

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

GEORGE DE ATH, University of Exeter, United Kingdom

RICHARD M. EVERSON, University of Exeter, United Kingdom

ALMA A. M. RAHAT, Swansea University, United Kingdom

JONATHAN E. FIELDSSEND, University of Exeter, United Kingdom

The performance of acquisition functions for Bayesian optimisation is investigated in terms of the Pareto front between exploration and exploitation. We show that Expected Improvement and the Upper Confidence Bound always select solutions to be expensively evaluated on the Pareto front, but Probability of Improvement is never guaranteed to do so and Weighted Expected Improvement does 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, 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.

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

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 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, e.g. [18, 28]. 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 at the location which 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 [2, 8, 13, 38], 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

Authors' addresses: George De Ath, g.de.ath@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; Jonathan E. Fieldsend, j.e.fieldsend@exeter.ac.uk, Department of Computer Science, University of Exeter, Exeter, United Kingdom.

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  $i = 1 \rightarrow M$  do
3:    $f_i \leftarrow f(\mathbf{x}_i)$  ▷ Expensively evaluate all initial samples
4: for  $i = M \rightarrow T$  do
5:    $\mathcal{GP} \leftarrow \text{TrainGP}(X, \mathbf{f})$  ▷ Train a  $\mathcal{GP}$  model
6:    $\mathbf{x}' \leftarrow \text{argmax}_{\mathbf{x} \in X} \alpha(\mathbf{x}, \mathcal{GP})$  ▷ Maximise infill criterion
7:    $f' \leftarrow f(\mathbf{x}')$  ▷ Expensively evaluate  $\mathbf{x}'$ 
8:    $X \leftarrow X \cup \{\mathbf{x}'\}$  ▷ Augment data
9:    $\mathbf{f} \leftarrow \mathbf{f} \cup \{f'\}$ 
10: return  $X, \mathbf{f}$ 

```

with the most promising (exploitative) value, but occasionally use an alternative solution selected at random from either the estimated Pareto set or the feasible space.

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. Extensive empirical evaluations on well-known test problems and a real world computational fluid dynamics optimisation are presented in Section 4.

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 [20] in the early 1960s, and later improved and popularised by Moćkus et al. [25] and Jones et al. [18]. A recent review of the topic can be found in [28].

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

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

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 [18]. It starts (line 1) with a space filling design (e.g. Latin hypercube sampling [23]) of the parameter space, constructed independent of the function space. The samples $X = \{\mathbf{x}_i\}_{i=1}^M$ from this initial design are then (expensively) evaluated with the function, $f_i = f(\mathbf{x}_i)$, 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 5). 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 (\mathcal{GP} s), 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 $\alpha(\mathbf{x})$ 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 6), 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 7 to 9). 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 [27]. In essence, a \mathcal{GP} is a collection of random variables, and any finite number of these have a joint Gaussian distribution [27]. 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 $\sigma^2(\mathbf{x})$:

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

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 [29]. 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 evaluation.¹

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

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 $\alpha(\mathbf{x}, \mathcal{D})$ 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 (\mathcal{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 $\alpha(\mathbf{x}, \mathcal{D})$ depends on \mathcal{D} and on the hyperparameters of the \mathcal{GP} , for clarity we suppress this dependence and write $\alpha(\mathbf{x})$. The predictive distribution (2) is Gaussian, with mean and variance given by (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

¹We use the L-BFGS algorithm with 10 restarts to estimate the hyper-parameters [12].

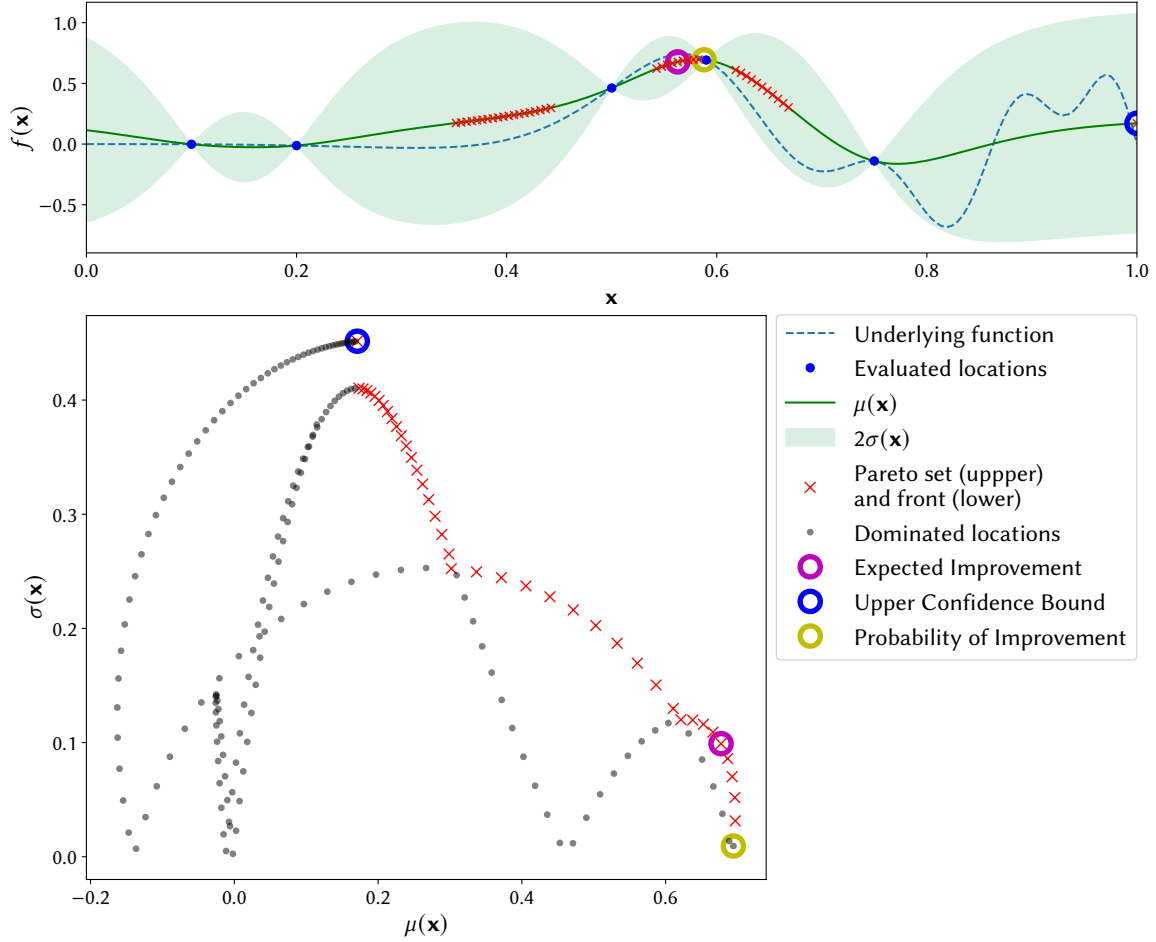


Fig. 1. Example Pareto front: *Top*: Gaussian Process approximation to a function (blue dashed curve) resulting from the 5 observations shown; mean $\mu(x)$ is shown in dark green and twice the posterior standard deviation $\sigma(x)$ is shown as the light green envelopes. *Bottom*: 200 samples uniformly 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.

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, we say that a location \mathbf{x} dominates \mathbf{x}' , written $\mathbf{x} < \mathbf{x}'$, iff $\mu(\mathbf{x}) > \mu(\mathbf{x}')$ and $\sigma(\mathbf{x}) > \sigma(\mathbf{x}')$. We present BO procedures that select solutions from the Pareto optimal set of locations which are not dominated by any other feasible locations:

$$\mathcal{P} = \{\mathbf{x} \in \mathcal{X} \mid \mathbf{x}' \not< \mathbf{x} \forall \mathbf{x}' \in \mathcal{X}\}. \quad (6)$$

Figure 1 illustrates the Pareto front, $\{(\mu(\mathbf{x}), \sigma(\mathbf{x})) \mid \mathbf{x} \in \mathcal{P}\}$, for a simple one-dimensional function. Note that the Pareto set is disjoint in \mathcal{X} and in (μ, σ) space. In this case, the locations maximising three popular acquisition functions, Expected Improvement, Upper Confidence Bound and Probability of Improvement, are elements of the Pareto set. The Pareto set identifies plausible locations for the next expensive evaluation of $f(x)$.




Fig. 2. Contours of upper confidence bound (UCB, $\beta = 1$), 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 $[0, 1]$.

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 [21] 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 [31]. The UCB acquisition function is a weighted sum of the mean prediction and uncertainty:

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

where $\beta \geq 0$ is the weight. The addition of a multiple of the uncertainty means that the criterion prefers locations where the mean is large (exploitation) or mean combined with the uncertainty is sufficiently large to warrant exploration.

When $\beta = 0$ 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 is large, the optimisation becomes purely exploratory, evaluating the location where the posterior uncertainty is largest. This is equivalent to reducing the overall entropy of the model [31]. Consequently, it may eventually locate the global optima, but the rate of convergence may be very slow.

Some authors suggest tuning β during the course of the optimisation [28]; indeed Srinivas et al.'s convergence proof depends on a particular schedule [31].

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, $\mathbf{x}' = \operatorname{argmax}_{\mathbf{x} \in \mathcal{S}} \alpha_{UCB}(\mathbf{x})$, 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 expected improvement (EI) is perhaps the most popular infill criterion and is very widely used. It was first proposed by Moćkus et al. [25], and further developed by Jones et al. [18]. Bull has shown that, under certain conditions, BO using EI is guaranteed to converge to the global optimum [3].

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 [18]:

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

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 which are very likely to be a little better than f^* with the exploration of others which 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 [18]. 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 [8, 30] have associated the term, $\sigma(\mathbf{x})s\Phi(s) = (\mu(\mathbf{x}) - f^*)\Phi(s)$, in (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. [30] 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 $1 - 2\omega < 0$ 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

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




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.

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 (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 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 0.9 ; cf Figure 2 for $\omega = 0.5$. We note that when $\omega = 0$, when the acquisition function might be expected to yield purely exploratory behaviour, that mean predictions close to f^* , but with large variance are preferred. When $\omega = 0.9$, which might be expected to result in purely exploitative search, the contours of α_{WEI} are not parallel with the μ direction.

2.2.4 Probability of Improvement. The Probability of Improvement (PI) is one of the earliest proposed infill criteria [20]. 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:

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

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 (20) shows that PI increases with uncertainty σ . However, in contrast to EI and UCB, when $\mu(\mathbf{x}) > f^*$ then (20) shows that PI decreases with σ ; that is, 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 [16]. 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 [16, 20, 22]. 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} (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 [31].

These observations lead us, like [2, 8, 13, 38], to 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. We note that while proofs of convergence for particular trade-offs between exploration and exploitation exist [3, 31], 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 that select $\mathbf{x}' = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x})$ are liable to become stuck at a local optimum, while purely exploratory searches that choose $\mathbf{x}' = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \sigma(\mathbf{x})$ will eventually locate the optimum, but only very slowly as even very unpromising locations where $\mu(\mathbf{x}) \ll f^*$ are visited in order to minimise the posterior variance/entropy everywhere.

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.

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

```

1: if rand() <  $\epsilon$  then
2:    $\tilde{\mathcal{P}} \leftarrow \text{MOOptimise}_{\mathbf{x} \in \mathcal{X}}(\mu(\mathbf{x}), \sigma(\mathbf{x}))$ 
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: Selection from feasible space.

```

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

```

3 ϵ -GREEDY BAYESIAN OPTIMISATION

Motivated by the success of ϵ -greedy schemes reinforcement learning [24, 32–34], we propose two novel BO infill criteria which use the ϵ -greedy methodology to select the next point for expensive evaluation. 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. This replaces line 6 in standard BO, i.e. Algorithm 1. The approximate Pareto set of model predictions is found using a standard evolutionary optimiser (MOOptimise); here we use NSGA-II [7]. 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 it selects a random location 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 in [3], 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 well-known benchmark functions with a range of domain sizes and dimensionality; see Table 1 for details. Their performance is compared to the infill criteria discussed in Section 2.2, namely Expected Improvement (EI), Upper Confidence Bound (UCB) and Probability of Improvement (PI). We also compare performance with the following alternative strategies that select locations on the Pareto front:

Exploit	Pure exploitation: $\mathbf{x}' = \text{argmax}_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x})$
Explore	Pure exploration: $\mathbf{x}' = \text{argmax}_{\mathbf{x} \in \mathcal{X}} \sigma(\mathbf{x})$
PFRandom	Selection of a solution on the approximated Pareto front at random: $\mathbf{x}' = \text{randomChoice}(\tilde{\mathcal{P}})$

These were selected for evaluation as they represent extreme examples of solution selection strategies employable using the approximated Pareto front. In addition, we compare the performance of all the infill criteria with the quasi-random search produced by max-min Latin Hypercube Sampling (LHS, [23]).

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

Name	Domain	d	Name	Domain	d
WangFreitas [35]	[0, 1]	1	logSixHumpCamel [†]	$[-3, 2] \times [3, 2]$	2
Branin [†]	$[-5, 0] \times [10, 15]$	2	modHartman6 [†]	$[0, 1]^d$	6
BraninForrester [9]	$[-5, 0] \times [10, 15]$	2	logGSobol [10]	$[-5, 5]^d$	10
Cosines [11]	$[0, 0] \times [5, 5]$	2	logRosenbrock [†]	$[-5, 10]^d$	10
logGoldsteinPrice [†]	$[-2, -2] \times [2, 2]$	2	logStyblinskiTang [†]	$[-5, 5]^d$	10

Table 1. Functions used in these experiments, along with their domain and dimensionality, d . Formulae can be found as cited or at <http://www.sfu.ca/~ssurjano/optimization.html> for those labelled with [†].

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	2.01	6.60×10^{-2}	2.01×10^{-1}	2.54×10^{-1}	8.42×10^{-3}	1.11×10^{-2}	6.56×10^{-2}	1.99×10^{-2}	1.77×10^{-1}	2.78×10^{-2}
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)		modHartman6 (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	5.42	7.64×10^{-1}	4.54×10^{-2}	1.77×10^{-2}	8.28	8.43×10^{-1}	7.92	4.46×10^{-1}	2.72	1.89×10^{-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}

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. The method with the lowest median performance is shown in red, with those with statistically equivalent performance are shown in blue.

statistical performance measures to be used, each optimisation was repeated 51 times. The same sets of initial samples were used for each method’s runs to allow for paired statistical comparisons to be carried out between the methods. In all experiments a value of $\epsilon = 0.1$ was used for both ϵ -PF and ϵ -RS. The UCB algorithm was run with a value of β as defined for continuous functions in Theorem 2 of [31] with $a = b = 1$. All acquisition functions were optimised with NSGA-II [7], apart from PI which was optimised with DIRECT [17]; in both cases the optimiser had a budget of $5000d$ evaluations. For NSGA-II, we set the parameters to commonly used values: the population size was $100d$, the number of generations was 50, 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 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 minimum median f^* on each function is highlighted in red, and those which are statistically equivalent to the best method according to a one-sided paired Wilcoxon signed-rank test [19] with Holm-Bonferroni correction [14] ($p \geq 0.05$), are shown in blue.

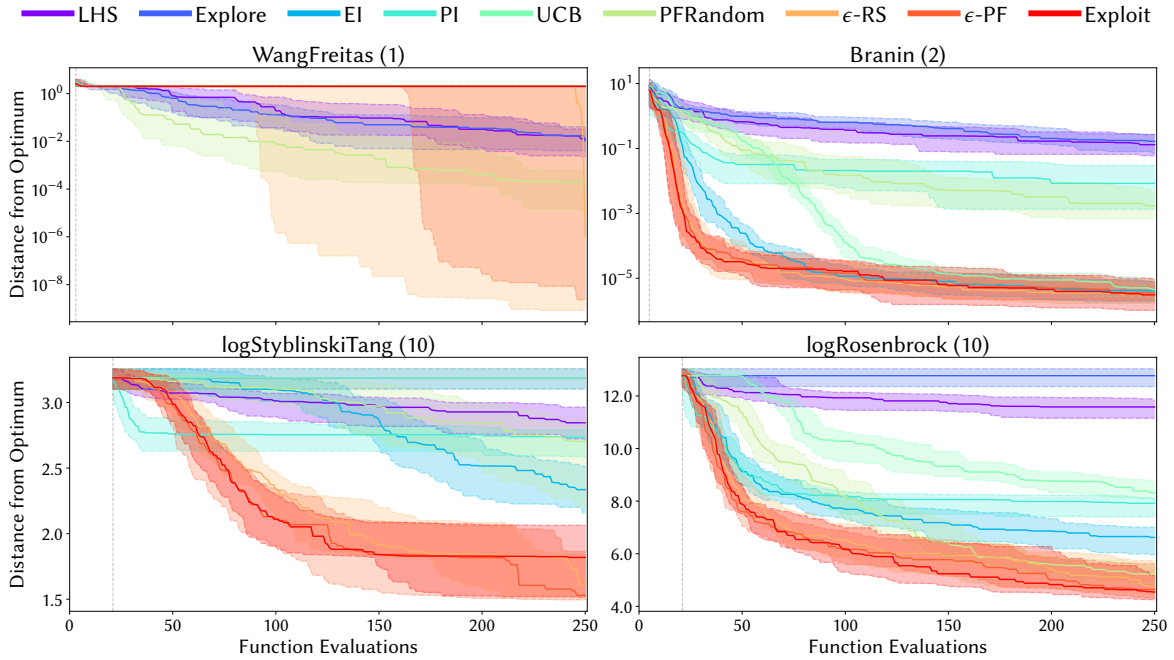


Fig. 4. Illustrative convergence plots for four benchmark problems. Each plot shows the median difference between the best function value seen and the true optimum, with shading representing the interquartile range across the 51 runs. The dashed vertical line indicates the end of the initial LHS phase.

Figure 4 shows the convergence of the various algorithms on illustrative four test problems in $d = 1, 2$ and 10 dimensions. Note that the benchmarks are couched as minimisation problems. Convergence plots for all the benchmark problems, together with Python code to generate the figures, are available as supplementary material².

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 and BraninForrester problems. The performance of PI is particularly poor, which may be due to its tendency to select solutions with low uncertainty when the mean prediction improves on f^* ; see Section 2.2.4. 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 [35] which contains a large local optimum and a narrow global optimum that is surrounded by plateaus; see supplementary material for a plot. Convergence on this problem is shown in Figure 4, which demonstrates how LHS sampling and a purely exploratory

²<http://www.github.com/XXXXX/XXXXX/>

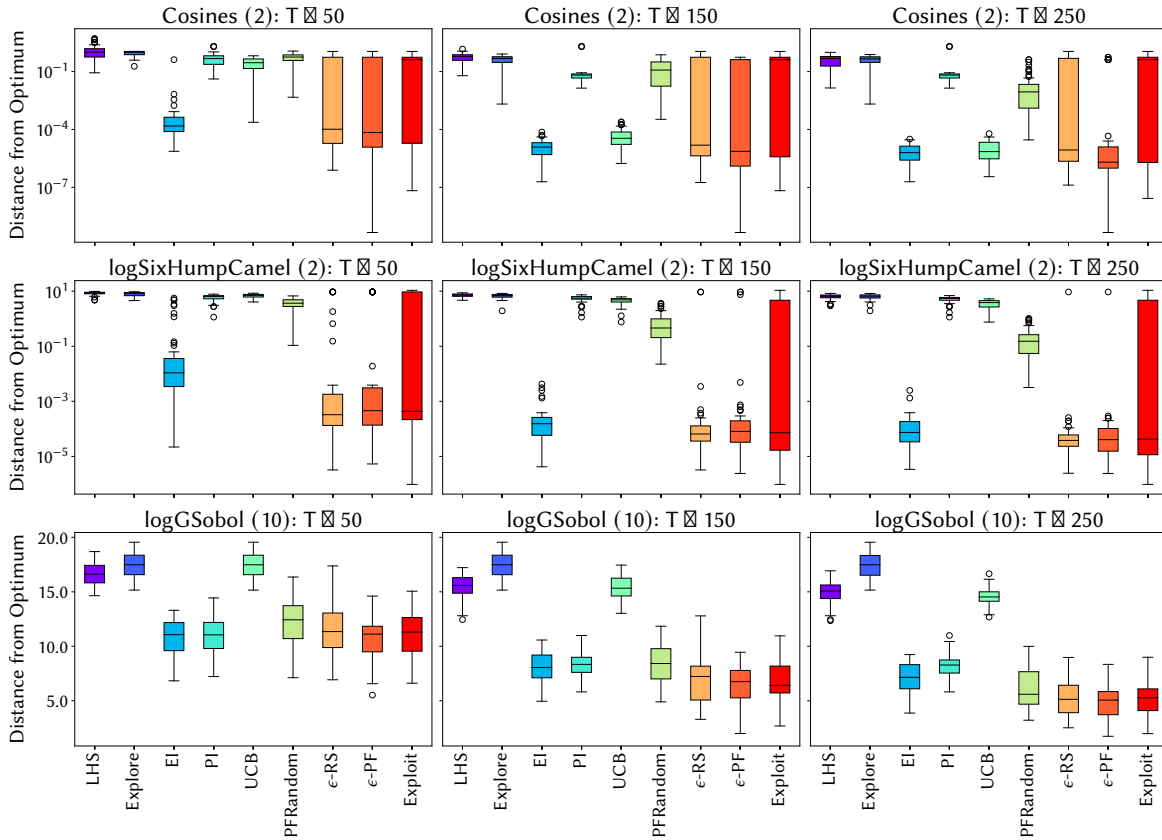


Fig. 5. Distribution of the best-seen function values after 50 (left), 150 (centre) and 250 (right) function evaluations on three benchmark problems.

strategy (Explore) converge slowly towards the optimum, while Exploit fails to find the vicinity of the optimum in any case. On the other hand, 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 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 in Figure 4, 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

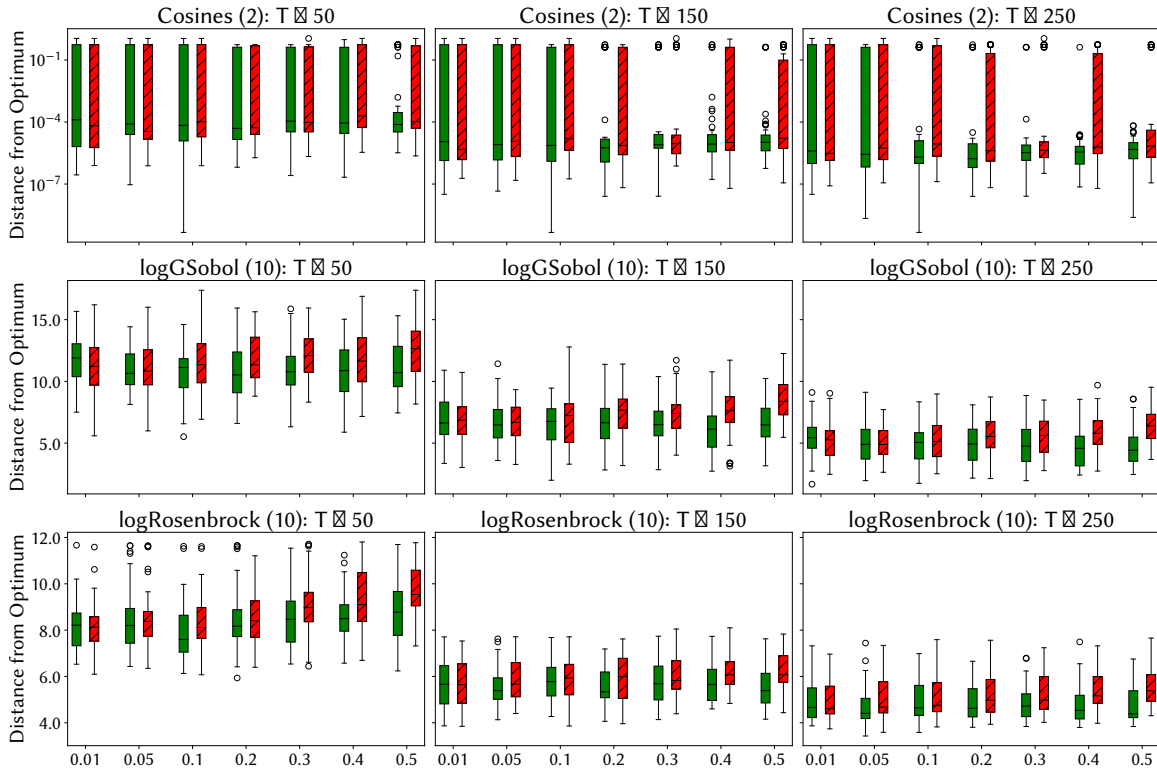


Fig. 6. Comparison of ϵ -PF (green) and ϵ -RS (red, hatched) for different values of ϵ (horizontal axis) after 50 (left), 150 (centre) and 250 (right) function evaluations.

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 used, as the convergence proofs in [31] rely on β 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 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 practise.

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 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. 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, with larger values giving 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.




Fig. 7. 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.

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 7, the PitzDaily test problem [5] involves optimising the shape of a pipe in order 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 7 and as described in detail by [5], we represent the wall geometry in terms of a Catmull-Clark sub-division curve, whose control points comprise the decision variables. Here there are 5 control points, resulting in a 10-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 8, 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 [26].

We observe that good solutions typically replace the step shown in Figure 7 with a slope, as illustrated in the two solutions in Figure 9. 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 [26]. 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 [6].




Fig. 8. Distribution of the best-seen function values after 50 (left), 150 (centre) and 250 (right) function evaluations on the real-world PitzDaily test problem.




Fig. 9. The streamlines for two solutions: the local optimum identified by Nilsson et al. [26] (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.

4.2 Real-World Application: Active Learning for Robot Pushing

Following Wang and Jegelka [37] and Jiang et al. [15], we optimise the control parameters for two active learning robot pushing problems [36]. In the first problem, illustrated in Figure 10, a robot hand (rectangle) is given the task of pushing an object (circle) towards 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 10, two robots (blue and green rectangles) in the same arena have to push their respective objects (circles) towards 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




Fig. 10. Two robot pushing tasks. `PUSH4` (*left*): a robot hand (blue rectangle) pushes the object (circle) towards a target 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 towards 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 towards 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.

a joint learning task. We treat this as a minimisation problem: the 8 parameters determining the robots’ paths are to be optimised to minimise the combined distance of the objects to their targets.

Like Wang and Jegelka [37], 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 [1] for a broader discussion on problem generators and generalisable results). Likewise, in `PUSH8` the object’s initial locations were fixed as shown in Figure 10 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, in order 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 [4] algorithm. 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 11 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

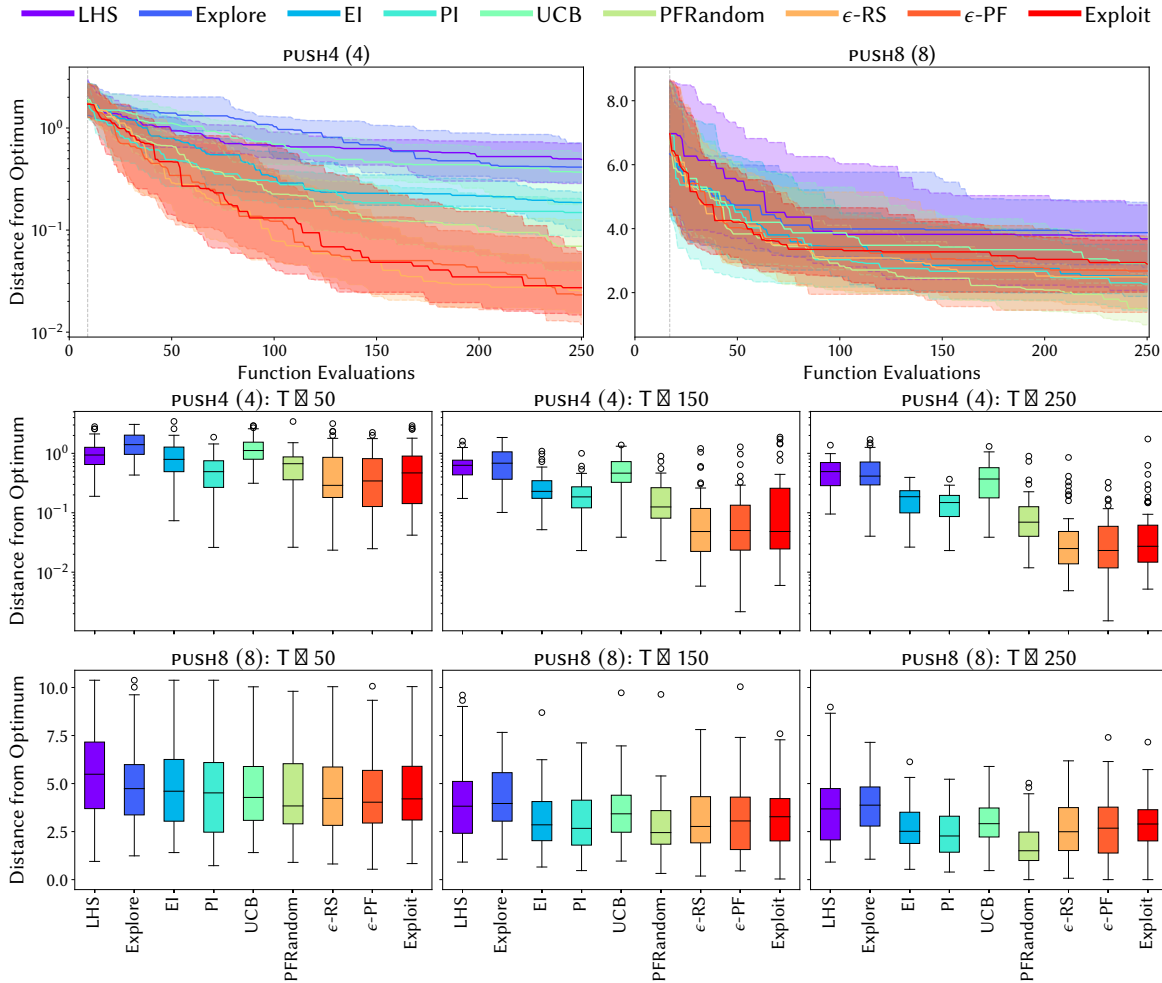


Fig. 11. Illustrative convergence plots for the two robot pushing problems (upper) and the distribution of the best-seen function values after 50 (left), 150 (centre), and 250 (right) evaluations for both problems.

also performed well in the 10-dimensional PitzDaily optimisation. We speculate that the PFRandom, which selects from the entire Pareto front at each iteration, owes its good performance to the 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.

ACKNOWLEDGMENTS

We thank Dr Steven Daniels for helping us prepare Figure 9. This work was supported by Innovate UK grant number 104400.

REFERENCES

- [1] Thomas Bartz-Beielstein. 2015. How to create generalizable results. In *Springer Handbook of Computational Intelligence*, Janusz Kacprzyk and Witold Pedrycz (Eds.). Springer, Berlin, Heidelberg, 1127–1142.
- [2] Bernd Bischl, Simon Wessing, Nadja Bauer, Klaus Friedrichs, and Claus Weihs. 2014. MOI-MBO: Multiobjective infill for parallel model-based optimization. In *International Conference on Learning and Intelligent Optimization*. Springer, 173–186.
- [3] Adam D. Bull. 2011. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research* 12, Oct (2011), 2879–2904.
- [4] 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.
- [5] 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.
- [6] 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. *International Journal of Computational Fluid Dynamics* (2019), 1–16.
- [7] 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.
- [8] 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. *Journal of Global Optimization* 61, 4 (2015), 677–694.
- [9] Alexander I. J. Forrester, Andras Sobester, and Andy J. Keane. 2008. *Engineering Design via Surrogate Modelling - A Practical Guide*. Wiley.
- [10] 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.
- [11] 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.
- [12] GPy. since 2012. GPy: A Gaussian process framework in Python. <http://github.com/SheffieldML/GPy>.
- [13] Carla Grobler, Schalk Kok, and Daniel N Wilke. 2017. Simple Intuitive Multi-objective Parallelization of Efficient Global Optimization: SIMPLE-EGO. In *World Congress of Structural and Multidisciplinary Optimisation*. Springer, 205–220.
- [14] Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70.
- [15] Shali Jiang, Henry Chai, Javier González, and Roman Garnett. 2019. Efficient nonmyopic Bayesian optimization and quadrature. arXiv:arXiv:1909.04568
- [16] Donald R. Jones. 2001. A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization* 21, 4 (2001), 345–383.

- [17] Donald R. Jones, Cary D. Perttunen, and Bruce E. Stuckman. 1993. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications* 79, 1 (1993), 157–181.
- [18] 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.
- [19] 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.
- [20] Harold J. Kushner. 1964. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal Basic Engineering* 86, 1 (1964), 97–106.
- [21] Tze Leung Lai and Herbert Robbins. 1985. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics* 6, 1 (1985), 4–22.
- [22] Daniel James Lizotte. 2008. *Practical Bayesian optimization*. Ph.D. Dissertation. University of Alberta.
- [23] 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.
- [24] 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.
- [25] Jonas Moćkus, Vytautas Tiešis, and Antanas Žilinskas. 1978. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization* 2, 1 (1978), 117–129.
- [26] 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.
- [27] Carl Edward Rasmussen and Christopher K. I. Williams. 2006. *Gaussian processes for machine learning*. The MIT Press, Boston, MA.
- [28] 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.
- [29] 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.
- [30] András Söbester, Stephen J. Leary, and Andy J. Keane. 2005. On the Design of Optimization Strategies Based on Global Response Surface Approximation Models. *Journal of Global Optimization* 33 (2005), 31–59.
- [31] 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.
- [32] Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA.
- [33] Michel Tokic. 2010. Adaptive ϵ -greedy exploration in reinforcement learning based on value differences. In *Annual Conference on Artificial Intelligence*. Springer, 203–210.
- [34] 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.
- [35] Ziyu Wang and Nando de Freitas. 2014. Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters. arXiv:arXiv:1406.7758
- [36] 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.
- [37] 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.
- [38] Antanas Žilinskas and James Calvin. 2019. Bi-objective decision making in global optimization based on statistical models. *Journal of Global Optimization* 74, 4 (2019), 599–609.

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

GEORGE DE ATH, University of Exeter, United Kingdom

RICHARD M. EVERSON, University of Exeter, United Kingdom

ALMA A. M. RAHAT, Swansea University, United Kingdom

JONATHAN E. FIELDSSEND, University of Exeter, United Kingdom

A SYNTHETIC FUNCTION DETAILS

In the following section we give the formulae of each of the 10 synthetic functions optimised in this work. Where functions have been modified from their standard form, we label the original functions as $g(\mathbf{x})$ and minimised function as $f(\mathbf{x})$.

A.1 WangFreitas

$$g(x) = 2 \exp\left(-\frac{1}{2} \left(\frac{x-a}{\theta_1}\right)^2\right) + 4 \exp\left(-\frac{1}{2} \left(\frac{x-b}{\theta_2}\right)^2\right) \quad (1)$$

$$f(x) = -g(x), \quad (2)$$

where $a = 0.1$, $b = 0.9$, $\theta_1 = 0.1$ and $\theta_2 = 0.01$.

A.2 Branin

$$f(\mathbf{x}) = a(x_2 - bx_1^2 + cx_1 - r)^2 + s(1-t) \cos(x_1) + s, \quad (3)$$

where $a = 1$, $b = \frac{5.1}{4\pi^2}$, $c = \frac{5}{\pi}$, $r = 6$, $s = 10$, $t = \frac{1}{8\pi}$ and x_i refers to the i -th element of vector \mathbf{x} .

A.3 BraninForrester

$$f(\mathbf{x}) = a(x_2 - bx_1^2 + cx_1 - r)^2 + s(1-t) \cos(x_1) + s + 5x_1, \quad (4)$$

where $a = 1$, $b = \frac{5.1}{4\pi^2}$, $c = \frac{5}{\pi}$, $r = 6$, $s = 10$, and $t = \frac{1}{8\pi}$.

A.4 Cosines

$$g(\mathbf{x}) = 1 - \sum_{i=1}^2 \left[(1.6x_i - 0.5)^2 - 0.3 \cos(3\pi(1.6x_i - 0.5)) \right] \quad (5)$$

$$f(\mathbf{x}) = -g(\mathbf{x}). \quad (6)$$

A.5 logGoldsteinPrice

$$g(\mathbf{x}) = (1 + (x_1 + x_2 + 1)^2 (19 - 14x_1 + 3x_1^2 - 14x_2 + 6x_1x_2 + 3x_2^2)) \quad (7)$$

$$\times (30 + (2x_1 - 3x_2)^2 (18 - 32x_1 + 12x_1^2 + 48x_2 - 36x_1x_2 + 27x_2^2))$$

$$f(\mathbf{x}) = \log(g(\mathbf{x})). \quad (8)$$

Authors' addresses: George De Ath, g.de.ath@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; Jonathan E. Fieldsend, j.e.fieldsend@exeter.ac.uk, Department of Computer Science, University of Exeter, Exeter, United Kingdom.

A.6 logSixHumpCamel

$$g(\mathbf{x}) = (4 - 2.1x_1^2 + \frac{x_1^4}{3})x_1^2 + x_1x_2 + (-4 + 4x_2^2)x_2^2 \quad (9)$$

$$f(\mathbf{x}) = \log(g(\mathbf{x}) + a + b), \quad (10)$$

where $a = 1.0316$ and $b = 10^{-4}$. Note that because $g(\mathbf{x})$ has a minimum value of -1.0316 , we add a plus a small constant (b) to avoid taking the logarithm of a negative number; this does not change the function's landscape.

A.7 modHartman6

$$g(\mathbf{x}) = - \sum_{i=1}^4 \alpha_i \exp\left(- \sum_{j=1}^6 A_{ij} (x_j - P_{ij})^2\right) \quad (11)$$

$$f(\mathbf{x}) = - \log(-g(\mathbf{x})) \quad (12)$$

where

$$\alpha = (1.0, 1.2, 3.0, 3.2)^T \quad (13)$$

$$\mathbf{A} = \begin{pmatrix} 10 & 3 & 17 & 3.50 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{pmatrix} \quad (14)$$

$$\mathbf{P} = 10^{-4} \begin{pmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{pmatrix}. \quad (15)$$

A.8 logGSobol

$$g(\mathbf{x}) = \prod_{i=1}^D \frac{4x_i - 1}{2} \quad (16)$$

$$f(\mathbf{x}) = \log(g(\mathbf{x})), \quad (17)$$

where $a = 1$ and $D = 10$.

A.9 logRosenbrock

$$g(\mathbf{x}) = \sum_{i=1}^{D-1} [100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2] \quad (18)$$

$$f(\mathbf{x}) = \log(g(\mathbf{x}) + 0.5), \quad (19)$$

where $D = 10$. Note, similarly to logSixHumpCamel, because $g(\mathbf{x})$ has a minimum value of 0, we add a value to ensure it is always positive.

A.10 logStyblinskiTang

$$g(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^D (x_i^4 - 16x_i^2 + 5x_i) \quad (20)$$

$$f(\mathbf{x}) = \log(g(\mathbf{x}) + 40D), \quad (21)$$

where $D = 10$. Again, because $g(\mathbf{x})$ has a minimum value of $-39.16599D$, we add $40D$ to it to ensure it is always positive.

B THE LANDSCAPE OF THE WANGFREITAS TEST PROBLEM




Fig. 1. The WangFreitas test problem. The blue line shows the true function, the green solid line shows the mean prediction of a \mathcal{GP} model trained on the red crosses, and the green areas depict the uncertainty (twice the standard deviation).

Figure 1 shows an illustration of the test problem (Equation 1) proposed by Wang and de Freitas [3]. It has one local optimum and a global optimum. The global optimum has a narrow basin surrounded by vast flat regions. Therefore it is easy for the model to become overconfident about the flatness in the vicinity of the optimum with no data identifying the basin, and mislead the search away from it. Consequently, methods with high exploration do well in solving this problem.

C FULL EXPERIMENTAL RESULTS

In this section we show the results table for the PitzDaily and robot pushing problems, and the convergence and box plots for all test problems evaluated in this work.

C.1 PitzDaily Results Table

Method	PitzDaily (10)	
	Median	MAD
Uniform	9.58×10^{-2}	3.52×10^{-3}
Explore	8.82×10^{-2}	4.82×10^{-3}
EI	8.42×10^{-2}	1.43×10^{-3}
PI	9.74×10^{-2}	6.49×10^{-3}
UCB	8.55×10^{-2}	2.96×10^{-3}
PFRandom	8.36×10^{-2}	9.72×10^{-4}
ϵ -RS	8.49×10^{-2}	2.68×10^{-3}
ϵ -PF	8.45×10^{-2}	2.44×10^{-3}
Exploit	8.40×10^{-2}	1.82×10^{-3}

Table 1. Median absolute distance (left) and median absolute deviation from the median (MAD, right) from the optimum after 250 function evaluations, across the 51 runs. The method with the lowest median performance is shown in red, with those with statistically equivalent performance are shown in blue.

The full results of the optimisation runs on the PitzDaily test problem are shown in Table 1. It shows the median difference between the estimated optimum and the true optimum over the 51 repeated experiments, together with the median absolute deviation from the median (MAD). The method with the minimum median on each function is highlighted in red, and those which are statistically equivalent to the best method according to a one-sided paired Wilcoxon signed-rank test [2] with Holm-Bonferroni correction [1] ($p \geq 0.05$), are shown in blue.

C.2 Robot Pushing Results Table

Method	PUSH4 (4)		PUSH8 (8)	
	Median	MAD	Median	MAD
LHS	4.93×10^{-1}	3.08×10^{-1}	3.68	2.18
Explore	4.14×10^{-1}	2.41×10^{-1}	3.88	1.44
EI	1.86×10^{-1}	1.05×10^{-1}	2.52	1.07
PI	1.49×10^{-1}	8.91×10^{-2}	2.27	1.39
UCB	3.70×10^{-1}	2.90×10^{-1}	2.91	1.19
PFRandom	6.95×10^{-2}	6.71×10^{-2}	1.50	1.07
ϵ -RS	2.50×10^{-2}	2.17×10^{-2}	2.49	1.56
ϵ -PF	2.32×10^{-2}	2.47×10^{-2}	2.68	1.80
Exploit	2.73×10^{-2}	2.51×10^{-2}	2.89	1.23

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. The method with the lowest median performance is shown in red, with those with statistically equivalent performance are shown in blue.

The full results of the optimisation runs on the PUSH4 and PUSH8 test problems are shown in Table 2. It shows the median difference between the estimated optimum and the true optimum over the 51 repeated experiments, together with the median absolute deviation from the median (MAD). The method with the minimum median on each function is highlighted in red, and those which are statistically equivalent to the best method according to a one-sided paired Wilcoxon signed-rank test [2] with Holm-Bonferroni correction [1] ($p \geq 0.05$), are shown in blue.

C.3 Convergence Histories and Boxplots

In this section we display the full set of results for the experimental evaluations carried out in this paper. Each figure shows the convergence of each algorithm on the respective test problem (top), snapshots of their performance at 50, 150, and 250 function evaluations (centre), and the comparative performance between ϵ -PF (green) and ϵ -RS (red, hatched) for increasing values of ϵ (lower).

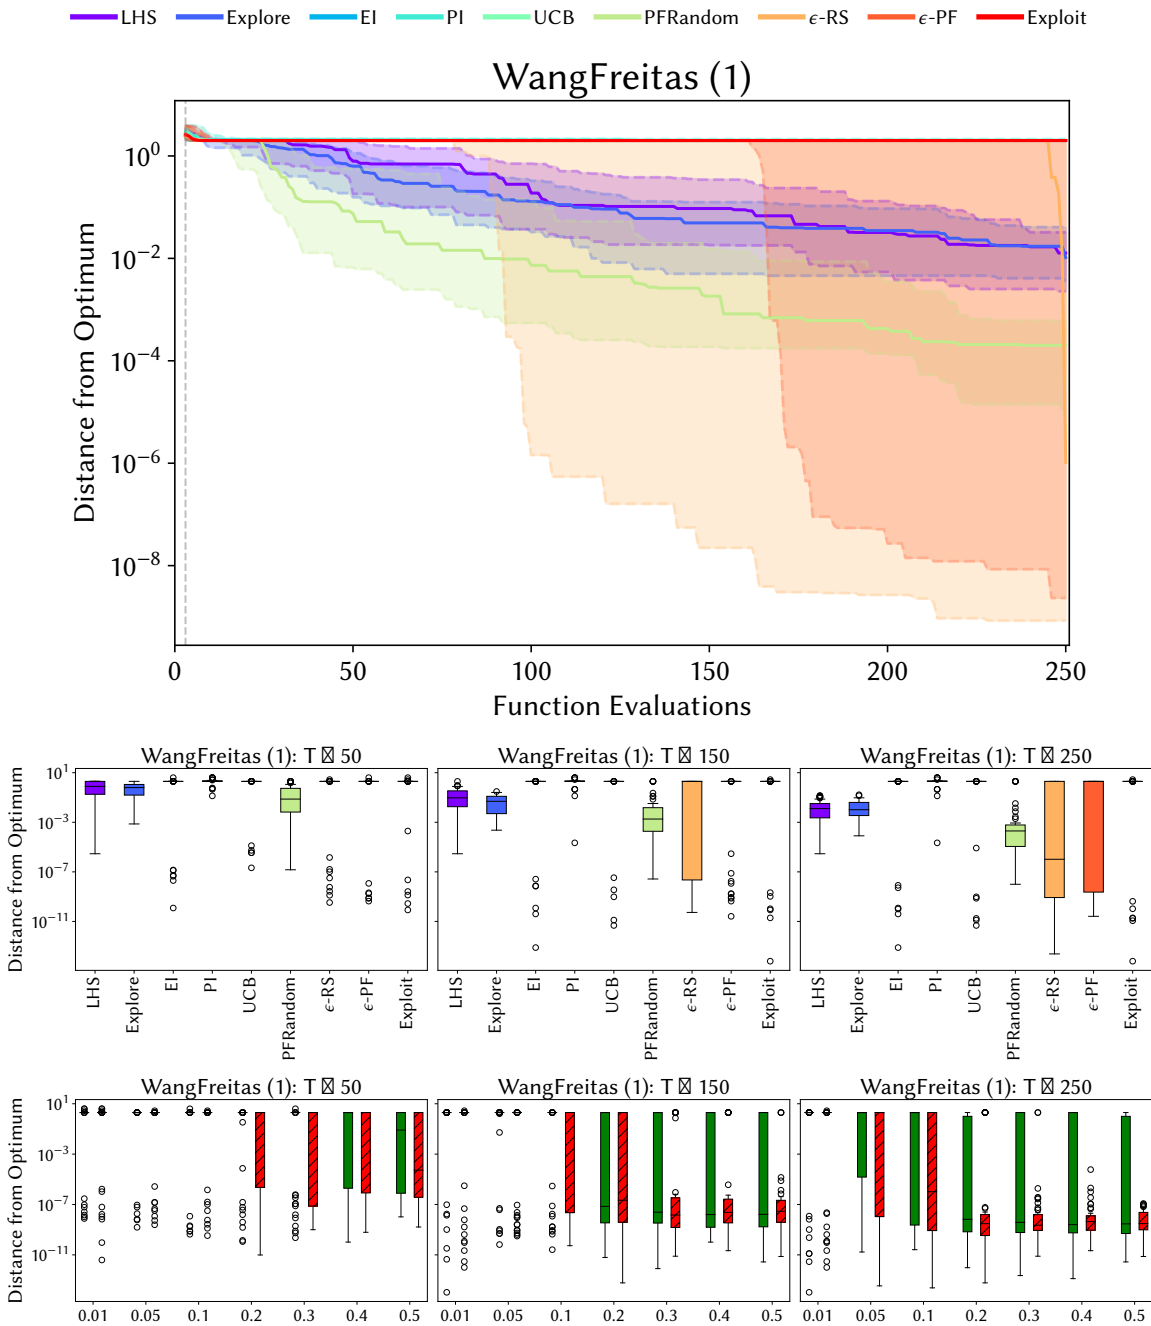


Fig. 2. Results for the one-dimensional WangFreitas test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

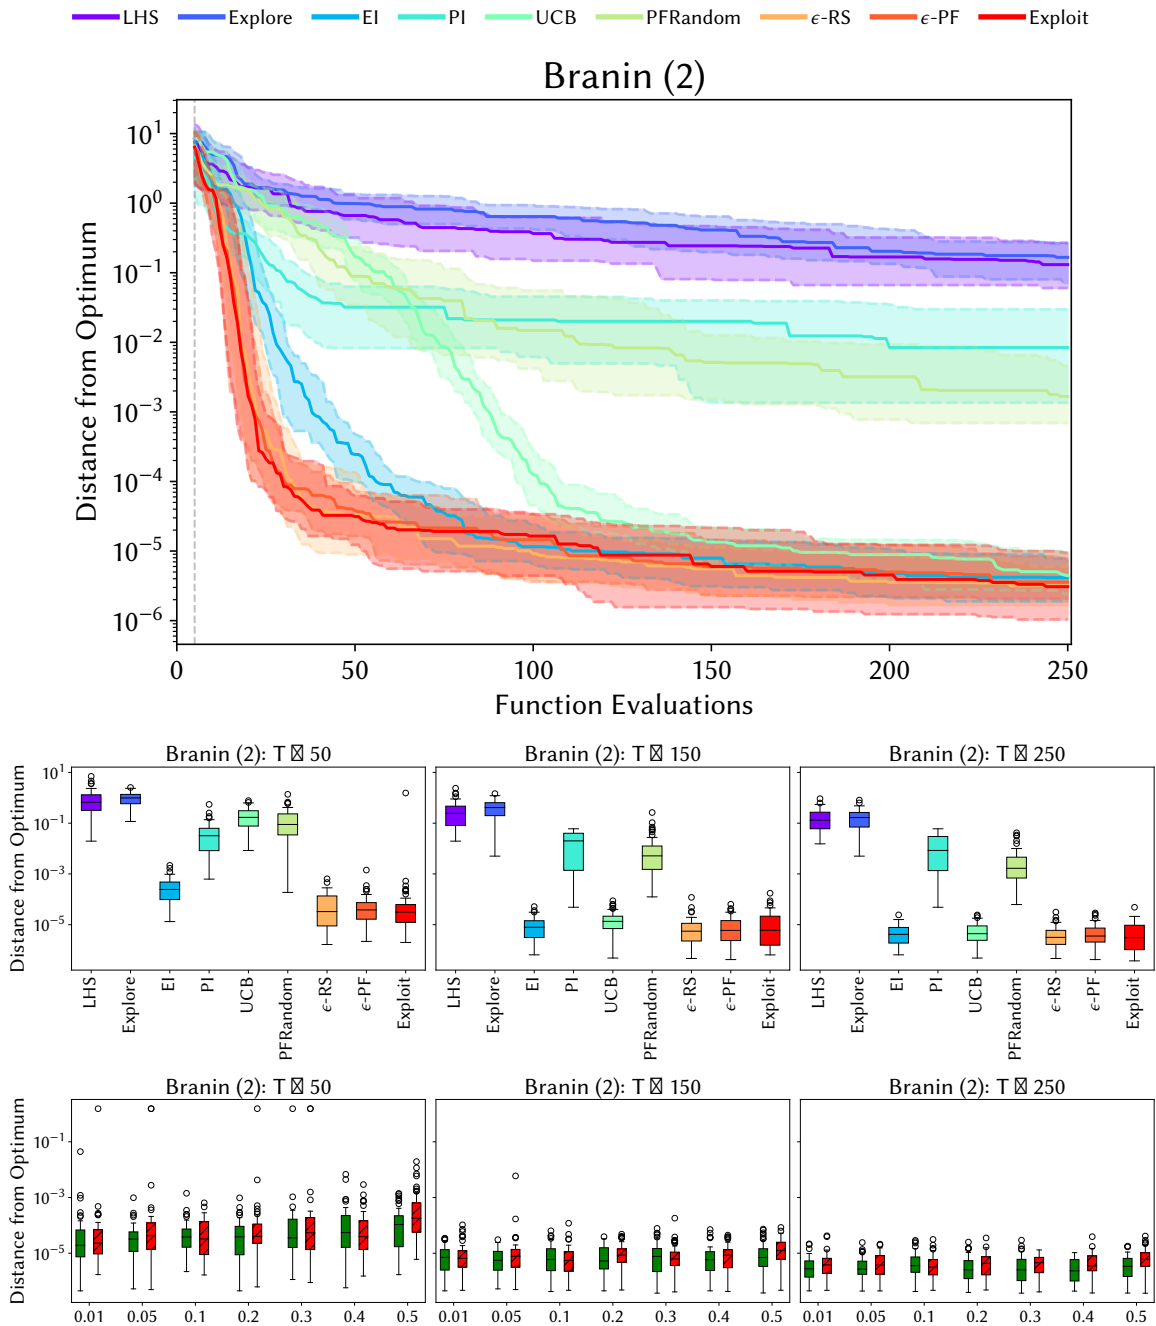


Fig. 3. Results for the two-dimensional Branin test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

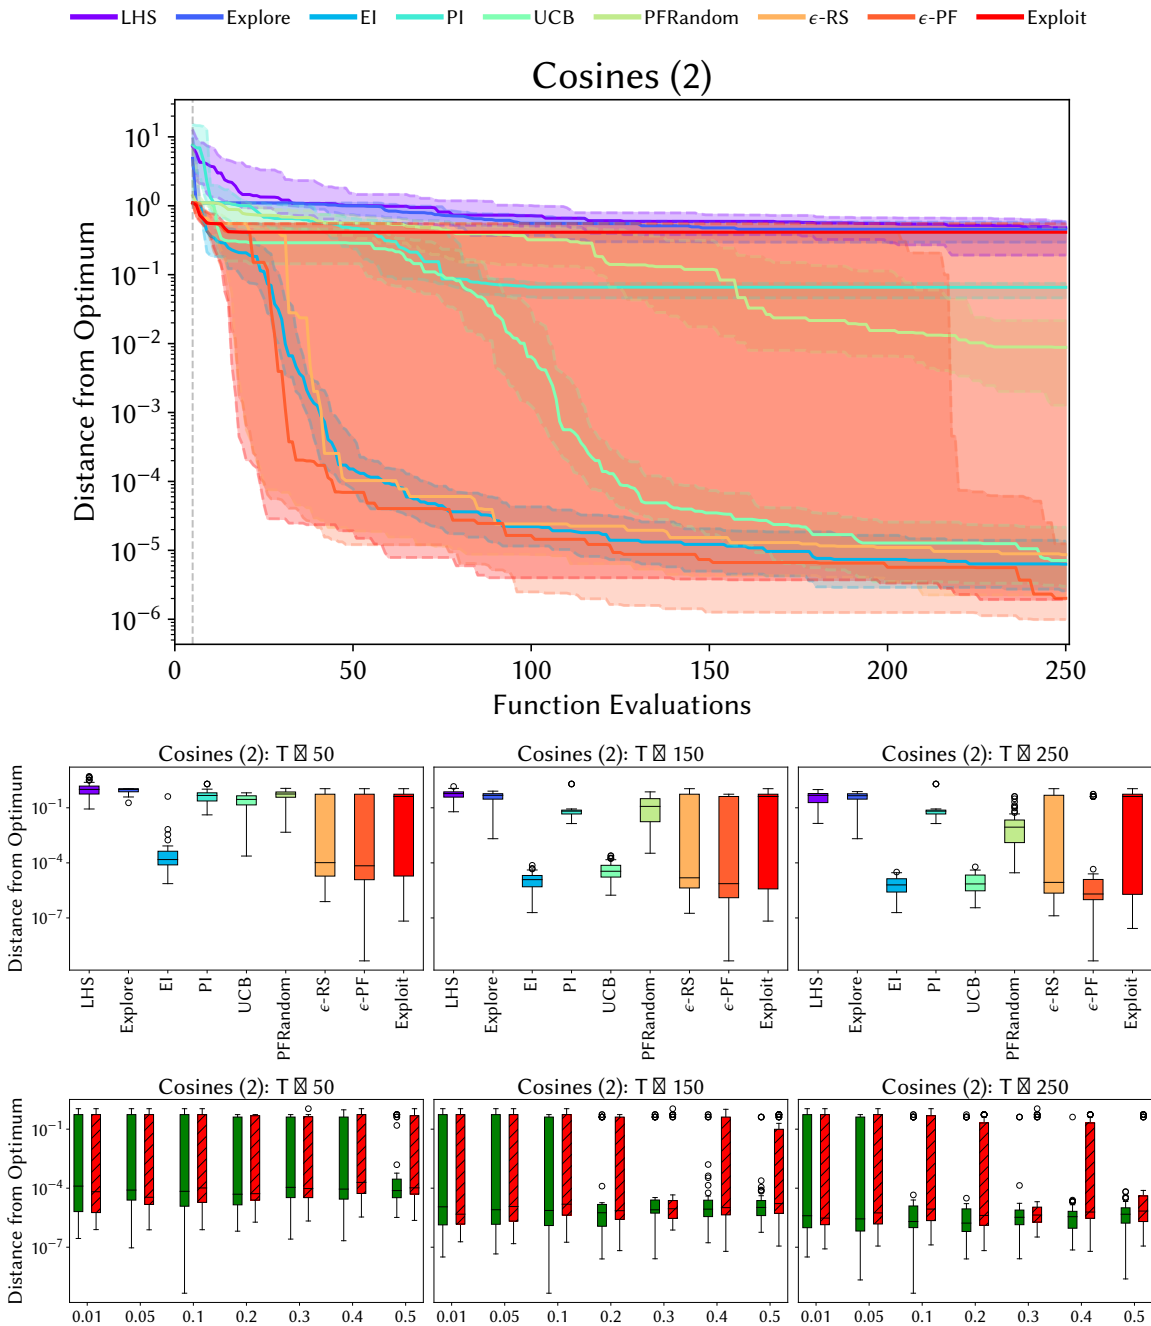


Fig. 4. Results for the two-dimensional Cosines test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

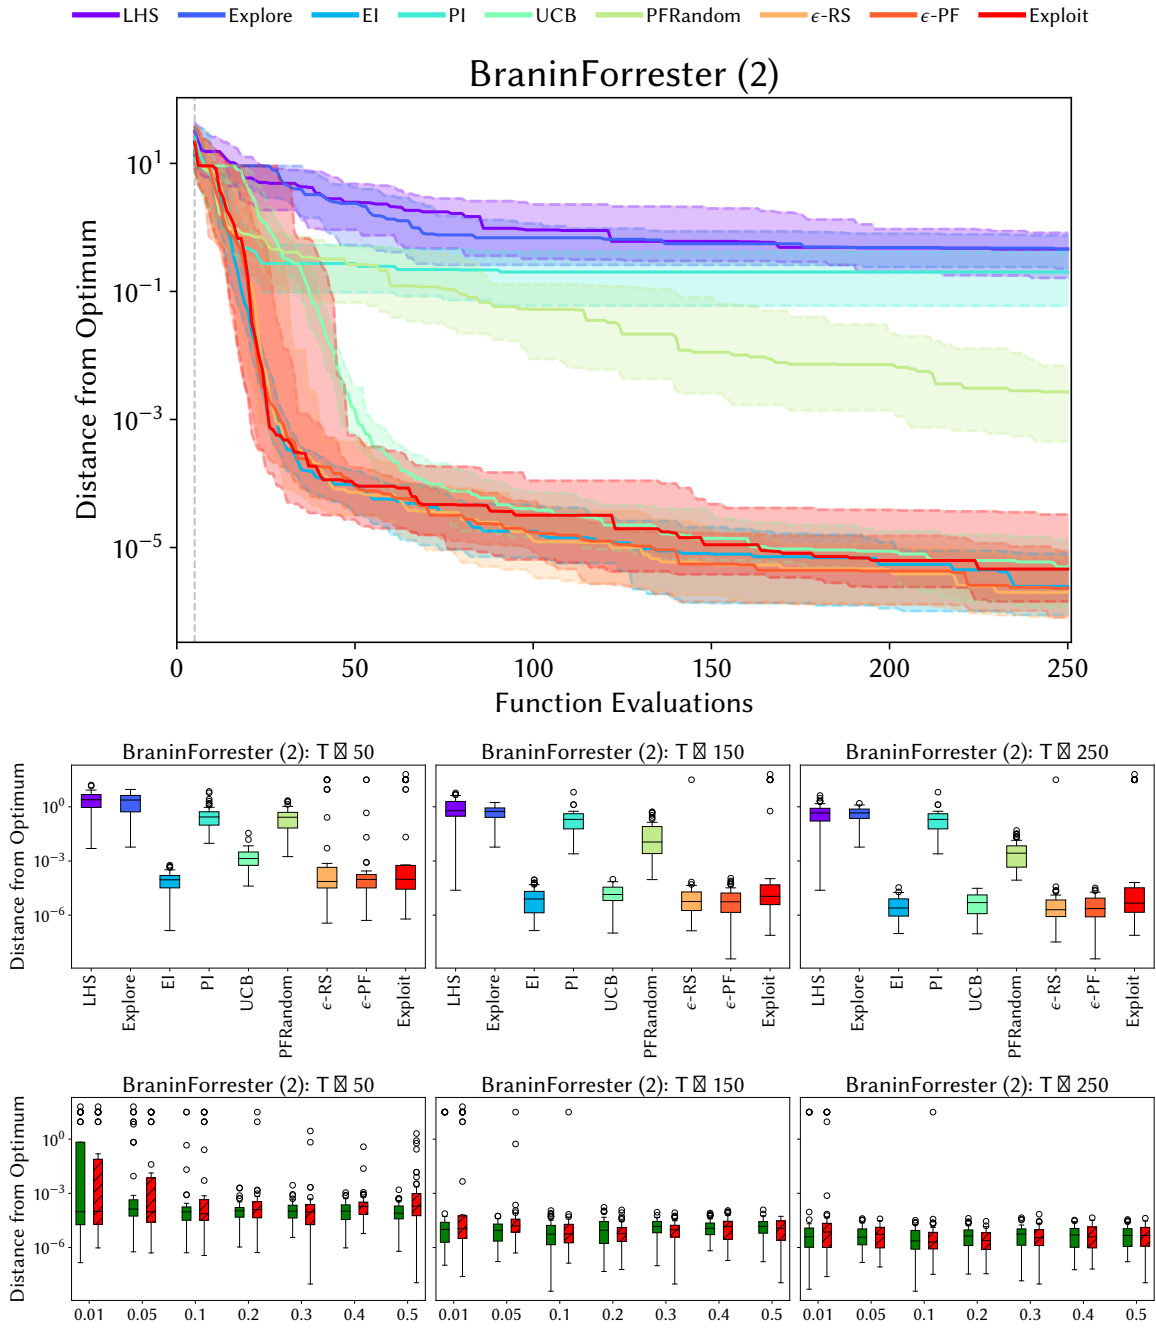


Fig. 5. Results for the two-dimensional BraninForrester test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

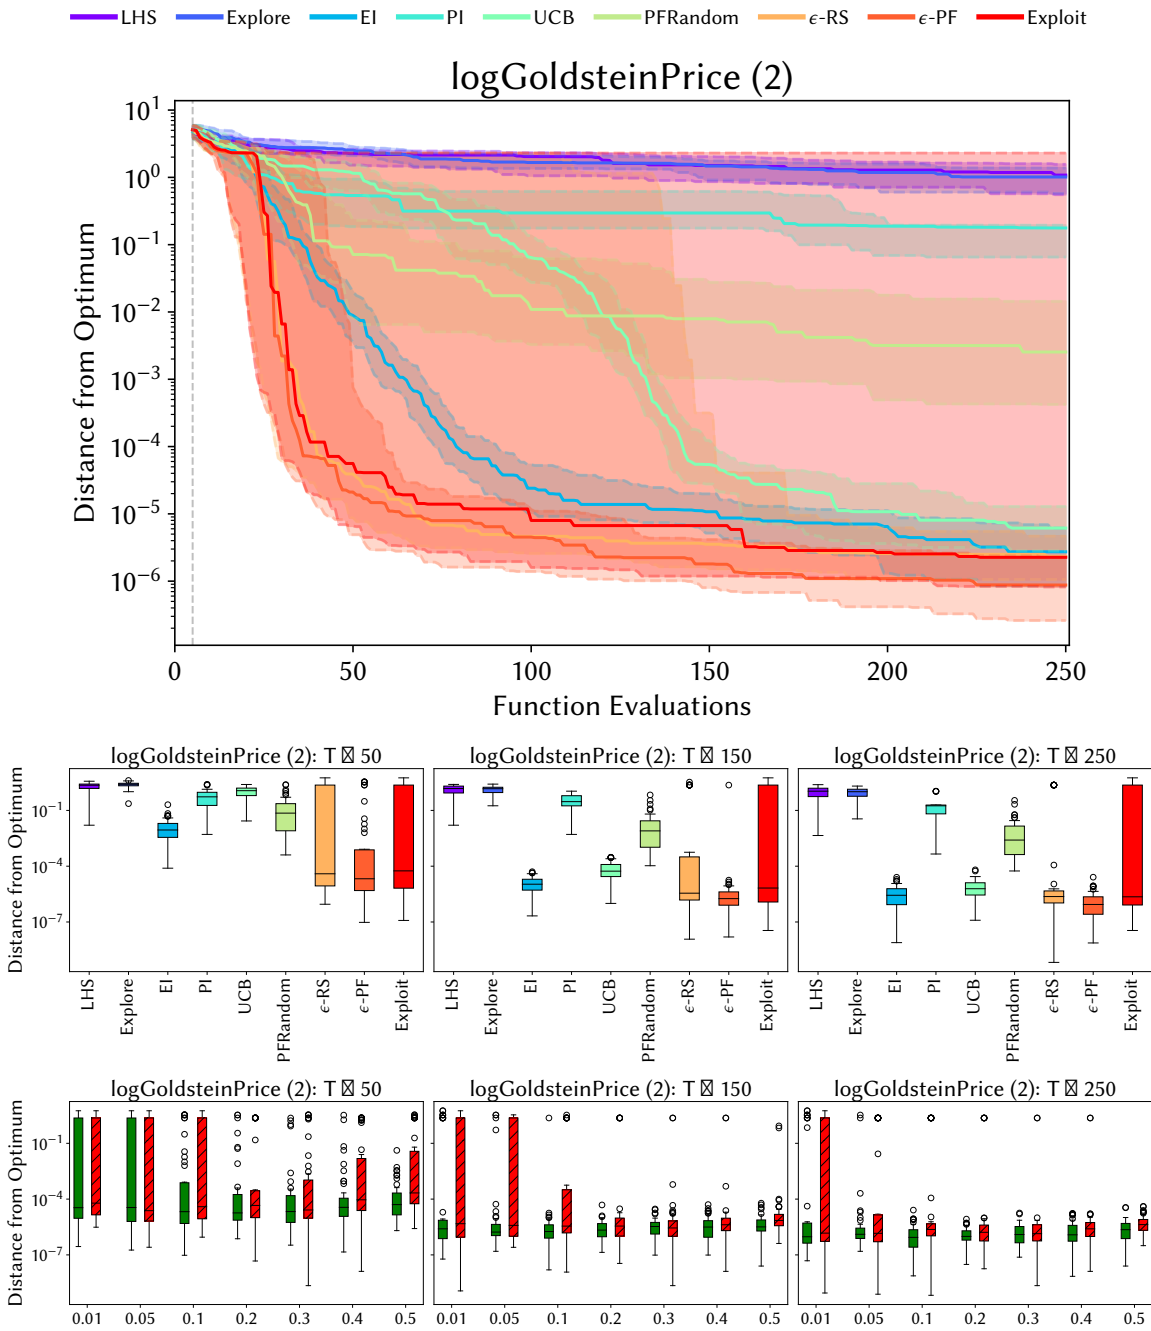


Fig. 6. Results for the two-dimensional logGoldsteinPrice test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

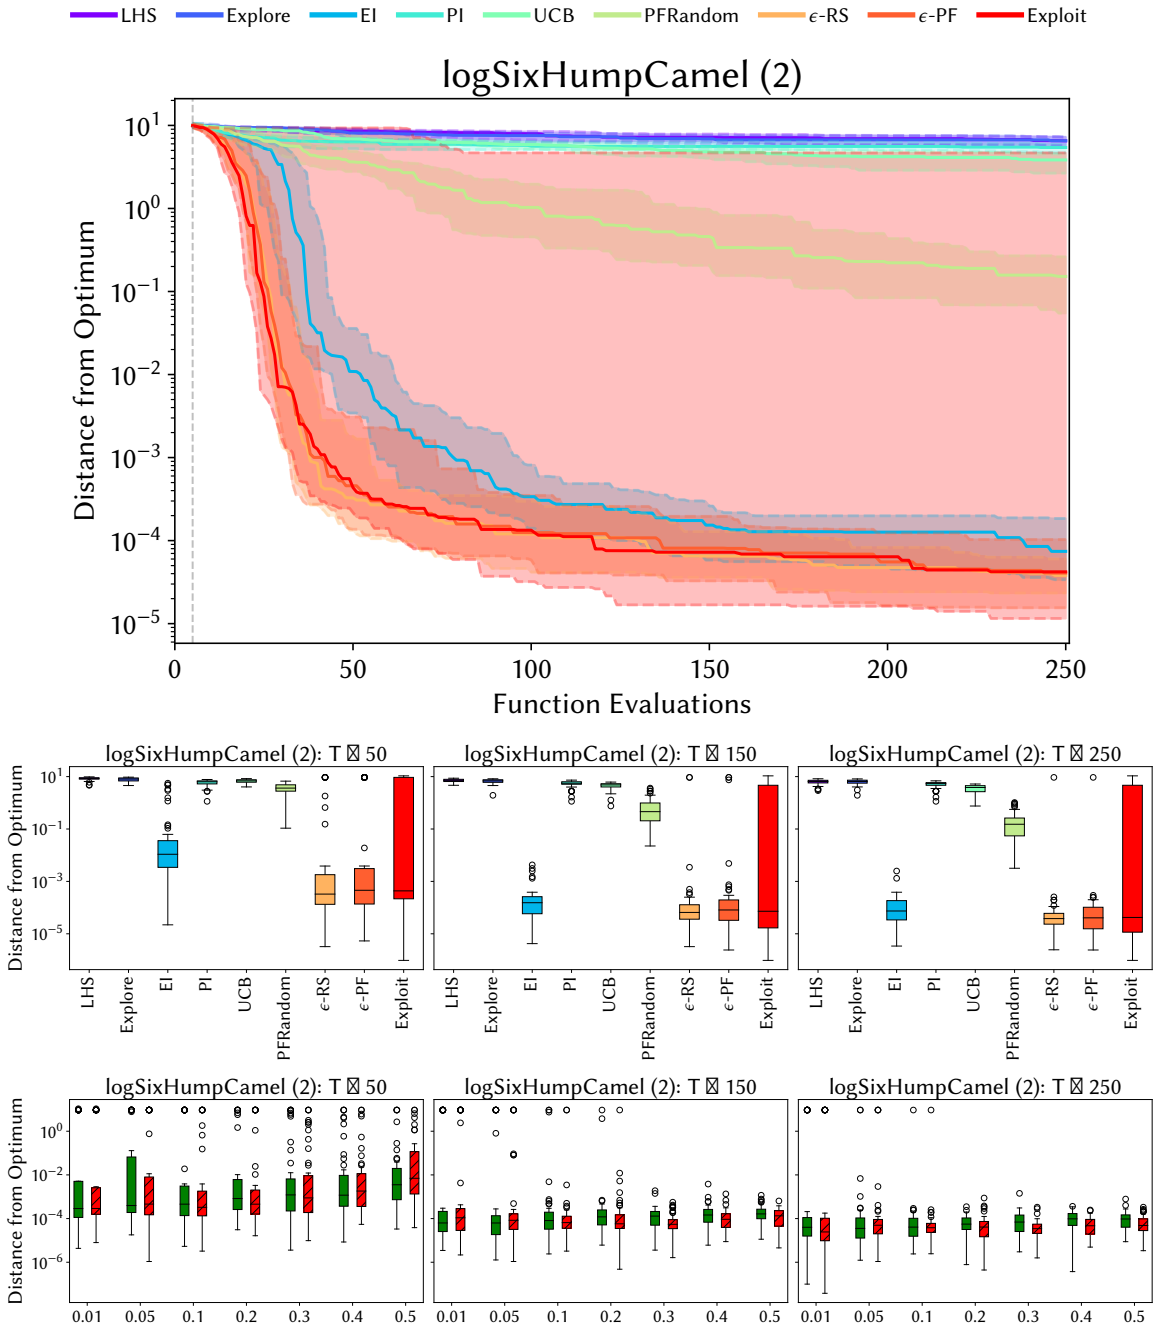


Fig. 7. Results for the two-dimensional logSixHumpCamel test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

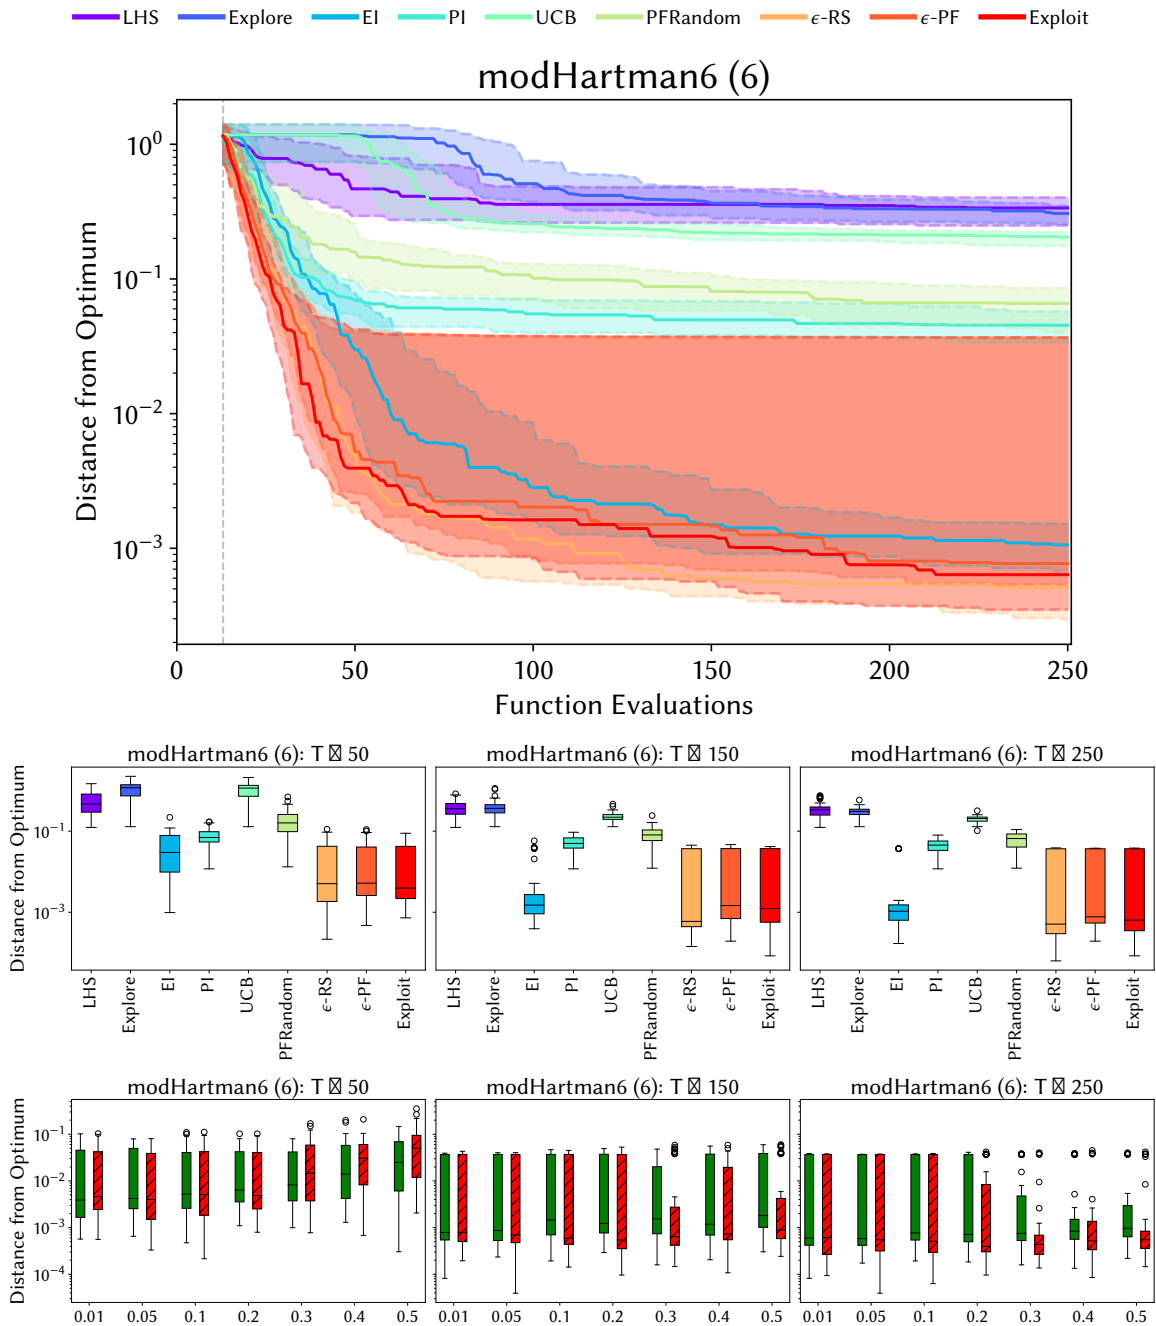


Fig. 8. Results for the six-dimensional modHartman6 test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

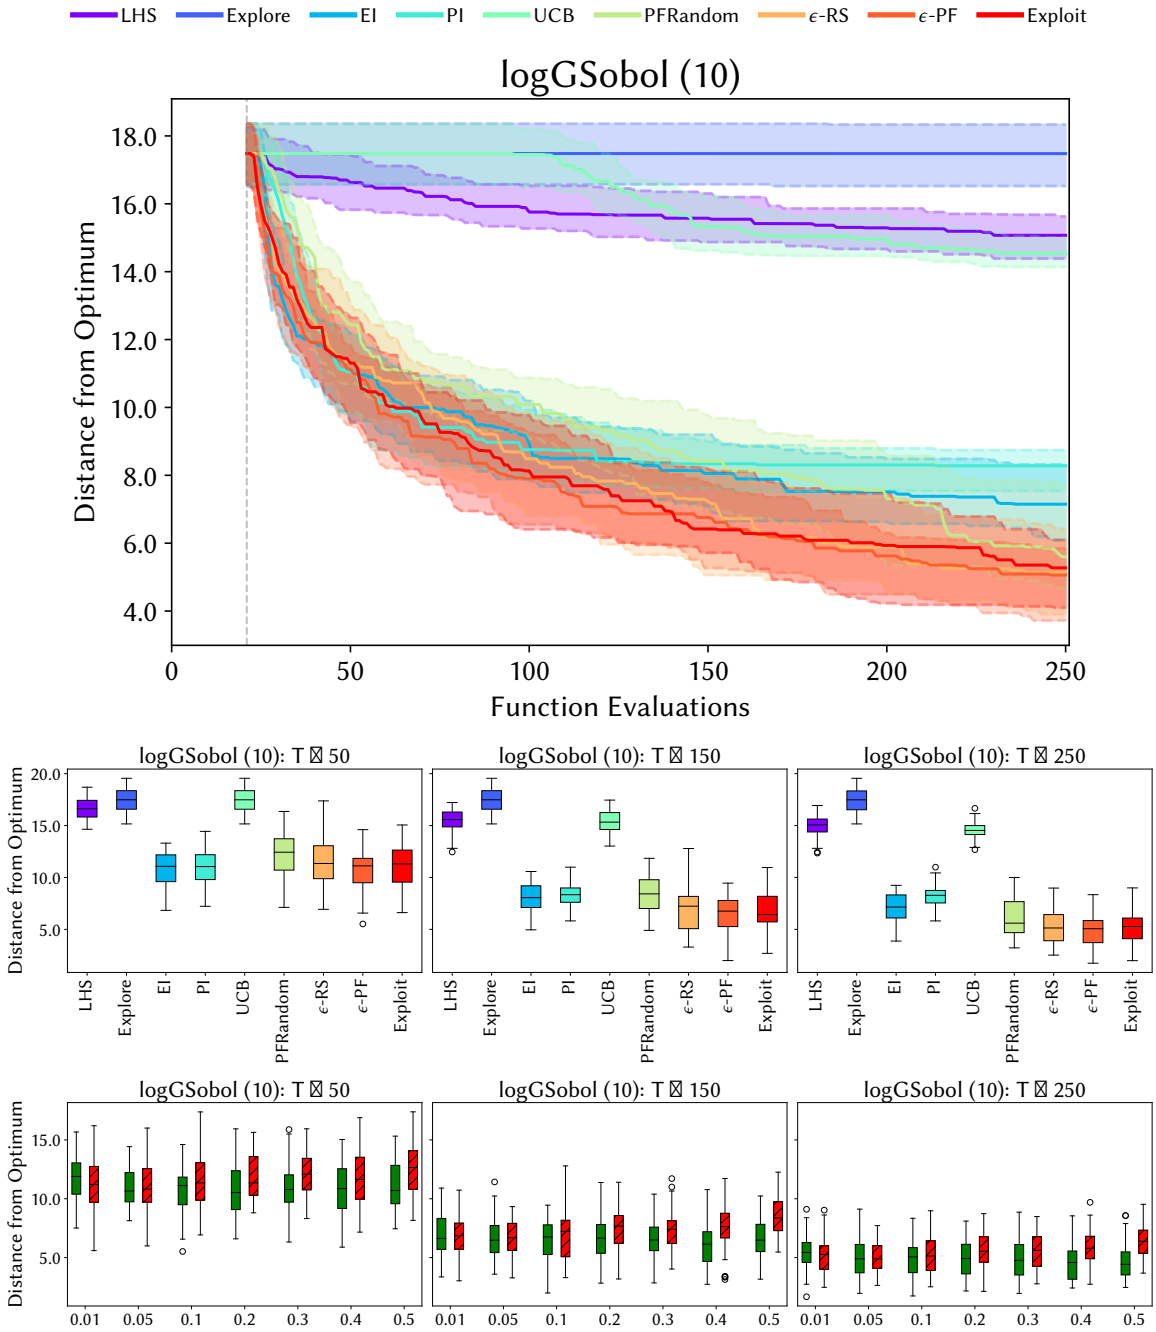


Fig. 9. Results for the ten-dimensional logGSobol test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

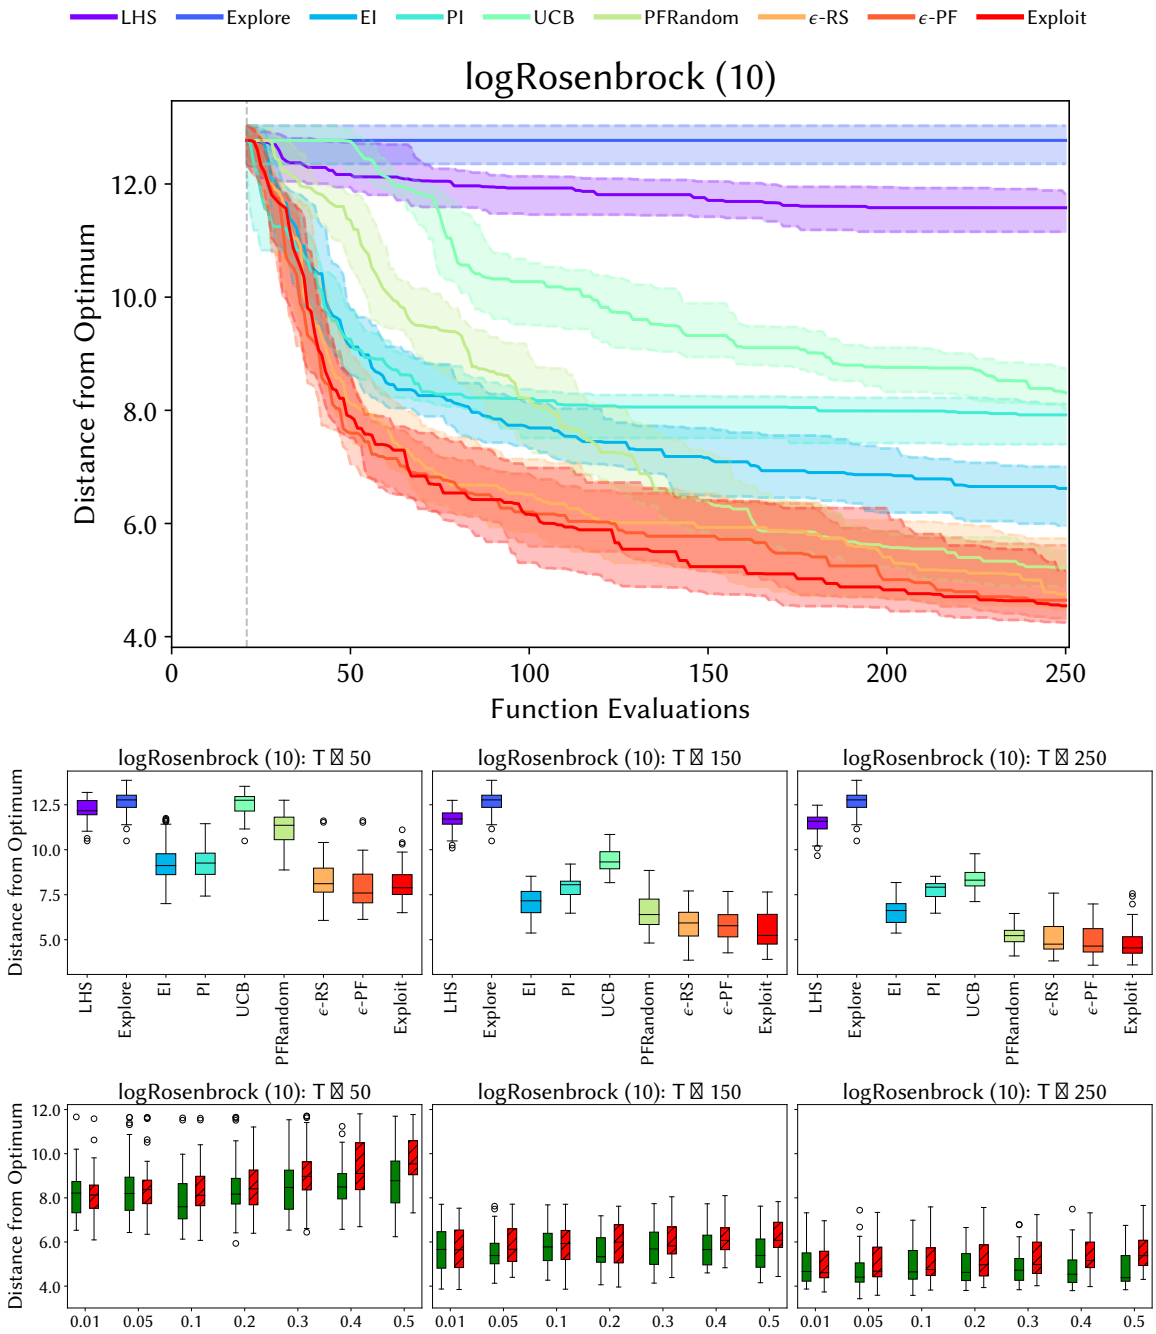


Fig. 10. Results for the ten-dimensional logRosenbrock test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

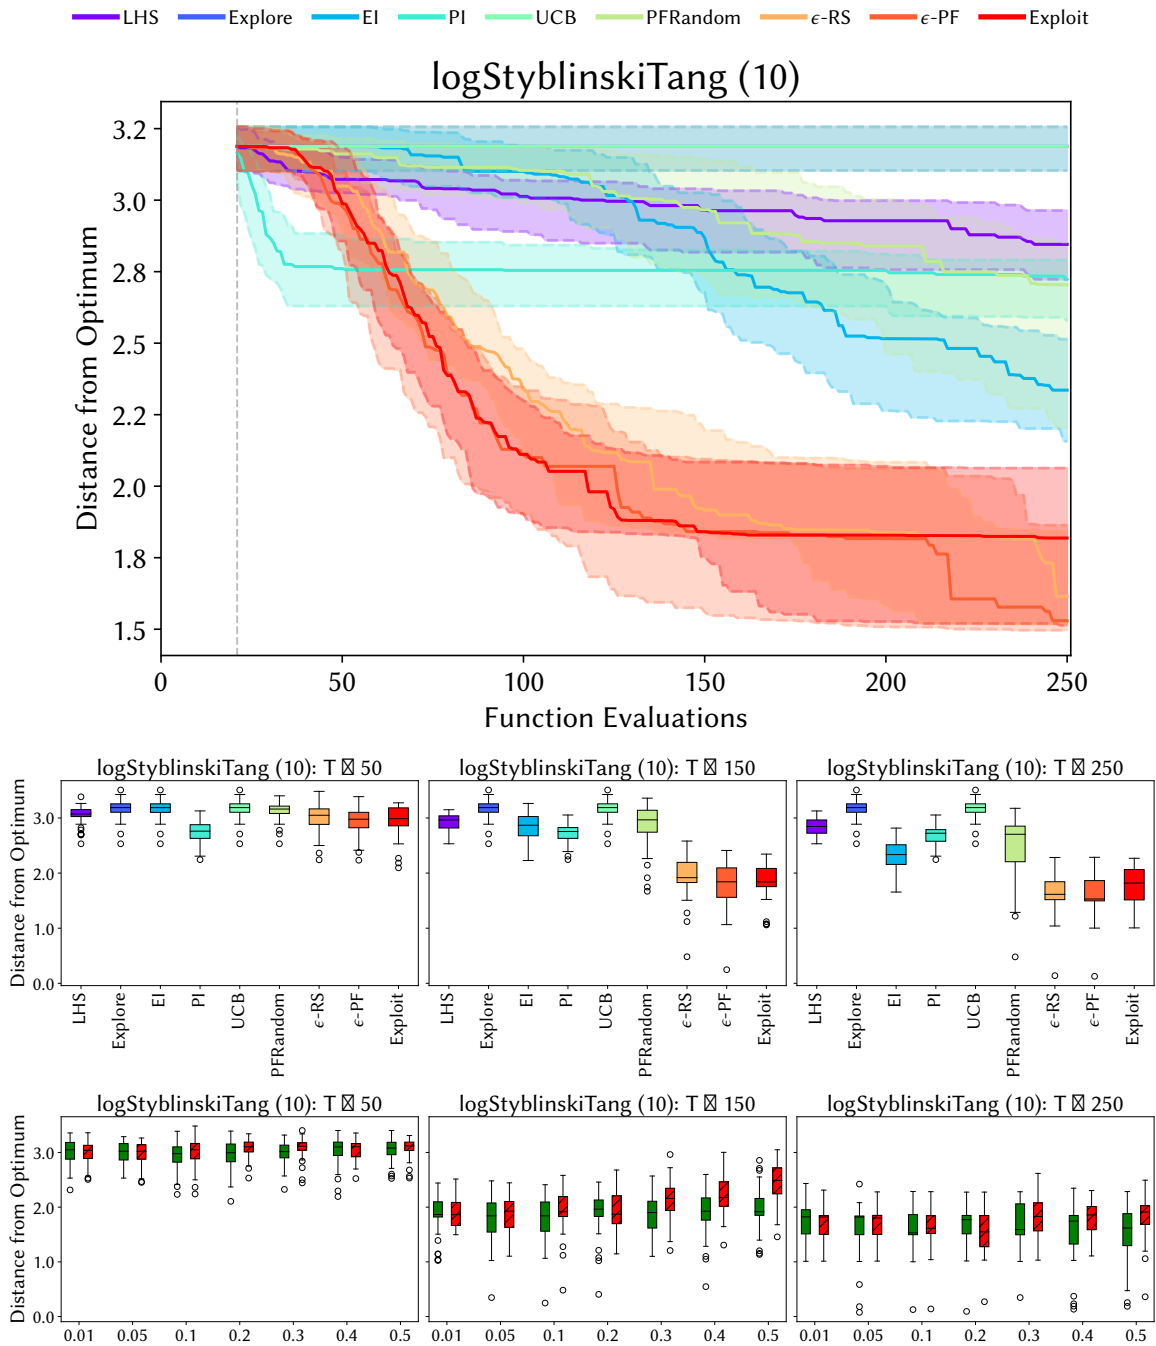


Fig. 11. Results for the ten-dimensional $\log\text{StyblinskiTang}$ test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between $\epsilon\text{-PF}$ (green) and $\epsilon\text{-RS}$ (red, hatched) for different values of ϵ (horizontal axis) after 50 (left), 150 (centre) and 250 (right) function evaluations.

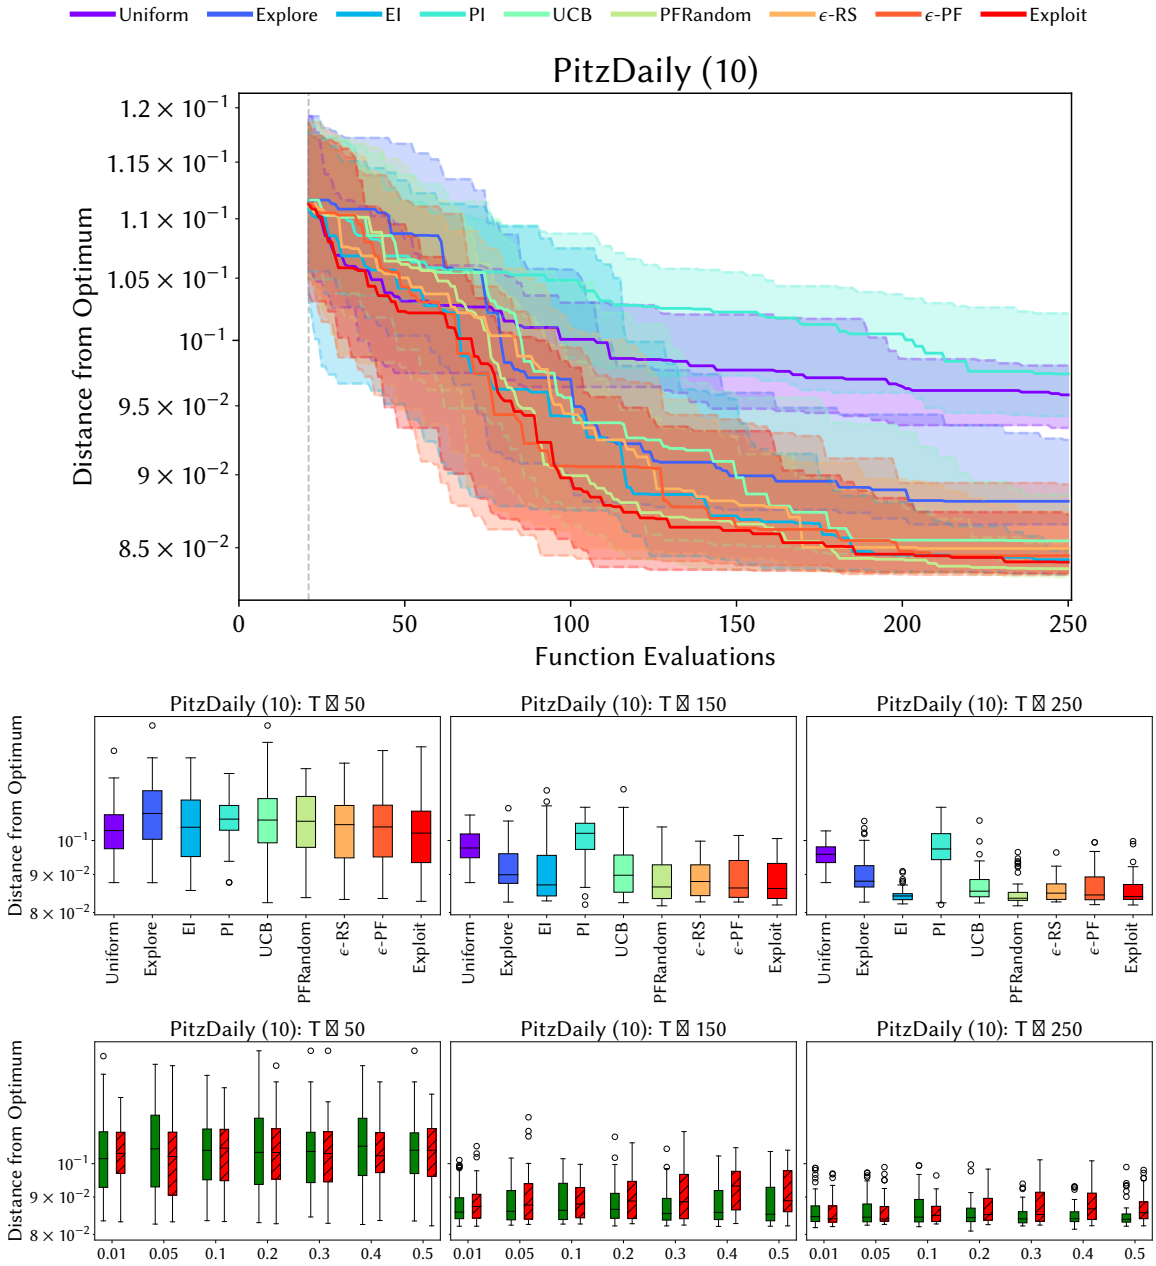


Fig. 12. Results for the ten-dimensional real-world PitzDaily test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

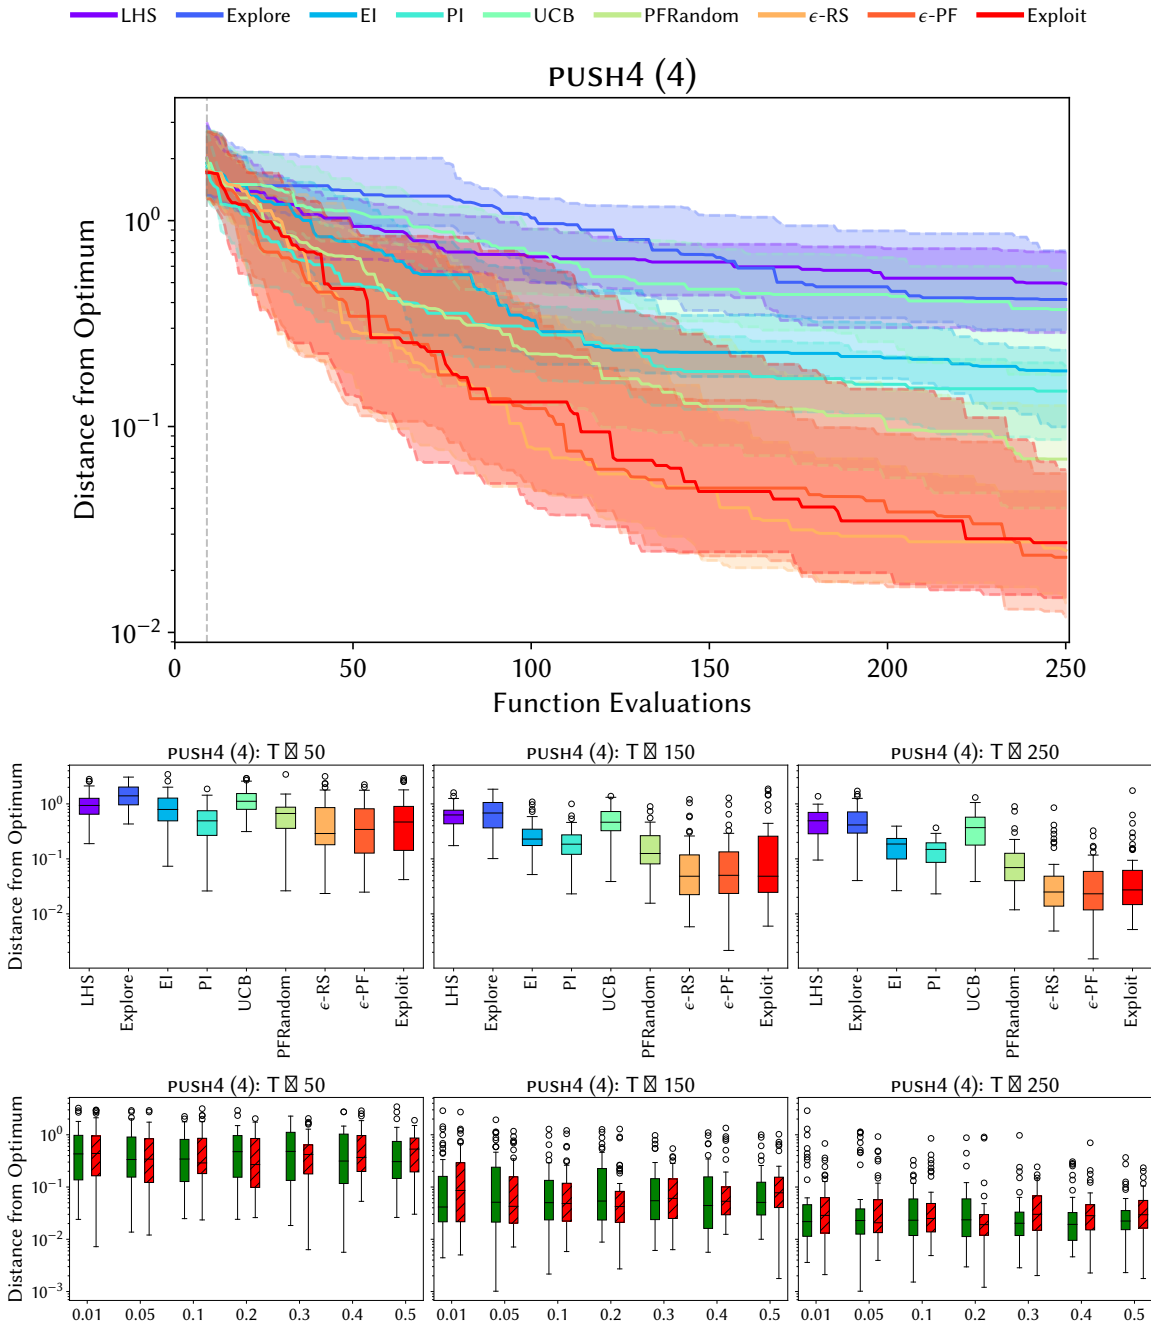


Fig. 13. Results for the four-dimensional real-world push4 test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

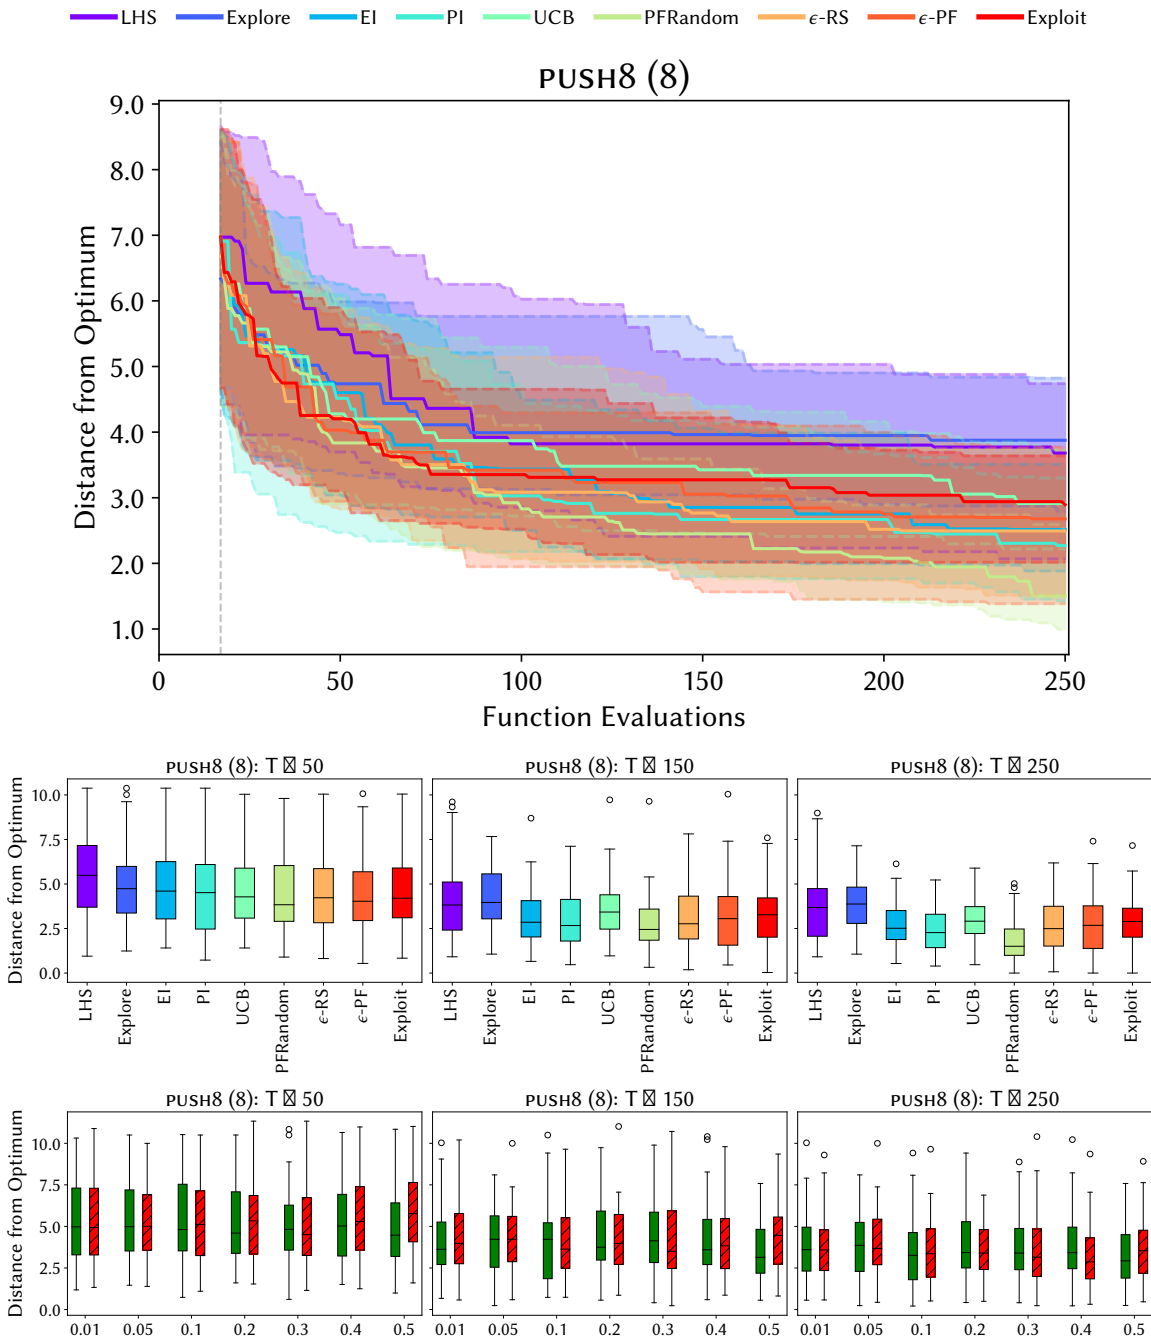


Fig. 14. Results for the eight-dimensional real-world push8 test problem. The convergence histories for each algorithm are shown in the upper figure, where the shaded regions correspond to the interquartile range. The central figure shows the distribution of best seen function evaluations after 50 (left), 150 (centre) and 250 (right) function evaluations have occurred. The lower figure shows a comparison between 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.

REFERENCES

- [1] Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70.
- [2] 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.
- [3] Ziyu Wang and Nando de Freitas. 2014. Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters. arXiv:arXiv:1406.7758