] \quad (12)$$

where $\mathcal{X} = [-5, 10]^6$. The function has a global minimum at $\mathbf{x}^* = [1, \dots, 1]^T$ with $f^* = f(\mathbf{x}^*) = 0$.

While ϵ -greedy TS can be applied to any stationary covariance function listed in Table 1, we restrict our experiments to the ARD SE covariance function, assuming no specific prior knowledge about the characteristics of the objective functions. We call the DACE toolbox [35] to find optimal GP hyperparameters for each minimization problem. We set $\sigma_n = 10^{-3}$ for z -score standardized output observations and generate 10^3 spectral points to approximate the sample paths of GP posteriors. To minimize the generated sample paths, we use the DIRECT algorithm [48] whose output solution is set as the initial point of an interior-point optimizer. Parameters for the DIRECT algorithm including the function tolerance, maximum number of function evaluations, maximum number of iterations, and maximum number of rectangle divisions are 10^{-9} , 10^3d , 10^4d , and 10^4d , respectively. Parameters for the interior-point optimizer including the function tolerance, optimality tolerance, maximum number of iterations, maximum number of function evaluations, and termination tolerance on \mathbf{x} are 10^{-12} , 10^{-12} , 200, 5×10^3 , and 10^{-12} , respectively. We carry out all experiments for the benchmark functions using the Carya Cluster of the University of Houston. We compare the optimization results obtained from ϵ -greedy TS with those from other BO methods, including EI, the lower confidence bound (LCB), the averaging TS, and the generic TS.

To determine an appropriate value of N_s in Eq. (10) which trade-offs between the exploitation and computational cost, our preliminary experiments investigate how changes in N_s affect the optimization results for a fixed value of ϵ . We observe that when the number of sample paths reaches a sufficiently large value of $N_s = 50$, the solutions exhibit less sensitivity to changes in N_s .

To ensure fair comparisons, we randomly generate 100 initial datasets for each problem using Latin hypercube sampling [49], and start each BO method from each of these datasets. In each BO iteration, we record the best-found value of observation error $\log(y_{\min} - f^*)$ and the corresponding solution vector. Here y_{\min} and f^* represent the best observation of the objective function found in each iteration and its true minimum value, respectively. We set the number of initial observations and the number of BO iterations at $N = 5d$ and $K = 50$ for the 2d Ackley function, and $N = 10d$ and $K = 200$ for the 6d Rosenbrock function.

Figures 3(a) and 3(b) show the medians and interquartile ranges obtained from 100 runs of each BO method for the 2d Ackley and 6d Rosenbrock functions, respectively. Figures 3(c) and 3(d) show the medians and interquartile ranges of final solutions. In both cases, the optimization results from ϵ -greedy TS for an appropriate ϵ are better than those from one of its two extremes and competitive with the results from the other. In addition, ϵ -greedy TS for $\epsilon = 0.5$ can provide the best objective values among those from the considered BO methods.

While the optimization results in Fig. 3 suggest that it is safe to set $\epsilon = 0.5$ to balance exploitation and exploration, an optimal ϵ may depend on the nature of the problem of interest. A problem that requires more exploitation for its optimal solution (e.g., the 2d Ackley function) may benefit from a smaller ϵ . A problem with more exploration (e.g., the 6d Rosenbrock function) might be better addressed with a larger ϵ . Yet determining whether a black-box objective function leans toward exploitation or exploration for its extrema is often elusive. This, ironically, underscores the rationale behind employing ϵ -greedy TS to prevent our search from overly favoring either exploitation or exploration. In ϵ -TS for multi-armed bandits, ϵ can be set as $\epsilon = 1/N_a$ for achieving simultaneous minimax and asymptotic optimality [31], where N_a is the number of arms to select. This, however, does not apply to




Fig. 3 Performance of EI, LCB, averaging TS, generic TS, and ϵ -greedy TS methods for the 2d Ackley and 6d Rosenbrock functions. Optimization histories for (a) the 2d Ackley function and (b) the 6d Rosenbrock function. Medians and interquartile ranges of final solutions from 100 runs of each BO method for (c) the 2d Ackley function and (d) the 6d Rosenbrock function.

BO because the number of new solution points, which are similar to arms in multi-armed bandits, is uncountable. It is also worth noting that any endeavors to find an optimal value of ϵ for a particular problem should correspond to a specific accuracy level of the initial GP model because lowering the model fidelity always encourages exploration, which is irrespective of the selection of ϵ .

As discussed in Sec. 1, increasing ϵ value results in more exploration of the ϵ -greedy strategy, and this holds true for ϵ -greedy TS. An increase in ϵ encourages the algorithm to explore more unseen regions of the input variable space, which is confirmed by observing how the algorithm, with varying ϵ values, selects new solution points for the Ackley function, as shown in Fig. 4.

We also investigate the effect of varying ϵ on the runtime for selecting a new solution point. Figure 5 shows approximate distributions of runtime for selecting a new solution point by the sample-average TS with $N_s = 50$, ϵ -greedy TS with different ϵ values and $N_s = 50$, and the generic TS for the 2d Ackley and 6d Rosenbrock functions. There are two distinct clusters of the runtime values for ϵ -greedy TS. One cluster aligns with the runtime of the generic TS and the other corresponds to the runtime of the sample-average TS. In addition, the average runtimes for each BO trial of the sample-average TS, ϵ -greedy TS with $\epsilon = 0.1, 0.5, 0.9$, and the generic TS are 124, 102, 64, 18, and 7 s for the 2d Ackley function, respectively. Those for the 6d Rosenbrock functions are 1377, 1188, 809, 229, and 111 s, respectively. We see that increasing ϵ reduces the runtime of each BO trial. This can be attributed to the fact that a larger ϵ requires more exploration of calling the generic TS, which is cheaper than the sample-average TS.

4.2 Cantilever Beam. In this example, we use ϵ -greedy TS to identify the constitutive parameters for simulating the behavior of a steel cantilever beam under a static cyclic loading history RH1, as shown in Fig. 6. The web and flange of the beam have different material parameters. In the experiment conducted by Yamada and

Jiao [50], forced vertical displacement was applied at the free end of the beam. The deflection angle θ was defined as the ratio of the vertical tip displacement Δ mm to the beam length $L = 800$ mm, i.e., $\theta = \Delta/L$. Values of moment reaction M about the x -axis at the beam support were measured.

Let \mathbf{x} represent the vector of parameters underlying a constitutive model for the cyclic behavior of the beam, for which the detailed description is given in Sec. 4.2.1. Let M_t^s represent the simulated value of the moment reaction at the t th time-step of a discretized loading history of RH1 consisting of T steps, and M_t^m the corresponding measured value. We formulate the following misfit function $f(\mathbf{x})$ for finding an optimal set of \mathbf{x} [51]:

$$f(\mathbf{x}) = \sqrt{\frac{1}{T} \sum_{t=1}^T [M_t^s(\mathbf{x}) - M_t^m]^2} \quad (13)$$

To evaluate $f(\mathbf{x})$, we model the beam using ABAQUS 2022. We reduce the sensitivity of simulation estimates to the finite element (FE) mesh density by generating a fine mesh of 4960 nodes and 3510 linear hexahedral elements of type C3D8, see Fig. 6. The maximum increment size for each loading history is set as 0.01 s.

4.2.1 Nonlinear Combined Isotropic/Kinematic Hardening. Following Ref. [46], we capture the cyclic elastoplastic behavior of the beam using the nonlinear combined isotropic/kinematic hardening model [52]. This model combines the properties of both isotropic and nonlinear kinematic hardening to describe the relationship between strain and stress states at each time instant. This relationship is established based on the material status that, either elastic or plastic, is detected by the following von Mises yield condition:

$$F = \|\xi\| - \sqrt{\frac{2}{3}} \sigma_y \leq 0 \quad (14)$$




Fig. 4 Initial and added points for 2d Ackley function with $N_s = 50$ and different values of ϵ : (a) averaging TS, $\epsilon = 0$, (b) ϵ -greedy TS, $\epsilon = 0.1$, (c) ϵ -greedy TS, $\epsilon = 0.9$, and (d) generic TS, $\epsilon = 1$




Fig. 5 Approximate distributions of runtime for selecting a new solution point with different ϵ values and $N_s = 50$: (a) 2d Ackley and (b) 6d Rosenbrock

where $\xi = \text{dev}[\sigma] - \text{dev}[\alpha]$ is the shifted-stress tensor, $\text{dev}[\cdot]$ the deviatoric part of $[\cdot]$, $\|\cdot\|$ the L^2 -norm of the tensor, σ the point stress tensor, α the back-stress tensor, and σ_y the yield stress.

The isotropic hardening increases the size of the yield surface $F = 0$ during the evolution of plastic deformation, while fixing the shape and location of this surface. In this case, ξ in Eq. (14) does not involve the back-stress tensor α , resulting in an isotropic yield surface of the stress. Since structural steels exhibit a saturation point of the stress at large deformation, the isotropic hardening for them models the increment of yield surface size using the following monotonically increasing nonlinear function [53]:

$$\sigma_y = \sigma_{y,0} + Q_\infty [1 - \exp(-b\epsilon_{\text{eq}}^{\text{p}})] \quad (15)$$

where $\sigma_{y,0}$ is the initial yield stress, Q_∞ the difference of the stress saturation, $\sigma_{y,0}$, b the isotropic saturation rate, and $\epsilon_{\text{eq}}^{\text{p}}$ the current equivalent plastic strain determined based on its previous state and the associated plastic strain rate $\dot{\epsilon}_{\text{eq}}^{\text{p}}$.

Unlike the isotropic hardening, the kinematic hardening does not change the size and shape of the yield surface during the evolution of plastic deformation. Instead, it relocates the center of the yield surface by a rigid translation in the evolution direction of the plastic strain. This allows the model to capture the Bauschinger effect. In pure kinematic hardening, the yield function in Eq. (14) has σ_y fixed at $\sigma_{y,0}$ and the back-stress tensor α defined as the superposition of n_k back-stress components, such that [54]

$$\alpha = \sum_{k=1}^{n_k} \alpha_k \quad (16)$$

Here the evolution of the k th component α_k is described by a nonlinear kinematic hardening rule, as [55]

$$\dot{\alpha}_k = \sqrt{\frac{2}{3}} C_k \dot{\epsilon}_{\text{eq}}^{\text{p}} \mathbf{n} - \gamma_k \dot{\epsilon}_{\text{eq}}^{\text{p}} \alpha_k \quad (17)$$




Fig. 6 The cantilever beam, its FE mesh, and loading history RH1 for cyclic test [46].

Table 2 Material parameter intervals for the cantilever beam

Component	Parameter	Lower bound	Upper bound
Web	E (GPa)	175.05	—
	ν	0.3	—
	$\sigma_{y,0}$ (MPa)	300	340
	Q_∞ (MPa)	10	100
	b	5	30
	C_1 (MPa)	2×10^3	10^4
	γ_1	10	200
Flange	E (GPa)	175.05	—
	ν	0.3	—
	$\sigma_{y,0}$ (MPa)	270	300
	Q_∞ (MPa)	10	100
	b	5	30
	C_1 (MPa)	2×10^3	10^4
	γ_1	10	200

where $\mathbf{n} = \boldsymbol{\xi} / \|\boldsymbol{\xi}\|$, and C_k and γ_k the translation and relaxation rates of the back-stress component k , respectively.

The nonlinear combined isotropic/kinematic hardening model makes use of Eqs. (15)–(17) for checking the material status at each time instant. If we use one back-stress component for this model, its underlying parameters read $\mathbf{x} = [E, \nu, Q_\infty, b, \sigma_{y,0}, C_1, \gamma_1]^T$,

where Young's modulus E and Poisson's ratio ν are incorporated in the isotropic elastic tensor.

4.2.2 Constitutive Parameters by Bayesian Optimization. During the minimization of misfit function $f(\mathbf{x})$, we keep Young's modulus and Poisson's ratio for the web and flange of the beam




Fig. 7 Performance of EI, LCB, averaging TS, generic TS, and ϵ -greedy TS methods for the cantilever beam: (a) optimization histories and (b) medians and interquartile ranges of final solutions from ten runs of each BO method




Fig. 8 Measured and simulated M - θ curves of the cantilever beam under different cyclic loading histories. The simulated M - θ curves correspond to the best parameter set identified from ε -TS and the experimental results from RH1: (a) RH1, (b) RH2, and (c) RH3. Here experimental data obtained from RH2 and RH3 are unseen by the identification process.

as constants, set at $E = 175.05$ GPa and $\nu = 0.3$ [50]. As a result, we identify ten parameters for the beam of which five parameters are associated with the web and the remaining five with the flange. The permissible range for each parameter is taken from Ref. [46] and listed in Table 2. To initiate BO, we generate ten distinct initial datasets, half containing 50 data points each and the remainder 100 data points each. With a total of 100 BO iterations planned, we execute the optimization process using the Carya Cluster housed at the University of Houston. This cluster hosts 9984 Intel CPU cores and 327680 Nvidia GPU cores integrated within 188 compute and 20 GPU nodes. This infrastructure allows us to complete a costly beam simulation in 3 min.

Figure 7 shows the optimization histories and variations in the final solution obtained from different BO methods for the cantilever beam, where f_{\min} represents the best observation of $f(\mathbf{x})$ found in each BO iteration. The identification results obtained from ε -greedy TS with $\varepsilon = 0.1$ and 0.5 are comparable to or better than that from the sample-average TS, and are better than that from the generic TS. More notably, ε -greedy TS for $\varepsilon = 0.5$ provides the best set of identified parameters.

We use the best set of parameters identified from ε -greedy TS to predict the M - θ curves of the beam under RH1 and other two cyclic loading histories RH2 and RH3. The experimental data associated with RH2 and RH3 were not fed to the parameter identification process. The agreement observed between the measured and simulated M - θ curves for each loading history, as shown in Fig. 8, confirms the reliable prediction performance achieved with the identified parameters.

5 Conclusions

We introduced ε -greedy TS to optimizing costly simulation-based objective functions. The method addresses the classic exploitation-exploration dilemma by randomly switching between its two extremes, namely the generic TS and the sample-average TS. Our empirical findings reveal that ε -greedy TS with an appropriate ε value robustly achieves the better performance of its two extremes across the benchmark problems, and can provide the best result in a practical engineering application.

While several ε -greedy algorithms and the generic TS are guaranteed to converge eventually [4,30], we look forward to a theoretical analysis to elucidate the convergence properties of ε -greedy TS. The regret bound analysis of ε -TS [31] for multi-armed bandits can serve as a starting point for this endeavor. We are also interested in exploring extensions of ε -greedy TS to high-dimensional settings where good approximations of the GP posteriors and efficient optimization of TS acquisition functions in high dimensions are of interest. For this, subspace-based approaches [56,57] or trust-region methods [58] present viable strategies. Additionally, the application of ε -greedy TS to non-Gaussian likelihood settings and the evolution of ε during the optimization process are still open problems.

The former may be addressed using a Markov chain Monte Carlo algorithm (see, e.g., Refs. [59,60]) to sample a large set of function values or to sample the weight values β in Eq. (8) from a non-Gaussian posterior.

Acknowledgment

We thank the University of Houston for providing startup fund to support this research. We thank Professor Satoshi Yamada at the University of Tokyo and Professor Makoto Ohsaki at Kyoto University for providing experimental results of the cantilever beam used in Sec. 4.2.

Conflict of Interest

There are no conflicts of interest.

Data Availability Statement

Data provided by a third party listed in Acknowledgment.

References

- [1] Snoek, J., Larochelle, H., and Adams, R. P., 2012, "Practical Bayesian Optimization of Machine Learning Algorithms," *Advances in Neural Information Processing Systems*, Vol. 25, Dec. 3–8, Curran Associates, Inc., Lake Tahoe, NV, pp. 2951–2959.
- [2] Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N., 2016, "Taking the Human out of the Loop: A Review of Bayesian Optimization," *Proc. IEEE*, **104**(1), pp. 148–175.
- [3] Frazier, P. I., 2018, "Bayesian Optimization," *Recent Advances in Optimization and Modeling of Contemporary Problems*, INFORMS Tutorials in Operations Research, INFORMS, pp. 255–278.
- [4] Garnett, R., 2023, *Bayesian Optimization*, Cambridge University Press, Cambridge, UK.
- [5] Do, B., and Zhang, R., 2023, "Multi-fidelity Bayesian Optimization in Engineering Design," arXiv, <https://arxiv.org/abs/2311.13050>.
- [6] Zhang, R., and Alemazkoor, N., 2024, "Multi-fidelity Machine Learning for Uncertainty Quantification and Optimization," *J. Machine Learn. Model. Comput.*, **5**(4), pp. 77–94.
- [7] Karandikar, J., Chaudhuri, A., No, T., Smith, S., and Schmitz, T., 2022, "Bayesian Optimization for Inverse Calibration of Expensive Computer Models: A Case Study for Johnson-Cook Model in Machining," *Manuf. Lett.*, **32**, pp. 32–38.
- [8] Kuhn, J., Spitz, J., Sonnweber-Ribic, P., Schneider, M., and Böhlke, T., 2022, "Identifying Material Parameters in Crystal Plasticity by Bayesian Optimization," *Optim. Eng.*, **23**(3), pp. 1489–1523.
- [9] Tran, A., Tran, M., and Wang, Y., 2019, "Constrained Mixed-Integer Gaussian Mixture Bayesian Optimization and Its Applications in Designing Fractal and Auxetic Metamaterials," *Struct. Multidiscipl. Optim.*, **59**(6), pp. 2131–2154.
- [10] Zhang, Y., Apley, D. W., and Chen, W., 2020, "Bayesian Optimization for Materials Design With Mixed Quantitative and Qualitative Variables," *Sci. Rep.*, **10**(1), p. 4924.
- [11] Zheng, H., Xie, F., Ji, T., Zhu, Z., and Zheng, Y., 2020, "Multifidelity Kinematic Parameter Optimization of a Flapping Airfoil," *Phys. Rev. E*, **101**(1), p. 013107.

- [12] Greenhill, S., Rana, S., Gupta, S., Vellanki, P., and Venkatesh, S., 2020, "Bayesian Optimization for Adaptive Experimental Design: A Review," *IEEE Access*, **8**, pp. 13937–13948.
- [13] Roussel, R., Hanuka, A., and Edelen, A., 2021, "Multiobjective Bayesian Optimization for Online Accelerator Tuning," *Phys. Rev. Accel. Beams*, **24**(6), p. 062801.
- [14] Hennig, P., Osborne, M. A., and Kersting, H. P., 2022, *Probabilistic Numerics: Computation as Machine Learning*, Cambridge University Press, Cambridge, UK.
- [15] Villemonteix, J., Vazquez, E., and Walter, E., 2009, "An Informational Approach to the Global Optimization of Expensive-to-Evaluate Functions," *J. Global Optim.*, **44**(4), pp. 509–534.
- [16] Hennig, P., and Schuler, C. J., 2012, "Entropy Search for Information-Efficient Global Optimization," *J. Mach. Learn. Res.*, **13**(6), pp. 1809–1837.
- [17] Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z., 2014, "Predictive Entropy Search for Efficient Global Optimization of Black-Box Functions," *Advances in Neural Information Processing Systems*, Vol. 27, Dec. 8–13, Curran Associates, Inc., Montreal, Canada, pp. 918–926.
- [18] Wang, Z., and Jegelka, S., 2017, "Max-Value Entropy Search for Efficient Bayesian Optimization," *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70, Aug. 6–11, PMLR, Sydney, pp. 3627–3635.
- [19] Jones, D. R., Schonlau, M., and Welch, W. J., 1998, "Efficient Global Optimization of Expensive Black-Box Functions," *J. Global Optim.*, **13**(4), pp. 455–492.
- [20] Söbester, A., Leary, S. J., and Keane, A. J., 2005, "On the Design of Optimization Strategies Based on Global Response Surface Approximation Models," *J. Global Optim.*, **33**(1), pp. 31–59.
- [21] Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M., 2010, "Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design," *Proceedings of the 27th International Conference on Machine Learning*, Vol. 13, June 21–24, Omnipress, Haifa, pp. 1015–1022.
- [22] Frazier, P. I., Powell, W. B., and Dayanik, S., 2008, "A Knowledge-Gradient Policy for Sequential Information Collection," *SIAM J. Control Optim.*, **47**(5), pp. 2410–2439.
- [23] Blanchard, A., and Sapsis, T., 2021, "Bayesian Optimization With Output-Weighted Optimal Sampling," *J. Comput. Phys.*, **425**, p. 109901.
- [24] Thompson, W. R., 1933, "On the Likelihood That One Unknown Probability Exceeds Another in View of the Evidence of Two Samples," *Biometrika*, **25**(3/4), pp. 285–294.
- [25] Chapelle, O., and Li, L., 2011, "An Empirical Evaluation of Thompson Sampling," *Advances in Neural Information Processing Systems*, Vol. 24, Dec. 12–17, Curran Associates, Inc., Granada, Spain, pp. 2249–2257.
- [26] Russo, D. J., and Van Roy, B., 2014, "Learning to Optimize Via Posterior Sampling," *Math. Oper. Res.*, **39**(4), pp. 1221–1243.
- [27] Russo, D. J., Roy, B. V., Kazerouni, A., Osband, I., and Wen, Z., 2018, "A Tutorial on Thompson Sampling," *Found. Trends Mach. Learn.*, **11**(1), pp. 1–96.
- [28] Kandasamy, K., Krishnamurthy, A., Schneider, J., and Póczos, B., 2018, "Parallelised Bayesian Optimisation Via Thompson Sampling," *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, Vol. 84, Apr. 9–11, PMLR, Playa Blanca, Spain, pp. 133–142.
- [29] Sutton, R. S., and Barto, A. G., 2018, *Reinforcement Learning: An Introduction*, The MIT Press, Cambridge, MA.
- [30] De Ath, G., Everson, R. M., Rahat, A. A. M., and Fieldsend, J. E., 2021, "Greed is Good: Exploration and Exploitation Trade-Offs in Bayesian Optimisation," *ACM Trans. Evolutionary Learn. Optim.*, **1**(1), pp. 1–22.
- [31] Jin, T., Yang, X., Xiao, X., and Xu, P., 2023, "Thompson Sampling With Less Exploration is Fast and Optimal," *Proceedings of the 40th International Conference on Machine Learning*, *Proceedings of Machine Learning Research*, Vol. 202, July 23–29, PMLR, Honolulu, HI, pp. 15239–15261.
- [32] Rasmussen, C. E., and Williams, C. K. I., 2006, *Gaussian Processes for Machine Learning*, The MIT Press, Cambridge, MA.
- [33] Bishop, C., 2006, *Pattern Recognition and Machine Learning*, Springer, Berkeley, CA.
- [34] MacKay, D. J., 2003, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, Cambridge, UK.
- [35] Lophaven, S. N., Nielsen, H. B., and Søndergaard, J., 2002, "DACE-A Matlab Kriging Toolbox," Lyngby, Denmark, Technical Report IMMTR-2002-12.
- [36] Rasmussen, C. E., and Nickisch, H., 2010, "Gaussian Processes for Machine Learning (GPML) Toolbox," *J. Mach. Learn. Res.*, **11**(100), pp. 3011–3015.
- [37] Vanhatalo, J., Riihimäki, J., Hartikainen, J., Jylänki, P., Tolvanen, V., and Vehtari, A., 2013, "GPstuff: Bayesian Modeling With Gaussian Processes," *J. Mach. Learn. Res.*, **14**(1), pp. 1175–1179.
- [38] Neumann, M., Huang, S., Marthaler, D. E., and Kersting, K., 2015, "pyGPs – A Python Library for Gaussian Process Regression and Classification," *J. Mach. Learn. Res.*, **16**(80), pp. 2611–2616.
- [39] de G. Matthews, A. G., van der Wilk, M., Nickson, T., Fujii, K., Boukouvalas, A., León-Villagrà, P., Ghahramani, Z., and Hensman, J., 2017, "GPflow: A Gaussian Process Library Using TensorFlow," *J. Mach. Learn. Res.*, **18**(40), pp. 1–6.
- [40] Riutort-Mayol, G., Bürkner, P.-C., Andersen, M. R., Solin, A., and Vehtari, A., 2022, "Practical Hilbert Space Approximate Bayesian Gaussian Processes for Probabilistic Programming," *Stat. Comput.*, **33**(1), p. 17.
- [41] Rahimi, A., and Recht, B., 2007, "Random Features for Large-Scale Kernel Machines," *Advances in Neural Information Processing Systems*, Vol. 20, Dec. 3–6, Curran Associates, Inc., Vancouver, pp. 1177–1184.
- [42] Wendland, H., 2004, *Scattered Data Approximation*, Vol. 17, Cambridge University Press, Cambridge, UK.
- [43] Wilson, J., Borovitskiy, V., Terenin, A., Mostowsky, P., and Deisenroth, M., 2020, "Efficiently Sampling Functions From Gaussian Process Posteriors," *Proceedings of the 37th International Conference on Machine Learning*, *Proceedings of Machine Learning Research*, Vol. 119, July 13–18, PMLR, Virtual, Online, pp. 10292–10302.
- [44] Balandat, M., Karrer, B., Jiang, D., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E., 2020, "BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization," *Advances in Neural Information Processing Systems*, Vol. 33, Dec. 6–12, Curran Associates, Inc., Virtual, Online, pp. 21524–21538.
- [45] Adebisi, T., Do, B., and Zhang, R., 2024, "Optimizing Posterior Samples for Bayesian Optimization Via Rootfinding," arXiv.
- [46] Do, B., and Ohsaki, M., 2022, "Proximal-Exploration Multi-objective Bayesian Optimization for Inverse Identification of Cyclic Constitutive Law of Structural Steels," *Struct. Multidiscipl. Optim.*, **65**(7), p. 199.
- [47] Surjanovic, S., and Bingham, D., 2013, "Virtual Library of Simulation Experiments: Test Functions and Datasets."
- [48] Finkel, D. E., and Kelley, C. T., 2006, "Additive Scaling and the DIRECT Algorithm," *J. Global Optim.*, **36**(4), pp. 597–608.
- [49] Forrester, A. I. J., Söbester, A., and Keane, A., 2008, *Engineering Design Via Surrogate Modelling: A Practical Guide*, John Wiley & Sons, West Sussex, UK.
- [50] Yamada, S., and Jiao, Y., 2016, "A Concise Hysteretic Model of Structural Steel Considering the Bauschinger Effect," *Int. J. Steel Struct.*, **16**(3), pp. 671–683.
- [51] Ohsaki, M., Miyamura, T., and Zhang, J. Y., 2016, "A Piecewise Linear Isotropic-Kinematic Hardening Model With Semi-Implicit Rules for Cyclic Loading and Its Parameter Identification," *Comput. Model. Eng. Sci.*, **111**(4), pp. 303–333.
- [52] Lemaitre, J., and Chaboche, J.-L., 1994, *Mechanics of Solid Materials*, Cambridge University Press, Cambridge, UK.
- [53] Voce, E., 1948, "The Relationship Between Stress and Strain for Homogeneous Deformation," *J. Inst. Metals*, **74**, pp. 537–562.
- [54] Chaboche, J. L., and Rousselier, G., 1983, "On the Plastic and Viscoplastic Constitutive Equations—Part I: Rules Developed With Internal Variable Concept," *ASME J. Pressure Vessel. Technol.*, **105**(2), pp. 153–158.
- [55] Armstrong, P. J., and Frederick, C. O., 1966, "A Mathematical Representation of the Multiaxial Bauschinger Effect," Report RD/B/N731.
- [56] Nayebi, A., Munteanu, A., and Poloczek, M., 2019, "A Framework for Bayesian Optimization in Embedded Subspaces," *Proceedings of the 36th International Conference on Machine Learning*, Vol. 97, June 9–15, PMLR, Long Beach, CA, pp. 4752–4761.
- [57] Zhang, R., Mak, S., and Dunson, D., 2022, "Gaussian Process Subspace Prediction for Model Reduction," *SIAM J. Sci. Comput.*, **44**(3), pp. A1428–A1449.
- [58] Eriksson, D., Pearce, M., Gardner, J., Turner, R. D., and Poloczek, M., 2019, "Scalable Global Optimization Via Local Bayesian Optimization," *Advances in Neural Information Processing Systems*, Vol. 32, Dec. 8–14, Curran Associates, Inc., Vancouver, pp. 5496–5507.
- [59] Mazumdar, E., Pacchiano, A., Ma, Y., Jordan, M., and Bartlett, P., 2020, "On Approximate Thompson Sampling With Langevin Algorithms," *Proceedings of the 37th International Conference on Machine Learning*, *Proceedings of Machine Learning Research*, Vol. 119, July 13–18, PMLR, Virtual, Online, pp. 6797–6807.
- [60] Zheng, H., Deng, W., Moya, C., and Lin, G., 2024, "Accelerating Approximate Thompson Sampling With Underdamped Langevin Monte Carlo," *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics*, *Proceedings of Machine Learning Research*, Vol. 238, May 2–4, PMLR, Palau de Congressos, Spain, pp. 2611–2619.