



## OPEN ACCESS

## EDITED BY

Man Fai Leung,  
Anglia Ruskin University,  
United Kingdom

## REVIEWED BY

Xin-Lei Zhang,  
Institute of Mechanics (CAS), China  
Edgar O. Reséndiz-Flores,  
Tecnológico Nacional de México/IT  
Saltillo, Mexico

## \*CORRESPONDENCE

Mike Diessner  
m.diessner2@newcastle.ac.uk  
Kevin Wilson  
kevin.wilson@newcastle.ac.uk  
Richard D. Whalley  
richard.whalley@newcastle.ac.uk

## SPECIALTY SECTION

This article was submitted to  
Optimization,  
a section of the journal  
Frontiers in Applied Mathematics and  
Statistics

RECEIVED 21 October 2022

ACCEPTED 11 November 2022

PUBLISHED 08 December 2022

## CITATION

Diessner M, O'Connor J, Wynn A,  
Laizet S, Guan Y, Wilson K and  
Whalley RD (2022) Investigating  
Bayesian optimization for  
expensive-to-evaluate black box  
functions: Application in fluid  
dynamics.  
*Front. Appl. Math. Stat.* 8:1076296.  
doi: 10.3389/fams.2022.1076296

## COPYRIGHT

© 2022 Diessner, O'Connor, Wynn,  
Laizet, Guan, Wilson and Whalley. This  
is an open-access article distributed  
under the terms of the [Creative  
Commons Attribution License \(CC BY\)](#).  
The use, distribution or reproduction  
in other forums is permitted, provided  
the original author(s) and the copyright  
owner(s) are credited and that the  
original publication in this journal is  
cited, in accordance with accepted  
academic practice. No use, distribution  
or reproduction is permitted which  
does not comply with these terms.

# Investigating Bayesian optimization for expensive-to-evaluate black box functions: Application in fluid dynamics

Mike Diessner<sup>1\*</sup>, Joseph O'Connor<sup>2</sup>, Andrew Wynn<sup>2</sup>,  
Sylvain Laizet<sup>2</sup>, Yu Guan<sup>1</sup>, Kevin Wilson<sup>3\*</sup> and  
Richard D. Whalley<sup>4\*</sup>

<sup>1</sup>School of Computing, Newcastle University, Newcastle upon Tyne, United Kingdom, <sup>2</sup>Department of Aeronautics, Imperial College London, London, United Kingdom, <sup>3</sup>School of Mathematics, Statistics and Physics, Newcastle University, Newcastle upon Tyne, United Kingdom, <sup>4</sup>School of Engineering, Newcastle University, Newcastle upon Tyne, United Kingdom

Bayesian optimization (BO) provides an effective method to optimize expensive-to-evaluate black box functions. It has been widely applied to problems in many fields, including notably in computer science, e.g., in machine learning to optimize hyperparameters of neural networks, and in engineering, e.g., in fluid dynamics to optimize control strategies that maximize drag reduction. This paper empirically studies and compares the performance and the robustness of common BO algorithms on a range of synthetic test functions to provide general guidance on the design of BO algorithms for specific problems. It investigates the choice of acquisition function, the effect of different numbers of training samples, the exact and Monte Carlo (MC) based calculation of acquisition functions, and both single-point and multi-point optimization. The test functions considered cover a wide selection of challenges and therefore serve as an ideal test bed to understand the performance of BO to specific challenges, and in general. To illustrate how these findings can be used to inform a Bayesian optimization setup tailored to a specific problem, two simulations in the area of computational fluid dynamics (CFD) are optimized, giving evidence that suitable solutions can be found in a small number of evaluations of the objective function for complex, real problems. The results of our investigation can similarly be applied to other areas, such as machine learning and physical experiments, where objective functions are expensive to evaluate and their mathematical expressions are unknown.

## KEYWORDS

Bayesian optimization, Gaussian Process, black box function, computer simulation, fluid dynamics, turbulent drag reduction

## 1. Introduction

Expensive black box functions are functions characterized by high evaluation costs and unknown or non-existent mathematical expressions. These functions cannot be optimized with conventional methods as they often require gradient information and large numbers of function evaluations. Bayesian optimization (BO), however, is a specialized optimization strategy developed for this specific scenario. It can be used to locate global optima of objective functions that are expensive to evaluate and whose mathematical expressions are unknown [1, 2].

Recently, it has been applied successfully in the data-intensive field of computational fluid dynamics (CFD). Most notably, Talmnikar et al. [3] developed a Bayesian optimization framework for the parallel optimization of large eddy simulations. They used this framework (i) on a one-dimensional problem to determine the wave speed of a traveling wave that maximizes the skin-friction drag reduction in a turbulent channel flow and (ii) on a four-dimensional problem to find an efficient design for the trailing edge of a turbine blade that minimizes the turbulent heat transfer and pressure loss. For the former, Bayesian optimization was able to locate a wave speed to generate a skin-friction drag reduction of 60%, while for the latter within 35 objective function evaluations a design was found that reduced the heat transfer by 17% and the pressure loss by 21%. Mahfoze et al. [4] utilized Bayesian optimization on a four-dimensional problem to locate optimal low-amplitude wall-normal blowing strategies to reduce the skin-friction drag of a turbulent boundary-layer with a net power saving of up to 5%, within 20 optimization evaluations. Morita et al. [5] considered three CFD problems: the first two problems concerned the shape optimization of a cavity and a channel flow. The third problem optimized the hyperparameters of a spoiler-ice model. Lastly, Nabaie and Fukagata [6] maximized the skin-friction drag reduction in a turbulent channel flow by optimizing the velocity amplitude and the phase speed of a traveling wave-like wall deformation. They achieved a maximum drag reduction of 60.5%.

Apart from these experiments, Bayesian optimization has also been used to tune hyperparameters of statistical models. For example, Wu et al. [7] optimized the hyperparameters of a selection of models, such as random forests, convolutional neural networks, recurrent neural networks, and multi-grained cascade forests. In these low-dimensional cases (two or three parameters) the model performance increased and the time to find the solutions was reduced compared to manual search for each model.

While these examples show that Bayesian optimization performs well on a wide selection of problems, such as computer simulations and hyperparameter tuning of machine learning models, there is no extensive study on the many different types of Bayesian optimization algorithms in the literature. In particular, there is a research gap regarding the performance and the robustness of Bayesian optimization when applied to distinct challenges. This paper aims to address this gap by considering a wide range of algorithms and a wide range of problems with increasing levels of complexity that are, consequently, increasingly difficult to solve. The comparison and the thorough analysis of these algorithms can inform the design of Bayesian optimization algorithms and allows them to be tailored to the unique problem at hand. To show how the findings can guide the setup of Bayesian optimization, two novel simulations in the area of CFD are optimized, demonstrating that Bayesian optimization can find suitable solutions to complex problems in a small number of function evaluations. The results of these experiments show that Bayesian optimization is indeed capable of finding promising solutions for expensive-to-evaluate black box functions. In these specific cases, solutions were found that (a) maximize the drag reduction over a flat plate and (b) achieve drag reduction and net-energy savings (NES) simultaneously. For these experiments, the Newcastle University Bayesian Optimization (NUBO) framework, currently being developed to optimize fluid flow problems, is utilized.

The paper is structured as follows. In Section 2, the Bayesian optimization algorithm, in particular the underlying Gaussian Process (GP) and the different classes of acquisition functions, are reviewed. Section 3.1 discusses the synthetic benchmark functions and their advantages over other types of common benchmarking methods. In Section 3.2, four different sets of simulations are presented. These experiments study—Section 3.2.1 multiple analytical single-point acquisition functions, Section 3.2.2 varying numbers of initial training points, Section 3.2.3 acquisition functions utilizing Monte Carlo sampling, and Section 3.2.4 various multi-point acquisition functions. Finally, Section 3.3 discusses the main findings and relates them back to the problems presented in the introduction and in general. These findings are then used to inform the algorithm used to optimize the setup of two CFD experiments in Section 4. Lastly, a brief conclusion is drawn in Section 5.

---

Abbreviations: AUC, area under the curve; BO, Bayesian optimization; BUCB, Batched Upper Confidence Bound; CDF, cumulative distribution function; CFD, computational fluid dynamics; CL, Constant Liar; EI, Expected Improvement; ES, Entropy Search; GDR, global drag reduction; GP, Gaussian process; GP-UCB, Gaussian Process Upper Confidence Bound; MC, Monte Carlo; MES, Max-value Entropy Search; MEMS, Micro-Electro-Mechanical-Systems; MLE, maximum likelihood estimation; NES, net-energy savings; NUBO, Newcastle University Bayesian Optimization; PDF, probability density function; PES, Predictive Entropy Search; PI, Probability of Improvement; UCB, Upper Confidence Bound.

## 2. Bayesian optimization

Bayesian optimization aims to find the global optimum of an expensive-to-evaluate objective function, whose mathematical expression is unknown or does not exist, in a minimum number of evaluations. It first fits a surrogate model, most commonly a GP, to some initial training data. The GP reflects the current belief about the objective function and is used to compute heuristics, called acquisition functions. When optimized, these functions suggest the candidate point that should be evaluated next from the objective function. After the response for the new candidate point is computed from the objective function, it is added to the training data and the process is repeated until a satisfying solution is found or a predefined evaluation budget is exhausted [1]. The following sections provide an overview of the two main components of Bayesian optimization, the GP and the acquisition functions.

### 2.1. Gaussian Process

Consider an objective function  $f: \mathcal{X} \rightarrow \mathbb{R}$ , where  $\mathcal{X} \subset \mathbb{R}^d$  is a  $d$ -dimensional design space, that allows the evaluation of one  $d$ -dimensional point  $\mathbf{x}_i$  or multiple  $d$ -dimensional points  $\mathbf{X}$  within the design space  $\mathcal{X}$  yielding one or multiple observations  $f(\mathbf{x}_i) = y_i$  and  $f(\mathbf{X}) = \mathbf{y}$ , respectively. A popular choice for the surrogate model to represent such an objective function is a GP. A GP is a stochastic process for which any finite set of points can be represented by a multivariate Normal distribution. This so called prior is defined by a mean function  $\mu_0: \mathcal{X} \rightarrow \mathbb{R}$  and a positive definite covariance function or kernel  $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ . These induce a mean vector  $\mathbf{m}_i := \mu_0(\mathbf{x}_i)$  and a covariance matrix  $\mathbf{K}_{i,j} := k(\mathbf{x}_i, \mathbf{x}_j)$  that define the prior distribution of the GP as

$$\mathbf{f} \mid \mathbf{X} \sim \mathcal{N}(\mathbf{m}, \mathbf{K}) \quad (1)$$

where  $\mathbf{f} \mid \mathbf{X}$  are the unknown responses from the objective function given a collection of points  $\mathbf{x}_{1:n}$ . Given a set of observed design points  $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$  and a new point also called a candidate point  $\mathbf{x}$ , the posterior predictive distribution of  $\mathbf{x}$  can be computed:

$$Y \mid \mathcal{D}_n \sim \mathcal{N}(\mu_n(\mathbf{x}), \sigma_n^2(\mathbf{x})) \quad (2)$$

with the posterior mean and variance

$$\mu_n(\mathbf{x}) = \mu_0(\mathbf{x}) + \mathbf{k}(\mathbf{x})^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} (\mathbf{y} - \mathbf{m}) \quad (3)$$

$$\sigma_n^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x})^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}) \quad (4)$$

where  $\mathbf{K} = k(\mathbf{x}_{1:n}, \mathbf{x}_{1:n})$  is a square matrix and  $\mathbf{k}(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}_{1:n})$  is a vector containing the covariances between the new point  $\mathbf{x}$  and all design points  $\mathbf{x}_{1:n}$  [1, 2, 8].

For the experiments presented in this paper, the zero mean function is implemented as it fulfills the theoretical properties required to compute the variable hyperparameter  $\beta$  for the Gaussian Process Upper Confidence Bound (GP-UCB) algorithm as derived in Srinivas et al. [9] (see Section II.B.). While there are many covariance kernels used in the literature, this paper implements the Matérn kernel with  $\nu = 5/2$  with an individual length scale parameter for each input dimension as proposed by Snoek et al. [8] for practical optimization problems, as other kernels such as the squared-exponential kernel also known as the radial basis function kernel can be too smooth for applied problems such as physical experiments. This kernel requires an additional parameter, the output scale  $\tau^2$ , to allow the kernel to be scaled above a certain threshold by multiplying the output scale with the posterior variance  $\tau^2 \sigma_n^2(\mathbf{x})$  [2]. A nugget term  $\nu$  is added to the diagonal of the covariance matrix  $\mathbf{K}_{i,j} := k(\mathbf{x}_i, \mathbf{x}_j) + \nu I$  to take possible errors, measurement and other, into consideration where  $I$  is the identity matrix. Nuggets have been shown to improve the computational stability and the performance of the algorithm such as the coverage and robustness of the results when using sparse data [10]. Additionally, Gramacy and Lee [11] give a detailed reasoning why nuggets should be added to deterministic problems that exceeds solely technical advantages. For example, while there might not be a measurement error in a deterministic simulation, the simulation itself is biased as it only approximates reality. In such cases, it makes sense to include a model to take possible biases into account. The hyperparameters of the GP, in this case the nugget, the output scale and the length scales of the Matérn kernel, can be estimated directly from the data, for example using maximum likelihood estimation (MLE) or maximum a posteriori (MAP) estimation [1, 12]. Algorithm 1 in Section 2 of the [Supplementary material](#) shows how the hyperparameters can be estimated via MLE by maximizing the log marginal likelihood of the GP.

### 2.2. Acquisition functions

Acquisition functions utilize the posterior distribution of the GP, that is the surrogate model representing the objective function, to propose the next point to sample from the objective function. Shahriari et al. [1] group acquisition functions into four categories: improvement-based policies, optimistic policies, information-based policies, and portfolios. This section presents some popular representatives of each group, which we shall focus on in this paper. Pseudocode for all acquisition functions can be found in Section 2 of the [Supplementary material](#).

Improvement-based policies such as probability of improvement (PI) and expected improvement (EI) propose points that are better than a specified target, for example the best point evaluated so far, with a high probability. Given the

best point so far  $\mathbf{x}^*$  and its response value  $y^* = f(\mathbf{x}^*)$ , PI (5), and EI (6) can be computed as:

$$\alpha_{PI}(\mathbf{x}; \mathcal{D}_n) = \Phi\left(\frac{z}{\sigma_n(\mathbf{x})}\right) \quad (5)$$

$$\alpha_{EI}(\mathbf{x}; \mathcal{D}_n) = z\Phi\left(\frac{z}{\sigma_n(\mathbf{x})}\right) + \sigma_n(\mathbf{x})\phi\left(\frac{z}{\sigma_n(\mathbf{x})}\right) \quad (6)$$

respectively, where  $z = \mu_n(\mathbf{x}) - y^*$  and  $\Phi$  and  $\phi$  are the cumulative distribution function (CDF) and the PDF of the standard Normal distribution.

The Upper Confidence Bound (UCB) acquisition function is an optimistic policy. It is optimistic with regards to the uncertainty (variance) of the GP, that is UCB assumes the uncertainty to be true. In contrast to the improvement-based methods, UCB (7) has a tuneable hyperparameter  $\beta_n$  that balances the exploration-exploitation trade-off. This trade-off describes the decision the Bayesian optimization algorithm has to make at each iteration. The algorithm could either explore and select new points in an area with high uncertainty or it could exploit and select points in areas where the GP makes a high prediction [1]. Thus, choosing how to set the trade-off parameter  $\beta_n$  is an important decision. Common strategies for setting  $\beta_n$  include fixing it to a particular value or varying it, as in the GP-UCB algorithm [9].

$$\alpha_{UCB}(\mathbf{x}; \mathcal{D}_n) = \mu_n(\mathbf{x}) + \sqrt{\beta_n \sigma_n(\mathbf{x})} \quad (7)$$

Portfolios consider multiple acquisition functions at each step and choose the best-performing one. The Hedge algorithm outlined by Brochu et al. [13] is used in this research article. It requires the computation and optimization of all acquisition functions in the portfolio at each Bayesian optimization iteration and the assessment of their proposed new point compared to the posterior of the updated GP in retrospect. This means that portfolios are computationally costly. However, because the performance of different acquisition functions varies for each iteration, i.e., for some iterations, one acquisition function would be optimal, while another function would be preferred in a different iteration, the Hedge algorithm should in theory select the best acquisition function at each iteration, yielding a better solution than when just considering one individual acquisition function.

In addition to these analytical acquisition functions, Monte Carlo (MC) sampling can be implemented. While this is an approximation of the analytical functions, Monte Carlo sampling does not require the oftentimes strenuous explicit computations involved with the analytical functions, especially when considering multi-point approaches. Information-based policies, in particular those using entropy, aim to find a candidate point  $\mathbf{x}$  that reduces the entropy of the posterior distribution  $p(\mathbf{x}|\mathcal{D}_n)$ . While Entropy Search (ES) and Predictive

Entropy Search (PES) are computationally expensive, Wang and Jegelka [14] and Takeno et al. [15] introduce Max-value Entropy Search (MES), a method that uses information about simple to compute maximal response values instead of costly to compute entropies.

Furthermore, the analytical acquisition functions can be used with Monte Carlo sampling by reparameterizing Equations (5)–(7) [16]. Then, they can be computed by sampling from the posterior distribution of the GP instead of computing the acquisition functions directly:

$$\alpha_{MC PI}(\mathbf{x}; \mathcal{D}_n) = \sigma\left(\frac{\max(\mu_n(\mathbf{x}) + \mathbf{Lz}) - y^*}{\tau}\right) \quad (8)$$

$$\alpha_{MC EI}(\mathbf{x}; \mathcal{D}_n) = \max(0, \max(\mu_n(\mathbf{x}) + \mathbf{Lz}) - y^*) \quad (9)$$

$$\alpha_{MC UCB}(\mathbf{x}; \mathcal{D}_n) = \max(\mu_n(\mathbf{x}) + \sqrt{\beta_n/2}|\mathbf{Lz}|) \quad (10)$$

where  $\mathbf{z}$  is a vector containing samples from a Standard Normal distribution  $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$ , and  $\mathbf{L}$  is the lower triangular matrix of the Cholesky decomposition of the covariance matrix  $\mathbf{K} = \mathbf{LL}^T$ . The softmax function  $\sigma$  with the so-called temperature parameter  $\tau$  enables the reparameterization of PI by approximating necessary gradients that are otherwise almost entirely equal to 0. As  $\tau \rightarrow 0$ , this approximation turns precise as described in Section 2 of Wilson et al. [16].

It is straightforward to extend the reparameterizations of the analytical acquisition functions to multi-point methods that propose more than one point at each iteration of the Bayesian optimization algorithm. In this case, samples are taken from a joint distribution (including training points and new candidate points) that is optimized with respect to the new candidate points. This optimization can be performed jointly, i.e., optimizing the full batch of points simultaneously, or in a sequential manner where a new joint distribution is computed for each point in the batch including the previously computed batch points, but only optimizing it with respect to the newest candidate point [17]. While analytical functions cannot be extended naturally to the multi-point case, there exist some frameworks that allow the computation of batches. The Constant Liar (CL) framework uses EI in combination with a proxy target, usually the minimal or maximal response value encountered so far. After one iteration of optimizing EI, the lie is taken as the true response for the newly computed candidate point and the process is repeated until all points in the batch are computed. At this point, the true responses of the candidate points are evaluated via the objective function and the next batch is computed [18]. Similarly, the Gaussian Process Batched Upper Confidence Bound (GP-BUCB) approach leverages the fact that the variance of the multivariate Normal distribution can be updated based solely on the inputs of a new candidate point. However, to update the posterior mean, the response value is required. Gaussian Process Batched Upper Confidence Bound updates the posterior variance after each new candidate

point, while the posterior mean stays the same for the full batch. After all points of a batch are computed, the response values are gathered from the objective function and the next batch is computed [19].

## 3. Synthetic test functions

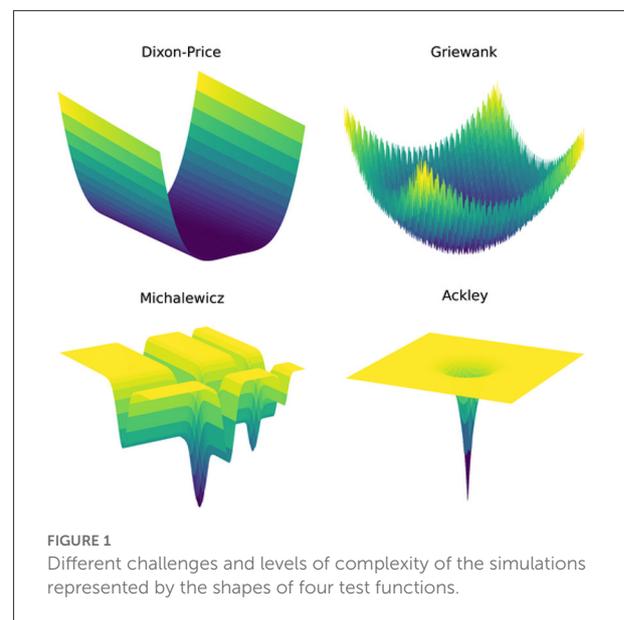
### 3.1. Functions

While there are various methods of benchmarking the performance of optimization algorithms, such as sampling objective functions from Gaussian Processes or optimizing parameters of models [8, 13, 14], this paper focuses on using synthetic test functions. Synthetic test functions have the main advantages that their global optima and their underlying shape are known. Thus, when testing the algorithms outlined in Section 2.2 on these functions, the distance from their results to the true optimum can be evaluated. For other benchmarking methods, this might not be possible as, for example, in the case of optimizing the hyperparameters of a model, the true optimum, i.e., the solution optimizing the performance of the model, is rarely known. Hence, results can only be discussed relative to other methods, ignorant of the knowledge of how close any method actually is to the true optimum. Furthermore, knowing the shape of the test function is advantageous as it provides information on how challenging the test function is. For example, knowing that a test function is smooth and has a single optimum indicates that it is less complex and thus less challenging than a test function that possesses multiple local optima, in each of which the algorithm potentially could get stuck.

This paper assesses the Bayesian optimization algorithms on eight different synthetic test functions spanning a wide range of challenges. Table 1 presents the test functions and gives details on the number of input dimensions, the shape, and the number of optima to give an idea about the complexity of the individual functions, while Figure 1 shows some of the functions that can be presented in three-dimensional space and indicates the increasing complexity of the functions. Detailed mathematical definitions for all test functions can be found in Section 3 of the Supplementary material. The number of input dimensions for the test functions was chosen to be equal to or slightly greater than the number in the simulations recently published in the fluid dynamics community, which is typically three to eight [3–6]. The 10D Sphere and 10D Dixon-Price functions are less challenging problems with a high degree of smoothness and only one optimum. They are therefore considered for a higher number of input dimensions. The 8D Griewank function adds a layer of complexity by introducing oscillatory properties. The 6D Hartmann function increases the level of difficulty further through its multi-modality. It has six optima with only one global optimum.

TABLE 1 Overview of the seven synthetic test functions.

Test function	Number of input dimensions D	Shape	Number of optima
Sphere	10	Bowl-shaped	1
Dixon-Price	10	Valley-shaped	1
Griewank	8	Oscillatory	1
Hartmann	6	Multi-modal	6
Noisy Hartmann	6	Noisy	6
Michalewicz	5	Steep edges	120
Ackley	6	Mostly flat	1



This function is also considered in two noisy variants to simulate typical measurement uncertainty encountered during experiments in fluid dynamics. In particular, the standard deviation of the added Gaussian noise is calculated so that it represents the measurement errors of state-of-the-art Micro-Electro-Mechanical-Systems (MEMS) sensors which directly measure time-resolved skin-friction drag in turbulent air flows (e.g., the flow over an aircraft); that is, 1.4–2.4% in an experimental setting [20]. Factoring in the range of the 6D Hartmann function, the corresponding standard deviations for the Gaussian noise, taken as a 99% confidence interval, are therefore 0.0155 and 0.0266, respectively. The most complex test functions are the 5D Michalewicz and the 6D Ackley functions (this paper considers a modified Ackley function with  $a = 20.0$ ,  $b = 0.5$ , and  $c = 0.0$ ). While the former has 120 optima and steep ridges, the latter is mostly flat with a single global minimum in the center of the space. The gradient of the function

close to the optimum and the large flat areas represent a great level of difficulty, as illustrated in the bottom row of [Figure 1](#).

## 3.2. Results

This section focuses on the results for four of the eight test functions that were considered in this paper. The results for the other four test functions, as well as more extensive tables, can be found in the [Supplementary material](#). The 8D Griewank function, the 6D Hartmann function in variations without noise and with the high noise level and the 6D Ackley function are examined here, as they represent increasing levels of complexity and come with unique challenges as illustrated in Section 3.1. Mirroring real world applications, a budget for the number of evaluations was imposed on the test functions: 200 evaluations for the Griewank and Hartmann functions and 500 for the more complex Ackley function. If not otherwise stated, the number of initial training points is equivalent to five points per input dimension of the given test function, i.e., 40 training points for the Griewank function and 30 training points for the Hartmann and Ackley functions. For each of the methods in the following sections, 50 different optimization runs were computed to investigate how robust and sensitive the methods are to varying initial training points. Each run was initialized with a different set of initial training points sampled from a Maximin Latin Hypercube Design [21]. However, the points were identical for all methods for a specific test function. All experiments were run on container instances on the cloud with the same specifications (two CPU cores and 8 GB of memory) to make runs comparable.

Overall, four different sets of experiments are considered. First, analytical single-point acquisition functions are compared. Second, the effects of a varying number of initial training points are investigated. Third, analytical methods are compared with Monte Carlo methods and, lastly, multi-point or batched methods are compared to the single-point results. As a baseline of performance, space filling designs exhausting the full evaluation budgets were sampled from a Maximin Latin Hypercube.

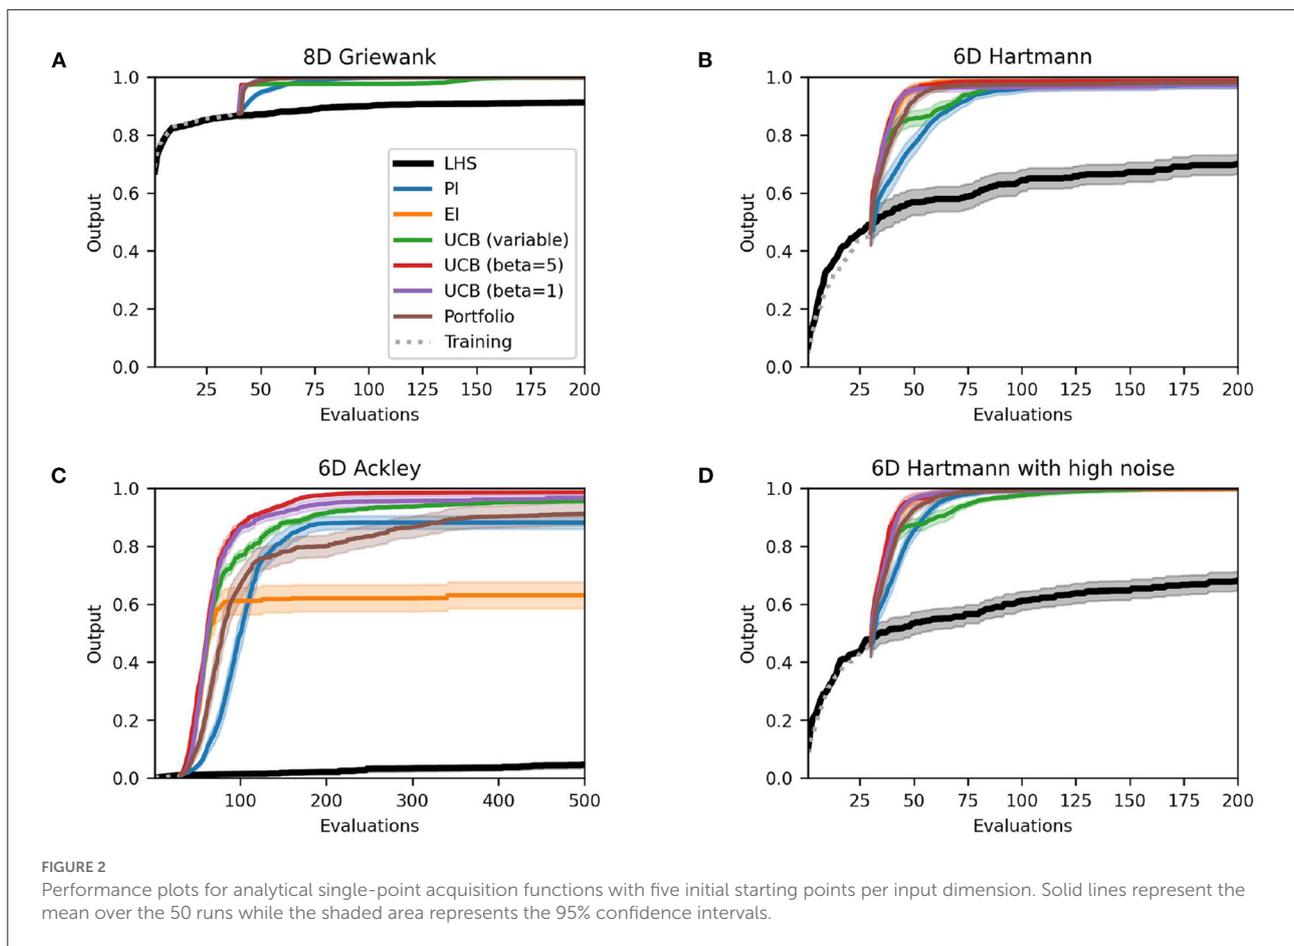
### 3.2.1. Analytical single-point acquisition functions

Section 2.2 describes four different groups of acquisition functions: improvement-based, optimistic, portfolios, and entropy-based. In this section, representative functions from the first three groups are tested on the synthetic test functions outlined above. The entropy-based approach is considered in Section 3.2.3. The focus lies on analytical single-point acquisition functions as they are widely used and thus present a natural starting point. Overall, seven different methods are considered: PI, EI, UCB with a variable and fixed  $\beta$  (5.0 and 1.0)

and a Hedge portfolio that combines the PI, EI, and variable UCB acquisition functions. For a detailed discussion of these functions see Section 2.2.

[Figure 2](#) presents the best solutions, defined by the output value closest to the global optimum, found so far at each evaluation. The outputs are normalized to the unit range where 1.0 represents the global optimum. Most methods perform very well on the Griewank function, all reaching 1.00, and both variations of the Hartmann function with and without added noise (all reaching 1.00 and  $>0.97$ , respectively). The acquisition functions all find the optimum or a solution fairly close to the optimum within the allocated evaluation budget. However, PI and variable UCB typically take more evaluations to find a solution close to the optimum. All acquisition functions perform noticeably better than the Latin Hypercube benchmark and also exhibit much less variation over the 50 runs, as indicated by the 95% confidence intervals. There is little difference between the Hartmann function with and without added noise, indeed the results are almost identical. The Ackley function on the other hand is more challenging. While the UCB methods still perform very well (all  $> 0.96$ ) the performance of the portfolio and PI decrease slightly (0.91 and 0.88, respectively) and EI performs considerably worse (0.63) with much more variability between runs.

Similar conclusions can be drawn from [Table 2](#), which presents the area under the curve (AUC) of each method. The AUC indicates how quickly the individual methods find solutions near the optimum. A perfect score (1.0) would indicate that the algorithm finds the optimum perfectly at the first iteration. The lower the score (a) the further the algorithm is away from the optimum and (b) the more evaluations are required for the algorithm to find promising solutions. While all methods score at least 0.98 for the Griewank function with a standard error of at most 0.01, the scores worsen slightly for the Hartmann function and significantly for the Ackley function. In particular, the AUC of EI for the Ackley function is very poor (0.59) and exhibits a high degree of variability between the 50 runs (standard error: 0.15). At the opposite end of the performance spectrum lie the optimistic methods, i.e., the UCB with variable and fixed  $\beta$ , that perform very well on all test functions with low standard errors. These results suggest that while the choice of acquisition function is less relevant for simple and moderately complex objective functions such as the Griewank or Hartmann functions, it is instrumental for solving challenging problems such as the Ackley function, especially if they are characterized by large flat areas. Here, the optimistic acquisition functions are advantageous and should be preferred over the Hedge portfolio and improvement-based approaches. It should be noted that this may only be true for the specific portfolio defined previously. Implementing a different collection of acquisition functions could yield different results.



### 3.2.2. Varying number of initial training points

The training points used to initialize the Bayesian optimization algorithm directly effect the surrogate model, i.e., the GP, that in turn is a representation of the objective function. A larger number of training points yields a GP that will typically represent the objective function more closely as it incorporates more points and thus more information. However, the more evaluations of the total budget are allocated for these initial training points, the fewer points can be evaluated as part of the Bayesian optimization algorithm. This trade-off indicates that the number of training points (and by extension their selection) is an important choice in Bayesian optimization and should be considered thoroughly. This section explores this trade-off by taking the same experimental setup as the previous section but varying the number of initial training points. Overall, setups with one, five, and ten initial points per input dimension of the objective function are considered.

Figures 3, 4 depict the performance plots for the case with one and ten training points per dimension, respectively. If the number of points is reduced to one point per dimension, the individual methods find solutions that are virtually identical

to the results with five training points per dimension (see [Supplementary material](#)). Expected improvement still performs much worse than other methods for the Ackley function. Furthermore, most methods seem to find their best solution in a comparable or just slightly higher number of evaluations than before, as the AUC values in [Table 3](#) show. This is expected, as using fewer initial training points means that the Bayesian optimization algorithm has less information at the earlier iterations than when using more initial training points. However, the difference in mean AUC is small and the results suggest that Bayesian optimization makes up for this lack of information quickly. PI and the Hedge portfolio have the highest decrease in performance on average for the Hartmann function with the mean AUC decreasing by 0.09 and 0.06, respectively. The variability between runs of the Hedge algorithm also rises, as indicated by a AUC standard error that is 0.08 higher. PI and the portfolio also perform worse for the Noisy Hartmann function, where the mean AUC decreased by 0.05 for both methods. Intuitively, this makes sense as the surrogate model includes less information than before and thus it takes more evaluations to find a good solution. In early iterations, the individual methods

TABLE 2 Averaged area under the curve with standard error for analytical single-point acquisition functions with five initial training points per input dimension.

Method	Griewank	Hartmann	Noisy Hartmann	Ackley
PI	0.99 (± 0.00)	0.91 (± 0.04)	0.95 (± 0.02)	0.76 (± 0.08)
EI	1.00 (± 0.00)	0.97 (± 0.02)	0.97 (± 0.02)	0.59 (± 0.15)
UCB (variable)	0.98 (± 0.01)	0.94 (± 0.03)	0.95 (± 0.03)	0.85 (± 0.01)
UCB ( $\beta = 5$ )	1.00 (± 0.00)	0.97 (± 0.02)	0.98 (± 0.02)	0.90 (± 0.02)
UCB ( $\beta = 1$ )	1.00 (± 0.00)	0.95 (± 0.04)	0.97 (± 0.02)	0.88 (± 0.05)
Hedge	1.00 (± 0.00)	0.95 (± 0.03)	0.97 (± 0.02)	0.77 (± 0.11)

deviate from one-another more than when five training points per dimension are used.

If the number of starting points is increased to 10 points per dimension, there is essentially no change to the case with just five starting points per dimension for the Griewank and the two Hartmann functions. For the Ackley function, however, the performance in terms of both the best solution found and the AUC worsen (Table 4) for EI and, most significantly, for PI. The average best solution for the latter decreases by 0.37 to only 0.51 while EI decreases by 0.10–0.53. Both policies struggle with the large area of the test function that gives the same response value and is hence flat (see Section 3.1). The optimistic policies and the portfolio perform much better and no real change is noticeable to results of Section 3.2.1.

These results suggest that choosing a larger number of training points to initialize the Bayesian optimization algorithm cannot necessarily be equated with better performance and solutions. This is particularly true considering that there was no improved performance when increasing from five to ten points per input dimension. On the other hand, reducing the number of training points did not yield results that were much worse. Overall, a similar picture as in the previous section emerges: While all methods perform well on simpler problems, optimistic policies achieve the best results on the more challenging problems independent of the number of starting points. For the test functions we considered, five training points per dimension appeared to be sufficient, with no discernible improvement when moving to 10 training points, and a small loss of performance when reducing to one training point per dimension.

### 3.2.3. Monte Carlo single-point acquisition functions

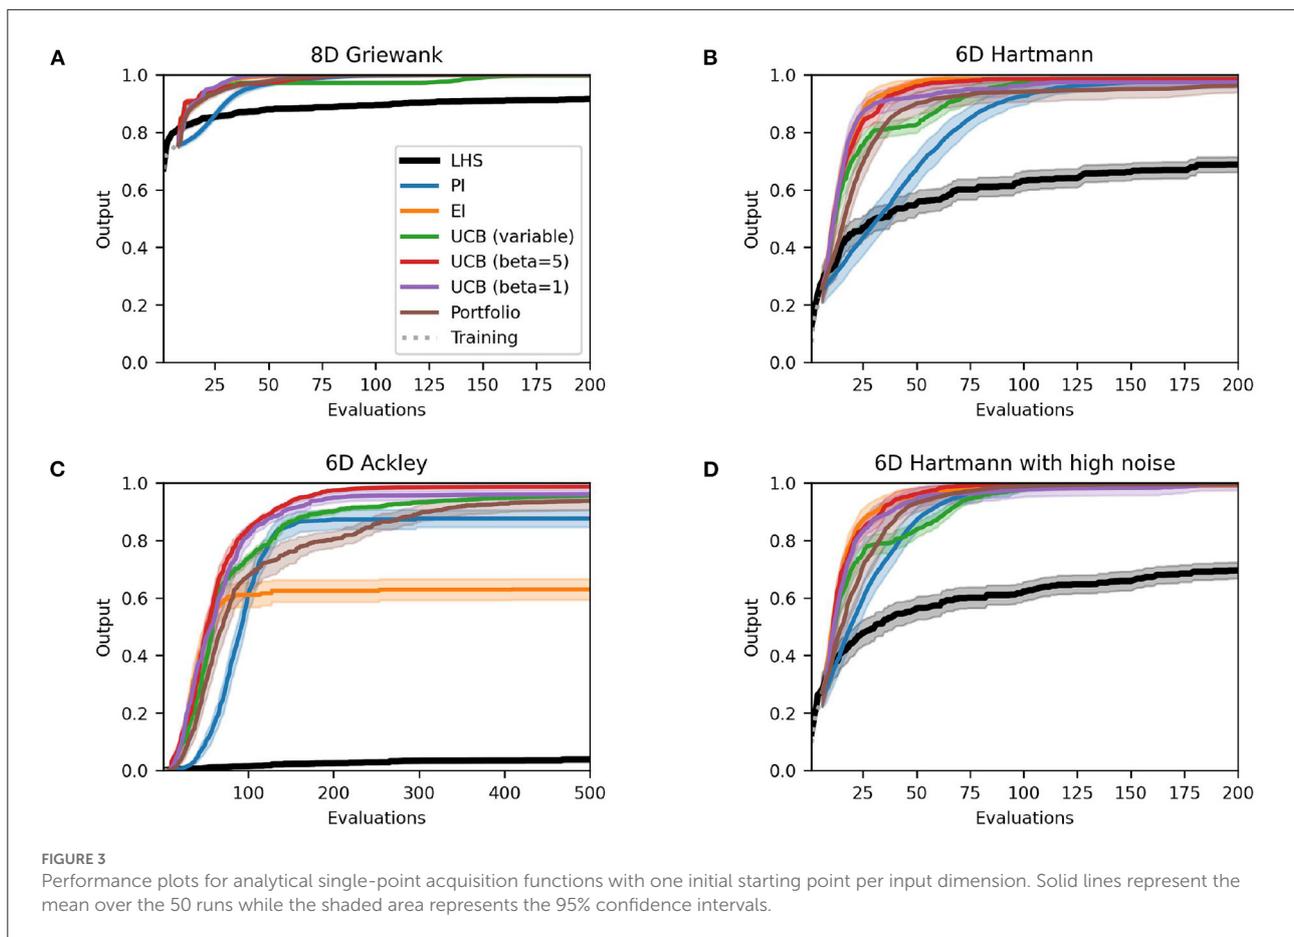
The previous experiments considered analytical acquisition functions. This section goes one step further and assesses the Monte Carlo approach outlined in Section 2.2. As not all acquisition functions can be rewritten to suit such an approach, the experiments are restricted to PI, EI, and UCB with variable and fixed  $\beta$ . Additionally, MES is introduced as a new method.

For the Griewank and both Hartmann functions all results are almost identical to the analytical case except for the variability between runs which increased slightly for some of the methods. However, Figure 5 clearly shows that the performance for PI and EI tested on the Ackley function decreased significantly. The average best solutions decreased by 0.68 and 0.56 and the average AUC (Table 5) decreased by 0.59 and 0.54, respectively. Table 5 shows that MES performs well on the Griewank and Hartmann functions, all reaching an AUC of above 0.97 with low standard errors. When it comes to the more complex Ackley function, however, MES performs much worse. With a mean AUC of 0.49 its performance ranks below the optimistic acquisition functions but still above the improvement-based methods. Earlier sections showed that improvement-based policies (in particular EI) perform poorly on the Ackley function when using analytical acquisition functions, but the results from this comparison show that their performance suffers even more severely when using the Monte Carlo approach. One reason for this could be that the Monte Carlo variants are essentially an approximation of the analytical acquisition functions as discussed in Section 2.2. However, there seems to be little to no change when using optimistic policies. This suggests that, similar to previous sections, optimistic policies should be preferred when optimizing complex and challenging objective functions with large flat areas using Monte Carlo acquisition functions.

### 3.2.4. Multi-point acquisition functions

The previous sections focused on single-point approaches, where each iteration of the Bayesian optimization algorithm yields one new point that is sampled from the objective function before the next iteration is started. While this makes sense for objective functions that can be evaluated quickly or do not allow parallel evaluations, it might slow down the optimization process needlessly when objective functions take a long time to evaluate and allow parallel evaluations. Thus, this section implements multi-point acquisition functions that propose a batch of points at each iteration, which are then evaluated simultaneously before the next iteration.

Section 2.2 outlined how Monte Carlo approaches using the reparameterization trick can be extended to compute batches naturally. Max-value Entropy Search does not use reparameterization and is thus not naturally extendable to the multi-point setting. Hence, we considered the same acquisition

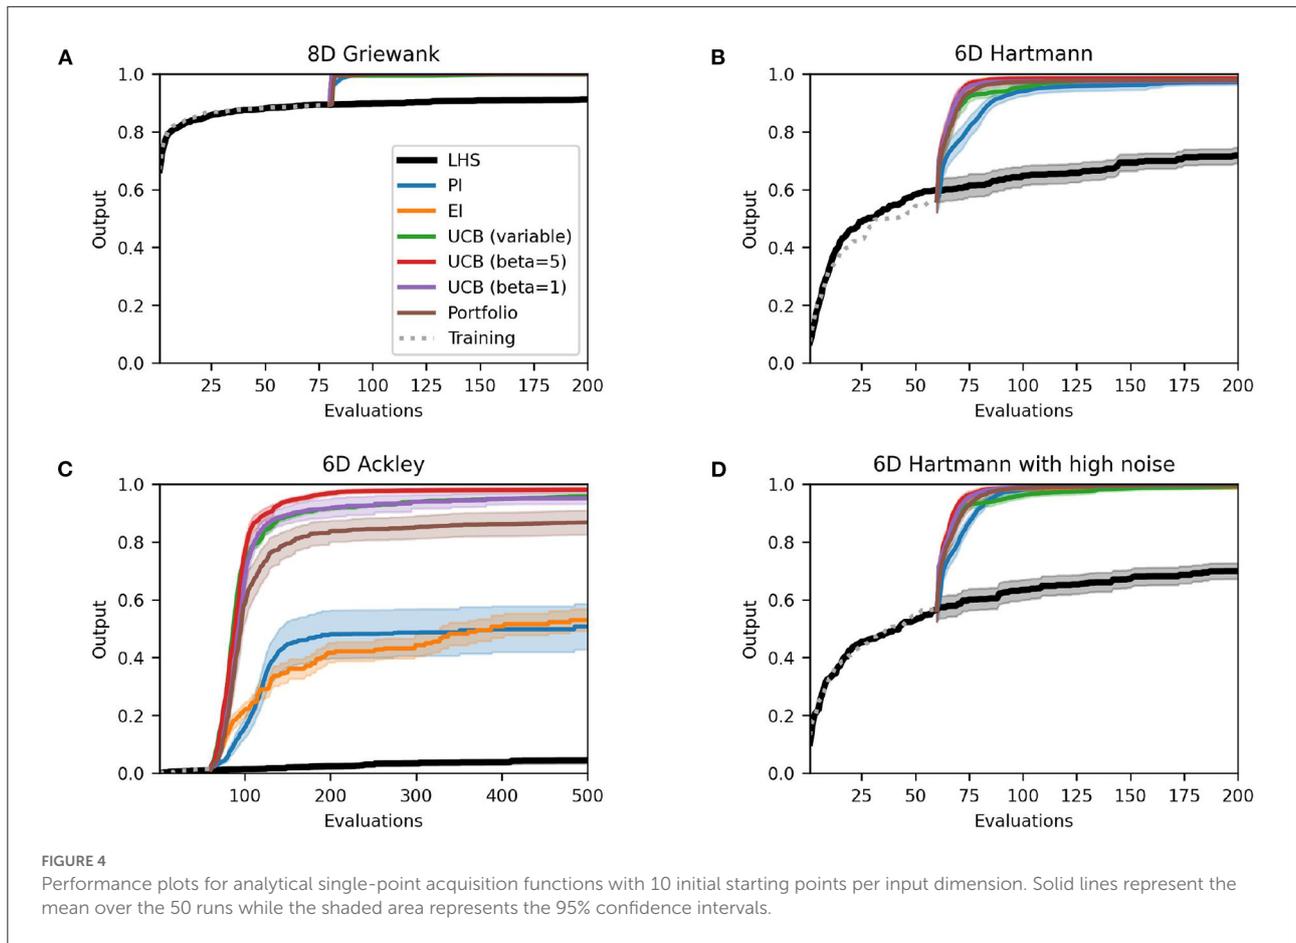


functions as in the previous section but this time for a batch size of five points. Each of the acquisition functions is optimized with two different methods, sequentially and jointly. The latter optimizes the acquisition function based on a joint distribution that includes training points and new points and optimizes all new points in a single step. The former recomputes the joint distribution each time a new point is found and only optimizes with respect to the newest point. For example, for a batch of five points this process is repeated five times [17]. This approach, also known as greedy optimization, might be preferable and can yield better results [22]. While analytical functions cannot be extended to the multi-point case as easily as the Monte Carlo evaluations, there are some frameworks that allow the computation of batches. This section considers the Constant Liar approach with a lie equivalent to the minimal (CL min) and maximal (CL max) value so far [18] and the GP-BUCB algorithm, an extension of the UCB function [19]. For more details see Section 2.2.

Figure 6 shows the results of some multi-point acquisition functions. They essentially find identical solutions, on average, to the single-point approach for the Griewank function and Hartmann function with and without noise. Only selected

methods are shown in the plot. Particularly, two variations of the UCB approach are not included as they are almost identical to the UCB with  $\beta = 1$  (see [Supplementary material](#) for results of all methods). While there are no changes to the AUC for the Griewank function, there are some differences for the Hartmann functions that suggest that methods require a different number of evaluations to find the best value, as Table 6 shows. Although the joint Monte Carlo approach for PI and UCB (variable and  $\beta = 5$ ) have a lower mean AUC, the decrease is less severe than for the GP-BUCB methods that all decrease by 0.08–0.10. The other methods perform comparably to the single-point case and the variability between the runs of the sequential and joint Monte Carlo UCB with  $\beta = 1$  even decreases. For the noisy Hartmann test function, the results are very similar to the Hartmann function without noise. In general, over all experiments there was no real drop in performance that can be declared as a result of adding noise up to 2.4% of the overall objective function range.

Most methods find similar best solutions to the Ackley test function as their single-point counterparts. However, there are some changes in the improvement-based policies: the sequential Monte Carlo EI and PI, and CL max, find solutions that are



**TABLE 3** Averaged area under the curve with standard error for analytical single-point acquisition functions with one initial training point per input dimension.

Method	Griewank	Hartmann	Noisy Hartmann	Ackley
PI	0.97 (± 0.01)	0.82 (± 0.09)	0.90 (± 0.06)	0.73 (± 0.11)
EI	0.99 (± 0.00)	0.95 (± 0.03)	0.95 (± 0.03)	0.58 (± 0.13)
UCB (variable)	0.97 (± 0.01)	0.91 (± 0.04)	0.91 (± 0.03)	0.82 (± 0.01)
UCB ( $\beta = 5$ )	0.98 (± 0.00)	0.94 (± 0.03)	0.95 (± 0.03)	0.87 (± 0.02)
UCB ( $\beta = 1$ )	0.99 (± 0.00)	0.93 (± 0.08)	0.93 (± 0.09)	0.85 (± 0.06)
Hedge	0.98 (± 0.01)	0.89 (± 0.11)	0.92 (± 0.05)	0.76 (± 0.09)

**TABLE 4** Averaged area under the curve with standard error for analytical single-point acquisition functions with 10 initial training points per input dimension.

Method	Griewank	Hartmann	Noisy Hartmann	Ackley
PI	1.00 (± 0.00)	0.93 (± 0.05)	0.96 (± 0.02)	0.43 (± 0.25)
EI	1.00 (± 0.00)	0.97 (± 0.02)	0.97 (± 0.02)	0.41 (± 0.07)
UCB (variable)	1.00 (± 0.00)	0.95 (± 0.03)	0.96 (± 0.03)	0.87 (± 0.01)
UCB ( $\beta = 5$ )	1.00 (± 0.00)	0.97 (± 0.02)	0.98 (± 0.02)	0.91 (± 0.03)
UCB ( $\beta = 1$ )	1.00 (± 0.00)	0.96 (± 0.03)	0.98 (± 0.03)	0.86 (± 0.08)
Hedge	1.00 (± 0.00)	0.96 (± 0.03)	0.97 (± 0.02)	0.78 (± 0.13)

better than before (by 0.33, 0.04, and 0.11, respectively). At the opposite end, CL min and joint Monte Carlo PI provide inferior solutions in the batched case (0.13 and 0.08, respectively).

Looking at how quickly the individual methods find good values on average, i.e., the AUC, it is clear that the analytical multi-point and joint Monte Carlo methods perform worse than

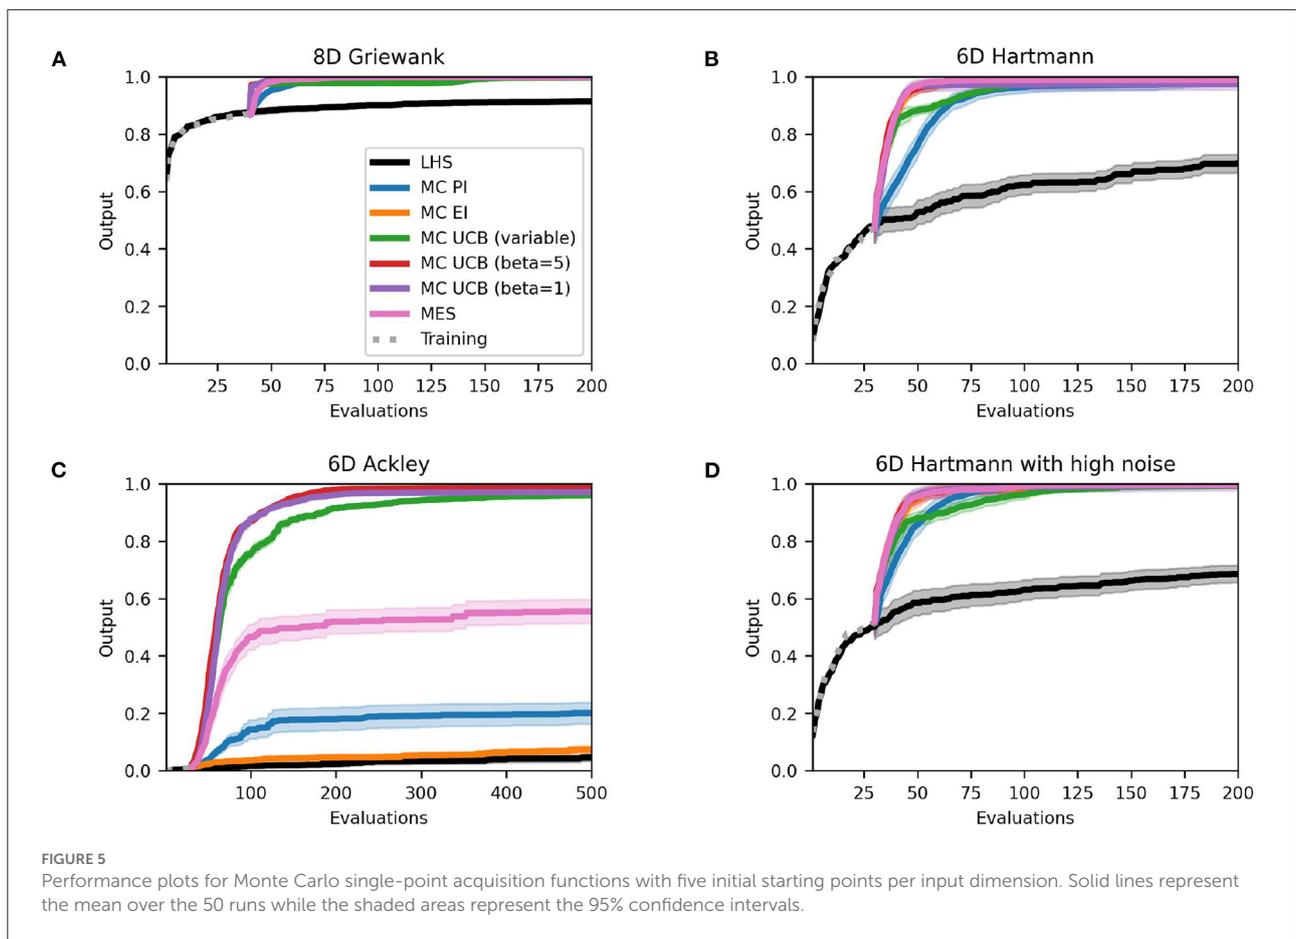


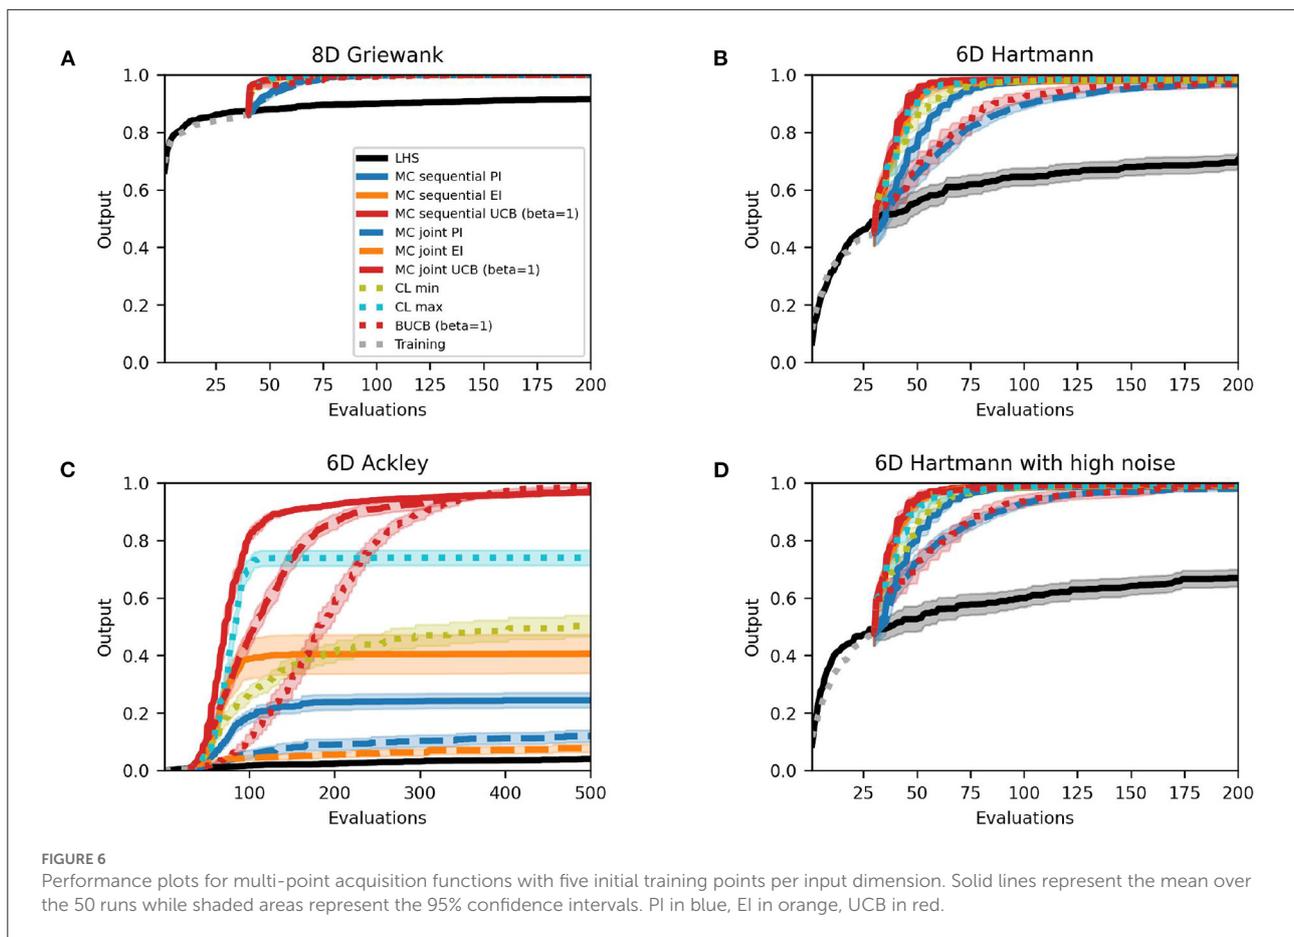
TABLE 5 Averaged area under the curve with standard error for Monte Carlo single-point acquisition functions with five initial training points per input dimension.

Method	Griewank	Hartmann	Noisy Hartmann	Ackley
MC PI	0.99 (± 0.00)	0.92 (± 0.06)	0.95 (± 0.04)	0.17 (± 0.12)
MC EI	1.00 (± 0.00)	0.97 (± 0.02)	0.97 (± 0.04)	0.05 (± 0.03)
MC UCB (variable)	0.98 (± 0.01)	0.95 (± 0.02)	0.95 (± 0.04)	0.86 (± 0.01)
MC UCB ( $\beta = 5$ )	1.00 (± 0.00)	0.97 (± 0.02)	0.98 (± 0.03)	0.91 (± 0.01)
MC UCB ( $\beta = 1$ )	1.00 (± 0.00)	0.96 (± 0.07)	0.97 (± 0.07)	0.89 (± 0.03)
MES	0.99 (± 0.00)	0.97 (± 0.02)	0.97 (± 0.03)	0.49 (± 0.12)

the single-point implementations for the Ackley function: the AUC of all analytical multi-point methods worsen by 0.19–0.24, except for CL max that improved by 0.08. Similarly, all joint

methods provide poorer performance than in the sequential single-point optimization where UCB with  $\beta = 1$  sees the largest drop of 0.11. Expected improvement is the exception and stays about the same. The biggest improvement is provided by the sequential Monte Carlo EI with an increase of 0.32. However, this improvement is mainly caused by the very poor performance of the single-point Monte Carlo EI acquisition function. Furthermore, this approach comes with a higher variability, as the standard error of the AUC rises by 0.19.

In terms of the best solutions found, multi-point methods present a good alternative to single-point acquisition functions. They generally find best solutions comparable to the single-point approach, but for a slightly larger number of objective function evaluations. However, while it requires more evaluations, the multi-point approach would still be faster when computing batches in parallel. Multi-point acquisition functions will, therefore, be most beneficial for expensive to evaluate objective functions that can be evaluated in parallel. With some exceptions, this benefit requires a higher number of evaluations until solutions comparable to single-point methods are found. Overall, the optimistic methods, combined with sequential optimization, appear to be favorable over the rest, as the red lines in Figure 6 clearly show.



### 3.3. Discussion

Six main conclusions can be drawn from the simulation results presented in detail above. This section discusses these findings and relates them back to the applied problems that motivated this work, i.e., optimizing experiments in engineering, particularly in fluid dynamics, and tuning hyperparameters of neural networks and other statistical models.

The first findings concern the choice of acquisition functions related to the complexity of the problem. The results show that this choice is less important when optimizing simpler objective functions. Improvement-based, optimistic, and information-based acquisition functions, as well as portfolios, performed well on a wide range of synthetic test functions with up to 10 input dimensions (see [Supplementary material](#) online for more examples to reinforce this result.). However, for more complex functions, such as the Ackley function, the optimistic methods performed significantly better than the rest and thus should be favored. For the Bayesian optimization algorithm, this means that all acquisition functions considered in this paper can yield good results. But the choice of acquisition function must be considered more carefully with increasing complexity

of the objective function. This indicates the importance of expert knowledge when applying Bayesian optimization to a specific problem, such as a drag reduction problem in fluid dynamics. Basing the choice of acquisition function on specific knowledge about the suspected complexity of the objective function could potentially increase the performance of the Bayesian optimization algorithm significantly. In cases where no expert knowledge or other information about the objective function is accessible, results suggest that the optimistic methods are a good starting point.

Secondly, the results suggest that the number of initial training points is not critical in achieving good performance from the algorithm. In fact, there is only a little difference when increasing the number of starting points from one point per input dimension to 10 points per dimension. While the performance of the acquisition functions differs initially for the former case, they still find comparable results to the runs with five or ten points per input dimension over all iterations. This means that Bayesian optimization explores the space efficiently even when provided with only a few training points. Deciding on the number of starting points is an important decision in problems where evaluating a point is expensive. For example,

TABLE 6 Averaged area under the curve with standard error for multi-point acquisition functions with five initial training points per input dimension.

Type	Method	Griewank	Hartmann	Noisy Hartmann	Ackley
Sequential Monte Carlo	PI	0.99 (± 0.00)	0.92 (± 0.03)	0.94 (± 0.02)	0.22 (± 0.08)
	EI	1.00 (± 0.00)	0.95 (± 0.03)	0.97 (± 0.02)	0.37 (± 0.22)
	UCB (variable)	0.98 (± 0.01)	0.94 (± 0.03)	0.94 (± 0.03)	0.84 (± 0.01)
	UCB ( $\beta = 5$ )	0.99 (± 0.00)	0.95 (± 0.02)	0.96 (± 0.02)	0.88 (± 0.02)
	UCB ( $\beta = 1$ )	1.00 (± 0.00)	0.96 (± 0.02)	0.97 (± 0.02)	0.86 (± 0.02)
Joint Monte Carlo	PI	0.99 (± 0.00)	0.86 (± 0.05)	0.89 (± 0.04)	0.09 (± 0.06)
	EI	1.00 (± 0.00)	0.95 (± 0.03)	0.96 (± 0.02)	0.06 (± 0.03)
	UCB (variable)	0.97 (± 0.02)	0.89 (± 0.04)	0.90 (± 0.04)	0.81 (± 0.02)
	UCB ( $\beta = 5$ )	0.99 (± 0.00)	0.94 (± 0.03)	0.95 (± 0.03)	0.87 (± 0.02)
	UCB ( $\beta = 1$ )	1.00 (± 0.00)	0.96 (± 0.02)	0.97 (± 0.02)	0.78 (± 0.06)
Analytical	CL min	0.99 (± 0.00)	0.94 (± 0.02)	0.95 (± 0.02)	0.40 (± 0.08)
	CL max	1.00 (± 0.00)	0.95 (± 0.03)	0.96 (± 0.03)	0.67 (± 0.09)
	BUCB (variable)	0.98 (± 0.02)	0.85 (± 0.05)	0.85 (± 0.05)	0.65 (± 0.03)
	BUCB ( $\beta = 5$ )	0.98 (± 0.00)	0.88 (± 0.05)	0.90 (± 0.05)	0.66 (± 0.05)
	BUCB ( $\beta = 1$ )	0.99 (± 0.00)	0.88 (± 0.06)	0.90 (± 0.06)	0.65 (± 0.05)

when evaluating a set of hyperparameters for a complex model such as a deep neural network, the model has to be trained anew for each set of hyperparameters, racking up costs in time and computing resources. In most cases, the problem boils down to dividing a predefined budget into two parts: evaluations to initialize the algorithm and evaluations for points proposed by the Bayesian optimization algorithm. The decision of how many points to allocate for the training data is important as using too many could mean that the Bayesian optimization algorithm does not have sufficient evaluations available to find a good solution, i.e., the budget is exhausted before a promising solution is found. The simulations in this paper suggest that this decision might not be as complex as it initially seems, as only a few training points are necessary for the algorithm to achieve good results. Taking this approach of using only a small number of evaluations

for training points saves more evaluations for the optimization itself, thus increasing the chances of finding a good solution.

The third finding regards the information-based acquisition functions. Through the range of different simulations, PI and EI failed to find good results for the Ackley function. The reason for this is the large area of the parameter space where all response values are identical and thus the response surface is flat. As discussed in Section 2.2, improvement-based acquisition functions propose a point that is most likely to improve upon the previous best point. With a flat function like the Ackley function it is likely that all of the initial starting points fall into the flat area (especially in higher dimensions where the flat area grows even quicker). The posterior mean of the GP will then be very flat, and will in turn predict that the underlying objective function is flat as well. This leads to a very flat acquisition

function, as most input points will have a small likelihood of improving upon the previous best points. Gramacy [2] mentions problems when optimizing a flat EI acquisition function that result in the evaluation of points that are not optimal, which is especially problematic when the shape of the objective function is not known and flat regions cannot be ruled out *a priori*. If such properties are possible in a particular applied problem, the results suggest that using a different acquisition function, such as an optimistic policy, achieves better results. This again highlights the importance of expert knowledge for applications to specific problems.

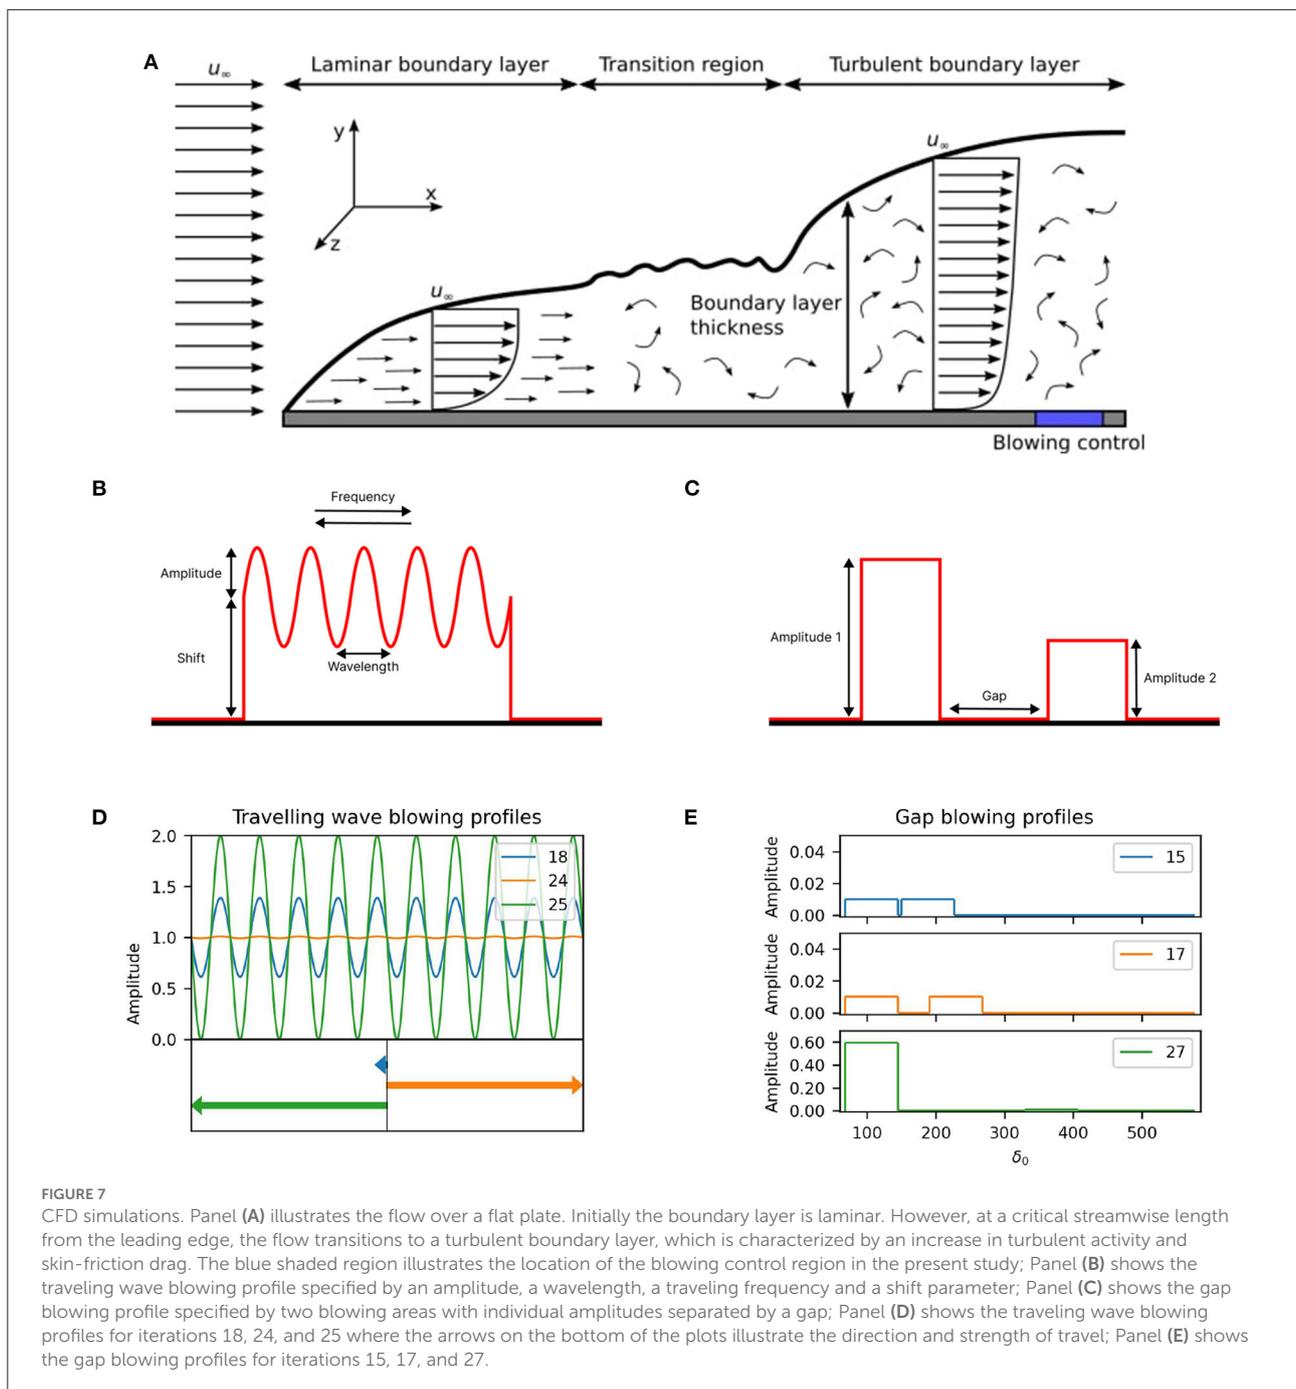
Fourthly, the simulations show that Monte Carlo acquisition functions yield comparable results to analytical functions. These functions use Monte Carlo sampling to compute the acquisition functions, instead of analytically solving them. Utilizing sampling to compute a function that can also be solved analytically might not appear useful at first glance as it is essentially just an approximation of the analytical results. However, this approach makes it straightforward to compute batches of candidate points (as outlined in Section 2.2), which presents the foundation for the next finding.

The fifth finding suggests that multi-point acquisition functions perform comparably to single-point approaches. All methods found good solutions across the range of considered problems, with the exception of improvement-based methods for the Ackley function, as discussed previously. Multi-point approaches are particularly beneficial for cases in which the objective function is expensive to evaluate and allows parallel evaluations [e.g., in high-fidelity turbulence resolving simulations [4]]. For these problems, the total time required to conclude the full experiment can be reduced as multiple points are evaluated in parallel each time. However, more evaluations in total might be required to achieve results comparable to the single-point approach. This method is particularly advantageous for applications involving simulations that can be evaluated in parallel, e.g., the CFD simulations mentioned in the introduction.

Lastly, the results showed no decrease in performance when adding noise to the objective function, in this case, the 6D Hartmann function. When optimizing the function with low and high noise (corresponding to the measurement error range of MEMS sensors to mirror the circumstances of an applied example in fluid dynamics) the same solutions were found, in a comparable number of evaluations, as for the deterministic case. This result can be attributed to the nugget that is added to the deterministic and noisy cases, as discussed in Section 2.1. These results are promising, as they show that Bayesian optimization can handle noisy objective functions just as efficiently as deterministic functions. This enables the use of Bayesian optimization for physical experiments where measurements cannot be taken without

noise, but also for simulations where other noise can be introduced unknowingly. Consider, for example, a physical experiment in which the drag created by air blowing over a surface can be reduced by blowing air upwards orthogonally to the same surface using multiple actuators (for further details refer to the application Section 4). As each actuator has a high, if not infinite, number of settings, finding the optimal overall blowing strategy is a complex problem that is a prime candidate for the use of Bayesian optimization. However, the drag cannot be measured without noise as the measurement errors associated with the MEMS sensors introduced in Section 3 show. As it is impossible to collect noise-free data in such circumstances, and taking the same measurement twice would yield different results, it is critical that Bayesian optimization performs equally well on these problems.

While these findings show that the use of Bayesian optimization to optimize expensive black box functions is promising, some limitations should be noted. Firstly, the acquisition functions considered in this paper represent only a subset of those available in the literature. The general results inferred from this selection might not extend to all acquisition functions. Secondly, the dimensionality of the test functions was selected to be no greater than 10. While Bayesian optimization is generally considered to work best in this range, and even up to 20 input parameters, it would appear unlikely that these results would hold in a much higher-dimensional space [23]. While this is an ongoing area of research, there are multiple algorithms that attempt to make Bayesian optimization viable for higher dimensions. A popular approach is to find a low-dimensional embedding of the high-dimensional input space. Examples of this approach are REMBO [24], HeSBO [25], and ALEBO [26]. Another approach is to leverage any possible underlying additive structure of the objective function with a Bayesian optimization algorithm based on additive GPs, such as the Add-GP-UCB algorithm [27]. Other promising approaches are TuRBO [28], which attempts to focus on promising areas of the parameter space via local optimization, and SAASBO [29], which utilizes a fully Bayesian GP that aims to identify sparse axis-aligned subspaces to extend Bayesian optimization to higher dimensions. Thirdly, the conclusions drawn in this paper are based on the synthetic test functions chosen and their underlying behavior. When encountering objective functions with shapes and challenges that are not similar to those considered in this paper, the findings might not hold. Lastly, the noise added to the Hartmann function could be too low to represent every possible experiment. It is possible that the measurement error, or other sources of noise, from an experiment or simulation, are too large for Bayesian optimization to work effectively. More investigation is needed to find the maximal noise levels that Bayesian optimization can tolerate while still performing well.



### 4. Application

In this section, we apply the findings from the investigation on synthetic test functions in Section 3 to the selection and design of Bayesian optimization algorithms to specific simulations, in this case in the area of CFD. Consider high-fidelity simulations involving the turbulent flow over a flat plate, as illustrated in Figure 7. Initially, the flow within the boundary layer is laminar. However, after a critical streamwise length,

the flow transitions into a turbulent regime, characterized by an increase in turbulence activity, and thereby an increase in skin-friction drag. This setup mimics the flow encountered on many vehicles, e.g., the flow over the wing of an aircraft, a high-speed train, or the hull of a ship. The skin-friction drag is a resistive force that opposes the motion of any moving vehicle and is typically responsible for more than half of the vehicle’s energy consumption. To place this into context, just a 3% reduction in the skin-friction drag force experienced by

a long-range commercial aircraft would save around £1.2 M in jet fuel per aircraft per year, and prevent the annual release of 3,000 tons of carbon dioxide [30]. There are currently over 25,000 aircraft in active service around the world. Yet despite this significance, the complexity of turbulence has prevented the realization of any functional and economical system to reduce the effects of turbulent skin-friction drag on any transportation vehicle. This in part, is due to our inability to find optimum solutions in parameter space to control the turbulence effectively and efficiently. The aim of these simulations is to minimize the turbulent skin-friction drag by utilizing active control, via actuators, which are seen as a key upstream technology approach for the aerospace sector, that allows us to either blow fluid away from the plate or suck fluid toward the plate. These actuators are located toward the beginning of the plate, but after the transition region, where the flow is fully turbulent. An averaged skin-friction drag coefficient is then measured along the plate from the blowing region. Within the blowing region a very large number of blowing setups are possible. For example, a simple setup could be a 1D problem where fluid is blown away from the plate uniformly over the entire blowing region, with a constant amplitude. In this case, the only parameter to optimize would be the amplitude of the blowing. However, many more complex setups are possible [e.g., see [4]]. These numerical simulations are a prime candidate for Bayesian optimization as they fulfill the characteristics of an expensive-to-evaluate black box function: the underlying mathematical expression is unknown and each evaluation of the objective function is expensive. Indeed, one evaluation (a high-fidelity simulation with converged statistics) can take up to 12 h on thousands of CPU cores and requires the use of specialist software, since the full turbulence activity covering the flat plate must be simulated in order to correctly resolve the quantity of interest (i.e., skin-friction drag). To perform these numerical simulations, the open-source flow solver Xcompact3D [31] is utilized on the UK supercomputing facility ARCHER2. This section optimizes two blowing profiles that follow this setup: a traveling wave with four degrees of freedom and a gap problem with three degrees of freedom, where two blowing areas with individual amplitudes are separated by a gap.

The computational setup consists of a laminar Blasius solution at the inlet, a convective condition at the outlet, a homogenous Neumann condition in the far-field, and periodic conditions in the spanwise direction. The domain dimensions are  $L_x \times L_y \times L_z = 750\delta_0 \times 80\delta_0 \times 30\delta_0$ , where  $\delta_0$  is the boundary-layer thickness at the inlet. The Reynolds number at the inlet is  $Re_{\delta_0} = 1,250$ , based on the boundary-layer thickness (i.e., the thin layer of fluid above the plate where the flow velocity is reduced due to the presence of the plate) and free-stream velocity ( $u_\infty$ , i.e., the speed of the moving vehicle) at the inlet of the simulation domain. This corresponds to a momentum Reynolds number of  $Re_{\theta_0} \approx 169$  to 2,025, for the canonical case (no control), from the inlet to the outlet. The

mesh size is chosen to be  $n_x \times n_y \times n_z = 1,537 \times 257 \times 128$ , with a uniform spacing in the streamwise (x-axis in Figure 7A) and spanwise (z-axis) directions and non-uniform spacing in the wall-normal direction (y-axis) to properly resolve the near-wall effects. This results in a mesh resolution of  $\Delta x^+ = 31$ ,  $0.54 \leq \Delta y^+ \leq 705$ , and  $\Delta z^+ = 15$  in viscous (inner) units, where the inner scaling is with respect to the friction velocity (i.e., scaled by the wall-shear stress generated by the skin-friction drag force) at the start of the control region for the canonical case. The control region extends from  $x = 68\delta_0$  to  $145\delta_0$  in the streamwise direction, corresponding to a Reynolds number range of  $Re_\theta \approx 479$ –703, for the canonical case. To accelerate the transition-to-turbulence, a random volume forcing approach [32], located at  $x = 3.5\delta_0$ , is used to trip the boundary layer.

Figure 7 depicts the two blowing profiles in question, and defines the parameters for the optimization. The traveling wave with four degrees of freedom, given in Figure 7B, is a wave defined by an amplitude in the range 0.01–1.00% of the overall free-stream velocity and a wavelength between 0.00 and 0.02 (inner scaling). The angular frequency, restricted to values between -0.25 and 0.25 (inner scaling), allows the wave to travel up- and down-stream. Lastly, a shift parameter displaces the wave vertically up and down. This parameter is restricted to values between -1.00 and 1.00% of the free-stream velocity. The blowing turns into suction for cases where the blowing profile is negative. The gap configuration with three degrees of freedom, illustrated in Figure 7C, includes two blowing regions with individual amplitudes in the range 0.01–1.00% of the overall free-stream velocity and a gap restricted to between  $5\delta_0$  and  $355\delta_0$ . While the aim for both problems is the maximization of the global drag reduction (GDR), defined as the globally averaged skin-friction drag reduction with respect to the canonical case, the gap problem with three degrees of freedom also considers the energy consumption of the actuators and tries to find profiles that reduce the drag while also achieving NES. Net-energy savings is achieved when the energy used by the blowing device is smaller than the energy saved by the drag reduction (much of the energy expenditure in aerodynamics/hydrodynamics applications is used to overcome the skin-friction drag). In this work, NES is calculated following the approach of Mahfoze et al. [4], where the input power for the blowing device is estimated from real-world experimental data and a relationship between input power and blowing velocity is derived [see Mahfoze et al. [4] for more details].

The Bayesian optimization algorithm used for both problems is informed by the results of the synthetic test functions from Section 3. For the surrogate model a GP with a zero mean function and a Matérn kernel with  $\nu = 5/2$  was defined and its hyperparameters were estimated from the training data using MLE. While CFD simulations are expensive, they allow for parallel evaluations. This is done by running, concurrently, multiple simulations, and combining the results once all simulations are completed. Therefore, these setups

TABLE 7 Results for traveling wave experiment.

Evaluation	Amplitude	Shift	Wavenumber	Frequency	GDR (in %)
1	0.50	-0.74	0.02	-0.22	-25.71
2	0.96	0.50	0.00	0.20	12.26
3	0.85	0.67	0.01	0.23	15.91
4	0.09	-0.48	0.02	-0.07	-15.63
5	0.42	-0.16	0.01	-0.02	-4.67
6	0.62	0.89	0.00	-0.05	19.69
7	0.44	-0.50	0.00	-0.22	-16.73
8	0.19	0.31	0.01	-0.19	8.33
9	0.02	-0.34	0.00	0.17	-10.94
10	0.68	-0.98	0.01	-0.12	-34.70
11	0.31	0.62	0.01	0.12	14.86
12	0.73	0.85	0.01	-0.15	19.52
13	0.34	0.16	0.02	0.08	4.28
14	0.81	0.06	0.02	0.04	1.94
15	0.90	-0.02	0.01	0.01	-1.22
16	0.24	-0.77	0.01	0.15	-26.98
17	0.84	1.00	0.01	0.04	21.43
<b>18</b>	<b>0.39</b>	<b>1.00</b>	<b>0.01</b>	<b>-0.00</b>	<b>22.07</b>
19	1.00	1.00	0.00	-0.25	21.77
20	1.00	1.00	0.02	0.00	21.83
21	0.48	1.00	0.02	0.13	21.90
<b>22</b>	<b>0.01</b>	<b>1.00</b>	<b>0.02</b>	<b>-0.25</b>	<b>22.06</b>
23	0.01	1.00	0.00	-0.25	21.88
<b>24</b>	<b>0.01</b>	<b>1.00</b>	<b>0.02</b>	<b>0.25</b>	<b>22.19</b>
<b>25</b>	<b>1.00</b>	<b>1.00</b>	<b>0.02</b>	<b>-0.25</b>	<b>22.03</b>
26	0.01	1.00	0.00	0.25	21.95
<b>27</b>	<b>0.01</b>	<b>1.00</b>	<b>0.02</b>	<b>0.01</b>	<b>22.00</b>
<b>28</b>	<b>0.42</b>	<b>1.00</b>	<b>0.01</b>	<b>-0.25</b>	<b>22.07</b>
<b>29</b>	<b>0.01</b>	<b>1.00</b>	<b>0.01</b>	<b>0.25</b>	<b>22.04</b>
<b>30</b>	<b>0.01</b>	<b>1.00</b>	<b>0.01</b>	<b>0.01</b>	<b>22.17</b>
31	0.43	0.79	0.00	-0.22	18.45
32	1.00	1.00	0.00	0.25	21.88
33	0.01	1.00	0.01	-0.25	21.93
34	0.47	1.00	0.02	-0.25	21.75
35	1.00	1.00	0.02	0.25	21.83
36	0.83	0.94	0.02	-0.02	21.02

Horizontal line separates initial training points and points suggested by Bayesian optimization. Points with global drag reduction over 22% in bold.

are well-suited for the multi-point approach presented in Section 3.2.4 that, as our investigations showed, has no real disadvantage compared to the single-point approach. Based on previous work, for example Mahfoze et al. [4], the possibility that the underlying objective function is characterized by large flat areas similar to the 6D Ackley function cannot be ruled out. Thus, an acquisition function that allows the selection of batches and performs well even when encountering flat areas should be implemented. The sequential Monte Carlo Upper Confidence

Bound acquisition function with the trade-off hyperparameter  $\beta = 1$  yielded very good results for the Ackley function as well as all other test function and is thus chosen with a batch size of four to optimize the CFD problems in this section. Section 3.2.2 showed that Bayesian optimization is able to find good solutions even with a relatively small number of initial training data points. Hence, for both problems, four points per input dimension were randomly selected from a Maximin Latin Hypercube [21].

TABLE 8 Results for gap experiment.

Evaluation	Amplitude 1	Amplitude 2	Gap	GDR (in %)	NES (in %)
1	0.97	0.55	110.35	32.86	-4.31
2	0.28	0.12	191.95	10.19	-1.31
3	0.87	0.89	225.60	34.75	-7.41
4	0.40	0.41	289.40	17.75	-3.36
5	0.13	0.69	241.42	16.14	-4.34
6	0.71	0.25	133.84	22.47	-1.66
7	0.66	0.80	352.52	26.74	-7.97
8	0.53	0.64	179.61	25.28	-3.18
9	0.83	0.03	318.82	19.83	-0.97
10	0.19	0.96	51.96	25.34	-3.53
11	0.04	0.30	92.16	8.72	-0.96
12	0.49	0.45	31.67	22.86	-0.87
13	0.27	0.01	5.00	7.72	-0.28
14	1.00	0.01	5.00	22.04	-2.84
15	<b>0.01</b>	<b>0.01</b>	<b>5.00</b>	<b>0.91</b>	<b>0.25</b>
16	0.15	0.38	5.00	13.91	-0.77
17	<b>0.01</b>	<b>0.01</b>	<b>45.54</b>	<b>0.98</b>	<b>0.32</b>
18	<b>0.01</b>	<b>0.01</b>	<b>355.00</b>	<b>0.72</b>	<b>0.06</b>
19	0.01	0.10	5.00	3.01	-0.44
20	0.47	0.01	355.00	11.84	-0.39
21	<b>0.01</b>	<b>0.01</b>	<b>131.32</b>	<b>0.77</b>	<b>0.11</b>
22	0.01	0.01	237.00	0.62	-0.04
23	0.01	0.63	5.00	15.33	-0.15
24	<b>0.01</b>	<b>0.01</b>	<b>77.50</b>	<b>0.80</b>	<b>0.13</b>
25	<b>0.01</b>	<b>0.01</b>	<b>5.00</b>	<b>0.91</b>	<b>0.25</b>
26	0.28	0.63	5.00	21.72	-1.24
27	0.60	0.01	184.71	14.83	-0.01
28	0.38	0.67	5.00	24.86	-1.31

Horizontal line separates initial training points and points suggested by Bayesian optimization. Points with net-energy savings in bold.

Table 7 present the results of the 16 training points plus 20 points (or five batches of four points), selected using Bayesian optimization, by maximizing the GDR, for the traveling wave problem. While the highest GDR of the initial training data was 19.69%, Bayesian optimization improves upon this value with each evaluated batch, finding multiple strategies that achieve a GDR above 22%, with the best strategy, from batch 2, giving a GDR of 22.19%. Three blowing profiles, including the overall best solution found, are depicted in Figure 7D. Overall, the shift parameter seems to be the main driver behind the drag reduction, as almost all strategies selected by Bayesian optimization implement the upper limit of this parameter, independent of the other parameter values.

While blowing at a high amplitude yields increased skin-friction drag reduction, it also consumes more energy. The second experiment addresses this point, and accounts for the energy consumption, following Mahfoze et al. [4], when

optimizing the blowing profile. Table 8 provides the results for 12 initial training points and four batches, suggested via Bayesian optimization, by maximizing the NES. The initial strategies selected by the Latin Hypercube did not find a solution that achieved both GDR and NES. Bayesian optimization was able to find multiple strategies that achieved both, in which the NES and the GDR were relatively small (0.11–0.32% and 0.77–0.98%). However, the algorithm also found one strategy with NES of -0.01% and a GDR of 14.83%. While this strategy did not achieve NES, it did not increase overall energy use, and a small increase in the efficiency of the actuators could yield NES with a considerable GDR. Compared to the high intensity blowing for the traveling wave, in this experiment the amplitudes are clustered toward the lower end of the parameter space. This is a result of the objective to optimize NES, which penalizes high-velocity blowing due to its increased power requirements. Figure 7E illustrates this by providing the blowing profiles of three solutions; the two solutions with the

greatest NES and the solution with a high GDR, as previously described.

## 5. Conclusion

In this paper, Bayesian optimization algorithms, implemented with different types of acquisition functions, were benchmarked regarding their performance and their robustness on synthetic test functions inspired by applications in engineering and machine learning. Synthetic test functions have the advantage that their shape and their global optimum are known. This allows the algorithms to be evaluated on (a) how close their best solutions are to the global optimum and (b) how well they perform on specific challenges such as oscillating functions or functions with steep edges. This evaluation can indicate advantages and shortcomings of the individual acquisition functions and can inform researchers on the best approach to choose for their specific problem.

Four sets of comparisons were conducted in this research article. First, analytical single-point acquisition functions were compared to each other. Second, the effect of varying the number of initial training data points was investigated. Third, the analytical approach was contrasted with acquisition functions based on Monte Carlo sampling. Fourth, the single-point approach was compared to the multi-point or batched approach.

From these experiments six main conclusions could be drawn: (i) While all acquisition functions performed well on simple test functions, optimistic policies, such as the UCB, dealt best with challenging problems. (ii) Varying the number of initial training data points did not have a significant effect on the performance of the individual methods. (iii) Improvement-based acquisition functions struggled with flat test functions. (iv) The results showed that Monte Carlo acquisition functions and multi-point acquisition functions present a good alternative to the widely used analytical single-point methods. (v) The multi-point approach is particularly advantageous when the objective function takes a long time to evaluate and allows parallel evaluations. (vi) Bayesian optimization performs equally well on noisy objective functions.

Finally, two experiments in CFD were taken as illustrative examples on how the findings of this paper can be used to guide the design of a Bayesian optimization algorithm and tailor it to unique problems. In detail, a multi-point approach was employed that used the Monte Carlo Upper Confidence Bound acquisition function allowing multiple points to be evaluated in parallel with concurrent simulations. For both experiments, Bayesian optimization was able to improve upon the training points straight away and find solutions that implement GDR for the traveling wave and GDR as

well as NES for the experiment where two blowing areas were separated by a gap. However, the effects of the second experiment remained relatively small but these optimization studies are nevertheless encouraging in the quest to design a robust and efficient control strategy to reduce drag around moving vehicles.

## Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://github.com/mikediessner/investigating-BO>.

## Author contributions

MD: methodology, software, validation, formal analysis, investigation, writing—original draft, writing—review and editing, and visualization. JO'C: methodology, software, investigation, and writing—review and editing. AW, SL, KW, and RW: conceptualization, writing—review and editing, and supervision. YG: conceptualization and supervision. All authors contributed to the article and approved the submitted version.

## Funding

The work has been supported by the Engineering and Physical Sciences Research Council (EPSRC) under (grant numbers EP/T020946/1 and EP/T021144/1) and the EPSRC Center for Doctoral Training in Cloud Computing for Big Data under (grant number EP/L015358/1).

## Acknowledgments

The authors would also like to thank the UK Turbulence Consortium (EP/R029326/1) for the computational time made available on the [ARCHER2 UK National Supercomputing Service](#).

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those

of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## References

1. Shahriari B, Swersky K, Wang Z, Adams RP, De Freitas N. Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE*. (2015) 104:148–75. doi: 10.1109/JPROC.2015.2494218
2. Gramacy RB. *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. Boca Raton, FL: Chapman and Hall/CRC (2020).
3. Talnikar C, Blonigan P, Bodart J, Wang Q. Parallel optimization for large Eddy simulations. *arXiv Preprint*. arXiv:14108859. (2014). doi: 10.48550/arXiv.1410.8859
4. Mahfoze O, Moody A, Wynn A, Whalley R, Laizet S. Reducing the skin-friction drag of a turbulent boundary-layer flow with low-amplitude wall-normal blowing within a Bayesian optimization framework. *Phys Rev Fluids*. (2019) 4:094601. doi: 10.1103/PhysRevFluids.4.094601
5. Morita Y, Rezaeiravesh S, Tabatabaei N, Vinuesa R, Fukagata K, Schlatter P. Applying Bayesian optimization with Gaussian process regression to computational fluid dynamics problems. *J Comput Phys*. (2022) 449:110788. doi: 10.1016/j.jcp.2021.110788
6. Nabaee Y, Fukagata K. Bayesian optimization of traveling wave-like wall deformation for friction drag reduction in turbulent channel flow. *J Fluid Sci Technol*. (2021) 16:JFST0024. doi: 10.1299/jfst.2021jfst0024
7. Wu J, Chen XY, Zhang H, Xiong LD, Lei H, Deng SH. Hyperparameter optimization for machine learning models based on Bayesian optimization. *J Electr Sci Technol*. (2019) 17:26–40. doi: 10.11989/JEST.1674-862X.80904120
8. Snoek J, Larochelle H, Adams RP. Practical Bayesian optimization of machine learning algorithms. In: Pereira F, Burges CJ, Bottou L, Weinberger KQ, editors. *Advances in Neural Information Processing Systems 25*. Red Hook, NY: Curran Associates, Inc. (2012) 2951–9.
9. Srinivas N, Krause A, Kakade SM, Seeger M. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv Preprint*. arXiv:09123995. (2009). doi: 10.48550/arXiv.0912.3995
10. Andrianakis I, Challenor PG. The effect of the nugget on Gaussian process emulators of computer models. *Comput Statist Data Anal*. (2012) 56:4215–28. doi: 10.1016/j.csda.2012.04.020
11. Gramacy RB, Lee HK. Cases for the nugget in modeling computer experiments. *Stat Comput*. (2012) 22:713–22. doi: 10.1007/s11222-010-9224-x
12. Rasmussen CE, Williams CKI. *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press (2006).
13. Brochu E, Hoffman M, De Freitas N. Portfolio allocation for Bayesian optimization. In: *UAI'11: Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*. Arlington, VA: Citeseer (2011) pp. 327–36.
14. Wang Z, Jegelka S. Max-value entropy search for efficient Bayesian optimization. In: *Proceedings of the 34th International Conference on Machine Learning, PMLR*. Sydney, NSW (2017) pp. 3627–35.
15. Takeno S, Fukuoka H, Tsukada Y, Koyama T, Shiga M, Takeuchi I, et al. Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In: *Proceedings of the 37th International Conference on Machine Learning, PMLR*. Virtual Event (2020). p. 9334–45.
16. Wilson JT, Moriconi R, Hutter F, Deisenroth MP. The reparameterization trick for acquisition functions. *arXiv Preprint*. arXiv:171200424 (2017). doi: 10.48550/arXiv.1712.00424
17. Balandat M, Karrer B, Jiang D, Daulton S, Letham B, Wilson AG, et al. BoTorch: a framework for efficient Monte-Carlo Bayesian optimization. In: *33rd*

## Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fams.2022.1076296/full#supplementary-material>

- Annual Conference on Neural Information Processing Systems. NeurIPS 2019*. Vancouver, BC (2020). pp. 21524–38.
18. Ginsbourger D, Riche RL, Carraro L. Kriging is well-suited to parallelize optimization. In: Tenne, Y., Goh, CK, editors. *Computational Intelligence in Expensive Optimization Problems*. Vol. 2. Berlin; Heidelberg: Springer (2010). pp. 131–62.
  19. Desautels T, Krause A, Burdick JW. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *J Mach Learn Res*. (2014) 15:4053–103. Available online at: <https://jmlr.org/papers/v15/desautels14a.html>
  20. Ebrahimzade N, Portoles J, Wilkes M, Cumpson P, Whalley RD. Optical MEMS sensors for instantaneous wall-shear stress measurements in turbulent boundary-layer flows. In: *12th Symposium on Turbulence and Shear Flow Phenomena (TSFP)*. Osaka. (2022). pp. 1–6.
  21. Husslage BG, Rennen G, Van Dam ER, Den Hertog D. Space-filling Latin hypercube designs for computer experiments. *Optim Eng*. (2011) 12:611–30. doi: 10.1007/s11081-010-9129-8
  22. Wilson J, Hutter F, Deisenroth M. Maximizing acquisition functions for Bayesian optimization. In: Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, editors. *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*. Montreal, QC (2018).
  23. Moriconi R, Deisenroth MP, Sesh Kumar K. High-dimensional Bayesian optimization using low-dimensional feature spaces. *Mach Learn*. (2020) 109:1925–43. doi: 10.1007/s10994-020-05899-z
  24. Wang Z, Hutter F, Zoghi M, Matheson D, De Freitas N. Bayesian optimization in a billion dimensions via random embeddings. *J Artif Intell Res*. (2016) 55:361–87. doi: 10.1613/jair.4806
  25. Nayebi A, Munteanu A, Poloczek M. A framework for Bayesian optimization in embedded subspaces. In: *International Conference on Machine Learning, PMLR*. Long Beach, CA (2019) pp. 4752–61.
  26. Letham B, Calandra R, Rai A, Bakshy E. Re-examining linear embeddings for high-dimensional Bayesian optimization. In: Larochelle H, Ranzato MA, Hadsell R, Balcan M-F, Lin H-T, editors. *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*. Virtual. (2022) 1546–58.
  27. Kandasamy K, Schneider J, Póczos B. High dimensional Bayesian optimisation and bandits via additive models. In: *Proceedings of the 32nd International Conference on Machine Learning, PMLR*. Lille. (2015) pp. 295–304.
  28. Eriksson D, Pearce M, Gardner J, Turner RD, Poloczek M. Scalable global optimization via local bayesian optimization. In: *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*. Vancouver, BC (2019).
  29. Eriksson D, Jankowiak M. High-dimensional Bayesian optimization with sparse axis-aligned subspaces. In: *Proceedings of the Thirty-Seventh Uncertainty in Artificial Intelligence (UAI 2021)*. Virtual. pp. 493–503.
  30. Bushnell DM. *Viscous Drag Reduction in Boundary Layers*. Vol. 123. Reston, VA: American Institute of Aeronautics and Astronautics (1990).
  31. Bartholomew P, Deskos G, Frantz RA, Schuch FN, Lamballais E, Laizet S. Xcompact3D: an open-source framework for solving turbulence problems on a Cartesian mesh. *SoftwareX*. (2020) 12:100550. doi: 10.1016/j.softx.2020.100550
  32. Schlatter P, Örlü R. Assessment of direct numerical simulation data of turbulent boundary layers. *J Fluid Mech*. (2010) 659:116–26. doi: 10.1017/S0022112010003113