Greed is Good: Exploration and Exploitation Trade-offs in Bayesian Optimisation

GEORGE DE ATH and RICHARD M. EVERSON, University of Exeter, United Kingdom
ALMA A. M. RAHAT, Swansea University, United Kingdom
JONATHAN E. FIELDSSEND, University of Exeter, United Kingdom

1

The performance of acquisition functions for Bayesian optimisation to locate the global optimum of continuous functions is investigated in terms of the Pareto front between exploration and exploitation. We show that Expected Improvement (EI) and the Upper Confidence Bound (UCB) always select solutions to be expensively evaluated on the Pareto front, but Probability of Improvement is not guaranteed to do so and Weighted Expected Improvement does so only for a restricted range of weights.

We introduce two novel ϵ -greedy acquisition functions. Extensive empirical evaluation of these together with random search, purely exploratory, and purely exploitative search on 10 benchmark problems in 1 to 10 dimensions shows that ϵ -greedy algorithms are generally at least as effective as conventional acquisition functions (e.g., EI and UCB), particularly with a limited budget. In higher dimensions, ϵ -greedy approaches are shown to have improved performance over conventional approaches. These results are borne out on a real-world computational fluid dynamics optimisation problem and a robotics active learning problem. Our analysis and experiments suggest that the most effective strategy, particularly in higher dimensions, is to be mostly greedy, occasionally selecting a random exploratory solution.

CCS Concepts: • **Computing methodologies** → **Optimization algorithms**; • **Theory of computation** → **Optimization with randomized search heuristics**; **Nonconvex optimization**;

Additional Key Words and Phrases: Bayesian optimisation, acquisition function, infill criteria, ϵ -greedy, exploration-exploitation trade-off

ACM Reference format:

George De Ath, Richard M. Everson, Alma A. M. Rahat, and Jonathan E. Fieldsend. 2021. Greed is Good: Exploration and Exploitation Trade-offs in Bayesian Optimisation. *ACM Trans. Evol. Learn. Optim.* 1, 1, Article 1 (April 2021), 22 pages.

<https://doi.org/10.1145/3425501>

1 INTRODUCTION

Global function optimisers search for the minimum or maximum of a function by querying its value at selected locations. All optimisers must therefore balance exploiting knowledge of the

This work was supported by Innovate UK Grant No. 104400.

Authors' addresses: G. De Ath, R. M. Everson, and J. E. Fieldsend, Department of Computer Science, Innovation Centre I, University of Exeter, Rennes Drive, EX4 4RN; emails: {g.de.ath, r.m.everson, j.e.fieldsend}@exeter.ac.uk; A. A. M. Rahat, Department of Computer Science, Swansea University, Room 112, Computational Foundry, Bay Campus, Swansea University, Fabian Way, Swansea, SA1 8EN; email: a.a.m.rahat@swansea.ac.uk.

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.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2688-299X/2021/04-ART1 \$15.00

<https://doi.org/10.1145/3425501>

function gained from the evaluations thus far with exploring other regions in which the landscape is unknown and might hold better solutions. This balance is particularly acute when a limited budget of function evaluations is available, as is often the case in practical problems; see, e.g., Jones et al. [1998] and Shahriari et al. [2016]. Bayesian optimisation is an effective form of surrogate-assisted optimisation in which a probabilistic model of the function is constructed from the evaluations made so far. The location at which the function is next (expensively) evaluated is chosen as the location that maximises an *acquisition function*, which makes the balance between exploration and exploitation explicit by combining the predicted function value at a location with the uncertainty in that prediction.

Here, we regard the balance between exploration and exploitation as itself a two-objective optimisation problem. We show that many, but not all, common acquisition functions effectively select from the Pareto front between objectives quantifying exploration and exploitation. In common with Bischl et al. [2014], Feng et al. [2015], Grobler et al. [2017], and Žilinskas and Calvin [2019], we propose choosing the next location to be expensively evaluated from the estimated Pareto set of solutions found by a two-objective evolutionary optimisation of the exploration and exploitation objectives. We compare the performance of various methods for selecting from the estimated Pareto front and propose two new ϵ -greedy schemes that usually choose the solutions with the most promising (exploitative) value, but occasionally use an alternative solution selected at random from either the estimated Pareto set or the entire feasible space.

Our main contributions can be summarised as follows:

- We present a unified analysis of common acquisition functions in terms of exploration and exploitation and give the first detailed analysis of weighted expected improvement.
- We investigate the use of the exploration-exploitation trade-off front in selecting the next location to expensively evaluate in Bayesian optimisation.
- We present two novel ϵ -greedy acquisition functions for Bayesian optimisation as well as other acquisition functions that use the exploration-exploitation trade-off front.
- These methods are empirically compared on a variety of synthetic test problems and two real-world applications.
- We demonstrate that the ϵ -greedy approaches are at least as effective as the conventional acquisition functions on lower-dimensional problems and become superior as the number of decision variables increases.

We begin in Section 2 by briefly reviewing Bayesian optimisation together with Gaussian processes—which are commonly used for surrogate modelling of the function. We pay particular attention to acquisition functions and the way in which they balance exploration and exploitation. The exploration-exploitation trade-off is viewed through the lens of multi-objective optimisation in Section 2.3, which leads to the proposed ϵ -greedy schemes in Section 3.1. Extensive empirical evaluations on well-known test problems are presented in Section 4, along with comparisons on a real-world computational fluid dynamics optimisation and robot active learning problem.

2 BAYESIAN OPTIMISATION

Bayesian optimisation (BO) is a particular method of surrogate-assisted optimisation. In practice, it has proved to be a very effective approach for single-objective expensive optimisation problems with limited budget on the number of true function evaluations. It was first proposed by Kushner [1964] in the early 1960s, and later improved and popularised by Moćkus et al. [1978] and Jones et al. [1998]. A recent review of the topic can be found in Shahriari et al. [2016].

Without loss of generality, the optimisation problem may be expressed as

$$\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (1)$$

ALGORITHM 1: Standard Bayesian optimisation.**Inputs:** M : Number of initial samples T : Budget on the number of expensive evaluations**Steps:**

```

1:  $X \leftarrow \text{LatinHypercubeSampling}(X, M)$  ▷ Generate initial samples
2: for  $t = 1 \rightarrow M$  do
3:    $f_t \leftarrow f(\mathbf{x}_t)$  ▷ Expensively evaluate all initial samples
4:  $\mathcal{D} \leftarrow \{(X, \mathbf{f})\}$ 
5: for  $t = M + 1 \rightarrow T$  do
6:    $\theta \leftarrow \text{TrainGP}(\mathcal{D})$  ▷ Train a  $\mathcal{GP}$  model
7:    $\mathbf{x}' \leftarrow \arg\max_{\mathbf{x} \in X} \alpha(\mathbf{x}, \mathcal{D}, \theta)$  ▷ Maximise infill criterion
8:    $f' \leftarrow f(\mathbf{x}')$  ▷ Expensively evaluate  $\mathbf{x}'$ 
9:    $X \leftarrow X \cup \{\mathbf{x}'\}$  ▷ Augment data
10:   $\mathbf{f} \leftarrow \mathbf{f} \cup \{f'\}$ 
11:   $\mathcal{D} \leftarrow \{(X, \mathbf{f})\}$ 
12: return  $\mathcal{D}$ 

```

where $X \subset \mathbb{R}^d$ is the feasible space and $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Algorithm 1 outlines the standard Bayesian optimisation procedure. In essence, it is a global search strategy that sequentially samples the design space at likely locations of the global optimum taking into account not only the predictions of the surrogate model but also the uncertainty inherent in modelling the unknown function to be optimised [Jones et al. 1998]. It starts (line 1) with a space filling design (e.g., Latin hypercube sampling [McKay et al. 2000]) of the parameter space, constructed independent of the function space. The samples $X = \{\mathbf{x}_t\}_{t=1}^M$ from this initial design are then (expensively) evaluated with the function, $f_t = f(\mathbf{x}_t)$, to construct a training dataset from which the surrogate model may be learned. We denote the vector of evaluated samples by $\mathbf{f}$. Then, at each iteration of the main part of the algorithm, a regression model is trained using the function evaluations obtained thus far (line 6). In Bayesian optimisation the regression model is used to predict the most likely value of $f(\mathbf{x})$ at new locations but also the uncertainty in the model estimate. In common with most work on Bayesian optimisation, we use Gaussian process models (GPs), which subsume Kriging models, as regressors; these are described in Section 2.1. The choice of where to next evaluate f is made by finding the location that maximises an *acquisition function* or *infill criterion* $\alpha(\mathbf{x}, \mathcal{D}, \theta)$, which balances exploitation of good regions of design space found thus far with the exploration of promising regions indicated by the uncertainty in the surrogate's prediction. Various common infill criteria are discussed and analysed from a multi-objective point of view in Section 2.2. The design maximising the infill criterion, $\mathbf{x}'$ is often found by an evolutionary algorithm (line 7), which is able to repeatedly evaluate the computationally cheap infill criterion. Finally, $f(\mathbf{x}')$ is expensively evaluated and the training data $(X, \mathbf{f})$ augmented with $\mathbf{x}'$ and $f(\mathbf{x}')$ (lines 8 to 10). The process is repeated until the budget is exhausted.

2.1 Modelling with Gaussian Processes

Gaussian processes are commonly used to construct a surrogate model of $f(\mathbf{x})$, and we therefore briefly describe them here; a comprehensive introduction may be found in Rasmussen and Williams [2006]. In essence, a GP is a collection of random variables, and any finite number of these have a joint Gaussian distribution [Rasmussen and Williams 2006]. With data comprising $f(\mathbf{x})$ evaluated at M locations $\mathcal{D} = \{(\mathbf{x}_m, f_m \triangleq f(\mathbf{x}_m))\}_{m=1}^M$, the predictive probability for f at $\mathbf{x}$ is

a Gaussian distribution with mean $\mu(\mathbf{x})$ and variance $\sigma^2(\mathbf{x})$:

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

where the mean and variance are

$$\mu(\mathbf{x}) = \boldsymbol{\kappa}(\mathbf{x}, X)K^{-1}\mathbf{f}, \quad (3)$$

$$\sigma^2(\mathbf{x}) = \kappa(\mathbf{x}, \mathbf{x}) - \boldsymbol{\kappa}(\mathbf{x}, X)^{\top}K^{-1}\boldsymbol{\kappa}(X, \mathbf{x}). \quad (4)$$

Here $X \in \mathbb{R}^{M \times d}$ is the matrix of design locations and $\mathbf{f} \in \mathbb{R}^M$ is the corresponding vector of the true function evaluations; thus $\mathcal{D} = \{(X, \mathbf{f})\}$. The covariance matrix $K \in \mathbb{R}^{M \times M}$ represents the covariance function $\kappa(\mathbf{x}, \mathbf{x}'; \theta)$ evaluated for each pair of observations and $\boldsymbol{\kappa}(\mathbf{x}, X) \in \mathbb{R}^M$ is the vector of covariances between $\mathbf{x}$ and each of the observations; θ denotes the kernel hyperparameters. We use a flexible class of covariance functions embodied in the Matérn 5/2 kernel, as recommended for modelling realistic functions [Snoek et al. 2012]. Although it is beneficial to marginalise θ with respect to a prior distribution, here we follow standard practise and fix on a single value of the hyperparameters by maximising the log likelihood each time the data is augmented by a new expensive evaluation¹:

$$\log p(\mathcal{D} | \theta) = -\frac{1}{2} \log |K| - \frac{1}{2} \mathbf{f}^{\top} K^{-1} \mathbf{f} - \frac{M}{2} \log(2\pi). \quad (5)$$

Henceforth, we omit θ for notational simplicity, and assume that these are set by maximum likelihood estimates.

2.2 Infill Criteria and Multi-objective Optimisation

An infill criterion or acquisition function $\alpha(\mathbf{x}, \mathcal{D}, \theta)$ is a measure of quality that enables us to decide which locations $\mathbf{x}$ are promising and consequently where to expensively evaluate f . It is based on the prediction $p(f | \mathbf{x}, \mathcal{D})$ from the surrogate (GP) model that represents our belief about the unknown function f at a decision vector $\mathbf{x}$ based on the M observations $\mathcal{D}$. Although $\alpha(\mathbf{x}, \mathcal{D}, \theta)$ depends on $\mathcal{D}$ and on the hyperparameters (θ) of the GP, for clarity we suppress this dependence and write $\alpha(\mathbf{x})$. The predictive distribution Equation (2) is Gaussian, with mean and variance given by Equations (3) and (4). The predicted mean and uncertainty enable an infill criterion to strike a balance between myopic exploitation (concentrating on regions where the mean prediction $\mu(\mathbf{x})$ is large) and global exploration (concentrating on regions where the uncertainty $\sigma(\mathbf{x})$ about f is large). Since, in general both exploitation and exploration are desirable, we may view these as competing criteria: a location $\mathbf{x}$ that is both more exploitative and more exploratory than an alternative $\mathbf{x}'$ is to be preferred over $\mathbf{x}'$. Using the notation of multi-objective optimisation, a location $\mathbf{x}$ dominates $\mathbf{x}'$, written $\mathbf{x} > \mathbf{x}'$, iff $\mu(\mathbf{x}) \geq \mu(\mathbf{x}')$ and $\sigma(\mathbf{x}) \geq \sigma(\mathbf{x}')$ and they are not equal on both. We present BO procedures that select solutions from the Pareto optimal set of locations, namely, those that are not dominated by any other feasible locations:

$$\mathcal{P} = \{\mathbf{x} \in \mathcal{X} | \mathbf{x}' \not> \mathbf{x} \forall \mathbf{x}' \in \mathcal{X}\}, \quad (6)$$

where $\mathbf{x}' \not> \mathbf{x}$ indicates that $\mathbf{x}'$ does not dominate $\mathbf{x}$.

Figure 1 illustrates the approximate Pareto front, $\{(\mu(\mathbf{x}), \sigma(\mathbf{x})) | \mathbf{x} \in \mathcal{P}\}$, for a simple one-dimensional function. Note that the Pareto set is disjoint in $\mathcal{X}$ and in (μ, σ) space. The locations maximising three popular acquisition functions, Expected Improvement (EI), Upper Confidence

¹We use the L-BFGS algorithm with 10 restarts to estimate the hyper-parameters [GPY 2012].

Fig. 1. Example Pareto front: *Top*: Gaussian Process approximation to a function (blue dashed curve) resulting from the five observations shown; mean $\mu(x)$ is shown in dark green and twice the posterior standard deviation $\sigma(x)$ is shown as the light green envelopes. *Bottom*: 200 samples uniformly spaced in $\mathcal{X}$ plotted in μ, σ space. The non-dominated locations forming the Pareto front are shown in red and their locations marked above. Locations maximising the Expected Improvement, Upper Confidence Bound and Probability of Improvement acquisition functions are marked in both plots.

Bound (UCB), and Probability of Improvement (PI) are highlighted. The maximisers of EI and UCB are elements of the Pareto set, whereas the maximiser of PI is not.

We now present some of the most popular acquisition functions used in BO, and discuss how they achieve a balance between exploration and exploitation.

2.2.1 Upper Confidence Bound. An optimistic policy, first proposed by Lai and Robbins [1985] is to overestimate the mean with added uncertainty: this is known as the upper confidence bound infill criterion (UCB). A proof of convergence under appropriate assumptions is given in Srinivas et al. [2010]. The UCB acquisition function is a weighted sum of the mean prediction and uncertainty:

$$\alpha_{UCB}(\mathbf{x}) = \mu(\mathbf{x}) + \sqrt{\beta_t} \sigma(\mathbf{x}), \quad (7)$$

where $\sqrt{\beta_t} \geq 0$ is the weight, which generally depends upon the number of function evaluations, t . The addition of a multiple of the uncertainty means that the criterion prefers locations where the

Fig. 2. Contours of upper confidence bound (UCB, $\beta_t = 1$), expected improvement (EI) and probability of improvement (PI) as functions of predicted mean μ and uncertainty σ . Since the scale of α is immaterial, all three infill criteria have been mapped to $[0, 1]$.

mean is large (exploitation) or mean combined with the uncertainty is sufficiently large to warrant exploration.

When $\beta_t = 0$ UCB becomes a purely exploitative scheme and therefore the solution with the best predicted mean is evaluated expensively. Thus, it may rapidly converge to a local maximum prematurely. In contrast, when β_t is large, the optimisation becomes purely exploratory, evaluating the location where the posterior uncertainty (variance) is largest, which is equivalent to maximally reducing the overall predictive entropy of the model [Srinivas et al. 2010]. Consequently, it may eventually locate the global optima, but the rate of convergence may be very slow.

Some authors suggest tuning β_t during the course of the optimisation [Shahriari et al. 2016]; indeed, Srinivas et al.'s convergence proof depends on a particular schedule in which $\sqrt{\beta_t}$ increases like the logarithm of t , so that more weight is given to exploratory moves as the optimum is approached [Srinivas et al. 2010].

Clearly, UCB increases monotonically as either the mean prediction μ or the uncertainty σ increase; see Figure 2. Consequently, if a set $\mathcal{S}$ of candidate locations for expensive evaluation is available and α_{UCB} is used to select the location with maximum upper confidence bound, $\mathbf{x}' = \arg\max_{\mathbf{x} \in \mathcal{S}} \alpha_{UCB}(\mathbf{x})$, then $\mathbf{x}'$ is a member of the maximal non-dominated subset of $\mathcal{S}$; that is, there is no element of $\mathcal{S}$ that dominates $\mathbf{x}'$. We note however, that although UCB selects a non-dominated location, there will generally be other non-dominated locations that trade-off exploration and exploitation differently.

2.2.2 Expected Improvement. The EI is perhaps the most popular infill criterion and is very widely used. It was first proposed by Moćkus et al. [1978] and further developed by Jones et al. [1998]. Bull has shown that, under certain conditions, BO using EI is guaranteed to converge to the global optimum [Bull 2011].

EI is based on the positive predicted improvement over the best solution $f^* = \max_m \{f_m\}$ observed so far. If $\hat{f} = f(\mathbf{x})$ is an evaluation of f at $\mathbf{x}$, then the improvement is

$$I(\mathbf{x}, \hat{f}, f^*) = \max(\hat{f} - f^*, 0). \quad (8)$$

Then the expected improvement at $\mathbf{x}$ may be expressed as [Jones et al. 1998]

$$\begin{aligned} \alpha_{EI}(\mathbf{x}) &= \mathbb{E}[I(\mathbf{x}, f^*)] = \int_{-\infty}^{\infty} I(\mathbf{x}, \hat{f}, f^*) p(\hat{f} | \mathbf{x}, \mathcal{D}) d\hat{f} \\ &= \sigma(\mathbf{x}) (s\Phi(s) + \phi(s)), \end{aligned} \quad (9)$$

where $s = (\mu(\mathbf{x}) - f^*)/\sigma(\mathbf{x})$ is the predicted improvement at $\mathbf{x}$ normalised by the uncertainty, and $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard Gaussian probability density and cumulative density functions. The infill criterion is therefore the improvement averaged with respect to the posterior predictive probability of obtaining it. Thus, EI balances the exploitation of solutions that are very likely to be a little better than f^* with the exploration of others that may, with lower probability, turn out to be much better.

As illustrated in Figure 2, $\alpha_{EI}(\mathbf{x})$ is monotonic with respect to increase in both exploration, σ , and exploitation, μ . This can be seen by noting that

$$\frac{\partial \alpha_{EI}}{\partial \mu} = \Phi(s) \quad \text{and} \quad \frac{\partial \alpha_{EI}}{\partial \sigma} = \phi(s) \quad (10)$$

are both positive everywhere [Jones et al. 1998]. Consequently, like UCB, if the next location to be expensively evaluated is selected by maximising EI, the location will belong to the Pareto set maximally trading-off exploration and exploitation.

2.2.3 Weighted Expected Improvement. Some authors [Feng et al. 2015; Sóbester et al. 2005] have associated the term, $\sigma(\mathbf{x})s\Phi(s) = (\mu(\mathbf{x}) - f^*)\Phi(s)$, in Equation (9) with the exploitation inherent in adopting $\mathbf{x}$ as the next place to evaluate. Similarly, the term $\sigma(\mathbf{x})\phi(s)$ has been associated with the exploratory component. To control the balance between exploration and exploitation Sóbester et al. [2005] define an acquisition function that weights these two terms differently:

$$\alpha_{WEI}(\mathbf{x}, \omega) = \sigma(\mathbf{x}) [\omega s\Phi(s) + (1 - \omega)\phi(s)], \quad (11)$$

where $0 \leq \omega \leq 1$.

However, it turns out that if the next point for expensive evaluation is selected by maximising $\alpha_{WEI}(\mathbf{x})$ in some set $\mathcal{S}$ of candidate solutions, $\mathbf{x}' = \operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} \alpha_{WEI}(\mathbf{x}, \omega)$, then this only results in choosing $\mathbf{x}'$ in the maximal non-dominated set of $\mathcal{S}$ for a relatively small range of ω . This may be seen by considering the partial derivatives of $\alpha_{WEI}(\mathbf{x}, \omega)$. Without loss of generality, we take $f^* = 0$, so that $s = \mu/\sigma$. Then

$$\frac{\partial \alpha_{WEI}}{\partial \sigma} = -\omega s^2 \phi(s) + (1 - \omega)(\phi(s) - s\phi'(s)) \quad (12)$$

$$= [1 - \omega + (1 - 2\omega)s^2] \phi(s), \quad (13)$$

where we have used the fact that $\phi'(s) = -s\phi(s)$. Consequently, when $\omega \leq \frac{1}{2}$ the gradient $\frac{\partial \alpha_{WEI}}{\partial \sigma} > 0$ for all s . However, if $\omega > \frac{1}{2}$ so that $1 - 2\omega < 0$, then there are always regions where $s = \mu/\sigma$ is sufficiently large that $\frac{\partial \alpha_{WEI}}{\partial \sigma} < 0$. In this case, there are therefore regions of (μ, σ) space in which decreasing σ increases α_{WEI} , so $\operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} \alpha_{WEI}(\mathbf{x}, \omega)$ is not guaranteed to lie in the Pareto set.

The gradient in the μ direction is

$$\frac{\partial \alpha_{WEI}}{\partial \mu} = \omega \Phi(s) + (2\omega - 1)s\phi(s). \quad (14)$$

Requiring that the gradient is non-negative, so that α_{WEI} is non-decreasing with μ results in

$$\omega \geq (1 - 2\omega) \frac{s\phi(s)}{\Phi(s)}. \quad (15)$$

When $\omega > \frac{1}{2}$, it is straightforward to see that Equation (15) is always satisfied. The inequality is also always satisfied for all $s < 0$ when $\omega < \frac{1}{2}$. When $\omega < \frac{1}{2}$ and $s \geq 0$, the inequality may be

Fig. 3. Contours of weighted expected improvement as functions of the surrogate model's predicted mean μ and uncertainty σ for weights $\omega = 0, 0.1, 1$; Equation (11). In none of these cases is the $\mathbf{x}'$ maximising $\alpha_{WEI}(\mathbf{x}', \omega)$ guaranteed to lie in the Pareto set of maximally exploratory and exploitative solutions.

rewritten as

$$\frac{\omega}{1 - 2\omega} \geq \frac{s\phi(s)}{\Phi(s)}. \quad (16)$$

Defining

$$\gamma = \sup \frac{s\phi(s)}{\Phi(s)} \approx 0.295, \quad (17)$$

it can be seen that $\frac{\partial \alpha_{WEI}}{\partial \mu}$ is only non-negative everywhere if $\omega \geq \gamma / (2\gamma + 1) \approx 0.185$. It may therefore be concluded that when $\omega \in \left[\frac{\gamma}{(2\gamma + 1)}, \frac{1}{2} \right]$ maximising $\alpha_{WEI}(\mathbf{x}, \omega)$ results in the next location for expensive evaluation lying in the Pareto set of available solutions. However, this is not guaranteed for other values of ω . These results are illustrated in Figure 3, which shows α_{WEI} as a function of $\mu - f^*$ and σ for $\omega = 0, 0.1$, and 1 ; cf. Figure 2 for $\omega = 0.5$. The complicated nature of α_{WEI} is apparent when $\omega = 0.1$. When $\omega = 0$ the acquisition function might be expected to yield purely exploratory behaviour. However, in this case although locations with high variance are preferred over those with low variance *with the same* μ , the acquisition function guides the search toward locations with high variance but a mean prediction close to f^* . Purely exploitative behaviour might be expected when $\omega = 1$. In this case, the acquisition function is maximised for large μ and small σ , which implies that the location with the smaller σ will be preferred from two locations with the same μ . Consequently, although the acquisition function in this case encourages exploitation (preferring large μ) it discourages exploration (preferring small σ). This is in contrast to standard EI ($\omega = 0.5$, Figure 2), which prefers the high variance, more exploratory, location from two locations with the same μ .

2.2.4 Probability of Improvement. The Probability of Improvement (PI) is one of the earliest proposed infill criteria [Kushner 1964]. It is the probability that the prediction at a location $\mathbf{x}$ is greater than the best observed (expensively evaluated) function value f^* . As the predictive distribution is Gaussian, PI may be calculated in closed form:

$$\alpha_{PI}(\mathbf{x}) = p(f > f^* | \mathbf{x}, \mathcal{D}) = \Phi(s(\mathbf{x})). \quad (18)$$

Thus, $\alpha_{PI}(\mathbf{x})$ is the volume of the predictive distribution lying above f^* .

Since

$$\frac{\partial \alpha_{PI}}{\partial \mu} = \frac{1}{\sigma(\mathbf{x})} \phi(s(\mathbf{x})) \quad (19)$$

is positive for all $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$, PI is monotonically increasing with increasing mean prediction for fixed uncertainty. Thus, as might be expected, at fixed uncertainty, locations where the mean is predicted to be large are preferred. Interestingly, as Figure 2 illustrates, such a straightforward monotonic relationship does not exist with respect to uncertainty as shown by

$$\frac{\partial \alpha_{PI}}{\partial \sigma} = -s(\mathbf{x})\phi(s(\mathbf{x})). \quad (20)$$

When the improvement in the mean is negative $s(\mathbf{x}) < 0$ then Equation (20) shows that PI increases with uncertainty σ . However, in contrast to EI and UCB, when $\mu(\mathbf{x}) > f^*$ then Equation (20) shows that PI decreases with σ ; the locations with small uncertainty are preferred to those with high uncertainty. Therefore, the location $\mathbf{x}'$ selected by PI is not guaranteed to be a member of the maximal non-dominated set of candidates. In other words, there may be candidate locations $\mathbf{x}''$, which are more exploratory ($\sigma(\mathbf{x}'') > \sigma(\mathbf{x}')$) while having the same mean prediction ($\mu(\mathbf{x}'') = \mu(\mathbf{x}')$) as the $\mathbf{x}'$ selected by PI.

In practice, such behaviour leads to an overly exploitative scheme, see for example [Jones 2001]. To combat this exploitative nature, usually a higher target than the best observed value, f^* , is set for computing the probability of improvement. This often improves the performance of PI-based BO [Jones 2001; Kushner 1964; Lizotte 2008]. As Figure 2 shows, this can be attributed to the fact that solutions are evaluated as if their improvement were negative where the PI criterion encourages exploration as well as exploitation. Although this modification tends to improve performance, there is, however, no natural choice for a suitable high target.

2.3 Exploration and Exploitation Trade-off

As discussed above, the EI and UCB infill criteria select the next location to be expensively evaluated as one of the locations that are members of the maximal non-dominated set of available locations, namely, $\mathcal{P}$ [Equation (6)], the Pareto set resulting from simultaneous maximisation of $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$. PI only selects from $\mathcal{P}$ when $\mu(\mathbf{x}) < f^*$ and in practice an artificially high f^* is used to promote exploration. Note, however, that EI and UCB select from different regions of the Pareto set, balancing exploitation and exploration differently. Indeed, the proof of convergence for BO with UCB relies on varying the selection position along the Pareto front as the optimisation proceeds, becoming more exploratory in later stages [Srinivas et al. 2010].

Inspection of Figure 2 shows that EI is more exploitative than UCB in the sense that if the solutions available for selection all have the same upper confidence bound, that is they lie on a contour of α_{UCB} , then maximising α_{EI} will choose the most exploitative of them. Conversely, if the available solutions all have the same EI, then maximising α_{UCB} will choose the most exploratory.

3 UTILISING THE EXPLORATION VS. EXPLOITATION TRADE-OFF FRONT

Previous works [Bischl et al. 2014; Feng et al. 2015; Grobler et al. 2017] have used the exploration vs. exploitation (EE) front in a batch setting, in which multiple locations in the EE Pareto set are selected to be evaluated in parallel. Feng et al. [2015] use the two weighted components of α_{WEI} [Equation (11)] as the two objectives defining a trade-off front that is approximated via the use of a multi-objective evolutionary algorithm (MOEA). They select batches of q solutions to be expensively evaluated in parallel by choosing the two extremal solutions of the approximated Pareto set and the remaining $q - 2$ locations equally spread (in objective space) across the set. Grobler et al. [2017] replace the α_{WEI} formulation with a trade-off front consisting of the surrogate model's mean and variance, again using a MOEA to approximate the Pareto set. They select a batch of locations consisting of the two extremal solutions of the set, together with the location that maximises EI, and equally spaced solutions across the set. Bischl et al. [2014] consider the

maximisation of an additional objective, namely, the decision space distance to each solution's nearest neighbour, thus promoting exploration. They also limit the size of the MOEA population to be the batch size to avoid the problem of explicitly selecting a batch of locations from a large Pareto set.

The use of the EE front in the sequential setting is much less explored. However, Žilinskas and Calvin [2019] have recently highlighted the importance of visualising the EE front to better inform model selection and they recommend that future researchers should aim to exploit the EE front further.

Here, we focus on the sequential BO framework (recall Algorithm 1) and consider algorithms that select the next location for expensive evaluation from the entire Pareto set of feasible locations. Use of an efficient evolutionary multi-objective search algorithm means that finding an approximation $\tilde{\mathcal{P}}$ to $\mathcal{P}$ has about the same computational expense as maximising a scalar acquisition function such as EI or UCB directly. In this work the approximate Pareto set of model predictions is found using a standard evolutionary optimiser, NSGA-II [Deb et al. 2001].

We note that while proofs of convergence for particular trade-offs between exploration and exploitation exist [Bull 2011; Srinivas et al. 2010], it is clear that merely selecting locations for any fixed exploration-exploitation weighting are not guaranteed to converge. At the two extremes, purely exploitative schemes select $\mathbf{x}' = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x})$ and purely exploratory schemes select $\mathbf{x}' = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \sigma(\mathbf{x})$. The former are liable to become stuck at local optima, while the latter visits each location with the maximum posterior variance $\sigma^2(\mathbf{x})$, thus reducing the uncertainty of the model, as quantified by the entropy of the predictive posterior. This will lead to the eventual location of the optimum, but only very slowly as even very unpromising locations where $\mu(\mathbf{x}) \ll f^*$ are visited.

In Section 4, we evaluate the performance of the purely exploitative and exploratory strategies, denoted *Exploit* and *Explore*, respectively. Since all solutions in the Pareto set may be considered equally good and dominate all other feasible locations, we also consider the *PFRandom* algorithm, which selects a solution at random from $\tilde{\mathcal{P}}$ for the next expensive evaluation.

As discussed above, the maximally exploratory strategy will converge to the global optimum, but very slowly. At the other extreme of the Pareto front, a greedy, exploitative, strategy, while converging quickly, risks becoming stuck at a local optimum. In the next section, therefore, we seek to capitalise on the rapid convergence of the exploitative strategy while avoiding local minima by making occasional exploratory moves.

3.1 ϵ -Greedy Bayesian Optimisation

Motivated by the success of ϵ -greedy schemes in reinforcement learning [Mnih et al. 2015; Sutton and Barto 1998; Tokic 2010; Van Hasselt et al. 2016], we propose two novel BO acquisition functions that use the ϵ -greedy methodology to select the next point for expensive evaluation. Both methods mostly select the most exploitative solution, but differ in which exploratory solution is selected in a small proportion of steps.

The first method, which we denote ϵ -PF and is summarised in Algorithm 2a, usually selects the location $\mathbf{x}'$ with the most promising mean prediction from the surrogate model. In the remaining cases, with probability ϵ , it selects a random location from the approximate Pareto set $\tilde{\mathcal{P}}$, thus usually selecting a more exploratory $\mathbf{x}'$ instead of the most exploitative location available. The function `MOOptimise` denotes the use of a multi-objective optimiser to generate $\tilde{\mathcal{P}}$. This acquisition function replaces line 7 in standard BO, Algorithm 1.

The ϵ -RS scheme, summarised in Algorithm 2b, also usually selects $\mathbf{x}'$ with the most promising mean prediction from the surrogate. However, with probability ϵ a location is randomly selected

ALGORITHM 2: ϵ -greedy acquisition functions.**(2a)** ϵ -PF: Pareto front selection.

```

1: if rand() <  $\epsilon$  then
2:    $\tilde{\mathcal{P}} \leftarrow \text{M0Optimise}_{\mathbf{x} \in \mathcal{X}}(\mu(\mathbf{x}), \sigma(\mathbf{x}))$ 
3:    $\mathbf{x}' \leftarrow \text{randomChoice}(\tilde{\mathcal{P}})$ 
4: else
5:    $\mathbf{x}' \leftarrow \text{argmax}_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x})$ 

```

(2b) ϵ -RS: Random selection from feasible space.

```

1: if rand() <  $\epsilon$  then
2:    $\mathbf{x}' \leftarrow \text{randomChoice}(\mathcal{X})$ 
3: else
4:    $\mathbf{x}' \leftarrow \text{argmax}_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x})$ 

```

(hence the abbreviation ϵ -RS) from the entire feasible space $\mathcal{X}$. Selection of $\mathbf{x}'$ from $\tilde{\mathcal{P}}$ (ϵ -PF, Algorithm 2a) might be expected to be more effective than selecting $\mathbf{x}'$ from the entire feasible space (ϵ -RS, Algorithm 2b), because a selection from $\mathcal{X}$ is likely to be dominated by $\tilde{\mathcal{P}}$ and therefore is likely to be less exploratory and less exploitative.

We remark that these ϵ -greedy schemes are different to that proposed by Bull [2011], which greedily selects the location with maximum expected improvement with probability $1 - \epsilon$, and randomly chooses a location the remainder of the time. This is different from our proposals, because the Bull scheme greedily maximises EI rather than exploitation (μ).

4 EXPERIMENTAL EVALUATION

We investigate the performance of the two proposed ϵ -greedy methods, ϵ -PF and ϵ -RS, by evaluating them on ten benchmark functions with a range of domain sizes and dimensionality; see Table 1 for details. Note that the benchmarks are couched as minimisation problems. In common with other works [Jones et al. 1998; Schonlau 1997; Wagner and Wessing 2012; Wang et al. 2015], the functions prefixed with *log* are log-transformed, i.e., the logarithm of each observed values $\log(f(\mathbf{x}))$ is modelled rather than observed value $f(\mathbf{x})$ itself. Where the observations can be negative, a constant larger than minimum value of the function is added.² These functions are transformed in this *grey-box* fashion, using a small amount of prior information about the scales of the function, because we want the surrogate model to be as accurate as possible. As discussed in the seminal work of Jones et al. [1998], it is often possible to improve poorer surrogate model fits, as one typically observes with the untransformed functions, by using the log transformation. The equations defining each transformed function and optimisation results of all methods on the *untransformed* functions are available in the supplementary material. We discuss the differences in optimisation performance between the standard and log-transformed functions below.

We compare the two proposed methods to the purely exploitative and exploratory strategies, denoted *Exploit* and *Explore*, respectively, as well as random selection from the approximated Pareto front, *PFRandom*. Their performance is also compared to the infill criteria discussed in Section 2.2, namely, Expected Improvement (EI), Upper Confidence Bound (UCB), and Probability of Improvement (PI). In addition, we compare the performance of all the infill criteria with the quasi-random search produced by max-min Latin Hypercube Sampling (LHS) [McKay et al. 2000]. LHS is the generalisation of a Latin square, in which samples are placed in rows and columns of a square such that each sample resides in its own row and column. The max-min variant of LHS tries to maximise the minimum distance between each sample.

The methods were evaluated on the synthetic benchmark functions in Table 1, with a budget of 250 function evaluations that included $M = 2d$ initial LHS samples (Algorithm 1, Line 1). To allow statistical performance measures to be used, each optimisation was repeated 51 times. The same

²We have used prior information on the function's minimum value to choose the constant, but the actual value is immaterial, because the function observations are in any case standardised as part of the GP modelling.

Table 1. Functions Used in These Experiments, Along with Their Domain and Dimensionality, d

Name	Domain	d
WangFreitas [Wang and de Freitas 2014]	$[0, 1]$	1
Branin [†]	$[-5, 0] \times [10, 15]$	2
BraninForrester [Forrester et al. 2008]	$[-5, 0] \times [10, 15]$	2
Cosines [González et al. 2016b]	$[0, 0] \times [5, 5]$	2
logGoldsteinPrice [†]	$[-2, -2] \times [2, 2]$	2
logSixHumpCamel [†]	$[-3, 2] \times [3, 2]$	2
logHartmann6 [†]	$[0, 1]^d$	6
logGSobol [González et al. 2016a]	$[-5, 5]^d$	10
logRosenbrock [†]	$[-5, 10]^d$	10
logStyblinskiTang [†]	$[-5, 5]^d$	10

Formulae can be found as cited or at <http://www.sfu.ca/~ssurjano/optimization.html> for those labelled with [†]. Full details of all evaluated functions can also be found in the supplementary material.

sets of initial samples were used for each method's runs to allow for paired statistical comparisons between the methods to be carried out. In all experiments a value of $\epsilon = 0.1$ was used for both ϵ -PF and ϵ -RS. The UCB algorithm was run with β_t adjusted according to the schedule defined for continuous functions in Theorem 2 of Srinivas et al. [2010] with $a = b = 1$ and $\delta = 0.01$. All acquisition functions were optimised with NSGA-II [Deb et al. 2001], apart from PI, which was optimised following the common practise [Balandat et al. 2020; GPyOpt 2016] of uniformly sampling $\mathcal{X}$ and optimising the 10 most promising locations with L-BFGS-B [Byrd et al. 1995]. In both cases, the optimisation budget was $5,000d$ evaluations. The multi-start strategy was used to optimise PI, because, as shown in Section 2.2.4, the maximiser of PI may not lie in the Pareto set of $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$. For NSGA-II, we set the parameters to commonly used values: the population size was $100d$, the number of generations was 50 (100 generations lead to no significant improvement in performance), the crossover and mutation probabilities were 0.8 and $\frac{1}{d}$, respectively, and both the distribution indices for crossover and mutation were 20.

Table 2 shows the median regret, i.e., the median difference between the estimated optimum f^* and the true optimum over the 51 repeated experiments, together with the median absolute deviation from the median (MAD). The method with the lowest (best) median regret on 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 [Knowles et al. 2006] with Holm-Bonferroni correction [Holm 1979] ($p \geq 0.05$), are shown in light grey.

Figure 4 shows the convergence of the various algorithms on four illustrative test problems in $d = 1, 2$ and 10 dimensions. Convergence plots for all the benchmark problems are available in the supplementary material, and Python code to generate figures and reproduce all experiments is available online.³

As might be expected, Latin Hypercube Sampling (LHS) and purely exploratory search (Explore), which have roughly equivalent performance, are not the best methods on any of the test problems.

Perhaps surprisingly, none of the three well-known acquisition functions, EI, UCB, and PI, has the best median performance after 250 evaluations, although all three are statistically equivalent to the best method on $d = 2$ Cosines, and EI and UCB have good performance on the $d = 2$ Branin

³<https://github.com/georgedeath/egreedy/>.

Table 2. Median Absolute Distance (*Left*) and Median Absolute Deviation from the Median (MAD, *Right*) from the Optimum After 250 Function Evaluations Across the 51 Runs

Method	WangFreitas (1)		BraninForrester (2)		Branin (2)		Cosines (2)		logGoldsteinPrice (2)	
	Median	MAD	Median	MAD	Median	MAD	Median	MAD	Median	MAD
LHS	1.27×10^{-2}	1.80×10^{-2}	4.59×10^{-1}	4.73×10^{-1}	1.31×10^{-1}	1.33×10^{-1}	4.79×10^{-1}	2.71×10^{-1}	1.08	7.69×10^{-1}
Explore	1.04×10^{-2}	1.42×10^{-2}	4.58×10^{-1}	3.52×10^{-1}	1.66×10^{-1}	1.56×10^{-1}	4.56×10^{-1}	2.20×10^{-1}	1.01	5.50×10^{-1}
EI	2.00	6.91×10^{-11}	2.47×10^{-6}	3.23×10^{-6}	4.15×10^{-6}	3.76×10^{-6}	6.31×10^{-6}	7.68×10^{-6}	2.73×10^{-6}	3.34×10^{-6}
PI	2.06	8.24×10^{-2}	3.73×10^{-4}	3.70×10^{-4}	2.26×10^{-5}	3.22×10^{-5}	2.50×10^{-3}	3.18×10^{-3}	2.92×10^{-3}	4.32×10^{-3}
UCB	2.00	1.26×10^{-11}	4.96×10^{-6}	6.22×10^{-6}	4.42×10^{-6}	4.06×10^{-6}	7.12×10^{-6}	8.86×10^{-6}	6.15×10^{-6}	6.17×10^{-6}
PFRandom	2.00×10^{-4}	2.96×10^{-4}	2.70×10^{-3}	3.65×10^{-3}	1.67×10^{-3}	2.17×10^{-3}	8.82×10^{-3}	1.14×10^{-2}	2.54×10^{-3}	3.31×10^{-3}
ϵ -RS	1.04×10^{-6}	1.54×10^{-6}	2.00×10^{-6}	2.49×10^{-6}	3.17×10^{-6}	2.46×10^{-6}	8.66×10^{-6}	1.21×10^{-5}	2.33×10^{-6}	2.36×10^{-6}
ϵ -PF	2.00	3.72×10^{-11}	2.31×10^{-6}	3.01×10^{-6}	3.57×10^{-6}	3.13×10^{-6}	2.02×10^{-6}	2.52×10^{-6}	8.76×10^{-7}	1.08×10^{-6}
Exploit	2.00	6.00×10^{-9}	4.61×10^{-6}	6.04×10^{-6}	3.08×10^{-6}	3.29×10^{-6}	4.13×10^{-1}	6.12×10^{-1}	2.26×10^{-6}	2.90×10^{-6}
Method	logSixHumpCamel (2)		logHartmann6 (6)		logGSobol (10)		logRosenbrock (10)		logStyblinskiTang (10)	
	Median	MAD	Median	MAD	Median	MAD	Median	MAD	Median	MAD
LHS	6.52	1.10	3.37×10^{-1}	1.10×10^{-1}	1.51×10^1	9.03×10^{-1}	1.16×10^1	5.39×10^{-1}	2.85	1.77×10^{-1}
Explore	6.53	1.24	3.07×10^{-1}	6.85×10^{-2}	1.75×10^1	1.42	1.28×10^1	4.82×10^{-1}	3.19	1.13×10^{-1}
EI	7.42×10^{-5}	9.19×10^{-5}	1.06×10^{-3}	6.73×10^{-4}	7.15	1.58	6.62	6.58×10^{-1}	2.34	2.79×10^{-1}
PI	1.46×10^{-1}	1.58×10^{-1}	6.15×10^{-4}	7.69×10^{-4}	6.29	1.61	6.89	9.49×10^{-1}	2.29	2.37×10^{-1}
UCB	3.84	1.36	2.04×10^{-1}	3.21×10^{-2}	1.45×10^1	6.16×10^{-1}	8.31	5.90×10^{-1}	3.19	1.13×10^{-1}
PFRandom	1.52×10^{-1}	1.52×10^{-1}	6.57×10^{-2}	3.27×10^{-2}	5.60	1.73	5.23	4.98×10^{-1}	2.70	3.15×10^{-1}
ϵ -RS	3.81×10^{-5}	2.96×10^{-5}	5.09×10^{-4}	3.59×10^{-4}	5.13	1.86	4.75	7.85×10^{-1}	1.61	3.12×10^{-1}
ϵ -PF	4.06×10^{-5}	4.66×10^{-5}	7.71×10^{-4}	4.82×10^{-4}	5.06	1.37	4.64	6.25×10^{-1}	1.53	4.49×10^{-1}
Exploit	4.21×10^{-5}	4.95×10^{-5}	6.37×10^{-4}	5.82×10^{-4}	5.27	1.60	4.54	6.19×10^{-1}	1.82	3.71×10^{-1}

The method with the lowest median performance is shown in dark grey, with those with statistically equivalent performance are shown in light grey.

Fig. 4. Illustrative convergence plots for four benchmark problems. Each plot shows the median difference between the best function value seen and the true optimum (regret), with shading representing the interquartile range across the 51 runs. The dashed vertical line indicates the end of the initial LHS phase.

Fig. 5. Distribution of the best-seen function values after 50 (*left*), 150 (*centre*), and 250 (*right*) function evaluations on three benchmark problems.

and BraninForrester problems. In contrast, the ϵ -greedy algorithms ϵ -PF and ϵ -RS perform well across the range of problems, particularly on the higher-dimensional problems. Interestingly, Exploit, which always samples from the best mean surrogate prediction is competitive for most of the high-dimensional problems. This indicates one of the main conclusions of this work, namely, that as the dimension of decision space increases the approximate modelling of $f(\mathbf{x})$ is so poor that even adopting the modelled most-exploitative solution inherently leads to some (unintended) exploration.

While pure exploitation combined with fortuitous exploration appears to be a good strategy for many problems, introducing some deliberate exploration can be important. This is particularly apparent on the WangFreitas problem [Wang and de Freitas 2014], which contains a large local optimum and a narrow global optimum that is surrounded by plateaux; see supplementary material for a plot. Convergence on this problem is shown in Figure 4, which demonstrates how LHS sampling and a purely exploratory strategy (Explore) converge slowly toward the optimum, while Exploit fails to find the vicinity of the optimum in any case. However, the deliberate exploratory moves incorporated in both ϵ -greedy methods and PFRandom (random selection from the Pareto set) enable some of the runs to converge to the optimum. The ϵ -RS method, which makes exploratory moves from the entire feasible space, is most effective, although as discussed below, generally, we find ϵ -PF to be more effective.

Figure 5 shows the distribution of the best-seen function evaluations for each of the evaluated algorithms on three benchmark problems for budgets of $T = 50, 150$, and 250 function evaluations. Again, we see in the two-dimensional Cosines and logSixHumpCamel plots that driving

Fig. 6. Comparison of ϵ -PF (green) and ϵ -RS (red, hatched) for different values of ϵ (horizontal axis) after 50 (left), 150 (centre), and 250 (right) function evaluations.

the optimisation process solely by exploiting the surrogate's mean prediction can fail to correctly identify the optimum, because the model is inaccurate and may miss, for example, a small-scale optimum. When f is modelled poorly, then the mean function will not accurately represent the true function. However, as is the case with the logGSobol plot and indeed the other ten-dimensional functions, pure exploitation can provide a sufficient driver for optimisation, because the inaccurate and changing surrogate (as new evaluations become available) induces sufficient exploration. We note however, that the ϵ -greedy algorithms, incorporating deliberate exploration, offer more consistent performance.

A common trend apparent across the both Figures 4 and 5 is that EI tends to initially improve at a slower rate than the two ϵ -greedy methods, but then catches up to a greater or lesser extent after more function evaluations. This is well illustrated in the logSixHumpCamel plot in Figure 5 and also in the Branin and logRosenbrock plots in Figure 4. UCB performs poorly on the higher-dimensional functions. This may be due to the value of β_t used, as the convergence proofs in [Srinivas et al. 2010] rely on β_t increasing with the dimensionality of the problem, leading to over-exploration. One may argue that this can be overcome by simply using a smaller β_t value, set in some *ad hoc* manner. However, with no *a priori* knowledge as to how to select the parameter on a per-problem basis, we suggest that this is not a feasible strategy in practice.

How greedy? Choosing ϵ . Although the ϵ -greedy algorithms perform well in comparison with conventional acquisition functions, it is unclear what value of ϵ to choose, and indeed whether the exploratory moves should choose from the approximate Pareto front (ϵ -PF) or from the entire feasible space (ϵ -RS), which is marginally cheaper. Figure 6 illustrates the effect of ϵ on the performance of ϵ -PF (green) and ϵ -RS (red, hatched). As is clear from the Cosines problem, a larger

Fig. 7. A comparison of optimising the GSobol function with ϵ -PF (green) and ϵ -RS (red, hatched) for different values of ϵ (horizontal axis) after 50 (left), 150 (centre), and 250 (right) function evaluations.

value of ϵ may be required to avoid getting stuck, because the surrogate is not modelling the function well enough and needs a larger number of exploratory samples. However, there is very little change in performance with ϵ for the higher-dimensional decision spaces (e.g., logGSobol and logRosenbrock). As suggested above, we attribute this to the inaccurate surrogate modelling in higher dimensions, which leads to a large degree of random search irrespective of ϵ .

Interestingly, this is not the case for these functions without the log transformation (Rosenbrock and GSobol). Figure 7 shows the performance of ϵ -RS and ϵ -PF for different values of ϵ on the GSobol problem. As can be seen in the figure, increasing ϵ decreases the performance of ϵ -PF and increases the performance of ϵ -RS, in stark contrast to logGSobol in Figure 6. This indicates that the surrogate model is misleading the optimisation, because increasing the frequency of expensively evaluating random locations and decreasing the frequency of sampling from the Pareto front both improve optimisation. In this case, the log transformation enables more accurate modelling of the objective and thus more rapid optimisation.

Overall, setting $\epsilon = 0.1$ appears to be large enough to give good performance across all problems (see supplementary material for results on other problems), particularly for the ϵ -PF algorithm. Larger values give no real improvement in performance. Empirically it appears that ϵ -PF gives marginally better performance than ϵ -RS, as might be expected if the surrogate describes f well, as is the case in the later stages of optimisation. In this case, selection from the approximate Pareto front yields solutions that lie on the maximal trade-off between exploration and exploitation and may therefore be expected to yield the most information. However, in cases where the surrogate modelling is particularly poor throughout the entire optimisation run, as is the case in several of the test problems without log transformation, the increased stochasticity provided by ϵ -RS with larger values of ϵ appears useful in overcoming the misleading surrogate model.

Results on the black-box test problems. Here we briefly describe the optimisation results of the evaluated methods on the six test problems without log transformation—full results are available in the supplementary material. The ϵ -RS method is the best performing or statistically equivalent to the best performing method on all six of the benchmark problems, with ϵ -PF best or equivalent on five of the six. EI, PI, and Exploit were all the best or equivalent to the best performing on three of the six test problem. As noted above, ϵ -RS performs better than ϵ -PF on the higher-dimensional problems, with the two methods giving equivalent performance on the lower-dimensional problems. The main difference of the standard acquisition functions is that performance is closer to that of the ϵ -greedy methods than on the log-transformed functions. We attribute this to be a result of poorer surrogate modelling in the presence of a wide range of objective values so that the ϵ -greedy schemes are less able to exploit the model's mean predicted value $\mu(\mathbf{x})$. We reiterate here, however, that the performance of both ϵ -RS and ϵ -PF across the untransformed benchmark functions is still superior to the standard acquisition functions.

Fig. 8. PitzDaily test problem. Fluid enters on the left (Inflow), flows through the expanded pipe and leaves on the right (Outflow). The shape of the lower boundary is defined by a Catmull-Clark subdivision curve (green) controlled by the locations of control points ($\blacktriangle$). The curve is constrained to lie within the blue polygon by penalising the acquisition function for solutions that violate it.

Fig. 9. Distribution of the best-seen function values after 50 (left), 150 (centre), and 250 (right) function evaluations on the real-world PitzDaily test problem.

4.1 Real-world Application: Pipe Shape Optimisation

We also evaluate the range of acquisition functions on a real-world computational fluid dynamics optimisation problem. As illustrated in Figure 8, the PitzDaily test problem [Daniels et al. 2018] involves optimising the shape of a pipe to reduce the pressure loss between the inflow and outflow. Pressure loss is caused by a rapid expansion in the pipe (a backward-facing step), which forces the flow to separate at the edge of the step, creating a recirculation zone, before the flow re-attaches at some distance beyond the step. The goal of the optimisation is to discover the shape of the lower wall of the pipe that minimises the pressure loss, which is evaluated by running a computational fluid dynamics (CFD) simulation of the two-dimensional flow. Solution of the partial differential equations describing the flow means that each function evaluation takes about 60 seconds—which is sufficient for us to conduct multiple runs to enable statistical comparisons for this problem.

As shown in Figure 8 and as described in detail by Daniels et al. [2018], we represent the wall geometry in terms of a Catmull-Clark sub-division curve, whose control points comprise the decision variables. Here there are five control points, resulting in a ten-dimensional decision vector. The control points are constrained to reside within a polygon and, therefore, the initial locations used in each optimisation run are sampled from a uniform distribution, and those that reside outside the constrained region are discarded and new samples generated to replace them. Similarly, the optimisation runs are compared to uniformly sampling 250 locations rather than Latin hypercube sampling, and are denoted as Uniform in the following results.

Figure 9 shows random selection from the Pareto front (PFRandom) had the best median fitness after 250 function evaluations, but EI, ϵ -PF, ϵ -RS, and Exploit were all statistically equivalent. We remark that the optimum discovered outperforms that discovered by Nilsson et al. [2014].

Fig. 10. The streamlines for two solutions: the local optimum identified by Nilsson et al. [2014] (*upper*) and the best estimation of the global optimum from one of the runs using the Bayesian optimiser (*lower*). Colour indicates fluid speed (normalised units). Good solutions typically replace the backward step with a slope.

We observe that good solutions typically replace the step shown in Figure 8 with a slope, as illustrated by the two solutions shown in Figure 10. This improves the performance, because it reduces the size of the recirculation zone immediately following the increase in the tube's width. Generally, the size of the recirculation zone is reduced for shallower slopes, resulting in a reduced flow velocity (as the streamlines suggest) and increased frictional pressure recovery. However, such a shallow slope that the recirculation zone is completely removed (as found by an adjoint optimisation method) does not perform best [Nilsson et al. 2014]. The Bayesian optimiser consistently discovers a wall shape that results in a small recirculation zone that more effectively dampens the flow, resulting in a smaller pressure loss [Daniels et al. 2019].

4.2 Real-world Application: Active Learning for Robot Pushing

Following Wang and Jegelka [2017] and Jiang et al. [2019], we optimise the control parameters for two active learning robot pushing problems [Wang et al. 2018]. In the first problem, illustrated in Figure 11, a robot hand (rectangle) is given the task of pushing an object (circle) toward an unknown target location (cross). Once the robot has pushed the object it receives feedback in the form the distance of the object to the target. The robot's movement is constrained such that it can only travel in the direction of the object's initial location. Adjustable parameters are the robot's starting position, the orientation of its hand and the length of time it travels. This can therefore be viewed as minimisation problem in which these four parameters are optimised to minimise the distance of the object's final location to the target. We denote the resulting four-dimensional problem `PUSH4`.

In the second problem, `PUSH8`, shown in Figure 11, two robots (blue and green rectangles) in the same arena have to push their respective objects (circles) toward unknown targets (crosses). Their movements are constrained similarly to `PUSH4`, meaning that if they are initialised facing one another they will block each other's path. The final distances of each of the pushed objects to the corresponding target are summed and the total is used as the feedback for both robots, resulting in a joint learning task. We treat this as a minimisation problem: the eight parameters determining the robots' paths are to be optimised to minimise the combined distance of the objects to their targets.

Fig. 11. Two robot pushing tasks. PUSH4 (*left*): a robot hand (rectangle) pushes the object (circle) toward a target (cross) in an unknown location. As indicated by the arrows, the robot always travels in the direction of the object's initial location and only receives feedback in the form of the distance of the object, after pushing, to the target. PUSH8 (*right*): Similarly, two robots push their objects toward unknown target locations. Note that in PUSH4 the robot is likely to push the ball close to the target, because it is initially positioned well and has its hand orientated toward the object. In contrast, neither robot in PUSH8 is likely to push its object close to the target, because each begins in a worse location and is not orientated in a manner conducive to pushing.

Like Wang and Jegelka [2017], the object's initial location in PUSH4 is always the centre of the domain and the target location is changed on each optimisation run. Corresponding runs for each optimisation method used the same target location so that the runs were directly comparable. The targets' positions were selected by Latin hypercube sampling of 51 positions across the domain. We thus average over instances of the problem class, rather than repeatedly optimising the same function from different initialisations—this supports the assessment of results generalised to starting positions (see Bartz-Beielstein [2015] for a broader discussion on problem generators and generalisable results). Likewise, in PUSH8 the object's initial locations were fixed as shown in Figure 11 and each target's positions were generated in the same way as the PUSH4 targets. Target positions were paired such that the minimum distance between the targets for each problem instance was sufficient for the objects to be placed on the targets without overlapping. However, this does not mean that in each instance it is possible for the robots to actually push the objects to their targets, because the targets may be positioned so that the robots would block each other *en route* to their targets. Since this means that the optimum distance for some of these problem instances is not zero, to report the difference between the optimised function value and the optimum, we sought the global optimum of each problem instance by randomly sampling the feasible space with 10^5 sets of robot parameters and locally optimised the 100 best of these with the L-BFGS-B algorithm [Byrd et al. 1995]. In fact, several of the optimisation runs discovered better solutions than this procedure, and in these cases we used the resulting value as the estimate of the global optimum.

Figure 12 shows convergence histories and box plots summarising the performance of each of the tested methods after 50, 150, and 250 function evaluations. As these results show, in the four-dimensional PUSH4 problem, the exploitative methods outperform the EI, PI and UCB acquisition functions. The ϵ -PF method has the median approach to the optimum, but ϵ -RS and pure exploitation are statistically indistinguishable. In the harder PUSH8 problem all of the optimisers are still far from the optimum, even after 250 function evaluations. Only random selection from the Pareto front (PFRandom) is significantly better than any other method, and we note that PFRandom also performed well in the ten-dimensional PitzDaily optimisation. We speculate that PFRandom, which selects from the entire Pareto front at each iteration, owes its good performance to the

Fig. 12. Illustrative convergence plots for the two robot pushing problems (*upper*) and the distribution of the best-seen function values (*lower*) after 50 (*left*), 150 (*centre*), and 250 (*right*) evaluations for both problems.

additional exploration resulting from this strategy, allowing it to explore the complicated optimisation landscape. The **PUSH8** optimisation landscape is particularly rugged and difficult to approximate with Gaussian processes due to the abrupt changes in fitness occurring as the robots' paths intersect. However, we note that increasing exploration by increasing ϵ for the ϵ -PF and ϵ -RS methods does not significantly improve their performance. See the supplementary material for these results as well as for videos of the best solutions found to several of the problem instances evaluated.

5 CONCLUSION

How the balance between exploration and exploitation is chosen is clearer in Bayesian optimisation than in some stochastic optimisation algorithms. We have shown that the Expected Improvement and Upper Confidence Bound acquisition functions select solutions from the Pareto optimal trade-off between exploration and exploitation. However, the both the Weighted Expected Improvement (for ω not in the range (0.185, 0.5]) and Probability of Improvement function may choose dominated solutions. This may account for the poor empirical performance of the PI acquisition function.

Our analysis and experiments indicate that an effective strategy is to be mostly greedy, occasionally selecting a random exploratory solution. ϵ -greedy acquisition functions that select from either the Pareto front of maximally exploratory and exploitative solutions or the entire feasible space perform almost equivalently and the algorithms are not sensitive to the precise value of ϵ . The need for exploration via deliberate inclusion of exploratory moves turns out to be less important as the dimension of decision space increases and the purely exploitative method is fortuitously exploratory because of the low fidelity surrogate modelling; improving the quality of surrogate models in the face of the curse of dimensionality is an important topic of future research. While ϵ -greedy algorithms are trivially guaranteed to converge eventually, we look forward to theoretical results on the rate of convergence.

ACKNOWLEDGMENT

We thank Dr. Steven Daniels for helping us prepare Figure 10.

REFERENCES

- Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G. Wilson, and Eytan Bakshy. 2020. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. *Advances in Neural Information Processing Systems* 33 (2020), 21524–21538.
- Thomas Bartz-Beielstein. 2015. How to create generalizable results. In *Springer Handbook of Computational Intelligence*, Janusz Kacprzyk and Witold Pedrycz (Eds.). Springer, Berlin, 1127–1142.
- Bernd Bischl, Simon Wessing, Nadja Bauer, Klaus Friedrichs, and Claus Weihs. 2014. MOI-MBO: Multiobjective infill for parallel model-based optimization. In *Proceedings of the International Conference on Learning and Intelligent Optimization*. Springer, 173–186.
- Adam D. Bull. 2011. Convergence rates of efficient global optimization algorithms. *J. Mach. Learn. Res.* 12 (Oct. 2011), 2879–2904.
- Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyu Zhu. 1995. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* 16, 5 (1995), 1190–1208.
- 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 *Proceedings of the Conference on Parallel Problem Solving from Nature (PPSN'18)*. Springer, 296–307.
- Steven J. Daniels, Alma A. M. Rahat, Gavin R. Tabor, Jonathan E. Fieldsend, and Richard M. Everson. 2019. Automated shape optimisation of a plane asymmetric diffuser using combined Computational Fluid Dynamic simulations and multi-objective Bayesian methodology. *Int. J. Comput. Fluid Dynam.* 33, 6–7 (2019), 256–271.
- Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. 2001. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evolution. Comput.* 6, 2 (2001), 182–197.
- Zhiwei Feng, Qingbin Zhang, Qingfu Zhang, Qiangang Tang, Tao Yang, and Yang Ma. 2015. A multiobjective optimization-based framework to balance the global exploration and local exploitation in expensive optimization. *J. Global Optimiz.* 61, 4 (2015), 677–694.
- Alexander I. J. Forrester, Andras Sobester, and Andy J. Keane. 2008. *Engineering Design via Surrogate Modelling—A Practical Guide*. Wiley.
- Javier González, Zhenwen Dai, Philipp Hennig, and Neil Lawrence. 2016a. Batch Bayesian optimization via local penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, Vol. 51. PMLR, 648–657.
- Javier González, Michael Osborne, and Neil Lawrence. 2016b. GLASSES: Relieving the myopia of Bayesian optimisation. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, Vol. 51. PMLR, 790–799.
- GPy. 2012. GPy: A Gaussian process framework in Python. Retrieved from <http://github.com/SheffieldML/GPy>.
- GPyOpt. 2016. GPyOpt: A Bayesian Optimization framework in Python. Retrieved from <http://github.com/SheffieldML/GPyOpt>.
- Carla Grobler, Schalk Kok, and Daniel N. Wilke. 2017. Simple intuitive multi-objective Parallelization of efficient global optimization: SIMPLE-EGO. In *Proceedings of the World Congress of Structural and Multidisciplinary Optimisation*. Springer, 205–220.
- Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* 6, 2 (1979), 65–70.
- Shali Jiang, Henry Chai, Javier Gonzalez, and Roman Garnett. 2020. BINOCULARS for efficient, nonmyopic sequential experimental design. In *International Conference on Machine Learning*. PMLR, 4794–4803.

- Donald R. Jones. 2001. A taxonomy of global optimization methods based on response surfaces. *J. Global Optimiz.* 21, 4 (2001), 345–383.
- Donald R. Jones, Matthias Schonlau, and William J. Welch. 1998. Efficient global optimization of expensive black-box functions. *J. Global Optimiz.* 13, 4 (1998), 455–492.
- Joshua D. Knowles, Lothar Thiele, and Eckart Zitzler. 2006. *A Tutorial on the Performance Assessment of Stochastic Multi-objective Optimizers*. Technical Report TIK214, Computer Engineering and Networks Laboratory, ETH Zurich, Zurich, Switzerland.
- Harold J. Kushner. 1964. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *J. Basic Eng.* 86, 1 (1964), 97–106.
- Tze Leung Lai and Herbert Robbins. 1985. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* 6, 1 (1985), 4–22.
- Daniel J. Lizotte. 2008. *Practical Bayesian Optimization*. Ph.D. Dissertation, University of Alberta.
- Michael 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.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- Jonas Moćkus, Vytautas Tiešis, and Antanas Žilinskas. 1978. The application of Bayesian methods for seeking the extremum. *Towards Global Optimiz.* 2, 1 (1978), 117–129.
- 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.
- Carl Edward Rasmussen and Christopher K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. The MIT Press, Boston, MA.
- Matthias Schonlau. 1997. *Computer Experiments and Global Optimization*. Ph.D. Dissertation, University of Waterloo.
- 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.
- 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, 2951–2959.
- András Sóbester, Stephen J. Leary, and Andy J. Keane. 2005. On the design of optimization strategies based on global response surface approximation models. *J. Global Optimiz.* 33 (2005), 31–59.
- 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.
- Richard S. Sutton and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- Michel Tokic. 2010. Adaptive ϵ -greedy exploration in reinforcement learning based on value differences. In *Proceedings of the Annual Conference on Artificial Intelligence*. Springer, 203–210.
- Hado van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double Q-learning. In *Proceedings of the 13th AAAI Conference on Artificial Intelligence*. AAAI Press, 2094–2100.
- Tobias Wagner and Simon Wessing. 2012. On the effect of response transformations in sequential parameter optimization. *Evolution. Comput.* 20, 2 (2012), 229–248.
- Hao Wang, Thomas Bäck, and Michael T. M. Emmerich. 2015. Multi-point efficient global optimization using niching evolution strategy. In *EVOLVE—A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation VI*. Springer, 146–162.
- Ziyu Wang and Nando de Freitas. 2014. Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters. Retrieved from <https://arXiv:1406.7758>.
- 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.
- 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.
- Antanas Žilinskas and James Calvin. 2019. Bi-objective decision making in global optimization based on statistical models. *J. Global Optimiz.* 74, 4 (2019), 599–609.

Received November 2019; revised August 2020; accepted September 2020