

Hull-Shape Optimisation Using Adaptive Multi-Fidelity Kriging

Thomas Scholcz and Joy Klinkenberg
Maritime Research Institute Netherlands (MARIN)
THE NETHERLANDS

Keywords: hull-shape optimisation, SBDO, adaptive, multi-fidelity, Kriging, CFD, noise treatment

ABSTRACT

The paper presents and discusses the development and assessment of an active learning multi-fidelity Kriging method for military vehicle design within the NATO AVT-331 task group on “Goal-driven, multi-fidelity approaches for military vehicle system-level design.” This Simulation Based Design Optimisation (SBDO) method exploits the posterior correlation between the low and high-fidelity processes by defining an augmented expected improvement function. Using the method of activation coefficients, it is described how to obtain and implement the posterior correlation function in a practical way. Subsequently, the implementation is verified and single- and multi-fidelity methods are applied to analytical benchmarks and hull-shape optimisation of the bare-hull DTMB 5415 frigate. From these test cases it is concluded that a more stable convergence is obtained from the multi-fidelity method, when compared to the single-fidelity method using the same computational budget. Moreover, the multi-fidelity method may result in a computational speedup depending on the target error, noise levels, evaluations costs and correlation between the models. A significant computational speedup is observed for most cases in this study and the speedup improves with increasing problem dimensions.

1.0 INTRODUCTION

Simulation Based Design Optimisation (SBDO) requires a trade-off between computational cost and solver accuracy. When the cost per evaluation is of interest, surrogate models may be used to mitigate computational costs considerably, see e.g. [1, 2, 3, 4]. In the past few years, research has been focused on making these surrogate models more efficient using a combination of accurate high-fidelity solvers and cheap but less accurate low-fidelity solvers, also called *multi-fidelity surrogate models*. These methods have been investigated and applied in the maritime domain for many years now, see for example [5, 6, 7, 8, 9].

When replacing the actual response surface with a surrogate, the accuracy of the analysis is limited by the surrogate model accuracy. To be useful for design, this accuracy should be either known or adaptively improved to meet certain design objectives. The most common cases are: aiming for *global accuracy* and aiming for *optimisation*. In the latter case, the surrogate model is improved to gain efficiency in the optimisation algorithm. The simulation output is often corrupted with noise which may result from numerical convergence issues, discretisation errors and/or geometrical parametrisation. Actively improving a surrogate model to facilitate optimisation on noisy black-box functions is a challenge. This process is also called *active learning* because it uses the current state to decide at which location in the design space a new evaluation should be done to be most beneficial for the optimisation. This procedure can be extended to multiple levels of fidelity, each having a different cost and noise process. In this case, not only the location in the design space has to be chosen, also the fidelity level of the new evaluation is a choice that has to be made. Interesting applications in the maritime field are for example found in [10, 11, 12].

In this contribution, we implement and assess the performance of an active learning multi-fidelity method applied to numerical benchmarks and hull-shape optimisations of the bare-hull DTMB 5415 model. Using the adaptive multi-fidelity Kriging method described in [13], we explore practical implementation issues as well as performance in terms of computational speedup. The method exploits the posterior *correlation of*

the high and low-fidelity processes during optimisation and is the backbone of many other adaptive multi-fidelity methods such as described in [14, 15]. Special attention is given to the estimation of noise during the optimisation.

2.0 METHOD

The adaptive multi-fidelity Kriging method consists of the following steps:

- 1) Build an initial multi-fidelity Kriging model using initial Design of Experiments.
- 2) Find the next location in the design space and the next fidelity level of the evaluation that should be performed on this point with the goal to improve the efficiency of the method, using an *acquisition function*.
- 3) Update the multi-fidelity Kriging model using the new evaluation. Go to step 2.

In the following section we will describe the multi-fidelity method and how to implement the acquisition functions corresponding to the different fidelity levels. Next, the adaptive training method is described in more detail and a test case is used to verify correct implementation. Noise-free and noisy benchmark problems are subsequently used to investigate the performance of the method and in particular the noise treatment. Finally, the method is applied to hull optimisation of the DTMB 5415 frigate using the linear potential flow solver WARP (see, [16]) and using the RANS solver ReFRESCO (see, [17, 18]) with the objective to minimise resistance. In the latter case, the aft ship is optimised.

2.1 Multi-Fidelity Kriging

We follow the notation of Kennedy and O'Hagan, see [19]. The high fidelity process z_1 and low fidelity process z_2 are related by

$$z_1(\mathbf{x}) = z_2(\mathbf{x}) + \delta(x), \quad (1)$$

where the difference process $\delta(\mathbf{x})$ is modelled by an independent Gaussian process. In Equation 1 we do not multiply z_2 with a scale parameter ρ and we use the AVT-331 convention on the order of the fidelity levels (1=high fidelity, 2=low fidelity) which is different from the convention in [19]. The mean (E) and covariance (cov) of the posterior distribution are given by

$$E(f_l^p(\mathbf{x})) = \mathbf{h}_l(\mathbf{x})\boldsymbol{\beta} + \mathbf{t}_l(\mathbf{x})\mathbf{V}^{-1}(\mathbf{z} - \mathbf{H}\boldsymbol{\beta}) \quad (2)$$

$$\text{cov}\left(f_l^p(\mathbf{x}), f_l^p(\mathbf{x}')\right) = \text{cov}\left(f_{l'}(\mathbf{x}'), f_l(\mathbf{x})\right) - \left(\mathbf{h}_{l'}(\mathbf{x}')^T, \mathbf{t}_{l'}(\mathbf{x}')^T\right) \begin{pmatrix} \mathbf{0} & \mathbf{H}^T \\ \mathbf{H} & \mathbf{V} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{h}_l(\mathbf{x}) \\ \mathbf{t}_l(\mathbf{x}) \end{pmatrix}, \quad (3)$$

where $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2)$ are the high and low fidelity evaluation outputs on the designs-of-experiments D_1 and D_2 . In [20], the coefficients a_1 and a_2 are defined (deviating from Kennedy and O'Hagan) to compute the uncertainty of the low fidelity process and the difference process. Similarly, we define the activation coefficients $\mathbf{a} = (a_1, a_2)$ and $\mathbf{a}' = (a'_1, a'_2)$ to compute the posterior moments

$$\begin{aligned} E_{\mathbf{a}}(f_l^p(\mathbf{x})) &= \mathbf{h}_l(\mathbf{x})\boldsymbol{\beta} + \mathbf{t}_l(\mathbf{x})\mathbf{V}^{-1}(\mathbf{z} - \mathbf{H}\boldsymbol{\beta}) \\ \text{cov}_{\mathbf{a}, \mathbf{a}'}\left(f_l^p(\mathbf{x}), f_l^p(\mathbf{x}')\right) &= a'_1\sigma_1^2 + a'_2\sigma_2^2 - \mathbf{t}_{l'}(\mathbf{x}')^T\mathbf{V}^{-1}\mathbf{t}_l(\mathbf{x}) + \\ &\quad \left(\mathbf{h}'_l(\mathbf{x}) - \mathbf{t}'_l(\mathbf{x})\mathbf{V}^{-1}\mathbf{H}\right)^T \left(\mathbf{H}^T\mathbf{V}^{-1}\mathbf{H}\right)^{-1} \left(\mathbf{h}_l(\mathbf{x}) - \mathbf{t}_l(\mathbf{x})\mathbf{V}^{-1}\mathbf{H}\right) \end{aligned} \quad (4)$$

with

$$\begin{aligned}
\mathbf{h}_l(\mathbf{x}) &= (a_1, a_2), \\
\mathbf{h}_{l'}(\mathbf{x}) &= (a'_1, a'_2), \\
\mathbf{t}_l(\mathbf{x})^T &= (a_1\sigma_1^2 A_1(\mathbf{x}, D_1) + a_2\sigma_2^2 A_2(\mathbf{x}, D_1), a_2\sigma_2^2 A_2(\mathbf{x}, D_2)), \\
\mathbf{t}_{l'}(\mathbf{x})^T &= (a'_1\sigma_1^2 A_1(\mathbf{x}, D_1) + a'_2\sigma_2^2 A_2(\mathbf{x}, D_1), a'_2\sigma_2^2 A_2(\mathbf{x}, D_2)), \\
\mathbf{V} &= \begin{pmatrix} \sigma_1^2 A_1(D_1, D_1) + \sigma_2^2 \sigma_2^2 A_2(D_1, D_1) & \sigma_2^2 A_2(D_1, D_2) \\ \sigma_2^2 A_2(D_2, D_1) & \sigma_2^2 A_2(D_2, D_2) \end{pmatrix} + \begin{pmatrix} \epsilon_1^2 \mathbf{1}_{n_1} & \mathbf{0} \\ \mathbf{0} & \epsilon_2^2 \mathbf{1}_{n_2} \end{pmatrix}, \\
\mathbf{H} &= \begin{pmatrix} \mathbf{1}_{n_1} & \mathbf{1}_{n_1} \\ \mathbf{0} & \mathbf{1}_{n_2} \end{pmatrix}.
\end{aligned} \tag{5}$$

The hyper parameters $\beta = (\beta_1, \beta_2)$, σ_1 and σ_2 are determined from the statistical mean and variance of the observations. The matrices A_1 and A_2 are the correlation matrices defined by

$$A_l(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\sum_{d=1}^{N_d} \frac{|x_j^d - x_i^d|^2}{2\theta_{l,d}^2}\right), \tag{6}$$

where a different correlation length is allowed for each dimension. The correlation lengths are estimated using Maximum Likelihood Estimation (MLE). The noise parameters of the difference process ϵ_1 and the low fidelity process ϵ_2 are treated as hyper parameters (similar to the correlation lengths) in the MLE procedure. Each combination of activation coefficients results in certain output of the posterior moments, see Table 1. This will be referred to as the *method of activation coefficients* in the following.

Table 1: Meaning of Posterior Moments for Values of \mathbf{a} and \mathbf{a}' .

a_1	a_2	a'_1	a'_2	l	l'	Meaning of $E_{\mathbf{a}}(f_l^p(\mathbf{x}))$	Meaning of $\text{cov}_{\mathbf{a}, \mathbf{a}'}(f_l^p(\mathbf{x}), f_{l'}^p(\mathbf{x}))$
1	1	1	1	1	1	Mean of the high-fidelity process	Variance of the high-fidelity process
1	1	0	1	1	2	Mean of the high-fidelity process	Covariance of the low and high fidelity process
0	1	0	1	2	2	Mean of the low-fidelity process	Variance of the low-fidelity process
1	0	1	0	-	-	Mean of the difference process	Variance of the difference process

The posterior covariance of the low and high fidelity process and/or the moments of the difference process can be used in an adaptive refinement process. The adaptive multi-fidelity method from [20], for example, exploits the variance of the low fidelity process and the difference process whereas the covariance between the high and low fidelity process is used in the multi-fidelity Kriging methods described in [13, 14]. As mentioned earlier, the posterior covariance is used here to define acquisition functions for each level of fidelity.

2.2 Adaptive Training Method

The aim is to find the global minimum of a noisy high fidelity black-box function as efficient as possible. That is, for a given budget, the approximate solution to

$$\mathbf{x}^* = \underset{\mathbf{x} \in \mathcal{A}}{\text{argmin}} f_1(\mathbf{x}), \tag{7}$$

should be as close as possible to the actual minimum. Given a finite sequence of n approximations $\mathbf{x}_1 \dots \mathbf{x}_n$, the best approximation is defined by

$$\mathbf{x}^* = \underset{\mathbf{x} \in \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}}{\text{argmin}} E_{\mathbf{a}}(f_1^p(\mathbf{x})) \quad \text{with } \mathbf{a} = (1, 1). \tag{8}$$

To find the next sample, the expected improvement acquisition function

$$\begin{aligned}
 EI(\mathbf{x}) &= E_{\mathbf{a}} \left(\max(E_{\mathbf{a}}(f_1^p(\mathbf{x}^*)) - f_1^p(\mathbf{x}), 0) \right) \\
 &= (E_{\mathbf{a}}(f_1^p(\mathbf{x}^*)) - E_{\mathbf{a}}(f_1^p(\mathbf{x}))) \Phi \left(\frac{E_{\mathbf{a}}(f_1^p(\mathbf{x}^*)) - E_{\mathbf{a}}(f_1^p(\mathbf{x}))}{\sqrt{\text{Var}_{\mathbf{a}}(f_1^p(\mathbf{x}))}} \right) \\
 &\quad + \sqrt{\text{Var}_{\mathbf{a}}(f_1^p(\mathbf{x}))} \phi \left(\frac{E_{\mathbf{a}}(f_1^p(\mathbf{x}^*)) - E_{\mathbf{a}}(f_1^p(\mathbf{x}))}{\sqrt{\text{Var}_{\mathbf{a}}(f_1^p(\mathbf{x}))}} \right) \quad \text{with } \mathbf{a} = (1, 1),
 \end{aligned} \tag{9}$$

is used with Φ and ϕ the standard normal cumulative density and probability density functions respectively. To account for the different levels of fidelity, the improvement function is α ugmented, resulting in [13, 14]

$$EI_{\alpha}(\mathbf{x}, l) = EI(\mathbf{x})\alpha_1(\mathbf{x}, l)\alpha_2(\mathbf{x}, l)\alpha_3(\mathbf{x}, l). \tag{10}$$

The next point \mathbf{x}_{n+1} at fidelity level l_{n+1} is computed from

$$(\mathbf{x}_{n+1}, l_{n+1}) = \underset{\mathbf{x}, l}{\text{argmax}} EI_{\alpha}(\mathbf{x}, l), \tag{11}$$

with the factors α_1 , α_2 and α_3 given by

$$\alpha_1(\mathbf{x}, l) = \text{cov}_{\mathbf{a}, \mathbf{a}'} \left(f_1^p(\mathbf{x}), f_l^p(\mathbf{x}) \right) / \left(\sqrt{\text{Var}_{\mathbf{a}}(f_1^p(\mathbf{x}))} \sqrt{\text{Var}_{\mathbf{a}}(f_l^p(\mathbf{x}))} \right) \tag{12}$$

$$\alpha_2(\mathbf{x}, l) = 1 - \frac{\epsilon_l}{\sqrt{\text{Var}_{\mathbf{a}}(f_l^p(\mathbf{x}))} + \epsilon_l^2} \tag{13}$$

$$\alpha_3(\mathbf{x}, l) = \frac{c_1}{c_l}. \tag{14}$$

The optimisation in Equation 11 is performed using a restarted L-BFGS-B method, see [21]. The adaptive training method is stopped using a fixed computational budget where the cost of a high fidelity simulation equals 1.0 and the cost of a low fidelity simulation equals a fraction of the high fidelity cost.

2.2.1 1D Sasena Test Case

To verify a correct implementation the Sasena test case from [13, 22] is repeated. The high and low fidelity functions are given by

$$f_1(x) = -\sin(x) - \exp(x/100) + 10 \quad (0 < x < 10) \tag{15}$$

$$f_2(x) = -\sin(x) - \exp(x/100) + 10.3 + 0.03(x - 3)^2 \quad (0 < x < 10), \tag{16}$$

The initial design-of-experiments for the low fidelity function $x = \{0, 2, 4, 6, 8, 10\}$ is used and for the high fidelity function $x = \{3.5, 6.5\}$. Since there is no noise it holds that $\alpha_2(x, 1) = \alpha_2(x, 2) = 1$. The cost ratio is set to $\alpha_3(x, 2) = 4.0$.

The value of the expected improvement function EI_{α} as a well as the posterior correlation α_1 just before the 9th point was added is shown in Figure 1. The expected improvement function is more oscillatory than reported in [13] but shows good agreement in the relative values of the improvements across the different levels of fidelity. The posterior correlation α_1 is in very close agreement to the one reported in [13].

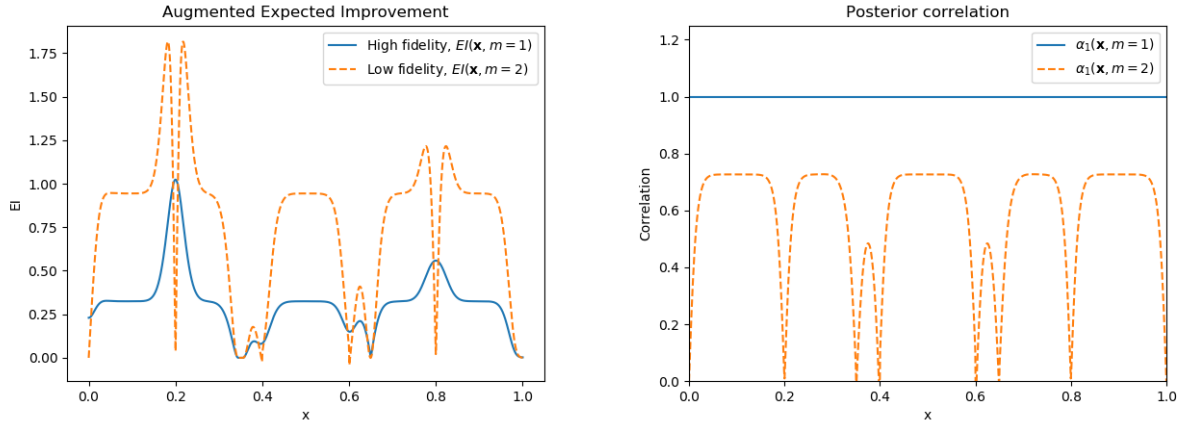


Figure 1: Sasena Test Case: Scaled Function EI_α (Left) and Posterior Correlation α_1 (Right) Just Before the 9th Point was Added.

3.0 ANALYTICAL OPTIMISATION BENCHMARKS

The analytical benchmarks used in this study are obtained from [23]. A subset of these benchmarks are listed in Table 2. For each test in the table, three functions with decreasing fidelity are given, with f_1 the function with the highest fidelity. In addition, the domain, the dimensions D , the analytical solutions $\hat{\mathbf{x}}$ and corresponding function values $f(\hat{\mathbf{x}})$ are given. Three error metrics are defined for these benchmark problems.

Table 2: Analytical Benchmarks.

Test	Formulation	Domain	D	$\hat{\mathbf{x}}$	$f(\hat{\mathbf{x}})$
P_1	$f_1(x) = (6x - 2)^2 \sin(12x - 4) + \mathcal{N}(0, \hat{\epsilon}_1)$	$x \in [0, 1]$	1	0.7572	-6.0207
	$f_2(x) = 0.75f_1(x) + 5(x - 0.5) - 2 + \mathcal{N}(0, \hat{\epsilon}_2)$				
	$f_3(x) = 0.5f_1(x) + 10(x - 0.5) - 5 + \mathcal{N}(0, \hat{\epsilon}_3)$				
P_2	$f_1(\mathbf{x}) = \sum_{j=1}^{D-1} (100(x_{j+1} - x_j^2)^2 + (1 - x_j)^2) + \mathcal{N}(0, \hat{\epsilon}_1)$	$\mathbf{x} \in [-2, 2]$	2	[1, ..., 1]	0.0000
	$f_2(\mathbf{x}) = \sum_{j=1}^{D-1} (50(x_{j+1} - x_j^2)^2 + (-2 - x_j)^2) - \sum_{j=1}^D 0.5x_j + \mathcal{N}(0, \hat{\epsilon}_2)$				
	$f_3(\mathbf{x}) = (f_1(\mathbf{x}) - 4 - \sum_{j=1}^D 0.5x_j) / (10 + \sum_{j=1}^D 0.25x_j) + \mathcal{N}(0, \hat{\epsilon}_3)$				

These characterize the normalized error in the design space, the objective function, and Euclidean distance in the normalized x-f hyperspace, respectively:

$$E_x \equiv \frac{\|\mathbf{x}^* - \hat{\mathbf{x}}\|}{\sqrt{D}} \tag{17}$$

$$E_f \equiv \frac{f(\mathbf{x}^*) - f(\hat{\mathbf{x}})}{R_1} \tag{18}$$

$$E_t \equiv \sqrt{\frac{E_x^2 + E_f^2}{2}}. \tag{19}$$

For each problem we study the convergence in these error metrics for the multi- and single-fidelity adaptive training method. The errors are plotted as a function of the consumed computational budget, which allows for a fair comparison between the methods.

Using Kriging, several convergence phases can be distinguished, each having their own characteristics. The convergence of a single-fidelity Kriging predictor has three phases: the uncorrelated, correlated and noise-inhibited phase, see [5]. During the uncorrelated phase the distance between samples is larger than the correlation range, making it appear uncorrelated for the Kriging predictor. This phase is characterised by a

large error and irregularity. As samples are added, the convergence moves into the correlated phase. This phase is characterised by a larger convergence rate due to the increased resolution and smaller distance between samples. As soon as the error drops below the noise level of the process, the noise-inhibited phase starts. The convergence rate of the noise-inhibited phase is typically much smaller than the convergence rate during the correlated phase. In case of a noise-free process, this phase results in stagnation of the convergence due to the finite precision of the underlying optimisation routines.

The multi-fidelity Kriging convergence has the same phases as the single-fidelity Kriging convergence, but the length (not necessarily the convergence rate) of the phases may differ. Because the multi-fidelity adaptive training method exploits cheap samples, it may have a shorter uncorrelated phase than the single-fidelity training method which results in a computational speedup. Another mechanism to achieve speedup is the effect of the (better) initial Kriging model on the error level. In the noise-inhibited phase however, the speedup becomes much smaller. Eventually, the speedup may disappear entirely.

The usefulness of a multi-fidelity method on practical hull-shape optimisation problems depends on the desired target error, correlation and evaluation costs of the models. When the target error is located in the correlated phase of both convergence processes and the correlation/costs of the models is beneficial, a speedup can be expected. When the target error is located in the uncorrelated or noise-inhibited phase however, it is highly uncertain what the speedup will be. In the following sections we will explore these characteristics using noise-free analytical benchmarks ($\hat{\epsilon}_l = 0$) and analytical benchmarks including artificial noise ($\hat{\epsilon}_l \neq 0$) for fidelity levels $l = 1, 2, 3$. The hat operator is used in the notation for the *artificial* (true) noise level in order to distinguish it from the *estimated* noise levels ϵ_l defined in Section 2.1.

3.1 Benchmarks Without Noise


Table 3 summarizes the initial design of experiments, budget, dimensions D , function costs c_l , function range R_1 and the noise levels $\hat{\epsilon}_l$ for each test case. The initial design is taken constant and given by a Central Composite Face-centered (CCF) design of experiments for both the initial high and low fidelity evaluations. The rationale behind this choice is that it provides a relatively coarse initial sampling plan, leaving more budget for the adaptive phase of the optimisation method. The function range is obtained from the analytical extremes reported in [23], using $R_1 = f_{\max} - f_{\min}$. Since there is no noise, we have $\hat{\epsilon}_l = 0$ for $l = 1, 2, 3$.

Table 3: Test Settings for Benchmarks Without Noise.


Test	Initial design	Budget	D	c_1	c_2	c_3	R_1	$\hat{\epsilon}_1$	$\hat{\epsilon}_2$	$\hat{\epsilon}_3$
P_1	CCF/CCF	100	1	1.0	0.5	0.1	2.1851E+1	0.0	0.0	0.0
P_2	CCF/CCF	200	2	1.0	0.5	0.1	3.6090E+1	0.0	0.0	0.0

Figure 2 shows convergence results for test P_1 and P_2 . The notation "L1-2" means that the multi-fidelity method uses high-fidelity level 1 in combination with low-fidelity level 2. Using the Forrester case, the convergence phases can be clearly distinguished and we see that L1-3 convergence results in a shorter uncorrelated phase, while the convergence rate during the correlated phase remains similar to the single-fidelity result. However, for L1-2 the correlation/costs of the models is not beneficial which results in no observed speedup. In this case, the convergence is close to the single-fidelity result but the behaviour is more erratic. Using the 2D Rosenbrock case, it is harder to distinguish the convergence phases. The convergence of L1-3 is found to be more stable than the others and reaches lower values for E_f during the optimisation. The differences in performance are, however, not very big. In all cases, the total error E_t is dominated by the normalised error in the design space E_x , which is orders of magnitude larger than the normalised error in the objective function.

Table 4 summarizes the results at the final Computational Cost (CC) of the simulations. The J value in the last column indicates how many high and low-fidelity evaluations were performed during the optimisations.



(a) Single-Fidelity and Multi-Fidelity Results Using the Forrester Test Case (P1).



(b) Single-Fidelity and Multi-Fidelity Results Using the 2D Rosenbrock Test Case (P2).

Figure 2: Convergence of E_x, E_f, E_t for Test P_1 and P_2 .

Table 4: Performance Comparisons at the Final Computational Cost.

Test	D	L	CC	$E_x\%$	$E_f\%$	$E_t\%$	\mathcal{J}
P_1	1	1	100	1.77E-4	4.00E-8	1.25E-4	100
	1	1-2	100	3.20E-4	3.33E-8	2.27E-4	92-16
	1	1-3	100	3.69E-4	3.40E-8	2.61E-4	95-50
P_2	2	1	200	1.41	3.71E-4	9.96E-1	200
	2	1-2	200	1.65	5.99E-5	1.16	193-14
	2	1-3	200	2.16	8.60E-5	1.52	198-20


3.2 Benchmarks Including Artificial Noise

Since many processes do not show smooth and clean behaviour for variations of the design parameters but contain some level of noise [4], it is interesting to find out how well this noise is quantified by the optimisation algorithm. The experimental setting in [24] is used here. A fixed initial DoE is used, in combination with random realisations of an artificial noise generator with known probability density function. The experiments are described in Table 5. The noise generator is a zero mean normal process with standard deviations given by $\hat{\epsilon}_1 = 0.025R_1$, $\hat{\epsilon}_2 = 0.05R_1$ and $\hat{\epsilon}_3 = 0.1R_1$ for test P_1 and $\hat{\epsilon}_1 = \frac{0.025}{500}R_1$, $\hat{\epsilon}_2 = \frac{0.05}{500}R_1$ and $\hat{\epsilon}_3 = \frac{0.1}{500}R_1$ for test P_2 , see also Table 2. The initial range R_1 is computed from the maximum and minimum high fidelity function value evaluated on the initial design of experiment. Note that the budgets and solver costs in this table are different from the budgets and solver costs in Table 3. The convergence metrics of the adaptive


Table 5: Test Settings for Benchmarks Including Artificial Noise.

Test	Initial design	Budget	D	c_1	c_2	c_3	R_1	$\hat{\epsilon}_1$	$\hat{\epsilon}_2$	$\hat{\epsilon}_3$
P_1	CCF/CCF	30	1	1.0	0.2	0.1	1.4904E+1	0.3730E-0	0.7460E-0	1.4920E-0
P_2	CCF/CCF	50	2	1.0	0.2	0.1	1.6080E+3	8.0400E-2	1.6080E-1	3.2160E-1
	CCF/CCF	65	5	1.0	0.2	0.1	2.0080E+3	1.0040E-1	2.0080E-2	4.0160E-1

training methods applied on the Forrester test case are given in Figure 3. The experiments were repeated for 50 different random seeds. In order to compute statistics, the errors were linearly interpolated on a equidistant grid with a resolution equal to the lowest fidelity cost. Subsequently, the median and interquartile



(a) Single Fidelity Result (Solid) and Multi-Fidelity Result Using Level 1 & 2 (Dashed). The Shaded Area Indicates the Interquartile Range.



(b) Single Fidelity Result (Solid) and Multi-Fidelity Result Using Level 1 & 3 (Dashed). The Shaded Area Indicates the Interquartile Range.

Figure 3: Convergence of E_x, E_f, E_t Using the 1D Forrester Test Case (P1).

range of the grid errors were computed from the grid values. It can be seen that fidelity level 2 results in a shorter uncorrelated range of the convergence, but the initial surrogate does not result in a speedup. Fidelity level 3 does not result in a speedup at all. Because of the relatively high noise level, this result is different from the result in Section 3.1. We also observe that the sensitivity of the L1-2 multi-fidelity error w.r.t noise is smaller, having a more narrow interquartile range than observed for the single fidelity error. In addition to these metrics, a fourth metric is used to measure the accuracy of the noise estimate during the adaptive training method:

$$E_n \equiv \left| \frac{\epsilon - \hat{\epsilon}}{\hat{\epsilon}} \right|. \quad (20)$$

It measures how close the MLE noise level is to the imposed artificial (True) noise level in the process. From Figure 4, we observe convergence up to a level of about 10% of the artificial noise level in both cases. Typically, the noise estimate is better for the high fidelity process and the difference process than for the low-fidelity process. This results from the fact that high fidelity evaluations are performed close together near the optimum while low-fidelity evaluations are performed in the initial phase of the multi-fidelity optimisation. In that way, the low-fidelity noise is "less visible" to the optimiser. The analysis is repeated for the 2D Rosenbrock test cases, see Figure 5. Fidelity level 2 does not shorten the uncorrelated phase of the multi-fidelity method, but results in a speedup by the initial surrogate. In the noise-inhibited phase, the convergence of the methods are very similar. Again, no significant speedup is observed using fidelity level 3, but convergence is observed in the noise-inhibited phase for both methods. A typical characteristic of this test case is that the optimum is located in a flat valley, resulting in orders magnitude difference in E_x compared to E_f . The convergence of the noise estimate is shown in Figure 6. The first phase demonstrates the struggle of the Kriging processes to distinguish noise from actual function trends during the uncorrelated phase. We also see that convergence of the low-fidelity noise estimate is not required to have convergence of the multi-fidelity method. The results for the Rosenbrock problem in 5 dimensions are shown in Figure 7. The budget is not sufficient to capture the noise-inhibited phase, but the end of the uncorrelated phase is visible. A large speedup due to the initial surrogate is achieved for both fidelity levels, but the L1-2 multi-fidelity method is

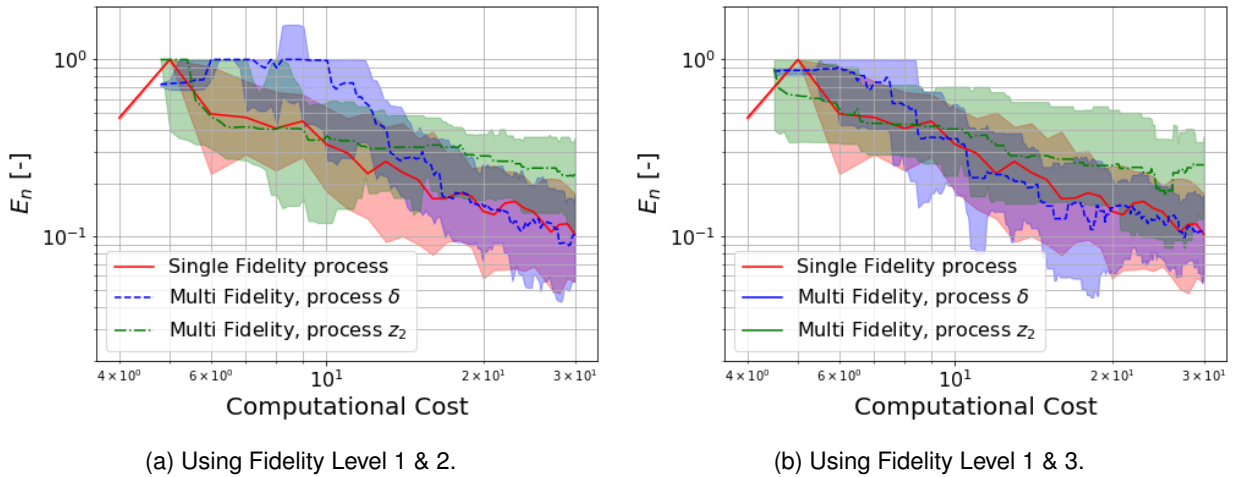


Figure 4: Convergence of E_n Using the 1D Forrester Test Case (P1).

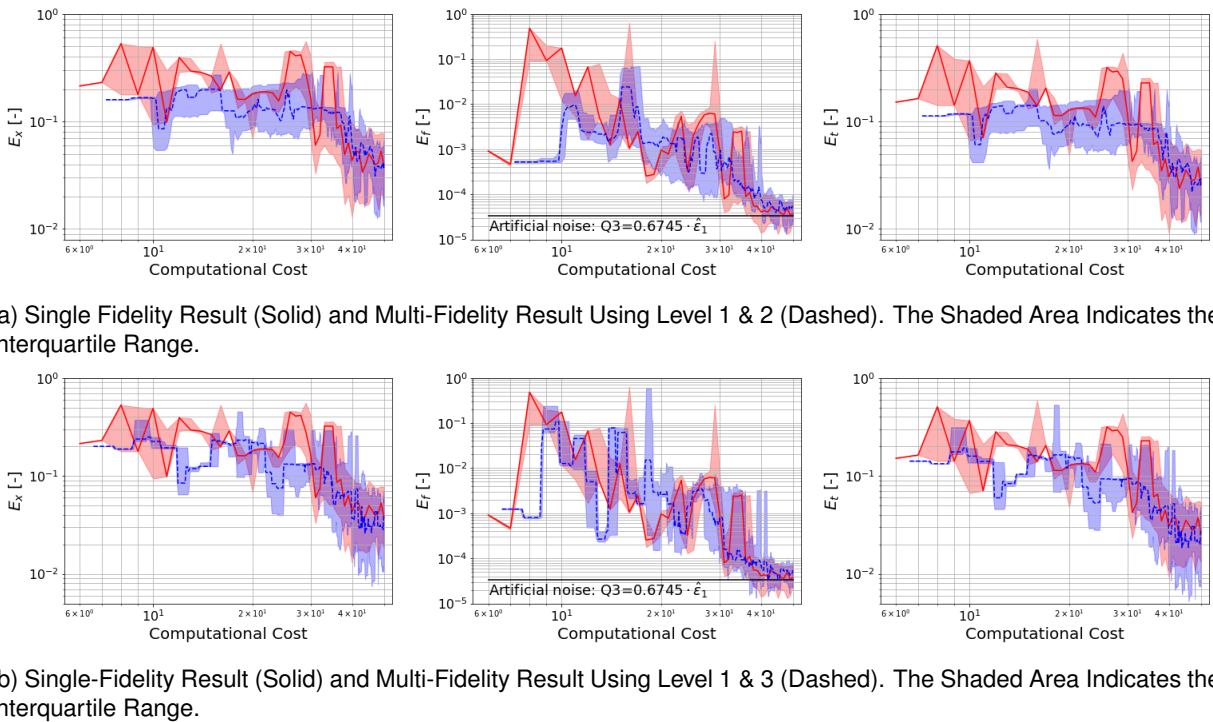


Figure 5: Convergence of E_x, E_f, E_t Using the 2D Rosenbrock Test Case (P2).

still more efficient than the L1-3 method. As computational cost increases the benefit of the multi-fidelity methods diminishes, to become zero in the noise-inhibited phase. Compared to the previous $D = 2$ case, we observe a larger speedup during the uncorrelated phase. The typical bumpy behaviour of the noise estimate during the uncorrelated phase is shown in Figure 8. While the multi-fidelity noise estimates already settle, the single-fidelity estimate is far from converged. The multi-fidelity convergence behaviour is found to be more stable than the single-fidelity convergence behaviour. This is an additional advantage of the adaptive multi-fidelity Kriging method. Table 6 summarizes the errors at the end of the computational budgets. As was pointed out, the speedup depends on the target error and in which phase of the Kriging convergence this target error is located. A large speedup may be present, which may vanish entirely in the noise inhibited convergence phase. End-budget error comparisons are therefore dangerous for conclusions on compu-

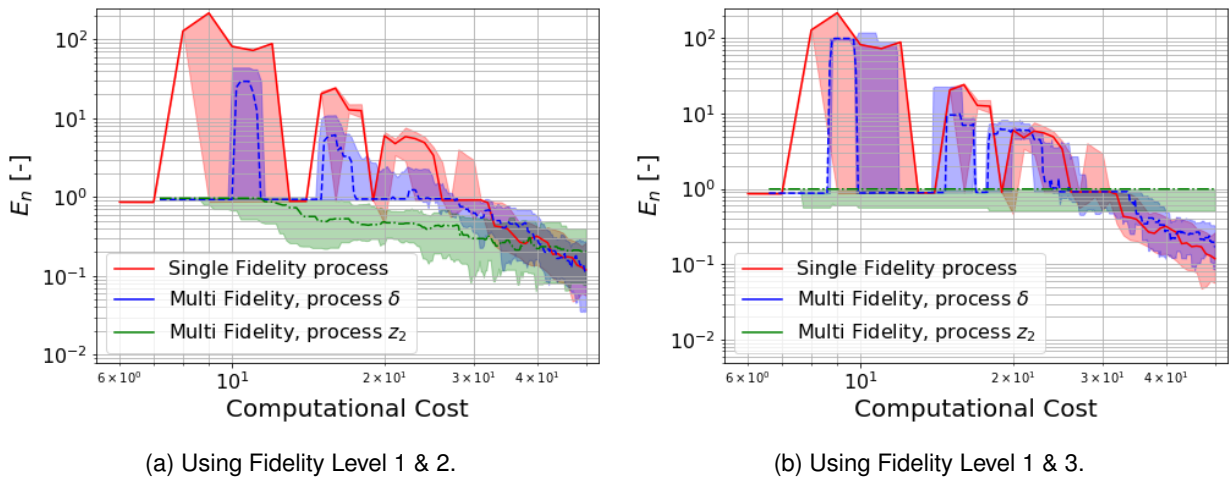




Figure 6: Convergence of E_n Using the 2D Rosenbrock Test Case (P2).



(a) Single Fidelity Result (Solid) and Multi-Fidelity Result Using Level 1 & 2 (Dashed). The Shaded Area Indicates the Interquartile Range.



(b) Single Fidelity Result (Solid) and Multi-Fidelity Result Using Level 1 & 3 (Dashed). The Shaded Area Indicates the Interquartile Range.

Figure 7: Convergence of E_x, E_f, E_t Using the 5D Rosenbrock Test Case (P2).

tational efficiency. For P_2 with $D = 5$ for example the simulation ends near the correlated phase but for P_2 with $D = 2$ the simulation ends in the noise inhibited phase. Clearly, this has an effect on how the errors of the multi-fidelity method compare to the errors of the single-fidelity method.

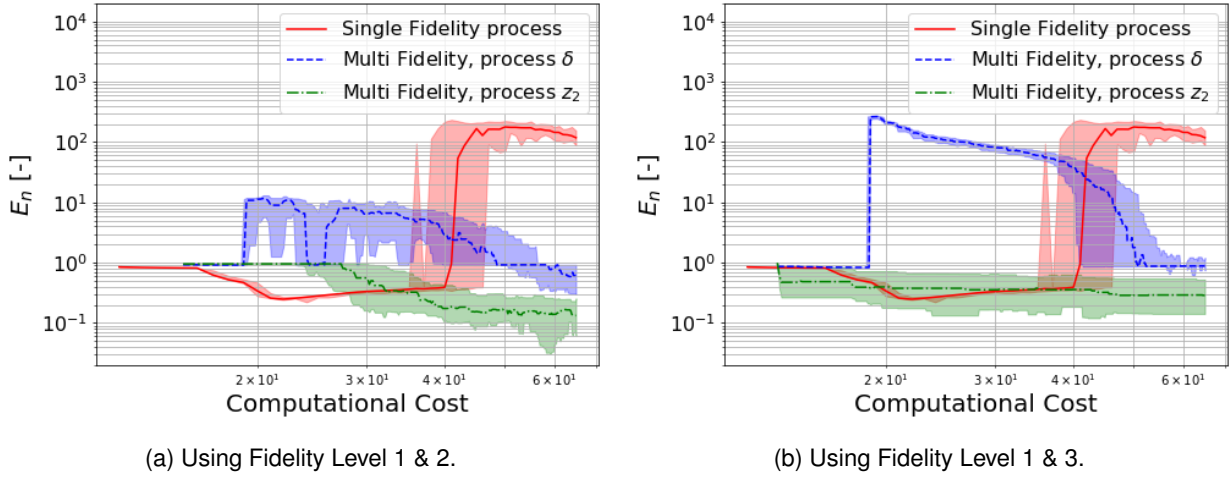


Figure 8: Convergence of E_n Using the 5D Rosenbrock Test Case (P2).

Table 6: Performance Comparisons at the Final Computational Cost.

Test	D	L	CC	median [E_x]%	median [E_f]%	median [E_t]%	median [E_n]%	median [J]
P_1	1	1	30	1.42	3.07	2.45	10.28	30
	1	1-2	30	0.54	2.12	1.59	10.22	27-15
	1	1-3	30	43.82	32.19	38.58	11.27	29-10
P_2	2	1	50	3.74	3.72E-3	2.64	11.85	50
	2	1-2	50	4.18	5.21E-3	2.96	10.85	43-35
	2	1-3	50	4.02	5.46E-3	2.84	17.17	50-9
	5	1	65	24.43	4.11	18.04	11791.72	65
	5	1-2	65	21.91	0.19	15.49	61.66	46-95
	5	1-3	65	21.72	0.22	15.36	88.44	63-26

4.0 DTMB 5415 HULL SHAPE OPTIMISATION

The DTMB 5415 model is used to demonstrate the method on an actual hull-shape optimisation problem. This is an open-to-public and early concept of the DDG-51, a USS Arleigh Burke-class destroyer.

4.1 Hull Shape Optimisation Using WARP

The model-scale resistance in calm water at $Fr = 0.28$ is taken as an optimisation objective. The optimisation problem is given by

$$\begin{aligned}
 &\text{Minimise} && R_T(\mathbf{x}) && \text{with } \mathbf{x} \in \mathbb{R}^D, \\
 &\text{subject to} && L_{pp}(\mathbf{x}) = L_{pp0}, \\
 &\text{and to} && \nabla(\mathbf{x}) = \nabla_0, \\
 &&& |\Delta B(\mathbf{x})| \leq 0.05B_0, \\
 &&& |\Delta T(\mathbf{x})| \leq 0.05T_0, \\
 &&& V(\mathbf{x}) \geq V_0, \\
 &&& x_i^l \leq x_i \leq x_i^u && i = 1, \dots, D,
 \end{aligned}$$

where L_{pp} denotes the length between perpendiculars, ∇ denotes the displacement, B denotes the Beam, T denotes the draft and V denotes the reserved volume of the sonar dome. A reduced design space is defined

using 14 orthogonal basis functions resulting from physics informed/augmented dimensionality reduction, see [25]. The modified hull shape is obtained from

$$g(\boldsymbol{\xi}, \mathbf{x}) = g_0(\boldsymbol{\xi}) + \sum_{i=1}^D x_i \psi_i(\boldsymbol{\xi}), \quad (21)$$

where $\boldsymbol{\xi}$ are geometrical coordinates, g_0 is the original geometry and ψ_i are orthonormal basis functions. The linear potential flow solver WARP, ([16]) is used to compute the resistance. The high-fidelity solve is obtained using 150x150 elements for the free-surface and 180x50 panels on the hull surface resulting in 16.5k elements. The low-fidelity solve is obtained by using 76x25 panels on the free-surface and 90x25 panels on the hull, resulting in 4.2k elements. For the work ratio, we have $c_1/c_2 = 14$. A CCF design of experiments is used for both the initial high- and low-fidelity sampling plans.

Since no analytical solution for this problem is available, the metrics used for this study are

$$\Delta_x = \frac{\|\mathbf{x}^* - \mathbf{x}_0\|}{\sqrt{D}} \quad \text{and} \quad \Delta_f = \frac{f(\mathbf{x}^*) - f(\mathbf{x}_0)}{f(\mathbf{x}_0)}. \quad (22)$$

The convergence results for the 2D problem are shown in Figure 9. Both the single-fidelity and multi-fidelity optimisations quickly find the minimum, but single-fidelity convergence is more erratic because it is less sure that the optimum is the actual global one. The multi-fidelity method, however, traded 4 high fidelity evaluations for 57 low-fidelity evaluations in the initial exploration phase, resulting in a more stable convergence.




Figure 9: Shape Optimisation on the 5415 Model with $D = 2$.

The convergence results for the 5D problem are shown in Figure 10. In this case the multi-fidelity method clearly outperforms the single-fidelity method. Because of very low expected improvement function values, the multi-fidelity method was stopped earlier at a budget of 45.96. At that time, 228 low-fidelity evaluations were done and 30 high-fidelity evaluation to reach a resistance reduction of 6.42%. The single-fidelity convergence however, behaves erratic and did not reach a stable phase while spending a budget of 350.

The results are summarised in Table 7. Relatively low noise levels are found for the single and multi-fidelity method but the estimated noise level for the 5-dimensional case is found to be higher than the estimated noise level for the 2-dimensional case. The cause of this phenomenon is unknown, but could be related to the convergence of the methods. The estimated single-fidelity noise for $D = 5$, for example, is about 10 times larger than for $D = 2$. Since the single-fidelity convergence for $D = 5$ is still in the uncorrelated phase, the noise level is typically large. The multi-fidelity methods converge faster, resulting in a smaller difference in estimated noise levels.

The waves and hull pressure of the optimised hull obtained with the multi-fidelity method with $D = 5$ ($\Delta f = -6.4\%$) is compared to the original hull in Figure 11. Compared to the original hull, the optimised hull has a softer forward shoulder and less pronounced stern waves.




Figure 10: Shape Optimisation on the 5415 Model with $D = 5$.

Table 7: Results of the 5415 Shape Optimisation Using WARP.

Case	D	L	CC	$\Delta_x\%$	$\Delta_f\%$	\mathcal{J}	$f_1(\mathbf{x}^*) [N]$	$\epsilon [N]$	x_1^*	x_2^*	x_3^*	x_4^*	x_5^*
Original hull	-	1	-	0.00	0.00	-	41.50	-	-	-	-	-	-
Single-fidelity	2	1	150	36.42	-3.18	150	40.19	3.34E-3	0.652	0.008	-	-	-
Multi-fidelity	2	1-2	150	36.96	-3.17	146-57	40.19	3.90E-3	0.664	0.002	-	-	-
Single-fidelity	5	1	350	66.06	+10.0	350	45.64	3.07E-2	0.624	0.742	0.805	0.568	0.521
Multi-fidelity	5	1-2	350	29.98	-6.42	30-228	38.84	7.05E-3	0.023	0.502	0.012	0.245	0.370

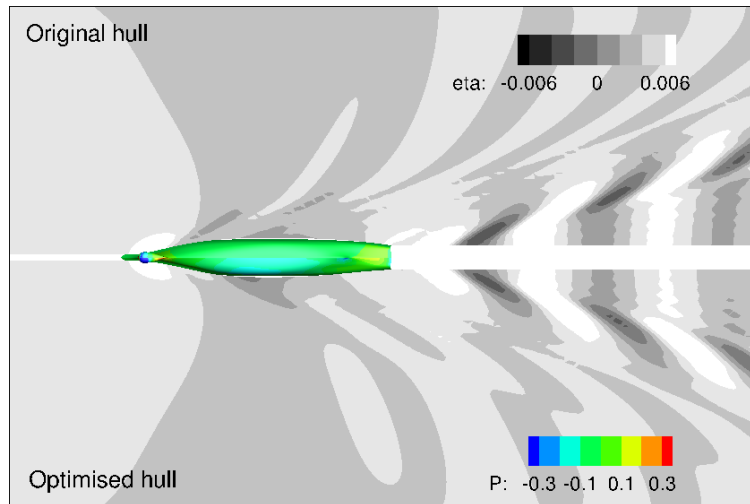


Figure 11: Flow around the original and optimised hull (multi-fidelity, $D = 5$, $\Delta f = -6.4\%$).

4.2 Hull Shape Optimisation Using ReFRESCO

A hull blending-technique is used to parameterise the hull shape. A small number of basis hull designs are chosen which define the scope of the deformations. Let H_0 denote the “original” hull shape. The other basis shapes are then denoted by H_1, H_2, \dots, H_D , each characterising a type of deformation. Each basis shape H_d is defined by N_c control points $\mathbf{c}_{d,i}$ of a B-spline surface. The parameterised control points are then obtained by linear interpolation on the control points of the basis designs:

$$\mathbf{c}_i(\mathbf{x}) = \mathbf{c}_i([x_1, x_2, \dots, x_N]) = \mathbf{c}_{0,i} + \sum_{d=1}^D x_d (\mathbf{c}_{d,i} - \mathbf{c}_{0,i}) \quad \text{for } i = 1, 2, \dots, N_c \quad (23)$$

Each linear interpolation coefficient represents a contribution of a certain shape deformation of the corresponding basis shape in the “hull-blend”. Using this definition, the design space is a N -dimensional unit

hypercube with the original hull design corresponding to $\mathbf{x} = \mathbf{0}$ and the other basis shapes corresponding to the corner points of this cube. The basis shapes used in this study are given in Figure 12.




Figure 12: Blending Parameterization Design Space Definition.

The resistance is computed using the RANS solver ReFRESCO (see, [17, 18]). The low-fidelity grid consists of approximately 1.7M cells, whereas the high-fidelity grid consists of approximately 15M cells, see Figure 13. The computational cost for the low-fidelity solution takes approximately 22 minutes using 240 cores on 10 nodes where each node has 2 Intel(R) Xeon(R) Gold 6126 CPU @ 2.60 Hz 12 core CPUs. Using the same resources, the high-fidelity computation takes about 165 minutes to complete (note that the computational cost includes the grid generation process). This results in a computational cost ratio $c_1/c_2 = 7.5$. A CCF design of experiments is used for both the initial high- and low-fidelity sampling plans. The rationale behind this choice is that it provides a relatively coarse initial sampling plan, leaving more budget for the adaptive phase of the optimisation method. A budget of 10 is given to the optimisation procedure.




Figure 13: Visualisation of the LF (Left) and HF (Right) Grids for ReFRESCO.

The optimisation results are collected in Table 8, where D denotes the dimensions of the design space, L the fidelity levels used, CC the computational budget, \mathcal{J} the number of evaluations, $f_1(\mathbf{x}^*)$ the resistance value at the optimum design \mathbf{x}^* , $Fric./Press.$ the friction and pressure resistance at the optimum, ϵ the estimated noise level and x_i^* the i^{th} element of the vector \mathbf{x}^* . The original hull has a full scale resistance of 43.06 N, where 31.21 N is due to friction and 11.84 N is due to pressure. The adaptive multi-fidelity method finds a slightly better design. A total resistance reduction is obtained of about 1.05% using the multi-fidelity method whereas the single-fidelity counterpart obtains 0.93%. The reduction of the pressure resistance is 3.72% and 3.21% respectively, while no significant reduction of the frictional resistance is observed. During the initial exploration phase of the multi-fidelity optimisation method, the algorithm decides to trade 2 high-fidelity evaluations for 18 low-fidelity evaluations (resulting in a final budget of 10.42).

The optimal hull shapes corresponding to Table 8 are shown in Figure 14. The optimal hull shapes are very similar but differ in the amount of V-shape applied to the aft-ship.

Table 8: Results of the 5415 Shape Optimisation Using ReFRESCO.

Case	D	L	CC	$\Delta_x\%$	$\Delta_f\%$	\mathcal{J}	$f_1(\mathbf{x}^*)$ [N]	Fric. [N]	Press. [N]	ϵ [N]	x_1^*	x_2^*	x_3^*
Original hull	-	1	-	0.00	0.00	-	43.06	31.21	11.84	-	0.000	0.000	0.000
Single-fidelity	3	1	10	65.82	-0.93	10	42.66	31.20	11.46	4.73E-7	0.266	1.000	0.478
Multi-fidelity	3	1-2	10	75.83	-1.05	8-18	42.61	31.20	11.40	4.90E-3	0.000	1.000	0.852




Figure 14: Optimal Hull Shapes Corresponding to Table 8.

Figure 15 shows the pressure coefficient on the 5415 hull, before and after the optimisation with the multi-fidelity method. It can be seen that the optimised hull has a better pressure recovery than the original hull. This results in a reduced pressure resistance and a small increase in the dynamic trim with 0.01 deg, bow down.




Figure 15: Pressure Around the 5415 Hull, Before and After the Optimisation.

5.0 CONCLUSIONS

An adaptive multi-fidelity fidelity Kriging method is implemented which is based on augmented expected improvement functions. Using the method of activation coefficients, it is shown how this active learning method can be applied in a practical way. A correct implementation of the augmented improvement functions is verified using a verification benchmark. The performance of the method is studied using analytical benchmarks and hull-shape optimisation problems.

Convergence of noise levels from Maximum Likelihood Estimation (MLE) is demonstrated using noise-free and noisy benchmarks. In the noise-free case, a low estimated noise level is observed which stagnates at a certain point during the optimisation. In the noisy case, a noise-inhibited convergence phase is observed.

From these experiments it is concluded that the adaptive multi-fidelity Kriging method results in a more stable convergence than the adaptive single fidelity Kriging method. Moreover, the errors from the multi-fidelity Kriging method are less sensitive to noise in the evaluations. A computational speed-up is not always achieved, but depends on the target error, noise levels, correlation and evaluation costs of the models. A significant speedup is observed for most cases in this study. Moreover, the speedup improves with increasing problem dimensions.

The method is applied to shape optimisation of the bare-hull DTMB 5415 frigate using the linear potential flow solver WARP and aft-ship optimisation using the RANS solver ReFRESKO with the objective to minimise resistance. In both cases it is beneficial to use the adaptive multi-fidelity Kriging method, given the problem setup. In case of the optimisation with WARP, we observe a more stable convergence when using the multi-fidelity method. In case of the aft-ship optimisation with the solver ReFRESKO, a better shape is found for a very tight computational budget equal to the cost of only 10 high-fidelity evaluations in a 3 dimensional design space. This multi-fidelity speedup is obtained using 8 high fidelity evaluations and 18 low-fidelity evaluations. Hence, for a fixed computational budget, a better ship can be designed by combining low- and high-fidelity computations during adaptive training than would otherwise be obtained by using high-fidelity computations only.

6.0 ACKNOWLEDGEMENTS

The present research is conducted in collaboration with “AVT-331 task group on Goal-driven, multi-fidelity approaches for military vehicle system-level design.” This research is funded by the Dutch Ministry of Economic Affairs.

7.0 REFERENCES

- [1] Huang, F., Wang, L., and Yang, C., “Hull Form Optimization for Reduced Drag and Improved Seakeeping Using a Surrogate-Based Method”, International Ocean and Polar Engineering Conference, June 2015. ISOPE-I-15-846.
- [2] Scholcz, T. P. and van Daalen, E. F. G., “Surrogate-Based Multi-Objective Optimisation for Powering and Seakeeping”, In: Okada, T., Suzuki, K., and Kawamura, Y. (editors), Practical Design of Ships and Other Floating Structures, pp. 477-492, Springer Singapore, Singapore, 2021.
- [3] Rotteveel, E., van der Ploeg, A., and Hekkenberg, R., “Optimization of ships in shallow water with viscous flow computations and surrogate modeling”, In: Nielsen, U. and Jensen et al, J. (editors), Proceedings of the 13th International Symposium on Practical Design of Ships and Other Floating Structures, pp. 1 - 8, PRADS Organising Committee, 2016.
- [4] Scholcz, T. P., Gornicz, T., and Veldhuis, C., “Multi-Objective Hull-Form Optimization Using Kriging on Noisy Computer Experiments”, MARINE VI: Proceedings of the VI International Conference on Computational Methods in Marine Engineering, pp. 1064-1077, 2015.
- [5] de Baar, J., “Stochastic Surrogates for Measurements and Computer Models of Fluids”, PhD thesis, Delft University of Technology, 2014.
- [6] de Baar, J., Roberts, S., Dwight, R., and Mallol, B., “Uncertainty Quantification for a Sailing Yacht Hull, Using Multi-Fidelity Kriging”, Computers & Fluids, Vol. 123, pp. 185-201, December 2015. doi:10.1016/j.compfluid.2015.10.004.
- [7] Gaggero, S., Vernengo, G., and Villa, D., “A Marine Propeller Design Method Based on Two-Fidelity Data Levels”, Applied Ocean Research, Vol. 123, 2022.

- [8] Scholcz, T. P., “Data driven uncertainty quantification for computational fluid dynamics based ship design”, Proceedings of the 8th International Conference on Computational Methods in Marine Engineering, MARINE 2019, pp. 309 - 320, 2019. 13-15 May 2019, Göteborg, CIMNE.
- [9] Raven, H. C. and Scholcz, T. P., “An assessment of multifidelity procedures for ship hull form optimisation”, Proceedings of the 8th International Conference on Computational Methods in Marine Engineering, MARINE 2019, pp. 189 - 200, 2019. 13-15 May 2019, Göteborg, CIMNE.
- [10] Bonfiglio, L., Perdikaris, P., Brizzolara, S., and Karniadakis, G. E., “Multi-Fidelity Optimization of Super-Cavitating Hydrofoils”, Computer Methods in Applied Mechanics and Engineering, Vol. 332, pp. 63-85, 2018.
- [11] Serani, A., Pellegrini, R., Broglia, R., Wackers, J., Visonneau, M., and Diez, M., “An Adaptive n - Fidelity Metamodel for Design and Operational-Uncertainty Space Exploration of Complex Industrial Problems”, Proceedings of the 8th International Conference on Computational Methods in Marine Engineering, MARINE 2019, 2019.
- [12] Liu, X., Zhao, W., and Wan, D., “Multi-Fidelity Co-Kriging Surrogate Model for Ship Hull Form Optimization”, Ocean Engineering, Vol. 243, p. 110239, 2022.
- [13] Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A., “Sequential Kriging Optimization Using Multiple-Fidelity Evaluations”, Structural and Multidisciplinary Optimization, Vol. 32, No. 5, pp. 369-382, 2006. doi:10.1007/s00158-005-0587-0.
- [14] Grassi, F., Manganini, G., Garraffa, M., and Mainini, L., “Resource Aware Multifidelity Active Learning for Efficient Optimization”, ArXiv, 2020. doi:10.48550/arXiv.2007.04674.
- [15] Zhang, Y., Han, Z.-H., and Zhang, K.-S., “Variable-Fidelity Expected Improvement Method for Efficient Global Optimization of Expensive Functions”, Structural and Multidisciplinary Optimization, Vol. 58, pp. 1431-1451, 2018.
- [16] Bassanini, P., Bulgarelli, U., Campana, E. F., and Lalli, F., “The Wave Resistance Problem in a Boundary Integral Formulation”, Surveys on Mathematics for Industry, Vol. 4, pp. 151-194, 1994.
- [17] Vaz, G., Jaouen, F., and Hoekstra, M., “Free-Surface Viscous Flow Computations: Validation of URANS Code FreSCo”, Proceedings of the 28th International Conference on Offshore Mechanics and Arctic Engineering, pp. 425-437, 2009. doi:10.1115/OMAE2009-79398.
- [18] Klaij, C. M. and Vuik, C., “SIMPLE-type preconditioners for cell-centered, colocated finite volume discretization of incompressible Reynolds-averaged Navier–Stokes equations”, International Journal for Numerical Methods in Fluids, Vol. 71, No. 7, pp. 830-849, 2013. doi:https://doi.org/10.1002/flid.3686.
- [19] Kennedy, M. and O’Hagan, A., “Predicting the Output From a Complex Computer Code When Fast Approximations Are Available”, Biometrika, Vol. 87, pp. 1-13, 2000.
- [20] de Baar, J. H. S., Leylek, Z., and Neely, A. J., “Efficient Uncertainty Quantification for an Axial Compressor, Using Adaptive Multi-Fidelity Kriging”, 23rd International Symposium on Air Breathing Engines (ISABE 2017): Economy, Efficiency and Environment, Volume 2, September 2017.
- [21] Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C., “A Limited Memory Algorithm for Bound Constrained Optimization”, SIAM J. Sci. Comput., Vol. 16, pp. 1190-1208, 1995.
- [22] Sasena, M. J., “Flexibility and Efficiency Enhancements for Constrained Global Design Optimization with Kriging Approximations”, PhD thesis, University of Michigan, 2012.

- [23] Mainini, L., Serani, A., Rumpfkeil, M. P., Minisci, E., Quagliarella, D., Pehlivan, H., Yildiz, S., Ficini, S., Pellegrini, R., Di Fiore, F., Bryson, D., Nikbay, M., Diez, M., and Beran, P., “Analytical Benchmark Problems for Multifidelity Optimization Methods”, arXiv, 2022. doi:10.48550/arXiv.2204.07867.
- [24] Pellegrini, R., Wackers, J., Broglia, R., Diez, M., Serani, A., and Visonneau, M., “A Multi-Fidelity Active Learning Method for Global Design Optimization Problems with Noisy Evaluations”, arXiv, 2022. doi:10.48550/arXiv.2202.06902.
- [25] Serani, A., Diez, M., Wackers, J., Visonneau, M., and Stern, F., “Stochastic Shape Optimization via Design-Space Augmented Dimensionality Reduction and RANS Computations”, AIAA Scitech 2019 Forum, p. 2218, 2019.