



HAL
open science

A multi-fidelity active learning method for global design optimization problems with noisy evaluations

Riccardo Pellegrini, Jeroen Wackers, Riccardo Broglia, Matteo Diez, Andrea Serani, Michel Visonneau

► **To cite this version:**

Riccardo Pellegrini, Jeroen Wackers, Riccardo Broglia, Matteo Diez, Andrea Serani, et al.. A multi-fidelity active learning method for global design optimization problems with noisy evaluations. *Engineering with Computers*, 2022, 10.1007/s00366-022-01728-0 . hal-03783173

HAL Id: hal-03783173

<https://hal.science/hal-03783173>

Submitted on 21 Sep 2022

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.

Engineering with Computers



A Multi-Fidelity Active Learning Method for Global Design Optimization Problems with Noisy Evaluations --Manuscript Draft--

Manuscript Number:					
Full Title:	A Multi-Fidelity Active Learning Method for Global Design Optimization Problems with Noisy Evaluations				
Article Type:	Original Article				
Corresponding Author:	Riccardo Pellegrini, Ph.D. Institute of Marine Engineering National Research Council: Istituto di Ingegneria del Mare Consiglio Nazionale delle Ricerche ITALY				
Corresponding Author Secondary Information:					
Corresponding Author's Institution:	Institute of Marine Engineering National Research Council: Istituto di Ingegneria del Mare Consiglio Nazionale delle Ricerche				
Corresponding Author's Secondary Institution:					
First Author:	Riccardo Pellegrini				
First Author Secondary Information:					
Order of Authors:	Riccardo Pellegrini Jeroen Wackers Riccardo Broglio Matteo Diez Andrea Serani Michel Visonneau				
Order of Authors Secondary Information:					
Funding Information:	<table border="1"><tr><td>Office of Naval Research (N62909-21-1-2042)</td><td>Dr. Matteo Diez</td></tr><tr><td>H2020 European Research Council (689074)</td><td>Dr. Riccardo Broglio</td></tr></table>	Office of Naval Research (N62909-21-1-2042)	Dr. Matteo Diez	H2020 European Research Council (689074)	Dr. Riccardo Broglio
Office of Naval Research (N62909-21-1-2042)	Dr. Matteo Diez				
H2020 European Research Council (689074)	Dr. Riccardo Broglio				
Abstract:	<p>A multi-fidelity (MF) active learning method is presented for design optimization problems characterized by noisy evaluations of the performance metrics. Namely, a generalized MF surrogate model is used for design-space exploration, exploiting an arbitrary number of hierarchical fidelity levels, i.e., performance evaluations coming from different models, solvers, or discretizations, characterized by different accuracy. The method is intended to accurately predict the design performance while reducing the computational effort required by simulation-driven design (SDD) to achieve the global optimum. The overall MF prediction is evaluated as a low-fidelity trained surrogate corrected with the surrogates of the errors between consecutive fidelity levels. Surrogates are based on stochastic radial basis functions (SRBF) with least squares regression and in-the-loop optimization of hyperparameters to deal with noisy training data. The method adaptively queries new training data, selecting both the design points and the required fidelity level via an active learning approach. This is based on the lower confidence bounding method, which combines performance prediction and associated uncertainty to select the most promising design regions. The fidelity levels are selected considering the benefit-cost ratio associated with their use in the training. The method's performance is assessed and discussed using four analytical tests and three SDD problems based on computational fluid dynamics (CFD) simulations, namely the shape optimization of a NACA hydrofoil, the DTMB 5415 destroyer, and a roll-on/roll-off passenger ferry. Fidelity levels are provided by both</p>				

	<p>adaptive grid refinement and multi-grid resolution approaches. Under the assumption of a limited budget of function evaluations, the proposed MF method shows better performance in comparison with the model trained by high-fidelity evaluations only.</p>
<p>Suggested Reviewers:</p>	<p>Domenico Quagliarella d.quagliarella@cira.it Expert in optimization and surrogate modeling</p> <hr/> <p>Umberto Iemma umberto.iemma@uniroma3.it Expert in optimization and multi-fidelity surrogate modeling</p> <hr/> <p>Markus Rumpfkeil mrumpfkeil1@udayton.edu Expert in multi-disciplinary design optimization</p> <hr/> <p>Philip Beran philip.beran@us.af.mil Expert in simulation based design and multi-fidelity methods</p> <hr/> <p>Thomas Scholcz T.P.Scholcz@marin.nl Expert in optimization and surrogate modeling</p> <hr/> <p>Luca Bonfiglio bonfi@mit.edu Expert in multi-fidelity optimization</p>



Click here to access/download
Supplementary Material
Cover_letter_EWCO-MF-SRBF.pdf



A Multi-Fidelity Active Learning Method for Global Design Optimization Problems with Noisy Evaluations

Riccardo Pellegrini^{1,*,\dagger} · Jeroen Wackers^{2,*} · Riccardo Brogna¹ · Matteo Diez¹ · Andrea Serani¹ · Michel Visonneau²

Received: date / Accepted: date

Abstract A multi-fidelity (MF) active learning method is presented for design optimization problems characterized by noisy evaluations of the performance metrics. Namely, a generalized MF surrogate model is used for design-space exploration, exploiting an arbitrary number of hierarchical fidelity levels, i.e., performance evaluations coming from different models, solvers, or discretizations, characterized by different accuracy. The method is intended to accurately predict the design performance while reducing the computational effort required by simulation-driven design (SDD) to achieve the global optimum. The overall MF prediction is evaluated as a low-fidelity trained surrogate corrected with the surrogates of the errors between consecutive fidelity levels. Surrogates are based on stochastic radial basis functions (SRBF) with least squares regression and in-the-loop optimization of hyperparameters to deal with noisy training data. The method adaptively queries new training data, selecting both the design points and the required fidelity level via an active learning approach. This is based on the lower confidence bounding method, which combines performance prediction and associated uncertainty to select the most promising design regions. The fidelity levels are selected considering the benefit-cost ratio associated with their use in the training. The method's performance is assessed and discussed using four analytical tests and three SDD problems based on computational fluid dynamics

(CFD) simulations, namely the shape optimization of a NACA hydrofoil, the DTMB 5415 destroyer, and a roll-on/roll-off passenger ferry. Fidelity levels are provided by both adaptive grid refinement and multi-grid resolution approaches. Under the assumption of a limited budget of function evaluations, the proposed MF method shows better performance in comparison with the model trained by high-fidelity evaluations only.

Keywords Multi-fidelity optimization · Active learning · Radial basis functions · RANS · Adaptive grid refinement · Multi-grid resolution

1 Introduction

The quest for ever-innovative engineering products has motivated the development of highly-accurate performance analysis tools, often prime-principle based and therefore able to span a vast variety of design solutions and operating conditions. The need for fast, effective, and ingenious design decisions has stimulated the integration of these tools with efficient global optimization approaches into simulation-driven design (SDD) [18] architectures. Naval architecture and ocean engineering are no exceptions, experiencing an outstanding development of computational tools to assess complex physical phenomena at different spatial and temporal scales [29, 41, 4], along with their integration in global SDD optimization approaches [10, 39].

These tools are generally computationally expensive, making the design- and operational-space exploration a technological and algorithmic challenge, thus preventing their widespread use in the industrial community. Recent work in the context of computational fluid dynamics (CFD) optimization showed

¹CNR-INM, National Research Council–Institute of Marine Engineering, Rome, Italy

²LHEEA Lab, Ecole Centrale de Nantes, CNRS UMR 6598, France

*Drs. Riccardo Pellegrini and Jeroen Wackers equally contributed to this paper.

[†]Corresponding author. E-mail: riccardo.pellegrini@inm.cnr.it

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 how a hull-form optimization based on accurate un-
 2 steady Reynolds Averaged Navier-Stokes (URANS)
 3 solutions under realistic stochastic conditions may re-
 4 quire up to 500k CPU hours on high performance
 5 computing (HPC) systems, even if computational cost
 6 reduction methods are used [39]. Similarly, an accu-
 7 rate URANS-based statistically significant evaluation
 8 of ship maneuvering performance in irregular waves
 9 may require up to 1M CPU hours on HPC systems
 10 [34]. Furthermore, multidisciplinary design optimiza-
 11 tion (MDO) problems, involving several interconnected
 12 disciplines, usually require additional computational
 13 efforts especially for global optimization problems [20]
 14 and stochastic conditions [11]. Achieving global and
 15 stochastic MDO still remains a challenge from the
 16 computational viewpoint [24].

17 To reduce the computational cost of SDD pro-
 18 cedures, supervised machine learning methods in the
 19 form of surrogate models have been developed and suc-
 20 cessfully applied in several engineering domains [42].
 21 With these methods, an approximate and easy to eval-
 22 uate model of expensive computations is constructed
 23 based on a limited number of simulations. The design-
 24 space exploration and optimization are then performed
 25 on the surrogate model, increasing the efficiency of the
 26 overall SDD process.

27 The performance of surrogate models is problem-
 28 dependent and determined by several concurrent fac-
 29 tors, such as the presence of nonlinearities, the problem
 30 dimensionality, and the approach used for the training
 31 [22]. Therefore, in the last decade, the research has
 32 moved from static to function-adaptive approaches,
 33 also known as dynamic (adaptive) surrogate models
 34 [43], which are able to improve their fitting capability
 35 by active learning. Namely, the design of experiments
 36 used for the surrogate model training is not defined
 37 a priori but dynamically updated, exploiting the in-
 38 formation about the problem that becomes available
 39 during the analysis process. Training points are dy-
 40 namically added based on active learning criteria, with
 41 the aim of producing globally-accurate representations
 42 of the desired functions with as few training points as
 43 possible [38]. Unfortunately, the active learning process
 44 is generally affected by the computational-output noise
 45 [22], which is often unavoidable when large systems of
 46 nonlinear partial differential equations are numerically
 47 solved. Active learning methods may react to noise by
 48 adding many training points in noisy regions, rather
 49 than selecting new points in unseen regions [47]. This
 50 may deteriorate the model quality/efficiency and needs
 51 to be carefully considered.

52 In addition to dynamic/adaptive surrogate
 53 models and with the aim of reducing further the com-

putational cost associated to SDD, multi-fidelity (MF)
 approximation methods have been developed, aiming
 at combining the accuracy of high-fidelity solvers with
 the computational cost of low-fidelity solvers [3]. Thus,
 MF surrogate models use mainly low-fidelity simula-
 tions and only few high-fidelity simulations are used
 to preserve the model accuracy. Additive and/or mul-
 tiplicative correction methods, also known as “bridge
 functions” [17], can be used to build MF surrogate
 models. Several surrogate models have been used in
 the literature with MF data, such as non-intrusive
 polynomial chaos [27], co-kriging [2] and stochastic ra-
 dial basis functions (SRBF) [38]. In general, different
 fidelity levels may be obtained by varying the physical
 model, the grid size, and/or combining experimental
 data with simulations. The active learning process is
 extended to MF analysis via the automatic selection of
 both the desired training point and fidelity level [38].
 Most MF methods use two-fidelity levels, nevertheless
 the approach can be generalized to an arbitrary num-
 ber of fidelity levels, as shown in [37]. The use of MF
 models with noisy training data and the assessment
 of the effect of the noise associated with each fidelity
 level are still little discussed and require a rigorous
 assessment.

The objective of the present work is to present
 a generalized adaptive MF surrogate model for global
 design optimization of complex industrial problems af-
 fected by noisy performance evaluations. The proposed
 MF method advances the authors’ previous work by
 combining an arbitrary number of fidelity levels [37]
 with noise reduction and in-the-loop optimization of
 the MF surrogate model [47]. SRBF [43] are used in
 combination with an active learning method based on
 the lower confidence bounding (LCB), combining ob-
 jective function prediction and associated uncertainty
 [6,38]. Furthermore, the present formulation fully ex-
 ploits the potential of simulation methods that natu-
 rally produce results spanning a range of fidelity levels:
 i.e. URANS simulations with adaptive grid refinement
 [44] and multi-grid resolution [4].

The performance of the proposed MF method
 is assessed using four analytical test problems and
 three SDD problems pertaining to the minimization of:
 (1) the drag-coefficient of a NACA hydrofoil, (2) the
 calm-water resistance of the DTMB 5415 ship model,
 and (3) the calm-water resistance/payload ratio of a
 roll-on/roll-off passengers (RoPax) ferry, under the
 assumption of a limited budget of function evalua-
 tions. CFD computations are based on two URANS
 solvers: ISIS-CFD [29], developed at Ecole Centrale
 de Nantes/CNRS and integrated in the FINE/Marine
 simulation suite from Cadence Design Systems, for the

1 NACA hydrofoil and the DTMB 5415; and χ navis [7, 8, 4], developed at CNR-INM, for the RoPax ferry. In
2 ISIS-CFD, mesh deformation and adaptive grid re-
3 finement are adopted to allow the automatic shape
4 deformation. The fidelity levels are defined by the grid
5 refinement ratio. In χ navis, the mesh deformation is
6 initially computed on the hull-surface grid and then
7 propagated in the volume mesh. The different fidelities
8 are the multi-grid levels. Problems are solved with a
9 number of fidelity levels between one and four.

10 The remainder of this paper is organized as fol-
11 lows. Section 2 introduces the generalized multi-fidelity
12 active learning method. Section 3 describes the analyt-
13 ical benchmarks and the SDD problems. CFD solvers
14 and numerical setup for the solutions of the SDD prob-
15 lems are provided in Sections 4 and 5, respectively.
16 The numerical results are discussed in Section 6 and,
17 finally, conclusions and an outlook on future work are
18 presented in Section 7.

2 Multi-fidelity Surrogate Modeling

24 Consider an objective function $f(\mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^D$ is
25 the design vector of dimension D . Let the true function
26 $f(\mathbf{x})$ be assessed by numerical simulations $s_l(\mathbf{x})$ with
27 different fidelity levels l , which are considered to have
28 random noise:
29

$$s_l(\mathbf{x}) \equiv f_l(\mathbf{x}) + \mathcal{N}_l(\mathbf{x}) \quad \text{with } l = 1, \dots, N, \quad (1)$$

30 where $s_1(\mathbf{x})$ denotes the highest-fidelity level, $s_N(\mathbf{x})$ is
31 the lowest-fidelity, and $\{s_l(\mathbf{x})\}_{l=2}^{N-1}$ are the intermediate-
32 fidelity levels. $f_l(\mathbf{x})$ is the hypothetical simulation re-
33 sponse without noise. The simulation noise for each
34 fidelity level \mathcal{N}_l is considered as realizations of zero-
35 mean uncorrelated random variables. This noise will be
36 (partially) removed in the surrogate models. A multi-
37 fidelity approximation $\hat{f}(\mathbf{x})$ of $f(\mathbf{x})$ can then be built
38 by hierarchical superposition of the lowest-fidelity sur-
39rogate model $\tilde{f}_N(\mathbf{x})$ and the inter-level errors $\tilde{\varepsilon}_l(\mathbf{x})$
40 as

$$f(\mathbf{x}) \approx \hat{f}(\mathbf{x}) = \tilde{f}_N(\mathbf{x}) + \sum_{l=1}^{N-1} \tilde{\varepsilon}_l(\mathbf{x}), \quad (2)$$

41 For the l -th fidelity level the available simu-
42 lation data is defined as $\mathcal{T}_l = \{(\mathbf{x}_j^\top, s_l(\mathbf{x}_j))\}_{j=1}^{\mathcal{J}_l}$, with
43 \mathcal{J}_l the training set size. The resulting inter-level error
44 training set, used to realize the error surrogate $\tilde{\varepsilon}_l(\mathbf{x})$, is
45 defined as $\mathcal{E}_l = \{(\mathbf{x}_j^\top, \varepsilon_l(\mathbf{x}_j))\}_{j=1}^{\mathcal{J}_l}$, with

$$\varepsilon_l(\mathbf{x}_j) = s_l(\mathbf{x}_j) - \hat{f}_{l+1}(\mathbf{x}_j), \quad (3)$$




Fig. 1: Example of MF surrogate model with $N = 2$ (without noise).

where the generical l -th multi-fidelity surrogate is:

$$\hat{f}_l(\mathbf{x}) = \tilde{f}_N(\mathbf{x}) + \sum_{i=l}^{N-1} \tilde{\varepsilon}_i(\mathbf{x}). \quad (4)$$

The choice of \hat{f}_{l+1} instead of s_{l+1} is based on the idea
that a significant amount of noise can be present in
the sampling data and that the surrogate models effec-
tively filter this noise. Thus, in the presence of noise,
the $(l+1)$ -th surrogate model is a better representation
of the $(l+1)$ -th response than the actual simulations.
However, s_{l+1} is available: to ensure the most effective
training process and considering the nature of the CFD
solvers used (using adaptive and/or multi grids), lower-
fidelity simulations are added in all the points \mathbf{x}_j where
a high-fidelity point is simulated. If there is no noise,
the surrogate model $\hat{f}_{l+1}(\mathbf{x}_j)$ interpolates the data in
all the training points, so the error data in the training
points $\varepsilon_l(\mathbf{x}_j)$ is exact and equal to $s_l(\mathbf{x}_j) - s_{l+1}(\mathbf{x}_j)$.

Finally, given a surrogate modeling that pro-
vides both the prediction and the associated uncer-
tainty, and assuming that the uncertainties associated
with the lowest-fidelity level $U_{\tilde{f}_N}$ and the inter-level er-
rors $U_{\tilde{\varepsilon}_l}$ are uncorrelated, the multi-fidelity prediction
uncertainty $U_{\hat{f}}$ reads

$$U_{\hat{f}}(\mathbf{x}) = \sqrt{U_{\tilde{f}_N}^2(\mathbf{x}) + \sum_{l=1}^{N-1} U_{\tilde{\varepsilon}_l}^2(\mathbf{x})}. \quad (5)$$

An example with two fidelities, without noise, is shown
in Figure 1.

In the present work, the term uncertainty al-
ways refers to the surrogate model prediction uncer-
tainty (see Eq. 5), whereas the noise is associated with
the objective function evaluation and intrinsically re-
lated to the fidelity level: higher fidelities tend to be
less noisy.




Fig. 2: Example of the active learning method using one fidelity without noise: (left) shows the initial surrogate model with the associated prediction uncertainty and training set; (right) shows the position of the new training point and the new surrogate model prediction and its uncertainty.

2.1 Active Learning

The multi-fidelity surrogate model is dynamically updated by adding new training points following a two-step procedure:

1. Identify the new training point \mathbf{x}^* ;
2. Defining $\beta_l = c_l/c_1$, where c_l is the computational cost associated to the l -th level and c_1 the computational cost of the highest-fidelity, $\mathbf{U} \equiv \{U_{\hat{\varepsilon}_1}/\beta_1, \dots, U_{\hat{\varepsilon}_{N-1}}/\beta_{N-1}, U_{\hat{f}_N}/\beta_N\}$ as the surrogate model prediction uncertainty vector, and $l^* = \text{maxloc}(\mathbf{U})$, add the new training point to the l^* -th training set \mathcal{T}_{l^*} (as well as to the lower-fidelity sets from $l^* + 1$ up to $l^* = N$).

The identification of the new training points is based on the LCB sampling method [6] with equal weight for the function value and the prediction uncertainty (see Figure 2). It aims to find points with large prediction uncertainty and small objective function value. Accordingly, LCB identifies new training points by solving

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

where $\hat{f}(\mathbf{x}) - U_{\hat{f}}(\mathbf{x})$ represents the so-called acquisition function for the active learning process.

2.2 Least Square Regression via In-the-loop Optimization

The surrogate model predictions $\tilde{f}(\mathbf{x})$ are computed as the expected value (EV) over a stochastic ensemble of Radial Basis Function (RBF) surrogate models, defined by a stochastic tuning parameter, $\tau \sim \text{unif}[1, 3]$:

$$\tilde{f}(\mathbf{x}) = \text{EV} [g(\mathbf{x}, \tau)]_{\tau}, \quad (7)$$

with

$$g(\mathbf{x}, \tau) = \sum_{j=1}^{\mathcal{K}^*} w_j \|\mathbf{x} - \mathbf{c}_j\|^\tau, \quad (8)$$

where w_j are unknown coefficients, $\|\cdot\|$ is the Euclidean norm, and \mathbf{c}_j are the RBF centers, whose coordinates are defined via k -means clustering [23] of the training point coordinates in the design space. The uncertainty $U_{\tilde{f}}(\mathbf{x})$ associated with the SRBF surrogate model prediction is quantified by the 95%-confidence interval of $g(\mathbf{x}, \tau)$, evaluated using a Monte Carlo sampling over τ [43]. Noise reduction is achieved by choosing a number of SRBF centers \mathcal{K} less than the number of training points \mathcal{J} . Hence, w_j are determined with least squares regression by solving

$$\mathbf{w} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{s}, \quad (9)$$

where $\mathbf{w} = \{w_j\}$, $a_{ij} = \|\mathbf{x}_i - \mathbf{c}_j\|^\tau$, and $\{(\mathbf{x}_i, s(\mathbf{x}_i))\}_{i=1}^{\mathcal{J}} \in \mathcal{T}$. The optimal number of SRBF centers (\mathcal{K}^*) is defined by minimizing a leave-one-out cross-validation (LOOCV) metric [13,21]. Let $\tilde{h}(\mathbf{x})$ be a surrogate model trained by all points but the i -th point using \mathcal{K} centers, then \mathcal{K}^* is defined as:

$$\mathcal{K}^* = \underset{\mathcal{K}}{\text{argmin}} (\text{RMSE}), \quad (10)$$

where the root mean squared error (RMSE) is defined as

$$\text{RMSE} = \sqrt{\frac{1}{\mathcal{J}} \sum_{i=1}^{\mathcal{J}} (s(\mathbf{x}_i) - \tilde{h}(\mathbf{x}_i))^2}. \quad (11)$$

To avoid abrupt changes in the surrogate model prediction from one iteration to the next one, during the active learning procedure the search for \mathcal{K}^* can be constrained. In the present work, $\mathcal{K}_{k-1}^* - 2 < \mathcal{K}_k^* < \mathcal{K}_{k-1}^* +$

2, with k the active learning iteration. An example with $\mathcal{J} = 6$ and $\mathcal{K}^* = 3$ is shown in Figure 3.

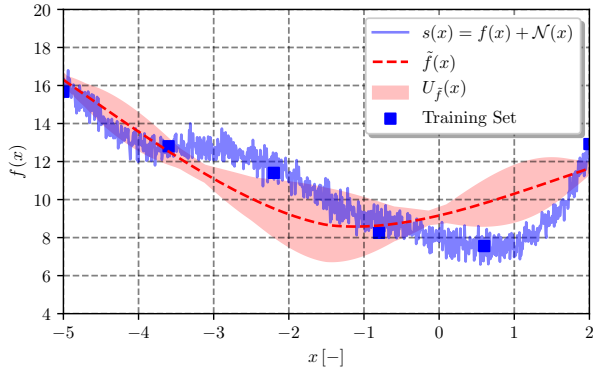


Fig. 3: Example of least squares regression by SRBF.

The function which gives the lowest EV for the LOOCV metric is the exact objective function $f(\mathbf{x})$. Since f and \tilde{h} are deterministic functions, the EV of an error-squared term in the RMSE measure can be expanded as

$$\begin{aligned} \text{EV} \left[\left(s(\mathbf{x}_i) - \tilde{h}(\mathbf{x}_i) \right)^2 \right] &= \\ \text{EV} \left[\left(s(\mathbf{x}_i) - f(\mathbf{x}_i) + f(\mathbf{x}_i) - \tilde{h}(\mathbf{x}_i) \right)^2 \right] &= \\ = \text{EV} \left[\left(s(\mathbf{x}_i) - f(\mathbf{x}_i) \right)^2 \right] + & \quad (12) \\ 2\text{EV} \left[\left(s(\mathbf{x}_i) - f(\mathbf{x}_i) \right) \left(f(\mathbf{x}_i) - \tilde{h}(\mathbf{x}_i) \right) \right] + & \\ \left(f(\mathbf{x}_i) - \tilde{h}(\mathbf{x}_i) \right)^2. & \end{aligned}$$

Since the noise has a zero mean, the EV in the middle term vanishes. Thus, the overall expected value for the error is minimized if $\tilde{h}(\mathbf{x}_i) = f(\mathbf{x}_i)$. Therefore, the LOOCV criterion in Eq. 10 is a suitable measure of the quality of a surrogate model function.

3 Optimization Problems

The assessment of the multi-fidelity machine learning method is based on four analytical test and three CFD-based design optimization problems, with design space dimensions ranging from $D = 1$ to 10. Problems are solved with a number of fidelity levels N ranging from 1 to 4. These are reasonable numbers in multi-fidelity shape optimization, where several grid resolution levels and/or physical models may be considered.

The initial training set for each problem is set as $2D + 1$ points for each fidelity, i.e. the domain center




Fig. 4: Analytical test problem P_1 .

and the centers of the domain boundaries. Details are provided in the following subsections.

A deterministic single-objective formulation of the particle swarm optimization algorithm [36], is used for the surrogate model-based optimizations, as well as for the solution of the minimization sampling problem of Eq. 6. The active learning is performed with a fixed budget of function evaluations: considering a normalized computational cost of a highest-fidelity evaluation (equal to 1), the overall computational cost CC is proportional to the training set sizes \mathcal{J}_l and is defined as:

$$CC = \sum_{l=1}^N \beta_l \mathcal{J}_l. \quad (13)$$

3.1 Analytical Test Problems

Four analytical test problems are used to assess the MF method, namely a MF version of the Forrester (P1), Griewank (P2), Rosenbrock (P3), and shifted-rotated Rastrigin (P4) functions. Table 1 summarizes the test problem equations, domain size, dimensions, reference global optimum position $\tilde{\mathbf{x}}$ and value $f(\tilde{\mathbf{x}})$, and finally the bibliographical reference. Most problems have $N = 3$.

Artificial noise is added to the analytical problems. For each fidelity l a normal distributed noise $\eta_l \in \mathcal{N}(0, \sigma_l)$ is defined, where $\sigma_l = \{0.025, 0.05, 0.10\} R_1$, with R_1 the function range of the highest-fidelity. Figure 4 shows an example of the artificial noise applied to the analytical function P_1 as semi-transparent scatter.

The analytical test problems are selected to be representative of complex real-world problems and challenging from the optimization viewpoint. They are characterized by several local minima and adjacent regions of strong gradients and flatness, as shown in Figs. 4 and 5. Furthermore, the different fidelity levels

Table 1: Analytical benchmark problems

Test	Formulation	Domain	D	$\bar{\mathbf{x}}$	$f(\bar{\mathbf{x}})$	Reference
P_1	$f_1(x) = [(6x - 2)^2] \sin(12x - 4) + \eta_1$ $f_2(x) = 0.75f_1(x) + 5(x - 0.5) - 2 + \eta_2$ $f_3(x) = 0.5f_1(x) + 10(x - 0.5) - 5 + \eta_3$	$x \in [0, 1]$	1	0.7572	-6.0207	[16] [28] [15]
P_2	$f_1(\mathbf{x}) = \sum_{j=1}^D x_j^2/25 - \prod_{j=1}^D \cos(x_j/\sqrt{j}) + 1 + \eta_1$ $f_2(\mathbf{x}) = -\prod_{j=1}^D \cos(x_j/\sqrt{j}) + 1 + \eta_2$ $f_3(\mathbf{x}) = \sum_{j=1}^D x_j^2/20 - \prod_{j=1}^D \cos(x_j/\sqrt{j+1}) - 1 + \eta_3$	$\mathbf{x} \in [-6, 5]$	2	[0, 0]	0.0000	[1] [15] [15]
P_3	$f_1(\mathbf{x}) = \sum_{j=1}^{D-1} [100(x_{j+1} - x_j^2)^2 + (1 - x_j)^2] + \eta_1$ $f_2(\mathbf{x}) = \sum_{j=1}^{D-1} [50(x_{j+1} - x_j^2)^2 + (-2 - x_j)^2] - \sum_{j=1}^D 0.5x_j + \eta_2$ $f_3(\mathbf{x}) = (f_1(\mathbf{x}) - 4 - \sum_{j=1}^D 0.5x_j)/(10 + \sum_{j=1}^D 0.25x_j) + \eta_3$	$\mathbf{x} \in [-2, 2]$	2,5,10	[1, ..., 1]	0.0000	[32] [15] [32]
P_4	$f_1(\mathbf{z}) = \sum_{j=1}^D (z_j^2 + 1 - \cos(10\pi z_j))$ $f_i(\mathbf{z}) = f_1(\mathbf{z}) + e_r(\mathbf{z}, \phi_i) + \eta_i, \quad i = 2, \dots, N$ $e_r(\mathbf{z}, \phi_i) = \sum_{j=1}^D a(\phi_i) \cos^2 \omega(\phi_i) z_j + b(\phi_i) + \pi, \quad i = 2, \dots, N$ with $a(\phi_i) = \Theta(\phi_i), \omega(\phi_i) = 10\pi\Theta(\phi_i), b(\phi_i) = 0.5\pi\Theta(\phi_i)$ and $\Theta(\phi_i) = 1 - 0.0001\phi_i$ with $\phi = \{10000, 5000, 2500\}$	$\mathbf{x} \in [-0.1, 0.2]$	2,5,10	[0.1, ..., 0.1]	0.0000	[48] [48] [48] [48] [48] -

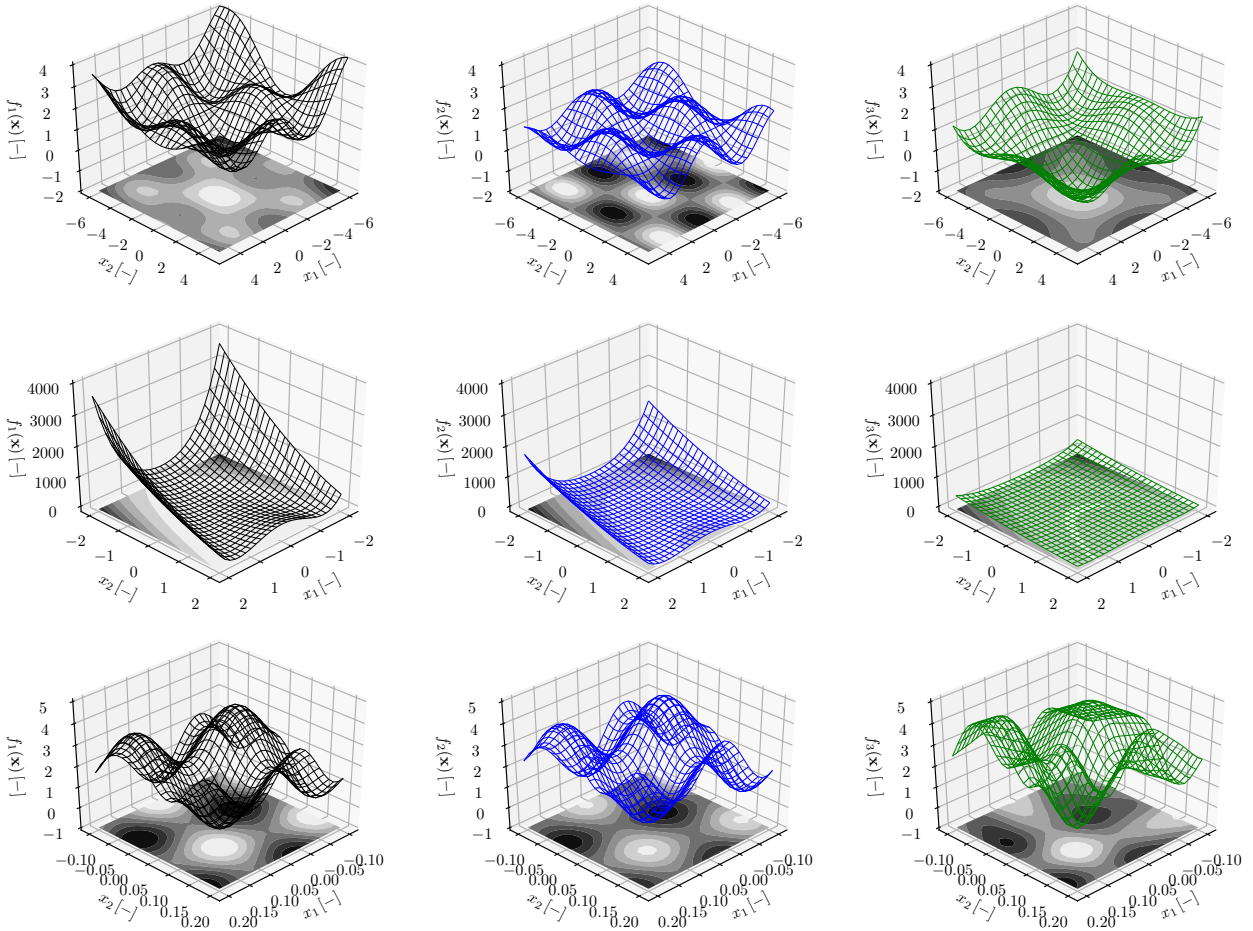


Fig. 5: Analytical test problem: from top to bottom P_2 , P_3 , and P_4 with $D = 2$ (without noise); from left to right f_1 , f_2 , and f_3 (high-, medium-, and low-fidelity).

are modeled to never have a coincident global minimum among them. Finally, P_3 and P_4 are arbitrarily scalable in the number of dimensions, and P_4 also in the number of fidelities.

3.2 NACA Hydrofoil

This problem addresses the drag coefficient minimization of a NACA four-digit airfoil. The following mini-




Fig. 6: NACA 4-digit hydrofoil.




Fig. 7: NACA hydrofoil computational grids (G) for ISIS-CFD; from left to right G1 (12.8k cells), G2 (5.7k cells), and G3 (3.6k cells).

mization problem is solved

$$\begin{aligned} & \text{minimize } f(\mathbf{x}) = C_D(\mathbf{x}) \\ & \text{subject to } C_L(\mathbf{x}) = 0.6 \\ & \text{and to } \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \end{aligned} \quad (14)$$

where \mathbf{x} is the design variable vector, C_D and C_L are respectively the drag and lift coefficient. The equality constraint on the lift coefficient is necessary in order to compare different geometries at the same lift force (equal to the weight of the object), since the drag depends strongly on the lift.

The simulation conditions are: velocity $U = 10$ m/s, chord $c = 1$ m, fluid density $\rho = 1,026$ kg/m³, with a chord based Reynolds number $Re = 8.41 \cdot 10^6$.

The hydrofoil shape (see Figure 6) is defined by the general equation for four-digit NACA foils [26]. The upper (y_u) and lower (y_l) hydrofoil surfaces are computed as

$$\begin{cases} \xi_u = \xi - y_t \sin \theta \\ \xi_l = \xi + y_t \sin \theta \\ y_u = y_c + y_t \cos \theta \\ y_l = y_c - y_t \cos \theta \end{cases} \quad (15)$$

with $\theta = \arctan(dy_c/d\xi)$ and

$$y_c = \begin{cases} \frac{m}{p^2} \left[2p \frac{\xi}{c} - \left(\frac{\xi}{c} \right)^2 \right], & 0 \leq \xi < pc \\ \frac{m}{(1-p)^2} \left[(1-2p) + 2p \frac{\xi}{c} - \left(\frac{\xi}{c} \right)^2 \right], & pc \leq \xi \leq c \end{cases}$$

(16)

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:

$$y_t = 5t(0.2969\sqrt{\xi} - 0.1260\xi + 0.3516\xi^2 + 0.2843\xi^3 - 0.1015\xi^4). \quad (17)$$

In this work, three design spaces are defined. For $D = 1$, $\mathbf{x} = \{m\}$ with $m \in [0.025, 0.065]$; the thickness and maximum camber position are fixed at $t = 0.030$ and $p = 0.4$. For $D = 2$ the design variables vector is defined as $\mathbf{x} = \{t, m\}$ with $t \in [0.030, 0.120]$ and m, p like for $D = 1$. Finally, for $D = 3$ the design variables are $\mathbf{x} = \{p, t, m\}$ with $p \in [0.25, 0.70]$ and m, t like for $D = 2$.

Tests are run with one, two, and three fidelity levels ($N = 1, 2, 3$). The optimization budget is fixed at $CC = 45$ for all design spaces.

3.3 DTMB 5415 Model

The SDD problem concerns the hull-shape optimization of the DTMB 5415 destroyer considering the free water surface. The shape of the DTMB 5415 destroyer is optimized for minimal resistance (R_T). The optimization

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

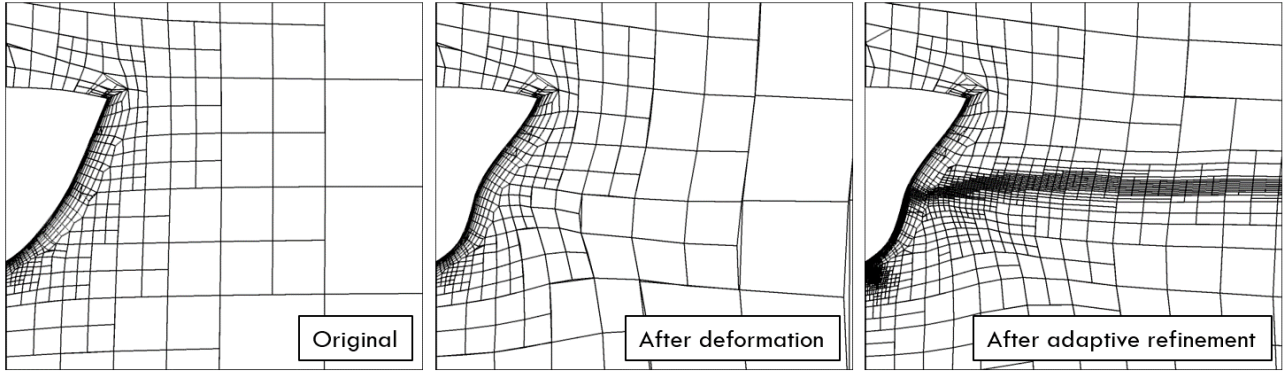


Fig. 8: The different steps of the ISIS-CFD grid creation for the DTMB 5415 optimization.

problem reads

$$\begin{aligned} & \text{minimize } f(\mathbf{x}) = R_T(\mathbf{x}) \\ & \text{subject to } L_{pp}(\mathbf{x}) = L_{pp,0} \\ & \text{and to } \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \end{aligned} \quad (18)$$

where $L_{pp,0} = 5.72$ m (model scale) is the original length between perpendiculars. The ship is at even keel, with Froude number $Fr = 0.30$ and $Re = 1.18 \cdot 10^7$. The L_{pp} constraint is automatically satisfied by the shape modification method.

The modified geometries (\mathbf{g}) are produced by the linear superposition of D orthonormal basis functions (ψ) on the original geometry (\mathbf{g}_0), as follows

$$\mathbf{g}(\boldsymbol{\xi}, \mathbf{x}) = \mathbf{g}_0(\boldsymbol{\xi}) + \delta(\boldsymbol{\xi}, \mathbf{x}), \quad (19)$$

with

$$\delta(\boldsymbol{\xi}, \mathbf{x}) = \sum_{k=1}^D x_k \psi_k(\boldsymbol{\xi}), \quad (20)$$

where $\boldsymbol{\xi}$ are the geometry Cartesian coordinates, whereas $-1.25 \leq \{x_k\}_{k=1}^D \leq 1.25$ and $\{\psi_k\}_{k=1}^D$ are the reduced design variables and the eigenfunctions, respectively, provided by the design-space augmented dimensionality reduction procedure described in [33]. Details about the original design space definition can be found in [35]. In this work, two design variables are used.

The optimization is performed with $N = 1, 2$, and 3 fidelity levels. For the initial sample plane (only for this problem), CFD simulations for all fidelities were run in the center of the domain, and with each design variable at either -1 or $+1$.

3.4 RoPax Ferry

The optimization of the RoPax ferry pertains to the minimization of the resistance over the ship displace-

ment (∇):

$$\begin{aligned} & \text{minimize } f(\mathbf{x}) = R_T(\mathbf{x})/\nabla(\mathbf{x}) \\ & \text{subject to } \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}. \end{aligned} \quad (21)$$

The design variable vector is defined as two physical parameters of the aftship: $\mathbf{x} = \{\text{ABL}, \text{DF}\}$, with the aft-body length $\text{ABL} \in [0.3, 0.61315]$ and the draught factor $\text{DF} \in [0.8, 1.2]$, respectively. The original ship hull coordinates are in the domain center.

The analysis is performed for a straight-ahead advancement, with the ship at even keel condition. The operational speed is 19kn (at full scale). Computations are performed at model scale (scale factor $\lambda = 27.14$), with $Fr = 0.245$ and $Re = 1.017 \cdot 10^7$, which corresponds to a water density $\rho = 998.2$ kg/m³, kinematic viscosity $\nu = 1.105 \cdot 10^{-6}$ m²/s, and gravitational acceleration $g = 9.81$ m/s. Free water surface effects are considered.

The parametric geometry of the RoPax is realized with the computer-aided design environment integrated in the CAESES software by FRIENDSHIP SYSTEMS AG. The deformation of the hull surface is obtained by imposing the design variable values into the parametric model in CAESES. A surface grid of the RoPax ferry (i.e. the grid discretizing the hull surface) provides the displacement of the nodes on the hull surface.

The next step is the interpolation of the deformation vector from the surface grid to the volume grid. This is done in two steps: (1) the deformation of the hull surface is interpolated from the CAESES surface grid onto the patches on the hull surface of the hydrodynamic volume grid (the interpolation is performed using a system of RBFs); (2) the deformation of the hull surface is propagated in the volume grid (vertices are moved along coordinate lines normal to the surface, with the displacement of the nodes decaying with the distance).

The optimization is performed with $N = 4$ fidelity levels.

4 CFD solvers

CFD simulations for the NACA hydrofoil and the DTMB 5415 are performed with the URANS solver ISIS-CFD developed at ECN – CNRS [29], available in the FINE/Marine computing suite from Cadence Design Systems. The hydrodynamics performance of the RoPax is assessed by the URANS code χ navis developed at CNR-INM [7, 8, 4].

4.1 ISIS-CFD

ISIS-CFD is an incompressible unstructured finite-volume solver for multifluid flow. The velocity field is obtained from the momentum conservation equations and the pressure field is extracted from the mass conservation constraint transformed into a pressure equation. These equations are similar to the Rhie and Chow SIMPLE method [30] but have been adapted for flows with discontinuous density fields. Free-surface flow is simulated with a conservation equation for the volume fraction of water, discretized with specific compressive discretization schemes. The method features sophisticated turbulence models, such as an anisotropic EASM model and DES models.

The unstructured discretization is face-based. While all unknown state variables are cell-centered, the systems of equations used in the implicit time stepping procedure are constructed face by face. Therefore, cells with an arbitrary number of arbitrarily-shaped constitutive faces are accepted. The code is fully parallel using the message passing interface (MPI) protocol. A detailed description of the solver is given by [29]. Information on the interface-capturing scheme can also be found in [46].

Computational grids are created through adaptive grid refinement [44, 45], to optimize the efficiency of the solver and to simplify the automatic creation of suitable grids. The adaptive grid refinement method adjusts the computational grid locally, during the computation, by dividing the cells of an original coarse grid. 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, which gives a crude indication of the local truncation errors. The grid is refined until the dimensions $\mathbf{d}_{p,j}$ ($j = 1, 2, 3$) of each hexahedral cell p satisfy

$$\|\mathcal{C}_p \mathbf{d}_{p,j}\| = T_r. \quad (22)$$

The refinement criterion based on the second derivatives of the flow is not very sensitive to grid refinement [45], so the cell sizes everywhere are proportional to the constant threshold T_r .

For the MF optimization, grid adaptation is used to take into account the need for several fidelities. The interest of this procedure is that different fidelity results can be obtained by running the same simulations and simply changing the threshold T_r . Thus, it is straightforward to automate the MF simulations.

4.2 χ navis

χ navis is a general purpose URANS solver based on a finite volume discretization, with variables co-located at cell centers. Turbulent stresses are taken into account by the Boussinesq hypothesis; several turbulence models (both algebraic and differential) are implemented. The free surface is taken into account through a single-phase level set algorithm [7, 4].

In order to treat complex geometries or bodies in relative motion, the numerical algorithm is discretized on a block-structured grid with partial overlap, possibly in relative motion [9, 49]. This approach makes domain discretization and quality control of the computational grid much easier than with similar discretization techniques implemented on structured grids with adjacent blocks. Unlike standard multi-block approaches, grid connections and overlaps are not trivial and have to be calculated in the pre-processing phase. The coarse/fine grain parallelization of the RANS code is obtained by distributing the structured blocks among available distributed and/or shared memory processors (nodes); shared memory capability (threads) is used mainly for do-loop parallelization. Pre-processing tools, which allow an automatic subdivision of structured blocks and their distribution among the processors, are used for load balancing. The communication between the processors for the coarse grain parallelization is obtained using the standard MPI library, whereas the fine grain parallelization (shared memory) is achieved through the open message passing library. The efficiency of the parallel code has been examined in earlier research, showing satisfactory results in terms of acceleration for different test cases [5].

The solver uses a full multi grid–full approximation scheme (FMG–FAS), with an arbitrary number of grid levels. In the FMG–FAS approximation procedure, the solution is computed on the coarsest grid level first. Then, it is approximated on the next finer grid and the solution is iterated by exploiting all the coarser grid levels available with a V-Cycle. The process is repeated up to the finest grid level. Thus, the




Fig. 9: RoPax Ferry grids (bow detail): from left to right G1, G2, G3, and G4.

procedure provides multi-fidelity data without any additional computational cost. More details on the code implementation and applications can be found in [14, 7, 8, 4].

5 Problem Setups

The setups for each CFD-based design optimization problem are described in the following subsections.

5.1 NACA Hydrofoil

The computational 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. The hydrofoil surface is treated with a wall law and $y^+ = 60$ for the first layer. Turbulence is modeled with the standard $k - \omega$ SST model [25]. To obtain the same lift for all geometries (see Eq. 14), the angle of incidence α for the hydrofoil is adjusted dynamically during the simulations.

Up to three fidelity levels are used. The initial computational grid has 2,654 cells, the refinement threshold value T_r is set equal to 0.1, 0.2, and 0.4 from highest- to lowest-fidelity. This results in a cell size ratio of 4 : 1 between the refined fine and coarse grids. The final grids (G) have about 12.8k, 5.7k, and 3.7k cells, respectively (see Figure 7). Highest- to lowest-fidelity simulations require about 17, 9, and 5 minutes, respectively, of wall-clock time to converge. The resulting computational cost ratios are about $\beta_2 = 0.5$ and $\beta_3 = 0.3$.

5.2 DTMB 5415 Model

Simulations of the DTMB 5415 are performed on half geometries. The domain runs from $1.5L_{pp}$ in front of the

bow to $3L_{pp}$ behind the stern, up to $2L_{pp}$ laterally, and from $-1.5L_{pp}$ to $0.5L_{pp}$ vertically. Dirichlet conditions on the velocity are imposed on the inflow and side faces, pressure is imposed on the top, bottom, and outflow side. The hull is treated with a wall law and $y^+ = 60$ for the first layer. Turbulence is modelled with $k - \omega$ SST.

The grids for the simulation of different geometries are obtained through grid deformation. Each simulation starts from the same original grid (see Figure 8a). The grid is divided in layers around the hull. For each geometry $\mathbf{g}(\boldsymbol{\xi}, \mathbf{x})$, the displacement of the hull faces with respect to $\mathbf{g}_0(\boldsymbol{\xi})$ is propagated through these layers [12]. The displacements are multiplied with a weighting factor which goes from 1 on the hull to 0 on the outer boundaries, so that the latter are not deformed (see Figure 8b). The original grid is coarse, since deforming these is easier and safer than for fine grids. The final grid, including all the refinement at the free surface, is created using adaptive refinement (see Figure 8c).

The initial grid has 130k cells. The thresholds for the simulations with different fidelities are $T_r = 0.0145, 0.0072, \text{ and } 0.0036$ from coarse to fine. This implies a 4 : 1 cell size ratio between the coarsest and finest grids and results in approximately 240k, 860k, and 3.4M cells respectively. On a 20-core workstation the computations take about 1.2 hours, 4 hours, and 19 hours each. The resulting computational cost ratios are about $\beta_2 = 0.21$ and $\beta_3 = 0.06$.

5.3 RoPax Ferry

The computational grid is composed by 54 blocks, for a total of about 5.4M cells for the finest grid (only half of the domain is discretized); the domain extends to $2L_{pp}$ in front of the hull, $3L_{pp}$ behind, and $1.5L_{pp}$ on the side; a depth of $2L_{pp}$ is imposed.

The Spalart-Allmaras turbulence model is used [40]. Wall-functions are not adopted, therefore $y^+ \leq 1$

is ensured at the wall. On solid walls, the velocity is set equal to zero and a zero normal gradient is enforced on the pressure field; at the (fictitious) inflow boundary, the velocity is set to the undisturbed flow value and the pressure is extrapolated from the inside; the dynamic pressure is set to zero at the outflow, whereas the velocity is extrapolated from inner points. On the top boundary, which remains always in the air region, fluid dynamic quantities are extrapolated from inside.

Four grid levels are used (see Figure 9, from coarser to finer: G4, G3, G2, and G1), each obtained from the next finer grid with a refinement ratio equal to 2, resulting in $\beta_2 = 0.125$, $\beta_3 = 0.0156$, and $\beta_4 = 0.002$.

6 Numerical Results

The optimization results are assessed by three error metrics [36]. Knowing the position of the global optimum $\tilde{\mathbf{x}}$, these metrics characterize the normalized error in the design space, the objective function space, and Euclidean distance in the normalized $\mathbf{x} - f(\mathbf{x})$ hyperspace, respectively:

$$E_x \equiv \|\mathbf{x}^* - \tilde{\mathbf{x}}\|, \quad (23)$$

$$E_f \equiv \frac{f(\mathbf{x}^*) - f(\tilde{\mathbf{x}})}{R_1}, \quad (24)$$

$$E_t \equiv \sqrt{\frac{E_x^2 + E_f^2}{2}}, \quad (25)$$

where \mathbf{x} is the array of the design variables (normalized in unit hypercube), \mathbf{x}^* is the location of the optimum of the approximation to f , and R_1 is the range of the highest-fidelity level computed considering the initial training set¹. The error metrics E_x and E_f evaluate design and goal accuracy, whereas the aggregate metric E_t evaluates accuracy balancing the performance quantification of the method when optima are in very flat or very peaky portions of the design space. It may be noted that Eq. 24 uses an evaluation of the objective with the highest fidelity level at the point \mathbf{x}^* identified by the surrogate as the global optimum.

In the absence of a reference optimum (as often occur in SDD problems), a different set of design-sensitive metrics are employed. These metrics quantify design point location and objective function, respectively:

$$\Delta_x \equiv \frac{\|\mathbf{x}^* - \mathbf{x}_0\|}{\sqrt{N}} \quad (26)$$

¹ Only for the P_3 problem R_1 is reduced by a factor of 500, since the benchmark has a variation of three orders of magnitude in the considered domain yielding a noise almost impossible to filter out.

$$\Delta_f \equiv \frac{f(\mathbf{x}^*) - f(\mathbf{x}_0)}{f(\mathbf{x}_0)}, \quad (27)$$

where \mathbf{x}_0 is the original objective function value, meaning that Δ_x evaluate the distance of the global optimum position from the original design in the design variable space, whereas Δ_f provides the objective function variation with respect to the parent design. Additionally, the prediction error is used to quantify the error of the surrogate model in predicting the minimum value:

$$E_p = \frac{\hat{f}(\mathbf{x}^*) - f(\mathbf{x}^*)}{R_1}, \quad (28)$$

Whenever different fidelities are tested, the single fidelity surrogate model is based on the highest-fidelity level available, the two-fidelity surrogate model is based on the highest- and the lowest-fidelities, whereas an intermediate-fidelity levels are added for three- and higher fidelity surrogate models.

6.1 Analytical Test Problems

The computational cost of the analytical test problems is negligible, therefore an artificial computational cost is defined as $\beta_1 = 1$, $\beta_2 = 0.2$, and $\beta_3 = 0.1$. The performance of the method is assessed using $N = 1, 2, 3$ fidelity levels. A computational budget equal to $40 + 5D$ is used. Since the noise is synthetically added to the analytical functions by a numerical generator of random numbers, a statistical analysis [15] is performed varying the seed of the random number generator. For each problem, problem size, and number of fidelity 50 repetitions are performed. The statistics of the three metrics (for the maximum computational cost) are reported in Table 2 and discussed using box plots. The box plot shows the q_1 , q_2 (median), and q_3 quartiles, while the lower and upper whiskers are given by $q_1 - 1.5\text{IQR}$ and $q_3 + 1.5\text{IQR}$, respectively, with $\text{IQR} = q_3 - q_1$ the interquartile range.

Figure 10 shows the results for the problem P_1 . For all the metrics the lowest median is achieved with $N = 3$. Furthermore, as the number of fidelities increases the IQR decreases. Finally, with $N = 3$ the median is always closer to q_1 than to q_3 .

For the problem P_2 (see Fig. 11), overall the MF ($N = 2, 3$) surrogate models achieve better results than the HF ($N = 1$) surrogate. Specifically, the median of the MF surrogates is lower than the HF surrogate and also the IQR is smaller. Finally, the results for $N = 2$ and $N = 3$ are very close, with $N = 3$ achieving slightly better results than $N = 2$, see Table 2.

Figure 12 shows the metrics for the P_3 problem with $D = 2, 5, 10$ (from top to bottom). For $D = 2$ the




Fig. 10: Numerical results, analytical test problem P_1 ; from left to right E_x , E_f , and E_t .




Fig. 11: Numerical results, analytical test problem P_2 ; from left to right E_x , E_f , and E_t

$N = 1, 2$ surrogates have negligible IQR, this is due to the fact that they use many high-fidelity function evaluations (namely 50 and 45, respectively) and that the noise for the high-fidelity function is small. Therefore, the active learning method is only slightly disturbed by the training set noise and performs an almost equal sampling for each repetition. Differently, when $N = 3$ a non-negligible IQR is shown since the lowest-fidelity function has a smaller range and the active learning method is therefore more affected by the training set noise. For $D = 2$ the HF ($N = 1$) surrogate achieves the best performance. For $D = 5$ the IQR is significant for all the surrogates, in particular the $N = 3$ shows the largest IQR but also the minimum median value for all the metrics. Finally, for $D = 10$ the MF surrogate with $N = 3$ shows the smaller median value for E_x and E_t and the smaller IQR for all the metrics.

Finally, Fig. 13 shows the error metrics for the P_4 problem with $D = 2, 5, 10$ (from top to bottom, respectively). For $D = 2$ and $D = 5$ the MF surrogate with $N = 2$ achieves the lowest E_t error, whereas the smallest IQR is achieved by $N = 1, 3$ surrogates for $D = 2, 5$, respectively. For $D = 10$ the MF ($N = 2, 3$)

surrogates achieve the smallest E_t values, where $N = 3$ shows a larger IQR than $N = 2$.

Table 2 summarizes the median of the surrogate models performance for the analytical problems. Overall, the MF ($N = 2, 3$) surrogate models outperform the HF ($N = 1$) surrogate model especially when $D > 2$. Furthermore, the MF surrogate model with $N = 3$ outperforms the MF surrogate model with $N = 2$ in most of the cases.

6.2 NACA Hydrofoil

For this problem a reference global minimum is available that is $\tilde{\mathbf{x}} = [0.3776, 0.3333, 0.0000]$ with $f(\tilde{\mathbf{x}}) = 7.2116E - 3$.

Figures 14 and 15 show the global surrogate model prediction with $D = 1, 2$, respectively, at the final iteration of the active learning procedure. The active learning method is able to identify the global minimum region for all cases ($N = 1, 2, 3$). For $D = 1$, using a single-fidelity surrogate model (see Figure 14, left) the noise in the CFD outputs is negligible. Differently, the use of two fidelities (see Figure 14, center) introduces a significant amount of noise (due to the low-fidelity

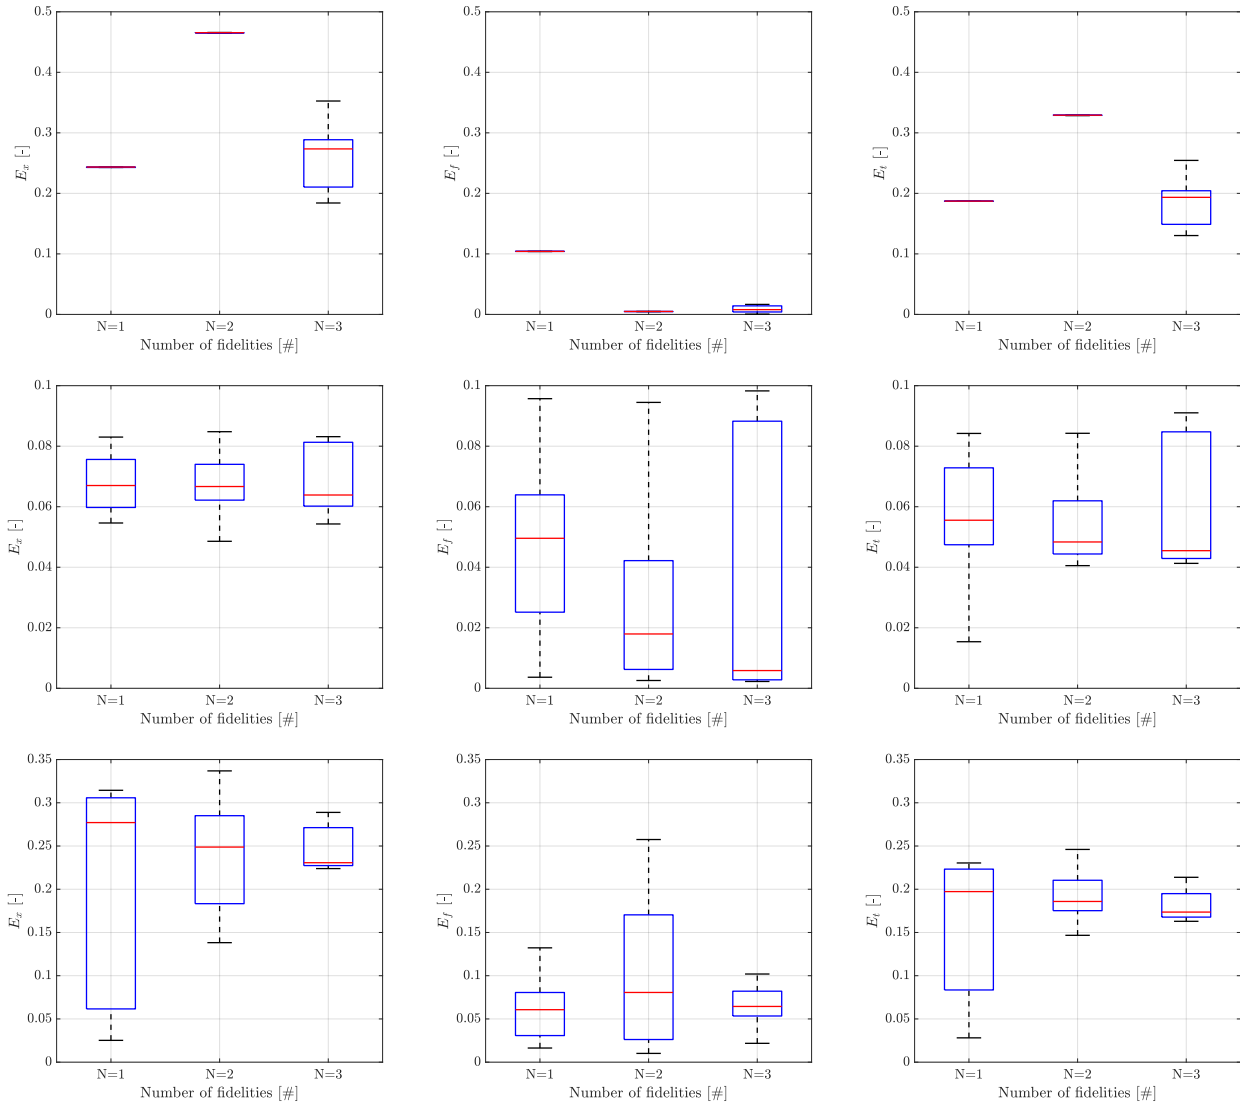


Fig. 12: Numerical results, analytical test problem P_3 ; from left to right E_x , E_f , and E_t , from top to bottom $D = 2, 5, 10$.

CFD outputs), negatively affecting the MF prediction. Finally, with an intermediate-fidelity level (see Figure 14, right) the noise is still present, but it is filtered out more effectively by the addition of a less noisy intermediate fidelity model. The improvement of the MF surrogate model prediction when adding fidelities can be associated with the different number of optimal RBF centers (see Eq. 10) used for the lowest fidelity: $\mathcal{K}^* = 20$ and 13, for $N = 2$ and 3, respectively. A similar behavior is found for $D = 2$ (see Fig. 15).

Table 3 summarizes the results for the NACA hydrofoil optimization problem. The error metrics for $D = 1, 2, 3$ are comparable. Although the lowest errors are generally achieved by the HF ($N = 1$) surrogate model, the MF ($N = 2, 3$) surrogate models achieve

very similar values although the numerical noise in the MF training sets is significantly higher than in the HF training set. This shows that the proposed approach is effective in filtering out the numerical noise. Furthermore, $N = 3$ always outperforms $N = 2$.

An analysis of the noise behavior in the simulations is given in Figure 16, which shows the evolution of the drag for a systematic variation of the foil thickness over a small range. The drag on adaptively refined grids is compared with a systematic series of grids created directly by the Hexpress grid generator. These results confirm that for this case, the noise is mostly due to the adaptive refinement procedure, since the grids without adaptation produce a much smoother behavior. These oscillations in the forces are related to small changes in

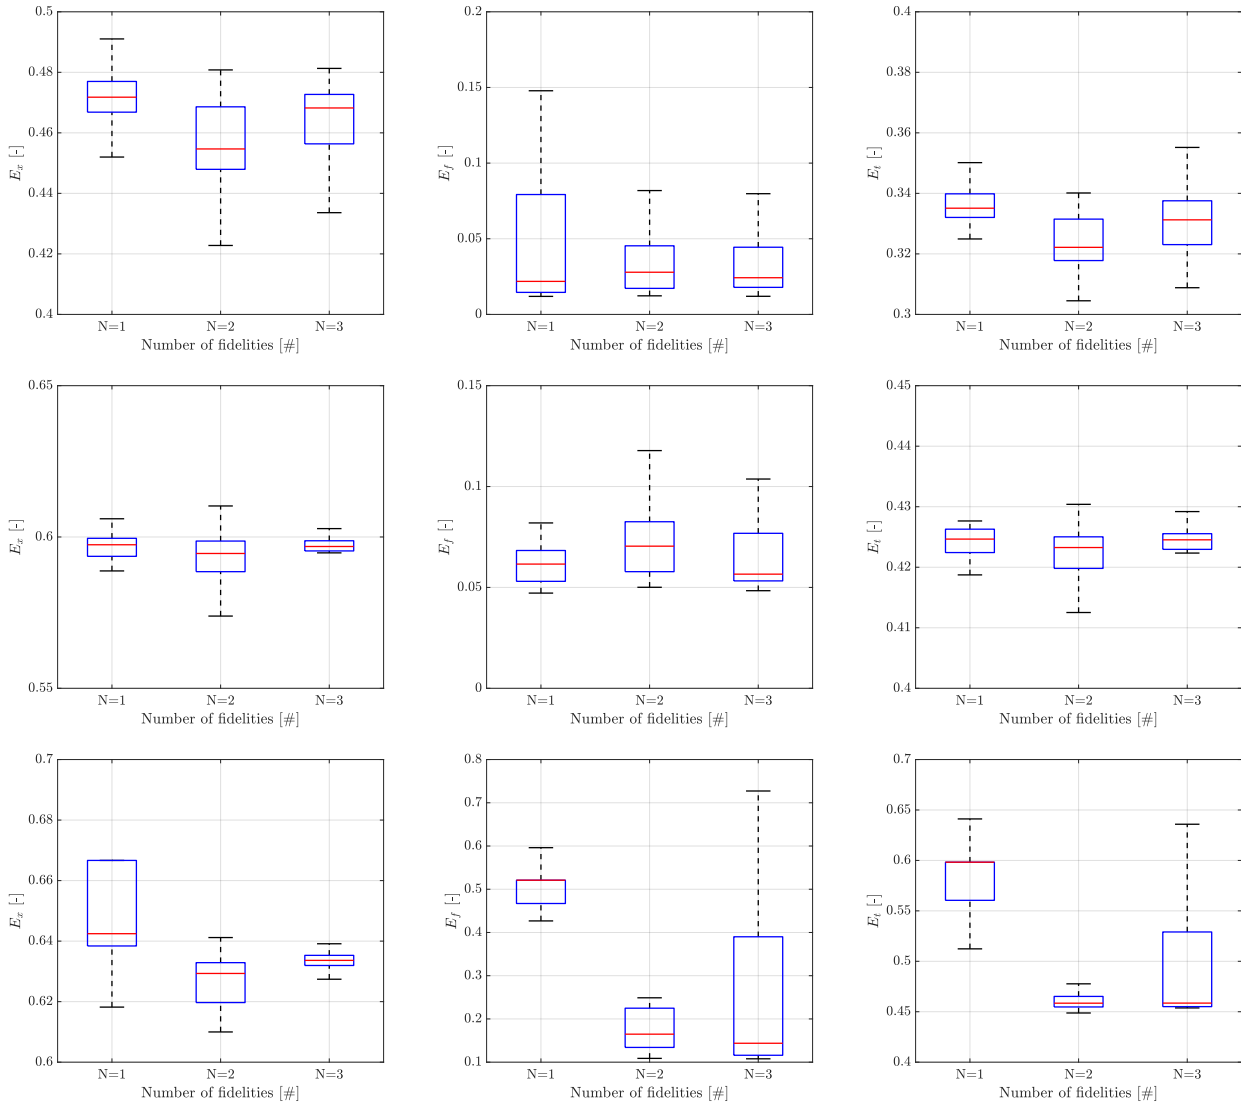


Fig. 13: Numerical results, analytical test problem P_4 ; from left to right E_f , E_x , and E_t , from top to bottom $D = 2, 5, 10$.




Fig. 14: NACA hydrofoil surrogate models and associated prediction uncertainty for $D = 1$; from left to right $N = 1, 2, 3$.

the topology of the adapted grids (i.e. cells, especially in the boundary layers and at the leading edge, which

are either refined or not depending on small changes in the hydrofoil geometry). Also, the dynamic adjustment

Table 2: Analytical test, summary of the results.

Test	D	N	CC	median[E_x]%	median[E_f]%	median[E_t]%	median[\mathcal{J}_1]	median[\mathcal{J}_2]	median[\mathcal{J}_3]
P_1	1	1	45	50.5	33.9	45.1	45	-	-
	1	2	45	2.87	2.70	2.79	27	-	183
	1	3	45	0.98	0.34	0.73	10	85	182
P_2	2	1	50	4.42	6.27	5.59	50	-	-
	2	2	50	2.68	2.25	2.56	26	-	235
	2	3	50	2.65	2.08	2.38	8	113	184
P_3	2	1	50	24.3	10.4	18.7	50	-	-
	2	2	50	46.55	0.49	32.9	45	-	55
	2	3	50	27.3	0.79	19.3	20	101	101
	5	1	65	6.70	4.96	5.55	65	-	-
	5	2	65	6.67	1.80	4.83	52	-	138
	5	3	65	6.39	0.59	4.54	15	165	167
	10	1	90	27.7	6.08	19.7	90	-	-
	10	2	90	24.9	8.07	18.6	72	-	182
	10	3	90	23.1	6.45	17.3	44	137	199
P_4	2	1	50	47.2	2.18	33.5	50	-	-
	2	2	50	45.5	2.79	32.2	9	-	410
	2	3	50	46.8	2.42	33.1	7	34	357
	5	1	65	59.74	6.16	42.5	65	-	-
	5	2	65	59.45	7.05	42.3	16	-	490
	5	3	65	59.7	5.66	42.4	15	25	441
	10	1	90	64.2	52.1	59.8	90	-	-
	10	2	90	62.9	17.7	45.8	27	-	620
	10	3	90	63.3	14.1	45.8	26	40	554

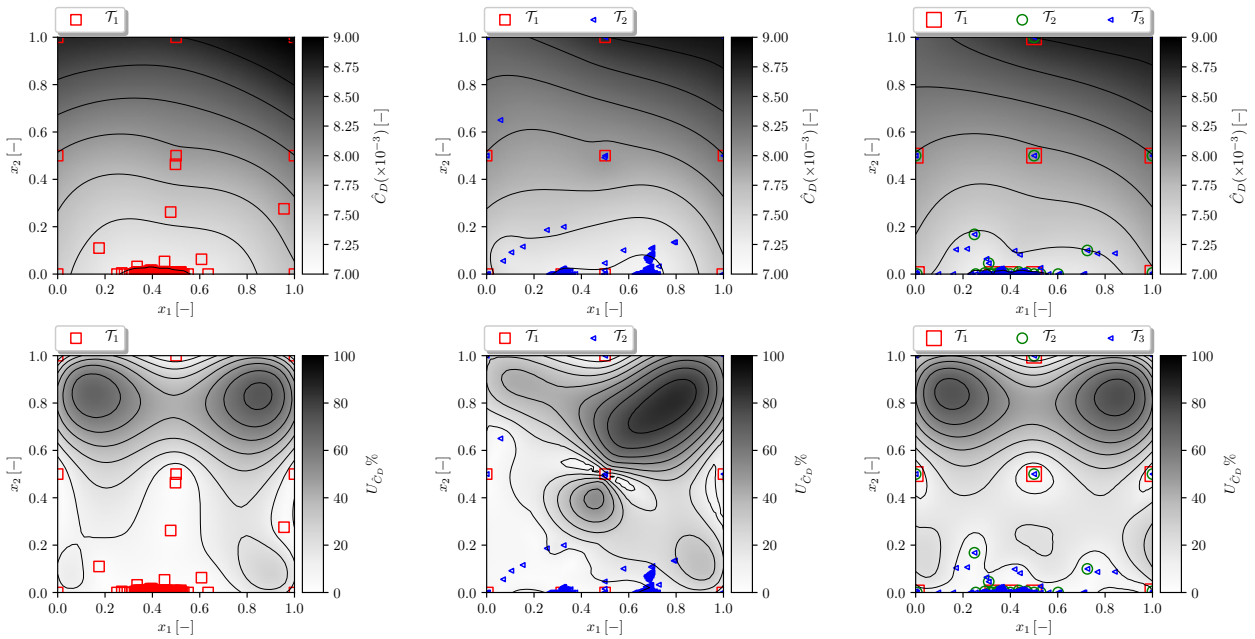


Fig. 15: NACA hydrofoil surrogate models (top) and associated prediction uncertainty (bottom) for $D = 2$: from left to right $N = 1, 2, 3$.

of the foil angle of attack depends, within the tolerances of the algorithm, on the history of the forces. Thus, if the history of the grid refinement is different, the con-

verged angle of attack may vary, even if the final grid topology is the same. This makes it difficult to identify a single cause for the noise.

However, the noise is proportional to the numerical errors of the simulations: it is most pronounced on the coarsest grid and disappears rapidly as the adapted grids become finer. Also, based on a rough estimation of the grid convergence for the two series, the adapted grids produce similar accuracy as non-adapted grids with four to five times more cells, thanks to the excellent capturing of the flow around the leading edge on the adapted grids (see Figure 8), that is essential to obtain the right forces for a lifting hydrofoil. Thus, the main benefit of the adaptive refinement here is a gain in efficiency for the CFD simulation.

Finally, even in the Hexpress-only grids, jumps in the drag can be observed when the topology of the grids changes. However, these jumps are less frequent than for the adaptively refined grids and thus, much harder to predict. It may actually be preferable to have noise everywhere, which can be filtered with a procedure like the one described in this paper, rather than apparently smooth CFD results but with some unpredictable local jumps.

Figure 17 shows the surrogate models with $N = 1, 2, 3$ and the associated prediction uncertainty for the DTMB 5415 optimization problem. The MF method with $N = 3$ identifies two minimum regions in the neighborhood of $(0.15, 0.75)$ and $(0.65, 0.2)$. However, inspection of the actual highest-fidelity data reveals that the first optimum may be a numerical artifact of the surrogate model. Nevertheless, the active learning procedure correctly ignores the upper-right and lower-left corners, where the prediction uncertainty is high but the objective function value is high too. The optimization results are summarized in Table 4. It may be noted that for $N = 2$ not all the available budget is used, since the active learning process was prematurely terminated. The numerical noise of the low-fidelity evaluations was forcing the sampling to cluster samples in previously sampled areas without improving the surrogate model, which made further simulations useless. For $N = 1$ the lowest prediction error is achieved, this is expected since the high-fidelity evaluations are less affected by numerical noise. However, for the MF with $N = 3$, the introduction of an intermediate fidelity level has significantly reduced the prediction error of the surrogate model. Finally, the largest reduction of the total resistance is achieved with $N = 3$. The optimal hull shape is compared to the original in Figure 18, whereas Figure 19 shows the original and optimized wetted area. It is worth noting that the optimized hull has a completely dry stern.

6.3 DTMB 5415 Model

Figure 20 (left) compares the pressure distribution along the optimal and the original hull surfaces. The optimized hull has a stronger pressure gradient along the hull, but the low pressure zone is mostly perpendicular to the flow direction, so it has little influence on the drag. Figure 20 (right) shows the wave elevation of the original and optimized hull, which indicates the main reason for the resistance reduction. The optimal geometry has a bulge behind the stern which creates a second bow wave, out of phase with the first one. The two waves cancel, which produces a flattened free-surface in comparison with the original hull. This indicates that the single-speed optimum shape is dependent on the Froude number Fr : since the wavelengths change with Fr , they do not cancel in off-design conditions.

Table 3: NACA hydrofoil optimization problem, summary of the results.

D	N	CC	$E_x\%$	$E_f\%$	$E_t\%$	\mathcal{J}_1	\mathcal{J}_2	\mathcal{J}_3
1	1	45	0.10	0.14	0.17	45	-	-
1	2	45	0.07	0.11	0.13	14	-	103
1	3	45	0.04	0.11	0.12	7	10	104
2	1	45	0.00	0.03	0.03	45	-	-
2	2	45	0.27	0.95	0.99	6	-	130
2	3	45	0.02	0.08	0.08	7	19	96
3	1	45	0.05	0.41	0.41	45	-	-
3	2	45	0.21	0.59	0.63	14	-	104
3	3	45	0.09	0.52	0.53	9	14	95

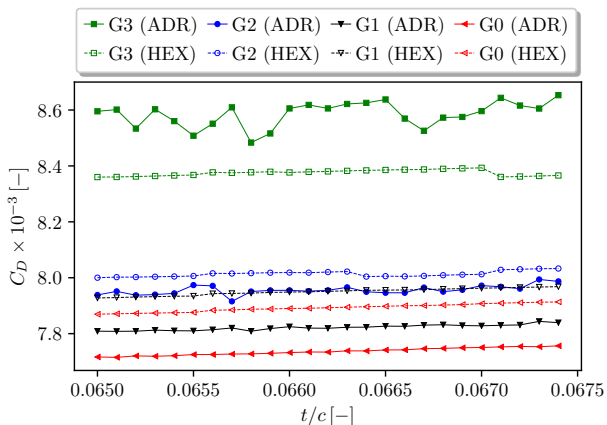


Fig. 16: NACA hydrofoil: drag coefficient as a function of t for $c = 0.0475$ and $p = 0.475$. The drawn lines are adapted grids (ADR, G3=3.6k, G2=5.4k, G1=11.8k, and G0=36.3k cells), the dashed lines are Hexpress-only grids (HEX, G3=14.4k, G2=24.5k, G1=37.4k, and G0=52.8k cells).

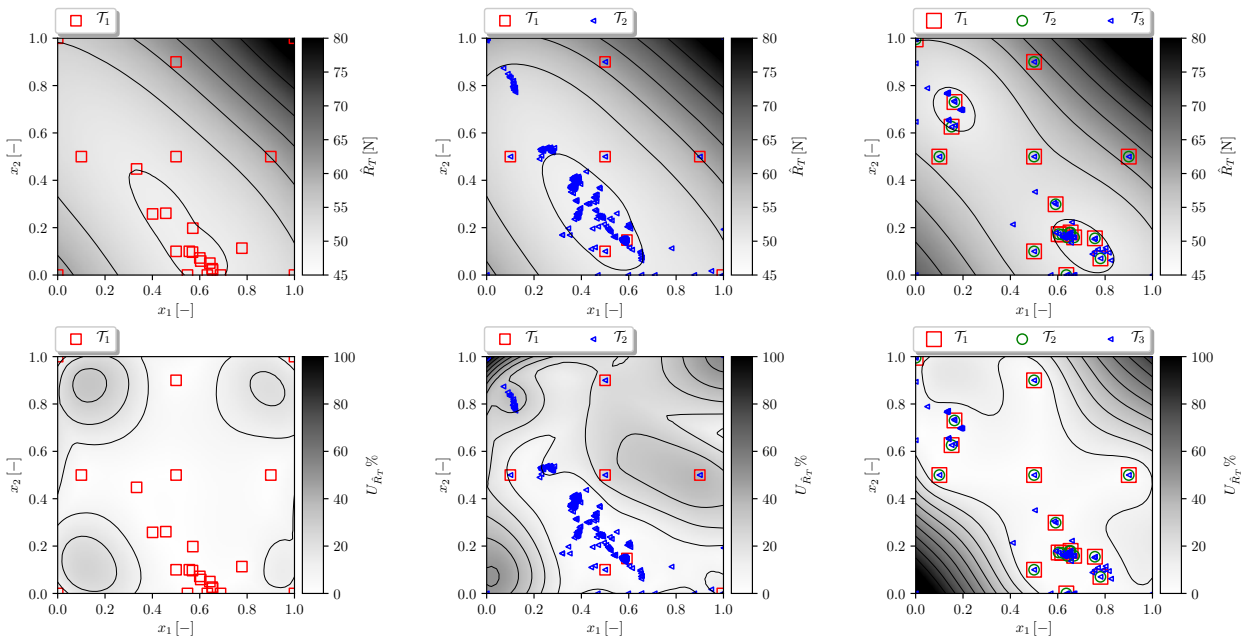


Fig. 17: DTMB 5415 surrogate models prediction (top) and associated uncertainty (bottom); from left to right $N = 1, 2, 3$.

Table 4: DTMB 5415 and RoPax SDD problems, summary of the results.

Problem	D	N	CC	x_1	x_2	Δ_x %	Δ_f %	$ E_p $ %	\mathcal{J}_1	\mathcal{J}_2	\mathcal{J}_3	\mathcal{J}_4
DTMB 5415	2	1	24.0	0.5728	0.0828	29.9	-3.9	0.71	24	-	-	-
	2	2	20.5	0.4244	0.3066	14.7	-3.5	9.07	7	-	225	-
	2	3	23.5	0.5506	0.1330	26.2	-4.5	1.73	16	18	62	-
RoPax	2	4	9.38	0.9158	1.0000	46.0	-12.7	10.8	8	9	10	50




Fig. 18: DTMB 5415, comparison of hull shapes.

Like for the NACA hydrofoil, a sensitivity analysis is performed to assess the noise in the CFD data for the DTMB 5415. Figure 21 shows the resistance variation for a small change in x_2 at fixed $x_1 = 0$ (i.e. close to the optimum), for the three fidelities used, confirming the noise presence in the CFD data. As for the hydrofoil problem, there is less noise for the medium

grid than for the coarse grid. In the fine-grid solutions, the noise appears to be even less. However, the curve has a discontinuous slope at $x_2 = -0.825$.

Moreover, for this more complex case, while the noise is still related to topology changes in the adapted grids, it is possible to identify distinct sources for the noise. For the lowest-fidelity data, the noise is mostly




Fig. 19: DTMB 5415, mass-fraction on the hull-surface; original (top) and optimized (bottom), along with a stern detail (right).




Fig. 20: DTMB 5415 comparison of the original (top) and optimized (bottom) hull shapes; hull-surface pressure (left) and wave elevation pattern (right).




Fig. 21: DTMB 5415 resistance sensitivity to x_2 variation ($x_1 = 0$).

associated with modifications in the position and the smearing of the wave system. For example, Figure 22 (left) shows random fluctuations in the thickness and the position of the volume fraction discontinuity in the first trough and the shoulder wave. Since the amount of wetted surface has a large influence on the drag, these changes in the waves explain the oscillations in

the forces. Figure 22 (center and right) show that this perturbation behaves like the noise for the hydrofoil: the finer the grid, the smaller the oscillations. The stern wave crest in Figure 23 shows similar oscillations. Also the gradient change in the fine grid solutions is explained by the stern wave (see Figure 23, right): a topology change occurs at $x_2 = -0.825$ close to the stern, where a double wave ridge with a small wetted patch on the transom is replaced by a single, shallower ridge with less wetting. On the coarse grid, this effect is absent (see Figure 23, left).

The transom is oriented normal to the flow, so even a small change in wetted surface here has a large influence on the resistance. This may imply that the optimum for this case is located around the transition from partially wetted to dry transom flow – just like the optimum for the NACA hydrofoil occurs when the incoming flow is aligned with the camber line. However, the effect likely depends on the choice of the optimization problem: for a higher Froude number, the transom would remain dry for all geometries.

This topology change in the highest-fidelity simulations, which probably occurs in a more or less




Fig. 22: Bow wave: from top to bottom $x_2 = -0.85, -0.825, -0.8$, from left to right G3, G2, G1.




Fig. 23: Stern wave pattern: from top to bottom $x_2 = -0.85, -0.825, -0.8, -0.8$, from left to right G3, G2, G1.

chaotic way depending on small changes in the mesh, is an explanation of why the LCB sampling requests a large number of highest-fidelity data points close to the optimum. While this may seem wasteful, it is required to filter the noise in the highest-fidelity data, so this behavior indicates that the algorithm adapts itself to the requirements of the data.

6.4 RoPax Ferry

As for the DTMB 5415 SDD problem, no reference solution is available for the RoPax SDD one, therefore the surrogate performance is assessed by Eq. 26 and 27.

Figure 24 shows the MF surrogate model (with $N = 4$) and the associated prediction uncertainty. The method identifies two minimum regions in the neighborhood of $(1.0, 1.0)$ and $(0.9, 0.75)$. The active learning is strictly focused on the global minimum region and the

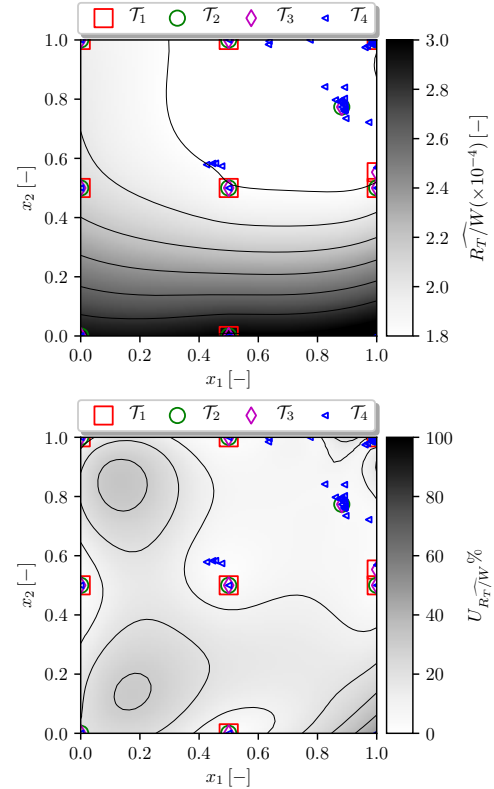


Fig. 24: RoPax problem with $N = 4$; MF prediction (top) and associated uncertainty (bottom).




Fig. 25: Original (black lines) and optimized (blue-dashed lines) shapes of the RoPax ferry with the free-surface line (FS).

overall surrogate model prediction uncertainty is low. Table 4 summarizes the results of the SDD procedure.

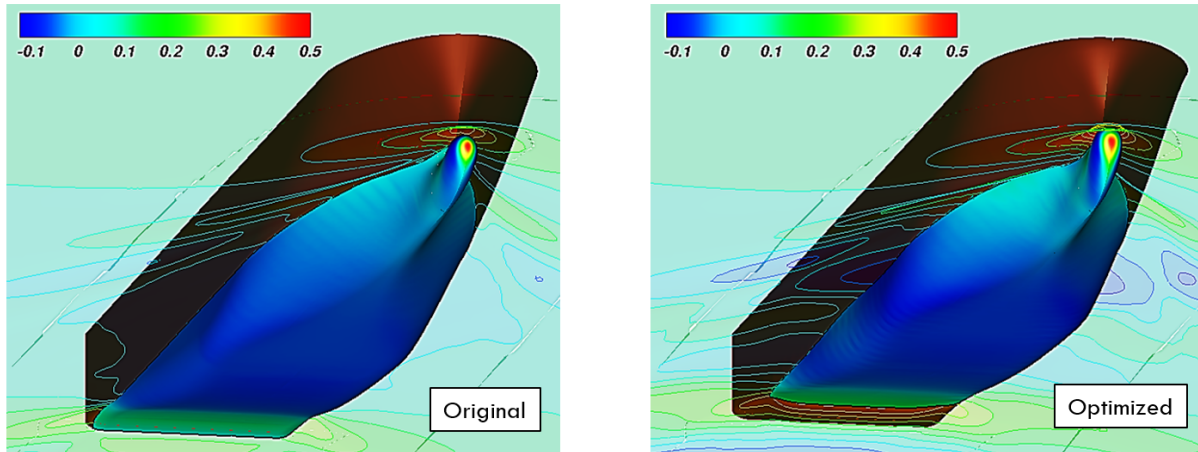


Fig. 26: RoPax ferry, non-dimensional hull-surface pressure and wave patterns. Colors represent non-dimensional pressure levels.

The MF surrogate model provides a prediction error at the minimum close to 10% and an objective function improvement equal to 12.7%.

The original and the optimized hull stations are compared in Figure 25: the optimized hull is characterized by a reduction of the submergence of the stern region and a less pronounced bulbous bow. As a consequence, surface pressure fields and wave patterns of the ship advancing straight ahead are significantly different in comparison to the original. In particular, the optimized hull shows a dry stern vault, with a reduction of the wetted area as a consequence (see Figure 26). Furthermore, the wave pattern of the optimized shape highlights a less pronounced wave throat and a stronger rooster tail, as shown in Figure 27. The bow wave of the

due to the reduced length and the increased height of the bulbous bow (see Figure 25).

7 Conclusions and Future Work

A method was presented and discussed based on a generalized multi-fidelity (MF) surrogate model and active learning, for design-space exploration and global optimization with a limited budget of function evaluations. The method is able to leverage an arbitrary number of fidelity levels and it is intended to keep a high effectiveness also in presence of training data affected by numerical noise, as often occur with CFD-based evaluations.

The proposed method is based on stochastic radial basis functions (SRBF) with least squares regression and in-the-loop optimization of the surrogate hyperparameters. The MF approximation is obtained with an additive correction of a low-fidelity trained surrogate with the surrogate of the errors between consecutive fidelities. The surrogate model is dynamically updated by an active learning procedure that automatically queries new samples at a specified fidelity. The sampling method is based on the LCB approach and considers the benefit-cost ratio associated to using different fidelity sources.

The method was tested on a set of four analytical test problems, a NACA hydrofoil optimization, the DTMB 5415 hull-form optimization, and the hull-form optimization of a RoPax ferry. The assessment of the MF surrogate model performance was performed based on its ability to identify the optimum position and the error between the verified predicted minimum and the real global optimum (when available).




Fig. 27: RoPax ferry non-dimensional wave patterns. Colors represent non-dimensional wave elevation.

optimized hull is similar to the original, with the crest for the optimized shape slightly higher and retarded

Based on the current test cases, the following conclusions can be drawn:

1. The use of a MF surrogate model is more convenient than a surrogate model based on HF only. In the analytical tests the MF surrogate model overall achieves lower errors than the HF surrogate. In most of the cases, the use of three fidelities ($N = 3$) provides lower errors than the use of two fidelities. The advantage of using MF surrogate models as opposed to HF surrogates is more significant when the problem dimensionality (D) increases. Based on the aggregate metric E_t , the MF surrogate models always outperformed the HF surrogate model for $D = 5$ and $D = 10$. Considering the NACA hydrofoil problem the MF surrogate model with three fidelities achieves lower error values than with two fidelities. The reason for this is that the extra fidelity levels add robustness, providing more reliable noise filtering. Increasing the number of fidelity levels could potentially provide even better performance. Nevertheless, this should be carefully assessed on a case by case basis considering the trade-off between accuracy and computational cost associated to additional fidelity levels. Indeed, cheaper models may lead to significant computational cost reduction but may be too inaccurate, thus misleading the MF method. Differently, fidelity levels with higher accuracy may have a quite high computational cost, thus not increasing significantly the overall efficiency.
2. The active learning process of the MF surrogate allows for a wider exploration of the design space, compared to using the HF only, as SDD problems show. This is because, in general, the lower the fidelity, the higher the associated noise (see e.g. Figs 16 and 21). Therefore, the surrogate model prediction uncertainty associated with the lowest fidelity is higher and distributed over a wider region of the design space. For this reason, the combination of LCB with the MF method provides a wide exploration of the whole design space (see Figs 14 and 15).
3. The proposed active learning method tends to cluster training points in the most noisy regions of the design space (see e.g. Figs 14, 15, and 17). This is due to the use of LCB (see Eq. 6) with a regressive model for noisy data. The prediction uncertainty associated to this model does not vanish at the training points. As a consequence, training point neighborhoods may be identified as promising regions to add new points through Eq. 6. This improves the noise identification in noisy regions, but at the same time it prevents a wider exploration of the design

space with more evenly distributed training points. A proper trade-off between noise identification and design space exploration should therefore be carefully addressed.

4. Even if the adaptive grid refinement and the multi-grid resolution methods cannot be directly compared, they perfectly fit the MF method philosophy and formulation. They represent a good example of simulation methods that naturally produce results spanning a range of fidelity levels and therefore represent a natural fit for MF methods.

Future work includes the assessment of alternative active learning approaches, through different definitions of the prediction uncertainty in the presence of noise, in order to address the question of training points clustering. Finally, the possibility of selecting the number of RBF centers via clustering metrics, such as the within-cluster sums of squares [19] or the silhouette [31], will be addressed as an alternative to the leave-one-out cross-validation metrics.

Acknowledgements CNR-INM is partially supported by the Office of Naval Research through NICOP grant N62909-21-1-2042, administered by Dr. Woei-Min Lin, Dr. Elena McCarthy, and Dr. Salahuddin Ahmed of the Office of Naval Research and Office of Naval Research Global, and the EU funded project HOLISHIP (HOListic optimisation of SHIP design and operation for life cycle), grant agreement N. 689074. The development of the methodology is conducted in collaboration with the NATO STO AVT task group on "Goal-driven, multi-fidelity approaches for military vehicle system-level design" (AVT-331).

References

1. Abdullah, J.M., Ahmed, T.: Fitness dependent optimizer: inspired by the bee swarming reproductive process. *IEEE Access* **7**, 43473–43486 (2019)
2. Baar, J.d., Roberts, S., Dwight, R., Mallol, B.: Uncertainty quantification for a sailing yacht hull, using multi-fidelity kriging. *Computers & Fluids* **123**, 185 – 201 (2015)
3. Beran, P.S., Bryson, D., Thelen, A.S., Diez, M., Serani, A.: Comparison of multi-fidelity approaches for military vehicle design. In: *AIAA AVIATION 2020 FORUM*, p. 3158 (2020)
4. Brogna, R., Durante, D.: Accurate prediction of complex free surface flow around a high speed craft using a single-phase level set method. *Computational Mechanics* **62**(3), 421–437 (2018)
5. Brogna, R., Zaghi, S., Muscari, R., Salvatore, F.: Enabling hydrodynamics solver for efficient parallel simulations. In: *2014 International Conference on High Performance Computing & Simulation (HPCS)*, pp. 803–810. *IEEE* (2014)
6. Cox, D.D., John, S.: A statistical method for global optimization. In: *[Proceedings] 1992 IEEE International Conference on Systems, Man, and Cybernetics*, pp. 1241–1246. *IEEE* (1992)

7. Di Mascio, A., Broglia, R., Muscari, R.: On the application of the one-phase level set method for naval hydrodynamic flows. *Computer & Fluids* **36**(5), 868–886 (2007)
8. Di Mascio, A., Broglia, R., Muscari, R.: Prediction of hydrodynamic coefficients of ship hulls by high-order godunov-type methods. *Journal of Marine Science and Technology* **14**(1), 19–29 (2009)
9. Di Mascio, A., Muscari, R., Broglia, R.: An Overlapping Grids Approach for Moving Bodies Problems. In: *Proceedings of 16th Int. Offshore and Polar Engineering Conference*. San Francisco, California (USA) (2006)
10. Diez, M., Campana, E.F., Stern, F.: Stochastic optimization methods for ship resistance and operational efficiency via CFD. *Structural and Multidisciplinary Optimization* **57**(2), 735–758 (2018)
11. Diez, M., Peri, D., Fasano, G., Campana, E.F.: Hydroelastic optimization of a keel fin of a sailing boat: a multidisciplinary robust formulation for ship design. *Structural and Multidisciplinary Optimization* **46**(4), 613–625 (2012)
12. Durand, M.: Light and flexible fluid/structure interaction, application to sailing boats. *Theses, Ecole Centrale de Nantes (ECN)* (2012). URL <https://hal.archives-ouvertes.fr/tel-01203748>
13. Fasshauer, G.E., Zhang, J.G.: On choosing “optimal” shape parameters for rbf approximation. *Numerical Algorithms* **45**(1-4), 345–368 (2007)
14. Favini, B., Broglia, R., Di Mascio, A.: Multi-grid Acceleration of Second Order ENO Schemes from Low Subsonic to High Supersonic Flows. *Int. J. Num. Meth. Fluids* **23**, 589–606 (1996)
15. Ficini, S., Iemma, U., Pellegrini, R., Serani, A., Diez, M.: Assessing the performance of an adaptive multi-fidelity gaussian process with noisy training data: A statistical analysis. In: *AIAA AVIATION 2021 FORUM*, p. 3098 (2021)
16. Forrester, A.I., Sóbester, A., Keane, A.J.: Multi-fidelity optimization via surrogate modelling. *Proceedings of the royal society A: mathematical, physical and engineering sciences* **463**(2088), 3251–3269 (2007)
17. Han, Z.H., Görtz, S., Zimmermann, R.: Improving variable-fidelity surrogate modeling via gradient-enhanced kriging and a generalized hybrid bridge function. *Aerospace Science and Technology* **25**(1), 177 – 189 (2013)
18. Karlberg, M., Löfstrand, M., Sandberg, S., Lundin, M.: State of the art in simulation-driven design. *International Journal of Product Development* **18**(1), 68–87 (2013)
19. Ketchen, D.J., Shook, C.L.: The application of cluster analysis in strategic management research: an analysis and critique. *Strategic Management Journal* **17**(6), 441–458 (1996)
20. Leotardi, C., Serani, A., Iemma, U., Campana, E.F., Diez, M.: A variable-accuracy metamodel-based architecture for global mdo under uncertainty. *Structural and Multidisciplinary Optimization* **54**(3), 573–593 (2016)
21. Li, X., Gao, W., Gu, L., Gong, C., Jing, Z., Su, H.: A cooperative radial basis function method for variable-fidelity surrogate modeling. *Structural and Multidisciplinary Optimization* **56**(5), 1077–1092 (2017)
22. Liu, H., Ong, Y.S., Cai, J.: A survey of adaptive sampling for global metamodeling in support of simulation-based complex engineering design. *Structural and Multidisciplinary Optimization* **57**(1), 393–416 (2018)
23. Lloyd, S.: Least squares quantization in PCM. *IEEE transactions on information theory* **28**(2), 129–137 (1982)
24. Martins, J.R., Lambe, A.B.: Multidisciplinary design optimization: a survey of architectures. *AIAA journal* **51**(9), 2049–2075 (2013)
25. Menter, F.R.: Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal* **32**(8), 1598–1605 (1994)
26. Moran, J.: An introduction to theoretical and computational aerodynamics. Courier Corporation (2003)
27. Ng, L.W.T., Eldred, M.: 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* (2012)
28. Park, C., Haftka, R.T., Kim, N.H.: Remarks on multi-fidelity surrogates. *Structural and Multidisciplinary Optimization* **55**(3), 1029–1050 (2017)
29. Queutey, P., Visonneau, M.: An interface capturing method for free-surface hydrodynamic flows. *Computers & Fluids* **36**(9), 1481–1510 (2007)
30. Rhie, C.M., Chow, W.L.: A numerical study of the turbulent flow past an isolated airfoil with trailing edge separation. *AIAA Journal* **17**, 1525–1532 (1983)
31. Rousseeuw, P.J.: Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* **20**, 53 – 65 (1987)
32. Rumpfkeil, M.P., Beran, P.S.: Multi-fidelity, gradient-enhanced, and locally optimized sparse polynomial chaos and kriging surrogate models applied to benchmark problems. In: *AIAA Scitech 2020 Forum*, p. 0677 (2020)
33. Serani, A., Diez, M.: 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, USA, June 25-29 (2018)
34. Serani, A., Diez, M., van Walree, F., Stern, F.: URANS analysis of a free-running destroyer sailing in irregular stern-quartering waves at sea state 7. *Ocean Engineering* **237**, 109600 (2021)
35. Serani, A., Fasano, G., Liuzzi, G., Lucidi, S., Iemma, U., Campana, E.F., Stern, F., Diez, M.: Ship hydrodynamic optimization by local hybridization of deterministic derivative-free global algorithms. *Applied Ocean Research* **59**, 115 – 128 (2016)
36. Serani, A., Leotardi, C., Iemma, U., Campana, E.F., Fasano, G., Diez, M.: Parameter selection in synchronous and asynchronous deterministic particle swarm optimization for ship hydrodynamics problems. *Applied Soft Computing* **49**, 313 – 334 (2016)
37. Serani, A., Pellegrini, R., Broglia, R., Wackers, J., Visonneau, M., Diez, M.: An adaptive N-fidelity metamodel for design and operational-uncertainty space exploration of complex industrial problems. In: *Proceedings of the 8th International Conference on Computational Methods in Marine Engineering (MARINE 2019)*, pp. 177–188 (2019)
38. Serani, A., Pellegrini, R., Wackers, J., Jeanson, C.E., Queutey, P., Visonneau, M., Diez, M.: Adaptive multi-fidelity sampling for CFD-based optimisation via radial basis function metamodels. *International Journal of Computational Fluid Dynamics* **33**(6-7), 237–255 (2019)
39. Serani, A., Stern, F., Campana, E.F., Diez, M.: Hull-form stochastic optimization via computational-cost reduction methods. *Engineering with Computers* pp. 1–25 (2021)
40. Spalart, P.R., Allmaras, S.R.: A one-equation turbulence model for aerodynamic flows. *AIAA Paper 92-0439* (1991)

41. Stern, F., Wang, Z., Yang, J., Sadat-Hosseini, H., Mousaviraad, M., Bhushan, S., Diez, M., Sung-Hwan, Y., Wu, P.C., Yeon, S.M., et al.: Recent progress in cfd for naval architecture and ocean engineering. *Journal of Hydrodynamics* **27**(1), 1–23 (2015)
42. Viana, F.A.C., Simpson, T.W., Balabanov, V., Vasilli, T.: Special section on multidisciplinary design optimization: Metamodeling in multidisciplinary design optimization: How far have we really come? *AIAA Journal* **52**(4), 670–690 (2014)
43. Volpi, S., Diez, M., Gaul, N., Song, H., Iemma, U., Choi, K.K., Campana, E.F., Stern, F.: Development and validation of a dynamic metamodel based on stochastic radial basis functions and uncertainty quantification. *Structural and Multidisciplinary Optimization* **51**(2), 347–368 (2015)
44. Wackers, J., Deng, G., Guilmineau, E., Leroyer, A., Queutey, P., Visonneau, M.: Combined refinement criteria for anisotropic grid refinement in free-surface flow simulation. *Computers and Fluids* **92**, 209 – 222 (2014)
45. Wackers, J., Deng, G., Guilmineau, E., Leroyer, A., Queutey, P., Visonneau, M., Palmieri, A., Liverani, A.: Can adaptive grid refinement produce grid-independent solutions for incompressible flows? *Journal of Computational Physics* **344**, 364 – 380 (2017)
46. Wackers, J., Koren, B., Raven, H.C., Ploeg, A.v.d., Starke, A.R., Deng, G.B., Queutey, P., Visonneau, M., Hino, T., Ohashi, K.: Free-surface viscous flow solution methods for ship hydrodynamics. *Archives of Computational Methods in Engineering* **18**, 1–41 (2011)
47. Wackers, J., Pellegrini, R., Serani, A., Diez, M., Visonneau, M.: Adaptive multifidelity shape optimization based on noisy CFD data. In: *Proceedings of the 2019 International Conference on Adaptive Modeling and Simulation (ADMOS 2019)*. El Campello (Alicante), Spain, 27-29 May (2019)
48. Wang, H., Jin, Y., Doherty, J.: A generic test suite for evolutionary multifidelity optimization. *IEEE Transactions on Evolutionary Computation* **22**(6), 836–850 (2017)
49. Zaghi, S., Di Mascio, A., Broglia, R., Muscari, R.: Application of dynamic overlapping grids to the simulation of the flow around a fully- appended submarine. *Mathematics and Computers in Simulation* **116**, 75–88 (2015). DOI <http://dx.doi.org/10.1016/j.matcom.2014.11.003>