

ϵ -shotgun: ϵ -greedy Batch Bayesian Optimisation

George De Ath
g.de.ath@exeter.ac.uk

Department of Computer Science
University of Exeter
Exeter, United Kingdom

Jonathan E. Fieldsend
j.e.fieldsend@exeter.ac.uk

Department of Computer Science
University of Exeter
Exeter, United Kingdom

Richard M. Everson
r.m.everson@exeter.ac.uk

Department of Computer Science
University of Exeter
Exeter, United Kingdom

Alma A. M. Rahat
a.a.m.rahat@swansea.ac.uk

Department of Computer Science
Swansea University
Swansea, United Kingdom

ABSTRACT

Bayesian optimisation is a popular surrogate model-based approach for optimising expensive black-box functions. Given a surrogate model, the next location to expensively evaluate is chosen via maximisation of a cheap-to-query acquisition function. We present an ϵ -greedy procedure for Bayesian optimisation in batch settings in which the black-box function can be evaluated multiple times in parallel. Our ϵ -shotgun algorithm leverages the model's prediction, uncertainty, and the approximated rate of change of the landscape to determine the spread of batch solutions to be distributed around a putative location. The initial target location is selected either in an exploitative fashion on the mean prediction, or – with probability ϵ – from elsewhere in the design space. This results in locations that are more densely sampled in regions where the function is changing rapidly and in locations predicted to be good (i.e. close to predicted optima), with more scattered samples in regions where the function is flatter and/or of poorer quality. We empirically evaluate the ϵ -shotgun methods on a range of synthetic functions and two real-world problems, finding that they perform at least as well as state-of-the-art batch methods and in many cases exceed their performance.

CCS CONCEPTS

• **Theory of computation** → **Gaussian processes; Mathematical optimization**; • **Computing methodologies** → **Modeling and simulation; Optimization algorithms**.

KEYWORDS

Bayesian optimisation, Batch, Parallel, Exploitation, ϵ -greedy, Infill criteria, Acquisition function

ACM Reference Format:

George De Ath, Richard M. Everson, Jonathan E. Fieldsend, and Alma A. M. Rahat. 2020. ϵ -shotgun: ϵ -greedy Batch Bayesian Optimisation. In *Genetic and Evolutionary Computation Conference (GECCO '20)*, July 8–12, 2020, Cancún, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3377930.3390154>

1 INTRODUCTION

Global optimisation of non-convex and black-box functions is a common task in many real-world problems. These include hyperparameter tuning of machine learning algorithms [36], drug discovery [21], analog circuit design [28], mechanical engineering design [6, 8, 31] and general algorithm configuration [23]. Bayesian optimisation (BO) has become a popular approach for optimising expensive, black-box functions that have no closed-form expression or derivative information [35, 36]. It employs a probabilistic surrogate model of a function using available function evaluations. The location at which the function is next expensively evaluated is chosen as the location that maximises an acquisition function (or infill criterion) that balances exploration and exploitation.

In real-world problems it is often possible to run multiple experiments in parallel by using modern hardware capabilities to expensively evaluate several locations at once. When optimising machine learning algorithms, for example, multiple model configurations can be evaluated in parallel across many processor cores on one or multiple machines [4, 26]. Consequently, this has led to the development of batch (or parallel) BO algorithms, which use acquisition functions to select q locations to be evaluated at each iteration. Clearly, a strictly serial evaluation makes the best overall use of the available CPU time because each new location to be evaluated is selected with the maximum available information. Parallel evaluation, however, holds the promise of substantially reducing the wall-clock time to locate the optimum.

The selection of a good set of locations to evaluate at each batch iteration is a non-trivial problem. In sequential BO, techniques which favour greedy exploitation of the surrogate model have been shown to be preferable to the more traditional acquisition functions [10, 33]. De Ath et al. [10], for example, show that using an ϵ -greedy strategy of exploiting the surrogate model the majority of the time and, with probability ϵ (where $\epsilon \approx 0.1$), randomly selecting a location to explore yields superior optimisation performance on a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '20, July 8–12, 2020, Cancún, Mexico

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7128-5/20/07...\$15.00

<https://doi.org/10.1145/3377930.3390154>

variety of synthetic and real-world problems. Consequently, in this work we investigate ϵ -greedy methods in the batch BO setting.

We present ϵ -shotgun, a novel approach to batch BO, which uses an ϵ -greedy strategy for selecting the first location \mathbf{x}'_1 in a batch, and then samples the remaining $q - 1$ points from a normal distribution centred on \mathbf{x}'_1 , with a scale parameter determined by the surrogate model's posterior mean and variance at \mathbf{x}'_1 and the magnitude of the gradient in the vicinity of \mathbf{x}'_1 . This embodies maximum exploitation of the surrogate model the majority of the time by virtue of the choice of \mathbf{x}'_1 . The remaining $q - 1$ locations may be exploratory or exploitative depending on the characteristics of the local landscape. Larger regions of decision space will be sampled when \mathbf{x}'_1 is surrounded by a relatively flat landscape, while denser sampling will occur where \mathbf{x}'_1 is in a locally steeper region, such as the landscape around a local (or global) optimum.

Our contributions can be summarised as follows:

- We present ϵ -shotgun, a new batch Bayesian optimisation approach based on the ϵ -greedy strategy of exploiting the surrogate model.
- We empirically compare a range of state-of-the-art batch Bayesian optimisers across a variety of synthetic test problems and two real-world applications
- We empirically show that the ϵ -shotgun approaches are equal to or better than several state-of-the-art batch BO methods on a wide range of problems.

We begin in Section 2 by briefly reviewing Bayesian optimisation along with Gaussian processes (the surrogate model generally used in BO) and common acquisition functions. Batch BO and the algorithm archetypes used for selecting the batch locations are then reviewed in Section 2.2, which leads to the proposed ϵ -shotgun approach in Section 3. Empirical evaluation on well-known test problems and two real-world applications are presented in Section 4. We finish with concluding remarks in Section 5.

2 BAYESIAN OPTIMISATION

Our goal is to minimise a black-box function $f : \mathcal{X} \mapsto \mathbb{R}$, defined on a compact domain $\mathcal{X} \subset \mathbb{R}^d$. The function itself is unknown, but we have access to the results of its evaluations $f(\mathbf{x})$ at any location $\mathbf{x} \in \mathcal{X}$. We are particularly interested in cases where the evaluations are expensive, either in terms of time or money or both, and we seek to minimise f in either as few evaluations as possible to incur as little cost as possible or for a fixed budget.

2.1 Sequential Bayesian Optimisation

Bayesian Optimisation (BO), also known as Efficient Global Optimisation, is a global search strategy that sequentially samples design space at locations that are likely contain the global optimum, taking into account the predictions of the surrogate model and their associated uncertainty [25]. It starts by generating M initial sample locations $\{\mathbf{x}_i\}_{i=1}^M$ with a space filling algorithm, typically Latin hypercube sampling [29], and expensively evaluates them with the function, $f_i = f(\mathbf{x}_i)$. This collected set of observations forms the dataset with which the surrogate model is initially trained. Following model training, and at each iteration of BO, the next location for expensive evaluation is selected according to an acquisition function (or infill criterion). These usually combine the surrogate

model's prediction and prediction uncertainty of the design space to balance the exploitation of promising solutions (those with good predicted values) and those solutions with high uncertainty. The location \mathbf{x}' maximising this criterion is used as the next point to be expensively evaluated. The dataset is augmented with \mathbf{x}' and $f(\mathbf{x}')$ and the process is repeated until the budget is exhausted. The value of the global minimum f_{\min} is estimated to be the best function evaluation seen during the optimisation run, i.e. $f^* = \min_i \{f_i\}$.

2.1.1 Gaussian Processes. Gaussian processes (GP) are a popular and versatile choice of surrogate model for $f(\mathbf{x})$, due to their strengths in function approximation and uncertainty quantification [32]. A GP is a collection of random variables, and any finite number of these are jointly Gaussian distributed. A GP prior over f can be defined as $\mathcal{GP}(m(\mathbf{x}), \kappa(\mathbf{x}, \mathbf{x}' | \theta))$ where $m(\mathbf{x})$ is the mean function, $\kappa(\cdot, \cdot)$ is the kernel function (also known as a covariance function) and θ are the hyperparameters of the kernel. Given data consisting of $f(\mathbf{x})$ evaluated at M sampled locations $\mathcal{D} = \{(\mathbf{x}_i, f_i \triangleq f(\mathbf{x}_i))\}_{i=1}^M$, the posterior estimate of f at location \mathbf{x} is a Gaussian distribution:

$$p(f(\mathbf{x}) | \mathbf{x}, \mathcal{D}, \theta) = \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x})) \quad (1)$$

with mean and variance

$$\mu(\mathbf{x} | \mathcal{D}, \theta) = \kappa(\mathbf{x}, X)K^{-1}\mathbf{f} \quad (2)$$

$$\sigma^2(\mathbf{x} | \mathcal{D}, \theta) = \kappa(\mathbf{x}, \mathbf{x}) - \kappa(\mathbf{x}, X)^\top K^{-1}\kappa(X, \mathbf{x}), \quad (3)$$

where $X \in \mathbb{R}^{M \times d}$ is matrix of input locations in each row and $\mathbf{f} \in \mathbb{R}^M$ is the corresponding vector of true function evaluations $\{f_1, f_2, \dots, f_M\}$. The matrix $K \in \mathbb{R}^{M \times M}$ contains the kernel evaluated at each pair of observations, and $\kappa(\mathbf{x}, X)$ is the M -dimensional vector whose elements are $[\kappa(\mathbf{x}, X)]_i = \kappa(\mathbf{x}, \mathbf{x}_i)$. Kernel hyperparameters θ are learnt via maximising the log likelihood:

$$\log p(\mathcal{D} | \theta) = -\frac{1}{2} \log \|K\| - \frac{1}{2} \mathbf{f}^\top K^{-1} \mathbf{f} - \frac{M}{2} \log(2\pi). \quad (4)$$

For notational simplicity, we drop explicit dependencies on the data \mathcal{D} and kernel hyperparameters θ from now on.

2.1.2 Acquisition Functions. An acquisition function $\alpha(\mathbf{x})$ is used to measure the anticipated quality of expensively evaluating f at any given location \mathbf{x} : the location that maximises the acquisition function is chosen as the next location for expensive evaluation. While this strategy may appear merely to transfer the problem of optimising $f(\mathbf{x})$ to a maximisation of $\alpha(\mathbf{x})$, the acquisition function is cheap to evaluate so the location of its global optimum can be cheaply estimated using an evolutionary algorithm.

Acquisition functions attempt, either implicitly or explicitly, to balance the trade-off between maximally exploiting the surrogate model, i.e. selecting a location with the best predicted value, and maximally exploring the model, i.e. selecting the location with the most uncertainty. Perhaps the two most widespread acquisition functions, are Expected Improvement (EI) [25] and Upper Confidence Bound (UCB) [37]. EI measures the positive predicted improvement over the best solution observed so far and UCB is a weighted sum of the surrogate model's mean prediction $\mu(\mathbf{x})$ and uncertainty $\sigma^2(\mathbf{x})$. These were both shown [10] to be monotonic with respect to increases in both $\mu(\mathbf{x})$ and $\sigma^2(\mathbf{x})$ and that the solutions that maximise them both belong to the Pareto set of locations

which maximally trade-off exploitation (minimising $\mu(\mathbf{x})$) and exploration (maximising $\sigma^2(\mathbf{x})$).

Recently, ϵ -greedy approaches have been successfully used as acquisition functions [10]. These select a maximally exploitative solution, $\mathbf{x}' = \operatorname{argmin}_{\mathbf{x}} \mu(\mathbf{x})$ with probability $1 - \epsilon$ and select a random solution with probability ϵ . De Ath et al. [10] present two methods for selecting the random solution, either uniformly from \mathcal{X} or from the approximate Pareto set of solutions of the surrogate model's mean prediction and variance. They showed that ϵ -greedy approaches are particularly effective on higher dimensional problems, and that performing pure exploitation (i.e. $\epsilon = 0$) is competitive with the best-performing methods. This result was recently confirmed by Rehbach et al. [33], who empirically show that solely using the surrogate model's predicted value performs better than EI on most problems with a dimensionality of 5 or more.

2.2 Batch Bayesian Optimisation

In batch Bayesian optimisation (BBO) the goal is to select a batch $\mathcal{X}' = \{\mathbf{x}'_1, \dots, \mathbf{x}'_q\}$ of q promising locations to expensively evaluate in parallel. One of the earliest BBO approaches, the qEI method of Ginsbourger et al. [15], generalised the sequential EI acquisition function to a batch setting in which all q batch locations are jointly estimated. However, it is not analytically tractable to compute qEI, even for small batch sizes [17]. Although a fast approximation to qEI does exist [5], it is not faster than naive Monte Carlo approximations for larger batch sizes. More recently, Wang et al. [38] proposed a more efficient algorithm to estimate the gradient of qEI, but the approach still results in having to optimise in a $d \times q$ dimensional space for each set of batch locations. Two other methods that jointly optimise the batch of locations, the parallel predictive entropy search [34] and the parallel knowledge gradient method [42], have also been shown to scale poorly as batch size increases [9].

Consequently, iteratively selecting the batch sample locations has become the prevailing methodology. One such strategy is to attempt to ensure that different locations are selected for the batch by, for each of the q locations, sampling a realisation from the surrogate model posterior (Thompson Sampling) and minimising it [26]. However, this relies on there being sufficient uncertainty in the model to allow for the realisations to have different minima [11]. De Palma et al. [11] proposed sampling from a distribution of acquisition functions, or rather from the distribution of hyperparameters that control the acquisition function's behaviour, such as the trade-off between exploration and exploitation in UCB.

Instead of relying on the stochasticity of either the surrogate model or acquisition function hyperparameters, another group of methods penalise the regions from which a batch point has already been selected; thus they are less likely (or unable to) select from nearby locations. A well-known heuristic to achieve this is to *hallucinate* the results of pending evaluations [2, 13, 16]. In this set of methods, the first batch location is selected by optimising an acquisition function and then subsequent locations are chosen by incorporating the predicted outcome of the already-selected batch locations into the surrogate model and optimising the acquisition function over the new model. The popular Kriging Believer method [16] uses the surrogate's mean prediction as the hallucinated value,

which reduces the model's posterior uncertainty to zero at the hallucinated locations without affecting the posterior mean.

An alternative to penalising the surrogate model via hallucination is to penalise an acquisition function in a region around the selected batch points [1, 17]. In these methods, the first point \mathbf{x}'_1 in a batch is selected via maximisation of a sequential acquisition function $\alpha(\mathbf{x})$, e.g. EI; the subsequent $q - 1$ locations are chosen by iteratively maximising a penalised version of the sequential acquisition function:

$$\mathbf{x}'_i = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \left\{ \alpha(\mathbf{x}) \prod_{j=1}^{i-1} \varphi(\mathbf{x} | \mathbf{x}'_j) \right\}, \quad i = 2, \dots, q \quad (5)$$

where $\varphi(\mathbf{x} | \mathbf{x}_j)$ are local penalisers centred at \mathbf{x}_j . These penalise a region around \mathbf{x}_j with decreasing penalisation as the distance from \mathbf{x}_j increases; for example [17] use a squared exponential function. The length scale over which the penalisation is significant is set by

$$r_j = \frac{|\mu(\mathbf{x}'_j) - f_{\min}|}{L} + \gamma \frac{\sigma^2(\mathbf{x}'_j)}{L}, \quad (6)$$

where f_{\min} is equal to the global minimum of the function, L is a valid Lipschitz constant expressing how rapidly f can change with \mathbf{x} , and $\gamma \geq 0$ weights the importance of the uncertainty about \mathbf{x}_j .

In practise, the true value of f_{\min} is unknown and therefore the best seen value so far, $f^* = \min_i \{f_i\}_{i=1}^M$, is used in lieu. It can be shown [17] that $L_{\nabla} = \max_{\mathbf{x} \in \mathcal{X}} \|\nabla f(\mathbf{x})\|$ is a valid Lipschitz constant and González et al. [17] approximate this using the surrogate model's mean prediction: $\tilde{L} = \max_{\mathbf{x} \in \mathcal{X}} \|\nabla \mu(\mathbf{x})\|$, resulting in a global estimate of the largest gradient in the model that is fixed for all selected batch locations. Alvi et al. [1] argue that this under-penalises flatter regions of space in which the estimated gradient of the function is much smaller, and therefore they calculate a different value of \tilde{L} for each selected batch location, estimating it within a length-scale of each location. González et al. [17] set $\gamma = 0$ and focus only on the difference between the predicted value of \mathbf{x}_j and the global optimum, whereas Alvi et al. [1] let $\gamma = 1$ to also include prediction uncertainty. This penalty shrinks as the predicted value of \mathbf{x}_j approaches the global minimum and also as the largest local (or global) gradient of the model increases.

Motivated by the success of the sequential exploitative and ϵ -greedy approaches [10, 33], we invert the local penalisation strategy and, instead, present a method that samples from *within* the region that would usually be penalised (6). We empirically show that this approach out-performs recent BBO methods on a range of synthetic functions and two real-world problems.

3 ϵ -SHOTGUN BBO

Motivated by the recent success of ϵ -greedy methods, we extend the two sequential methods of De Ath et al. [10] to the batch setting. We use an ϵ -greedy acquisition function to generate the first batch location \mathbf{x}'_1 and then sample the remaining locations from a normal distribution centred on \mathbf{x}'_1 , with a standard deviation given by:

$$r = \frac{|\mu(\mathbf{x}'_1) - f_{\min}|}{L} + \gamma \frac{\sigma^2(\mathbf{x}'_1)}{L}. \quad (7)$$

Sampling in this manner creates a scattered set of batch points around \mathbf{x}'_1 , akin to a shotgun blast, whose approximate spread is

Algorithm 1 ϵ -shotgun query point selection for BBO

Inputs:
 q : Batch size
 ϵ : Proportion of the time to explore
 l : Kernel length scale

```

1: if rand() <  $\epsilon$  then
2:   if Using Pareto front selection then                                ▶  $\epsilon$ S-PF
3:      $\tilde{\mathcal{P}} \leftarrow \text{MOOptimise}_{\mathbf{x} \in \mathcal{X}}(\mu(\mathbf{x}), \sigma^2(\mathbf{x}))$ 
4:      $\mathbf{x}'_1 \leftarrow \text{randomChoice}(\tilde{\mathcal{P}})$ 
5:   else                                                                ▶  $\epsilon$ S-RS
6:      $\mathbf{x}'_1 \leftarrow \text{randomChoice}(\mathcal{X})$ 
7:   else
8:      $\mathbf{x}'_1 \leftarrow \underset{\mathbf{x} \in \mathcal{X}}{\text{argmin}} \mu(\mathbf{x})$ 
9:      $\tilde{L} = \max_{\mathbf{x} \in [\mathbf{x}'_1 - l, \mathbf{x}'_1 + l]^d} \|\mu_{\nabla}(\mathbf{x})\|$  ▶ Largest gradient; centred on  $\mathbf{x}'_1$ 
10:     $f^* = \min \{f_i\}_{i=1}^M$  ▶ Best seen function value
11:     $r = \frac{|\mu(\mathbf{x}'_1) - f^*|}{\tilde{L}} + \gamma \frac{\sigma(\mathbf{x}'_1)}{\tilde{L}}$ 
12:     $\mathcal{X}' = \{\mathbf{x}'_1\}$ 
13:    while  $|\mathcal{X}'| < q$  do
14:       $\mathbf{x}' \sim \mathcal{N}(\mathbf{x}'_1, r^2 \mathbf{I})$ 
15:      if  $\mathbf{x}' \in \mathcal{X}'$  then
16:         $\mathcal{X}' \leftarrow \mathcal{X}' \cup \{\mathbf{x}'\}$ 
17:    return  $\mathcal{X}'$ 
    
```

determined by the amount of model uncertainty of \mathbf{x}'_1 , its predicted value relative to the best seen function evaluation, and by the steepest gradient within its vicinity.

We describe two alternative strategies employing this idea, which are summarised in Algorithm 1. The first method, which we call ϵ -shotgun with Pareto front selection (ϵ S-PF), selects with probability $1 - \epsilon$ the location \mathbf{x}'_1 with the most promising mean prediction from the surrogate model (line 8). In the remaining cases it selects a random element from the approximate Pareto set $\tilde{\mathcal{P}}$, which is found using an evolutionary multi-objective optimiser (lines 3 and 4).

Following the selection of the \mathbf{x}'_1 , ϵ S-PF samples the remaining $q - 1$ locations from a normal distribution $\mathcal{N}(\mathbf{x}'_1, r^2 \mathbf{I})$ centred on \mathbf{x}'_1 with a standard deviation equal to the radius (7) of penalisation used in [1] (lines 9 to 16). We conservatively estimate the global optimum f_{\min} to be the best function evaluation seen so far, i.e. $f^* = \min_i \{f_i\}$. The localised Lipschitz constant, similarly to Alvi et al. [1], is estimated to be $\tilde{L} = \max_{\mathbf{x} \in \mathcal{H}} \|\nabla \mu(\mathbf{x})\|$, where \mathcal{H} is a hypercube, centred on \mathbf{x}'_1 with side lengths of twice the length-scale of the surrogate model's kernel. This allows for the local gradient to influence the size of the sampling region. If the local gradient is small then it is beneficial to sample over a wide region to learn more about the structure of f . Conversely, a steeper local gradient would indicate that the modelled function is changing rapidly and therefore sampling more densely (due to a larger \tilde{L}) is required to accurately model f and guide the search.

Figure 1 shows three example locations of \mathbf{x}'_1 for a surrogate model and their corresponding probability density functions from which samples would be drawn. The blue circle is located at the minimum of the modelled function, and its corresponding sampling

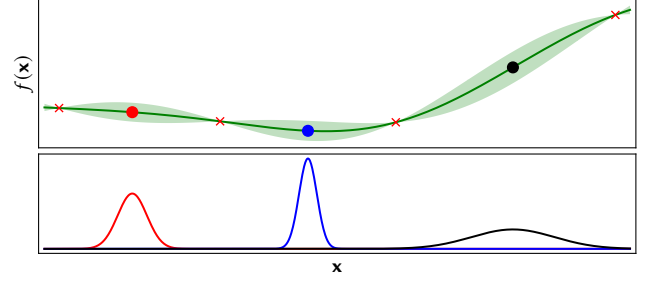


Figure 1: ϵ -shotgun selection example. The upper panel shows the predicted mean and uncertainty (green) of an unknown function sampled at four locations (red crosses). The lower panel shows the pdfs of $\mathcal{N}(\mathbf{x}'_1, r^2)$, centred on the three correspondingly coloured locations of \mathbf{x}'_1 in the upper figure.

radius is relatively small because $|\mu(\mathbf{x}'_1) - f^*|$ is small, the local gradient is small, and the predicted uncertainty is relatively large, resulting in a fairly sharp distribution to sample from. The red circle corresponds to a location with a sampling radius that is larger than the previous point, because, while the difference between the modelled function and f^* and the predicted uncertainty is similar to the blue location, the local gradient is smaller. Lastly, the black location has a similar model uncertainty to the red location, but has a much larger $|\mu(\mathbf{x}'_1) - f^*|$ resulting in a much wider pdf, even when taking into account the larger local gradient.

The second strategy, ϵ -shotgun with random selection (ϵ S-RS), is identical to ϵ S-PF except that, with probability ϵ , it selects \mathbf{x}'_1 at a random location in the entire feasible space (line 6), instead of a location from the approximate Pareto set. It might be expected that ϵ S-PF would outperform ϵ S-RS because it selects locations that are more informative to the optimisation process, because they are non-dominated with respect to their predicted value and uncertainty. However, since De Ath et al. [10] show that sequential ϵ -greedy methods based on selection from the Pareto set have only marginally better performance than purely random selection, we include ϵ S-RS to assess whether this is true in the batch setting.

4 EXPERIMENTAL EVALUATION

We investigate the performance of the two proposed ϵ -shotgun methods, ϵ S-PF and ϵ S-RS on ten well-known benchmark functions with a differing dimensionality and two real-world applications in the form of an active learning problem for robot pushing and pipe shape optimisation. Full results of all experimental evaluations are available in the supplementary material.

Following the reported success of purely exploitative methods [10, 33], we also compare ϵ S-PF and ϵ S-RS to the purely exploitative ϵ -shotgun method without any random point selection (i.e. $\epsilon = 0$), which we denote ϵ S-0. We also compare the batch ϵ -shotgun methods to five BBO methods representative of different styles of batch optimisation as discussed in Section 2.2: Two acquisition function-based penalisation methods: the popular Local Penalisation (LP) method [17], which uses soft penalisation ($\varphi(\mathbf{x}_j | \mathbf{x}_j) > 0$), and the more recent PLAYBOOK [1] method that uses hard local penalisation ($\varphi(\mathbf{x}_j | \mathbf{x}_j) = 0$). Kriging Believer (KB), which penalises

Name	d	Name	d
WangFreitas [39]	1	logSixHumpCamel [†]	2
Branin [†]	2	modHartman6 [†]	6
BraninForrester [14]	2	logGSobol [17]	10
Cosines [18]	2	logRosenbrock [†]	10
logGoldsteinPrice [†]	2	logStyblinskiTang [†]	10

Table 1: Synthetic functions used and their dimensionality d . Formulae can be found as cited or at <http://www.sfu.ca/~ssurjano/optimization.html> for those labelled with [†].

by hallucinating the already-selected batch points [16]; and the Thompson sampling (TS) method of Kandasamy et al. [26], which minimises a realisation of the modelled function from the surrogate. Lastly, we include qEI [15], which jointly estimates the location of the q batch members. All methods were implemented in Python using the same packages¹, apart from the local penalisers of LP and PLAyBOOK which used the PLAyBOOK implementation.²

A zero-mean Gaussian process surrogate model with an isotropic Matérn 5/2 kernel was used in all the experiments. The kernel was selected due to its widespread usage and recommended use for modelling realistic functions [36]. The models were initially trained on $2d$ observations generated by maximin Latin hypercube sampling [29], with each optimisation run repeated 51 times with different initialisations. The same sets of initial batch locations were common across all methods to enable statistical comparison. At each iteration, before batch point selection, the hyperparameters of the GP were optimised by maximising the log likelihood (4) with L-BFGS-B [3] using 10 restarts [19].

The LP, PLAyBOOK and KB methods all used the EI acquisition function. For each location selected in LP and PLAyBOOK, we followed the authors' guidelines [1] and uniformly sampled the acquisition function at 3000 locations, selecting the best location after locally optimising (with L-BFGS-B) the best 5. For the other methods, a maximum budget of $10000d$ acquisition function evaluations was used in conjunction with L-BFGS-B for functions with $d = 1$ and for $d \geq 2$ we used CMA-ES using the standard bi-population strategy [20] and (up to) 9 restarts. The approximate Pareto set $\tilde{\mathcal{P}}$ of non-dominated locations (in terms of $\mu(\mathbf{x})$ and $\sigma^2(\mathbf{x})$) in ϵ S-PF was found using NSGA-II [12] with a $100d$ population size, d^{-1} mutation rate, 0.8 crossover rate, and crossover and mutation distribution indices of $\eta_c = \eta_m = 20$. For both ϵ S-RS and ϵ S-PF we took $\epsilon = 0.1$.

4.1 Synthetic Experiments

The methods were evaluated on the 10 synthetic benchmark functions in Table 1 with batch sizes $q \in \{2, 5, 10, 20\}$ and a fixed budget of 200 function evaluations. Table 2 shows, for a batch size of $q = 10$, the median difference (over 51 repeated experiments) between the estimated optimum f^* and true optimum, as well as the median absolute deviation from the median (MAD), a robust measure of dispersion. The method with the minimum median f^* for each function is highlighted in dark grey, and those that are statistically equivalent to the best method according to a one-sided, paired

Wilcoxon signed-rank test [27] with Holm-Bonferroni correction [22] ($p \geq 0.05$), are shown in light grey. Note that tabulated results for all batch sizes are available in the supplementary material.

Figure 2 shows the convergence plots of the various algorithms on six test problems for $q \in \{5, 10, 20\}$. As might be expected, qEI tends to perform worse as q increases, which we suspect is linked to the dimensionality of the qEI acquisition function being $d \times q$. For the 10-dimensional functions and $q = 20$, this requires global optimisation in a 200-dimensional space, a far from trivial task. The Thompson sampling method (TS) relies upon there being sufficient stochasticity in the surrogate model to select batch locations that are well distributed in space. If there is too much or too little variation in the locations selected, as appears to be the case in these results, the batch will be selected in either locations with poor mean values (too much variation) or all at the same location (too little variation). Similar performance results for TS are also shown in [1].

As shown in the convergence plots and Table 2, the ϵ -shotgun batch algorithms, ϵ S-RS and ϵ S-PF, both performed well across the range of synthetic problems for all batch sizes. ϵ S-0, which always samples at and around the surrogate's best mean prediction, also performed well across the majority of synthetic functions and was statistically equivalent to ϵ S-PF on all functions with a batch size of $q = 10$. This indicates that fully exploiting the model at each iteration and learning about the best mean prediction's local landscape (via sampling its local neighbourhood) is a sound strategy and mirrors the findings, that being greedy is good, of Rehbach et al. [33] and De Ath et al. [10] in the sequential setting.

Interestingly, on the modHartman6 function in particular (Figure 2, lower-left), $q = 20$ led to better median f^* than for $q = 5$, even though there were 4 times fewer batches (10 instead of 40) and therefore the surrogate model was fitted far fewer times. This indicates that the model poorly estimated the underlying function, thus misleading the optimisation process. However, the expected trend prevails: an increase in q generally led to a decrease in the median f^* as well as a decrease in the rate of convergence.

The acquisition penalisation-based methods, LP and PLAyBOOK, performed similarly, with LP slightly ahead of PLAyBOOK. The dominating factor setting the penalisation radii (6) in both methods is the Lipschitz constant, which was estimated as being the largest value of $\|\nabla\mu(\mathbf{x})\|$ over the whole problem domain for LP and locally for PLAyBOOK. Since the global Lipschitz constant will always be at least as large as a local one, it is perhaps unsurprising that LP performs better, because a larger constant corresponds to a smaller radius of penalisation, meaning that the batch points will be, on average, closer together and, therefore, more similar to the better-performing ϵ -shotgun batch methods.

Convergence plots for ϵ S-RS and ϵ S-PF have a well-defined step-like appearance for several test functions, which is particularly visible in the plots with larger batch sizes. This is a consequence of the batch selection process because the first location in the batch \mathbf{x}'_1 minimises the surrogate model's mean function (recall Algorithm 1, line 8). It does, however, imply that the sequential ϵ -greedy strategy is driving the optimisation process as the subsequent evaluations in the batch generally show little improvement over $f(\mathbf{x}'_1)$. This also means that the locations sampled around \mathbf{x}'_1 are useful because they improve the surrogate model accuracy, allowing the surrogate's mean prediction to drive the optimisation.

¹Implementation available: <https://github.com/georgedeath/eshotgun>

²<https://github.com/a5a/asynchronous-BO>

Method	WangFreitas (1)		BraninForrester (2)		Branin (2)		Cosines (2)		logGoldsteinPrice (2)	
	Median	MAD	Median	MAD	Median	MAD	Median	MAD	Median	MAD
LP	2.00	3.08×10^{-9}	2.61×10^{-5}	3.86×10^{-5}	9.25×10^{-6}	1.23×10^{-5}	1.09×10^{-3}	1.54×10^{-3}	5.26×10^{-4}	5.86×10^{-4}
PLAyBOOK	2.00	4.76×10^{-10}	1.25×10^{-4}	1.81×10^{-4}	1.79×10^{-5}	2.51×10^{-5}	3.70×10^{-3}	4.35×10^{-3}	6.48×10^{-4}	8.50×10^{-4}
KB	2.00	1.19×10^{-9}	2.31×10^{-3}	3.32×10^{-3}	3.03×10^{-5}	3.28×10^{-5}	1.09×10^{-3}	1.41×10^{-3}	4.75×10^{-2}	5.92×10^{-2}
qEI	1.12×10^{-7}	1.55×10^{-7}	5.83×10^{-6}	7.37×10^{-6}	7.84×10^{-6}	6.94×10^{-6}	9.49×10^{-5}	1.35×10^{-4}	1.82×10^{-4}	1.89×10^{-4}
TS	2.00	3.02×10^{-8}	4.58×10^{-4}	4.77×10^{-4}	1.94×10^{-4}	2.15×10^{-4}	1.28×10^{-3}	1.14×10^{-3}	1.78×10^{-3}	1.59×10^{-3}
ϵ S-RS (0.1)	2.00	3.87×10^{-11}	6.07×10^{-7}	7.70×10^{-7}	1.51×10^{-6}	1.60×10^{-6}	1.07×10^{-6}	1.28×10^{-6}	6.65×10^{-7}	8.96×10^{-7}
ϵ S-PF (0.1)	2.00	1.66×10^{-12}	1.20×10^{-6}	1.77×10^{-6}	1.91×10^{-6}	1.89×10^{-6}	4.21×10^{-7}	5.64×10^{-7}	3.27×10^{-7}	4.51×10^{-7}
ϵ S-0	2.00	1.18×10^{-11}	9.89×10^{-7}	1.28×10^{-6}	1.70×10^{-6}	1.83×10^{-6}	4.12×10^{-7}	4.66×10^{-7}	3.23×10^{-7}	4.62×10^{-7}

Method	logSixHumpCamel (2)		modHartman6 (6)		logGSobol (10)		logRosenbrock (10)		logStyblinskiTang (10)	
	Median	MAD	Median	MAD	Median	MAD	Median	MAD	Median	MAD
LP	2.22×10^{-1}	2.59×10^{-1}	8.25×10^{-4}	1.04×10^{-3}	7.58	1.96	5.97	8.36×10^{-1}	2.07	3.76×10^{-1}
PLAyBOOK	1.88×10^{-1}	2.41×10^{-1}	2.11×10^{-3}	2.56×10^{-3}	9.82	1.45	5.98	1.48	2.28	2.89×10^{-1}
KB	4.72	1.35	7.33×10^{-3}	7.64×10^{-3}	7.21	1.65	5.29	2.19	1.96	3.08×10^{-1}
qEI	1.49×10^{-1}	1.09×10^{-1}	1.44×10^{-2}	9.65×10^{-3}	9.84	2.11	7.94	5.00×10^{-1}	2.28	2.04×10^{-1}
TS	1.15	6.79×10^{-1}	3.94×10^{-2}	1.60×10^{-2}	1.03×10^1	7.49×10^{-1}	8.48	4.55×10^{-1}	2.87	1.07×10^{-1}
ϵ S-RS (0.1)	1.38×10^{-3}	2.04×10^{-3}	3.08×10^{-4}	3.87×10^{-4}	8.07	2.53	5.03	1.58	2.05	3.64×10^{-1}
ϵ S-PF (0.1)	3.90×10^{-4}	5.71×10^{-4}	3.09×10^{-4}	3.09×10^{-4}	8.19	1.88	4.61	1.43	1.81	4.56×10^{-1}
ϵ S-0	1.15×10^{-3}	1.70×10^{-3}	4.24×10^{-4}	4.90×10^{-4}	7.40	2.23	4.45	1.44	1.81	3.55×10^{-1}

Table 2: Optimisation results with a batch size of $q = 10$. Median absolute distance from the optimum (left) and median absolute deviation from the median (MAD, right) after 20 batches (200 function evaluations) across the 51 runs. The method with the lowest median performance is shown in dark grey, with those with statistically equivalent performance shown in light grey.

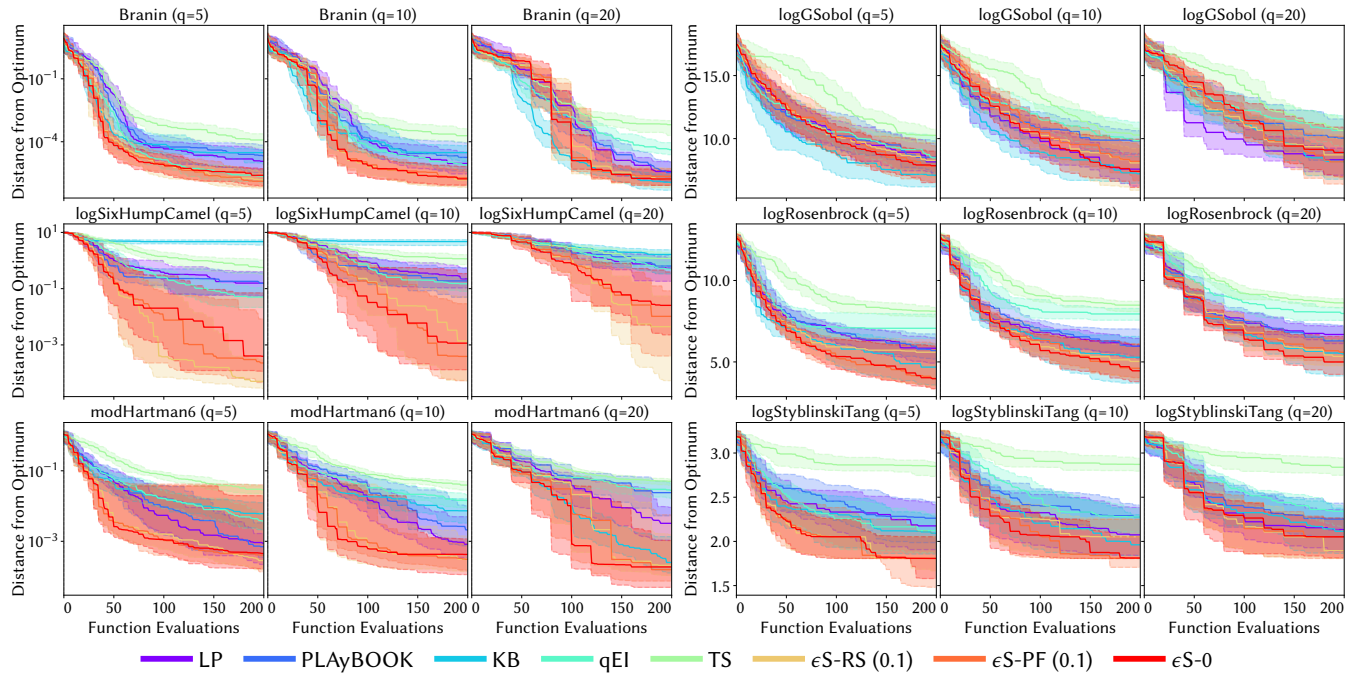


Figure 2: Illustrative convergence plots for six benchmark problems and three batch sizes $q \in \{5, 10, 20\}$. Each plot shows the median difference between the best function value seen f^* and the true optimum f_{\min} , with shading representing the interquartile range across the 51 runs.

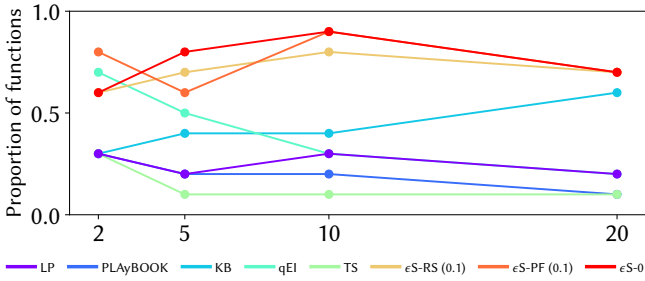


Figure 3: Synthetic function optimisation summary. Symbols correspond to the proportion of times that a method is best or statistically equivalent to the best method across the 10 synthetic functions for $q \in \{2, 5, 10, 20\}$.

Figure 3 summarises the performance of each of the 7 evaluated methods for the 4 batch sizes. We note that the ϵ -shotgun methods are consistently the best or statistically indistinguishable from the best performing methods across the set of benchmark functions across all batch sizes. Interestingly, the older Kriging Believer [16], that penalises the surrogate model’s variance around selected batch points, performed better than the newer, acquisition-based penalisers, particularly for larger batch sizes. The increase in relative performance may be related to the particular acquisition function used because, as shown in [10], EI weights improvements over the current f^* much more highly than increases in variance. This may lead to the variance penalisation in KB having a smaller radius of effect than the penalisation in EI-space by the LP and PLAyBOOK methods, resulting in KB sampling locations closer together, in a more similar fashion to the ϵ -shotgun methods.

As shown in Figure 3, for the ϵ -shotgun-based algorithms, there is little to differentiate overall between selecting a location at random from either the Pareto front (ϵ S-PF) or uniformly across the feasible space (ϵ S-RS). However, ϵ S-PF appears to be marginally better on lower-dimensional functions, most likely due to the surrogate model better describing the overall structure of the modelled function. Conversely, ϵ S-RS is slightly better on higher-dimensional functions because the modelled function with naturally be of a poorer quality and therefore relying solely on it, without sufficient stochasticity, could hinder the optimisation process.

Pure exploitation, i.e. $\epsilon = 0$, the ϵ S-0 method, leads to state-of-the-art performance across for many problems and dimensionalities. However, for problems in which a large amount of exploration is needed in order to locate a deceptive optimum, the ϵ -shotgun methods with little exploration are unable to escape local minima or expend enough of their optimisation budget exploring the landscape. This is particularly apparent on the WangFreitas problem [39], which has a large, shallow local minimum and a narrow, deep global minimum surrounded by plateaus; see [39] for a plot. Figure 4 compares the performance of the ϵ -shotgun methods for different ϵ on this problem. Note how ϵ S-RS (green) is able to more consistently find the global optimum for smaller values of ϵ , because, unlike ϵ S-PF (red), it is not constrained to only select non-dominated areas of decision space. ϵ S-PF, on the other hand, is consistently misled by the surrogate model’s incorrect estimation of the function and therefore fails to correctly optimise the function even when $\epsilon = 0.5$.

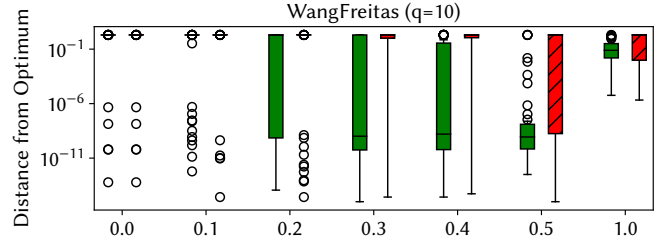


Figure 4: Distribution of $|f_{\min} - f^*|$ after 200 function evaluations, taken over 51 runs, for ϵ S-RS (green) and ϵ S-PF (red, hatched) for different values of ϵ (horizontal axis) on the WangFreitas test problem with a batch size of $q = 10$.

4.2 Active Learning for Robot Pushing

Following [10, 24, 41], we optimise the control parameters for two active learning robot pushing problems [40]; see [10] for diagrams. In the first problem, PUSH4, a robot is required to push an object towards an unknown target location. It receives the object-target distance once it has finished pushing. Its movement is constrained such that it can only travel in the direction towards the object’s initial location. The parameters to be optimised are the robot’s initial location, the orientation of its hand and for how long it travels. Thus optimising the values of the four parameters to reduce the final object-target distance can be cast as a minimisation problem. The object’s initial location is always set to be the centre of the domain [10, 41], but the target location is changed in each optimisation run, with these common across methods to ensure fair comparison. The performance of an optimisation algorithm is thus averaged over problem instances rather than the same function with different initialisations, as with the synthetic functions previously.

Similarly, in the second problem PUSH8, two robots push their own objects towards unknown targets, with the complication that they may block each other’s path. The 8 parameters controlling both robots can be optimised to minimise the summed final object-target distances. Initial object locations were fixed and target’s positions were generated randomly, while ensuring that each pair of target positions allowed each object to touch its target without the objects overlapping. However, in some problem instances it is not possible for both robots to push their objects to their targets because the objects may be positioned such that the robots need to cross each other’s paths. In order to report the difference between f^* and the true optimum we estimate the optimum of each problem instance by randomly sampling in the feasible space with 10^5 sets of robot parameters and locally optimising the best 100 of these using L-BFGS-B. We therefore report the difference between the algorithm’s optimum and this estimated global optimum.

Figure 5 shows the convergence for $q \in \{5, 10, 20\}$. In PUSH4 the ϵ -shotgun methods have statistically equivalent performance, but outperform other methods for $q = 5$ and $q = 10$. KB also has statistically similar performance to the exploitative methods when $q = 20$, echoing its efficacy on the synthetic problems in which it also comparatively improved with increasing q .

In the harder PUSH8, all methods were statistically equivalent with $q = 10$, while TS and KB were worse with $q = 5$; they were also worse, along with LP for $q = 20$. ϵ S-RS and ϵ S-PF consistently

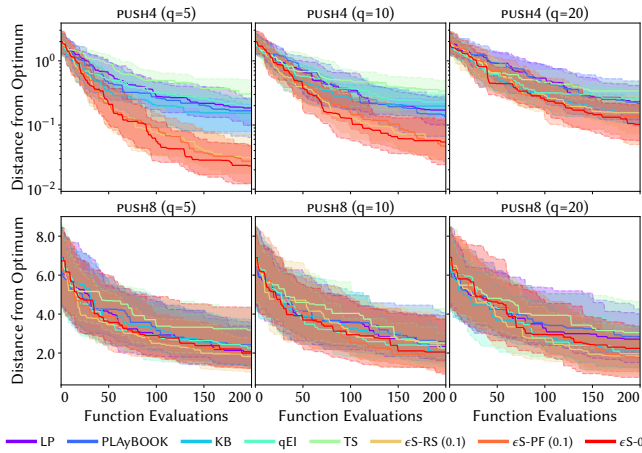


Figure 5: Convergence plots for the robot pushing problems (rows) over three batch sizes $q \in \{5, 10, 20\}$ (columns). Each plot shows the median value of $|f^* - f_{\min}|$, with shading representing the interquartile range across the 51 runs.

had the lowest median fitnesses, although other techniques were statistically equivalent. Interestingly, and echoing the results on the synthetic modHartman6 function, ϵ S-PF had a lower median fitness value for $q = 10$ and $q = 20$ than for the smaller batch sizes.

4.3 Pipe Shape Optimisation

We also evaluated the BBO methods on a real-world computational fluid dynamics (CFD) design problem. The PitzDaily test problem [7], involves reducing the pressure loss along a pipe of different inflow and outflow diameters by optimising the pipe’s internal shape. The optimisation aims to find the shape of the lower wall of the pipe that minimises the pressure loss. The loss is evaluated by running a CFD mesh generation and partial differential equation simulation of the two-dimensional flow. Each function evaluation takes between 60s and 90s, depending on mesh complexity.

The pipe’s lower wall geometry is represented by a Catmull-Clark sub-division curve, whose control points comprise the decision variables. We use 5 control points, resulting in a 10-dimensional decision vector. The control points are constrained to lie within a polygon and the initial locations used in the optimisation runs are uniformly sampled in this constrained domain. Similarly, for the batch optimisation methods themselves, we take the naive approach of rejection sampling when optimising acquisition functions for the KB, qEI and the ϵ -shotgun approaches; we do not consider locations that violate the constraints for LP, PLAyBOOK and TS.

Convergence plots of the flow loss with $q \in \{5, 10, 20\}$ are shown in Figure 6. The Kriging Believer (KB) consistently optimised the problem well, although ϵ S-RS, ϵ S-PF, ϵ S-0 and LP were statistically equivalent with batch sizes $q = 10$ and $q = 20$. The rates at which KB reaches its best flow losses, however, were superior to the other methods on the PitzDaily problem. In the $q = 10$ case, KB reaches close to its best flow losses after only 5 batches. We note that all methods were able to discover pipe shape configurations that led to flow losses that were better than 0.0903 found by an adjoint (local,

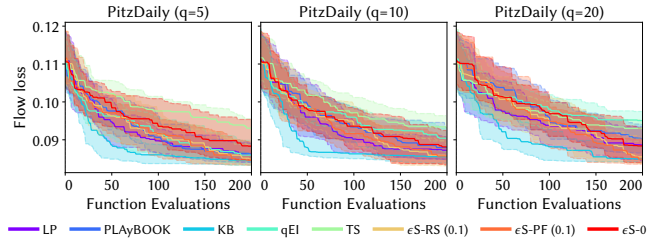


Figure 6: Convergence plots for the PitzDaily problem with $q \in \{5, 10, 20\}$. Each plot shows the median best seen flow loss, with shading representing the IQR across the 51 runs.

gradient-based) optimisation method [30], although TS, qEI, and PLAyBOOK were not able to reach a median flow loss of lower than the adjoint solution for all batch sizes.

5 CONCLUSIONS AND FUTURE WORK

Our novel ϵ -shotgun method, which uses an ϵ -greedy acquisition function to select the first batch location and samples the remaining locations around it, is both conceptually simple and computationally efficient because, unlike many other batch methods, only one global optimisation run is needed to select the batch locations. The method is competitive with state-of-the-art BBO algorithms and better than them in several cases. We attribute this to the exploitative nature of the first batch point selected together with benefit of learning an accurate function model with the remainder of the batch.

Pure exploitation ($\epsilon = 0$: ϵ S-0) led to good performance on the majority of problems because the surrogate model poorly estimates the true function, particularly on higher-dimensional functions, thus inducing enough exploration. However, in the case of degenerate functions, e.g. WangFreitas, the surrogate model is too poor to optimise well, requiring a larger ϵ to promote exploration. We have found that $\epsilon = 0.1$ works well.

Future research questions revolve around how best to select the locations of the $q - 1$ batch points. Although not described in detail here, we also investigated an alternative ϵ -shotgun approach of always selecting the first batch location as the surrogate model’s best mean prediction and dividing the remaining samples into two groups. One group, selected with probability $1 - \epsilon$ for each location, was used identically to the greedy shotgun approach of sampling around the first batch location and the second group (probability ϵ) were randomly sampled in space or on the approximated Pareto front. This approach gave similar results to ϵ -shotgun.

One possible extension is use Latin hypercube sampling instead of drawing random samples to better spread out the batch locations, making the best use of each function evaluation. In addition, it may be beneficial to tailor the sampling covariance to sample less/more densely in directions that are flatter/steeper in decision space.

Lastly, the function evaluation may take different times depending on location and computational hardware, resulting in some evaluations finishing before others. Current research, therefore, focuses on extending the ϵ -shotgun method to asynchronous BBO.

ACKNOWLEDGMENTS

This work was supported by Innovate UK [grant number 104400].

REFERENCES

- [1] Ahsan Alvi, Binxin Ru, Jan-Peter Calliess, Stephen Roberts, and Michael A. Osborne. 2019. Asynchronous Batch Bayesian Optimisation with Improved Local Penalisation. In *Proceedings of the 36th International Conference on Machine Learning*, Vol. 97. PMLR, 253–262.
- [2] Javad Azimi, Alan Fern, and Xiaoli Z. Fern. 2010. Batch Bayesian Optimization via Simulation Matching. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 109–117.
- [3] Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16, 5 (1995), 1190–1208.
- [4] Yutian Chen, Aja Huang, Ziyu Wang, Ioannis Antonoglou, Julian Schrittwieser, David Silver, and Nando de Freitas. 2018. Bayesian Optimization in AlphaGo. arXiv:1812.06855
- [5] Clément Chevalier and David Ginsbourger. 2013. Fast computation of the multi-points expected improvement with applications in batch selection. In *International Conference on Learning and Intelligent Optimization*. Springer, 59–69.
- [6] T. Chugh, K. Sindhya, K. Miettinen, Yaochu Jin, T. Kratky, and P. Makkonen. 2017. Surrogate-assisted evolutionary multiobjective shape optimization of an air intake ventilation system. In *IEEE Congress on Evolutionary Computation*. IEEE, 1541–1548.
- [7] Steven J. Daniels, Alma A. M. Rahat, Richard M. Everson, Gavin R. Tabor, and Jonathan E. Fieldsend. 2018. A Suite of Computationally Expensive Shape Optimisation Problems Using Computational Fluid Dynamics. In *Parallel Problem Solving from Nature – PPSN XV*. Springer, 296–307.
- [8] S. J. Daniels, A. A. M. Rahat, G. R. Tabor, J. E. Fieldsend, and R. M. Everson. 2019. Automated shape optimisation of a plane asymmetric diffuser using combined Computational Fluid Dynamic simulations and multi-objective Bayesian methodology. *International Journal of Computational Fluid Dynamics* 33, 6–7 (2019), 256–271.
- [9] Erik A. Daxberger and Bryan Kian Hsiang Low. 2017. Distributed Batch Gaussian Process Optimization. In *Proceedings of the 34th International Conference on Machine Learning*. PMLR, 951–960.
- [10] George De Ath, Richard M. Everson, Alma A. M. Rahat, and Jonathan E. Fieldsend. 2019. Greed is Good: Exploration and Exploitation Trade-offs in Bayesian Optimisation. arXiv:1911.12809
- [11] Alessandro De Palma, Celestine Mandler-Dünner, Thomas Parnell, Andreea Anghel, and Haralampos Pozidis. 2019. Sampling Acquisition Functions for Batch Bayesian Optimization. arXiv:1903.09434
- [12] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. 2001. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2001), 182–197.
- [13] Thomas Desautels, Andreas Krause, and Joel W. Burdick. 2014. Parallelizing Exploration-Exploitation Tradeoffs in Gaussian Process Bandit Optimization. *Journal of Machine Learning Research* 15 (2014), 4053–4103.
- [14] Alexander I. J. Forrester, Andras Sobester, and Andy J. Keane. 2008. *Engineering Design via Surrogate Modelling - A Practical Guide*. Wiley.
- [15] David Ginsbourger, Rodolphe Le Riche, and Laurent Carraro. 2008. A Multi-points Criterion for Deterministic Parallel Global Optimization based on Gaussian Processes. <https://hal.archives-ouvertes.fr/hal-00260579/>
- [16] David Ginsbourger, Rodolphe Le Riche, and Laurent Carraro. 2010. Kriging is well-suited to parallelize optimization. In *Computational intelligence in expensive optimization problems*. Springer, 131–162.
- [17] Javier González, Zhenwen Dai, Philipp Hennig, and Neil Lawrence. 2016. Batch Bayesian optimization via local penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, Vol. 51. PMLR, 648–657.
- [18] Javier González, Michael Osborne, and Neil Lawrence. 2016. GLASSES: Relieving the myopia of Bayesian optimisation. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, Vol. 51. PMLR, 790–799.
- [19] GPy. since 2012. GPy: A Gaussian process framework in Python. <http://github.com/SheffieldML/GPy>.
- [20] Nikolaus Hansen. 2009. Benchmarking a BI-population CMA-ES on the BBOB-2009 function testbed. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 2389–2396.
- [21] José Miguel Hernández-Lobato, James Requeima, Edward O. Pyzer-Knapp, and Alán Aspuru-Guzik. 2017. Parallel and Distributed Thompson Sampling for Large-scale Accelerated Exploration of Chemical Space. In *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70. PMLR, 1470–1479.
- [22] Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70.
- [23] Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown. 2011. Sequential Model-Based Optimization for General Algorithm Configuration. In *Learning and Intelligent Optimization*, Carlos A. Coello Coello (Ed.). Springer, 507–523.
- [24] Shali Jiang, Henry Chai, Javier González, and Roman Garnett. 2019. Efficient nonmyopic Bayesian optimization and quadrature. arXiv:1909.04568
- [25] Donald R. Jones, Matthias Schonlau, and William J. Welch. 1998. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization* 13, 4 (1998), 455–492.
- [26] Kirthevasan Kandasamy, Akshay Krishnamurthy, Jeff Schneider, and Barnabas Poczos. 2018. Parallelised Bayesian Optimisation via Thompson Sampling. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, Vol. 84. PMLR, 133–142.
- [27] Joshua D. Knowles, Lothar Thiele, and Eckart Zitzler. 2006. *A Tutorial on the Performance Assessment of Stochastic Multiobjective Optimizers*. Technical Report TIK214. Computer Engineering and Networks Laboratory, ETH Zurich, Zurich, Switzerland.
- [28] Wenlong Lyu, Fan Yang, Changhao Yan, Dian Zhou, and Xuan Zeng. 2018. Batch Bayesian Optimization via Multi-objective Acquisition Ensemble for Automated Analog Circuit Design. In *Proceedings of the 35th International Conference on Machine Learning*, Vol. 80. PMLR, 3306–3314.
- [29] Micheal D. McKay, Richard J. Beckman, and William J. Conover. 2000. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 42, 1 (2000), 55–61.
- [30] Ulf Nilsson, Daniel Lindblad, and Olivier Petit. 2014. *Description of adjointShapeOptimizationFoam and how to implement new objective functions*. Technical Report. Chalmers University of Technology, Gothenburg, Sweden.
- [31] Alma A.M. Rahat, Chunlin Wang, Richard M. Everson, and Jonathan E. Fieldsend. 2018. Data-driven multi-objective optimisation of coal-fired boiler combustion systems". *Applied Energy* 229 (2018), 446–458.
- [32] Carl Edward Rasmussen and Christopher K. I. Williams. 2006. *Gaussian processes for machine learning*. The MIT Press, Boston, MA.
- [33] Frederik Rehbach, Martin Zaefferer, Boris Naujoks, and Thomas Bartz-Beielstein. 2020. Expected Improvement versus Predicted Value in Surrogate-Based Optimization. arXiv:2001.02957
- [34] Amar Shah and Zoubin Ghahramani. 2015. Parallel Predictive Entropy Search for Batch Global Optimization of Expensive Objective Functions. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 3330–3338.
- [35] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. 2016. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* 104, 1 (2016), 148–175.
- [36] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2951–2959.
- [37] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. 2010. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*. Omnipress, 1015–1022.
- [38] Jialei Wang, Scott C. Clark, Eric Liu, and Peter I. Frazier. 2016. Parallel Bayesian Global Optimization of Expensive Functions. arXiv:1602.05149
- [39] Ziyu Wang and Nando de Freitas. 2014. Theoretical analysis of Bayesian optimization with unknown Gaussian process hyper-parameters. arXiv:1406.7758
- [40] Zi Wang, Caelan Reed Garrett, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. 2018. Active Model Learning and Diverse Action Sampling for Task and Motion Planning. In *Proceedings of the International Conference on Intelligent Robots and Systems*. IEEE, 4107–4114.
- [41] Zi Wang and Stefanie Jegelka. 2017. Max-value entropy search for efficient Bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning*. PMLR, 3627–3635.
- [42] Jian Wu and Peter Frazier. 2016. The Parallel Knowledge Gradient Method for Batch Bayesian Optimization. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 3126–3134.