



HAL
open science

Adaptive multi-fidelity sampling for CFD-based optimisation via radial basis function metamodels

Andrea Serani, Riccardo Pellegrini, Jeroen Wackers, Charles-Edouard Jeanson, Patrick Queutey, Michel Visonneau, Matteo Diez

► To cite this version:

Andrea Serani, Riccardo Pellegrini, Jeroen Wackers, Charles-Edouard Jeanson, Patrick Queutey, et al.. Adaptive multi-fidelity sampling for CFD-based optimisation via radial basis function metamodels. *International Journal of Computational Fluid Dynamics*, 2019, 33 (6-7), pp.1-19. 10.1080/10618562.2019.1683164 . hal-02391342

HAL Id: hal-02391342

<https://hal.science/hal-02391342>

Submitted on 28 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Adaptive multi-fidelity sampling for CFD-based optimisation via radial basis function metamodels

A. Serani^a, R. Pellegrini^a, J. Wackers^b, C.-E. Jeanson^b, P. Queutey^b, M. Visonneau^b and M. Diez^a

^aCNR-INM, National Research Council-Institute of Marine Engineering, Rome, Italy

^bLHEEA Lab, Ecole Centrale de Nantes, Nantes, France

The paper presents a study on four adaptive sampling methods of a multi-fidelity (MF) metamodel, based on stochastic radial basis functions (RBF), for global design optimisation based on expensive CFD computer simulations and adaptive grid refinement. The MF metamodel is built as the sum of a low-fidelity-trained metamodel and an error metamodel, based on the difference between high- and low-fidelity simulations. The MF metamodel is adaptively refined using dynamic sampling criteria, based on the prediction uncertainty in combination with the objective optimum and the computational cost of high- and low-fidelity evaluations. The adaptive sampling methods are demonstrated by four analytical benchmark and two design optimisation problems, pertaining to the resistance reduction of a NACA hydrofoil and a destroyer-type vessel. The performance of the adaptive sampling methods is assessed via objective function convergence.

1. Introduction

Fluid-dynamic shape design of complex industrial systems like aerial, ground, and water-borne vehicles demands the use of high-fidelity numerical solvers with large computational grids to assess accurately the design performance and make sound design decisions. The latter can be achieved by combining computational fluid dynamics (CFD) analysis with a shape/design modification tool (CAD) and a minimisation algorithm into an automatic simulation-based design optimisation (SBDO). The optimisation algorithm may require a large number of function evaluations to converge to the final solution, especially if a global optimum is desired. Therefore, the resulting computational cost of the SBDO process could become very high, making SBDO hardly affordable for most users and projects, for which computational resources and time are usually limited.

To reduce the computational cost of the SBDO process, metamodeling methods have been developed and successfully applied in several engineering field (Viana et al. 2014). Among other metamodels, radial basis functions (RBF) methods have demonstrated their accuracy and efficiency in engineering design (Jin, Chen, and Simpson 2001) along

with their ease of implementation. The performance of metamodels is problem-dependent and determined by several concurrent issues, such as the presence of nonlinearities, the problem dimensionality, the noisy or smooth behaviour of the function, and the approach used for its training (Liu, Ong, and Cai 2018).

The research of accurate and efficient methods for metamodel-based analysis and optimisation has recently moved from standard (or static) to function-adaptive approaches, also known as dynamic metamodels. In addition to auto tuning, a dynamic metamodel is able to improve its fitting capability by adaptive sampling or training. The design of experiments (DoE) used for metamodel training is not defined a priori but dynamically updated, exploiting the information that becomes available during the analysis process. The purpose of performing an adaptive DoE is to add training points anywhere it is most useful, so as to use a relatively low number of function evaluations to represent the function. For global exploration of design/uncertainty spaces (such as in global optimisation and uncertainty quantification), a dynamic RBF (DRBF) has been formulated by Volpi et al. (2015), based on the maximum prediction

uncertainty, and compared to dynamic Kriging (Zhao, Choi, and Lee 2011). Similarly, the non-linearity of the response and the largest distance among training-set samples have been used by Mackman and Allen (2010) for RBF. To improve the efficiency of global optimisation procedures, sequential and parallel infill criteria have been used balancing the exploration and exploitation of the design space: the expected improvement (Jones, Schonlau, and Welch 1998) has been proposed to update Kriging model; multiple points have been added on Kriging model based on approximate computation of the probability of improvement (Viana and Haftka 2010); a multi-objective infill criteria based on the estimated variance and the response value has been proposed by Yi, Kwon, and Choi (2014); similarly, a parallel sampling criterion for DRBF has been proposed by Diez et al. (2019) based on the identification of the global minimum of the objective function and maximum of the prediction uncertainty.

In addition to dynamic metamodels, multi-fidelity (MF, or variable-fidelity) approximation methods have been developed with the aim of combining the accuracy of high-fidelity solvers/evaluations with the computational cost of low-fidelity solvers/evaluations. Thus, MF metamodels are trained with a combination of high-fidelity (accurate, expensive) simulations and low-fidelity (approximate, less expensive) simulations. Combining metamodeling methods with MF approximations potentially leads to a further reduction of the computational cost of the SBDO procedure. Additive and/or multiplicative correction methods, also known as ‘bridge functions’ or ‘scaling functions’ (Han, Görtz, and Zimmermann 2013), are used to build MF metamodels, using high- and low-fidelity evaluations (Ng and Eldred 2012; Zheng et al. 2013; Pellegrini et al. 2016). Several metamodels have been used in the literature with MF data, such as non-intrusive polynomial chaos (Ng and Eldred 2012), co-kriging (Baar et al. 2015) and RBF (Pellegrini et al. 2016). In SBDO based on CFD computations, high- and low-fidelity evaluations may be obtained by varying the physical model, the size of the computational grid, the computational time step, the convergence level of the simulations, and/or combining experimental data with numerical simulations (Kuya et al. 2011).

Adaptive sampling methods have been recently combined with MF metamodels: Huang et al. (2006)

propose an augmented expected improvement for multi-fidelity Kriging; Pellegrini et al. (2016) use a DRBF, sequentially adding low-fidelity or both high- and low-fidelity samples, depending on which fidelity has the greatest prediction uncertainty and on a parameter based on the computational cost; Liu et al. (2016) propose a sampling method based on the propagation of the low-fidelity variance to the MF metamodel, updating only the fidelity that provides the maximum reduction of the MF uncertainty; Cai et al. (2017) sample both high- and low-fidelity at each iteration, based on the cross-validation error and a Voronoi partition.

The objective of the present work is to assess the performance of a multi-fidelity DRBF (MF-DRBF) metamodel for global optimisation, conditional to four adaptive sampling methods. These are based on: (i) the maximum prediction uncertainty, (ii) a multi-fidelity version of the expected improvement, (iii) the maximum prediction uncertainty and the objective function through an aggregated merit factor, (iv) multi-criteria sampling based on the solution of a multi-objective problem considering the maximum prediction uncertainty and the objective function value (Diez et al. 2019).

The sampling methods for the MF-DRBF metamodel are demonstrated by four multi-modal analytical benchmark problems with one and two dimensions. Their performance is assessed by studying the convergence of the optimisation procedure and the normalised root mean square error of the prediction. Finally, the sampling methods are applied to two CFD-based design problems pertaining to the minimisation of (a) the drag coefficient of a four-digit NACA hydrofoil at Reynolds number equal to $8.41 \cdot 10^6$ and (b) the total resistance of the DTMB 5415, an early and open-to-public version of the DDG-51 (a USS Arleigh Burke-class destroyer-type vessel), in calm water at Froude number equal to 0.3. CFD simulations are based on the Reynolds-averaged Navier-Stokes equation solver ISIS-CFD, developed at Ecole Centrale de Nantes/CNRS and integrated in the FINE/Marine simulation suite from NUMECA Int. The CFD-solver fidelity is varied by using two computational grid levels, defined by an adaptive grid refinement technique (Wackers et al. 2017). Optimisations are performed with a deterministic version of the particle swarm optimisation (DPSO, Serani et al. 2016b) algorithm.

2. Stochastic radial basis functions

Consider an objective function $f(\mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^N$ is the design variable vector and N the design space dimension. Let the true function value be known in a number J of training points \mathbf{x}_j with associated objective function values $f(\mathbf{x}_j)$. The metamodel prediction $\tilde{f}(\mathbf{x})$ is computed as the expected value (EV) of a prediction ensemble obtained considering a stochastic tuning parameter in the RBF kernel, e.g. $\tau \sim \text{unif}[1, 3]$:

$$\tilde{f}(\mathbf{x}) = \text{EV} [g(\mathbf{x}, \tau)]_\tau, \quad \text{with}$$

$$g(\mathbf{x}, \tau) = \sum_{j=1}^J w_j \|\mathbf{x} - \mathbf{x}_j\|^\tau, \quad (1)$$

where w_j are unknown coefficients, $\|\cdot\|$ is the Euclidean norm. The coefficients w_j are determined enforcing exact interpolation at the training points $g(\mathbf{x}_j, \tau) = f(\mathbf{x}_j)$ by solving $\mathbf{A}\mathbf{w} = \mathbf{f}$, with $\mathbf{w} = \{w_j\}$, $a_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|^\tau$ and $\mathbf{f} = \{f(\mathbf{x}_j)\}$.

The uncertainty $U_{\tilde{f}}(\mathbf{x})$ associated with the stochastic RBF metamodel prediction is quantified by the 95%-confidence interval of $g(\mathbf{x}, \tau)$, evaluated using a Monte Carlo sampling over τ (Volpi et al. 2015).

3. Multi-fidelity metamodel

Similarly to Kennedy and O'Hagan (2000), the MF prediction $\hat{f}(\mathbf{x})$ is defined by an additive correction to a low-fidelity trained metamodel $\tilde{f}_L(\mathbf{x})$ as

$$\hat{f}(\mathbf{x}) = \tilde{f}_L(\mathbf{x}) + \tilde{\varepsilon}(\mathbf{x}), \quad (2)$$

where the correction is provided by the metamodel of the error (or discrepancy) $\tilde{\varepsilon}(\mathbf{x})$, defined as the difference between high- (HF) and low-fidelity (LF) evaluations (f_H and f_L):

$$\varepsilon(\mathbf{x}) = f_H(\mathbf{x}) - f_L(\mathbf{x}). \quad (3)$$

The training set for \tilde{f}_L is denoted by \mathcal{L} , whereas the training set for $\tilde{\varepsilon}$ is denoted by $\mathcal{E} \subseteq \mathcal{L}$. Assuming (for the sake of simplicity and ease of implementation) that the uncertainties associated with the low-fidelity and error metamodels ($U_{\tilde{f}_L}$ and $U_{\tilde{\varepsilon}}$ respectively) are uncorrelated, the uncertainty of the MF prediction is

$$U_{\hat{f}}(\mathbf{x}) = \sqrt{U_{\tilde{f}_L}^2(\mathbf{x}) + U_{\tilde{\varepsilon}}^2(\mathbf{x})}. \quad (4)$$

A sketch of the MF metamodel concept and notation is shown in Figure 1. Note that ' \wedge ' indicates MF approximation and ' \sim ' denotes the RBF prediction.

4. Adaptive sampling methods

The MF metamodel is dynamically updated by adding new training points following a two-step procedure (see Figure 2): (1) the coordinates of the new training point \mathbf{x}^* are chosen using one of the adaptive sampling methods described in the following subsections; (2) once \mathbf{x}^* is identified, either \mathcal{L} or \mathcal{E} are refined, based on the following statement

$$\text{If } U_{\tilde{f}_L}(\mathbf{x}^*) \geq \beta U_{\tilde{\varepsilon}}(\mathbf{x}^*), \quad \text{add } \{\mathbf{x}^*, f_L(\mathbf{x}^*)\} \text{ to } \mathcal{L},$$

$$\text{else, add } \{\mathbf{x}^*, f_L(\mathbf{x}^*)\} \text{ to } \mathcal{L} \quad \text{and} \quad \{\mathbf{x}^*, \varepsilon(\mathbf{x}^*)\} \text{ to } \mathcal{E}, \quad (5)$$

where $\beta \in [0, 1)$ is the ratio between the LF and HF computational costs. In the first case, only a low-fidelity evaluation is performed, whereas the second case requires both low- and high-fidelity evaluations at the same point \mathbf{x}^* .

4.1. Maximum uncertainty

Maximum-uncertainty adaptive sampling (MUAS) has been discussed for multi-fidelity problems in Pellegrini et al. (2016). This method identifies a new training point (see Figure 3) by solving the single-objective maximisation problem

$$\mathbf{x}^* = \underset{\mathbf{x}}{\text{argmax}} [U_{\hat{f}}(\mathbf{x})]. \quad (6)$$

4.2. Multi-fidelity expected improvement

An extension of the expected improvement (Jones, Schonlau, and Welch 1998) for multi-fidelity RBF applications (MFEI) has been proposed by Pellegrini et al. (2018). It is defined as

$$\text{MFEI}(\mathbf{x}) = \text{EV} \left[\max(f_{\min} - g_L(\mathbf{x}, \tau) + g_\varepsilon(\mathbf{x}, \tau), 0) \right]_\tau,$$

$$\text{with } f_{\min} = \min[\hat{f}(\mathbf{x})], \quad (7)$$

where $g_L(\mathbf{x}, \tau)$ and $g_\varepsilon(\mathbf{x}, \tau)$ are RBF interpolation of f_L and ε , following Eq. (1). Therefore, the MFEI is the expected value of potential reduction considering stochastic multi-fidelity RBF and minima provided by \hat{f} at the current iteration of the sampling process.

The MFEI adaptive sampling identifies a new training point (see Figure 4) by solving the single-objective maximisation problem

$$\mathbf{x}^* = \underset{\mathbf{x}}{\text{argmax}} [\text{MFEI}(\mathbf{x})]. \quad (8)$$




Figure 1. Multi-fidelity metamodel concept and notation.




Figure 2. Updating scheme for the adaptive multi-fidelity metamodel.




Figure 3. MUAS sampling strategy.




Figure 4. MFEI sampling strategy.




Figure 5. ACAS sampling strategy.

4.3. Aggregate criteria

An aggregate-criteria adaptive sampling (ACAS) method is defined with the aim of finding points with large uncertainty and small objective function value. Accordingly, ACAS identifies a new training point (see Figure 5) by solving the single-objective minimisation

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} [\hat{f}(\mathbf{x}) - U_{\hat{f}}(\mathbf{x})]. \quad (9)$$

4.4. Multi-criteria

Similarly to the aggregate criteria, a multi-criteria adaptive sampling (MCAS, Yi, Kwon, and Choi 2014; Diez et al. 2019) identifies N_p new training points (see Figure 6) considering both the MF prediction and the associated uncertainty (multi-criteria infill sampling). Specifically, a multi-objective optimisation problem addressing the minimisation of the objective function and the maximisation of the prediction uncertainty is solved

$$\begin{aligned} & \text{minimise } \hat{f}(\mathbf{x}) \quad \text{and maximise } U_{\hat{f}}(\mathbf{x}), \\ & \text{subject to } U_{\hat{f}}(\mathbf{x}) > U_{\hat{f}}^*, \end{aligned} \quad (10)$$

where $U_{\hat{f}}^* = \gamma R$, with $R = \sup[f_H(\mathbf{x})] - \inf[f_H(\mathbf{x})] < \infty$, is a constraint parameter, defined in the view of the fact that: (a) sampling too close to available training points does not add useful information to the analysis, (b) as the distance between training points decreases, the matrix \mathbf{A} may become ill-conditioned, and (c) the uncertainty at the training points is zero, i.e.

$$\lim_{|\mathbf{x} - \mathbf{x}^*| \rightarrow 0} U_{\hat{f}}(\mathbf{x}) = 0 \quad (11)$$

The non-dominated solution set obtained by the solution of Equation (10) is down-sampled in order to identify N_p equally spaced points along the curvilinear




Figure 6. MCAS down-sampling strategy: green zone is defined to avoid over-fitting in the neighbourhood of the minimum.

coordinate defined by the Pareto frontier (see Figures 6 and 7). N_p depends on the number of non-dominated solutions and typical values go from 2 to 10, whereas γ is generally set equal to $(0.1 \div 1)\%R$.

5. Optimisation problems

The assessment of the adaptive sampling methods for the MF-DRBF is based on four analytical benchmarks and two CFD-based design optimisation problems. A deterministic single-objective formulation of the particle swarm optimisation (DPSO) algorithm (Serani et al. 2016b), is used for the metamodel-based optimisations, as well as for the solution of the minimisation/maximisation sampling problems of Equations (6), (8), and (9). A multi-objective extension of DPSO as presented by Pellegrini et al. (2017) is used for the solution of Equation (10).

5.1. Analytical benchmark problems

Four analytical benchmark problems (SK1, SSFYY2, MLF1, Far1) are selected from Huband et al. (2006). These problems have dimensionality ranging from one




Figure 7. MCAS sampling strategy.

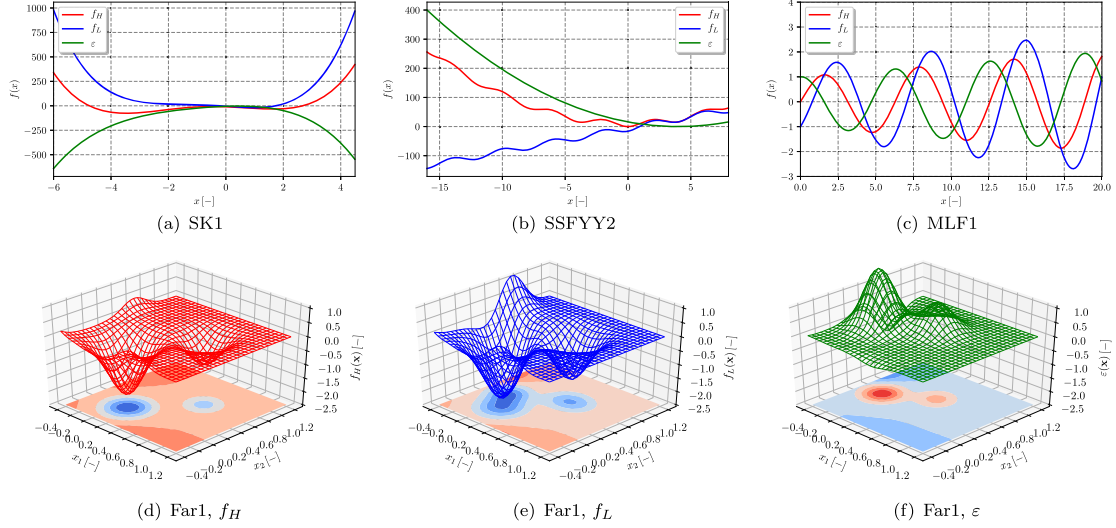


Figure 8. Analytical benchmark problems.

to two and contain two objective functions. To simulate a MF environment, the first objective function is considered as the high-fidelity objective (f_H), whereas the second objective is considered as the error function (ε). Therefore, the low-fidelity evaluation is computed as $f_L = f_H - \varepsilon$. The analytical benchmark problems are shown in Figure 8 and details are provided

in Table 1. Despite their low dimensionality, these benchmark problems are challenging for the construction of both single- and multi-fidelity metamodels, since they are generally multimodal with several minima. Furthermore, also the error functions are multimodal. Moreover, high- and low-fidelity functions are defined to show quite different behaviours, as could

Table 1. Analytical benchmark problems.

Name	Problem	Domain
SK1	$f_H(x) = x^4 + 3x^3 - 10x^2 - 10x - 10$ $\varepsilon(x) = -0.5x^4 - 2x^3 - 10x^2 + 10x - 5$	$[-6, 4.5]$
SSFYY2	$f_H(x) = 10 + x^2 - 10 \cos(x\pi/2)$ $\varepsilon(x) = (x - 4)^2$	$[-16, 8]$
MLF1	$f_H(x) = (1 + x/20) \sin(x)$ $\varepsilon(x) = (1 + x/20) \cos(x)$	$[0, 20]$
Far1	$f_H(\mathbf{x}) = -2e^{(15-(x_1-0.1)^2-x_2^2)} - e^{(20-(x_1-0.6)^2-(x_2-0.6)^2)}$ $+ e^{(20-(x_1+0.6)^2-(x_2-0.6)^2)} + e^{(20-(x_1-0.6)^2-(x_2+0.6)^2)}$ $+ e^{(20-(x_1+0.6)^2-(x_2+0.6)^2)}$ $\varepsilon(\mathbf{x}) = 2e^{(20-(x_1+0.1)^2-(x_2-0.3)^2)} + e^{(20-(x_1-0.4)^2-(x_2-0.6)^2)}$ $- e^{(20-(x_1+0.5)^2-(x_2-0.7)^2)} - e^{(20-(x_1-0.5)^2-(x_2+0.7)^2)}$ $+ e^{(20-(x_1+0.4)^2-(x_2+0.8)^2)}$	$[-0.4, 1.2]$




Figure 9. NACA 4-digit hydrofoil.

(moderately) happen when different physical models (such as potential flow and RANS) are used to solve the same fluid dynamic problem.

5.2. CFD-based design optimisation problems

Two CFD-based design optimisation problems are considered, namely the drag coefficient reduction of a NACA hydrofoil and the total resistance reduction of the DTMB 5415 model (a destroyer-type vessel). Details of the optimisation problems are provided in the following subsections, whereas details on the CFD solver and simulations are provided in Section 6.

5.2.1. NACA hydrofoil

The following minimisation problem is solved

$$\begin{aligned} & \text{minimise} && C_D(\mathbf{x}) \\ & \text{subject to} && C_L(\mathbf{x}) = 0.6 \\ & && \text{and to } \mathbf{l} \leq \mathbf{x} \leq \mathbf{u} \end{aligned} \quad (12)$$

where \mathbf{x} is the design variable vector, C_D and C_L are respectively the drag and lift coefficient of a four-digit NACA hydrofoil, and \mathbf{l} and \mathbf{u} are the lower and upper bound of \mathbf{x} . The equality constraint on the lift coefficient is necessary for the optimisation of lifting hydrofoils, in order to compare different geometries at the same lift force, rather than the same angle of attack. The goal of the foil is to create a specific lift (typically equal to the weight of the object), whereas the drag depends strongly on the lift and the foil geometry. Section 6.3.1 explains how to handle the equality constraint on the lift coefficient.

The hydrofoil shape (see Figure 9) is defined by the general equation for four-digit NACA foils. The upper (y_u) and lower (y_l) hydrofoil surfaces are computed as

$$\begin{aligned} \xi_u &= \xi - y_t \sin \theta \\ \xi_l &= \xi + y_t \sin \theta \\ y_u &= y_c + y_t \cos \theta \end{aligned}$$

$$\begin{aligned} y_l &= y_c - y_t \cos \theta \quad \text{with} \\ y_c &= \begin{cases} \frac{m}{p^2} \left[2p \frac{\xi}{c} - \left(\frac{\xi}{c} \right)^2 \right], & 0 \leq \xi < pc \\ \frac{m}{(1-p)^2} \\ \times \left[(1-2p) + 2p \frac{\xi}{c} - \left(\frac{\xi}{c} \right)^2 \right], & pc \leq \xi \leq c \end{cases} \end{aligned} \quad (13)$$

where ξ is the position along the chord, c the chord length, y_c the mean camber line, p the location of the maximum camber, m the maximum camber value, t the maximum thickness, and y_t the half thickness distribution given by

$$\begin{aligned} y_t &= 5t \left(0.2969\sqrt{\xi} - 0.1260\xi \right. \\ &\quad \left. - 0.3516\xi^2 + 0.2843\xi^3 - 0.1015\xi^4 \right) \end{aligned} \quad (14)$$

In this work, the design variables vector is defined as $\mathbf{x} = \{t, m\}$ with $t \in [0.030, 0.120]$ and $m \in [0.025, 0.070]$. The maximum camber position is fixed at $p = 0.4$. The simulation conditions are: velocity $U = 10$ m/s, chord $c = 1$ m, fluid density $\rho = 1,026$ kg/m³, and Reynolds number $Re = 8.41 \cdot 10^6$ based on the chord length.

5.2.2. DTMB 5415 model

The DTMB 5415 model (see Figure 10) is an open-to-public early concept of a USS Arleigh Burke-class destroyer, widely used for towing tank experiments (Irvine, Longo, and Stern 2008), CFD studies (Stern et al. 2001) and hull-form optimisation (Serani et al. 2016a).

The following minimisation problem is solved

$$\begin{aligned} & \text{minimise} && R_T(\mathbf{x}) \\ & \text{subject to} && \mathbf{l} \leq \mathbf{x} \leq \mathbf{u} \end{aligned} \quad (15)$$

where R_T is the calm water total resistance for free-surface flow at $Fr = 0.30$ and $Re = 1.18 \cdot 10^7$ based on




Figure 10. A geosym replica of the DTMB 5415 (CNR-INSEAN model 2340).

the length between perpendiculars, with the ship fixed in its design position (even keel).

The modified hull designs (\mathbf{g}) are produced by the linear superposition of N orthonormal basis functions ($\boldsymbol{\psi}$) on the original geometry (\mathbf{g}_0):

$$\begin{aligned} \mathbf{g}(\boldsymbol{\xi}, \mathbf{x}) &= \mathbf{g}_0(\boldsymbol{\xi}) + \delta(\boldsymbol{\xi}, \mathbf{x}), \quad \text{with} \\ \delta(\boldsymbol{\xi}, \mathbf{x}) &= \sum_{k=1}^N x_k \boldsymbol{\psi}_k(\boldsymbol{\xi}), \end{aligned} \quad (16)$$

where $\boldsymbol{\xi}$ are Cartesian coordinates, whereas $\{x_k\}_{k=1}^N$ and $\{\boldsymbol{\psi}_k\}_{k=1}^N$ are the design variables and the basis functions, provided by a design-space augmented dimensionality reduction (ADR) procedure described in Serani and Diez (2018). Herein, for the sake of simplicity and demonstration, $N = 2$ design variables are used with $x_k \in [-1.25, 1.25]$.

6. CFD solver and problem setups

CFD simulations are performed with the unstructured finite-volume Navier-Stokes solver ISIS-CFD developed at ECN – CNRS (Queutey and Visonneau 2007), available in the FINETM/Marine computing suite from NUMECA Int. Computational grids are created through adaptive grid refinement and mesh deformation to take into account the need for high and low fidelity, as well as the different geometries needed for shape optimisation. The following subsection describes the treatment of computational grids for ISIS-CFD and some methods necessary for both the CFD-based design optimisation problems: (a) hydrofoil dynamic positioning; (b) free-surface meshing improvement and (c) limiting refinement behind the stern for the vessel.

6.1. Adaptive grid refinement

The adaptive grid refinement method adjusts the computational grid locally, during the computation, by dividing the cells of an original coarse grid, to

improve the precision (Wackers et al. 2014, 2017). This method performs isotropic and anisotropic refinement of unstructured hexahedral meshes, by dividing the cells of the original grid into fine cells; this division can be repeated several times until the desired cell sizes are obtained. The decision where to refine comes from a refinement criterion, a tensor field $\mathcal{C}(x, y, z)$ computed from the flow. The tensor is based on the water surface position and on second derivatives of pressure and velocity. The mesh is refined until the dimensions $\mathbf{d}_{i,j}$ ($j = 1, 2, 3$) of each hexahedral cell i satisfy

$$\|\mathcal{C}_i \mathbf{d}_{i,j}\| = T_r \quad (17)$$

As shown by Wackers et al. (2017), the refinement criterion based on the second derivatives of the flow is not very sensitive to grid refinement, so the cell sizes everywhere are proportional to the constant threshold T_r . The cost of the mesh adaptation is moderate; the added computation time for refinement is less than 5% for a typical simulation.

For multi-fidelity optimisation, the interest of this procedure is that high- and low-fidelity results can be obtained by running the same simulation with two different thresholds T_r (see Figure 11).

6.2. Computational grid deformation

The unstructured hexahedral original grids for the adaptive refinement are generated using HEXPRESS from NUMECA Int. Since unstructured grids may be quite different for geometries that are nearly identical, numerical noise may appear in the simulation results: similar geometries may have different flow solutions. To prevent this, the simulations of all the candidate geometries are performed with the same original grid (see Figure 12(a)), which is deformed to fit each geometry (see Figure 12(b)). The deformation algorithm (Durand 2012) divides the grid in layers of cells around the geometry. For each geometry face, the displacement of the faces with respect




Figure 11. Coarse and fine computational grid example for the DTMB 5415 problem. The figure shows the mesh in the plane $z = 0$, coloured by the water volume fraction. (a) Coarse grid, $T_r = 0.0145$ (b) Fine grid, $T_r = 0.0072$




Figure 12. Example of grid deformation and adaptive refinement for the DTMB 5415 problem. (a) Original grid (b) After deformation (c) After refinement

to the parent one is propagated through these layers. Two smoothing mechanisms are applied: (i) the displacements are diffused over the faces of a single layer, so the grid deformation becomes more uniform; (ii) the displacements are multiplied with a weighting factor which goes from 1 on the geometry to 0 on the outer boundaries, so that the latter are not deformed.

An advantage of the combined deformation and adaptive refinement is that the grid deformation (see Figure 12(b)) can be performed on the coarse initial grid, instead of a fine grid where small errors in the placement of the nodes can lead to inverted cells. The final grid is then created using adaptive refinement (see Figure 12(c)). This procedure is more robust than attempting to deform fine grids. Furthermore, the adaptive refinement ensures that the cells are placed efficiently for capturing the flow around the object and in the wake, independently of the amount of mesh deformation.

6.3. NACA hydrofoil

The initial computational grid for both high- and low-fidelity has 2654 cells, the refinement threshold value T_r is set equal to 0.1 and 0.4 for high- and low-fidelity, respectively. This results in a cell size ratio of 4:1 between the refinement and coarse grids. The actual computational grids have 11 k and 3.6 k cells, respectively (see Figure 13(a,b)).

The domain runs from $11c$ in front of the leading edge to $16c$ behind the hydrofoil and from $-10c$ to $10c$ vertically. Dirichlet conditions on the velocity are imposed, except on the outflow side which has an imposed pressure condition. The hydrofoil surface is treated with a wall law, which is the engineering state-of-the-art for this type of simulation, with $y^+ = 60$ for the first layer. Turbulence is modelled with the standard $k - \omega$ SST model. The flow is computed by time-integrating the time-dependent flow equations towards a steady state; 4000 time steps are run for each




Figure 13. NACA hydrofoil computational grids. (a) Coarse grid, 3.6 k cells (b) Fine grid, 11 k cells

computation. The forces are then evaluated by integrating the pressure and skin friction over the profile and averaging this result over the final 500 time steps.

HF and LF simulations require, respectively, about 13 and 5 minutes of wall-clock time to converge. The resulting computational cost ratio is $\beta = 0.3$. Results for similar foils in Wackers et al. (2017) suggest a numerical uncertainty of around 2% for the fine grids.

A budget of 150 simulations is provided for the adaptive sampling methods, considering both HF and LF simulations. The initial training set for the problem is a set of $2N + 1$ points including the domain centre and min/max coordinates for each variable.

6.3.1. Dynamic positioning

To maintain a constant lift (see Equation (12)), the angle of incidence α for the hydrofoil is adjusted dynamically during the simulations. At regular intervals in a simulation, the difference between the target and the actual lift is evaluated. This difference is divided by the theoretical lift slope of 2D foils ($\Delta C_L = 2\pi \Delta\alpha$), to give a change in angle of attack $\Delta\alpha$. This change is applied over a few time steps, using grid deformation. Then, the flow is allowed to settle and another $\Delta\alpha$ is computed. The procedure converged to within 0.1% of the target lift in all cases. Contrary to the deformations to fit different geometries, the computational grid is deformed with an analytical weighting technique (Leroyer and Visonneau 2005) to accommodate the rotation.

6.4. DTMB 5415 model

The initial computational grid for both high- and low-fidelity has 130k cells, which allows to capture the base flow features. As for the NACA hydrofoil problem, a

4:1 cell size ratio between the HF ($T_r = 0.0145$) and LF ($T_r = 0.0036$) computational grid is selected. The actual computational grids have approximately 4.3 M and 260 k cells, respectively. This ratio appears ideal for the present problem: for higher ratios, either the HF computations are too costly or the LF grids are so coarse that they cannot capture the ship geometry.

Simulations are performed on the half geometry, in model scale with $L = 5.72$ m. The domain runs from $1.5 L$ (with L the vessel length) in front of the bow to $3 L$ behind the stern, up to $2 L$ laterally, and from $-1.5 L$ to $0.5 L$ vertically. Dirichlet conditions on the velocity are imposed on the inflow and side faces, pressure is imposed on the top, bottom, and outflow side. Once again, the hull is treated with a wall law, $y^+ = 60$ for the first layer, and turbulence is modelled with $k - \omega$ SST. The free-surface deformation for wavemaking is treated with a surface-capturing approach that resolves both water and air flow. As usual for free-surface flows, the flow is solved through time integration; 2000 time steps are used. The forces are averaged over the final 250 time steps.

On a 20-core workstation the HF and LF simulations take about 24 and 1.5 hours, respectively, corresponding to a computational cost ratio $\beta = 0.063$.

The initial training set for MF-DRBF is a set of $2N + 1$ points including the domain centre and with each single design variables at either $+1$ or -1 . Adaptive sampling is then run until 80 computations, HF and LF combined, are reached.

6.4.1. Improving the free-surface meshing

To provide adequate surface capturing even on a deformed computational grid, the fine grid around the free surface is created entirely by adaptive refinement (see Figure 12(c)); the original grid has no specific




Figure 14. Example of undisturbed free-surface grid (shown in x -cut planes) for DTMB 5415 problem, depending on the computational grid deformation algorithm (Red is water, blue is air). (a) Undeformed (b) Deformed with Durand (2012) weighting (c) Deformed with power 3.0 weighting

refinement around the free surface at all. This implies that the cells of the original mesh have to be divided many times by the grid adaptation. However, thanks to the anisotropic nature of the refinement algorithm (Wackers et al. 2014), only refinement normal to the free surface is requested. On undeformed grids (see Figure 14(a)) this works well: where the surface is at rest, the computational grid is refined only in vertical direction, resulting of good quality. When the computational grid is deformed (see Figure 14(b)) the cells of the original grid are stretched and rotated. As a consequence, the cells are no longer aligned with the free surface, so they are refined in more than one direction, which leads to unnecessary refinement and mediocre grid quality. This problem is alleviated by changing the weighting factor. Instead of the weighting defined by Durand (2012), which is based on the linear and

squared distance to the body, a weighting based on the distance to the power 3.0 is chosen. Since this weighting law goes to zero rapidly when the distance to the body increases, the grid far from the body is less deformed. As a result, the refined mesh is improved (see Figure 14(c)).

6.4.2. Limiting refinement behind the stern

When simulating ships with adaptive refinement, unrestricted application of the velocity/pressure refinement criterion leads to the accurate resolution of the near and the far wake (Wackers, Guilmineau, and Visonneau 2017). However, the far wake may not be required for drag evaluation. Therefore, tests were run where refinement is forbidden from a certain distance behind the stern; thus, the far-wake region retains a coarse mesh.




Figure 15. Computational grid refinement behind the stern for DTMB 5415 problem, conditional to the refinement limit. (a) $-0.05L$ (b) $-0.25L$ (c) $-1.75L$




Figure 16. Total resistance when horizontal refinement is forbidden beyond a limit behind the stern, for three geometries.

Figure 15 shows the grid at the undisturbed water level for three cases: the refinement is suppressed from 0.05 L, 0.25 L, and 1.75 L behind the stern respectively. To study the effect on the computed drag, such series were computed for two geometries created by varying the optimisation design variable x_2 , as well as the undisturbed case (see Figure 16). The results for the middle limit (0.25 L) are quite close to the largest one (1.75 L), with the biggest discrepancy occurring for $x_2 = -1.25$. The shortest limit (0.05 L) however, changes the drag significantly when the mesh is deformed to $x_2 = 1$, probably because the stern wave pattern is modified; the case $x_2 = 1$ has the steepest stern wave of the three. As a compromise between the computational cost and the accuracy of the total resistance estimation in all deformation cases, the resolution of the near wake (0.25 L) is chosen.

7. Numerical results

The following subsections present the results of the analytical benchmarks and the CFD-based design optimisation problems. For practical purposes, solving the Equations (6), (8)–(10), a distance-based penalisation is applied if the new candidate sample has a distance lower than 0.001 (considering a normalised design variable range) from other points of the training set. The penalisation aims at preventing sample clusterisation and ill-conditioning problems. The above minimum distance between training points is also used as stopping criterion. The number of samples (N_p) computed at once by the MCAS sampling method is set to $N_p = 4$. Furthermore, a threshold $U_f^* = 0.5\%R$ is used for the MCAS (see Equation (10)).

The computational cost is normalised with the HF cost: one HF evaluation costs 1, whereas the LF cost is equal to β , for a total cost equal to $N_H + \beta N_L$, where N_H and N_L are the numbers of high- and low-fidelity computations, respectively.

7.1. Analytical benchmark

Analytical benchmark results are assessed with two metrics: (1) the objective function value for the current optimum (high-fidelity evaluated) and (2) the normalised root mean square error (NMRSE) of the multi-fidelity prediction versus the analytical function (normalisation is based on the initial HF training set range). The NMRSE is computed on a set

of 50^N evenly-spaced points. The available computational budget is set equal to a maximum cost of $50N$. Due to the analytical nature of the test functions, the computational costs are not significant; a computational cost ratio $\beta = 0.1$ between low- and high-fidelity is artificially assumed.

Figure 17 shows the convergence of the metamodel-based optimisation towards the objective minimum and the corresponding NMRSE conditional to the adaptive sampling methods for problems SK1, SSFYY2, MLF1, and Far1. All methods achieve the minimum, excluding the MFEI sampling for the MLF1 problem (note that MFEI faces premature convergence since it has activated the clustering stopping criterion). It is worth noting that the ACAS shows better convergence to the minimum than other methods. Differently, considering the global approximation, the fastest and most significant NMRSE reduction is achieved by the MUAS. The MCAS has similar performance to MUAS.

Since ACAS shows the best performance from the optimisation viewpoint, a parametric analysis for ACAS only is conducted considering $\beta = \{0.05, 0.1, 0.2, 0.4, 0.8\}$. Moreover, a single-fidelity DRBF trained only by HF evaluations is used to assess and compare the MF-DRBF efficiency. The corresponding results for the analytical benchmarks are shown in Figure 18. The MF approach achieves faster convergence towards the minimum and lower NRMSE than using single HF-based metamodel for problems SK1, SSFYY2, and MLF1, regardless of β . For problem MLF1, the HF-based metamodel does not achieve the objective minimum. The Far1 problem shows comparable convergence towards the objective minimum of MF and HF-based DRBF, whereas in terms of NRMSE the HF-based DRBF achieved the lowest value.

7.2. NACA hydrofoil

The drag optimisation of the NACA four-digit hydrofoil Equation (12) is a two-parameter problem like Far1, but it only has a single, global minimum. Therefore, a convergence like for the SSFYY2 problem can be expected. A particular difficulty of this problem is that, while the minimum lies on the domain boundary, it is located close to high drag values which may prevent the sampling of points close by. The adaptive MF-DRBF-based optimisation results are compared to an optimal benchmark solution for Equation (12)

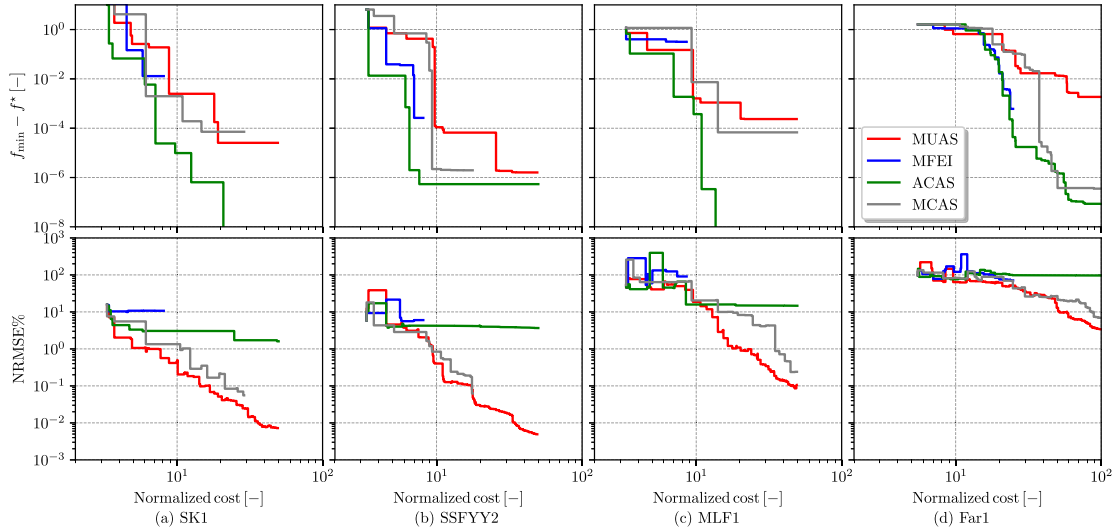


Figure 17. Analytical benchmark problems: adaptive sampling performance comparison for $\beta = 0.1$.




Figure 18. Analytical benchmark problems: ACAS performance conditional to β (only HF is used when $\beta = 1$).

identified by an earlier single-fidelity metamodel-based optimisation trained by 150 HF simulations (Ploé et al. 2017).

Figure 19 shows the convergences of the maximum prediction uncertainty, the drag coefficient minimum, and the corresponding design variables. MUAS and MCAS have achieved lower values of the maximum prediction uncertainty than MFEI and ACAS. On the contrary, only MFEI and ACAS have reached a minimum close to the benchmark. Overall, MFEI has provided the fastest convergence towards the minimum.

Figure 20 shows the MF-DRBF prediction and the corresponding training set at the final iteration of the four sampling strategies. The MUAS and MCAS

strategies have provided a global exploration of the domain. On the contrary, the MFEI and ACAS methods have clustered training points in a small region, close to the global minimum. All the sampling strategies have used a similar number of HF evaluations (see Table 2), with MUAS and MCAS spreading the HF evaluations over the whole design space. Furthermore, both MUAS and MCAS have requested HF training points in the design space corners. The MFEI and ACAS strategies have focused the HF evaluations only in the minimum region.

Furthermore, Figure 19 shows that the convergence of the uncertainty prediction of MUAS is noisy. This happens since MUAS aims at the minimisation of the maximum uncertainty, therefore clusters the training

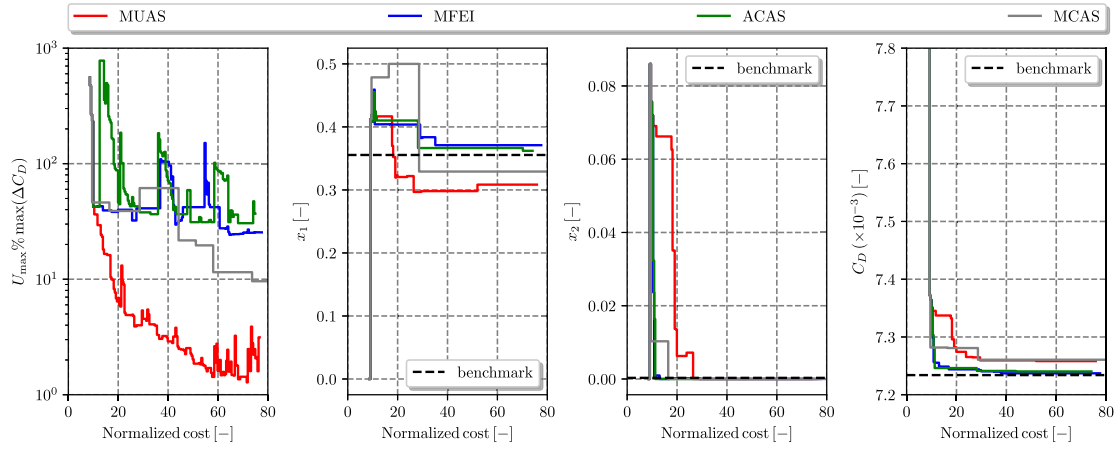


Figure 19. NACA hydrofoil problem: convergence of adaptive MF-DRBF maximum uncertainty, drag coefficient minimum, and corresponding coordinates x_1 and x_2 .




Figure 20. NACA hydrofoil problem: training sets and drag coefficient MF-DRBF prediction at the final iteration of the adaptive sampling procedures.

points in regions where the uncertainty is high (see Figure 20). However, the MF-DRBF uncertainty is affected by a small-scale noise found in the numerical simulations, in regions with clustered training points. (This noise is due to different geometries producing different grids, so two almost similar geometries may have different numerical errors in the solution. Another source is the small variations in lift coming

from the dynamic positioning.) Therefore, MUAS has added training points in regions where the uncertainty is due to the numerical noise and not to the global objective function shape. The effects of the numerical noise are evident in Figures 20(a). MUAS has clustered samples in the neighbourhood of $\mathbf{x} = \{4.0, 4.0\} \times 10^{-2}$ (see Figure 20(a)) mainly LF training points; this region has the highest numerical uncertainty in the

Table 2. Summary of the adaptive sampling methods performance on the NACA hydrofoil problem. The table provides both the metamodel-predicted \hat{C}_D and an actual simulated C_D in the computed optimum.

			minimum position		minimum value	
	N_H [-]	N_L [-]	x_1 [-]	x_2 [-]	\hat{C}_D [-]	C_D [-]
MUAS	14	136	3.8878E-2	3.0000E-2	7.1759E-3	7.2582E-3
MFEI	15	135	4.1691E-2	3.0000E-2	7.1545E-3	7.2371E-3
ACAS	12	138	4.1290E-2	3.0009E-2	7.1754E-3	7.2403E-3
MCAS	11	137	3.9810E-2	3.0000E-2	7.2816E-3	7.2606E-3
Benchmark	150		4.0993E-2	3.0027E-2		7.2340E-3




Figure 21. Hydrofoil optimisation: velocity contour for the optimal configuration identified with the sampling methods. (a) MUAS (b) MFEI (c) ACAS (d) MCAS

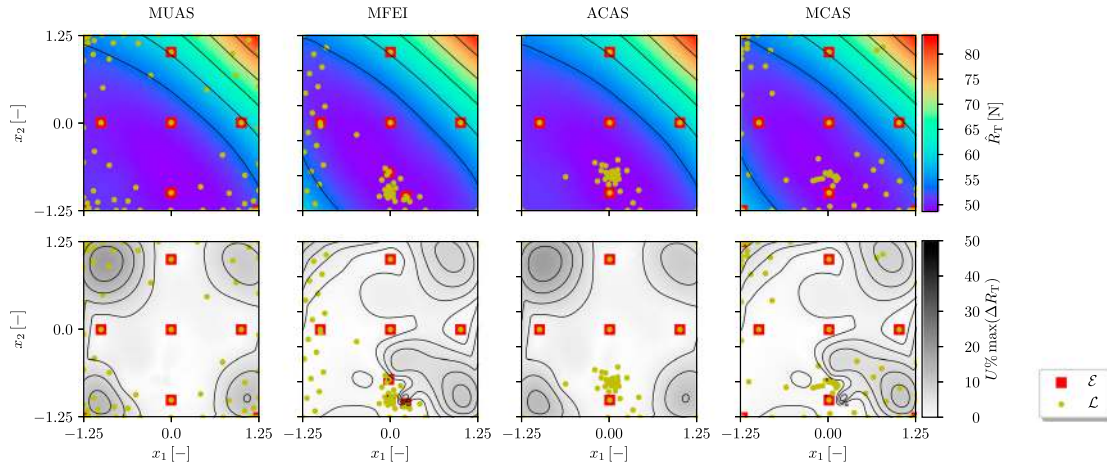


Figure 22. DTMB 5415 problem: training sets and total resistance MF-DRBF prediction at the final iteration of the adaptive sampling procedures.

simulations. The clustering observed for the MFEI, ACAS, and MCAS methods is around the MF-DRBF optimum, since they are based on the objective function value. For those methods, the LF points are accompanied by a less significant clusterisation of HF training points.

Finally, Figure 21 shows the velocity contours of the optimal NACA hydrofoil shape identified by the adaptive sampling strategies. The differences of the hydrofoil shapes are not evident, although the velocity contours are different near the leading edge.

The results of the adaptive sampling strategies are summarised in Table 2, showing that the MFEI method is the most effective adaptive sampling technique for the present problem, followed closely by ACAS.

7.3. DTMB 5415 model

Figure 22 shows the MF-DRBF prediction and the corresponding training set at the final iteration of the four sampling strategies. The figure shows that, even if the design optimisation is a challenging problem from

the CFD perspective, the objective function shape is actually simpler than for the previous problems, since it is unimodal and has a single minimum. Thus, all sampling methods produce similar MF-DRBF meta-models and the optima found are close. Nevertheless, the samples spreading are completely different. The MUAS has concentrated samples on the corners of the design space, following almost a symmetric pattern. Apart from the 5 initial training points, HF samples have been added in the bottom corners. The MFEI shows some erratic sampling; as noted before (see Figure 17) the expected improvement goes to zero at a given moment, meaning that the samples placement is stopped before exhausting the budget. However, with three HF samples lying close to the optimum (one hidden by LF samples in Figure 22) the error metamodel is locally reliable. The ACAS has concentrated all the samples around the optimum. Only one HF sample has been added, close to the optimum location (hidden behind the LF samples in Figure 22). As for the NACA problem, MCAS takes a middle road, adding samples around the optimum but also in the corners. Regarding HF samples, all additional samples are in

Table 3. Summary of the performance of the adaptive sampling strategies on the DTMB 5415 problem.

	N_H [-]	N_L [-]	minimum position		minimum value	
			x_1 [-]	x_2 [-]	\hat{R}_T [N]	R_T [N]
MUAS	8	72	0.4425E-2	-7.8258E-1	48.4044	48.5456
MFEI	7	73	0.3725E-2	-9.9753E-1	48.5508	48.6088
ACAS	6	74	-1.0500E-2	-7.4843E-1	48.3062	48.6236
MCAS	8	70	0.3750E-2	-7.3195E-1	48.4409	48.7344
Original			0.0	0.0		50.6465




Figure 23. Initial (top) and optimised (bottom) geometry and wavepattern (z-coordinate of the free-surface position).

the design space corners. Finally, considering the MF-DRBF uncertainty, only MUAS and ACAS achieve a uniform low uncertainty around the optimum.

Table 3 summarises the locations of the optimum and the optimised resistance for the sampling strategies. Apart from MFEI, the optimum locations lie close together. ACAS has achieved the lowest predicted total resistance. The actual resistance improvement is about 4% for all methods: the MUAS is best overall, whereas the MFEI has the lowest validation error (about 0.1%) with respect to the original DTMB 5415 geometry. This can be explained by the error metamodel. Since $\beta = 0.0625$ is low for this problem, few HF samples are added in the training set, meaning that the error metamodel is not defined very accurately close to the optimum, resulting in a discrepancy between the predicted and the actual optima. Since MFEI has three HF points around the optimum, the error metamodel is well defined in this region and the MF-DRBF prediction with MFEI sampling is the most reliable for this problem.

Figure 23 shows the initial and the MFEI optimal geometry (this is similar to the other optimised geometries). The optimum, compared to the original

geometry, has a more slender aftship which reduces the width of the stern wave, and a slight bulge aft of the bow. This bulge creates a second bow wave out of phase with the first one. By superposition, these two waves cancel so the total bow wave is reduced. Since this interference depends on the wave lengths, which are determined by the ship's velocity, the total resistance reduction is expected to be effective only around the target velocity.

8. Conclusions and future work

Four adaptive sampling methods for MF-DRBF metamodels have been presented. These are assessed by four analytical benchmarks and applied to two CFD-based design optimisation problems, pertaining to the resistance reduction of a NACA hydrofoil and a destroyer-type vessel (DTMB 5415). The MF approximation is obtained as the sum of a low-fidelity-trained metamodel and the metamodel of the error (difference) between high- and low-fidelity evaluations. The stochastic RBF provides the metamodel prediction and the associated uncertainty. The prediction uncertainty of both the low-fidelity and the error metamodel is

used for the adaptive refinement of the low-fidelity and the error training sets. The ratio of the computational cost of high- and low-fidelity evaluations affects the choice of the fidelity to sample.

The criteria for the refinement of the multi-fidelity training set are: (i) the minimisation of the maximum uncertainty of the MF metamodel prediction (MUAS), (ii) the maximisation of the multi-fidelity expected improvement (MFEI), (iii) the minimisation of an aggregated merit factor of the prediction uncertainty and the predicted objective function (ACAS), and (iv) a multi-objective optimisation aiming at maximising the prediction uncertainty and minimising the predicted objective function (MCAS). The MF-DRBF performance has been assessed in terms of convergence towards the minimum. In addition, MF-DRBF maximum uncertainty and NRMSE are used for further comparisons.

The analytical benchmarks results have shown that the ACAS (closely followed by the MFEI sampling) is the most effective method for the MF-DRBF, from the optimisation point of view. For this reason, a further analysis for the ACAS conditional to the computational-cost ratio (β) has been conducted. The ACAS strategy combined with the MF formulation has shown better performance than using a single HF metamodel, regardless of β .

The CFD-based design optimisation problems are challenging for the MF metamodel. The existence of numerical noise affects the RBF interpolation, resulting in large uncertainty of the MF prediction in noisy regions of the domain. Therefore, sampling methods that directly take into account the prediction uncertainty are 'trapped' in such regions. For both CFD-based problems ACAS and MFEI are found the most effective sampling strategies.

In order to overcome the difficulties arising from the presence of numerical noise (associated to CFD simulation residuals), future work will consider the use of regularisation approaches, approximation/regression methods (e.g. least square fit) as opposed to exact interpolation, as well as extensions to the use of Gaussian process. Preliminary studies by the authors can be found in Wackers et al. (2019) and Antognoli et al. (2019). Moreover, the possibility to generalise the MF concept to more than two fidelities (multi-level) is currently being investigated (Serani et al. 2019) and will be further developed in future studies.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

CNR-INM is grateful to Dr. Salahuddin Ahmed and Dr. Woei-Min Lin of the Office of Naval Research, for their support through NICOP grant N62909-18-1-2033.

ORCID

A. Serani  <http://orcid.org/0000-0002-8814-1793>

R. Pellegrini  <http://orcid.org/0000-0003-4953-3038>

M. Diez  <http://orcid.org/0000-0001-6113-7893>

References

- Antognoli, Luca, Marco Bibuli, Matteo Diez, Danilo Durante, Simone Ficini, Salvatore Marrone, Angelo Odetti, Ivan Santic, and Andrea Serani. 2019. "A Synergetic Design Study of a Passenger-Hydrofoil Flapped Surface: Experimental and Computational Fluid Dynamics, Optimization, and Control." In *Proceedings of the 8th International Conference on Computational Methods in Marine Engineering (Marine 2019)*, 334–345.
- Baar, Jouke de, Stephen Roberts, Richard Dwight, and Benoit Mallof. 2015. "Uncertainty Quantification for a Sailing Yacht Hull, Using Multi-fidelity Kriging." *Computers & Fluids* 123: 185–201.
- Cai, Xiwen, Haobo Qiu, Liang Gao, Li Wei, and Xinyu Shao. 2017. "Adaptive Radial-Basis-Function-Based Multifidelity Metamodeling for Expensive Black-Box Problems." *AIAA Journal* 55 (7): 2424–2436.
- Diez, Matteo, Silvia Volpi, Andrea Serani, Frederick Stern, and Emilio F. Campana. 2019. *Simulation-Based Design Optimization by Sequential Multi-criterion Adaptive Sampling and Dynamic Radial Basis Functions*, 213–228. Cham: Springer International Publishing.
- Durand, Mathieu. 2012. "Light and Flexible Fluid/Structure Interaction, Application to Sailing Boats." Theses, Ecole Centrale de Nantes (ECN). <https://hal.archives-ouvertes.fr/tel-01203748>.
- Han, Zhong-Hua, Stefan Görtz, and Ralf Zimmermann. 2013. "Improving Variable-fidelity Surrogate Modeling Via Gradient-enhanced Kriging and a Generalized Hybrid Bridge Function." *Aerospace Science and Technology* 25 (1): 177–189.
- Huang, Deng, Theodore T. Allen, William I. Notz, and R. Allen Miller. 2006. "Sequential Kriging Optimization Using Multiple-fidelity Evaluations." *Structural and Multidisciplinary Optimization* 32 (5): 369–382.
- Huband, S., P. Hingston, L. Barone, and L. While. 2006. "A Review of Multiobjective Test Problems and a Scalable Test Problem Toolkit." *IEEE Transactions on Evolutionary Computation* 10 (5): 477–506.
- Irvine Jr., M., J. Longo, and F. Stern. 2008. "Pitch and Heave Tests and Uncertainty Assessment for a Surface Combatant

- in Regular Head Waves.” *Journal of Ship Research* 52 (2): 146–163.
- Jin, Ruichen, Wei Chen, and Timothy W. Simpson. 2001. “Comparative Studies of Metamodelling Techniques Under Multiple Modelling Criteria.” *Structural and Multidisciplinary Optimization* 23 (1): 1–13.
- Jones, Donald R., Matthias Schonlau, and William J. Welch. 1998. “Efficient Global Optimization of Expensive Black-Box Functions.” *Journal of Global Optimization* 13 (4): 455–492.
- Kennedy, M. C., and A. O’Hagan. 2000. “Predicting the Output From a Complex Computer Code When Fast Approximations Are Available.” *Biometrika* 87 (1): 1–13.
- Kuya, Yuichi, Kenji Takeda, Xin Zhang, and Alexander I. J. Forrester. 2011. “Multifidelity Surrogate Modeling of Experimental and Computational Aerodynamic Data Sets.” *AIAA Journal* 49 (2): 289–298.
- Leroyer, A., and M. Visonneau. 2005. “Numerical Methods for RANSE Simulations of a Self-propelled Fish-like Body.” *Journal of Fluids and Structures* 20 (7): 975–991. Fluid-Plate Interactions.
- Liu, Haitao, Yew-Soon Ong, and Jianfei Cai. 2018. “A Survey of Adaptive Sampling for Global Metamodeling in Support of Simulation-based Complex Engineering Design.” *Structural and Multidisciplinary Optimization* 57 (1): 393–416.
- Liu, Y., Y. Shi, Q. Zhou, and R. Xiu. 2016. “A Sequential Sampling Strategy to Improve the Global Fidelity of Metamodels in Multi-level System Design.” *Structural and Multidisciplinary Optimization* 53 (6): 1295–1313.
- Mackman, T. J., and C. B. Allen. 2010. “Investigation of An Adaptive Sampling Method for Data Interpolation Using Radial Basis Functions.” *International Journal for Numerical Methods in Engineering* 83 (7): 915–938.
- Ng, Leo Wai-Tsun, and Michael Eldred. 2012. “Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation.” In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Structures, Structural Dynamics, and Materials and Co-located Conferences*.
- Pellegrini, R., U. Iemma, C. Leotardi, E. F. Campana, and M. Diez. 2016. “Multi-Fidelity Adaptive Global Metamodel of Expensive Computer Simulations.” In *2016 IEEE Congress on Evolutionary Computation (CEC)*, July, 4444–4451.
- Pellegrini, R., A. Serani, M. Diez, J. Wackers, P. Queutey, and M. Visonneau. 2018. “Adaptive Sampling Criteria for Multifidelity Metamodels in CFD-based Shape Optimization.” In *7th European Conference on Computational Fluid Dynamics (ECFD 7)*, Glasgow, UK: 11-15 June.
- Pellegrini, Riccardo, Andrea Serani, Cecilia Leotardi, Umberto Iemma, Emilio F. Campana, and Matteo Diez. 2017. “Formulation and Parameter Selection of Multi-objective Deterministic Particle Swarm for Simulation-based Optimization.” *Applied Soft Computing* 58: 714–731.
- Ploé, Patrick, Romain Lanos, Michel Visonneau, and Jeroen Wackers. 2017. “Bayesian Strategies for Simulation Based Optimisation and Response Surface Creation Using a Single Tool – Application to Hydrofoil Optimisation.” In *Proceedings of the 7th International Conference on Computational Methods in Marine Engineering (Marine 2017)*.
- Queutey, Patrick, and Michel Visonneau. 2007. “An Interface Capturing Method for Free-Surface Hydrodynamic Flows.” *Computers & Fluids* 36 (9): 1481–1510.
- Serani, Andrea, and Matteo Diez. 2018. “Shape Optimization under Stochastic Conditions by Design-Space Augmented Dimensionality Reduction.” In *19th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference (MA&O), AVIATION 2018*, Atlanta, GA, USA, June 25–29.
- Serani, Andrea, Giovanni Fasano, Giampaolo Liuzzi, Stefano Lucidi, Umberto Iemma, Emilio F. Campana, Frederick Stern, and Matteo Diez. 2016a. “Ship Hydrodynamic Optimization by Local Hybridization of Deterministic Derivative-free Global Algorithms.” *Applied Ocean Research* 59: 115–128.
- Serani, Andrea, Cecilia Leotardi, Umberto Iemma, Emilio Fortunato Campana, Giovanni Fasano, and Matteo Diez. 2016b. “Parameter Selection in Synchronous and Asynchronous Deterministic Particle Swarm Optimization for Ship Hydrodynamics Problems.” *Applied Soft Computing* 49: 313–334.
- Serani, Andrea, Riccardo Pellegrini, Riccardo Brogna, Jeroen Wackers, Michel Visonneau, and Matteo Diez. 2019. “An Adaptive N-Fidelity Metamodel for Design and Operational-Uncertainty Space Exploration of Complex Industriail Problems.” In *Proceedings of the 8th International Conference on Computational Methods in Marine Engineering (Marine 2019)*, 177–188.
- Stern, F., J. Longo, R. Penna, A. Olivieri, T. Ratcliffe, and H. Coleman. 2001. “International Collaboration on Benchmark CFD Validation Data for Surface Combatant DTMB Model 5415”. In *Twenty-Third Symposium on Naval Hydrodynamics*. Washington, DC: Office of Naval Research, Bassin d’Essais des Carenes, National Research Council.
- Viana, Felipe, and Raphael Haftka. 2010. “Surrogate-Based Optimization with Parallel Simulations using the Probability of Improvement.” In *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*.
- Viana, Felipe A. C., T. W. Simpson, V. Balabanov, and T. Vasilli. 2014. “Special Section on Multidisciplinary Design Optimization: Metamodeling in Multidisciplinary Design Optimization: How Far Have We Really Come?” *AIAA Journal* 52 (4): 670–690.
- Volpi, Silvia, Matteo Diez, Nicholas J. Gaul, Hyeongjin Song, Umberto Iemma, K. K. Choi, Emilio F. Campana, and Frederick Stern. 2015. “Development and Validation of a Dynamic Metamodel Based on Stochastic Radial Basis Functions and Uncertainty Quantification.” *Structural and Multidisciplinary Optimization* 51 (2): 347–368.
- Wackers, Jeroen, Ganbo Deng, Emmanuel Guilmineau, Alban Leroyer, Patrick Queutey, and Michel Visonneau. 2014. “Combined Refinement Criteria for Anisotropic Grid Refinement in Free-surface Flow Simulation.” *Computers and Fluids* 92: 209–222.

- Wackers, Jeroen, Ganbo Deng, Emmanuel Guilmineau, Alban Leroyer, Patrick Queutey, Michel Visonneau, Alexandro Palmieri, and Alfredo Liverani. 2017. "Can Adaptive Grid Refinement Produce Grid-independent Solutions for Incompressible Flows?" *Journal of Computational Physics* 344: 364–380.
- Wackers, Jeroen, Emmanuel Guilmineau, and Michel Visonneau. 2017. "Unsteady Behaviour in RANS Simulation of the JBC and KVLCC2." In *Proceedings of the 20th Numerical Towing Tank Symposium – NuTTS 2017*, Wageningen, The Netherlands.
- Wackers, Jeroen, Andrea Serani, Riccardo Pellegrini, Matteo Diez, and Michel Visonneau. 2019. "Adaptive Multifidelity Shape Optimization Based on Noisy CFD Data." In *Proceedings of the International Conference on Adaptive Modeling and Simulation ADMOS 2019*.
- Yi, Seulgi, Hyung I. Kwon, and Seongim Choi. 2014. *Efficient Global Optimization using a Multi-point and Multi-objective Infill Sampling Criteria*.
- Zhao, L., K. K. Choi, and I. Lee. 2011. "Metamodeling Method Using Dynamic Kriging for Design Optimization." *AIAA Journal* 49 (9): 2034–2046.
- Zheng, Jun, Xinyu Shao, Liang Gao, Ping Jiang, and Zilong Li. 2013. "A Hybrid Variable-fidelity Global Approximation Modelling Method Combining Tuned Radial Basis Function Base and Kriging Correction." *Journal of Engineering Design* 24 (8): 604–622.