ORIGINAL ARTICLE

Comparing multi-index stochastic collocation and multi-fidelity stochastic radial basis functions for forward uncertainty quantification of ship resistance

Chiara Piazzola[1] · Lorenzo Tamellini[1] · Riccardo Pellegrini[2] · Riccardo Broglia[2] · Andrea Serani[2] · Matteo Diez[2]

Abstract

This paper presents a comparison of two multi-fidelity methods for the forward uncertainty quantification of a naval engineering problem. Specifically, we consider the problem of quantifying the uncertainty of the hydrodynamic resistance of a roll-on/roll-off passenger ferry advancing in calm water and subject to two operational uncertainties (ship speed and payload). The first four statistical moments (mean, variance, skewness, and kurtosis), and the probability density function for such quantity of interest (QoI) are computed with two multi-fidelity methods, i.e., the Multi-Index Stochastic Collocation (MISC) and an adaptive multi-fidelity Stochastic Radial Basis Functions (SRBF). The QoI is evaluated via computational fluid dynamics simulations, which are performed with the in-house unsteady Reynolds-Averaged Navier–Stokes (RANS) multi-grid solver χnavis. The different fidelities employed by both methods are obtained by stopping the RANS solver at different grid levels of the multi-grid cycle. The performance of both methods are presented and discussed: in a nutshell, the findings suggest that, at least for the current implementation of both methods, MISC could be preferred whenever a limited computational budget is available, whereas for a larger computational budget SRBF seems to be preferable, thanks to its robustness to the numerical noise in the evaluations of the QoI.

Keywords Uncertainty quantification · Computational fluid dynamics · Finite volumes · Reynolds-averaged Navier–Stokes equations · Multi-index stochastic collocation · Multi-fidelity stochastic radial basis functions

✉ Chiara Piazzola
chiara.piazzola@imati.cnr.it

Lorenzo Tamellini
tamellini@imati.cnr.it

Riccardo Pellegrini
riccardo.pellegrini@inm.cnr.it

Riccardo Broglia
riccardo.broglia@cnr.it

Andrea Serani
andrea.serani@cnr.it

Matteo Diez
matteo.diez@cnr.it

[1] Consiglio Nazionale delle Ricerche-Istituto di Matematica Applicata e Tecnologie Informatiche "E. Magenes" (CNR-IMATI), Via Ferrata 5/A, 27100 Pavia, Italy

[2] Consiglio Nazionale delle Ricerche-Istituto di Ingegneria del Mare (CNR-INM), Via di Vallerano 139, 00128 Roma, Italy

1 Introduction

Aerial, ground, and water-born vehicles must perform, in general, under a variety of environmental and operating conditions and, therefore, their design analysis and optimization processes cannot avoid taking into account the stochasticity associated with environmental and operational parameters. An example is given by ships and their subsystems, which are required to operate under a variety of highly stochastic conditions, such as speed, payload, sea state, and wave heading [1]. In this context, the accurate prediction of relevant design metrics (i.e., resistance and powering requirements; seakeeping, maneuverability, and dynamic stability; structural response and failure) requires prime-principles-based high-fidelity computational tools (e.g., computational fluid/structural dynamics, CFD/CSD), especially for innovative configurations and off-design conditions. These tools are, however, generally computationally expensive, making the quantification of the relevant statistical indicators (through the use of many function evaluations)

a technological challenge. As an example, an accurate hull-form optimization based on unsteady Reynolds-Averaged Navier-Stokes (URANS) solvers under stochastic conditions may require up to 500K CPU hours on high performance computing (HPC) systems, even if computational cost reduction methods are used [1]. Similarly, a URANS-based statistically significant evaluation of ship maneuvering performance in irregular waves may require up to 1M CPU hours on HPC systems [2]. In this context, the use of efficient uncertainty quantification (UQ) methods is essential to make the design analysis and optimization processes affordable. The development and application of UQ methods for sea-vehicle problems were discussed in [3]. Moreover, the numerical UQ analysis of a high-speed catamaran was performed and discussed for calm water [4], regular [5] and irregular [6] waves conditions. An experimental UQ analysis was presented in [7] for validation purposes of the same model. The efficiency of the UQ methods is, in general, problem dependent and has to be carefully assessed. As an example, several UQ methods were compared for an airfoil benchmark problem in [8].

In general, there is by now a large consensus in the UQ and computational sciences communities on the fact that large-scale UQ analyses can only be performed by leveraging on multi-fidelity methodologies, i.e., methodologies that explore the bulk of the variability of the quantities of interest (QoI) of the simulation over coarse grids (or more generally, computationally inexpensive models with, e.g., simplified physics), and resort to querying high-fidelity models (e.g., refined grids or full-physics models) only sparingly, to correct the initial guess produced with the low-fidelity models, see e.g., [9]. Within this general framework, several approaches can be conceived, depending on the kind of fidelity models considered and on the strategy used to sample the parameter space (i.e., for what values of the uncertain parameters the different fidelity models should be queried/evaluated).

One large class of methods that has received increasing attention in this context is the family of multi-level/multi-index methods, due to its effectiveness and solid mathematical ground. The hierarchy of models considered by these methods is usually obtained by successive (most often—but not necessarily—dyadic) refinements of a computational grid. The multi-level/multi-index distinction arises from the number of hyper-parameters that are considered to control the overall discretization of the problem, i.e., how many hyper-parameters are used to determine the computational grids (e.g., one or multiple size parameters for the grid elements and/or time-stepping) and the number of samples from the parameter space to be solved on each grid (e.g., specified by a single number or by a tuple of different numbers along different directions in the parametric space).

Combining the above considerations with a specific sampling strategy over the parameter space results in different variations of the method. One approach is to use random/quasi random sampling methods: this leads to methods such as Multi-Level Monte Carlo [10, 11], Multi-Index Monte Carlo [12], Multi-Level/Multi-Index Quasi-Monte Carlo [13].

A different option is to resort to methods that build a polynomial approximation over the parameter space: methods such as Multi-Level Stochastic Collocation [14], Multi-Index Stochastic Collocation [15–18], Multi-Level Least-Squares polynomial approximation [19], etc. fall in this category. Note that the wording "Stochastic Collocation" is to be understood as a synonym of "sampling in the parametric space": it refers to the fact that the parameters of the problem can be seen as random (stochastic) variables, and sampling the parametric space can be seen as "collocating the approximation problem at points of the stochastic domain".

Another widely studied class of multi-level methods employs kernel-based surrogates such as hierarchical kriging [20], co-kriging [21], Gaussian processes [22], and radial-basis functions [23]. Additive, multiplicative, or hybrid correction methods, also known as "bridge functions" or "scaling functions" [24], are used to build multi-fidelity surrogates. Further efficiency of multi-fidelity surrogates is gained using dynamic/adaptive sampling strategies, for which the multi-fidelity design of experiments for the surrogate training is not defined a priori but dynamically updated, exploiting the information that becomes available during the training process. Training points are dynamically added with automatic selection of both their location and the desired fidelity level, with the aim of reducing the computational cost required to properly represent the function [23].

Moving away from the multi-level/multi-index paradigm, multi-fidelity methods that are based on different physical models rather than multiple discretizations have been proposed, e.g., in [25–29].

The objective of the present work is to assess and compare the use of two methods, one from each methodological family, for the forward UQ analysis of a naval engineering problem. Specifically, the performance of the Multi-Index Stochastic Collocation (MISC) and adaptive multi-fidelity Stochastic Radial Basis Functions (SRBF) methods is compared on: (i) an analytical test function and (ii) the forward UQ analysis of a roll-on/roll-off passenger (RoPax) ferry sailing in calm water with two operational uncertainties, specifically ship speed and draught, the latter being directly linked to the payload. The estimation of the expected value, variance, skewness, kurtosis, and of the probability density function (PDF) of the function and the hydrodynamic resistance of the RoPax, is presented and discussed. The test function considered in the analytical test is tailored to resemble the surrogate model of the naval engineering problem: the results of this preliminary test can then be considered as a

baseline for the assessment of the relative performances of the two methods, and help in interpreting the results of the naval test case. In the RoPax problem the hydrodynamic resistance for each value of speed and draught requested by MISC and SRBF is computed by the URANS equation solver χnavis [30–32], developed at CNR-INM. χnavis embeds a multi-grid approach for iterations acceleration, based on a sequence of grids obtained by derefining an initial fine grid. More specifically, in this work four grids are used, and leveraged by both MISC and SRBF to vary the fidelity of the simulations. Therefore, both MISC and SRBF are used as multi-index methods with only one component controlling the spatial discretization. Another relevant aspect is that χnavis is an iterative solver, and, as such, it stops as soon as a suitable norm of the residual drops below a prescribed tolerance. The fact that the RANS equations are not solved at machine precision introduces in practice some noise in the evaluation of the resistance, which needs to be dealt with during the computations of the UQ indicators (statistical moments, PDF, etc.).

A preliminary version of this work is available as proceedings of the AIAA Aviation 2020 Forum, see [33]. With respect to that version, the manuscript was significantly improved in many ways. First, the discussion on MISC is now focused on the construction of the surrogate model rather than on computing statistical moments, and in particular we added some (we believe) interesting considerations about the fact that the MISC surrogate model is *not* interpolatory, even when using nested points in the parametric space; to the best of the authors' knowledge, this fact was never mentioned in previous literature. Second, for the SRBF method applied to the RoPax UQ analysis, a methodological advancement is used. In the previous work interpolation was enforced for the early iterations. Then, when a certain number of training points was available, an optimization process was performed to automatically select the number and the position of centers of the SRBF, thus automatically selecting whether to perform interpolation or regression. Differently, in this work, the optimization process is performed since the first iteration, thus making the methodology fully adaptive. Finally, the numerical results section has been enriched by including an analytical test, by adding a reference solution for the naval problem (which is obtained by a sparse-grid sampling of the highest fidelity at our disposal), and by discussing a possible strategy to mitigate the impact of the RANS noise on the MISC framework.

The remainder of this paper is organized as follows. Section 2 introduces the general framework and notation for the UQ problem, and the two methodologies considered in this work; in particular, MISC is presented in Sect. 2.1, while SRBF is presented in Sect. 2.2. Section 3 presents the numerical results: a preliminary analytical test (see Sect. 3.2) and then the naval problem (see Sect. 3.3). Finally, a summary of the findings of the numerical tests and an outlook on future work is presented in Sect. 4.

2 Forward uncertainty quantification methods

Let us assume that we are interested in the outcome of a CFD simulation that depends on the value of N random/uncertain parameters collected in the vector $\mathbf{y} = [y_1, y_2, \ldots, y_N]$; we denote by $\Gamma \subseteq \mathbb{R}^N$ the set of all possible values of $\mathbf{y}$, and by $\rho(\mathbf{y})$ the PDF of $\mathbf{y}$ over Γ. The goal of a forward UQ analysis is to compute statistical indicators of the QoI, G, of such CFD simulation, to quantify its variability due to the uncertainties on $\mathbf{y}$. For instance, we might be interested in computing expected values and/or higher-order moments of G (in the numerical tests, we will report on mean, variance, skewness, and kurtosis, denoted by $\mathbb{E}[G]$, $\text{Var}[G]$, $\text{Skew}[G]$, and $\text{Kurt}[G]$, respectively), and the PDF of G, which completely describes its statistical variability.

This analysis is often performed by a sampling approach, i.e., the CFD simulation is run for several possible values of $\mathbf{y}$, and the corresponding results are post-processed to get the indicators of interest. For instance, the statistical moments can be approximated by weighted averages of the values obtained, while the PDF can be approximated by histograms or, e.g., kernel density methods [34, 35]. Clearly, these analyses require large datasets of evaluations of G: if computing a single instance of G requires a significant amount of computational time, obtaining the dataset can become prohibitively expensive. A possible workaround is then to replace the evaluations of G with the evaluations of a surrogate model, which is ideally a good approximation of the original G, cheap to evaluate and obtained by suitably combining together a relative small number of evaluations of G (less than what would be needed to perform the UQ analysis of the full model). The two methods that we consider in this work are both methods to construct such surrogate model, and in particular they leverage the fact that CFD simulations can be performed over multiple grid resolutions to further reduce the computational costs.

Before describing in detail each method, we need to introduce some notation. To this end, let us assume that the computational domain of our CFD simulation can be discretized by a grid with non-cubic hexahedral elements of the same size[1] and let us also assume for a moment that the level of refinement of the grid along each physical

---

[1] The assumption that all elements must be of the same size can be relaxed, but it is kept for simplicity of exposition.

direction can be specified by prescribing some integer values $\alpha_1, \alpha_2, \alpha_3$; to fix ideas, one can think, e.g., that the number of elements of the grid scales as $2^{\alpha_1} \times 2^{\alpha_2} \times 2^{\alpha_3}$, but this is not necessary. The three values of α_i are collected in a multi-index $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \alpha_3]$; prescribing the multi-index $\boldsymbol{\alpha}$ thus prescribes the computational grid to be generated. If this flexibility is not allowed by the grid-generator (or by the problem itself), it is possible to set $\alpha_1 = \alpha_2 = \alpha_3 = \alpha$, i.e., controlling the grid-generation by a single integer value α (this is actually the case for the RoPax ferry example considered in this work). The same philosophy applies both to single- and multi-patch grids, where in principle there could be up to three values α_i for each patch. In general, we assume that $\boldsymbol{\alpha}$ has d components, $\boldsymbol{\alpha} \in \mathbb{N}_+^d$. The QoI of the CFD simulation computed over the grid specified by $\boldsymbol{\alpha}$ is denoted by $G_{\boldsymbol{\alpha}}$; this could be, e.g., the full velocity field or a scalar quantity associated with it.

2.1 Multi-index stochastic collocation (MISC)

In this section, the MISC method is introduced. As already mentioned, the MISC method is a multi-fidelity method that falls under the umbrella of multi-index/multi-level methods: in particular, the single-fidelity models upon which MISC is built are global Lagrangian interpolants over Γ.

2.1.1 Tensorized Lagrangian interpolant operators

The first step to derive the MISC surrogate model is to select a sequence of collocation points for each uncertain parameter y_n, i.e., for each direction of Γ_n of Γ. For computational efficiency, these points should be chosen according to $\rho(\mathbf{y})$, and they should be of nested type (i.e., collocation grids of increasing refinement should be subset of one another). In the RoPax ferry example considered in this work, the uncertain parameters $\mathbf{y}$ can be modeled as uniform and independent random variables (see Sect. 3.3) for which we choose to employ Clenshaw–Curtis (CC) points, see, e.g., [36]. A set of K univariate CC points can be obtained as

$$t_K^{(j)} = \cos\left(\frac{(j-1)\pi}{K-1}\right), \quad 1 \leq j \leq K, \tag{1}$$

and two sets of CC points, with K_1 and K_2 points, are nested if $(K_2 - 1)/(K_1 - 1) = 2^\ell$ for some integer ℓ, see also below. Other nested alternatives for uniformly distributed parameters are Leja points [37, 38] and Gauss–Patterson points [39]. Next, we introduce the function

$$m(0) = 0, \ m(1) = 1, \ m(\beta_n) = 2^{\beta_n - 1} + 1 \text{ for } \beta_n \geq 2, \tag{2}$$

and denote by $\mathcal{T}_{n,\beta_n}$ the set of $m(\beta_n)$ CC points along y_n, i.e.,

$$\mathcal{T}_{n,\beta_n} = \left\{ y_{n,m(\beta_n)}^{(j_n)} \middle| j_n = 1, \ldots, m(\beta_n) \right\} \quad \text{for } n = 1, \ldots, N.$$

Note that this choice of m guarantees nestedness of two sets of CC points, i.e., $\mathcal{T}_{n,\beta_n} = \left\{ y_{n,m(\beta_n)}^{(j_n)} \middle| j_n = 1, \ldots, m(\beta_n) \right\}$ for $n = 1, \ldots, N$. if $\gamma \geq \beta$.

An N-dimensional interpolation grid can then be obtained by taking the Cartesian product of the N univariate sets just introduced. The number of collocation points in this grid is specified by a multi-index $\boldsymbol{\beta} \in \mathbb{N}_+^N$: such multi-index plays thus a similar role for the parametric domain Γ as the multi-index $\boldsymbol{\alpha}$ for the physical domain. We denote such tensor interpolation grid by $\mathcal{T}_{\boldsymbol{\beta}} = \bigotimes_{n=1}^{N} \mathcal{T}_{n,\beta_n}$ and its number of points by $M_{\boldsymbol{\beta}} = \prod_{n=1}^{N} m(\beta_n)$: using standard multi-index notation, they can be written as

$$\mathcal{T}_{\boldsymbol{\beta}} = \left\{ \mathbf{y}_{m(\boldsymbol{\beta})}^{(\mathbf{j})} \right\}_{\mathbf{j} \leq m(\boldsymbol{\beta})}, \quad \text{with} \quad \mathbf{y}_{m(\boldsymbol{\beta})}^{(\mathbf{j})} = \left[y_{1,m(\beta_1)}^{(j_1)}, \ldots, y_{N,m(\beta_N)}^{(j_N)} \right],$$

where $m(\boldsymbol{\beta}) = \left[m(\beta_1), m(\beta_2), \ldots, m(\beta_N) \right]$ and $\mathbf{j} \leq m(\boldsymbol{\beta})$ means that $j_n \leq m(\beta_n)$ for every $n = 1, \ldots, N$. For fixed $\boldsymbol{\alpha}$, the approximation of $G_{\boldsymbol{\alpha}}(\mathbf{y})$ based on global Lagrangian polynomials collocated at these grid points (single-fidelity approximation) has the following form

$$G_{\boldsymbol{\alpha}}(\mathbf{y}) \approx \mathcal{U}_{\boldsymbol{\alpha},\boldsymbol{\beta}}(\mathbf{y}) := \sum_{\mathbf{j} \leq m(\boldsymbol{\beta})} G_{\boldsymbol{\alpha}}\left(\mathbf{y}_{m(\boldsymbol{\beta})}^{(\mathbf{j})} \right) \mathcal{L}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}(\mathbf{y}), \tag{3}$$

where $\left\{ \mathcal{L}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}(\mathbf{y}) \right\}_{\mathbf{j} \leq m(\boldsymbol{\beta})}$ are N-variate Lagrange basis polynomials, defined as tensor products of univariate Lagrange polynomials, i.e.,

$$\mathcal{L}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}(\mathbf{y}) = \prod_{n=1}^{N} l_{n,m(\beta_n)}^{(j_n)}(y_n) \quad \text{with}$$

$$l_{n,m(\beta_n)}^{(j_n)}(y_n) = \prod_{k=1, k \neq j_n}^{m(\beta_n)} \frac{y_n - y_{n,m(\beta_n)}^{(k)}}{y_{n,m(\beta_n)}^{(k)} - y_{n,m(\beta_n)}^{(j_n)}}. \tag{4}$$

Naturally, the single-fidelity approximation $\mathcal{U}_{\boldsymbol{\alpha},\boldsymbol{\beta}}$ is more and more accurate the higher the number of collocations points in each direction. Hence, ideally one would choose both multi-indices $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ with large components, say $\boldsymbol{\alpha} = \boldsymbol{\alpha}^\star$ and $\boldsymbol{\beta} = \boldsymbol{\beta}^\star$, i.e., to consider many CFD simulations over a refined computational grid; however, this is typically infeasible due to the computational cost of a single CFD simulation.

2.1.2 MISC surrogate model

The above discussion on the costs of $\mathcal{U}_{\alpha^\star,\beta^\star}$ motivates the introduction of MISC. MISC is a multi-fidelity approximation method that replaces $\mathcal{U}_{\alpha^\star,\beta^\star}$ with a linear combination of multiple coarser $\mathcal{U}_{\alpha,\beta}$: as will be clearer later, the components of such linear combination are chosen obeying to the idea that whenever the spatial discretization α is refined, the order of the interpolation β is kept to a minimum and vice versa.

To build a MISC approximation, the so-called "detail operators" (univariate and multivariate) on the physical and parametric domains have to be introduced. They are defined as follows, with the understanding that $\mathcal{U}_{\alpha,\beta}(\mathbf{y}) = 0$ when at least one component of α or β is zero. In the following the dependence of the interpolation operator on the parameters $\mathbf{y}$ is omitted for sake of compactness. Thus, we denote by $\mathbf{e}_i$ the canonical multi-index, i.e., $(\mathbf{e}_i)_k = 1$ if $i = k$ and 0 otherwise, and define

Univariate physical detail:
$$\Delta_i^{\text{phys}}[\mathcal{U}_{\alpha,\beta}] = \mathcal{U}_{\alpha,\beta} - \mathcal{U}_{\alpha-\mathbf{e}_i,\beta} \text{ with } 1 \le i \le d;$$

Univariate parametric detail:
$$\Delta_i^{\text{param}}[\mathcal{U}_{\alpha,\beta}] = \mathcal{U}_{\alpha,\beta} - \mathcal{U}_{\alpha,\beta-\mathbf{e}_i} \text{ with } 1 \le i \le N;$$

Multivariate physical detail:
$$\boldsymbol{\Delta}^{\text{phys}}[\mathcal{U}_{\alpha,\beta}] = \bigotimes_{i=1}^{d} \Delta_i^{\text{phys}}[\mathcal{U}_{\alpha,\beta}];$$

Multivariate parametric detail:
$$\boldsymbol{\Delta}^{\text{param}}[\mathcal{U}_{\alpha,\beta}] = \bigotimes_{j=1}^{N} \Delta_j^{\text{param}}[\mathcal{U}_{\alpha,\beta}];$$

Mixed multivariate detail:
$$\boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{\alpha,\beta}] = \boldsymbol{\Delta}^{\text{param}}\left[\boldsymbol{\Delta}^{\text{phys}}[\mathcal{U}_{\alpha,\beta}]\right].$$

Observe that taking tensor products of univariate details amounts to composing their actions, i.e.,

$$\boldsymbol{\Delta}^{\text{phys}}[\mathcal{U}_{\alpha,\beta}] = \bigotimes_{i=1}^{d} \Delta_i^{\text{phys}}[\mathcal{U}_{\alpha,\beta}] = \Delta_1^{\text{phys}}\left[\cdots \left[\Delta_d^{\text{phys}}\left[\mathcal{U}_{\alpha,\beta}\right] \right] \right],$$

and analogously for the multivariate parametric detail operators $\boldsymbol{\Delta}^{\text{param}}[\mathcal{U}_{\alpha,\beta}]$. By replacing the univariate details with their definitions, we can then see that this implies that the multivariate operators can be evaluated by evaluating certain full-tensor approximations $\mathcal{U}_{\alpha,\beta}$ introduced in the previous subsection, and then taking linear combinations:

$$\begin{aligned}\boldsymbol{\Delta}^{\text{phys}}[\mathcal{U}_{\alpha,\beta}] &= \Delta_1^{\text{phys}}\left[\cdots \left[\Delta_d^{\text{phys}}\left[\mathcal{U}_{\alpha,\beta}\right] \right] \right] \\ &= \sum_{\mathbf{j}\in\{0,1\}^d} (-1)^{\|\mathbf{j}\|_1}\mathcal{U}_{\alpha-\mathbf{j},\beta};\end{aligned}$$
$$\boldsymbol{\Delta}^{\text{param}}[\mathcal{U}_{\alpha,\beta}] = \sum_{\mathbf{j}\in\{0,1\}^N} (-1)^{\|\mathbf{j}\|_1}\mathcal{U}_{\alpha,\beta-\mathbf{j}}.$$

The latter expressions are known as "combination-technique" formulations, and can be very useful for practical implementations. In particular, they allow to evaluate, e.g., $\boldsymbol{\Delta}^{\text{phys}}[\mathcal{U}_{\alpha,\beta}]$ by calling pre-existing softwares on different grids up to 2^d times in a "black-box" fashion. Analogously, evaluating $\boldsymbol{\Delta}^{\text{param}}[\mathcal{U}_{\alpha,\beta}]$ requires evaluating up to 2^N operators $\mathcal{U}_{\alpha,\beta}$ over different interpolation grids, and evaluating $\boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{\alpha,\beta}]$ requires evaluating up to 2^{d+N} operators $\mathcal{U}_{\alpha,\beta}$ over different parametric grids and physical grids. Observe that by introducing these detail operators a hierarchical decomposition of $\mathcal{U}_{\alpha,\beta}$ can be obtained; indeed, the following telescopic identity holds true:

$$\mathcal{U}_{\alpha,\beta} = \sum_{[\mathbf{i},\mathbf{j}]\le[\alpha,\beta]} \boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{\mathbf{i},\mathbf{j}}]. \tag{5}$$

As an example, the case of $d = N = 1$ (i.e., one-dimensional physical and parametric spaces) can be considered. Recalling that by definition $\mathcal{U}_{i,j} = 0$ when either $i = 0$ or $j = 0$, it can be seen that

$$\begin{aligned}\sum_{[i,j]\le[2,2]} \boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{i,j}] &= \boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{1,1}] + \boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{1,2}] + \boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{2,1}] \\ &\quad + \boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{2,2}] \\ &= \mathcal{U}_{1,1} + (\mathcal{U}_{1,2} - \mathcal{U}_{1,1}) + (\mathcal{U}_{2,1} - \mathcal{U}_{1,1}) \\ &\quad + (\mathcal{U}_{2,2} - \mathcal{U}_{2,1} - \mathcal{U}_{1,2} + \mathcal{U}_{1,1}) \\ &= \mathcal{U}_{2,2}.\end{aligned}$$

The crucial observation is that, under suitable regularity assumptions for $G(\mathbf{y})$ (see, e.g., [17, 18]), not all of the details in the hierarchical decomposition in Eq. (5) contribute equally to the approximation, i.e., some of them can be discarded and the resulting formula will retain good approximation properties at a fraction of the computational cost (roughly, the multi-indices to be discarded are those corresponding to "high-order" details, i.e., those for which $\|\alpha\|_1 + \|\beta\|_1$ is sufficiently large). Upon collecting the multi-indices $[\alpha, \beta]$ to be retained in the sum in a multi-index set $\Lambda \subset \mathbb{N}_+^{d+N}$, the MISC approximation of G can be introduced as

$$G(\mathbf{y}) \approx \mathcal{S}_\Lambda(\mathbf{y}) := \sum_{[\alpha,\beta]\in\Lambda} \boldsymbol{\Delta}^{\text{mix}}[\mathcal{U}_{\alpha,\beta}(\mathbf{y})]. \tag{6}$$

Fig. 1 Multi-index sets for the construction of the MISC approximation (in the case $d = N = 1$). **a** the gray set is downward closed, whereas adding the blue multi-index to it would result in a set not downward closed; **b** a downward closed set (in gray) and its margin (indices marked in red and blue). If Algorithm 1 reaches the gray set, it will next explore all indices marked in red (their addition to the gray set keeps the downward closedness property) but not those marked in blue. The red set is also known as "reduced margin"

(a)   **(b)**
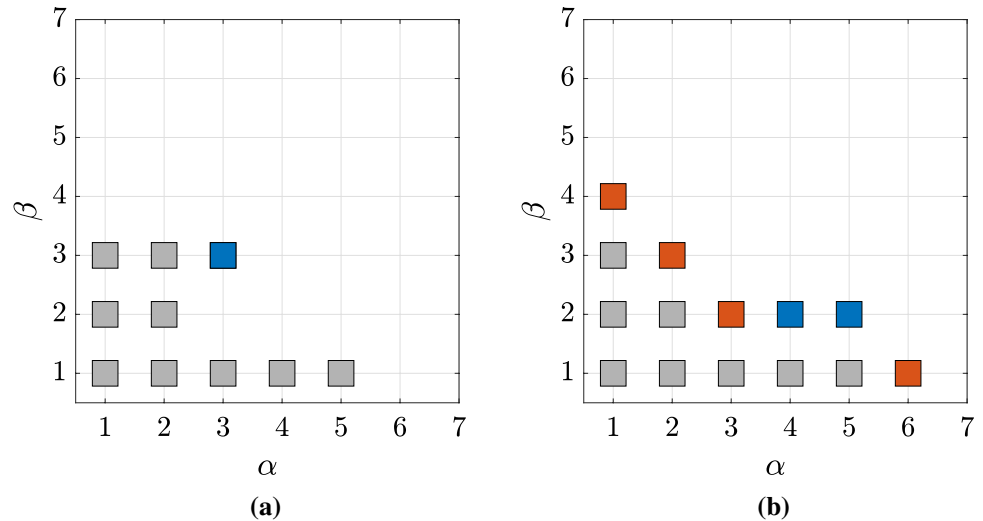
To obtain a meaningful expression, Λ should be chosen as downward closed, i.e., (see Fig. 1a)

$$\forall \mathbf{k} \in \Lambda, \quad \mathbf{k} - \mathbf{e}_j \in \Lambda \text{ for every } j = 1, \ldots, d+N \text{ such that } k_j > 1.$$

Clearly, the MISC formula in Eq. (6) has a combination-technique expression as well, which can be written in compact form as

$$S_\Lambda(\mathbf{y}) = \sum_{[\alpha, \beta] \in \Lambda} \Delta^{\text{mix}}[\mathcal{U}_{\alpha, \beta}(\mathbf{y})] = \sum_{[\alpha, \beta] \in \Lambda} c_{\alpha, \beta} \mathcal{U}_{\alpha, \beta}(\mathbf{y}), \qquad (7)$$

where the coefficients $c_{\alpha, \beta}$ are defined as

$$c_{\alpha, \beta} = \sum_{\substack{[\mathbf{i}, \mathbf{j}] \in \{0, 1\}^{d+N} \\ [\alpha + \mathbf{i}, \beta + \mathbf{j}] \in \Lambda}} (-1)^{\|[\mathbf{i}, \mathbf{j}]\|_1}.$$

This is the approximation formula which is used in our practical implementation of the MISC method, which shows our initial statement that the MISC evaluation is computed by evaluating full-tensor interpolation operators $\mathcal{U}_{\alpha, \beta}$ independently and combining them linearly, as specified by Eq. (7). Before going further, we remark a few important points:

1. The effectiveness of the MISC approximation depends on the choice of the multi-index set Λ. The optimal choice of Λ depends on the regularity assumptions on $G(\mathbf{y})$: in general, the result will be a method akin to classical multi-level schemes such as Multi-Level

Monte Carlo, where most of the statistical variability of the QoI is explored by solving many CFD simulations with coarse grids (large $\|\beta\|_1$ with small $\|\alpha\|_1$) and then the result is corrected with a few CFD simulations with refined grids (large $\|\alpha\|_1$ with small $\|\beta\|_1$). A practical adaptive algorithm to construct Λ is presented in Sect. 2.1.4, following the discussion in, e.g., [16]. Another option is to design Λ a-priori, by a careful analysis of the PDE at hand, see, e.g., [15, 17, 18].

2. MISC works well only if the levels are sufficiently separated, i.e., if the number of degrees of freedom of the computational grid (and the corresponding computational cost) grows significantly from one level to the next one: to fix ideas, one such case is if the number of elements in the grid scales, e.g., as $2^{\alpha_1} \times 2^{\alpha_2} \times 2^{\alpha_3}$, but not if, e.g., increasing α_1 to $\alpha_1 + 1$ adds only one element to the grid. If this separation does not hold, the cost of computing all the components in Eq. (7) would exceed the cost of the construction of a highly-refined single-fidelity surrogate model $\mathcal{U}_{\alpha^\star, \beta^\star}$.

3. The MISC surrogate model $S_\Lambda(\mathbf{y})$ is not interpolatory, even when nested nodes are used in the parametric space (as it is the case here). To illustrate this, let us consider as an example the case $d = 1, N = 2$, with $\Gamma = [-1, 1]^2$. We construct the MISC approximation based on the multi-index set $\Lambda = \{[1, 1\ 1], [1, 2\ 1], [2, 1\ 1]\}$: formula (7) results in

$$S_\Lambda(\mathbf{y}) = -\mathcal{U}_{[1, 1\ 1]}(\mathbf{y}) + \mathcal{U}_{[2, 1\ 1]}(\mathbf{y}) + \mathcal{U}_{[1, 2\ 1]}(\mathbf{y}),$$

where the first and the second operators are constant interpolants ($\beta = [1\ 1]$) whose value is equal to the first and second fidelity evaluated at the center of the parametric domain, $\mathbf{y}^{(C)} = [0,0]$, respectively; the third operator is instead an interpolant of degree two based on the value of the first fidelity evaluated at the following three CC bivariate points: $\mathbf{y}^{(L)} = [-1,0]$, $\mathbf{y}^{(C)} = [0,0]$, and $\mathbf{y}^{(R)} = [1,0]$. Then, evaluating the MISC approximation at, e.g., $\mathbf{y}^{(R)}$ results in

$$\mathcal{S}_\Lambda(\mathbf{y}^{(R)}) = -G_1(\mathbf{y}^{(C)}) + G_2(\mathbf{y}^{(C)}) + G_1(\mathbf{y}^{(R)}) \neq G_1(\mathbf{y}^{(R)}),$$

i.e., the value of $\mathcal{S}_\Lambda$ at $\mathbf{y}^{(R)}$ is different from the only model evaluation available at such point (which is $G_1(\mathbf{y}^{(R)})$). This is in contrast with the well-known property of single-fidelity sparse-grid surrogate models, for which the use of nested nodes over Γ guarantees that $\mathcal{S}_\Lambda(\mathbf{y}_{SG}) = G(\mathbf{y}_{SG})$ at the sparse-grid points $\mathbf{y}_{SG}$.

2.1.3 MISC quadrature

By taking the integral of the MISC surrogate model defined in Eq. (7) it is straightforward to obtain a quadrature formula $\mathcal{R}_\Lambda$ to approximate the expected value of G:

$$\mathbb{E}[G] \approx \mathbb{E}[\mathcal{S}_\Lambda] =: \mathcal{R}_\Lambda = \int_\Gamma \mathcal{S}_\Lambda(\mathbf{y})\rho(\mathbf{y})\,d\mathbf{y}$$
$$= \sum_{[\boldsymbol{\alpha},\boldsymbol{\beta}]\in\Lambda} c_{\boldsymbol{\alpha},\boldsymbol{\beta}} \int_\Gamma \mathcal{U}_{\boldsymbol{\alpha},\boldsymbol{\beta}}(\mathbf{y})\rho(\mathbf{y})\,d\mathbf{y}.$$

By recalling the definition of $\mathcal{U}_{\boldsymbol{\alpha},\boldsymbol{\beta}}$ given in Eq. (3) and of the multivariate Lagrange polynomials in Eq. (4), each of the integrals at the right-hand side of the previous formula can be rewritten in compact form as tensor quadrature operators, i.e.,

$$\mathcal{Q}_{\boldsymbol{\alpha},\boldsymbol{\beta}} := \int_\Gamma \mathcal{U}_{\boldsymbol{\alpha},\boldsymbol{\beta}}(\mathbf{y})\rho(\mathbf{y})\,d\mathbf{y}$$
$$= \sum_{\mathbf{j}\leq m(\boldsymbol{\beta})} G_{\boldsymbol{\alpha}}\left(\mathbf{y}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}\right) \left(\prod_{n=1}^N \int_{\Gamma_n} l_{n,m(\beta_n)}^{(j_n)}(y_n)\rho(y_n)\,dy_n\right)$$
$$= \sum_{\mathbf{j}\leq m(\boldsymbol{\beta})} G_{\boldsymbol{\alpha}}\left(\mathbf{y}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}\right) \left(\prod_{i=1}^N \omega_{n,m(\beta_n)}^{(j_n)}\right)$$
$$= \sum_{\mathbf{j}\leq m(\boldsymbol{\beta})} G_{\boldsymbol{\alpha}}\left(\mathbf{y}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}\right) \omega_{m(\boldsymbol{\beta})}^{(\mathbf{j})},$$

where $\omega_{n,m(\beta_n)}^{(j_n)}$, are the standard quadrature weights obtained by computing the integrals of the associated univariate Lagrange polynomials (available as analytical or tabulated values for most families of collocation points), and $\omega_{m(\boldsymbol{\beta})}^{(\mathbf{j})}$ are their multivariate counterparts. The quadrature formula $\mathcal{R}_\Lambda$ can then be understood as a linear combination of tensor quadrature operators, in complete analogy with the MISC surrogate model construction:

$$\mathcal{R}_\Lambda = \sum_{[\boldsymbol{\alpha},\boldsymbol{\beta}]\in\Lambda} c_{\boldsymbol{\alpha},\boldsymbol{\beta}} \mathcal{Q}_{\boldsymbol{\alpha},\boldsymbol{\beta}}. \tag{8}$$

Equivalently, one can also write

$$\mathcal{R}_\Lambda = \sum_{[\boldsymbol{\alpha},\boldsymbol{\beta}]\in\Lambda} \boldsymbol{\Delta}^{\mathrm{mix}}[\mathcal{Q}_{\boldsymbol{\alpha},\boldsymbol{\beta}}],$$

where the definition of $\boldsymbol{\Delta}^{\mathrm{mix}}[\mathcal{Q}_{\boldsymbol{\alpha},\boldsymbol{\beta}}]$ can be easily deduced by replacing the interpolation operators with quadrature operators in the definition of the detail operators given earlier in this section. Clearly, formula (8) easily generalizes to the computation of higher-order moments:

$$\mathbb{E}[G^r] \approx \sum_{[\boldsymbol{\alpha},\boldsymbol{\beta}]\in\Lambda} c_{\boldsymbol{\alpha},\boldsymbol{\beta}} \sum_{\mathbf{j}\leq m(\boldsymbol{\beta})} G_{\boldsymbol{\alpha}}^r\left(\mathbf{y}_{m(\boldsymbol{\beta})}^{(\mathbf{j})}\right)\omega_{m(\boldsymbol{\beta})}^{(\mathbf{j})}, \text{ with } r \geq 1. \tag{9}$$

However, this formula might not be the most effective approach to approximate $\mathbb{E}[G^r]$, especially in case of noisy evaluations of G; this aspect is discussed in more detail in Sect. 3.3.2.

We close the discussion on the MISC quadrature by remarking that the collocation points have a twofold use, i.e., they are both interpolation and quadrature points. This aspect significantly differentiates the MISC method from the approach based on radial basis functions presented in Sect. 2.2, where two distinct sets of points are considered: one for constructing a surrogate models ("training points"), and one for obtaining sample values of the surrogate models and deriving an estimate of expected value and higher moments of G ("quadrature points").

2.1.4 An adaptive algorithm for the multi-index set Λ

As already mentioned, the effectiveness of the MISC approximation depends on the choice of the multi-index set Λ: in this work such set is built with an adaptive algorithm, see [16]. We begin by introducing the following decomposition of the quadrature error

$$|\mathbb{E}[G] - \mathcal{R}_\Lambda| = \left| \mathbb{E}[G] - \sum_{[\alpha,\beta]\in\Lambda} \mathbf{\Delta}^{\mathrm{mix}}[\mathcal{Q}_{\alpha,\beta}] \right|$$

$$= \left| \sum_{[\alpha,\beta]\notin\Lambda} \mathbf{\Delta}^{\mathrm{mix}}[\mathcal{Q}_{\alpha,\beta}] \right| \le \sum_{[\alpha,\beta]\notin\Lambda} \left| \mathbf{\Delta}^{\mathrm{mix}}[\mathcal{Q}_{\alpha,\beta}] \right|$$

$$= \sum_{[\alpha,\beta]\notin\Lambda} \mathcal{E}^{\mathcal{R}}_{\alpha,\beta},$$

where $\mathcal{E}^{\mathcal{R}}_{\alpha,\beta} := \left| \mathbf{\Delta}^{\mathrm{mix}}[\mathcal{Q}_{\alpha,\beta}] \right|$. $\mathcal{E}^{\mathcal{R}}_{\alpha,\beta}$ thus represents the "error contribution" of $[\alpha,\beta]$, i.e., the reduction in the quadrature error due to having added $[\alpha,\beta]$ to the current index-set Λ. In practice, $\mathcal{E}^{\mathcal{R}}_{\alpha,\beta}$ can be conveniently computed by

$$\mathcal{E}^{\mathcal{R}}_{\alpha,\beta} = |\mathcal{R}_{\Lambda\cup[\alpha,\beta]} - \mathcal{R}_\Lambda| \tag{10}$$

for any Λ downward-closed set such that $\Lambda \cup [\alpha,\beta]$ is also downward closed. A similar quadrature-based error contribution is considered in [16], where a convex combination of the error in the computation of the mean and of the variance of the QoI is used. Another possibility is to introduce an error decomposition based on the point-wise accuracy of the surrogate model, following the same arguments above. The "error contribution" of $[\alpha,\beta]$ is then taken as

$$\mathcal{E}^{\mathcal{S}}_{\alpha,\beta} = \|\mathcal{S}_{\Lambda\cup[\alpha,\beta]} - \mathcal{S}_\Lambda\|_{L_\infty} \approx \max_{\mathbf{y}\in\mathcal{H}} \left| \mathcal{S}_{\Lambda\cup[\alpha,\beta]}(\mathbf{y}) - \mathcal{S}_\Lambda(\mathbf{y}) \right|, \tag{11}$$

where $\mathcal{H} \subset \Gamma$ is a suitable set of "testing points". Note that a similar criterion has been proposed also in the context of sparse-grid methods: different choices of $\mathcal{H}$ can be considered, depending whether nested or non-nested points are used (cf., e.g., [40, 41] and [42], respectively). In this work, we consider a set of 10000 random points (note that this operation is not expensive since it does not require evaluations of the full model).

Similarly to the "error contribution", the "work contribution" $\mathcal{W}_{\alpha,\beta}$ of $[\alpha,\beta]$ is defined as the work required to add $[\alpha,\beta]$ to the current index-set Λ. It is the product of the computational cost associated with the spatial grid identified by the multi-index α, denoted by $\mathrm{cost}(\alpha)$ (see details in Sect. 3.3.1, Eq. (29)), times the number of new evaluations of the PDE required by the multi-index β, i.e.,

$$\mathcal{W}_{\alpha,\beta} = \mathrm{cost}(\alpha) \prod_{n=1}^{N} (m(\beta_n) - m(\beta_n - 1)), \tag{12}$$

with m defined as in Eq. (2). Note that the expression above is based on the fact that the collocation points used here are nested.

We then introduce the so-called "profit" associated with the multi-index $[\alpha,\beta]$, which is defined in correspondence with the two choices of error contribution above as

$$P^{\mathcal{S}}_{\alpha,\beta} = \frac{\mathcal{E}^{\mathcal{S}}_{\alpha,\beta}}{\mathcal{W}_{\alpha,\beta}} \quad \text{or} \quad P^{\mathcal{R}}_{\alpha,\beta} = \frac{\mathcal{E}^{\mathcal{R}}_{\alpha,\beta}}{\mathcal{W}_{\alpha,\beta}}. \tag{13}$$

An effective strategy to build adaptively a MISC approximation can then be broadly described as follows: given the MISC approximation associated with a multi-index set Λ, a new MISC approximation is obtained by adding to Λ the multi-index $[\alpha,\beta] \notin \Lambda$ with the largest profit (either $P^{\mathcal{S}}$ or $P^{\mathcal{R}}$, depending on the goal of the simulation), such that $\Lambda \cup \{[\alpha,\beta]\}$ is downward closed. In practice the implementation reported in Algorithm 1 is used: it makes use of an auxiliary multi-index set, i.e., the margin of a multi-index set Λ, $\mathrm{Mar}(\Lambda)$, which is defined as the set of multi-indices that can be reached "within one step" from Λ (see Fig. 1b)

$$\mathrm{Mar}(\Lambda) = \{\mathbf{i} \in \mathbb{N}^{d+N} \text{ s.t. } \mathbf{i} = \mathbf{j} + \mathbf{e}_k$$
$$\text{for some } \mathbf{j} \in \Lambda \text{ and some } k \in \{1, \dots, d+N\}\}.$$

This algorithm was first proposed in the context of sparse-grids quadrature in [43] and its MISC implementation was first proposed in [16]. It is an a-posteriori algorithm and as such it determines the error contribution $\mathcal{E}^{\mathcal{S}}_{\alpha,\beta}, \mathcal{E}^{\mathcal{R}}_{\alpha,\beta}$ of $[\alpha,\beta]$ after having added $[\alpha,\beta]$ to the grid. Therefore, at the end of the algorithm we do not have $\mathcal{S}_\Lambda$, but actually $\mathcal{S}_J$, where J is the set of all indices whose profit has been computed, and clearly $\Lambda \subseteq J$: the richer approximation $\mathcal{S}_J$ is thus actually returned in practical implementations instead of $\mathcal{S}_\Lambda$. Finally, note that many stopping criteria can be considered (and possibly used simultaneously), which typically check that computational work, error contributions or profit estimator are below a desired threshold.

Algorithm 1: MISC implementation

Multi-Index Stochastic Collocation

$\Lambda = \{[\mathbf{1}, \mathbf{1}]\}$;
$J = \{[\mathbf{1}, \mathbf{1}]\}$;   // set of multi-indices whose profit has been computed, $\Lambda \subseteq J$
$L = \emptyset$;   // set of candidate multi-indices to be added to Λ
Compute MISC estimate $\mathcal{S}_J$ as in Eq. (7);
while *stopping criteria are not met* **do**
  for $\mathbf{j} \in Mar(\Lambda)$ **and** $\Lambda \cup \{\mathbf{j}\}$ *downward closed* **do**   // for short, $\mathbf{j} = [\boldsymbol{\alpha}, \boldsymbol{\beta}]$
    if $\mathbf{j} \notin J$ **then**
      Compute MISC estimate $\mathcal{S}_{J \cup \{\mathbf{j}\}}$ as in (7);
      Compute error contribution $\mathcal{E}_{\mathbf{j}}^{\mathcal{R}}$ (or $\mathcal{E}_{\mathbf{j}}^{\mathcal{S}}$) as in Eq. (10) (or Eq. (11));
      Compute work contribution $\mathcal{W}_{\mathbf{j}}$ as in Eq. (12);
      Compute profit $P_{\mathbf{j}}^{\mathcal{S}} = \mathcal{E}_{\mathbf{j}}^{\mathcal{S}}/\mathcal{W}_{\mathbf{j}}$ (or $P_{\mathbf{j}}^{\mathcal{R}} = \mathcal{E}_{\mathbf{j}}^{\mathcal{R}}/\mathcal{W}_{\mathbf{j}}$);
      $J = J \cup \{\mathbf{j}\}$;
      $L = L \cup \{\mathbf{j}\}$;
    end
  end
  Choose $\mathbf{i} \in L$ with the highest profit;
  $\Lambda = \Lambda \cup \{\mathbf{i}\}$;
  $L = L \setminus \{\mathbf{i}\}$;
end
return $\Lambda, J, \mathcal{S}_J$

2.2 Adaptive multi-fidelity stochastic radial basis functions (SRBF)

In this section, each of the components of the adaptive multi-fidelity SRBF surrogate model are discussed. In particular, we emphasize that here the word "stochastic" denotes not only the fact that we are sampling parameters that are affected by uncertainty, but also to the fact that the SRBF method treats one of its hyper-parameters as a random variable, as will be clear later on.

2.2.1 SRBF surrogate model

Given a training set $\mathcal{T} = \{(\mathbf{y}_i, G(\mathbf{y}_i))\}_{i=1}^{\mathcal{J}}$ and normalizing the uncertain parameters domain into a unit hypercube, the
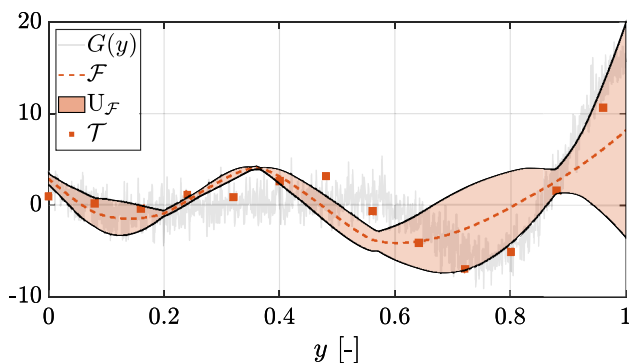
Fig. 2 SRBF example with least-squares regression
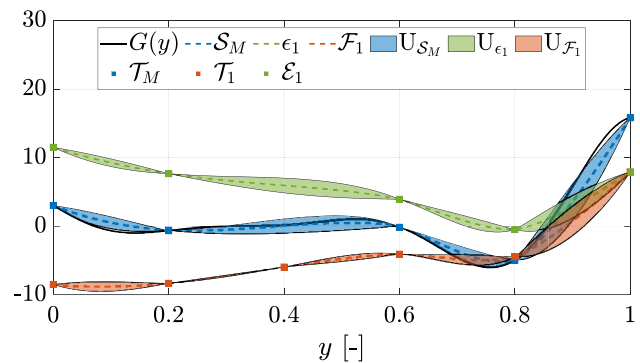
Fig. 3 Example of multi-fidelity surrogate with $M = 2$ and exact interpolation at the training points

RBF prediction is here based on a power function kernel and reads

$$f(\mathbf{y}, \tau) = \sum_{j=1}^{\mathcal{K}} w_j ||\mathbf{y} - \mathbf{c}_j||^{\tau}, \tag{14}$$

where w_j are unknown coefficients, $\mathbf{c}_j$ are $\mathcal{K}$ points in Γ called RBF centers, and $\tau \sim \text{unif}[\tau_{\min}, \tau_{\max}]$ is a stochastic tuning parameter that follows a uniform distribution. The range of τ is defined within $\tau_{\min} = 1$ and $\tau_{\max} = 3$, where $\tau = 1$ provides a polyharmonic spline of first order (linear kernel) [44] and $\tau = 3$ provides a polyharmonic spline of third order (cubic kernel) [45]. Note that the choice of the distribution for τ is arbitrary and, from a Bayesian viewpoint, this represents the degree of belief in the definition of the tuning parameter. The SRBF surrogate model $\mathcal{F}(\mathbf{y})$ is computed as the expected value (approximated by Monte Carlo) of f over τ [46]:

$$G(\mathbf{y}) \approx \mathcal{F}(\mathbf{y}) = \mathbb{E}_\tau \left[f(\mathbf{y}, \tau) \right] \approx \frac{1}{\Theta} \sum_{i=1}^{\Theta} f(\mathbf{y}, \tau_i), \tag{15}$$

where Θ is the number of samples for τ, here set equal to 1000. To give more flexibility to the method, the coordinates of the RBF centers $\mathbf{c}_j$ are not a-priori set to be coincident with the training points, but rather chosen by a k-means clustering algorithm applied to the training coordinates, see [47]. Several values of the number of centers $\mathcal{K} \leq \mathcal{J}$ are tested and their optimal number $\mathcal{K}^*$ is chosen by minimizing a leave-one-out cross-validation (LOOCV) metric, see [48]. In details, letting $g_{i,\mathcal{K}}(\mathbf{y})$, $i = 1, \ldots, \mathcal{J}$ be the surrogate models with $\mathcal{K}$ centers trained on the whole training set $\mathcal{T}$ but the i-th point, $\mathcal{K}^*$ is defined as:

$$\mathcal{K}^* = \underset{\mathcal{K} \in C}{\arg\min} \, \text{RMSE}(\mathcal{K}), \tag{16}$$

where $\mathcal{K} \leq \mathcal{J}, \mathcal{K} \in C \subset \mathbb{N}$ and $\text{RMSE}(\mathcal{K})$ is the root mean square error of the $\mathcal{J}$ leave-one-out models $g_{1,\mathcal{K}}, \ldots, g_{\mathcal{J},\mathcal{K}}$ at the point that is being left out for each $g_{i,\mathcal{K}}$:

$$\text{RMSE}(\mathcal{K}) = \sqrt{\frac{1}{\mathcal{J}} \sum_{i=1}^{\mathcal{J}} \left(G(\mathbf{y}_i) - g_{i,\mathcal{K}}(\mathbf{y}_i) \right)^2}, \quad \mathbf{y}_i \in \mathcal{T}. \tag{17}$$

Clearly, once the optimal number of centers $\mathcal{K}^*$ is chosen, the whole set of points is used for the construction of the final surrogate model. Whenever the number of RBF centers is lower than the training set size ($\mathcal{K} < \mathcal{J}$), the coefficients w_j in Eq. (14) are determined through a least-squares regression by solving

$$\mathbf{w} = \left(\mathbf{A}^{\mathsf{T}} \mathbf{A} \right)^{-1} \mathbf{A}^{\mathsf{T}} \mathbf{b}, \tag{18}$$

with $\mathbf{w} = [w_1, \ldots, w_{\mathcal{K}}]^{\mathsf{T}}$, $\mathbf{A}_{ij} = ||\mathbf{y}_i - \mathbf{c}_j||^{\tau}$, $1 \leq i \leq \mathcal{J}$, $1 \leq j \leq \mathcal{K}$ and $\mathbf{b} = [G(\mathbf{y}_1), \ldots, G(\mathbf{y}_{\mathcal{J}})]^{\mathsf{T}}$; otherwise when the optimal number of RBF centers equals the training set size, exact interpolation at the training points ($f(\mathbf{y}_i, \tau) = G(\mathbf{y}_i)$) is imposed and Eq. (18) reduces to

$$\mathbf{w} = \mathbf{A}^{-1}\mathbf{b}, \tag{19}$$

with $\mathbf{c}_j = \mathbf{y}_j$. Having less RBF centers than training points and employing the least-squares approximation in Eq. (18) to determine the coefficients w_j is particularly helpful when the training data are affected by noise. An example of least-squares regression is shown in Fig. 2.

The uncertainty $U_{\mathcal{F}}(\mathbf{y})$ associated with the SRBF surrogate model prediction is quantified by the 95%-confidence band of the cumulative density function (CDF) of $f(\mathbf{y}, \tau)$ with respect to τ for fixed $\mathbf{y}$ as follows

$$U_{\mathcal{F}}(\mathbf{y}) = \text{CDF}^{-1}(0.975; \mathbf{y}) - \text{CDF}^{-1}(0.025; \mathbf{y}), \tag{20}$$

with

$$\text{CDF}(\lambda; \mathbf{y}) \approx \frac{1}{\Theta} \sum_{i=1}^{\Theta} H[\lambda - f(\mathbf{y}, \tau_i)],$$

where $H(\cdot)$ is the Heaviside step function.

2.2.2 Multi-fidelity approach

In this section, we restrict to the case of the CFD grid generation being controlled by a scalar value α, i.e., the QoI computed with the α-th grid is denoted by $G_\alpha, \alpha = 1, \ldots, M$. The multi-fidelity approximation of G is adaptively built following the approach introduced in [49] and extended to noisy data in [50]. Extending the definition of the training set to an arbitrary number M of fidelity levels as $\{\mathcal{T}_\alpha\}_{\alpha=1}^{M}$, with each $\mathcal{T}_\alpha = \{ (\mathbf{y}_j, G_\alpha(\mathbf{y}_j)) \}_{j=1}^{\mathcal{J}_\alpha}$, the multi-fidelity approximation $\mathcal{S}_\alpha(\mathbf{y})$ of $G(\mathbf{y})$ reads

$$\mathcal{S}_\alpha(\mathbf{y}) := \mathcal{F}_1(\mathbf{y}) + \sum_{i=1}^{\alpha-1} \epsilon_i(\mathbf{y}), \tag{21}$$

where $\mathcal{F}_1$ is the single-fidelity surrogate model associated with the lowest-fidelity training set (constructed as in Eq. (15)), and $\epsilon_i(\mathbf{z})$ is the inter-level error surrogate with associated training set $\mathcal{E}_i = \{ (\mathbf{z}, \phi - \mathcal{S}_i(\mathbf{z})) \, | \, (\mathbf{z}, \phi) \in \mathcal{T}_{i+1} \}$. An example of the multi-fidelity approximation with two fidelities is shown in Fig. 3.

Algorithm 2: Adaptive multi-fidelity SRBF implementation

Multi-Fidelity SRBF for numerical quadrature

Define the initial training sets $\{\mathcal{T}_\alpha\}_{\alpha=1}^M$, with $\mathcal{T}_\alpha = \{(\mathbf{y}_j, G_\alpha(\mathbf{y}_j))\}_{j=1}^{\mathcal{J}_\alpha^0}$;

$t = 0$;

while *stopping criteria are not met* **do**   // t-th iteration of the adaptive sampling

 Perform LOOCV and find $\mathcal{K}_1^{*,t}$ as in Eq. (16) with:

 if $t = 0$ **then**

  $C = [1, \mathcal{J}_1^0]$;

 else

  $C = [\mathcal{K}_1^{*,t-1}, \mathcal{K}_1^{*,t-1} + 1]$;

 end

 Construct the SRBF surrogate $\mathcal{F}_1(\mathbf{y})$ as in Eq. (15) ;   // low-fidelity approximation

 Compute the prediction uncertainty $U_{\mathcal{F}_1}(\mathbf{y})$ as in Eq. (20) ;

 for $\alpha = 2, \ldots, M$ **do**   // evaluate surrogates of the inter-level errors

  Compute the inter-level errors $\mathcal{E}_\alpha$;

  Perform LOOCV and find $\mathcal{K}_\alpha^{*,t}$ as in Eq. (16) with:

  if $t = 0$ **then**

   $C = [1, \mathcal{J}_\alpha^0]$;

  else

   $C = [\mathcal{K}_\alpha^{*,t-1}, \mathcal{K}_\alpha^{*,t-1} + 1]$;

  end

  Construct the SRBF surrogates of the inter-level errors $\epsilon_\alpha(\mathbf{y})$ as in Eq. (15) ;

  Compute the prediction uncertainty $U_{\epsilon_\alpha}(\mathbf{y})$ as in Eq. (20) ;

  Construct the multi-fidelity SRBF surrogate $\mathcal{S}_\alpha(\mathbf{y})$ as in Eq. (21) ;

 end

 Construct the multi-fidelity approximation $\mathcal{S}_M(\mathbf{y})$ as in Eq. (22) ;   // MF approximation

 Compute the multi-fidelity prediction uncertainty $U_{\mathcal{S}_M}(\mathbf{y})$ as in Eq. (22) ;

 Set $\mathcal{T}_{\alpha 0} = \mathcal{T}_\alpha$ with $\alpha = 1, \ldots, M$;   // store current training sets

 Set $\mathcal{J}_{\alpha 0}^t = \mathcal{J}_\alpha^t$ with $\alpha = 1, \ldots, M$;   // store current training sets size

 for $j = 1, \ldots, p$ **do**   // perform parallel infill

  Find $\mathbf{y}_j^\star = \underset{\mathbf{y} \in \Gamma}{\arg\max}[U_{\mathcal{S}_M}(\mathbf{y})]$;

  Find $k_j = \mathrm{maxloc}\left[\mathbf{U}(\mathbf{y}_j^\star)\right]$ as in Eq. (24) ;

  Update the training sets $\{\mathcal{T}_\alpha\}_{\alpha=1}^{k_j} \cup \{(\mathbf{y}_j^\star, \mathcal{S}_\alpha(\mathbf{y}_j^\star))\}_{\alpha=1}^{k_j}$;   // consider exact prediction

  Update the training sets size $\{\mathcal{J}_\alpha^{t+j}\}_{\alpha=1}^{k_j} = \{\mathcal{J}_\alpha^t\}_{\alpha=1}^{k_j} + 1$;

 end

 Set $\mathcal{T}_\alpha = \mathcal{T}_{\alpha 0}$ with $\alpha = 1, \ldots, M$;   // restore previuos training sets

 Set $\mathcal{J}_\alpha^t = \mathcal{J}_{\alpha 0}^t$ with $\alpha = 1, \ldots, M$;   // restore previous training sets size

 for $j = 1, \ldots, p$ **do**   // perform new simulations

  Evaluate $\{G_\alpha(\mathbf{y}_j^\star)\}_{\alpha=1}^{k_j}$;

  Update the training sets $\{\mathcal{T}_\alpha\}_{\alpha=1}^{k_j} \cup \{(\mathbf{y}_j^\star, G_\alpha(\mathbf{y}_j^\star))\}_{\alpha=1}^{k_j}$;

  Update the training sets size $\{\mathcal{J}_\alpha^{t+j}\}_{\alpha=1}^{k_j} = \{\mathcal{J}_\alpha^t\}_{\alpha=1}^{k_j} + 1$;

 end

 $t = t + 1$;   // move to the next adaptive sampling iteration
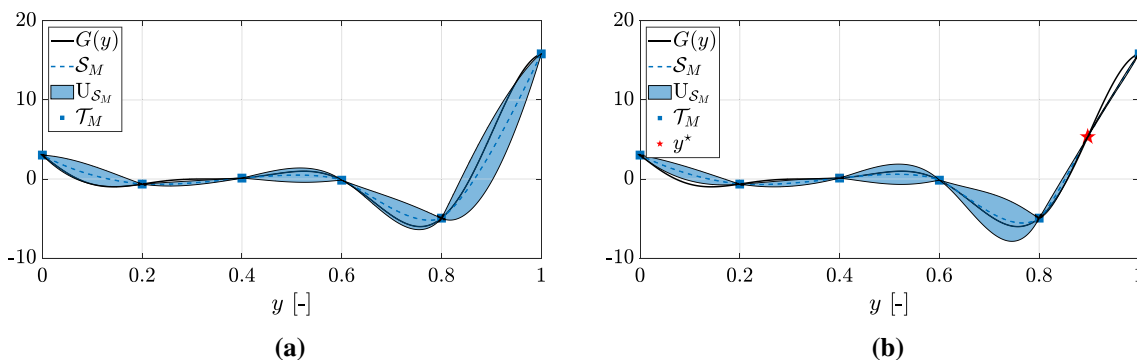
end

end

Fig. 4 Example of the adaptive sampling method using a single-fidelity training set consisting of noiseless evaluations: **a** shows the initial SRBF with the associated prediction uncertainty and training set; **b** shows the position of the new training point, the new SRBF prediction, and its uncertainty

Assuming that the uncertainty associated with the prediction of the lowest-fidelity $U_{\mathcal{F}_1}$ and inter-level errors U_{ϵ_i} as uncorrelated, the multi-fidelity approximation $\mathcal{S}_M(\mathbf{y})$ of $G(\mathbf{y})$ and its uncertainty $U_{\mathcal{S}_M}$ read

$$G(\mathbf{y}) \approx \mathcal{S}_M(\mathbf{y}) = \mathcal{F}_1(\mathbf{y}) + \sum_{i=1}^{M-1} \epsilon_i(\mathbf{y})$$

$$\text{and} \quad U_{\mathcal{S}_M}(\mathbf{y}) = \sqrt{U_{\mathcal{F}_1}^2(\mathbf{y}) + \sum_{i=1}^{M-1} U_{\epsilon_i}^2(\mathbf{y})}. \tag{22}$$

2.2.3 Adaptive sampling approach

Upon having evaluated $U_{\mathcal{S}_M}$, the multi-fidelity surrogate is then updated adding a new training point following a two-steps procedure: first, the coordinates of the new training point $\mathbf{y}^\star$ are identified based on the SRBF maximum uncertainty, see [23], solving the maximization problem:

$$\mathbf{y}^\star = \underset{\mathbf{y} \in \Gamma}{\operatorname{argmax}}[U_{\mathcal{S}_M}(\mathbf{y})]. \tag{23}$$

An example (with one fidelity only) is shown in Fig. 4. Secondly, once $\mathbf{y}^\star$ is identified, the training set/sets to be updated with the new training point $(\mathbf{y}^\star, G_\alpha(\mathbf{y}^\star))$ are $\mathcal{T}_\alpha$ with $\alpha = 1, \ldots, k$, where k is defined as

$$k = \operatorname{maxloc}\left[\mathbf{U}(\mathbf{y}^\star)\right] \quad \text{and}$$
$$\mathbf{U}(\mathbf{y}^\star) \equiv \{U_{\mathcal{F}_1}(\mathbf{y}^\star)/\gamma_1, U_{\epsilon_1}(\mathbf{y}^\star)/\gamma_2, ..., U_{\epsilon_{M-1}}(\mathbf{y}^\star)/\gamma_M\}, \tag{24}$$

with γ_α being the computational cost associated with the α-th level.

In the present work, the adaptive sampling procedure starts with five training points (for each fidelity level) located at the domain center and at the centers of each boundary of Γ. Furthermore, to avoid abrupt changes in the SRBF prediction from one iteration to the next one, the search for the optimal number of centers for the α-th fidelity $\mathcal{K}_\alpha^*$ can be constrained. Herein, at every adaptive sampling iteration, the problem in Eq. (16) is solved assuming $\mathcal{K}$ to be either equal to the number of centers at the previous iteration or incremented by 1, i.e., $\mathcal{C} = [\mathcal{K}_\alpha^{*,t-1}, \mathcal{K}_\alpha^{*,t-1} + 1]$, except for the first iteration where no constraint is imposed.

A deterministic version of the particle swarm optimization algorithm [51] is used for the solution of the optimization problem in Eq. (23).

The adaptive sampling is, therefore, inherently sequential (the uncertainty changes every time a new point is added), but this is sub-optimal whenever the numerical simulations can be performed with an hardware capable of running p simulations simultaneously. In this case, it would be ideal to identify p training points where the models G_α can be run in parallel, instead of running them one after the other. To this end, we follow a parallel-infill procedure, i.e. we perform p "guessing steps": the adaptive sampling procedure is repeated p times replacing the evaluations of the actual model G_α with the evaluations of the multi-fidelity models $\mathcal{S}_\alpha$. This replacement significantly speeds up the p steps, since the true models G_α are not evaluated at this stage. Upon doing these p guessing steps, the actual G_α are evaluated all at once (i.e., in parallel) at the p training points just obtained and these evaluations replace the corresponding multi-fidelity evaluations in the training set.

Finally, numerical quadrature is used on the multi-fidelity SRBF surrogate model to estimate the statistical moments of the QoI.

3 Numerical tests

In this section, two numerical tests are considered. First, the performance of the MISC and SRBF methods are compared on an analytical test, and then the main problem of this work, i.e., the RoPax ferry problem mentioned in the introduction, is discussed. In the analytical example, Taylor expansions of increasing order are considered as different fidelities to be employed, while in the RoPax problem the fidelities are obtained by using different grid refinements, as will be clearer later. Both problems consider uniformly distributed uncertain parameters. Before entering the detailed discussion of the two tests, an overview of the error metrics used to carry out the comparison is given in the following.

3.1 Error metrics

The performance of MISC and SRBF are assessed by comparing the convergence of both methods to a reference solution according to several error metrics. The specific choice of the reference solution (denoted below by G_{ref}) for each test will be detailed in the corresponding sections. In the following list, we use the symbol $\mathcal{S}$ for both the MISC and SRBF surrogate models for sake of compactness, i.e., $\mathcal{S} = \mathcal{S}_\Lambda$ for MISC (cf. Eq. (7)) and $\mathcal{S} = \mathcal{S}_M$ for SRBF (cf. Eq. (22)).

1. Relative error between the first four centered moments (mean, variance, skewness, and kurtosis) of the MISC/SRBF approximations and those of the reference solution:

$$err_i = \frac{|\text{Mom}_i[\mathcal{S}] - \text{Mom}_i[G_{\text{ref}}]|}{|\text{Mom}_i[G_{\text{ref}}]|}, \quad i = 1, \ldots, 4, \quad (25)$$

where $\text{Mom}_i[\mathcal{S}]$ denotes the MISC/SRBF approximation of the i-th centered moment (computed by the quadrature rule associated with MISC/SRBF) and $\text{Mom}_i[G_{\text{ref}}]$ the approximation of the i-th centered moment of the reference solution computed by a suitable quadrature rule (more details will be given later).

2. Relative error in discrete L_2 and L_∞ norm between the MISC/SRBF surrogates and the reference solution, i.e., sample mean square error and largest point-wise prediction error, respectively; the differences are evaluated at a set of $n = 10000$ random points $\mathbf{y}_i \in \Gamma$. In formulas,

$$err_{L_2} = \frac{\sqrt{\frac{1}{n} \sum_{i=1}^{n} \left[\mathcal{S}(\mathbf{y}_i) - G_{\text{ref}}(\mathbf{y}_i) \right]^2}}{\sqrt{\frac{1}{n} \sum_{i=1}^{n} G_{\text{ref}}(\mathbf{y}_i)^2}},$$

$$err_{L_\infty} = \frac{\max_{i=1,\ldots,n} \left(|\mathcal{S}(\mathbf{y}_i) - G_{\text{ref}}(\mathbf{y}_i)| \right)}{\max_{i=1,\ldots,n} G_{\text{ref}}(\mathbf{y}_i)}. \quad (26)$$

3. A visual comparison of the PDFs obtained by Matlab's *ksdensity* function, using again as input the 10000 points used before.

4. Convergence of the CDF approximated by MISC/SRBF to the CDF of the reference solution, as measured by the Kolmogorov–Smirnov (KS) test statistic. In details, we evaluate an approximation of the quantity

$$T = \sup_{t \in \text{range}[\mathcal{S}, G_{\text{ref}}]} |\text{CDF}^{\mathcal{S}}(t) - \text{CDF}^{G_{\text{ref}}}(t)|, \quad (27)$$

where $\text{range}[\mathcal{S}, G_{\text{ref}}]$ is the largest common range of values taken by $\mathcal{S}$ and G_{ref}, $\text{CDF}^{\mathcal{S}}$ and $\text{CDF}^{G_{\text{ref}}}$ are the empirical CDFs obtained by the set used before of 10000 random samples of the MISC/SRBF surrogate models and reference model, respectively. We then check that T converges to zero as the surrogate models get more and more accurate. The values of T reported in the next sections are obtained with the Matlab's *kstest2* function.

We emphasize that the adaptive criteria that drive the construction of the MISC and SRBF approximations need not match the error metrics above (compare them against Eqs. (13), (10) and (11) for MISC, and Eq. (23) for SRBF). It is actually interesting to investigate how MISC and SRBF converge when monitoring error norms that are not aligned with the adaptivity criteria.

3.2 Analytical problem

3.2.1 Formulation

As analytical test, a two-dimensional function is chosen. This function is designed to be representative of the RoPax problem: the input parameters $\mathbf{y}$ are independent and have a uniform distribution, and the function is non-linear, non-polynomial and monotonic. In details, it is defined as

$$G(\mathbf{y}) = G(y_1, y_2) = \sin\left(\frac{\exp(y_1 + y_2)}{5} \right),$$

with $\mathbf{y} \in [0, 1]^2$. To provide a range of fidelities $G_\alpha(\mathbf{y})$ for $G(\mathbf{y})$, Taylor expansions of order α of the argument of the $\sin(\cdot)$ function, that is $\frac{\exp(y_1 + y_2)}{5}$, are performed for $\alpha = 1, \ldots, 6$ in the neighborhood of $\mathbf{y} = (0, 0)$. The sixth-order Taylor expansion $G_6(\mathbf{y})$ is then considered as the highest-fidelity and the first order $G_1(\mathbf{y})$ as the lowest-fidelity. Figure 5 shows the true function $G(\mathbf{y})$ and its approximations $G_6(\mathbf{y})$ and $G_1(\mathbf{y})$. We mention in-passing that the sixth-order Taylor expansion is almost indistinguishable from the true function in the range of y_1, y_2 considered, whereas the low-fidelity function is significantly different and does not show any change in curvature. Note that the difference between the
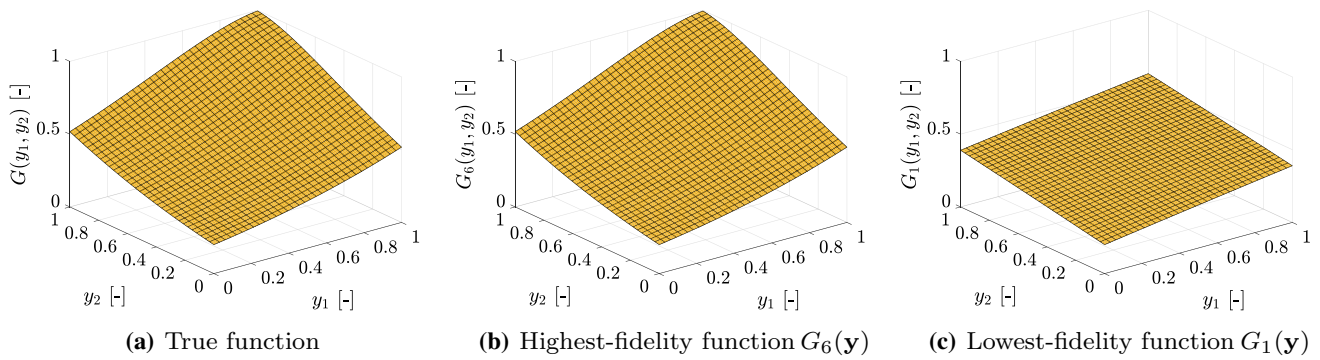
(a) True function **(b)** Highest-fidelity function $G_6(\mathbf{y})$ **(c)** Lowest-fidelity function $G_1(\mathbf{y})$

Fig. 5 Analytical problem, true function and highest- and lowest-fidelity approximations

(a) Results for the MISC method with profit $P^{\mathcal{R}}$.

(b) Results for the MISC method with profit $P^{\mathcal{S}}$. For ease of comparison, we overlap the results for MISC-$P^{\mathcal{R}}$ (gray dashed lines) in the plot on the right.
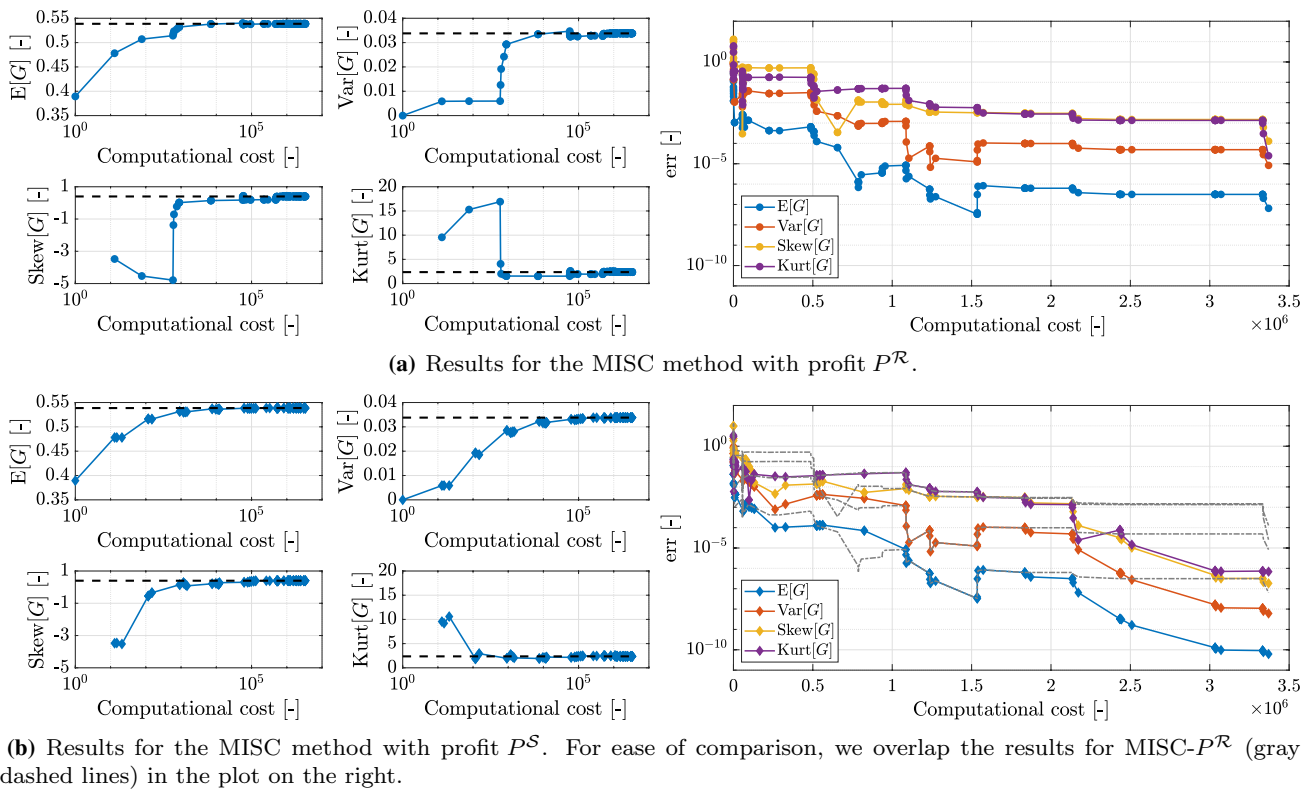
Fig. 6 Analytical problem, results for the MISC method. Left: convergence of the values of the first four centered moments. The black dashed line marks the reference value of the moments. Right: relative error of the moments (see Eq. (25))

sixth-order and the true function is actually irrelevant for our purposes since we never consider the true function in the numerical tests: errors are indeed computed with respect the sixth-order approximation, in analogy with PDE-based problems where no exact closed-formula solution is available, and the "ground-truth" is usually taken as a "refined-enough" solution.

A normalized computational cost is associated with each evaluation of the α-th Taylor expansion $G_\alpha(\mathbf{y})$ as

$$\text{cost}(\alpha) = 8^{\alpha-1}. \tag{28}$$

This choice is done to keep the analogy with the RoPax problem and will become clear later.

3.2.2 Numerical results

We start the discussion with the comparison of the MISC/SRBF estimates of the moments with the reference values. The reference values are computed by an accurate sparse-grids quadrature rule with $2^{15} + 1$ points where the reference surrogate model/function $G_{\text{ref}}(\mathbf{y})$ is the highest-fidelity
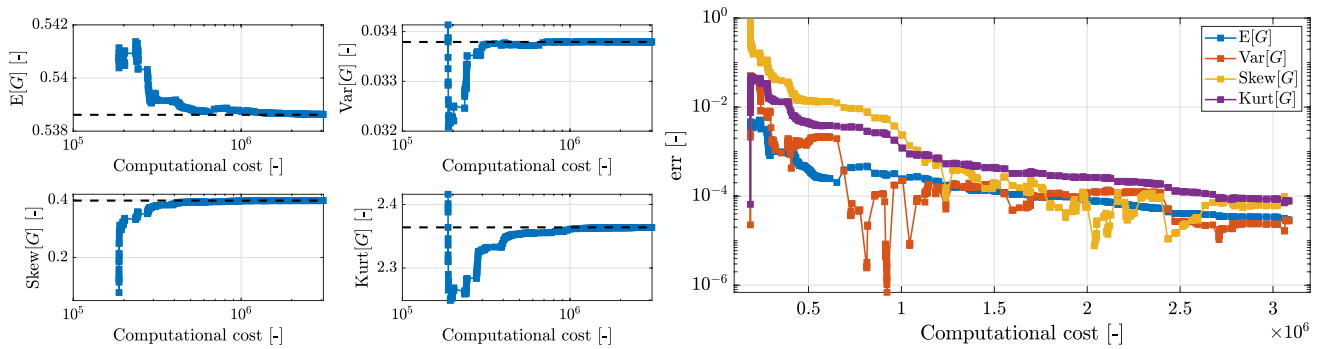
Fig. 7 Analytical problem, results for the SRBF method. Left: convergence of the values of the first four centered moments. The black dashed line marks the reference value of the moments. Right: relative error of the moments (see Eq. (25))
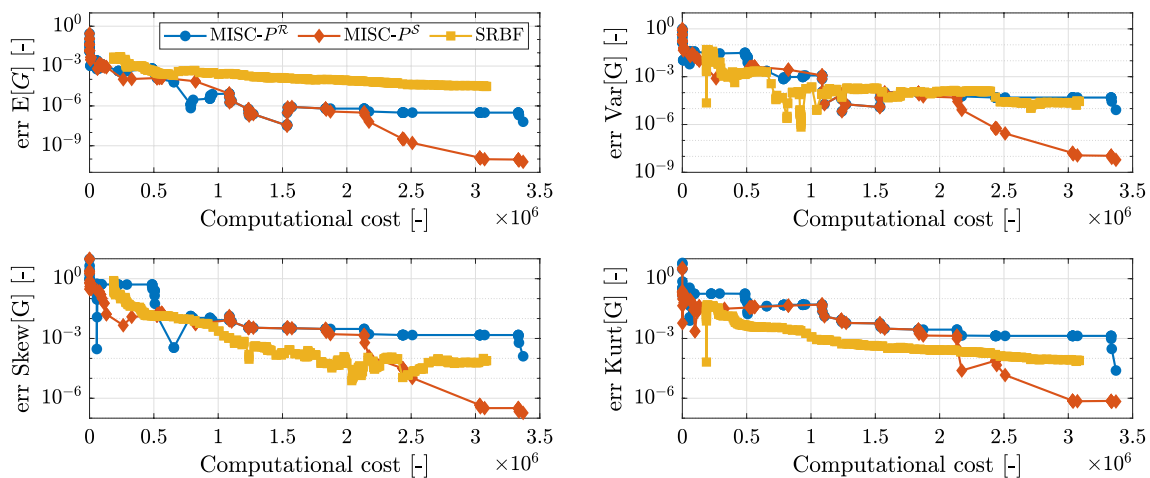
Fig. 8 Analytical problem, comparison of MISC and SRBF methods: relative errors of the moments. It is a compact visualization of the results of Figs. 6 and 7, where the results of MISC are reported only until a computational cost comparable with the one reached by the SRBF method

approximation $G_6(\mathbf{y})$. The calculations have been done using the Sparse Grids Matlab Kit[2] [52].

In Fig. 6 the convergence of the MISC estimates of the first four centered moments and their relative errors as defined in Eq. (25) are reported. The two variants of MISC (denoted by MISC-$P^{\mathcal{R}}$ and MISC-$P^{\mathcal{S}}$ in the following) introduced in Sect. 2.1 are tested, i.e., two type of profits (see Eq. (13)), $P^{\mathcal{R}}$ based on the quadrature error (see Eq. (10)) and $P^{\mathcal{S}}$ based on the point-wise accuracy of the surrogate model (see Eq. (11)), are considered. The results for the case of a quadrature-based profit $P^{\mathcal{R}}$ are displayed in the first set of plots (see Fig. 6a left). All the moments converge to the reference results marked with the black dashed line. In Fig. 6a right one can observe that the error is larger the higher the order of the moment. Remarkably, even if the adaptivity of the MISC method is driven by the improvement in the first-order moment, all the moments are estimated very well. The second set of plots

(see Fig. 6b) suggests that also the version of MISC driven by the accuracy of the surrogate model $P^{\mathcal{S}}$ is effective in the estimation of the moments. By comparing the two methods, one can observe that the latter one brings better results, as the error for all the moments is always smaller.

Figure 7 shows the convergence of the moments for SRBF: differently from MISC, all the moments have a quite similar convergence towards their reference values, with errors all converging within the same order of magnitude. The convergence is almost monotonic for the expected value and for the kurtosis whereas some oscillations can be observed for the variance and the skewness. Note that for this problem the SRBF method is based on interpolation at the training points (i.e., the weights are computed by solving Eq. (19)), whereas in the following RoPax example the regressive approach (i.e., solving Eq. (18)) is used, for reasons that will be clear later on.

Figure 8 shows that MISC and SRBF achieve comparable results in terms of the relative errors of the centered
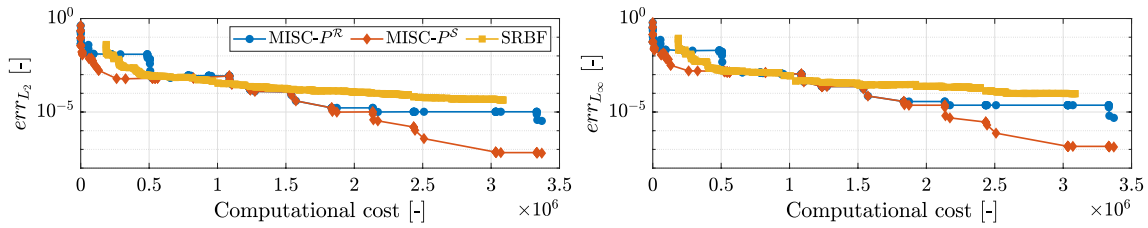
Fig. 9 Analytical problem, comparison of MISC and SRBF methods: relative error of the approximation of G in L_2 (left) and L_∞ norm (right) (see Eq. (26))
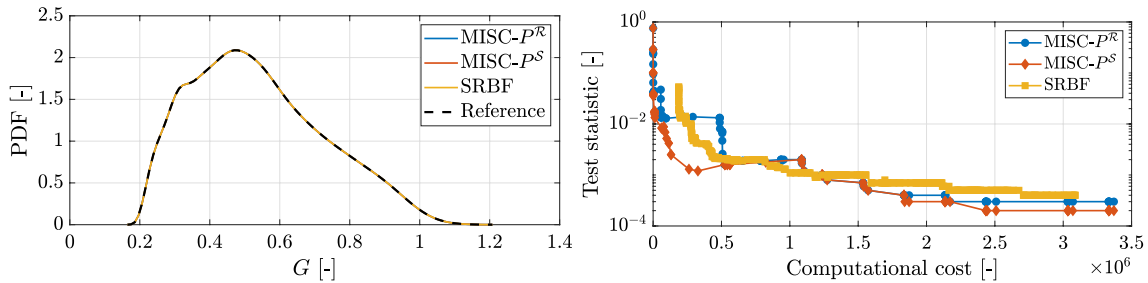
(a) PDFs at the final iteration (obtained by Matlab's *ksdensity* enforcing positive support since G takes positive values only)

(b) KS test statistic (see Eq. (27))

**Fig. 10** Analytical problem, comparison of MISC and SRBF methods

Fig. 11 Analytical problem, results for the MISC and SRBF method. Sampling points at the last iteration

(a) MISC with profit $P^{\mathcal{S}}$

(b) SRBF
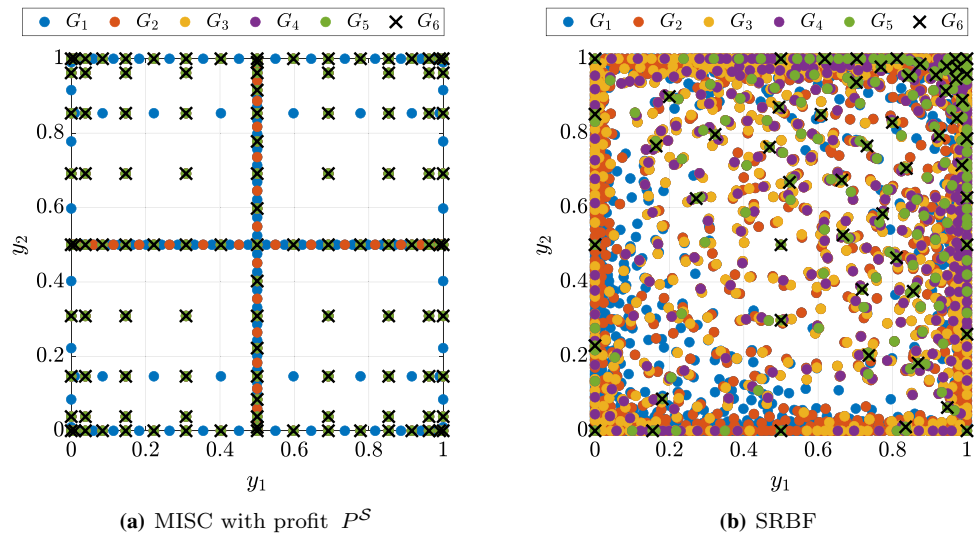
moments. Specifically, MISC-$P^{\mathcal{S}}$ performs better in the evaluation of all the moments, whereas SRBF performs slightly better than MISC-$P^{\mathcal{R}}$ in the evaluation of the skewness. It is also worth noting that SRBF starts with a higher computational cost in comparison with MISC, due to the fact that the initialization strategy requires to sample all the available fidelities.

The results for the L_2 and L_∞ norms of the MISC error (see Eq. (26)) of Fig. 9 are in agreement with the previous findings for the estimation of the moments. An improvement of about two orders of magnitude is observed in favor of the surrogate-based method MISC-$P^{\mathcal{S}}$ with respect to MISC-$P^{\mathcal{R}}$ in the final part of the convergence curve. When comparing these results with the convergence of the SRBF one can observe that both versions of MISC achieve better results than SRBF.

Fig. 12 RoPax ferry, hull form. Free surface level is reported for the nominal draught
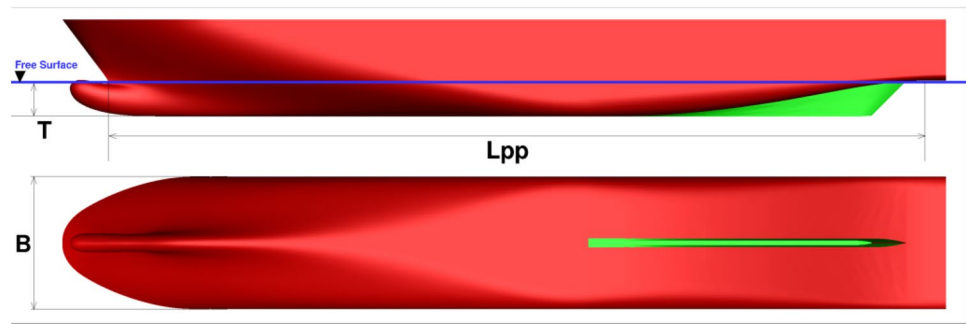
The comparison of the PDFs given in Fig. 10a shows a very good agreement of the MISC and SRBF results with the reference ones. In Fig. 10b the results of the KS test statistics (cf. Eq. (27)) are reported: both versions of MISC show a slightly better convergence of the test statistic.

Finally, it is worth looking at the sampling performed by the two methods. In Fig. 11 the samples selected by MISC-P^S and SRBF are displayed. In the first case (see Fig. 11a) the samples are well distributed over the domain in a symmetric way, with no preferential directions. The SRBF sampling is instead slightly more clustered in the regions where the high-fidelity function shows a larger curvature (cf. Figs. 5b and 11b). The sampling performed by MISC-P^R is not shown for brevity, as it is very similar to the one for MISC-P^S.

3.3 RoPax resistance problem

3.3.1 Formulation and CFD method

The main problem addressed in this work is the forward UQ analysis of the model-scale resistance (R_T) of a RoPax ferry in straight ahead advancement, subject to two operational uncertainties $\mathbf{y} = [U, T]$, namely the advancement speed (U) and the draught (T), uniformly distributed within the ranges summarized in Table 1. The choice of using two operational parameters allows visual investigation of the results while preserving all the main difficulties that arise when solving parametric approximation problems. Furthermore, the advancement speed and draught are two operational parameters with significant practical implications: for instance Froude and Reynolds numbers vary with the advancement speed; allowed payload and block coefficient vary with the draught, etc.

The RoPax ferry is characterized by a length between perpendicular at nominal draught (L_{PP}) of 162.85 m and a block coefficient $C_B = 0.5677$ (see Fig. 12). The parametric geometry of the RoPax is produced with the computer-aided design

environment integrated in the CAESES software, developed by FRIENDSHIP SYSTEMS AG, and made available in the framework of the H2020 EU Project Holiship[3]. The analysis is performed at model scale with a scale factor equal to 27.14. The main dimensions and the operative conditions are summarized in Table 1. The advancement speed ranges from 12 to 26 knots at full scale and the draught variation is $\pm 10\%$ of the nominal draught, which corresponds to a variation of about $\pm 15\%$ of the nominal displacement. The corresponding range in Froude number $Fr = U/\sqrt{gL_{PP}}$ is [0.154, 0.335], whereas the variation in Reynolds number (at model scale) is $Re = \rho U L_{PP}/\mu = U L_{PP}/\nu \in [6.423 \cdot 10^6, 1.392 \cdot 10^7]$, where $\rho = 998.2$ kg/m^3 is the water density, $\nu = \mu/\rho = 1.105 \cdot 10^{-6}$ m^2/s the kinematic viscosity and $g = 9.81$ m/s^2 the gravitational acceleration.

The hydrodynamics performance of the RoPax ferry is assessed by the RANS code χnavis developed at CNR-INM. The main features of the solver are summarized here; for more details, the interested reader is referred to [30–32, 53] and references therein. χnavis is based on a finite volume discretization of the RANS equations, with variables collocated at the cell centers. Turbulent stresses are related to the mean velocity gradients by the Boussinesq hypothesis; the turbulent viscosity is estimated by the Spalart–Allmaras turbulence model [54]. Wall functions are not adopted, therefore the wall distance $y^+ = 1$ is ensured on the no-slip wall. Free-surface effects are taken into account by a reliable single-phase level-set approach.

The computational domain extends to 2 L_{PP} in front of the hull, 3 L_{PP} behind, and 1.5 L_{PP} sideway; a depth of 2 L_{PP} is imposed (see Fig. 13a). On the solid walls (in red in the figure), the velocity is set equal to zero, whereas zero normal gradient is enforced on the pressure field; at the (fictitious) inflow boundary (in blue in Fig. 13a), the velocity is set to the undisturbed flow value and the pressure is extrapolated from inside; the dynamic pressure is set to zero at the outflow (in yellow), 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 (in purple). Taking advantage of the symmetry of the flow relative to the $y = 0$ plane, computations are performed

Table 1 Main geometrical details and operative conditions of the RoPax ferry (model scale 1 : 27.14)

Description	Symbol	Full scale	Model scale	Unit
Length between perpendiculars	L_{PP}	162.85	6.0	m
Beam	B	29.84	1.0993	m
Block coefficient	C_B	0.5677	0.5677	–
Nominal displacement	∇	19584.04	0.9996	m^3
Nominal draught	T_n	7.10	0.261660	m
Draught range	T	[6.391, 7.812]	[0.2355, 0.2878]	m
Speed range	U	[6.173, 13.376]	[1.185, 2.567]	m/s
Froude range	Fr	[0.154, 0.335]	[0.154, 0.335]	–
Reynolds range	Re	$[9.081 \cdot 10^8, 1.968 \cdot 10^9]$	$[6.423 \cdot 10^6, 1.392 \cdot 10^7]$	–

(a) Boundary conditions

(b) Computational grid
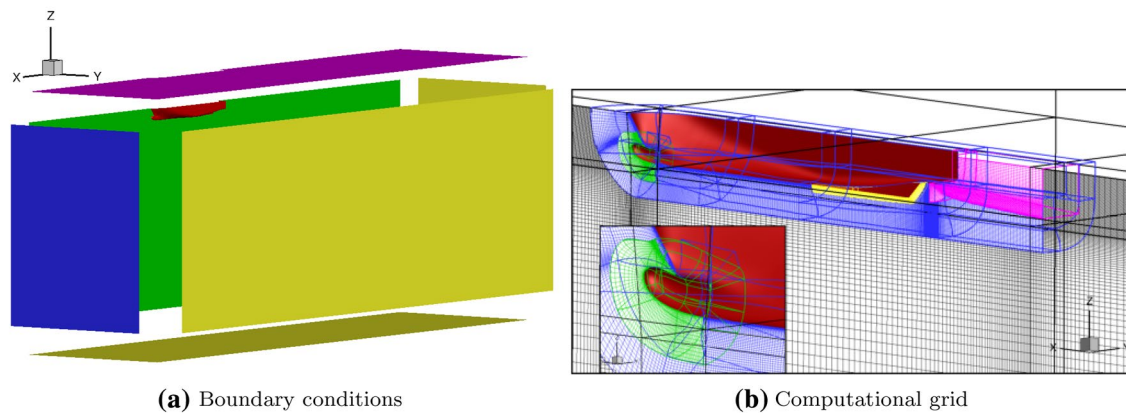
Fig. 13 RoPax ferry, numerical setup. Boundary conditions and computational grid
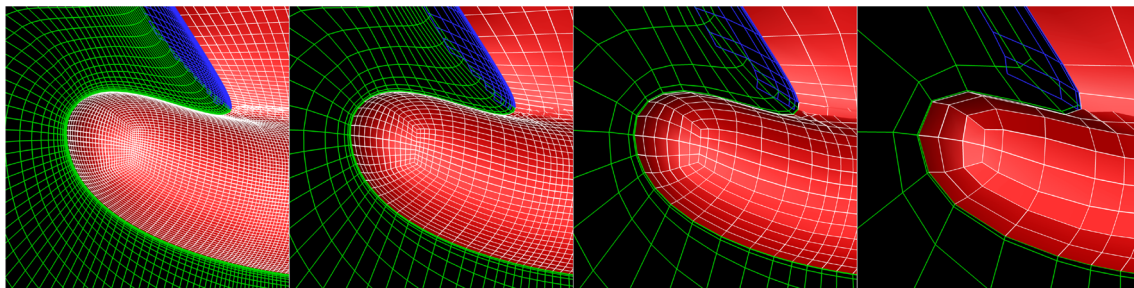
Fig. 14 RoPax grids, detail of the bow region, left to right: $\mathcal{M}_4$, $\mathcal{M}_3$, $\mathcal{M}_2$, $\mathcal{M}_1$

for half ship only, and the usual symmetry boundary conditions are enforced on the symmetry longitudinal plane (in green).

The computational grid is composed by 60 adjacent and partially overlapped blocks; Fig. 13b shows a particular of the block structures in the region around the ship hull and the computational grid on the symmetry plane. Taking the advantage of a Chimera overlapping approach, the grids around the skeg and around the bow are generated separately from the grid around the hull; a background Cartesian grid

is then built and the whole grid is assembled by means of an in-house overlapping grid pre-processor. The final grid counts for a total of about 5.5M control volumes for half the domain. The numerical solutions are computed by means of a full multi-grid–full approximation scheme (FMG–FAS), with four grid levels (from coarser to finer: $\mathcal{M}_1$, $\mathcal{M}_2$, $\mathcal{M}_3$, and $\mathcal{M}_4$), each obtained from the next finer grid with a coarsening ratio equal to 2, along each curvilinear direction. In the FMG–FAS approximation procedure, the solution is first computed on the coarsest grid level and then

Fig. 15 RoPax ferry, $\mathcal{M}_4$ CFD results in terms of non-dimensional wave pattern (left) and surface pressure (right) for: Fr = 0.193, $T = 3.9249 \cdot 10^{-2} L_{\text{PP}}$ and $T = 4.7971 \cdot 10^{-2} L_{\text{PP}}$, top row left and right; Fr = 0.335, $T = 3.9249 \cdot 10^{-2} L_{\text{PP}}$ and $T = 4.7971 \cdot 10^{-2} L_{\text{PP}}$, bottom row left and right
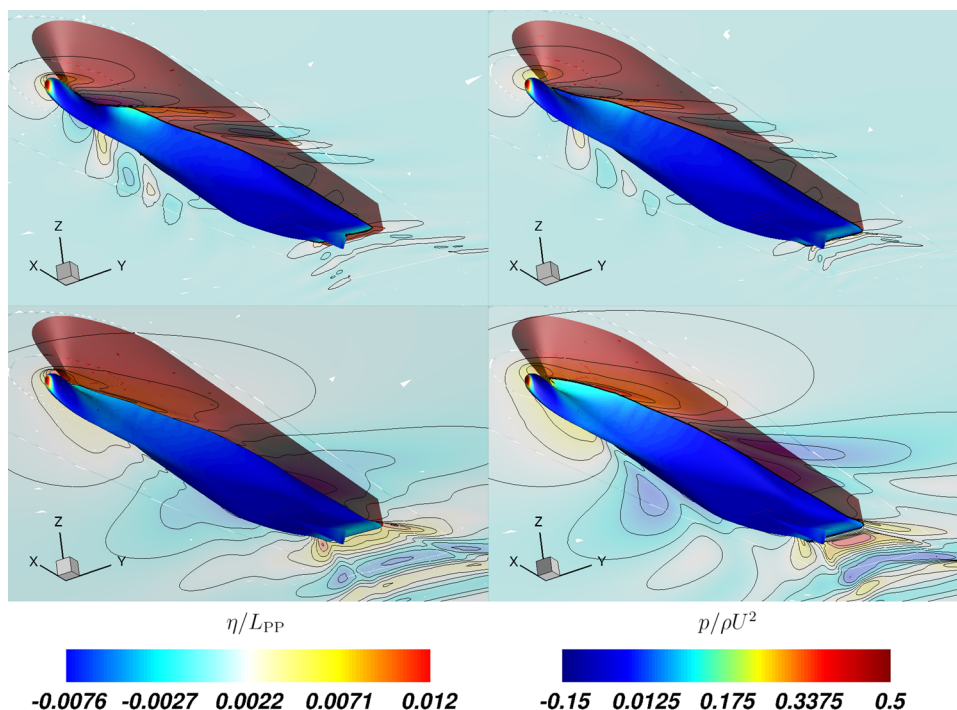
approximated on the next finer grid by exploiting all the coarser grid levels available with a V-Cycle. The process is repeated up to the finest grid level. For the present UQ problem the number of grid volumes is 5.5M for $\mathcal{M}_4$, 699K for $\mathcal{M}_3$, 87K for $\mathcal{M}_2$, and 11K for $\mathcal{M}_1$. To provide an idea about the different grid resolutions between the grid levels, Fig. 14 shows a particular of the grid at the bow region for $\mathcal{M}_4$, $\mathcal{M}_3$, $\mathcal{M}_2$ and $\mathcal{M}_1$ grids.

Since the grids are obtained by a dyadic derefinement, the following normalized computational costs can be assigned to each grid:

$$\text{cost}(\alpha) = 8^{\alpha-1}, \tag{29}$$

with $\alpha = 1, \ldots, 4$. In the FMG-FAS scheme the computation on the α-th grid level involves computations on all the coarser grids . However, with the estimation in Eq. (29), only the cost of the highest-fidelity level samples is taken into account, i.e., the computations on the coarser grids are considered negligible.

Fig. 15 shows an overview of the numerical solutions obtained for different conditions in terms of wave pattern and pressure on the hull surface; wave height (as elevation with respect to the unperturbed level) and surface pressure are reported in non-dimensional values, making the height non-dimensional with L_{PP} and the pressure with ρU^2. A clear, and obvious, Froude number dependency is seen for the wave patterns; at the lower speed shown, the free surface is weakly perturbed (note that the same color range has been

used for all the panels), whereas, at higher Froude, a clear Kelvin pattern is seen. Also, at higher speed, the formation of a well-defined transom wave system is observed, including the presence of the classical rooster tail. It is also worth to observe the influence of the draught on the wave system; in particular at the lowest speed and smallest draught reported, the rear part of the bulbous is partially dry (better seen in the enlarged views presented in Fig. 16). The region of very low pressure caused by the flow acceleration around the bow is obviously the cause. For all cases, the high pressure in the stagnation point at the bow prevents the bow to be outside the water, as it is at the rest conditions at least for the nominal and the smaller draughts (see Fig. 12). At the higher speed, the larger draught condition causes a stronger rooster tail system at the stern, with higher crest and trough.

Figure 17 shows the complete FMG–FAS cycle for the minimum and the maximum Froude numbers. The final evaluation of the R_T for each grid is performed averaging the R_T value among the last 100 iterations of the cycle. These are highlighted by the gray areas. Even if a general second-order convergence has been verified (not shown here for the sake of conciseness), it is evident that, although the FMG–FAS switches to a finer grid when the solver residual are lower than the defined convergence threshold (not shown here), the value of R_T is clearly not converged yet, at least for the coarsest grid level. This has been observed mostly for low Froude numbers. Therefore, the final value of R_T can significantly deviate for simulations performed on the same
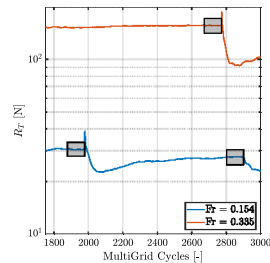
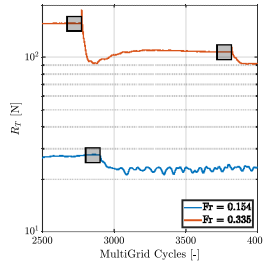**Fig. 16** RoPax ferry, enlarged view of the bow region as in Fig. 15

η/L_{PP}
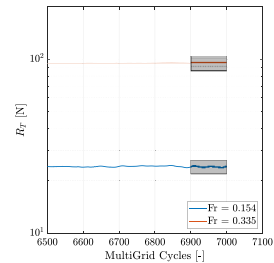
$p/\rho U^2$

-0.0076 -0.0027 0.0022 0.0071 0.012

-0.15 0.0125 0.175 0.3375 0.5
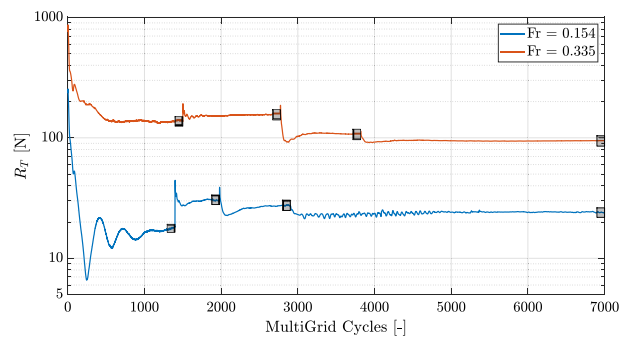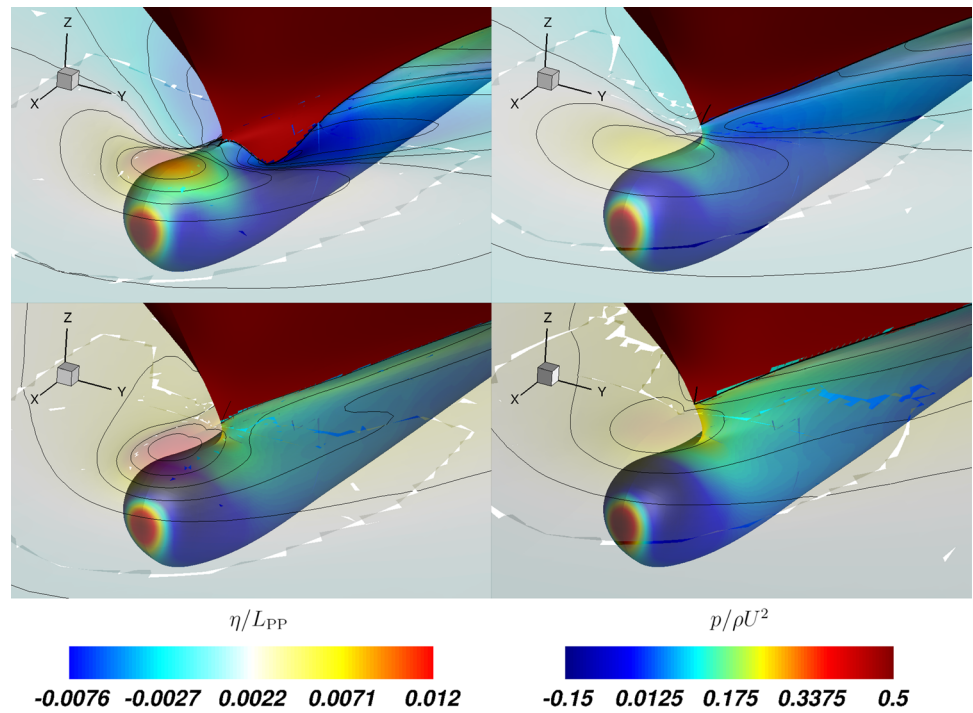
(a) Complete FMG–FAS cycle
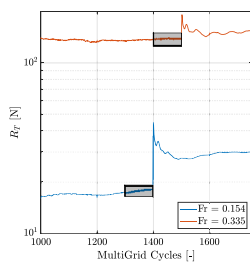
(b) Detail of the shift from $\mathcal{M}_1$ to $\mathcal{M}_2$

(c) Detail of the shift from $\mathcal{M}_2$ to $\mathcal{M}_3$

(d) Detail of the shift from $\mathcal{M}_3$ to $\mathcal{M}_4$

(e) Detail of the last iterations on $\mathcal{M}_4$

Fig. 17 RoPax problem, FMG–FAS cycle and highlight (gray areas) of the iterations considered to evaluate the R_T for each grid for the minimum and maximum Froude numbers

Fig. 18 RoPax problem, surrogate models

(a) Lowest-fidelity model

(b) Highest-fidelity model

(c) MISC model

(d) SRBF model
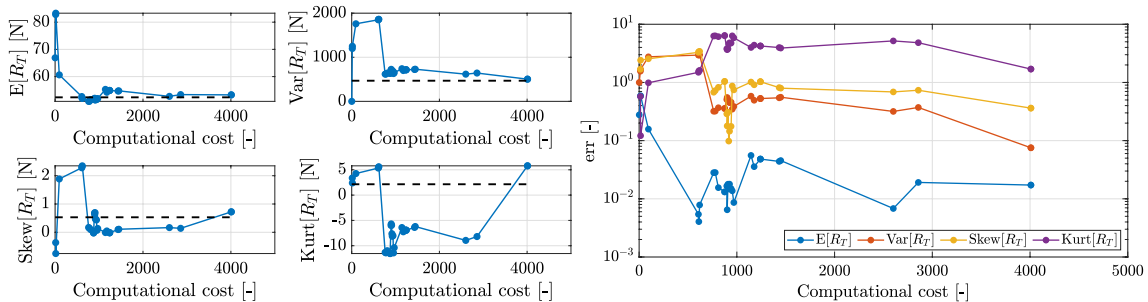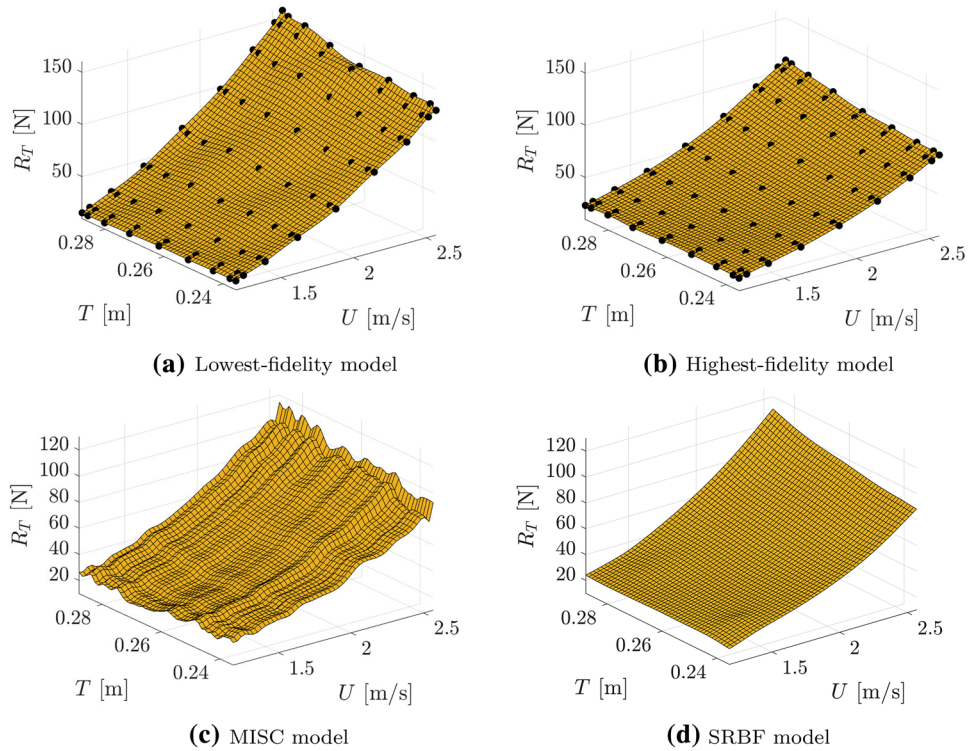
Fig. 19 RoPax problem, results for the MISC method. Left: convergence of the values of the first four centered moments. The black dashed line marks the reference value of the moments. Right: relative error of the moments (see Eq. (25))

grid but with slightly different conditions, thus producing evaluations affected by numerical noise. This will have an impact on the following UQ analysis.

3.3.2 Numerical results

Hereafter, a detailed comparison of the performance of the MISC and SRBF methods is provided. The reference surrogate model $G_{\text{ref}}(\mathbf{y})$ is obtained considering highest-fidelity simulations only. In details, an isotropic tensor grid consisting of 9×9 CC points (see Eq. (1)) is constructed over Γ and the corresponding simulations on the grid $\mathcal{M}_4$ are performed. The resulting surrogate model is an interpolatory model, based on global tensor Lagrange polynomials, which

is shown in Fig. 18b. Figure 18a shows instead the surrogate obtained with simulations on $\mathcal{M}_1$, at the same CC points. Notice that both surrogates are affected by the noise, and more specifically, the noise is more evident in the lowest-fidelity surface which is significantly less smooth than the highest-fidelity surrogate. Reference values for the centered moments are then computed applying the tensor quadrature formula associated with the CC points to the highest-fidelity simulations.

First, the performance of the MISC method is discussed. Only the results for the version of MISC with quadrature-based profits $P^{\mathcal{R}}$ are reported here, since this approach outperforms the version of MISC with surrogate-based profits $P^{\mathcal{S}}$, for reasons related with the presence of the numerical

Fig. 20 RoPax problem, surrogate models obtained with the MISC method

(a) Surrogate model for R_T

(b) Surrogate model for R_T^2

(c) Surrogate model for R_T^3

(d) Surrogate model for R_T^4
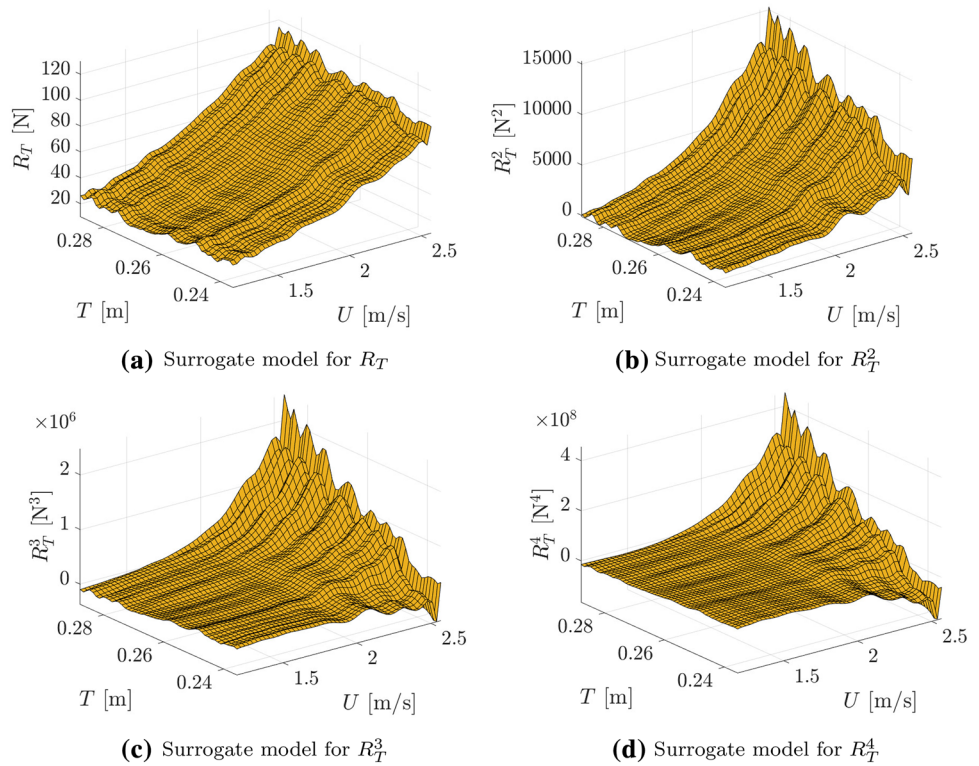
Fig. 21 RoPax problem, number of simulations per grid selected by the MISC method at each iteration

(a) Total number of simulations

(b) Number of new simulations
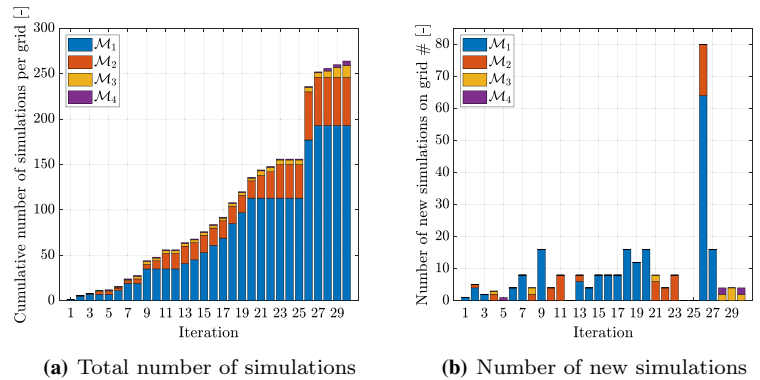
noise that will be clarified in a moment. In Fig. 19 the values of the approximations of the first four centered moments of R_T obtained with MISC at different computational costs are displayed on the left, while the relative errors are shown on the right. Upon inspection of these plots, we can conclude that the quality of the estimates decreases with increasing order of the moments. In particular, the expected value and the variance seem to converge reasonably well (although the estimate of the expected value seems to hit a temporary plateau, after having obtained a good estimate at a low computational cost), whereas the kurtosis is strongly underestimated and its approximation results to be very poor.

To explain this behavior, we have a closer look at the MISC quadrature formula (9). In particular, let us recall that

the computation of the first four centered moments implicitly uses surrogate models for r th powers of the quantity of interest R_T^r, $r = 1, \dots, 4$, (see Sect. 2.1). These models are displayed in Fig. 20. The first one, corresponding to $r = 1$, is quite rough: the surface shows an oscillatory behavior and the expected monotonicity of R_T with respect to U and T is destroyed. This is due to the already discussed presence of numerical noise, which particularly affects the low-fidelity simulations. Indeed, MISC intensively samples low-fidelities by construction, see Fig. 21a, where we report the evaluations allocated on each fidelity as the iterations proceed. In particular, most of the low-fidelity simulations are added from iteration 13 on (see Fig. 21b): this explains that the estimate of $\mathbb{E}[R_T]$ reaches reasonable values at early iterations,
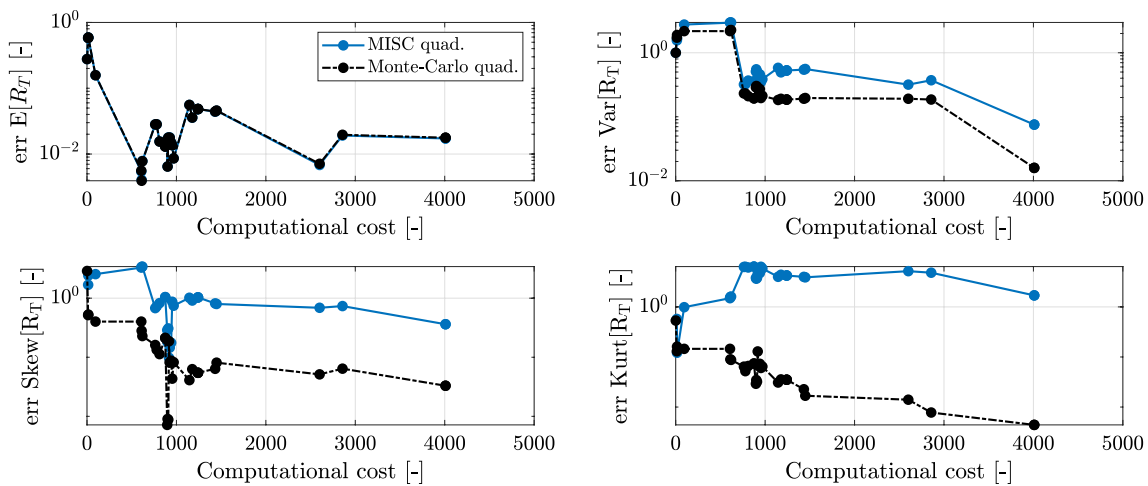
Fig. 22 RoPax problem, results for the MISC method: relative error of the moments (see Eq. (25)). A comparison between the results obtained with the MISC quadrature formula (9) and the Monte Carlo quadrature

i.e., when there is still a balance of low- and higher- fidelity simulations, and its convergence deteriorates later, i.e., when low-fidelity simulations are the majority. Given that the numerical noise introduces spurious oscillations in the surrogate model already for $r = 1$, such oscillations can only be amplified for $r > 1$, as can be observed in Fig. 20b,c,d. Hence, the computation of moments suffers more from the noise the higher the order.

This observation then suggests that a way to mitigate the impact of such oscillations in the computation of statistical moments is to employ a method that does not require higher-order surrogate models. In this work, we propose to compute such moments by taking Monte Carlo samples of the surrogate model of R_T, and approximate the moments from these values, with the usual sample formulas. The results reported in Fig. 22 have been obtained taking the average of 10 repetitions with 10000 samples (again, note that this computation is not expensive since it only requires evaluations of the MISC surrogate model) and are quite promising: the benefits increase for higher and higher-order moments, and in particular, the improvement in the estimate of the kurtosis is quite impressive; the results of MISC and Monte Carlo quadrature in the case $r = 1$ are instead substantially equivalent, which is to be expected since they both work with the same surrogate model. This strategy thus mitigates the impact of the noise on the computation of moments. However, it is not entirely satisfactory, since the choice of the number of samples to be employed is non trivial: on the one hand, we need a sufficiently large number of samples to ensure accuracy of the estimates, on the other hand, taking too many samples results in resolving the spurious oscillations. In other words, the chosen number of samples should give the best compromise between these two aspects, and some trial-and-error tuning, or some

accurate a priori analysis should be carried out. Deeper studies of this matter will be subject of future work.

Further, at this point it is also clear why the version of MISC with profit $P^{\mathcal{R}}$, i.e., based on the quadrature, gives better results than the version with profit $P^{\mathcal{S}}$, i.e. based on the quality of the surrogate model. Indeed, in case of noisy simulations, the adaptivity criterion based on the profit $P^{\mathcal{S}}$ results in capturing even more the spurious oscillations due to the numerical noise, since it is based on the direct pointwise evaluation of the noisy surrogate model.

Next, we move to SRBF. As already mentioned, in this application we use a regression approach to compute the weights of the surrogate model (i.e., solving Eq. (18)), motivated by the fact that the evaluations of the CFD solver are noisy as just discussed. The SRBF surrogate model at the last iteration of the adaptive sampling procedure is shown in Fig. 18d. The surface is smoother than the one produced by MISC (cf. Fig. 18c), although a small bump is present in the bottom part. This figure, thus, shows that SRBF is in general effective in filtering-out the numerical noise in the training set.

Figure 23 shows that SRBF spent about 50% of the final computational cost at the first iteration, then requiring simulations on the finest grids only at iterations 5 and 6. In all the other iterations mainly low-fidelity simulations are performed. This sampling behavior is due to the high values of prediction uncertainty that are found in the corners of the parametric domain, because the topology of the initial training leads to extrapolation in those zones. Such corner regions are those with the highest estimated prediction uncertainty, and the adaptive sampling procedure requires all the fidelities before exploring other regions.

Figure 24 shows the convergence of the first four centered moments of R_T and their relative errors obtained with SRBF. The method initially converges rapidly to the reference

(a) Total number of simulations

(b) Number of new simulations

(a) Convergence over the full range of available computational costs.
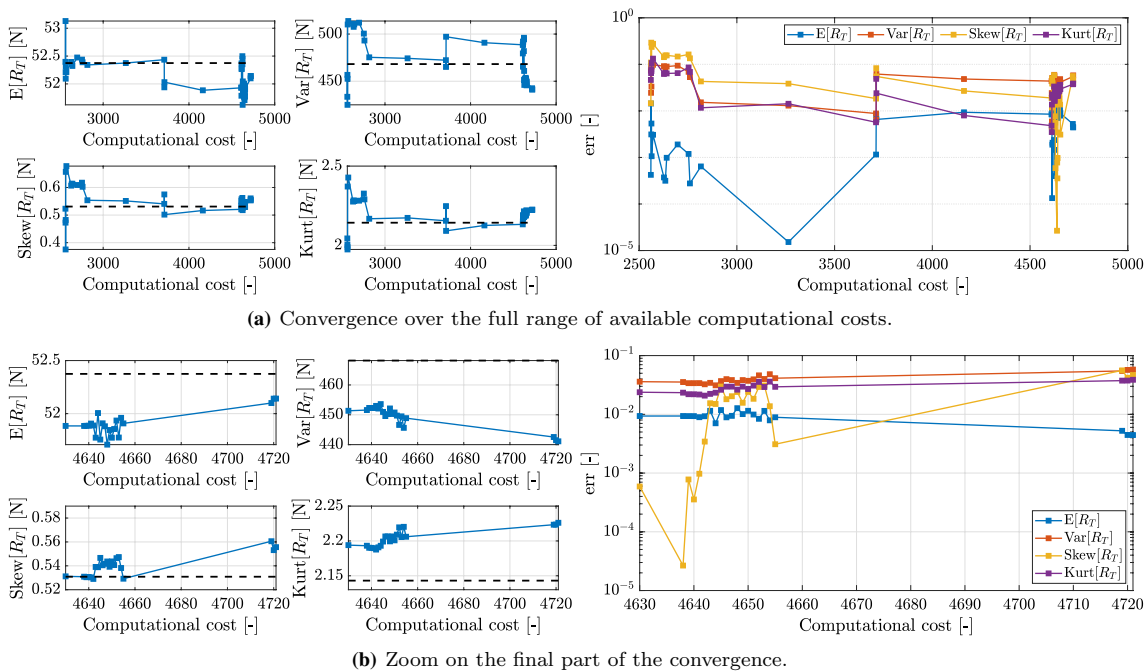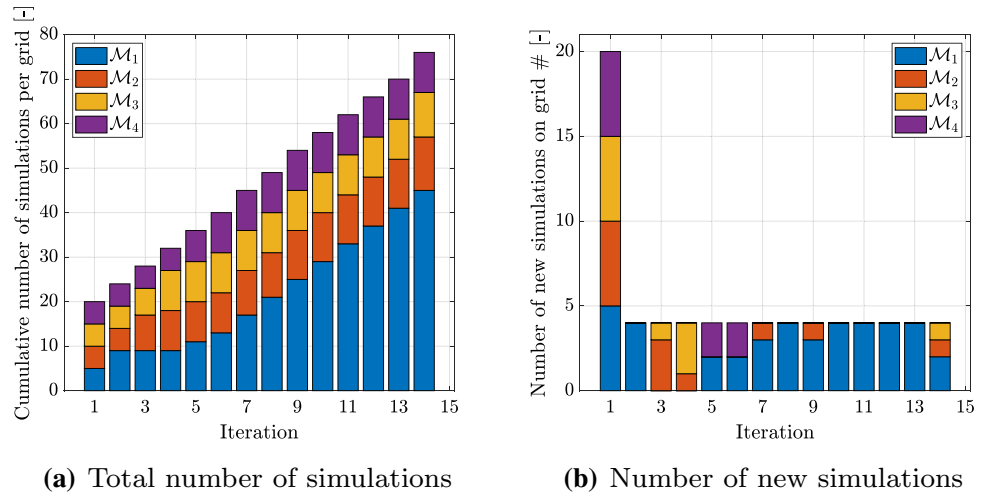
(b) Zoom on the final part of the convergence.

Fig. 24 RoPax problem, results for the SRBF method. Left: values of the first four centered moments. The black dashed line marks the reference value of the moments. Right: relative error of the moments (see Eq. (25))

values but then the metrics start oscillating. Similarly to the analytical problem, the errors of all the moments have a quite similar convergence. This is particularly evident in the last iterations of the adaptive sampling process, see Fig. 24a. Figure 24b shows the detail of the last iterations of the adaptive sampling. The oscillatory behavior is evident and mostly associated with the intensive sampling of the lowest fidelity happening in correspondence with computational costs between 4638 and 4655. In this range the expected value and the skewness oscillates more than the other moments,

indicating that the surrogate model is oscillating around the training data.

To conclude the discussion on the convergence of the moments, in Fig. 25 the convergence of the relative errors of the moments obtained with the two methods is compared (of course we consider MISC results when using Monte Carlo quadrature). Both MISC and SRBF achieve similar values: SRBF reach smaller errors in all moments but the kurtosis, but the convergence trend is bumpier than for MISC, and has a larger initial computational cost. Especially for SRBF,
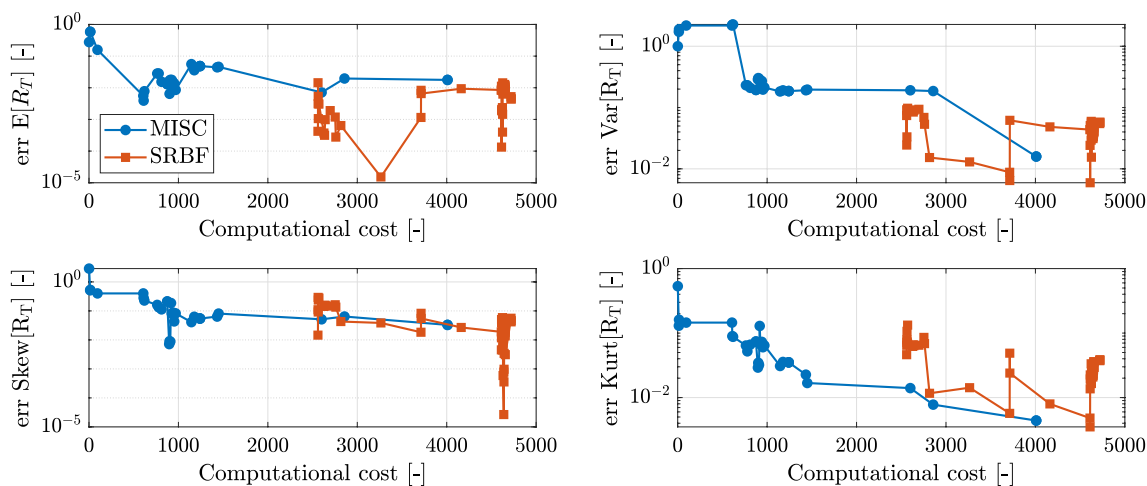
Fig. 25 RoPax problem, comparison of MISC and SRBF results: relative error of the moments. It is a compact visualization of the results of Fig. 22 and 24
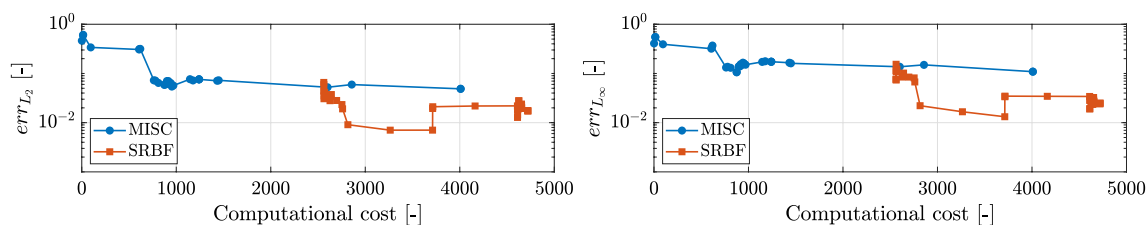
Fig. 26 RoPax problem, comparison of MISC and SRBF methods: relative error of the approximation R_T in L_2 (left) and L_∞ norm (right) (see Eq. (26))

the convergence of the moments is oscillating, nevertheless the oscillations fall within a range that can be considered reasonable from a practical viewpoint as it is comparable with the numerical uncertainties and/or noise of the solver.

In Fig. 26 the relative error in L_2 and L_∞ norms of the estimates of the advancement resistance are plotted. These metrics confirm that the MISC method reaches reasonable estimates with a low computational cost, whereas the SRBF method returns slightly better results but requires an higher computational cost. In the case of SRBF, it is worth noting that in the last iterations the relative errors of variance, skewness, and kurtosis increase whereas the L_2 metric decreases. This apparent contradiction is discussed comparing the convergence of the variance and the L_2 metric. Figure 27a and d show the convergence of the difference of the variances between the multi-fidelity surrogate model prediction R_T and the reference value R_T^*, i.e., of $\Delta\text{Var}[R_T] = \text{Var}[R_T] - \text{Var}[R_T^*]$, and of the $L_2^2(R_T)$ metric, along with the summands of their decompositions: $\Delta\text{Var}[R_T] = \mathbb{E}[R_T^2] - \mathbb{E}[R_T^{*2}] - (\mathbb{E}[R_T]^2 - \mathbb{E}[R_T^*]^2)$,

$$L_2^2(R_T) = \mathbb{E}\left[\left(R_T - R_T^*\right)^2\right] = \mathbb{E}[R_T^2] + \mathbb{E}[R_T^{*2}] - 2\mathbb{E}[R_T R_T^*].$$

To have $\Delta\text{Var}[R_T]$ going to zero it should happen that its two components, $\mathbb{E}[R_T^2] - \mathbb{E}[R_T^{*2}]$ and $\mathbb{E}[R_T]^2 - \mathbb{E}[R_T^*]^2$, go to zero remaining equal in size; however, this does not happen, see Fig. 27b and e. Conversely, the two components of the $L_2^2(R_T)$ metric, i.e., $\mathbb{E}[R_T^2] + \mathbb{E}[R_T^{*2}]$ and $2\mathbb{E}[R_T R_T^*]$, are always almost equal in size (see Fig. 27c and f), therefore the $L_2^2(R_T)$ metric goes to zero. Figure 27e shows a zoom on the last iterations of the $\Delta\text{Var}[R_T]$ convergence: the component $\mathbb{E}[R_T]^2 - \mathbb{E}[R_T^2]^2$ is converging to zero faster than $\mathbb{E}[R_T^2] - \mathbb{E}[R_T^{*2}]$, therefore their difference $\Delta\text{Var}[R_T]$ does not converge to zero.

Finally, in Fig. 28 the PDFs obtained with both methods and with the reference surrogate, as well as the results of the KS test statistic (cf. Eq. (27)), are plotted. Both MISC and SRBF predict well the position of the mode of the PDF and its magnitude and the tails of the distribution, with good agreement with the reference solution. In the range [40 70]N, the PDF obtained by the MISC method is more "wobbly", again due to the presence of noise in the evaluations of the solver, which corrupts the surrogate. Finally, the KS test statistic is seen to be convergent for both methods, implying convergence towards the reference CDF. Both convergences are not monotonic due to the influence of the noisy simulations.
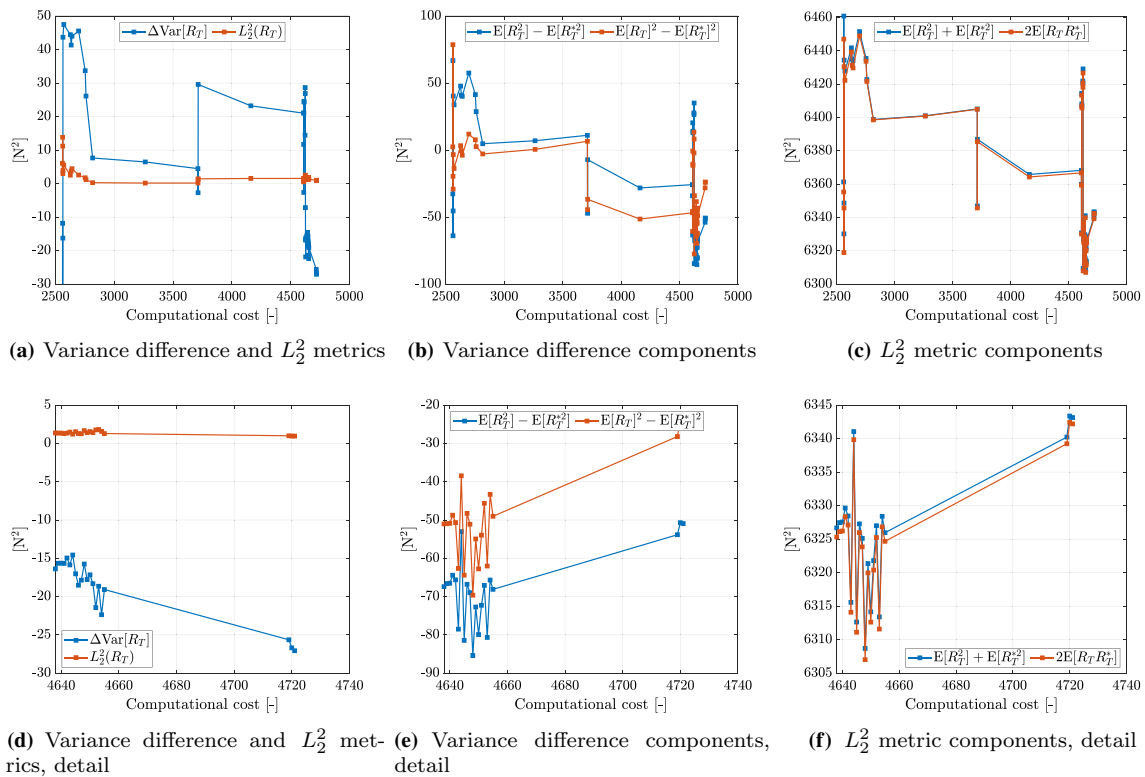
Fig. 27 RoPax problem, results for the SRBF method. Convergence of the variance difference and L_2^2 metrics and their components

4 Conclusions and future work

In this work, two multi-fidelity methods for forward UQ applications, MISC and SRBF, have been presented and applied to two examples to highlight their strengths and critical points. The first numerical test considered in this work is an analytical example, that served as benchmark for the results of the second test, which instead consists in a realistic application in naval engineering and it is more demanding for a number of reasons (noisy evaluations of the quantity of interest, large setup time and computational costs). For the former, the different fidelities considered are Taylor expansions of increasing order of a given function, while for the latter the fidelities are obtained by stopping the multi-grid computations in the RANS solver at different grid levels.

In detail, we have considered the a posteriori adaptive MISC method already presented in [16], with slight modifications on the profit computation, and we have highlighted in passing that MISC is not an interpolatory method, contrary to its single-fidelity counterpart (i.e., sparse grids); this detail was never previously discussed (up to the authors' knowledge) in the MISC literature. SRBF has been used as an interpolatory surrogate model for the analytical test problem and as a regressive surrogate model [22] for the RoPax problem.

For both tests, we have computed a number of error metrics for the quantity of interest (value of the function for the analytic test / advancement resistance for the ferry problem): convergence of the approximation of the first four centered moments, mean squared and maximum prediction errors over a set of samples, and convergence of the CDF (as measured by the Kolmogorov–Smirnov test statistic).

Overall, both MISC and SRBF confirmed to be viable multi-fidelity approaches to forward UQ problems. MISC has an edge in providing reasonable estimates of most statistical quantities at reduced computational cost, but is more sensitive to the noise in the evaluations of the quantities of interest: indeed, noise can strongly influence the adaptive selection process of the multi-indices, corrupt the estimates of, e.g., higher-order moments (e.g., skewness, kurtosis), and introduce artifacts in the estimation of the PDF of the quantities of interest. With respect to the first issue, a quadrature-based adaptive criterion is expected to be more robust than a criterion based on the pointwise accuracy of the surrogate model. A possible strategy to mitigate the second issue, that consists in computing higher-order statistical moments by taking Monte Carlo samples of the MISC surrogate model and then computing the moments from such set of values by sample formulas (sample variance/skewness/kurtosis), has been proposed but it is not entirely satisfactory, since it is
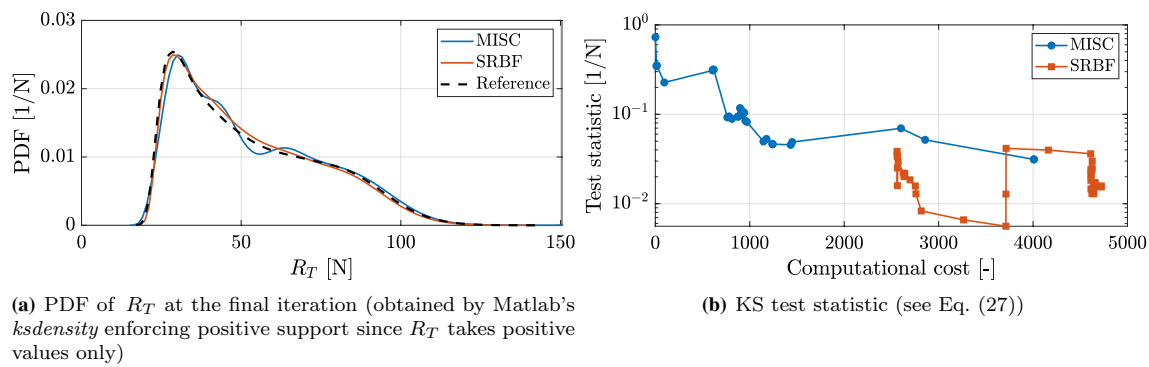
(a) PDF of R_T at the final iteration (obtained by Matlab's *ksdensity* enforcing positive support since R_T takes positive values only)

(b) KS test statistic (see Eq. (27))

Fig. 28 RoPax problem, comparison of MISC and SRBF methods

not clear how to choose an appropriate number of samples (enough to be accurate, not too many to avoid resolving the scale of the noise). This aspect deserves more investigations and is one of the subjects of future work. Another practical issue is caused by the non-monotonic behavior of the profits, where some indices with low profits shade useful neighbors, thus delaying the convergence of MISC. More robust strategies to explore the set of multi-indices, that blend the profit-based selection of indices with other criteria are also subject of future work; see, e.g., [40, 43], where this problem was discussed in the context of adaptive sparse-grids quadrature/interpolation.

Conversely, SRBF is less prone to the issue of noisy evaluations but on the other hand the current initialization strategy needs to sample all available fidelities, which results in a significantly larger initialization cost. Different initialization strategies are under investigation to reduce this gap: a possible approach would be, e.g., to build the initial training set using only one evaluation for each fidelity other than the lowest one, instead of $2N + 1$ as in the current implementation, see [55]. This would allow the adaptive sampling method to freely define the best training set for the higher fidelities.

Note that the discussion on computational costs in this work was set up in terms of sheer evaluation costs of the fidelities, i.e., neglecting the CPU-time related to the other operations required by MISC and SRBF. In details, MISC needs to keep track of profits, update the multi-index sets Λ and Mar(Λ), compute the combination technique coefficients in Eq. (7), generate the tensor grids corresponding to the multi-indices in Λ and Mar(Λ), and evaluate the associated multivariate Lagrange polynomials. Conversely, SRBF requires to solve the linear system (18) at each iteration (the dimension of the linear system grows at each iteration; furthermore, when the LOOCV procedure is performed, the linear system is solved for each tested value of the number of centers $\mathcal{K}$), as well as to solve the problem in Eq. (23).

An aspect worth investigating for both methods is how to incorporate available soft information (monotonicity, multimodality, etc.) on the physical nature of the problem in the construction of the surrogate models. In the particular RoPax problem considered here, the resistance is expected to be monotone increasing with respect to advancement speed and draught: such property could be preserved by employing, e.g., least-squares regressions with appropriate polynomial degrees and/or monotonic smoothing, see, e.g., [56].

Concerning the application to naval UQ, once fixed the current limitations of both methods, future research will address more complex test cases, such as a larger number of uncertain parameters and more realistic conditions, to take, e.g., regular/irregular waves into account.

References

1. Serani A, Stern F, Campana EF, Diez M (2021) Hull-form stochastic optimization via computational-cost reduction methods. Eng Comput 1–25

2. Serani A, Diez M, van Walree F, Stern F (2021) URANS analysis of a free-running destroyer sailing in irregular stern-quartering waves at sea state 7. Ocean Eng 237:109600

3. Stern F, Volpi S, Gaul NJ, Choi K, Diez M, Broglia R, Durante D, Campana EF , Iemma U(2017) Development and assessment of uncertainty quantification methods for ship hydrodynamics. In: 55th AIAA aerospace sciences meeting, p. 1654

4. Diez M, He W, Campana EF, Stern F (2014) Uncertainty quantification of delft catamaran resistance, sinkage and trim for variable froude number and geometry using metamodels, quadrature and Karhunen-Loève expansion. J Mar Sci Technol 19(2):143–169

5. He W, Diez M, Zou Z, Campana EF, Stern F (2013) URANS study of Delft catamaran total/added resistance, motions and slamming loads in head sea including irregular wave and uncertainty quantification for variable regular wave and geometry. Ocean Eng 74:189–217

6. Diez M, Broglia R, Durante D, Olivieri A, Campana EF, Stern F (2018) Statistical assessment and validation of experimental and computational ship response in irregular waves. J Verif Valid Uncertain Quantif 3(2):021004

7. Durante D, Broglia R, Diez M, Olivieri A, Campana E, Stern F (2020) Accurate experimental benchmark study of a catamaran in regular and irregular head waves including uncertainty quantification. Ocean Eng 195:106685

8. Quagliarella D, Serani A, Diez M, Pisaroni M, Leyland P, Montagliani L, Iemma U, Gaul NJ, Shin J, Wunsch D, Hirsch C, Choi K, Stern F (2019) Benchmarking uncertainty quantification methods using the NACA 2412 airfoil with geometrical and operational uncertainties. In: 57th AIAA Aerospace Sciences Meeting, SciTech 2019, p 3555

9. Beran PS, Bryson DE, Thelen AS, Diez M, Serani A (2020) Comparison of multi-fidelity approaches for military vehicle design, in: 21th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference (MA&O), AVIATION 2020, Virtual Event, June 15-19,

10. Giles MB (2008) Multilevel Monte Carlo path simulation. Oper Res 56(3):607–617

11. Cliffe K, Giles M, Scheichl R, Teckentrup A (2011) Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. Comput Vis Sci 14(1):3–15

12. Haji-Ali A-L, Nobile F, Tempone R (2015) Multi-index Monte Carlo: when sparsity meets sampling. Numerische Mathematik 1–40

13. Kuo FY, Schwab C, Sloan I (2015) Multi-level Quasi-Monte carlo finite element methods for a class of elliptic PDEs with random coefficients. Found Comput Math 15(2):411–449

14. Teckentrup AL, Jantsch P, Webster CG, Gunzburger M (2015) A multilevel stochastic collocation method for partial differential equations with random input data. SIAM/ASA J Uncertain Quantif 3(1):1046–1074

15. Beck J, Tamellini L, Tempone R (2019) IGA-based Multi-Index Stochastic Collocation for random PDEs on arbitrary domains. Comput Methods Appl Mech Eng 351:330–350

16. Jakeman JD, Eldred M, Geraci G, Gorodetsky A (2020) Adaptive multi-index collocation for uncertainty quantification and sensitivity analysis. Int J Numer Methods Eng 121(6):1314–1343

17. Haji-Ali A, Nobile F, Tamellini L, Tempone R (2016) Multi-index stochastic collocation for random PDEs. Comput Methods Appl Mech Eng 306:95–122

18. Haji-Ali A-L, Nobile F, Tamellini L, Tempone R (2016) Multi-index Stochastic Collocation convergence rates for random PDEs with parametric regularity. Found Comput Math 16(6):1555–1605

19. Haji-Ali A-L, Nobile F, Tempone R, Wolfers S (2020) Multi-level weighted least squares polynomial approximation. ESAIM: M2AN 54(2):649–677

20. Han Z-H, Görtz S (2012) Hierarchical kriging model for variable-fidelity surrogate modeling. AIAA J 50(9):1885–1896

21. Baar Jd, Roberts S, Dwight R, Mallol B (2015) Uncertainty quantification for a sailing yacht hull, using multi-fidelity kriging. Comput Fluids 123:185–201

22. Wackers J, Visonneau M, Pellegrini R, Ficini S, Serani A, Diez M (2020) Adaptive N-fidelity metamodels for noisy CFD data. In: 21th AIAA/ISSMO multidisciplinary analysis and optimization conference (MA&O), AVIATION 2020, Virtual Event, June 15–19,

23. Serani A, Pellegrini R, Wackers J, Jeanson C-J, Queutey P, Visonneau M, Diez M (2019) Adaptive multi-fidelity sampling for CFD-based optimization via radial basis functions metamodel. Int J Comput Fluid Dyn 33(6–7):237–255

24. Han Z-H, Görtz S, Zimmermann R (2013) Improving variable-fidelity surrogate modeling via gradient-enhanced kriging and a generalized hybrid bridge function. Aerosp Sci Technol 25(1):177–189

25. Peherstorfer B, Willcox K, Gunzburger M (2018) Survey of multifidelity methods in uncertainty propagation, inference, and optimization. SIAM Review 60(3):550–591

26. Gorodetsky A, Geraci G, Eldred MS, Jakeman J (2020) A generalized approximate control variate framework for multifidelity uncertainty quantification. J Comput Phys 408:109257

27. Pisaroni M, Nobile F, Leyland P (2017) A multilevel Monte Carlo evolutionary algorithm for robust aerodynamic shape design, in: 18th AIAA/ISSMO multidisciplinary analysis and optimization conference. Denver, Colorado

28. Pisaroni M, Nobile F, Leyland P (2017) A continuation multi level Monte Carlo (C-MLMC) method for uncertainty quantification in compressible inviscid aerodynamics. Comput Methods Appl Mech Eng 326:20–50

29. Geraci G, Eldred MS, Gorodetsky A, Jakeman J (2019) Recent advancements in Multilevel-Multifidelity techniques for forward UQ in the DARPA sequoia project, in: 57th AIAA Aerospace Sciences Meeting, SciTech 2019, p. 0722

30. Di Mascio A, Broglia R, Muscari R (2007) On the application of the single-phase level set method to naval hydrodynamic flows. Comput Fluids 36(5):868–886

31. Di Mascio A, Broglia R, Muscari R (2009) Prediction of hydrodynamic coefficients of ship hulls by high-order godunov-type methods. J Mar Sci Technol 14(1):19–29

32. Broglia R, Durante D (2018) Accurate prediction of complex free surface flow around a high speed craft using a single-phase level set method. Comput Mech 62(3):421–437

33. Piazzola C, Tamellini L, Pellegrini R, Broglia R, Serani A, Diez M (2020) Uncertainty quantification of ship resistance via multi-index stochastic collocation and radial basis function surrogates: A comparison, in: 21th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference (MA&O), AVIATION 2020, Virtual Event, June 15-19,

34. Rosenblatt M (1956) Remarks on some nonparametric estimates of a density function. Ann Math Stat 27(3):832–837

35. Parzen E (1962) On estimation of a probability density function and mode. Ann Math Stat 33(3):1065–1076

36. Trefethen LN (2008) Is Gauss quadrature better than Clenshaw-Curtis? SIAM Rev 50(1):67–87

37. Narayan A, Jakeman JD (2014) Adaptive leja sparse grid constructions for stochastic collocation and high-dimensional approximation. SIAM J Sci Comput 36(6):A2952–A2983

38. Nobile F, Tamellini L, Tempone R (2015) Comparison of Clenshaw-Curtis and Leja Quasi-Optimal Sparse Grids for the Approximation of Random PDEs. In: Kirby RM, Berzins M, Hesthaven JS (Eds.), Spectral and high order methods for partial differential equations-ICOSAHOM '14, Vol. 106 of lecture notes in computational science and engineering, Springer International Publishing, pp. 475–482

39. Patterson TNL (1968) The optimum addition of points to quadrature formulae. Math Comput 22: 847–856; addendum, ibid. 22 (104) C1–C11

40. Chkifa A, Cohen A, Schwab C (2014) High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs. Found Comput Math 14(4):601–633

41. Schillings C, Schwab C (2013) Sparse, adaptive Smolyak quadratures for Bayesian inverse problems. Inverse Probl 29(6):065011

42. Nobile F, Tamellini L, Tesei F, Tempone R (2016) An adaptive sparse grid algorithm for elliptic PDEs with lognormal diffusion coefficient. In: Garcke J, Pflüger D (Eds.) Sparse grids and applications– Stuttgart 2014, Vol. 109 of Lecture Notes in Computational Science and Engineering, Springer International Publishing Switzerland, pp. 191–220

43. Gerstner T, Griebel M (2003) Dimension-adaptive tensor-product quadrature. Computing 71(1):65–87

44. Gutmann HM (2001) A radial basis function method for global optimization. J Glob Opt 19(3):201–227

45. Forrester AIJ, Keane AJ (2009) Recent advances in surrogate-based optimization. Prog Aerosp Sci 45(1–3):50–79

46. Volpi S, Diez M, Gaul NJ, Song H, Iemma U, Choi KK, Campana EF, Stern F (2015) Development and validation of a dynamic metamodel based on stochastic radial basis functions and uncertainty quantification. Struct Multidiscip Opt 51(2):347–368

47. Lloyd S (1982) Least squares quantization in PCM. IEEE Trans Inf theory 28(2):129–137

48. Li X, Gao W, Gu L, Gong C, Jing Z, Su H (2017) A cooperative radial basis function method for variable-fidelity surrogate modeling. Struct Multidiscip Opt 56(5):1077–1092

49. Serani A, Pellegrini R, Broglia R, Wackers J, Visonneau M, Diez M (2019) An adaptive N-fidelity metamodel for design and operational-uncertainty space exploration of complex industrail problems. In: Proceedings of the 8th International Conference on Computational Methods in Marine Engineering (Marine 2019), pp. 177–188

50. Wackers J, Visonneau M, Serani A, Pellegrini R, Broglia R, Diez M (2020) Multi-fidelity machine learning from adaptive- and multi-grid RANS simulations. In: Proceedings of the 33rd Symposium on Naval Hydrodynamics, Osaka, Japan,

51. Serani A, Leotardi C, Iemma U, Campana EF, Fasano G, Diez M (2016) Parameter selection in synchronous and asynchronous deterministic particle swarm optimization for ship hydrodynamics problems. Appl Soft Comput 49:313–334

52. Bäck J, Nobile F, Tamellini L, Tempone R (2011) Stochastic spectral Galerkin and collocation methods for PDEs with random coefficients: a numerical comparison. In: Hesthaven J, Ronquist E (Eds.), Spectral and high order methods for partial differential equations, vol. 76 of lecture notes in computational science and engineering, Springer, pp. 43–62, selected papers from the ICO-SAHOM '09 conference, June 22-26, Trondheim, Norway

53. Broglia R, Zaghi S, Muscari R, Salvadore F (2014) Enabling hydrodynamics solver for efficient parallel simulations. In: 2014 International conference on high performance computing and simulation (HPCS), IEEE, pp. 803–810

54. Spalart P, Allmaras S (1992) A one-equation turbulence model for aerodynamic flows. In: 30th aerospace sciences meeting and exhibit, p. 439

55. Pellegrini R, Serani A, Diez M, Visonneau M, Wackers J (2021) Towards automatic parameter selection for multi-fidelity surrogate-based optimization. In: 9th Conference on computational methods in marine engineering (Marine 2021), Virtual conference, 2–4 June,

56. Dykstra RL, Robertson T (1982) An algorithm for isotonic regression for two or more independent variables. Ann Stat 10(3):708–716