arXiv:2104.04790v1 [cs.LG] 10 Apr 2021

What Makes an Effective Scalarising Function for Multi-Objective Bayesian Optimisation?

Clym Stock-Williams[1], Tinkle Chugh[2], Alma Rahat[3], and Wei Yu[4]

[1]TNO, Westerduinweg 4, 1755LE Petten, The Netherlands
Clym.Stock-Williams@tno.nl
[2]Department of Computer Science, University of Exeter, UK
t.chugh@exeter.ac.uk
[3]Department of Computer Science, Swansea University, UK
a.a.m.rahat@swansea.ac.uk
[4]Faculty of Aerospace Engineering, Delft University of Technology, The Netherlands
w.yu@tudelft.nl

Abstract

Performing multi-objective Bayesian optimisation by scalarising the objectives avoids computation of expensive multi-dimensional integral-based acquisition functions, instead allowing one-dimensional standard acquisition functions—such as Expected Improvement—to be applied. Here, two infill criteria based on hypervolume improvement—one recently-introduced and one novel—are compared with the multi-surrogate Expected Hypervolume Improvement. The reasons for the disparities in these methods' effectiveness in maximising the hypervolume of the acquired Pareto Front are investigated. In addition, the effect of the surrogate model mean function on exploration and exploitation is examined: careful choice of data normalisation is shown to be preferable to the exploration parameter commonly used with the Expected Improvement acquisition function. Finally, the effectiveness of all the methodological improvements defined here are demonstrated on a real-world problem: the optimisation of a wind turbine blade aerofoil for both aerodynamic performance and structural stiffness. With effective scalarisation, Bayesian optimisation finds a large number of new aerofoil shapes which strongly dominate standard designs.

Keywords:hypervolume, Expected Improvement, Gaussian process, infill criteria, scalarisation, mono-surrogate, aerofoil, airfoil

1 Introduction

The class of metaheuristic optimisation algorithms have proven to be powerful at finding global optima in cases where the derivative of the optimisation objective cannot be calculated directly [6]. However, these algorithms usually require many thousands of evaluations to approximate an optimal solution, due to their reliance on stochastic search procedures. In many situations where optimal design is desired (for example geometry optimisation for a complex engineering component [10]) it is impractical to evaluate these many candidate designs, due to time and/or financial constraints.

To alleviate this, Bayesian optimisation (BO) algorithms have been developed, first popularised for single-objective problems by Jones *et al.* [20].

At its heart, BO learns a surrogate function of the objective space from evaluations (often with uncertainty) of candidate solutions. This surrogate is used to acquire new candidate solutions to evaluate, by explicitly maximising a trade-off between exploitation of existing high-performing areas, and exploration of previously un-explored areas [29].

The surrogate is commonly a Gaussian process, conventionally constructed with a zero mean prior. In multi-objective optimisation problems there are typically two approaches to surrogate modelling. The most common is to build surrogates for each objective function (referred to here as the multi-surrogate approach) [13]. Another technique (here described as the mono-surrogate approach) is to build one surrogate after scalarising the multiple objectives, e.g. [21, 27]. A recent survey comparing different scalarising functions (or 'infill criteria') can be found in [8].

Mono-surrogate approaches are usually computationally faster—but under-perform—multi-surrogate approaches. This can be attributed to the fact that the scalarising functions can hide important details of the multi-objective function landscape, and thus reduce the effectiveness of the surrogate model in some problems. Here we aim to develop further insights into how to construct an effective mono-surrogate infill criterion. To do this, we consider one of the most successful infill criterion and propose a variant. We then compare performance (also against the standard multi-surrogate approach) on standard multi-objective test problems and the real-world optimisation problem of wind turbine aerofoil design.

In addition to the scalarising function, two key elements in Gaussian process surrogate modelling are the choice of mean and covariance (or kernel) functions. The choice of covariance function is the subject of active research. For example, a kernel search meta-optimisation loop has recently been shown to improve robustness to different objective function surfaces [23]. However, the literature is thin on the effect of mean functions on optimisation effectiveness; and in particular, there are none for BO when solving multi-objective problems. In this paper (taking inspiration from [24]) we first focus on this topic, primarily from a mono-surrogate BO approach.

To summarise, the contributions of this article are as follows:

1. New insights into the use of mono-surrogate approaches in Bayesian multi-objective optimisation;

2. Examination of the effectiveness of the mean function in mono-surrogate (or scalarisation) based approaches;

3. Demonstration of performance on standard and real-world multi-objective optimisation problems.

The rest of this paper is organised by those key contributions, as follows:

- Section 2 provides the necessary background to both single-objective and multi-objective BO;
- The role of controlling the surrogate model mean function is investigated in Section 3;
- Section 4 establishes the properties of two related mono-surrogate infill criteria; before
- Section 5 demonstrates their performance on standard test problems. Important conclusions are then drawn on what makes a good infill criterion. Following this, the approaches are applied to a real-world problem—wind turbine aerofoil design—to demonstrate that the expected performance improvements are realised.
- Finally, section 6 summarises the conclusions and highlights further opportunities for research.

2 Bayesian Optimisation: Background and Investigation

2.1 Overview of the Method for Single-Objective Optimisation Problems

We now define the BO method, for a minimisation problem. The desired outcome is to find:

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}\in\mathcal{X}} f(\boldsymbol{x})$$

where $\mathcal{X}$ is the space of feasible solutions, a constrained subset of $\mathbb{R}^D$. The search is initialised by generating n_{init} candidate solutions in $\mathcal{X}$, often created by means of a Latin Hypercube [22]. The optimisation loop then begins, commencing with the evaluation of $y_i = f(\boldsymbol{x}_i)$, where $i \in [1, n_{init}]$.

The surrogate model is first fitted to the data: usually a Gaussian process with a Matérn-5/2 covariance kernel $\kappa(\boldsymbol{x}, \boldsymbol{\theta})$ [31] with hyperparameters $\boldsymbol{\theta}$. This surrogate can then be used to predict function values with uncertainty at any other point $\boldsymbol{x_*}$ in the decision space:

$$p(y_*|\boldsymbol{x}_*, X, \boldsymbol{y}, \boldsymbol{\theta}) = \mathcal{N}\big(\mu(\boldsymbol{x}_*), \sigma^2(\boldsymbol{x}_*)\big), \tag{1}$$

where we can compute the mean and the variance of the Normal posterior predictive distribution using the following formulae:

$$\begin{aligned} \mu(\boldsymbol{x}_*) &= \kappa(\boldsymbol{x}_*, X)\kappa(X, X)^{-1}\boldsymbol{y}, \\ \sigma^2(\boldsymbol{x}_*) &= \kappa(\boldsymbol{x}_*, \boldsymbol{x}_*) - \kappa(\boldsymbol{x}_*, X)^T\kappa(X, X)^{-1}\kappa(X, \boldsymbol{x}_*). \end{aligned} \tag{2}$$

The next candidate solution is now selected through optimisation of an acquisition function applied to the surrogate's predictions. This acquisition function calculates the usefulness of any new point in objective space, based on the current information. One of the most common, effective and intuitive acquisition functions is the Expected Improvement (EI) over the current best-found function value y^* [20]:

$$\begin{aligned} I(\boldsymbol{x}) &= max(y^* - \hat{y}, 0) \\ \mathbb{E}[I(\boldsymbol{x})] &= \int_{-\infty}^{\infty} max(y^* - \hat{y}, 0)F(y)dy, \end{aligned}$$

where $\hat{y}$ is the best found fitness value. If we assume the Gaussian process surrogate from equation 1, and substitute $y = \mu(\boldsymbol{x}) + \sigma(\boldsymbol{x})\epsilon$ (where ϵ is the standard Normal distribution, with PDF ϕ and CDF Φ), this gives:

$$\mathbb{E}[I(\boldsymbol{x})] = (y^* - \mu)\,\Phi\left(\frac{y^* - \mu}{\sigma}\right) + \sigma\phi\left(\frac{y^* - \mu}{\sigma}\right) \tag{3}$$

Optimisation of the acquisition function is then performed (using another optimisation algorithm, such as CMA-ES or L-BFGS with multiple restarts [29]) to select one or more locations to evaluate next. The optimisation loop then continues until the allowed budget of evaluations is exhausted.

2.2 Multi-objective Bayesian Optimisation

The goal in multi-objective optimisation is to find a set of mutually non-dominated solutions $\mathcal{P}$, as close as possible to the global Pareto Front [14]. The area dominated by this set, relative to a reference point R, is called the hypervolume (HV).

Maximising this hypervolume is normally performed with the same fundamental process as that described in section 2.1. A new acquisition function is defined, which targets minimising the loss

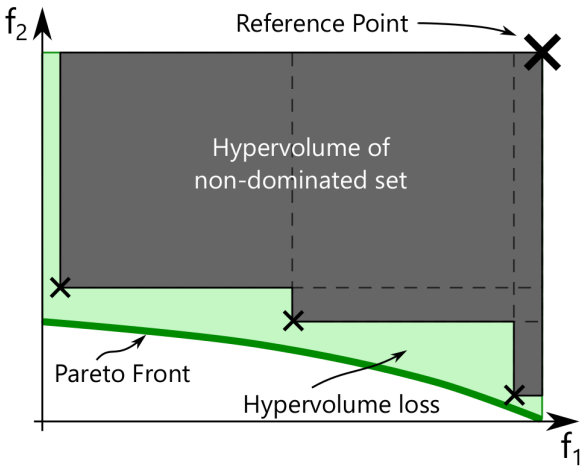

Figure 1: Schematic showing the hypervolume loss between a non-dominated set of solutions found by a multi-objective optimisation algorithm and the true Pareto Front.

between the global Pareto Front and the found non-dominated set, see Fig. 1. The direct extension of EI to multiple objectives is Expected Hypervolume Improvement (EHVI). Recent efforts have largely focused on making the hypervolume and expected hypervolume improvement calculations faster, although usually at the expense of assuming zero uncertainty on the evaluated solutions [17, 4, 9, 37, 36].

2.3 Motivation for Mono-Surrogate Infill Criteria

An alternative approach to defining a new acquisition function is to use a scalarising function (or 'infill criterion'). The reasons to choose this approach are:

1. the existing acquisition functions and software for single-objective optimisation can be used;
2. the expensive process of fitting multiple surrogate models can be avoided;
3. the very large number of expensive integrations required for EHVI and similar acquisition functions can be avoided.

The simplest version is to calculate a weighted mean of the objective function values for each solution, and then apply EI. This, however, has the unfortunate side-effect of directing the optimisation down a single vector in the objective space.

More sophisticated approaches have been developed in the BO literature [27] and it is now known that Pareto-compliance is a pre-requisite for effectiveness [40]. Nevertheless, this leaves open a large design space for such functions, and no principled approach to increase effectiveness is apparent in the previous literature proposing new scalarisation functions. Further, analyses of their effectiveness have been largely limited to empirical studies and plots of their theoretical properties in objective space.

3 Effect of Mean Function in Balancing Exploration and Exploitation

To the knowledge of the authors, the balance of exploitation versus exploration has previously focused entirely on the acquisition function [1, 39], neglecting the important role the prior mean takes in determining the predictive mean and variance.

Despite being poorly-reviewed in the literature, a 'jitter' parameter ζ is commonly included in equation (3) [3]:

$$\mathbb{E}[I(\boldsymbol{x})] = (y^* + \zeta - \mu)\,\Phi\left(\frac{y^* + \zeta - \mu}{\sigma}\right) + \sigma\phi\left(\frac{y^* + \zeta - \mu}{\sigma}\right) \tag{4}$$

The intention is clearly to enhance exploration by making the current best function evaluation slightly worse. The default value used in the Python package GPyOpt, for example, is $\zeta = 0.01$ [33].

We now compare the effects of changing this ζ parameter with varying the prior mean relative to the current best. As shown in Fig. 2, comparatively very large values of ζ are required to significantly change the exploration behaviour of the acquisition function. On the other hand, the relationship between the surrogate mean function value and the shape of the EI surface is intuitive and better-behaved. As the mean improves, the value already found becomes less surprising, thus exploration is preferred. Conversely, worsening the mean increases the surprise of the value found, making exploitation preferable. Note further the poor shape of the EI function surface as ζ becomes negative in the lower right plot of Fig. 2.

Returning to equations (1), (2), (3) and (4), the equivalence between ζ and a mean function increase of value λ can quickly be derived, for the situation in which only one solution has been acquired:

$$y^* + \zeta - \kappa(\boldsymbol{x}_*, \boldsymbol{x})\,\kappa(\boldsymbol{x}, \boldsymbol{x})^{-1}\,y^* = (y^* + \lambda) - \kappa(\boldsymbol{x}_*, \boldsymbol{x})\,\kappa(\boldsymbol{x}, \boldsymbol{x})^{-1}\,(y^* + \lambda)$$

$$\zeta = \lambda\left(1 - \kappa(\boldsymbol{x}_*, \boldsymbol{x})\,\kappa(\boldsymbol{x}, \boldsymbol{x})^{-1}\right) \tag{5}$$

The jitter method therefore does not apply a constant offset to the prior mean, but instead a less-intuitive offset that changes with distance in the decision space from the existing solutions.

Instead of altering the acquisition function, it is therefore recommended to focus on carefully choosing the prior mean (or, equivalently, normalising the measured function values for a zero-mean prior). The following approach is therefore suggested (given the usual situation where the global minimum function value of an optimisation objective is unknown). Each time a new (batch of) candidate solutions are desired, the archive of solution values is renormalised using equation 6 before fitting the zero-mean Gaussian process surrogate.

$$y\prime = (1 - \xi) \cdot \overline{y} + \xi \cdot y^* \tag{6}$$

Setting the exploration parameter $\xi = 0.0$ implies using the current mean $\overline{y}$ as the prior mean. Positive values promote more exploratory behaviour, while negative values promote exploitation. In the following sections, this approach will be applied.

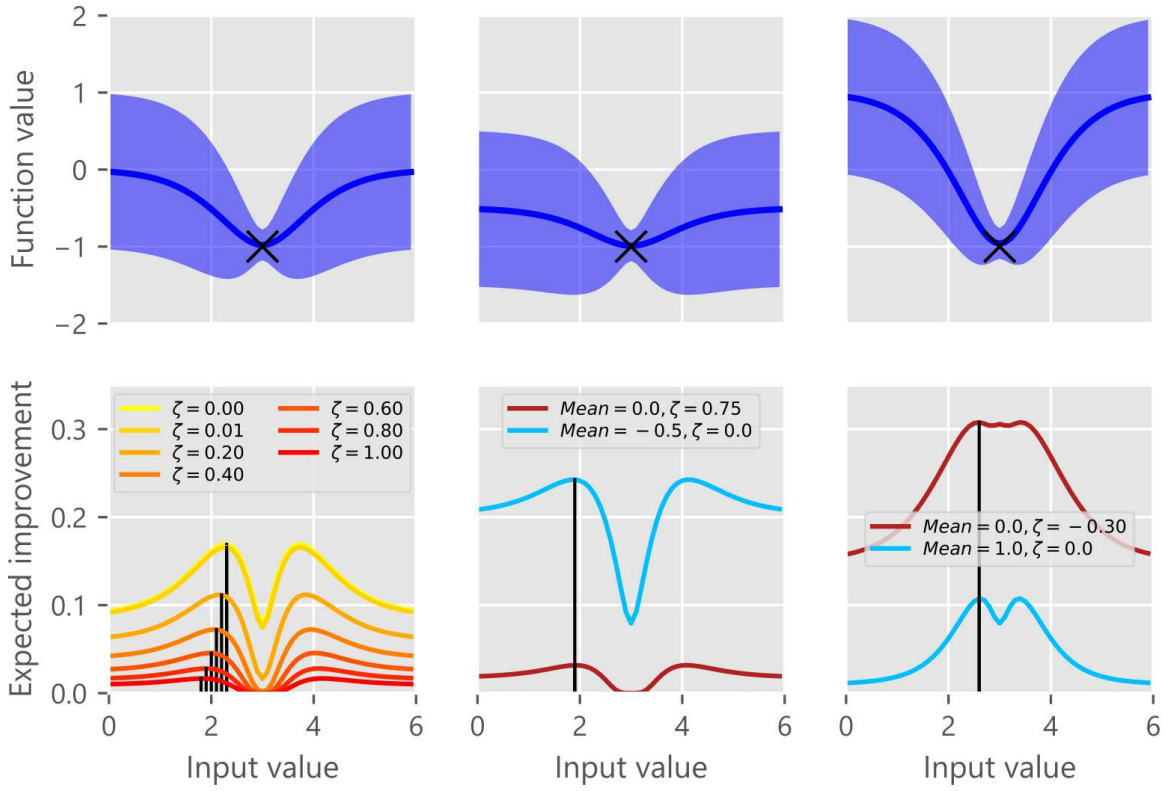

Figure 2: Comparison of the effect of the mean function and jitter parameter ζ on the exploration of the Expected Improvement (EI) acquisition function. Top row: Gaussian process surrogate with varying mean and constant kernel variance and lengthscale, provided evaluated solutions marked with a black 'x'. Bottom row: EI calculations. Left: when the prior mean is 1 standard deviation above the evaluated solution, the change with ζ of the optimum EI location is shown. Centre: when the mean is closer to the evaluated solution, more exploration is performed. Right: when the mean is worse, more exploitation is performed. For both centre and right lower plots, the red line shows the ζ parameter value required to obtain the same EI optimum when using the zero mean function value from the left plot.

4 Extending Hypervolume Improvement

4.1 Defining Mono-Surrogate Infill Criteria

We now define a novel infill criterion which fulfills the criteria from section 2.3. Instead of estimating the hypervolume improvement for predicted objective function values, we shall instead follow the approach in [26] and calculate the hypervolume improvement contributed by each existing solution evaluated with expensive objective functions. We store all the evaluated solutions in an Archive ($\mathcal{A}$).

An illustration for a bi-objective minimisation problem is shown in Fig. 3. For a non-dominated solution—see the left-hand schematic in Fig. 3—this shall be a positive number:

$$\mathrm{HVI}^{+}(\boldsymbol{x}, R) = \mathrm{HV}(\mathcal{A} \setminus \{\boldsymbol{x}\}, R) - \mathrm{HV}(\mathcal{A}, R). \tag{7}$$

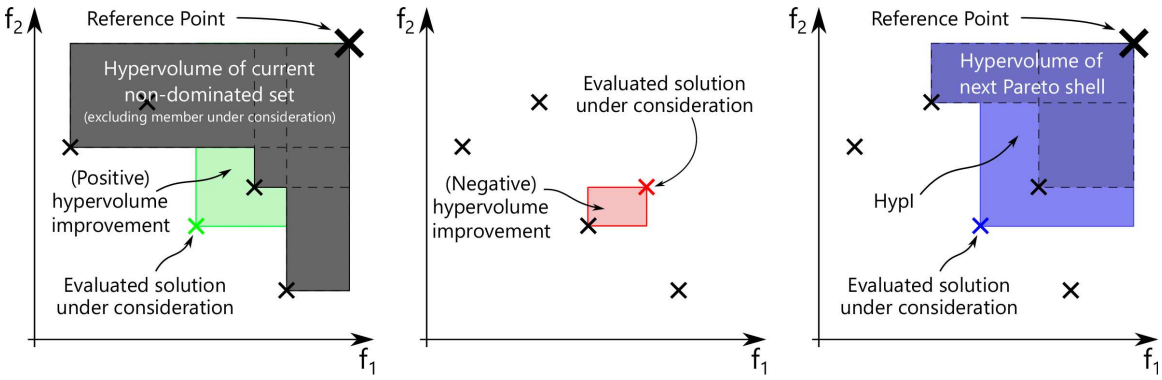

Figure 3: Schematic showing the calculation of Extended Hypervolume Improvement (xHVI) and HypI. Left: positive hypervolume improvement for the non-dominated solution marked in green; centre: negative hypervolume improvement for the dominated solution marked in red; right: HypI for the solution marked in blue.

Since we want to maximise the total hypervolume of the non-dominated set, this metric gives an intuitive measure of the contribution of each non-dominated solution, as well as penalising solutions which are close together in the objective space.

However, a zero value is assigned to dominated solutions, whereas we want a smoothly varying gradient over the decision space. We therefore estimate how poor dominated solutions are by calculating a 'negative hypervolume improvement', as shown in the centre schematic of Fig. 3:

$$\mathrm{HVI}^{-}(\boldsymbol{x}) = \mathrm{HV}(\mathcal{P}, \boldsymbol{x}) \tag{8}$$

This naturally assigns non-dominated solutions a zero value. Thus the infill criterion, Extended Hypervolume Improvement (xHVI), is defined as follows:

$$\mathrm{xHVI}(\boldsymbol{x}; R) = \mathrm{HVI}^{+}(\boldsymbol{x}; R) - \mathrm{HVI}^{-}(\boldsymbol{x}) \tag{9}$$

We shall compare the performance of this infill criterion with the recently-proposed HypI [27], shown in the right-hand schematic of Fig. 3. The definition of this requires that we assign each solution to a Pareto shell $\mathcal{P}_k$, where the second Pareto shell $\mathcal{P}_2$ is the set of mutually non-dominated solutions once all members of $\mathcal{P}_1 = \mathcal{P}$ are removed from the archive, and so on. Each solution is assigned the following value:

$$\mathrm{HypI}(\boldsymbol{x}, R) = \mathrm{HV}(\mathcal{P}_{k+1} \cup \{\boldsymbol{x}\}, R) \tag{10}$$

Conceptually, this is rather similar to xHVI. A solution's value increases with the amount of hypervolume it generates. The main difference apparent at this stage is that it does not penalise solutions by their proximity in objective space (note that this is performed by the acquisition function).

4.2 Illustration on bi-objective optimisation problem

The optimisation process outlined in section 2.1 is now elaborated in Fig. 4 for both multi- and mono-surrogate BO, illustrating similarities and differences using the ZDT3 test function [41] with $D = 2$ input dimensions and $M = 2$ objectives.

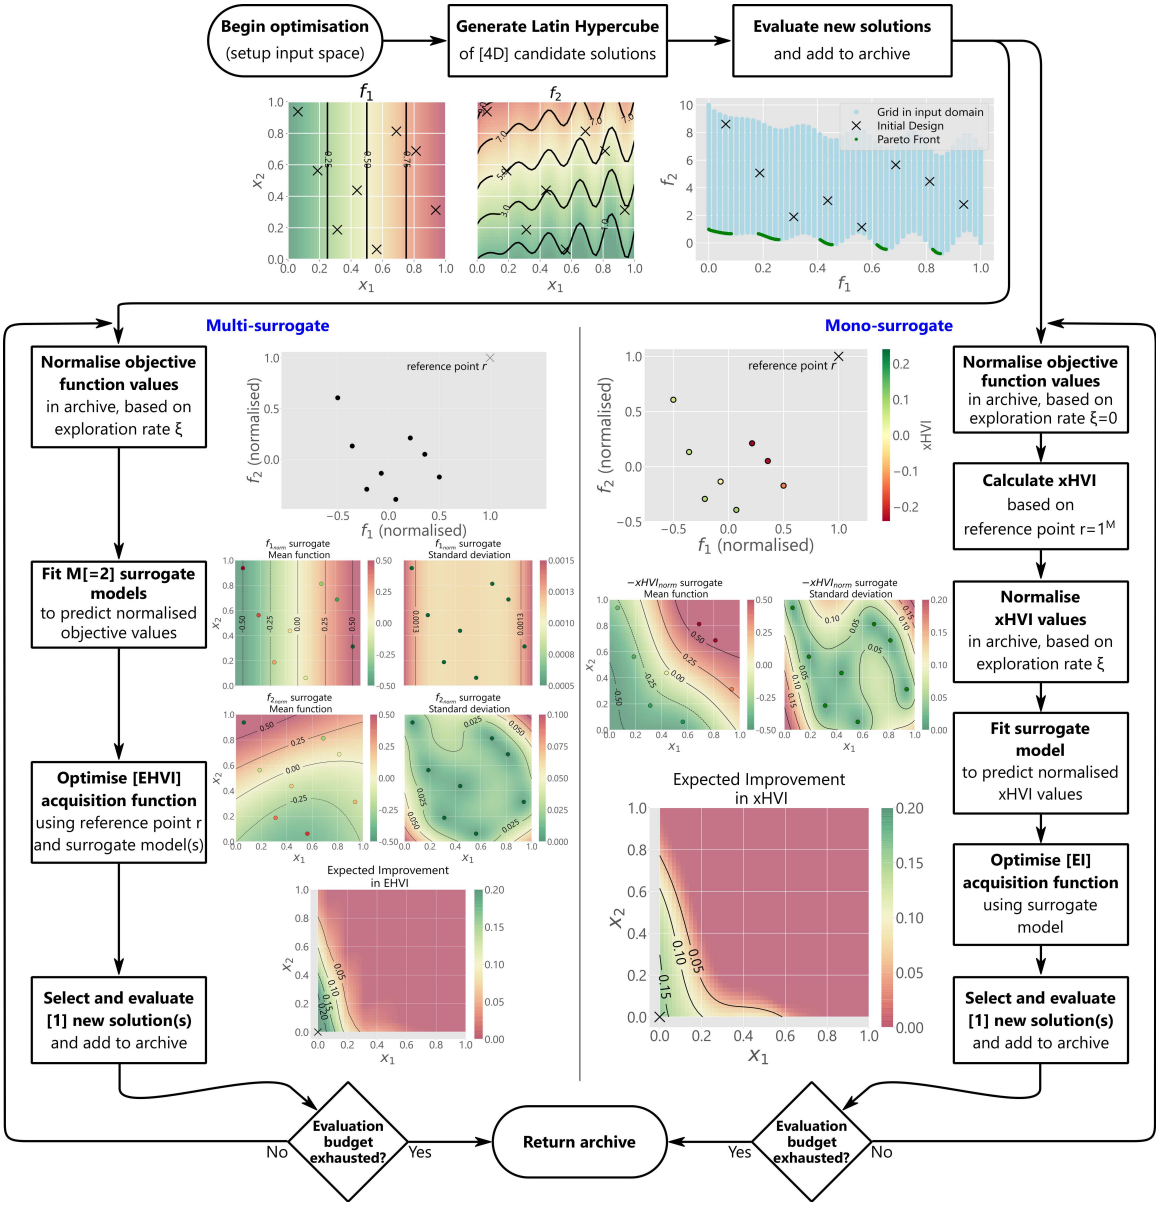

Figure 4: Process diagram comparing multi-objective Bayesian optimisation with (right hand side) and without (left hand side) the use of a mono-objective infill criterion. The illustrative examples shown use the ZDT3 test function with 2 input dimensions and 2 objectives. Values in square brackets refer to the choices made in section 5.1.

Initially, as shown on the top row of Fig. 4, a space-filling design is created. The exact balance of benefit in the proportion of the evaluation budget used up here and in the optimisation will depend on the properties of the function and the optimisation method chosen. With EHVI, for example, it is important to obtain a good initial surrogate fit to the objective functions. A Latin Hypercube of

size $4D = 8$ is used here, immediately finding two solutions reasonably close to the Pareto Front. A grid of 2500 solutions in the decision space is provided (in light blue) for reference.

With the standard BO process described in section 2.1, we proceed down the left hand side of the process diagram. The objective function values are first normalised as described in section 3, using $\xi = 0.0$ (i.e. the mean of each set of evaluated objective values is zero). Two independent surrogate models are fitted, one for each objective. As can be seen (referring back to the plots at the top of the figure for the true objective surfaces) the independence of f_2 with x_2 is captured, and the general shape of f_2 is also realised.

Finally, the Expected Hypervolume Improvement is calculated—the surface is shown in the lowest plot—using the reference point $r = \mathbb{1}^M$ and 1000 Monte Carlo samples. The acquisition function is optimised with L-BFGS, with empirical gradients, conducting 10 re-starts from initial locations which are chosen randomly close to the existing non-dominated set. The next solution to acquire is $[0.0, 0.0]$, which is immediately on the true Pareto Front.

Now, we turn our attention to the right hand side of Fig. 4, where the xHVI infill criterion is used. The first step is to calculate the xHVI for each evaluated solution. The top plot illustrates how the objective values are first normalised as for the multi-surrogate process. Then equations 7-9 are applied, based on a reference point $r = \mathbb{1}^M$. This reference point can also be fixed from knowledge of the lowest useful objective value.

Once this is achieved, only one surrogate is fitted. As can be seen in the middle plots, the shape of the surrogate fits well the shape of the non-dominated set. Optimisation of Expected Improvement (again with L-BFGS, with empirical gradients) can proceed as usual for single-objective BO. The next solution to acquire is in agreement with EHVI: $[0.0, 0.0]$.

5 Performance Analysis on Multi-Objective Problems

5.1 Performance on Standard Multi-objective Test Functions

The methodologies described in section 4.2 for EHVI-, xHVI- and HypI-based BO are now applied to functions from the well-known ZDT [41] and DTLZ [12] suites, as implemented in the DEAP Python toolbox [18]. $D = 10$ input dimensions and $M = 2$ objectives are used for all functions.

The evaluation budget is restricted to 300, which is reasonable for the expensive problems encountered in practice. A Latin Hypercube design of 40 candidates is provided, using the PyDOE Python toolbox. Thus 260 additional evaluations are available for BO, which is run sequentially.

Two free parameters have been identified for the optimisation processes:

1. The exploration parameter ξ from section 3; and
2. The value of the normalised reference point r from section 4.2.

It is likely that these hyperparameters can be tuned for every problem, however, one choice will now be made for all functions. First, ξ is set to 0.0. As discussed in 4.2, this corresponds to the commonly-assumed case where the evaluated objective function values are normalised to their mean. Second, the reference point r is set to $\mathbb{1}^M$.

Each optimisation is run 21 times, each iteration using a different initial random seed for the Latin Hypercube, in order to provide statistics on the outcomes of the optimisation. The same random seeds are used to initialise the xHVI, HypI and EHVI runs, to ensure a fair test of the optimisation algorithms' performance.

Many performance metrics have been proposed in literature [2], which analyse the convergence to, and spread along, the true Pareto Front achieved by the final non-dominated set, and the efficiency of the algorithm. We used the following performance metrics in this work:

- **Hypervolume (HV) or S-metric** - presented here as a percentage of the achievable hypervolume (see also Fig. 1). The reference point required for its calculation is chosen to be 10% greater than the true Pareto Front's nadir point.
- **Non-dominated set size** - this indicates the number of meaningful choices presented to the user of the optimisation.
- **Empirical attainment function (EAF)** - proposed by [16], this is the empirical cumulative distribution over the probability of a point in objective space being dominated by the obtained non-dominated set.

Since the evaluation budget is very limited, little attention is given to efficiency. Comparison of runtime on one problem for xHVI and EHVI is, however, provided for interest in Fig. 5.

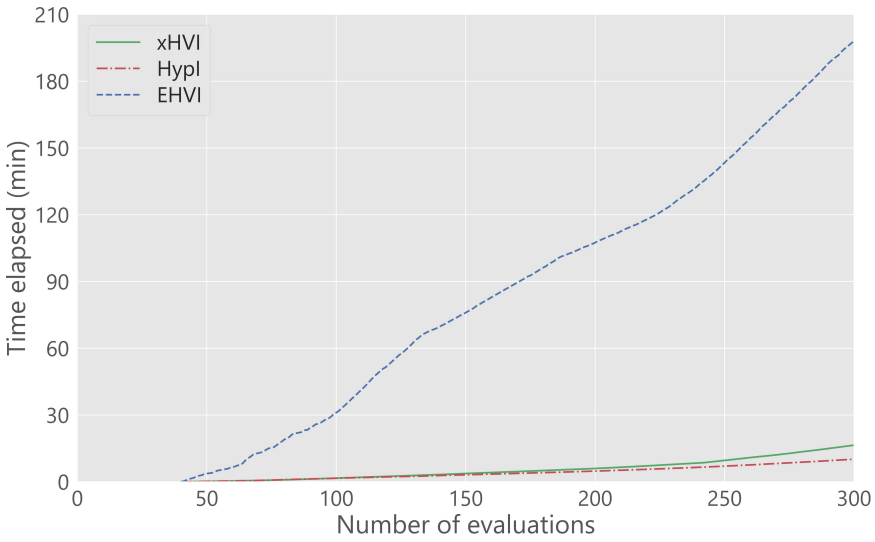

Figure 5: Comparison of computation times required for xHVI, HypI and EHVI Bayesian optimisation methods (using ZDT6) on a 4-core i7-6700HQ CPU.

5.2 Analysis

Beginning with Fig. 5, it is apparent that the runtime of EHVI—due the need to fit twice as many surrogate models, and conduct extensive Monte Carlo calculations of HVI to optimise the acquisition function—is more than 2.5 times that of xHVI or HypI. With an expensive evaluation function and batch optimisation this difference will become less significant.

The EAFs of HypI and xHVI are shown in Fig. 6 for ZDT and Fig. 7 for DTLZ problems. We also show the differences in EAF of HypI and xHVI in comparison to EHVI. Next to the EAFs, we present bivariate kernel density plots of (%) hypervolume and non-dominated set size. This information is now discussed and analysed.

It is clear from these figures that EHVI and HypI demonstrate impressive performance in reliably converging to the true Pareto Front on ZDT1, ZDT2, ZDT3, DTLZ2 and DTLZ7. Conversely, none of the algorithms are able to solve ZDT4, DTLZ1 or DTLZ3 adequately.

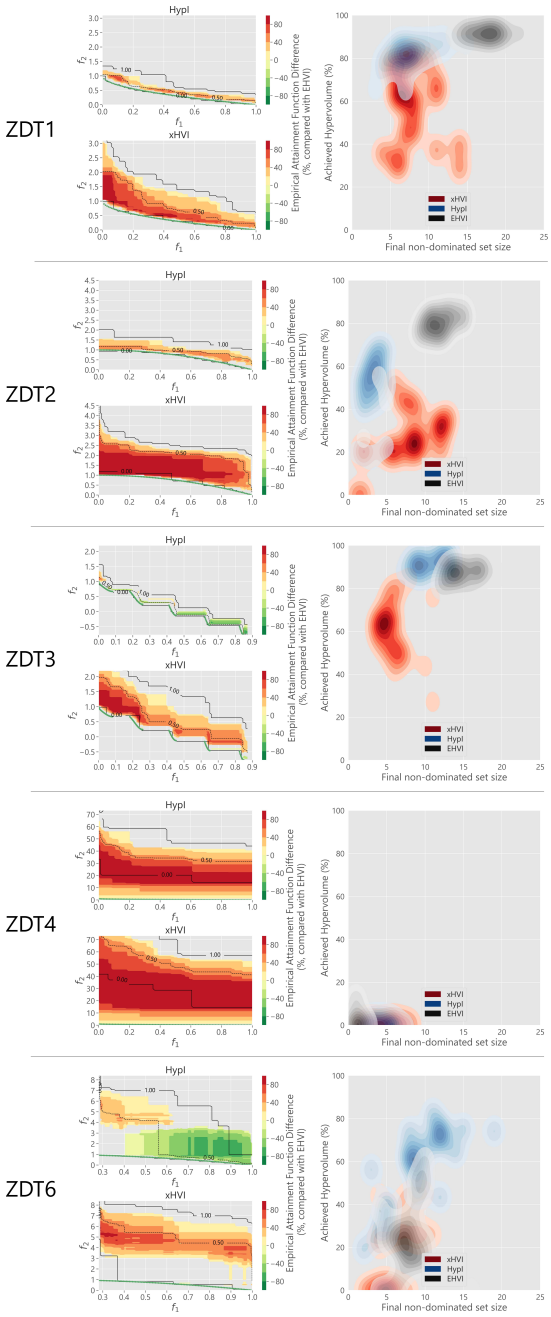

Figure 6: Comparison of 21 optimisations with 300 evaluations using the EHVI, xHVI and HypI Bayesian optimisation approaches, on five ZDT test functions each with 10 dimensions and 2 objectives. Left: Empirical attainment functions (EAFs) for the final non-dominated sets (black lines), the difference with the EAF of EHVI (colour), and the true Pareto Front (green). Right: Bivariate kernel density estimates of the % hypervolume achieved and non-dominated set size, for all three approaches.

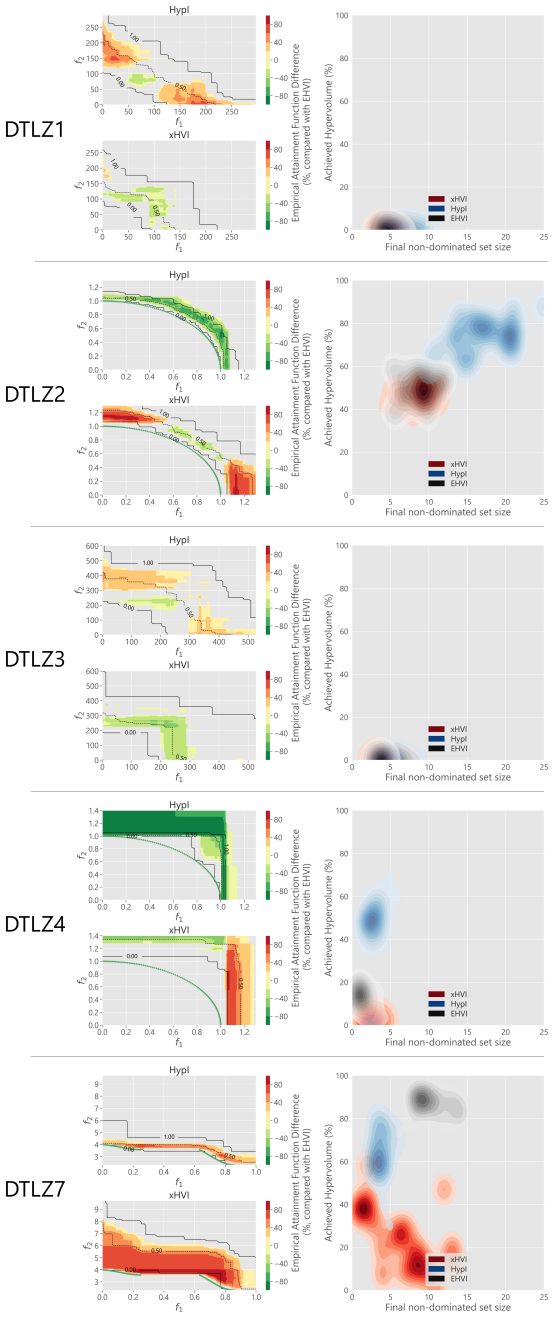

Figure 7: Comparison of 21 optimisations with 300 evaluations using the EHVI, xHVI and HypI Bayesian optimisation approaches, on five DTLZ test functions each with 10 dimensions and 2 objectives. Left: Empirical attainment functions (EAFs) for the final non-dominated sets (black lines), the difference with the EAF of EHVI (colour), and the true Pareto Front (green). Right: Bivariate kernel density estimates of the % hypervolume achieved and non-dominated set size, for all three approaches.

The colours in the left-hand objective space plots (and the vertical axis of the right-hand graphs) give an impression of the relative performance of EHVI and the mono-objective methods. To quantify these differences, it is also informative to (double-)integrate over the EAF differences to obtain a single numerical estimate of the relative performance of three methods; the results are shown in Table 1.

Test function	Integrated EAF difference HypI-EHVI	Integrated EAF difference xHVI-EHVI
ZDT1	-0.09	-0.50
ZDT2	-0.19	-1.2
ZDT3	0.04	-0.42
ZDT4	-25	-41
ZDT6	0.50	-1.2
DTLZ1	-5.5×10^3	2.0×10^3
DTLZ2	0.11	-0.13
DTLZ3	-1.7×10^4	1.5×10^4
DTLZ4	0.62	0.02
DTLZ7	-0.27	-1.6

Table 1: Numerical analysis of results shown in Figs. 6 and 7. The double integral over the difference in empirical attainment function (EAF) is calculated, comparing each scalarising function (HypI and xHVI) with the multi-surrogate EHVI. Negative numbers indicate a worse performance than EHVI overall; positive numbers the opposite.

The performance of an optimisation algorithm can itself be viewed as a multi-objective decision making problem. Turning to the kernel density estimate plots, two metrics for assessing performance are presented, in a manner such that dominance of one optimisation algorithm over another can easily be seen.

Maximising HV is of course the main measure of optimisation effectiveness. This metric has theoretically superior behaviour to others with which it shares a high correlation, such as inverted generational distance [35]. Any optimisation methods producing equal HV, however, can be further distinguished by the non-dominated set size. This metric—despite not distinguishing solutions which are close together in objective space—does provide a useful indicator of the amount of information available to the user of the optimisation results.

Comparing the performance of the algorithms themselves, xHVI is clearly a much less capable scalarising function than HypI. Determining the reasons for this requires us to return to the example from Fig. 4, and compare their surrogate models and acquisition surfaces, which are shown for two states during the optimisation process in Fig. 8.

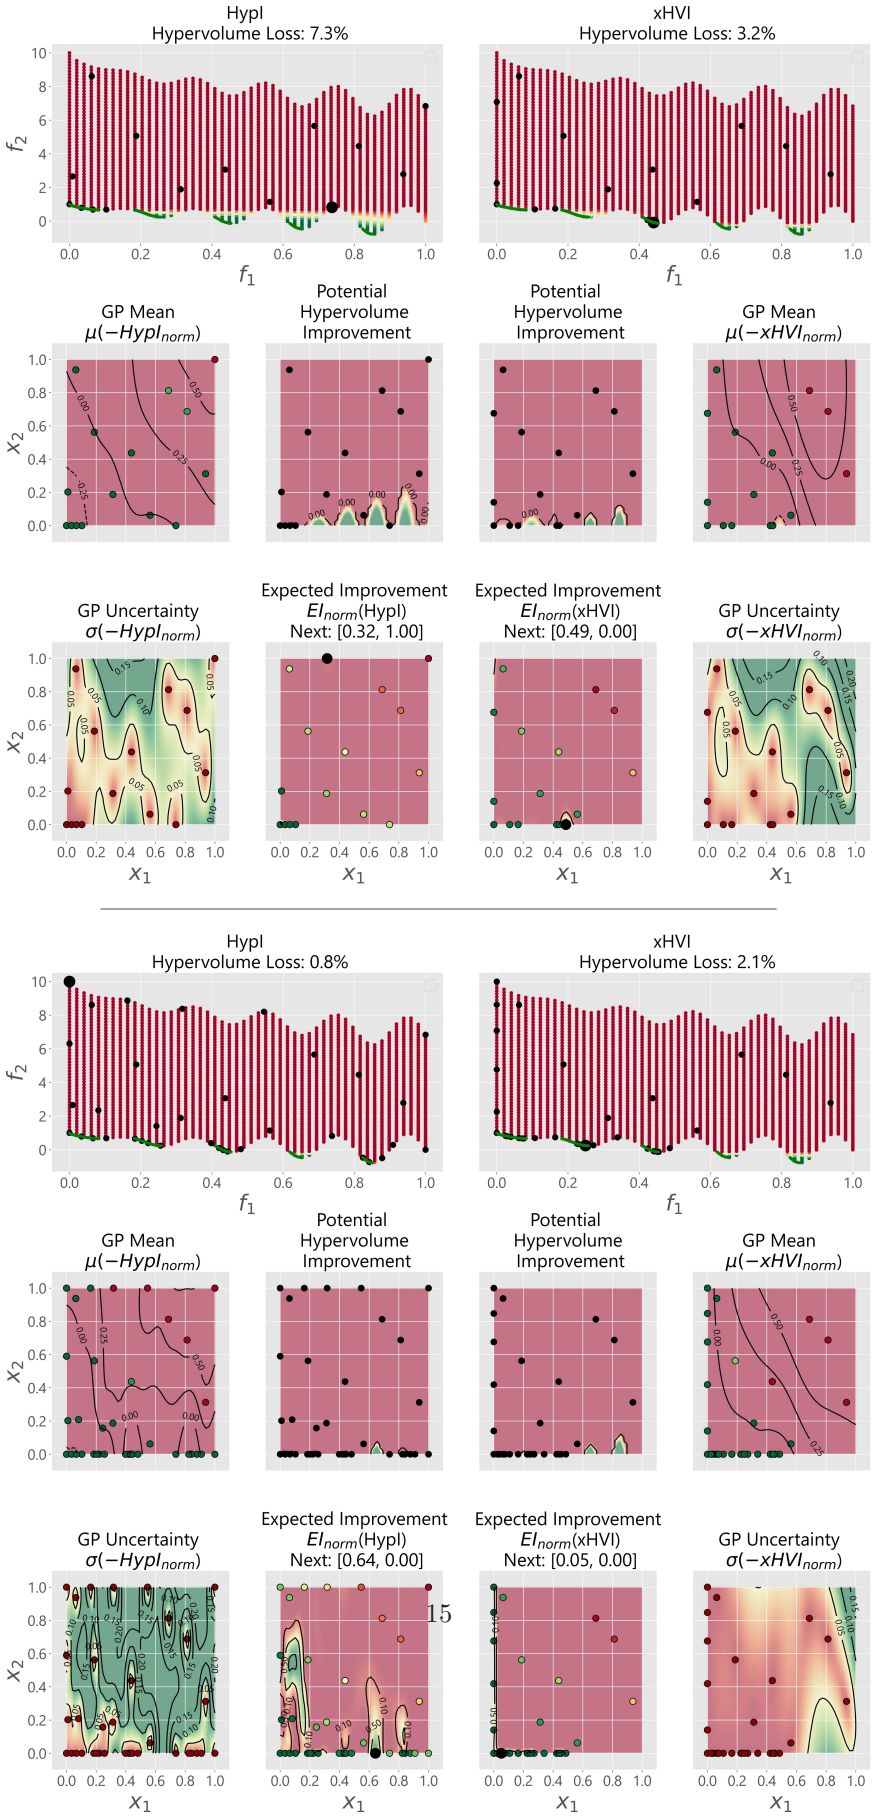

Figure 8: Illustration of optimisation for the xHVI and HypI infill criteria, using the ZDT3 test function with 2 input dimensions and 2 objectives. Top: after 15 acquired solutions; bottom: after 36 acquired solutions.

Each graph shows the currently acquired points, initially in the objective space, and then in decision space. "Potential Hypervolume Improvement" indicates the best possible locations for the optimisation algorithms to sample next. By comparison, the "Expected Improvement" graph beneath indicates the surface which is being optimised over to select the next point for acquisition (shown as a large dot). The mean and standard deviation of the Gaussian process complete the picture. They are coloured so that the two components of Equation 3 are shown equally, and green indicates that a particular location is better for acquisition.

The key driver behind the different behaviour of the two scalarising functions is their gradients. The gradient of HypI is much more consistent over the whole surface, whereas xHVI assigns much larger values of HVI^- to dominated solutions than the HVI^+ assigned to non-dominated solutions.

In the early stages, this works to allow xHVI to obtain a higher hypervolume. The two sections of the Pareto Front with lowest f_1 are found and exploited. However, the very large negative values assigned by xHVI to the locations with higher f_1 conspire to force xHVI never to explore that area, despite the high uncertainty. As a result, within another 20 acquisition steps HypI has obtained a much higher hypervolume, driven by the smaller GP lengthscales in x_1, which allow the periodic variation in Pareto Front in that dimension to be found.

5.3 Application to Wind Turbine Aerofoil Design

Optimisation of aerodynamic shapes is a common topic in the literature [30]. Wind turbine aerofoil design under uncertainty [7] and considering multiple objectives [15, 11] are also recent topics, using evolutionary computation. There has also been some attention to single-objective Bayesian optimisation of aerofoil shapes [19]. However, to the knowledge of the authors, there have been no Bayesian optimisation studies of aerofoils considering multiple objectives, particularly for wind turbines.

5.3.1 Aerofoil Geometry Definition

Several parameterisations of aerofoil shape have been used in the design [32, 25], some of which have serious limitations for optimisation, since constraints on their decision spaces can be difficult to set *a priori* such that realistic aerofoils are produced.

Bézier curves offer a highly controllable—while also simple—way of defining shapes, which can be made flexible to the limitations the designer wishes to impose. Referring to Fig. 9, we choose here to optimise an aerofoil with a fixed chord c and thickness $t = 0.18c$ (this is a realistic constraint given that an aerofoil is a 2D cross-section through a blade with continuously changing shape). Points A (the leading edge) and C (the trailing edge) are fixed, leaving the 13-dimensional decision space given in Table 2.

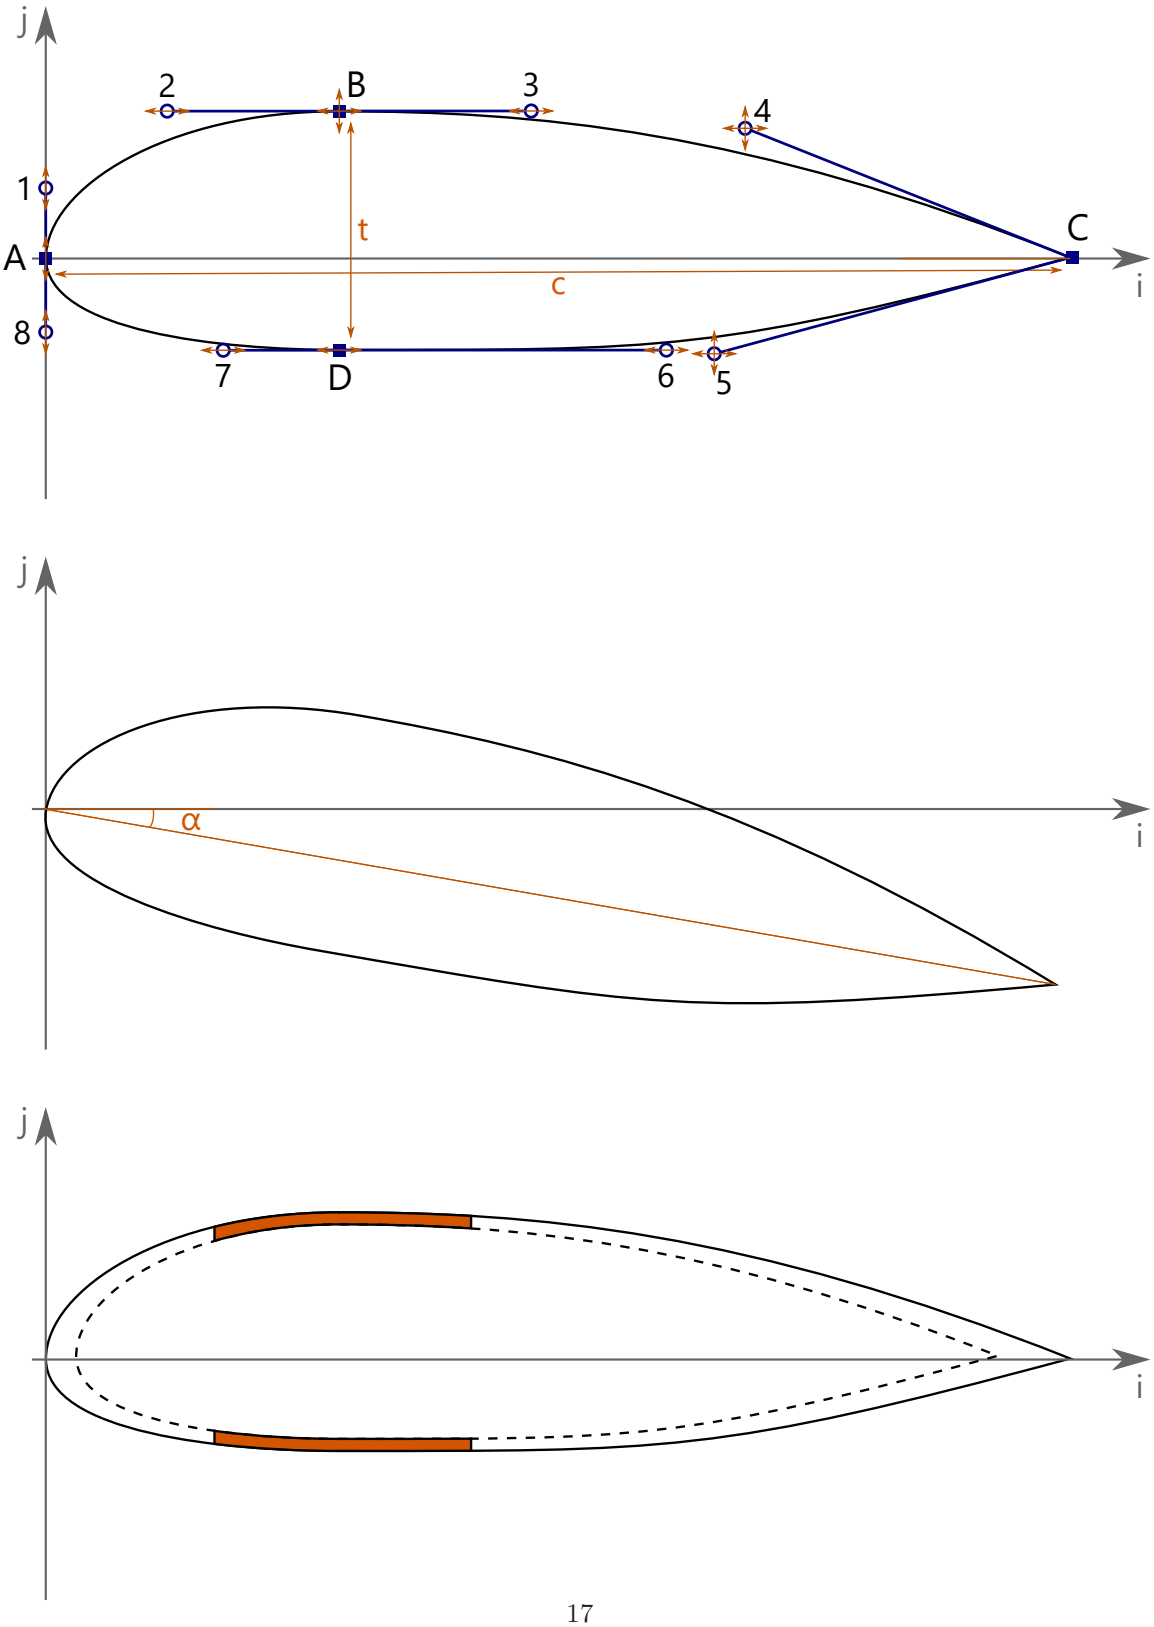

Figure 9: Top: Definition of a wind turbine blade aerofoil, using Bézier curves. The black curve is the aerofoil shape; the blue control points define the Bézier curves; and the orange arrows indicate optimisation parameters. Middle: Definition of the angle of attack α for aerodynamic evaluation. Bottom: Construction of a spar cap for evaluation of structural stiffness.

Decision variable	Description	Bounds
j_1	the vertical position of control point 1, relative to A and as a fraction of $j_B - j_A$	$[0.2, 0.9]$
j_B	the vertical position of B, relative to A and as a fraction of t	$[0.50, 0.95]$
i_B	the horizontal position of B, relative to A and as a fraction of c	$[0.15, 0.70]$
i_2	the horizontal position of control point 2, relative to B and as a fraction of $i_B - i_A$	$[0.1, 0.9]$
i_3	the horizontal position of control point 3, relative to B and as a fraction of $i_C - i_B$	$[0.2, 0.8]$
i_4	the horizontal position of control point 4, relative to C and as a fraction of $i_C - i_3$	$[0.05, 0.75]$
j_4	the vertical position of control point 4, relative to C and as a fraction of $j_B - j_C$	$[0.05, 0.95]$
i_5	the horijontal position of control point 5, relative to C and as a fraction of $i_4 - i_6$	$[0.05, 0.75]$
j_5	the vertical position of control point 5, relative to C and as a fraction of $j_4 - j_D$	$[0.05, 0.50]$
i_D	the horizontal position of D, relative to A and as a fraction of c	$[0.15, 0.70]$
i_6	the horizontal position of control point 6, relative to B and as a fraction of $i_C - i_D$	$[0.2, 0.8]$
i_7	the horizontal position of control point 7, relative to B and as a fraction of $i_D - i_A$	$[0.1, 0.9]$
j_8	the vertical position of control point 8, relative to A and as a fraction of $j_A - j_D$	$[0.2, 0.9]$

Table 2: Decision space for optimising the aerofoil shown in Fig. 9

5.3.2 Optimisation Objectives

Aerodynamic performance A wind turbine blade, of which the aerofoil is a cross-section, should maximise its extraction of energy from the wind. It is therefore logical to maximize also the aerofoil's local contribution to the power coefficient, which is equivalent to maximising the lift:drag ratio $\frac{c_l}{c_d}$ [5].

Stability of performance with variable angles of attack (the relative direction of the wind to the aerofoil axis, as shown in Fig. 9) is also important. Thus the aerodynamic performance objective used here is to maximize the integral of $\frac{c_l}{c_d}$ over a range of positive angles of attack α:

$$f_\alpha = \int_{\alpha=0}^{\alpha=10} \frac{c_l}{c_d} d\alpha \tag{11}$$

The aerofoils are here evaluated using Rfoil [28] (with Mach number 0, Reynolds number 6×10^6 and $N_{crit} = 9$).

Structural stiffness To create a survivable wind turbine blade, it is important that the aerofoil can withstand high loads, while keeping manufacturing costs low. In order to evaluate the structural performance realistically—while avoiding excessive computational cost—a simple spar cap is

assumed, as illustrated in Fig. 9. The spar cap width is chosen as $0.25c$, with thickness calculated by scaling the aerofoil shape by 90% and locating the centre of the spar cap at the same chord position as the upper crest (point B).

Structural performance is evaluated by calculating the bending stiffness of the spar cap in the flapwise direction, as follows:

$$f_s = \iint_{A_s} E(i,j) j^2 didj \tag{12}$$

where A_s is the area of the spar cap, and the elastic modulus E is assumed to be constant at 1 GPa.

5.3.3 Results and Comparison with Standard Aerofoils

Aerofoil shape optimisation was performed using the three BO approaches described earlier, given a budget of 800 evaluations, each repeated 5 times with different random seeds to test consistency. For comparison, six NACA and three DU aerofoils with the same thickness of 18% were also evaluated [34] (although DU08-W-180 dominates all other aerofoils). Results are shown in Fig. 10.

The most aerodynamically-efficient aerofoil found by HypI is shown in Fig. 11 and compared with DU08-W-180. The decision-maker will be interested in understanding the trade-off between the two objectives, thus Fig. 12 is included to present another aerofoil from the HypI optimisations, with almost identical aerodynamic performance to DU08-W-180, but 3.4 times the stiffness.

The main conclusions are:

- The new parameterisation proposed here allows for *a priori* constraints to be set on the desired aerofoil shape, largely avoiding unphysical designs.
- Standard aerofoil designs appear to have much less focus on structural stiffness than on aerodynamic performance.
- All three Bayesian Optimisation methods far outperform the standard aerofoils on these objectives.
- The multi-surrogate EHVI obtains stiffer aerofoils than either mono-surrogate, however, HypI obtains greater aerodynamic performance. (Not shown: EHVI's results also come at the expense of far longer runtimes, as previously illustrated in 5).
- Of the two scalarising functions, HypI outperforms xHVI. Compared with EHVI, HypI does exhibit some variability in outcome with initial conditions. This effect should reduce with a higher evaluation budget.

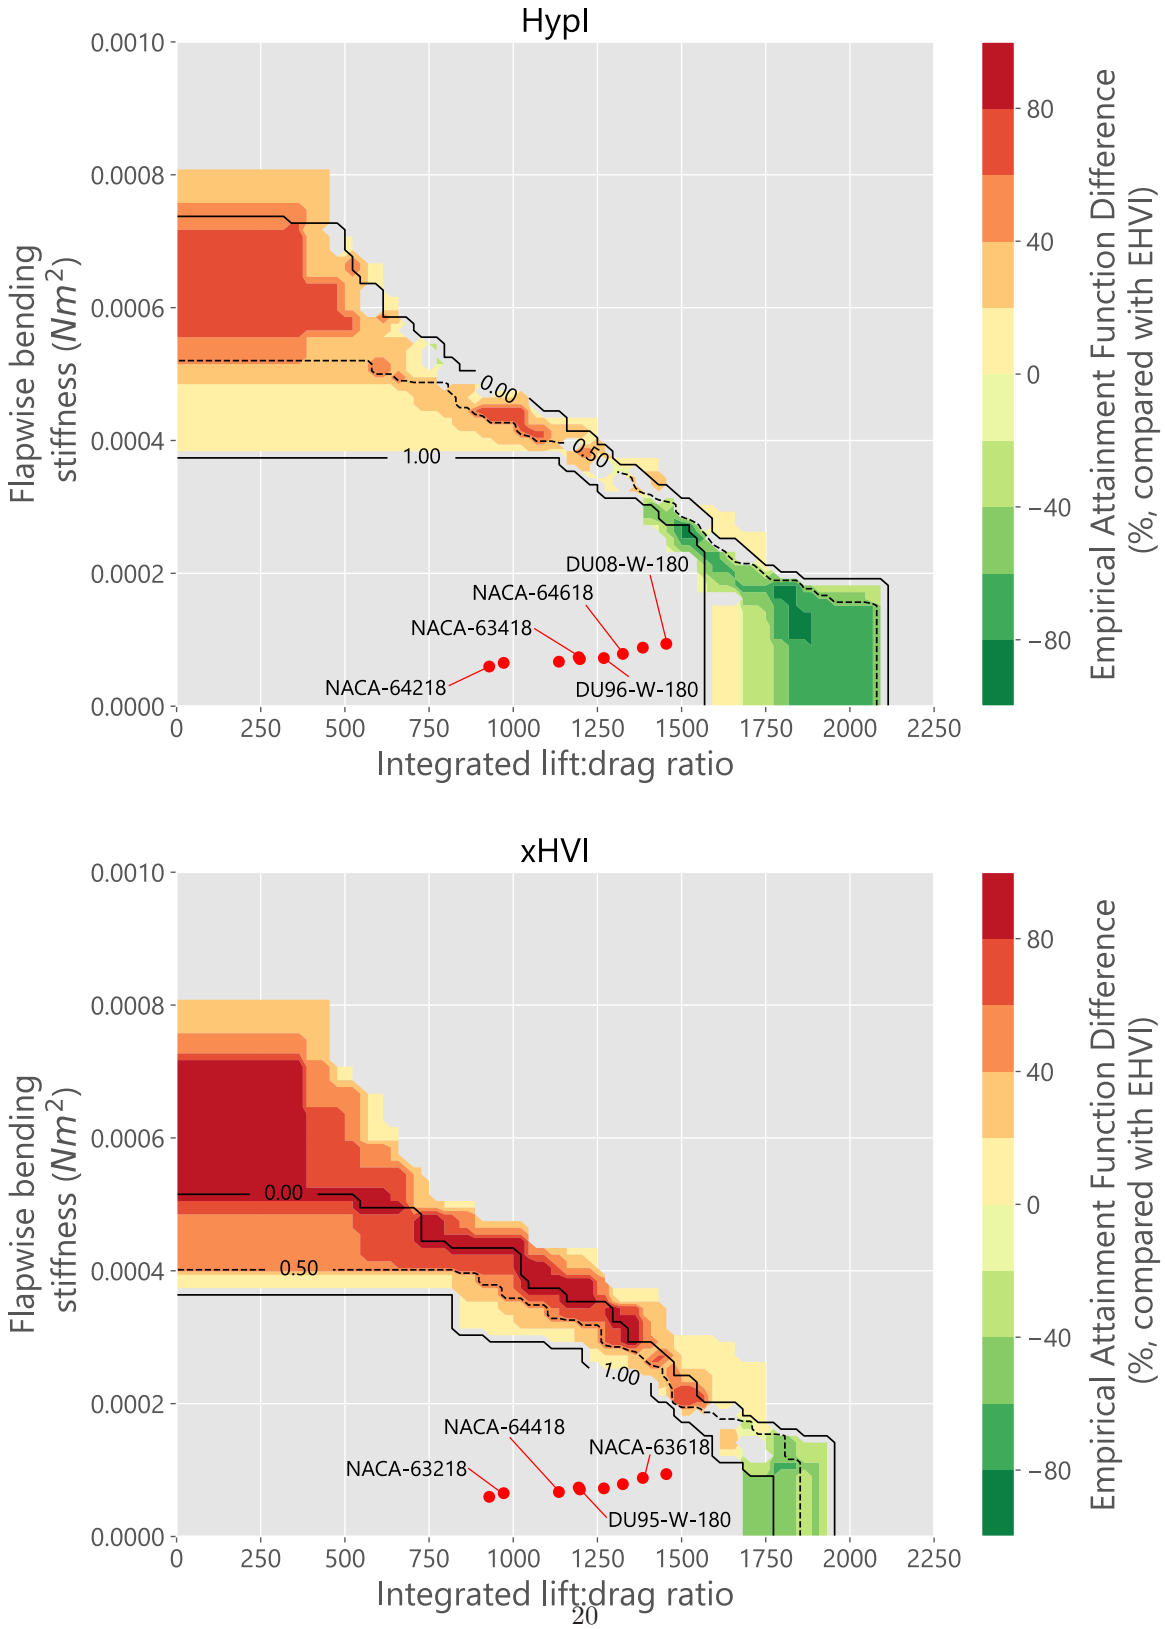

Figure 10: Comparison of results from Bayesian Optimisation of aerofoil shape after 800 evaluations, repeated 5 times using different random seeds. The black lines show the Empirical attainment functions for HypI (top) and xHVI (bottom); the colours indicate difference with EHVI. Performance results for standard aerofoils with the same thickness are also provided with the red dots.

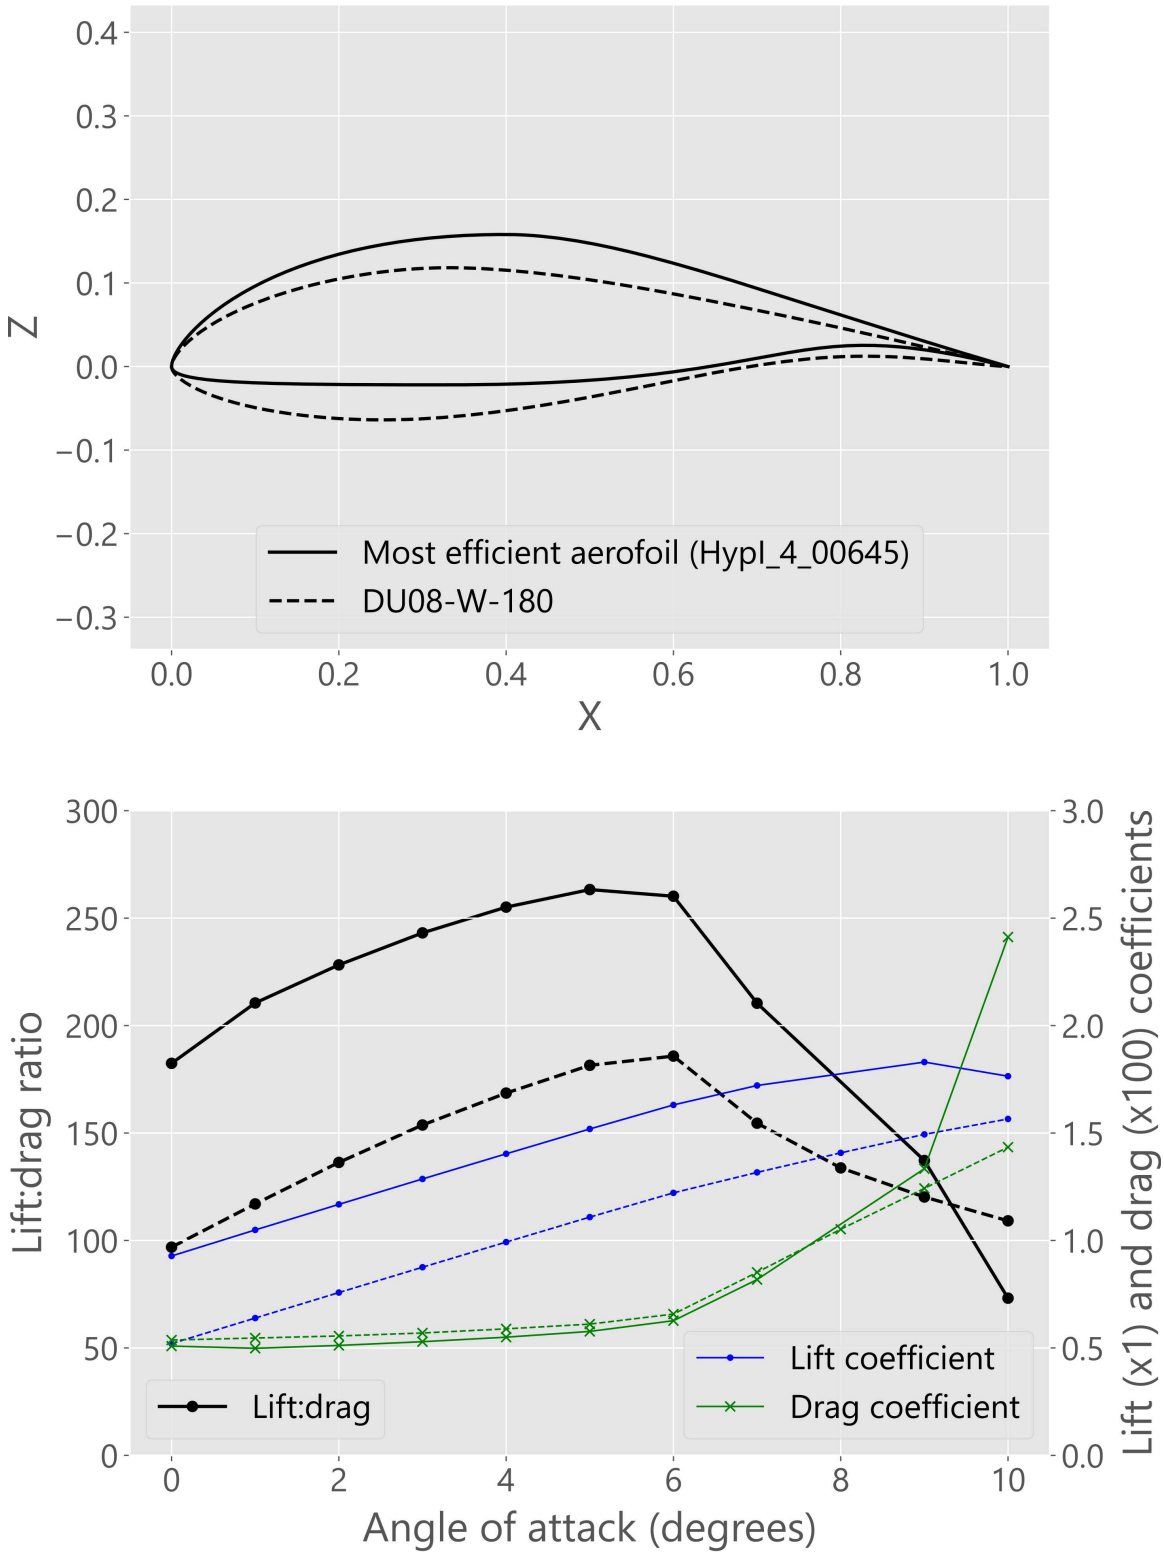

Figure 11: Comparison of the aerofoil with highest aerodynamic performance from the BO runs using HypI, with the standard aerofoil DU08-W-180.

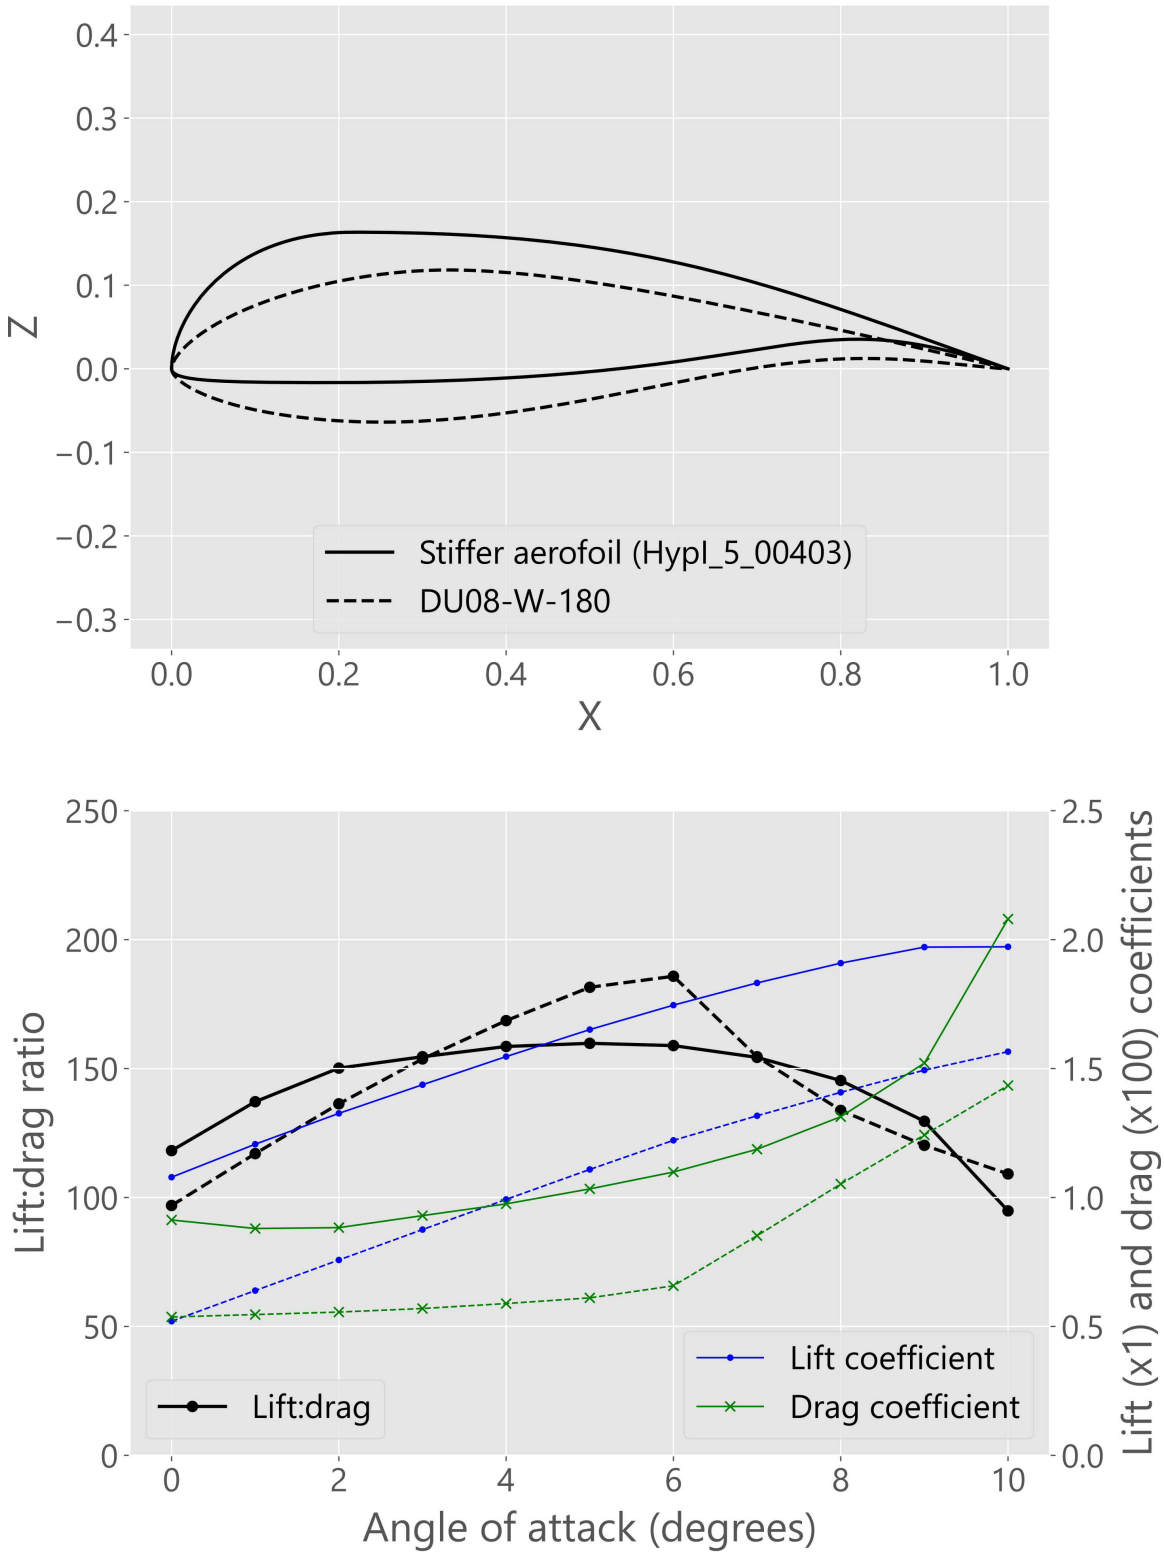

Figure 12: Comparison of the aerofoil found from the BO runs using HypI with the highest stiffness and equivalent aerodynamic performance to the standard aerofoil DU08-W-180.

6 Conclusions and Future Work

This study has shown that careful attention should be paid in Bayesian optimisation (BO) to the 'surprise' of objective values, relative to the mean and variance of the Gaussian process surrogate. This has an important effect on the exploration/exploitation trade-off. To assist with explicit management of this, a new exploration parameter based on normalisation of the data has been developed.

Taking this into account, an experimental investigation has then been performed on the effectiveness of two rather similar scalarising functions based on hypervolume (using the Expected Improvement acquisition function) and the multi-surrogate Expected Hypervolume Improvement acquisition function. The recently-introduced HypI outperforms the novel xHVI on every function. Further investigation identified that visualising the infill criterion in the objective space is not sufficient to understand its effectiveness. Consideration should rather be given to how well its value can be learned by the Gaussian process surrogate, and in particular whether the resulting surface will have a shallow and consistent gradient over both non-dominated and dominated solutions. This is desirable, to avoid areas of the search space remaining unexplored due to their extreme unattractiveness to the Expected Improvement acquisition function.

The practical usefulness of mono-objective infill criteria consists largely in speed and in re-using existing methods from single objective BO. While the xHVI infill criterion has been applied in parallel to a real wind turbine control problem [38], this paper applies xHVI, HypI and EHVI to the aerofoil design problem for wind turbine blades. A novel parameterisation for this task is introduced and shown to be valuable for *a priori* control over the decision space. HypI is demonstrated to be a powerful tool for this problem, obtaining a very large number of aerofoils which dominate standard aerofoils, using only 800 total evaluations.

There are two immediate areas for further investigation. First, this study has been performed using test functions without (aleatoric) evaluation uncertainty. It is important to now examine the influence of uncertainty on surrogate modelling of scalarising functions, since a Gaussian likelihood may not be a good assumption. Second, correlated objectives should be investigated, particularly as they have received as yet little attention in the multi-objective BO literature.

References

[1] Ath, G.D., Everson, R.M., Rahat, A.A.M., Fieldsend, J.E.: Greed is Good: Exploration and exploitation trade-offs in Bayesian optimisation. Submitted to ACM Transactions on Evolutionary Learning and Optimization (2019)

[2] Audet, C., Bigeon, J., Cartier, D., Digabel, S.L., Salomon, L.: Performance indicators in multiobjective optimization. techreport G-2018-90, GERAD, HEC Montreal, Canada (2018)

[3] Berk, J., Nguyen, V., Gupta, S., Rana, S., Venkatesh, S.: Exploration enhanced expected improvement for Bayesian optimization. In: Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 621–637. Springer (2018)

[4] Beume, N., Fonseca, C.M., López-Ibáñez, M., Paquete, L., Vahrenhold, J.: On the complexity of computing the hypervolume indicator. IEEE Transactions on Evolutionary Computation **13**(5), 1075–1082 (2009)

[5] Björck, A.: Airfoil design for horizontal axis wind turbines. In: 2nd IEA Symposium on Aerodynamics of Wind Turbines (1988)

[6] Boussaïd, I., Lepagnot, J., Siarry, P.: A survey on optimization metaheuristics. Information Sciences **237**, 82–117 (2013)

[7] Caboni, M., Minisci, E., Riccardi, A.: Aerodynamic Design Optimization of Wind Turbine Airfoils under Aleatory and Epistemic Uncertainty. Journal of Physics: Conf. Series **1037**(4) (2018)

[8] Chugh, T.: Scalarizing functions in Bayesian multiobjective optimization. In: IEEE Congress on Evolutionary Computation (2020)

[9] Couckuyt, I., Deschrijver, D., Dhaene, T.: Fast calculation of multiobjective probability of improvement and expected improvement criteria for Pareto optimization. Journal of Global Optimization **60**(3), 575–594 (2014)

[10] Daniels, S.J., Rahat, A.A.M., Tabor, G.R., Fieldsend, J.E., Everson, R.M.: Automated shape optimisation of a plane asymmetric diffuser using combined computational fluid dynamic simulations and multi-objective bayesian methodology. International Journal of Computational Fluid Dynamics **33**(6-7), 256–271 (2019). https://doi.org/10.1080/10618562.2019.1683165

[11] De Tavernier, D., Ferreira, C., van Bussel, G.: Airfoil optimisation for vertical-axis wind turbines with variable pitch. Wind Energy **22**(4), 547–562 (2019)

[12] Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable multi-objective optimization test problems. In: Proceedings of the 2002 Congress on Evolutionary Computation (2002)

[13] Emmerich, M., Yang, K., Deutz, A., Wang, H., Fonseca, C.M.: A multicriteria generalization of Bayesian global optimization. In: Pardalos, P.M., Zhigljavsky, A., Žilinskas, J. (eds.) Advances in Stochastic and Deterministic Global Optimization. Springer (2016)

[14] Emmerich, M.T.M., Deutz, A.H.: A tutorial on multiobjective optimization: fundamentals and evolutionary methods. Natural Computing **17**, 585—609 (2018)

[15] Ferreira, C.S., Geurts, B.: Aerofoil optimization for vertical-axis wind turbines. Wind Energy **18**, 1371–1385 (2015)

[16] Fonseca, C.M., Fleming, P.J.: On the performance assessment and comparison of stochastic multiobjective optimizers. In: Proceedings of the 1996 International Conference on Parallel Problem Solving from Nature. pp. 584–593 (1996)

[17] Fonseca, C.M., Paquete, L., López-Ibáñez, M.: An improved dimension-sweep algorithm for the hypervolume indicator. In: Proceedings of the 2006 IEEE Congress on Evolutionary Computation (2006)

[18] Fortin, F.A., De Rainville, F.M., Gardner, M.A., Parizeau, M., Gagné, C.: DEAP: Evolutionary algorithms made easy. Journal of Machine Learning Research **13**, 2171–2175 (2012)

[19] Jasrasaria, D., Pyzer-Knapp, E.O.: Dynamic control of explore/exploit trade-off in bayesian optimization. In: Proceedings of the Computing Conference 2018, London, UK (2018)

[20] Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. Journal of Global Optimization **13**, 455–492 (1998)

[21] Knowles, J.: Parego: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. IEEE Transactions on Evolutionary Computation **10**(1), 50–66 (2005)

[22] Koehler, J., Owen, A.: Design and Analysis of Experiments, chap. 9, pp. 295–303. Handbook of Statistics, North-Holland (1996)

[23] Malkomes, G., Garnett, R.: Automating Bayesian optimization with Bayesian optimization. In: Proceedings of the 32nd International Conference on Neural Information Processing. pp. 5988–5997 (2018)

[24] Maree, S.C., Alderliesten, T., Bosman, P.A.N.: Uncrowded hypervolume-based multi-objective optimization with gene-pool optimal mixing (2020), arXiv pre-print

[25] Masters, D.A., Taylor, N.J., Rendall, T., Allen, C.B., Poole, D.J.: Review of aerofoil parameterisation methods for aerodynamic shape optimisation. In: 53rd AIAA Aerospace Sciences Meeting (2015)

[26] Ponweiser, W., Wagner, T., Biermann, D., Vincze, M.: Multiobjective optimization on a limited budget of evaluations using model-assisted S-metric selection. In: Proceedings of the 10th International Conference on Parallel Problem Solving from Nature. pp. 784–794 (2008)

[27] Rahat, A.A.M., Everson, R.M., Fieldsend, J.E.: Alternative infill strategies for expensive multi-objective optimisation. In: Proceedings of the 2017 Genetic and Evolutionary Computation Conference (2017)

[28] van Rooij, R.P.J.O.M.: Modification of the boundary layer calculation in RFOIL for improved airfoil stall prediction. Tech. Rep. IW-96087R, TU Delft (1996)

[29] Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., de Freitas, N.: Taking the human out of the loop: A review of Bayesian optimization. Proceedings of the IEEE **104**(1), 148–175 (2016)

[30] Skinner, S., Zare-Behtash, H.: State-of-the-art in aerodynamic shape optimisation methods. Applied Soft Computing **62**, 933 – 962 (2018)

[31] Snoek, J., Larochelle, H., Adams, R.P.: Practical Bayesian optimization of machine learning algorithms. In: Proceedings of the 25th International Conference on Neural Information Processing Systems. vol. 2, pp. 2951–2959 (2012)

[32] Sripawadkul, V., Padulo, M., Guenov, M.: A comparison of airfoil shape parameterization techniques. In: 13th AIAA/ISSMO Mul-tidisciplinary Analysis Optimization Conference (2010)

[33] The GPyOpt authors: GPyOpt: A Bayesian optimization framework in Python. http://github.com/SheffieldML/GPyOpt (2016)

[34] Timmer, W., Bak, C.: Advances in Wind Turbine Blade Design and Materials, chap. 4, pp. 109–149. Woodhead Publishing Series in Energy (2013)

[35] Veldhuizen, D.A.V., Lamont, G.B.: Multiobjective evolutionary algorithm research: A history and analysis. techreport TR-98-03, Air Force Institute of Technology, Department of Electrical and Computer Engineering, Graduate School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, Ohio (1998)

[36] Yang, K., Emmerich, M., Deutz, A., Bäck, T.: Efficient computation of expected hypervolume improvement using box decomposition algorithms. Journal of Global Optimization **75**, 3–34 (2019)

[37] Yang, K., Emmerich, M., Deutz, A., Fonseca, C.M.: Computing 3-D expected hypervolume improvement and related integrals in asymptotically optimal time. In: Proceedings of the 2017 Conference on Evolutionary Multi-Criterion Optimization. pp. 685–700 (2017)

[38] Yu, W., Engels, W., Stock-Williams, C.: A comparison of multi-objective optimisation of two wind turbine controller designs. In: Proceedings of the 2020 Conference on The Science of Making Torque from Wind (2020)

[39] Zhan, D., Xing, H.: Expected improvement for expensive optimization: a review. Journal of Global Optimization (2020)

[40] Zitzler, E., Brockhoff, D., Thiele, L.: The hypervolume indicator revisited: On the design of Pareto-compliant indicators via weighted integration. In: Proceedings of the 2007 International Conference on Evolutionary Multi-Criterion Optimization. pp. 862–876 (2007)

[41] Zitzler, E., Deb, K., Thiele, L.: Comparison of multiobjective evolutionary algorithms: Empirical results. Evolutionary Computation **8**(2), 173–195 (2000)