

Uncertainty Quantification for a Sailing Yacht Hull, Using Multi-Fidelity Kriging

de Baar, Jouke; Roberts, S; Dwight, Richard; Mallol, B.

DOI

[10.1016/j.compfluid.2015.10.004](https://doi.org/10.1016/j.compfluid.2015.10.004)

Publication date

2015

Document Version

Accepted author manuscript

Published in

Computers & Fluids

Citation (APA)

de Baar, J., Roberts, S., Dwight, R., & Mallol, B. (2015). Uncertainty Quantification for a Sailing Yacht Hull, Using Multi-Fidelity Kriging. *Computers & Fluids*, 123, 185-201.
<https://doi.org/10.1016/j.compfluid.2015.10.004>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Uncertainty Quantification for a Sailing Yacht Hull, Using Multi-Fidelity Kriging

Jouke de Baar^a, Stephen Roberts^b, Richard Dwight^c, Benoit Mallol^d

Abstract

Uncertainty Quantification (UQ) for CFD-based ship design can require a large number of simulations, resulting in significant overall computational cost. Presently, we use an existing method, multi-fidelity Kriging, to reduce the number of simulations required for the UQ analysis of the performance of a sailing yacht hull, considering uncertainties in the tank blockage, mass and centre of gravity. We compare the UQ results with experimental values.

Keywords: Uncertainty Quantification, Multi-fidelity, Kriging, RANS, Free-surface

1. Introduction

Initial error analysis for marine applications focused on grid convergence and time discretisation errors [34, 40, 33]. Recently, as a result of these efforts (*i.e.* the reduction of discretisation errors) and due to a rapid increase of computer power, the focus of uncertainty analysis has been shifting from grid convergence and time discretisation errors towards stochastic uncertainty

^aAustralian National University (j.h.s.debaar@gmail.com)

^bAustralian National University

^cDelft University of Technology

^dNumeca, Brussels

quantification (UQ) [33, 6]. The main challenge, especially when considering multiple uncertain input parameters, is to decrease the number of simulations required to arrive at accurate UQ results [29, 30, 38]. A promising approach to this challenge are non-intrusive multi-fidelity methods, which combine a small number of expensive high-fidelity simulations with a larger number of less expensive low-fidelity simulations [36]. A recently developed example of such methods is multi-fidelity Kriging [17, 9]. An important asset of multi-fidelity Kriging is that — once the framework has been set up — it can not only be exploited for UQ, but also for parameter calibration and shape optimisation [16, 9].

An extensive overview of multi-fidelity methods is given in [36], we provide a brief discussion. In an early application, Haftka [12] presents a multiplacitive multi-fidelity analysis of a clamped beam. Mason et al. [26] augment results from three-dimensional simulations with results from two-dimensional simulations, and use polynomial regression to model a scale factor. Kennedy and O’Hagan [17] apply multi-fidelity Kriging to simulations of an oil reservoir. Forrester et al. [9] discuss the issue of sampling plans and apply multi-fidelity Kriging to shape optimisation of a transonic aircraft wing. Kuya et al. [23] use multi-fidelity Kriging to augment wind-tunnel data with CFD results. de Baar et al. [5] develop a fast way to estimate the hyperparameters of large data sets, which they incorporate into multi-fidelity Kriging to augment satellite data with F-16-acquired terrain elevation data. In a recent paper, Toal [35] presents a best practise for the application of multi-fidelity Kriging and applies it to the optimisation of a jet engine compressor and a transonic airfoil. Multi-level Monte Carlo simulations [1] are related to multi-

fidelity Kriging, however the difference correction [1, Eq. 5] acts directly on the UQ estimator instead of on the surrogate.

Presently, we investigate the free-surface flow around a sailing yacht hull over a range of velocities. We consider the uncertainties in the blockage, mass and centre of gravity, and use multi-fidelity Kriging to propagate these uncertainties to the resulting resistance, sinkage and pitch. We propose to reduce the cost of the uncertainty propagation by balancing a small number of high-fidelity fine grid simulations with a larger number of low-fidelity coarse grid simulations. In this case, we find that we can simplify multi-fidelity Kriging by setting the regression parameter $\rho = 1$, thus avoiding an additional hyperparameter estimate. We then validate the UQ results by comparing them to experimental results.

The purpose of the present work is to investigate and demonstrate the possibilities of using multi-fidelity Kriging-based UQ to analyse ship performance. In addition, we briefly demonstrate the flexibility of the approach by exploiting the multi-fidelity Kriging response for parameter calibration.

2. Methodology

2.1. Kriging

Kriging was originally developed as a method for spatial regression in geology [27] and meteorology [10]. A comprehensive overview of Kriging can be found in [2], while a detailed investigation of the origins of Kriging is made in [3]. An excellent introduction to Kriging in engineering applications is provided in [8]. A lucid derivation of Kriging in a Bayesian framework can be found in [39].

Along the lines of [39], consider a process with a normally distributed prior:

$$\mathbf{X} \sim \mathcal{N}(\mu, P), \quad (1)$$

with mean μ , covariance matrix P and a discrete number of process outputs \mathbf{x} — in the present application, the output \mathbf{x} is either the resistance, sinkage or pitch. For observations \mathbf{y} — in the present application, ‘observations’ are CFD simulations — assume a normally distributed and unbiased likelihood:

$$\mathbf{Y}|\mathbf{x} \sim \mathcal{N}(H\mathbf{x}, R), \quad (2)$$

with observation matrix H and error covariance matrix R . Here H is a matrix filled with zeros and ones which selects the observations from the process \mathbf{X} . Now, by applying Bayes’ Rule:

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x}) p(\mathbf{x})}{p(\mathbf{y})}, \quad (3)$$

the expectation of the process \mathbf{X} , conditional on a set of observations \mathbf{y} , is given by the Kriging predictor [39]:

$$\mathbb{E}(\mathbf{X}|\mathbf{y}) = \mu + K(\mathbf{y} - \mu), \quad (4)$$

and:

$$\text{cov}(\mathbf{X}|\mathbf{y}) = (I - KH)P, \quad (5)$$

with the gain matrix:

$$K = PH^T(R + HPH^T)^{-1}. \quad (6)$$

The Kriging predictor covariance (5) is not used in this paper.

An accurate Kriging prediction depends on the choice of the mean μ , the covariance matrix P and error covariance matrix R , contained in the prior and the likelihood, respectively. For P , we choose the following parameterisation:

$$P_{ij} = \sigma^2 \exp \left(- \sum_{k=1}^d \frac{|\xi_{j,k} - \xi_{i,k}|^2}{2\theta_k^2} \right). \quad (7)$$

Here we have introduced the input parameters $\boldsymbol{\xi}$, which form a d -dimensional input ‘parameter space’. The smoothness of the process is represented by the ‘hyperparameters’ $\boldsymbol{\theta}$, while σ^2 is the variance of the process. For the error covariance matrix, we choose:

$$R = \epsilon^2 I, \quad (8)$$

with ϵ small, representing an uncorrelated Gaussian ‘measurement’ uncertainty.

In the ideal case, we would select μ and the hyperparameter(s) $\boldsymbol{\theta}$ in (7) a priori. In most cases, however, we do not have any a priori information to base this selection on, so we choose the arithmetic mean of the observations for μ , and tune $\boldsymbol{\theta}$ to the observations \mathbf{y} instead, using a maximum likelihood estimate [25, 21, 20, 24]. This is equivalent to minimising the log-likelihood:

$$\log p(\boldsymbol{\theta}|\mathbf{y}) \propto \ln|A| + (\mathbf{y} - \mu)^T A^{-1}(\mathbf{y} - \mu), \quad (9)$$

with:

$$A = (R + HPH^T), \quad (10)$$

where P depends on θ , as given in (7), and where $|\cdot|$ is the determinant. Currently, we use a Nelder-Mead simplex direct search to minimise (9), the

computational cost of which is relatively small, given the limited number of observations.

As an example, Figure 1 shows the process:

$$x = 3 + \cos(5\xi^2) + \frac{1}{5} \sin(40\xi), \quad (11)$$

together with 6 samples and the Kriging prediction. In this case, the number of samples is too small to make a detailed prediction, i.e. we are under-sampling the process. One solution would be to add more samples, however, these samples can come at high computational cost. In the next section, we will see how we can augment the present data with a set of low-fidelity samples.

2.2. Kriging with multi-fidelity data

Multi-fidelity Kriging is a form of co-Kriging [2, page 138] [4, pages 234], and was originally developed in geostatistics, for explorations where measurements of different ores are available. A derivation of multi-fidelity Kriging is given in [17, 9]. Here we assume that we have a set of high-fidelity CFD simulations \mathbf{y}_{HF} , which are expensive to evaluate and provide an accurate output, as well as a set of low-fidelity simulations \mathbf{y}_{LF} , which are less expensive to evaluate but provide only an approximation of the output. The strategy is to augment a small number of high-fidelity simulations with a larger number of low-fidelity simulations.

Along the lines of [17, 9] let us assume that the high-fidelity process \mathbf{X}_{HF} is related to the low-fidelity process \mathbf{X}_{LF} as:

$$\mathbf{X}_{\text{HF}} = \rho \mathbf{X}_{\text{LF}} + \mathbf{X}_{\text{D}}, \quad (12)$$




Figure 1: Kriging prediction $E(\mathbf{x}|\mathbf{y})$ of the process x . The process is undersampled, and the prediction does not capture the higher frequencies.

where ρ is a regression parameter. We now have a new process \mathbf{X}_D that represents the difference between the low-fidelity and the high-fidelity output. This difference process \mathbf{X}_D will ‘correct’ the low-fidelity output.

In this particular case, provided that both the low-fidelity and the high-fidelity simulations are in the asymptotic range of grid convergence, the difference between the low-fidelity and the high-fidelity output is an effect of the grid error:

$$\begin{aligned} y_{\text{LF}} &= y_{\text{exact}} + Ch_{\text{LF}}^p + \text{H.O.T.} \\ y_{\text{HF}} &= y_{\text{exact}} + Ch_{\text{HF}}^p + \text{H.O.T.}, \end{aligned}$$

with gridsize h and grid convergence order p , such that:

$$y_{\text{HF}} \approx y_{\text{LF}} - Ch_{\text{LF}}^p. \tag{13}$$

Motivated by the similarity of (12) and (13), at the stage of hyperparameter estimation we will choose $\rho = 1$ for the present application, such that. Given that choice, the difference process \mathbf{x}_D is effectively a model of (minus) the grid convergence error Ch_{LF}^p of the low-fidelity simulation.

The key requirement for increasing the accuracy of the prediction, compared to regular Kriging, is that the difference process \mathbf{X}_D is smoother — thus has longer correlation ranges — than the low-fidelity and high-fidelity process. We assume that we satisfy this requirement when the high-fidelity and low-fidelity outputs are highly correlated.

To predict the multi-fidelity process:

$$\mathbf{X}_{\text{MF}} = \begin{pmatrix} \mathbf{X}_{\text{LF}} \\ \mathbf{X}_D \end{pmatrix}$$

our approach is to define a multi-fidelity prior covariance matrix P_{LF} and observation matrix H_{MF} and substitute them in the Kriging predictor (4–6). Because we are only interested in predicting the high-fidelity process, not in predicting the difference process, we also introduce a selector matrix H_{select} .

First, we assume that \mathbf{X}_{D} is independent of \mathbf{X}_{LF} , such that we have the multi-fidelity prior covariance matrix:

$$P_{\text{MF}} = \begin{pmatrix} P_{\text{LF}} & 0 \\ 0 & P_{\text{D}} \end{pmatrix}, \quad (14)$$

with P_{LF} and P_{D} the parameterised prior covariance matrices for the low-fidelity process \mathbf{X}_{LF} and the difference process \mathbf{X}_{D} , respectively.

As the second step, we write the multi-fidelity observation matrix:

$$H_{\text{MF}} = \begin{pmatrix} H_{\text{LF}} & 0 \\ \rho H_{\text{HF}} & H_{\text{HF}} \end{pmatrix}, \quad (15)$$

where the top row represents the observations of the low-fidelity process, while the bottom exploits (12) to observe the high-fidelity process. (It might seem counterintuitive to have ρH_{HF} , however, H_{HF} only acts to select observations. It can be checked that ρH_{HF} and H_{HF} have the same number of rows, and that substitution of (15) in (4) gives (19) and (20).) Note that $\mathbf{Y}_{\text{MF}}|\mathbf{x}_{\text{MF}} \sim \mathcal{N}(H_{\text{MF}}\mathbf{x}_{\text{MF}}, R)$ is now a compiled vector of low-fidelity and high-fidelity observations, and $\boldsymbol{\mu}_{\text{MF}}$ will be a vector of corresponding means.

Thirdly, as we are now considering a low-fidelity and a high-fidelity process, we should indicate that the multi-fidelity predictor predicts x_{HF} , the high-fidelity process, by defining:

$$H_{\text{select}} = \begin{pmatrix} \rho I & I \end{pmatrix}, \quad (16)$$

such that:

$$\begin{aligned}\mathbf{X}_{\text{HF}} &= H_{\text{select}}\mathbf{X}_{\text{MF}} \\ &= \rho \mathbf{X}_{\text{LF}} + \mathbf{X}_{\text{D}}.\end{aligned}$$

We can then substitute (14–15) into (4) to write the multi-fidelity Kriging predictor:

$$\mathbb{E}(\mathbf{X}_{\text{HF}}|\mathbf{y}_{\text{MF}}) = \boldsymbol{\mu}_{\text{HF}} + K_{\text{MF}}(\mathbf{y}_{\text{MF}} - \boldsymbol{\mu}_{\text{MF}}), \quad (17)$$

where:

$$\mathbf{y}_{\text{MF}} = \begin{pmatrix} \mathbf{y}_{\text{LF}} \\ \mathbf{y}_{\text{D}} \end{pmatrix} = \begin{pmatrix} \mathbf{y}_{\text{LF}} \\ \mathbf{y}_{\text{HF}} - \rho\mathbf{y}_{\text{LF}} \end{pmatrix}$$

and:

$$K_{\text{MF}} = H_{\text{select}}P_{\text{MF}}H_{\text{MF}}^T(R + H_{\text{MF}}P_{\text{MF}}H_{\text{MF}}^T)^{-1}, \quad (18)$$

with:

$$H_{\text{mf}}P_{\text{mf}}H_{\text{mf}}^T = \begin{pmatrix} H_{\text{LF}}P_{\text{LF}}H_{\text{LF}}^T & \rho H_{\text{LF}}P_{\text{LF}}H_{\text{HF}}^T \\ \rho H_{\text{HF}}P_{\text{LF}}H_{\text{LF}}^T & \rho^2 H_{\text{HF}}P_{\text{LF}}H_{\text{HF}}^T + H_{\text{HF}}P_{\text{D}}H_{\text{HF}}^T \end{pmatrix}, \quad (19)$$

and:

$$H_{\text{select}}P_{\text{MF}}H_{\text{MF}}^T = \begin{pmatrix} \rho H_x P_{\text{LF}}H_{\text{HF}}^T & \rho^2 P_{\text{LF}}H_{\text{HF}}^T + P_{\text{HF}}H_{\text{HF}}^T \end{pmatrix} \quad (20)$$

Motivated by (12) and (13) we choose $\rho = 1$, which results in a purely additive formulation of (12). We are then left with the problem of estimating the correlation ranges. We recall that the low-fidelity process and the difference process are uncorrelated, such that we can estimate θ_{D} and θ_{LF} independently. First, we choose θ_{LF} to minimise:

$$-\log p(\boldsymbol{\theta}_{\text{LF}}|\mathbf{y}_{\text{LF}}) \propto \ln|A_{\text{LF}}| + (\mathbf{y}_{\text{LF}} - \boldsymbol{\mu}_{\text{LF}})^T A_{\text{LF}}^{-1}(\mathbf{y}_{\text{LF}} - \boldsymbol{\mu}_{\text{LF}}), \quad (21)$$

with:

$$A_{\text{LF}} = R_{\text{LF}} + H_{\text{LF}}P_{\text{LF}}H_{\text{LF}}^T, \quad (22)$$

recalling that P_{LF} depends on θ_{LF} . Then we choose the parameter θ_{D} to minimise:

$$-\log p(\boldsymbol{\theta}_{\text{D}}|\mathbf{y}_{\text{D}}) \propto \ln|A_{\text{D}}| + (\mathbf{y}_{\text{D}} - \boldsymbol{\mu}_{\text{D}})^T A_{\text{D}}^{-1}(\mathbf{y}_{\text{D}} - \boldsymbol{\mu}_{\text{D}}), \quad (23)$$

with $\mathbf{y}_{\text{D}} = \mathbf{y}_{\text{HF}} - \rho\mathbf{y}_{\text{LF}}$ and:

$$A_{\text{D}} = R_{\text{D}} + H_{\text{HF}}P_{\text{D}}H_{\text{HF}}^T. \quad (24)$$

Returning to example (11), we now consider the high-fidelity and low-fidelity processes:

$$\mathbf{x}_{\text{HF}} = 3 + \cos(5\xi^2) + \frac{1}{5} \sin(40\xi), \quad (25)$$

$$\mathbf{x}_{\text{LF}} = \frac{1}{2} + \cos(5\xi^2) + \frac{1}{5} \sin(40\xi) + \frac{3}{2} \sin(3\xi). \quad (26)$$

In Figure 2 we augment the 6 high-fidelity solves with 46 low-fidelity solves, which improves the accuracy of the prediction significantly.

2.3. Uncertainty propagation

After obtaining the (multi-fidelity) Kriging predictor, we can apply a standard surrogate-based numerical integration approach to propagate the input uncertainties $p(\boldsymbol{\xi})$. This is illustrated in Figure 3.

Presently, because of the scalability to higher dimensions, we generate a sparse grid quadrature with nodes ξ_i and weights w_i [22]. We evaluate the predictor $E(x_i|\mathbf{y})$ on the sparse grid nodes to compute the output mean:

$$\begin{aligned} \mu_x &= \int p(\xi)x(\xi) \, \text{d} \xi, \\ &\approx \sum_i w_i p(\xi_i) E(x_i|\mathbf{y}), \end{aligned} \quad (27)$$




Figure 2: Multi-fidelity Kriging prediction $E(\mathbf{x}|\mathbf{y})$ of the process x . Although the low-fidelity data is inaccurate, it helps to resolve the higher frequencies of the process.




Figure 3: Multi-fidelity Kriging-based uncertainty propagation.

and variance:

$$\begin{aligned}
 \sigma_x^2 &= \int p(\xi) [x(\xi) - \mu_x]^2 \\
 &\approx \sum_i w_i p(\xi_i) [E(x_i|\mathbf{y}) - \mu_x]^2.
 \end{aligned} \tag{28}$$

It should be mentioned that this approach can be used to compute higher order statistics of the individual outputs, such as skewness and kurtosis, as well as the covariance of a pair of outputs.




Figure 4: Lines plan of the DTMB 5415 towing tank model, which is used to illustrate multi-fidelity Kriging. The lines plan shows the sections, buttocks, and waterlines, which are cuts parallel to the yz -, xz - and xy -planes, respectively, as well as two diagonal cuts.

Length on Waterline	5.72	m
Breadth	0.76	m
Draft	0.248	m
Displacement	0.549	m ³
Wetted Surface	4.786	m ²

Table 1: Hydrostatics of the DTMB 5415 towing tank model.

2.4. Two-parameter example: DTMB 5415

After introducing the multi-fidelity Kriging equations, let us consider a simple two-dimensional example. A well-known test case is hull number 5415 of the David Taylor Model Basin (DTMB). The lines plan of the DTMB 5415 is provided in Figure 4, while the hydrostatics are given in Table 1.

We use FINE/Marine [7, 37, 13], a free-surface RANS solver, to compute the resistance, sinkage and pitch (a detailed discussion of the solver can be

found in Section 3.4). For a velocity of 2.1 m s^{-1} (Froude number $\text{Fn} = 0.28$, Reynolds number $\text{Re} = 1 \times 10^7$), the result for different computational grids is shown in Figure 5, where the discretisation errors have been computed with respect to a reference grid of 2,184 cells. The resistance exhibits clear grid convergence, except for the coarsest grid, which we consider to be outside the asymptotic range. The sinkage also shows grid convergence, however, for the pitch, convergence is not monotonic. In the following, the low-fidelity and high-fidelity results are obtained on grids of 64,000 and 710,000 cells, respectively. These grids are indicated in the plot for the resistance in Figure 5. Typical CPU times on a Intel Core 2.80 GHz processor are 0.50 hours for the low-fidelity simulation and 4.5 hours for the high-fidelity simulation.

In this example, we consider the two input parameters $(\xi_1, \xi_2) = (u, x_{\text{CoG}})$. Figure 6 gives the Kriging prediction $E(\mathbf{x}|\mathbf{y}_{\text{HF}}, \mathbf{y}_{\text{LF}})$, where in the different plots \mathbf{X} is the output for the resistance, sinkage and pitch, respectively. The fourth plot provides the design-of-experiment that was used to obtain these results. The high-fidelity and low-fidelity sampling plans have been generated as independent LHS-designs with maximin-criteria. The plots provide a clear impression of the response of the ship to changes in velocity and trim. For example, we see that the resistance increases with the velocity, as expected. However, the resistance also increases when the centre-of-gravity is moved to the bow or the stern, which apparently results in sub-optimal trim of the hull.

To illustrate the multi-fidelity approach, let us consider Figures 7–9, for the resistance, sinkage and pitch, respectively. In each figure, the top-left plot is the reference response. The top-right plot shows the same response,




Figure 5: Grid convergence for the DTMB 5415 model. The open symbols are considered to be outside the asymptotic range. Note that the pitch does not show clear grid convergence.

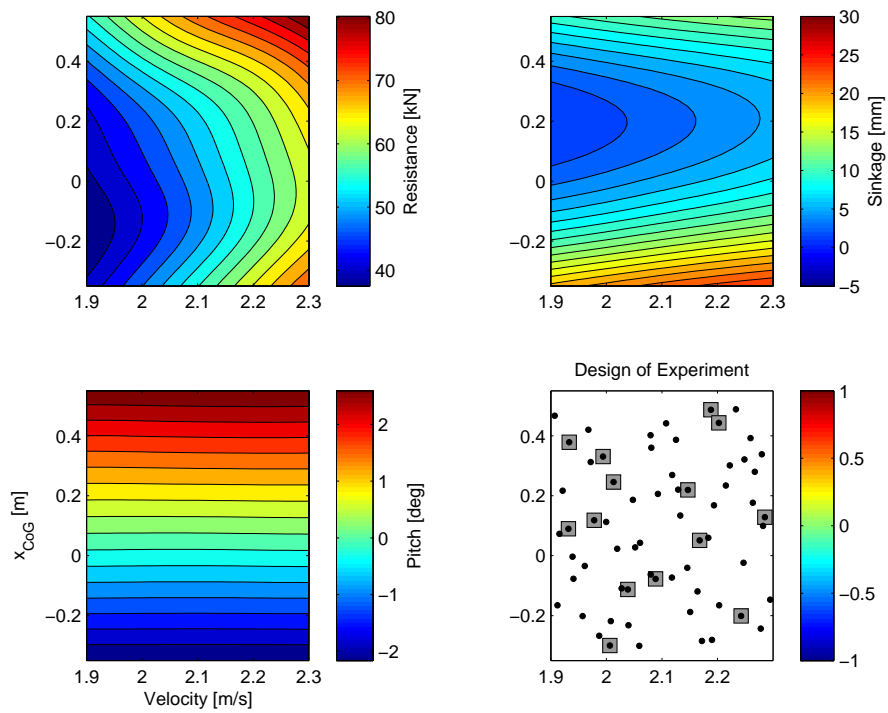


Figure 6: These plots show the response of the resistance, sinkage, and pitch to variations in velocity and centre of gravity. The bottom-right figure is the design-of-experiment, with grey squares for high-fidelity simulations and black dots for low-fidelity simulations.

but now based on only four high-fidelity CFD solves. Clearly, the number of observations is not enough to accurately resolve the shape of the response. When we add one high-fidelity solve, in the bottom-left plot, the response is still not accurate. However, at the same cost, we can add ten low-fidelity CFD solves instead. This results in the bottom-right response, which has a high level of accuracy.

The use of multi-fidelity data depends on a strong correlation between the low-fidelity and the high-fidelity results. For the DTMB 5415, this correlation is shown in Figure 10. Even if there is an overall offset, as in the case of the resistance, according to (12) the low-fidelity data can still improve results, as can be illustrated in a more formal study of the accuracy of the response for an increasing number of high-fidelity solves: In Figure 11, the red line shows the convergence of the root-mean-squared prediction error of the surrogate, with respect to the surrogate based on the full set of high-fidelity samples. The plotted error is the average error of 20 different LHS-designs. We should note that computing the error in this way is not perfect; in the ideal case we would consider a set of reference samples—which are not available, or use cross-validation—for which the current number of samples is too low. In Figure 11, increasing the number of high-fidelity solves decreases the prediction error. However, when for each high-fidelity solve we add two low-fidelity solves, the prediction error is reduced, at least for a low number of samples. Therefore, we define the lowfi-to-hifi ratio:

$$\frac{n_{\text{LF}}}{n_{\text{HF}}}. \quad (29)$$

We expect to find an optimum in this ratio; if we increase it to infinity, fully relying on low-fidelity solves, the prediction error is increased dramatically.




Figure 7: Four high-fidelity simulations are insufficient to get an accurate response for the resistance. Adding one more high-fidelity simulation doesn't improve the accuracy much. However, adding 10 low-fidelity simulations, at the same computational cost, improves the accuracy significantly.




Figure 8: As in Figure 7, four high-fidelity simulations are insufficient to get an accurate response for the sinkage. Adding one more high-fidelity simulation doesn't improve the accuracy much. However, adding 10 low-fidelity simulations, at the same computational cost, improves the accuracy significantly.




Figure 9: As in Figures 7 and 8, four high-fidelity simulations are insufficient to get an accurate response for the pitch. Adding one more high-fidelity simulation doesn't improve the accuracy much. However, adding 10 low-fidelity simulations, at the same computational cost, improves the accuracy significantly.

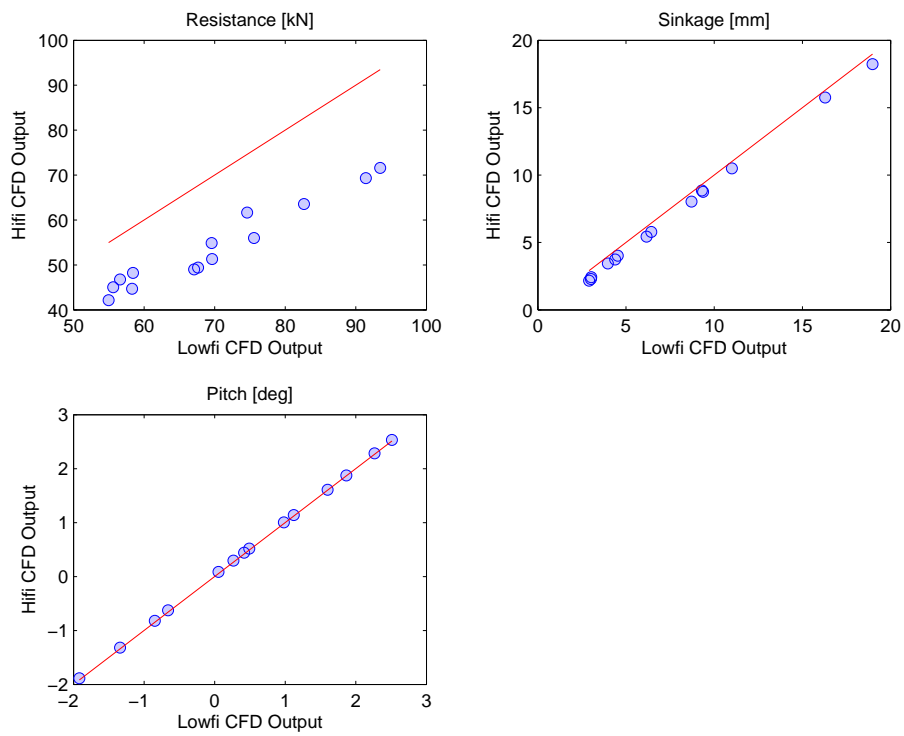


Figure 10: The low-fidelity results are highly correlated with the high-fidelity results. This is an important indication that we can use multi-fidelity Kriging.




Figure 11: Convergence of the response of the resistance, for an increasing cost. The prediction error is given for different lowfi-to-hifi ratios. For a low number of high-fidelity simulations, the low-fidelity simulations reduce the prediction error significantly.

To find this optimum, we consider the number of ‘solves-to-target’. This is the number of (high-fidelity cost-equivalent) solves that is required to reach a given target in the prediction error. Choosing a target accuracy of 4 N, which is roughly 5%, results in Figure 12: The dashed line shows the required solves-to-target without using low-fidelity solves. When we increase the lowfi-to-hifi ratio, the number of required solves-to-target drops. In this case, the minimum is around a ratio of 3. As the line rises steeper on the left than on the right, we choose to be on the safe side, and use a lowfi-to-hifi ratio of 5 in Section 3.

As a preview of the uncertainty quantification in Section 3, we assume that the centre of gravity x_{CoG} is normally distributed, with a standard deviation of 0.1 m, which represents a large uncertainty in the input parameter. Figure 13 shows the resulting uncertainty in the resistance. Clearly, even with this large standard deviation, the parameter uncertainty cannot explain the difference between the tank data and the CFD prediction. Instead, this difference is most likely caused by the spatial discretisation error, as is illustrated in Figure 5. As such, Figure 13 illustrates that before we consider parameter uncertainties, we should ensure that the discretisation error of the fine grid is sufficiently small.

Although presently we obtain low-fidelity results by computing on a coarser grid, there are other interesting possibilities. For example, one might use a RANS solver to obtain high-fidelity results and a potential flow solver to obtain the low-fidelity results.




Figure 12: The cost of reaching a certain target response accuracy, as a function of the low-to-high fidelity ratio. The dotted line represents the number of simulations that is required when we use only high-fidelity simulations.




Figure 13: Uncertainty propagation for the resistance of the DTMB 5415 does not explain the difference with the experimental results, because the grid convergence error is too large.

3. Results for Delft systematic yacht hull 25

In this section we focus on a more realistic UQ problem. We propagate the uncertainties in the velocity, mass and centre of gravity of a sailing yacht hull, and consider the resulting resistance, sinkage and pitch. In a CFD-based design process, experimental results would not be available. However, we have selected a case where towing tank data are available, and we compare the UQ results with these experimental data.

3.1. Problem definition

In 1973, as a joint research initiative, Delft University of Technology (Delft, The Netherlands) and the Massachusetts Institute of Technology (Boston, U.S.A.) started towing tank experiments on the Delft Systematic Yacht Hull Series (DSYHS) [18, 11]. The objective was to analyse the towing tank results of a series of hulls with systematic variations, in order to arrive at polynomial expressions for the hull performance and to make these polynomial expressions available for both design purposes and handicap racing [19].

Nowadays, the DSYHS consists of a range of 70 hulls, based on a number of parent models. We consider the second parent model, hull 25, which was designed by Cees van Tongeren of Van der Stadt and Partners (Zaandam, The Netherlands) [19]. The lines plan of hull 25 is shown in Figure 14, while the hydrostatics are given in Table 2.

The DSYHS models were tested in the number 1 towing tank at the Ship Hydromechanics Laboratory of Delft University of Technology. This tank has a length of 142 m, a width of 4.22 m and a water depth of 2.50 m. The tank




Figure 14: Lines plan of the DSYHS hull 25 towing tank model, showing sections, buttocks, waterlines and diagonal cuts.

Length on Waterline	2.00	m
Breadth	0.50	m
Draft	0.0928	m
Displacement	0.03698	m ³
Wetted Surface	0.671	m ²

Table 2: Hydrostatics of the DSHYS hull 25.

has a wave generator, but this was not used in the experiments presently considered. During the experiments, the hull is towed at a fixed speed while it is free to sink and pitch. Depending on the velocity, the Reynolds number is between $Re = 1 \times 10^6$ and 6×10^6 , however, the flow can be considered to be in the turbulent regime because turbulator strips were attached to the hull during all tests.

3.2. Input uncertainties

The Ship Hydromechanics Laboratory does not provide experimental uncertainties. We use two approaches to determine the uncertainties: (1) tolerances and indications are provided by the International Towing Tank Committee (ITTC) and (2) uncertainties estimated from well-documented experimental data of the DTMB 5415.

Based on expert opinion, we identify the three dominant uncertain inputs: blockage, mass and centre of gravity. These input uncertainties are the parameters to be changed during the uncertainty propagation. The blockage c corrects for the blockage effect in the tank, effectively by setting the free-stream velocity:

$$u_{\text{CFD}} = c u_{\text{Exp}} \quad (30)$$

in the CFD solver slightly higher (under normal conditions $c > 1$) than the carriage velocity u_{Exp} in the experiment. We use the first term of the Schuster equation [15] to estimate a blockage of: $c \sim \mathcal{N}(1.0038, 0.0007^2)$. However, we also have to take into account the relative uncertainty in the carriage velocity, which is given as 0.001 and 0.002, by [15] and [14], respectively. Incorporating this uncertainty in the carriage velocity into the uncertainty in the blockage, we estimate:

$$c \sim \mathcal{N}(1.004, 0.002^2). \quad (31)$$

The expected mass can be computed from the displacement and, using the uncertainty from [14], we estimate:

$$m \sim \mathcal{N}(36.95, 0.07^2) \text{ kg}. \quad (32)$$

The expected centre of gravity can be found from the centre of displacement of the hull, however, there is no indication of the uncertainty. Because, according to expert opinion, the centre of gravity can be a significant source of uncertainty, we estimate:

$$x_{\text{CoG}} \sim \mathcal{N}(-0.96, 0.05^2) \text{ m.} \quad (33)$$

3.3. Output uncertainties

There are three outputs: the resistance R , sinkage Δz and pitch α . Here, we mainly rely on the maximum likelihood estimate of the measurement uncertainty in the DTMB 5415 results reported in [31]. For the resistance, we find a measurement uncertainty of 0.6 N for the DTMB 5415 and of 0.2 N from [14]. We assume that, for a smaller ship, the uncertainty is closer to the ITTC tolerance, and estimate an uncertainty in the resistance of $\sigma_R = 0.2$ N. For the sinkage and pitch, we find an uncertainty of $\sigma_{\Delta z} = 0.13$ mm and $\sigma_\alpha = 0.05$ degrees, respectively, for the DTMB 5415, which we consider to be a good estimate for the present hull. It should be noted that the output uncertainties concern the measurement uncertainty in the validation data, and do not have an effect on the uncertainty propagation itself.

3.4. RANS simulations

The dedicated hydrodynamics solver FINE/Marine (Numeca, Brussels, Belgium) uses the ISIS code to solve the incompressible Reynolds Averaged Navier Stokes (RANS) equations [7, 37, 13]. The solver applies a finite volume spatial discretisation. An interface capturing approach is used to model the free-surface, the water-to-air interface is recovered from the volume frac-

tion [32]. We use a Menter Shear Stress Transport (SST) $k-\omega$ turbulence model [28].

We apply a zero slip boundary condition on the hull (using a wall-function), a pressure boundary condition on the ceiling and floor of the domain, a symmetry boundary condition on the symmetry plane and a prescribed velocity on the far-field boundaries.

The low-fidelity and high-fidelity grids are shown in Figure 15. Note the grid refinement near the hull, in the wake area and near the free surface. Typical CPU times on a Intel Core 2.80 GHz processor are 0.24 hours for the low-fidelity simulation and 2.2 hours for the high-fidelity simulation.




Figure 15: Grids with 42,000 (left) and 380,000 (right) cells, which are used for the low and high-fidelity simulations, respectively. Both grids show refinement near the hull and near the free-surface.

The iterative convergence of the solver is shown in Figure 16, which is used to verify that the results are fully converged for the number of specified iterations, as is necessary for this solver. The values for resistance, sinkage and pitch are averaged over the last 50 iterations. Grid convergence is shown in Figure 17, where the discretisation errors have been computed with respect to a reference grid of 1,256 cells. The coarsest grid is considered to be

outside the asymptotic range. The resistance shows an order of convergence of $p = 2.2$, which is close to the expected $p = 2$. The sinkage and pitch also show clear grid convergence, although the order of convergence is different. This might be caused by the change in geometry for the coarser grids.




Figure 16: To verify that the number of iterations is sufficient, we plot the iterative convergence. The CFD output is taken from the average of the last 50 iterations.

Figure 18 shows some typical solutions over a range of speeds. For a speed of $u_{\text{CFD}} = 0.990 \text{ m s}^{-1}$, which corresponds to a Froude number of $\text{Fn} = 0.22$, we see only minor waves. For a speed of $u_{\text{CFD}} = 1.585 \text{ m s}^{-1}$ ($\text{Fn} = 0.36$) we see a distinct wave pattern develop, as well as a significant bow wave. For a

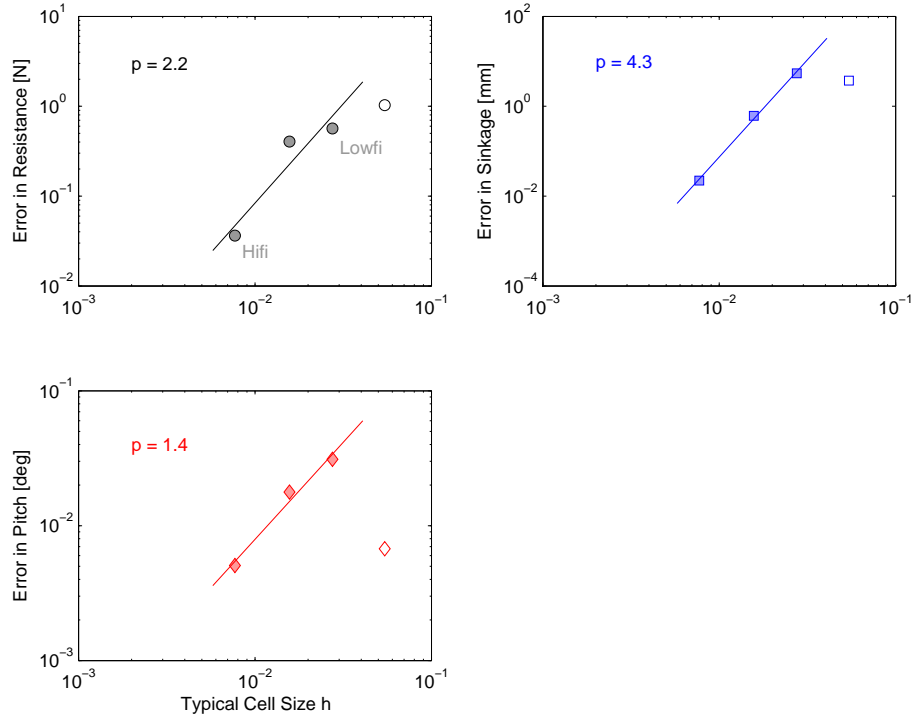


Figure 17: For further verification we plot the grid convergence. The open symbols, for the coarsest grid, are considered to be outside the asymptotic range.

speed of $u_{\text{CFD}} = 2.179 \text{ m s}^{-1}$ ($\text{Fn} = 0.49$) we are just below geometric hull speed, and the wavelength is close the hull length. Finally, at $u_{\text{CFD}} = 2.879 \text{ m s}^{-1}$ ($\text{Fn} = 0.65$) the hull is planing: The bow wave runs down the entire front half of the hull, the pitch of the hull has changed significantly and the wavelength is larger than the length of the ship.




Figure 18: Typical solutions for the free-surface, for increasing speed. The plots on the left show the wave pattern from the top (color scales differ for different speeds), while the plots on the right show the free-surface level on the hull.

3.5. Uncertainty quantification

The (u, m, x_{CoG}) space is sampled randomly, using a maximin-based optimal design. The design-of-experiment contains 24 high-fidelity and 120 low-fidelity solves. We assume this is a sufficient number of solves for an accurate surrogate, at present there is no a priori estimate of the number of required solves. The raw CFD output for the resistance, sinkage and pitch is shown in Figure 19. Note that, because the blockage only increases the velocity with a constant, we can sample over the full experimental velocity range and at the same time use the resulting response to propagate the uncertainty in the blockage.

Although the raw output still makes a random impression, Figure 20 shows that the outputs on the low-fidelity and high-fidelity outputs are highly correlated, as is required for the multi-fidelity approach. Figure 20 differs from Figure 10, the correlation plot for the DTMB 5415: for the DSYHS hull 25 the sinkage shows an offset, instead of the resistance. This is probably due to the difference in geometry: the DTMB 5415 has rectangular sections as well as a dome at the bow.

Figure 21 shows the result of the uncertainty propagation. After computing the multi-fidelity Kriging predictions, we apply (27–28) to exploit the multi-fidelity predictions and propagate the input uncertainties. The continuous lines show the resulting expected values for the outputs, while the shaded areas indicate the standard deviation.

For each output quantity, the symbols indicate the experimental results; it should be noted that the size of the symbols is larger than the experimental uncertainty. The CFD results are a valid prediction of the experimental

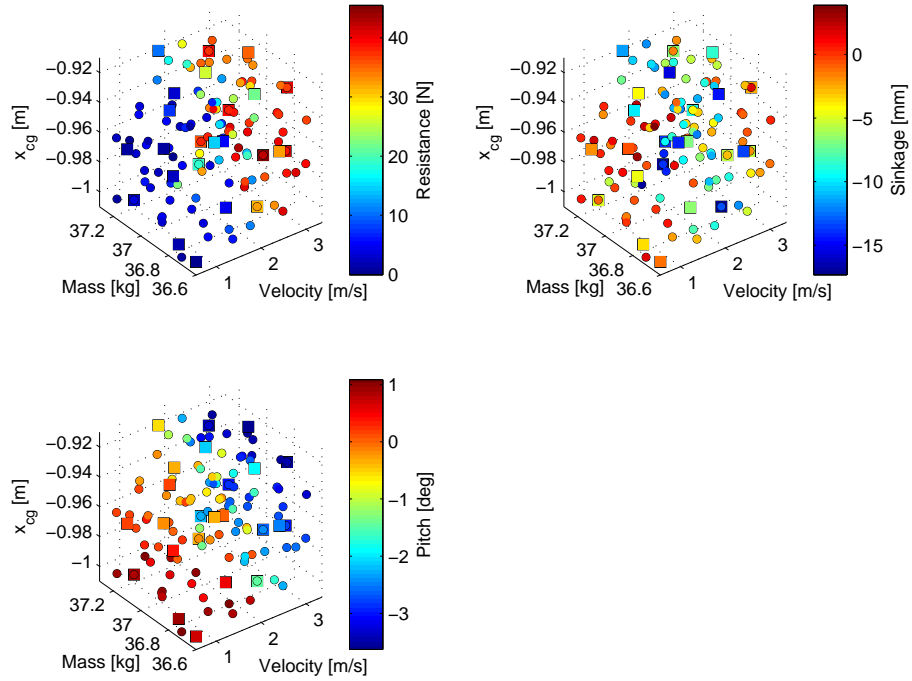


Figure 19: Raw CFD results on a random maximin sampling plan. The squares represent the high-fidelity simulations, while the smaller circles represent the low-fidelity simulations.

data for the lower velocities. However, the higher velocities show a discrepancy: The predicted resistance is lower than in the experiments, while the predicted sinkage is higher than in the experiments. Possible reasons for the discrepancies are: 1) for $F_n > 0.5$ the hull starts planing, 2) for $F_n > 0.5$ we might see ventilation and 3) the flow shows separation at the transom. All three phenomena are known to be hard to predict for a RANS solver. It is interesting to note that the uncertainty in the prediction of the resistance is small for increases exactly in this area.




Figure 20: In order to use multi-fidelity Kriging, it is important that the output of the low-fidelity and the high-fidelity simulations is sufficiently correlated.

3.6. Calibration

To emphasise the flexibility of the multi-fidelity Kriging response, we use it to calibrate the uncertain inputs. We do this only to illustrate that, once the response has been obtained, we can exploit it in many other ways. The objective of the calibration is to infer which values of blockage, mass and centre of gravity we should use for the CFD code to arrive at the best prediction of the experimental results.

Because the CFD code does not predict the experimental values that well




Figure 21: The result of the uncertainty quantification analysis, together with experimental results. The CFD results give a good prediction of the experimental data, except for $Fn > 0.5$, where phenomena like planing and flow separation are difficult to predict.

for the higher velocities, we will only include the experimental results for the lowest eight velocities (corresponding to a Froude number of $0.10 < Fn < 0.45$), and define this set of experimental data as $\mathbf{d} = (\mathbf{R}, \Delta\mathbf{z}, \alpha)$. We assume a uniform prior distribution $p(\boldsymbol{\xi})$, with $\boldsymbol{\xi} = (c, m, x_{CoG})$ the inputs blockage, mass and centre of gravity. We then evaluate Bayes' Rule to obtain:

$$p(\boldsymbol{\xi}|\mathbf{d}) = \frac{p(\mathbf{d}|\boldsymbol{\xi}) p(\boldsymbol{\xi})}{p(\mathbf{d})}. \quad (34)$$

Assuming that the experimental outputs are uncorrelated, we have the like-

likelihood:

$$\begin{aligned} p(\mathbf{d}|\boldsymbol{\xi}) &= p(\mathbf{R}, \mathbf{h}, \boldsymbol{\alpha}|\boldsymbol{\xi}) \\ &= p(\mathbf{R}|\boldsymbol{\xi}) p(\mathbf{h}|\boldsymbol{\xi}) p(\boldsymbol{\alpha}|\boldsymbol{\xi}). \end{aligned} \quad (35)$$

We take:

$$p(\mathbf{R}|\boldsymbol{\xi}) = \exp\left(-\frac{1}{2} \sum_{i=1}^8 \frac{[E(R|\mathbf{y})_{\boldsymbol{\xi}, u_{\text{Exp},i}} - R_i]^2}{\sigma_R^2}\right), \quad (36)$$

where $E(R|\mathbf{y})_{\boldsymbol{\xi}, u_{\text{Exp},i}}$ denotes the multi-fidelity kriging prediction of R , for a given input $\boldsymbol{\xi}$, at the carriage velocity $u_{\text{Exp},i}$ that corresponds to the measurement R_i . The distributions for $p(\Delta\mathbf{z}|\boldsymbol{\xi})$ and $p(\boldsymbol{\alpha}|\boldsymbol{\xi})$ are similar to (36).

As a result, we find the three-dimensional posterior probability density function $p(\boldsymbol{\xi}|\mathbf{d})$ for the input parameters. The marginals of this posterior PDF are shown in Figure 22. All calibrated inputs are close to the values provided, as indicated by the black lines. The setting for the mass and centre of gravity appear to be slightly correlated: If we increase one, we should probably decrease the other to obtain accurate results.

4. Conclusions

We have presented the multi-fidelity Kriging predictor equations, as well as a sparse grid approach to propagate parameter uncertainties through a multi-fidelity response. After illustrating the multi-fidelity method for a two-dimensional example of the DTMB 5415, we have made a full uncertainty quantification analysis of the DSYHS hull 25, comparing the CFD results with experimental results. Finally, to illustrate the possibilities of exploitation of the multi-fidelity response, we have calibrated the input parameters for the $0.10 < \text{Fn} < 0.45$ velocity range.




Figure 22: Input parameters after calibration. The black lines and black crosses indicate the uncalibrated mean values.

The multi-fidelity Kriging approach provides an accurate response for both the DTMB 5415 and DSYHS hull 25 test case. This approach can efficiently reduce the overall cost of an uncertainty quantification analysis. The recommend lowfi-to-hifi ratio is 5, however, further investigation is required.

For the lower velocities, corresponding to $Fn < 0.50$, the present uncertainty propagation results in a good prediction of the experimental results. However, for higher velocities the resistance is under-predicted, while the sinkage is over-predicted. Possible reasons are: 1) the hull is planing 2) there

might be ventilation 3) there is flow separation at the transom.

An important question remains: How many low-fidelity and high-fidelity samples are required to obtain an accurate surrogate, and how should we distribute these samples over our parameter space? As this depends on the shape of the low-fidelity and high-fidelity responses, it is unlikely that one would be able to provide a priori estimates. Instead, it might be possible to use an adaptive algorithm, which positions new samples strategically until the required accuracy has been achieved. This is an important topic for future research.

The multi-fidelity Kriging approach is a promising method to reduce the computational cost of UQ analyses. Future research should address the optimal lowfi-to-hifi ratio, include geometric uncertainties and investigate the scaling when the number of uncertain parameters increases.

References

- [1] Cliffe, K. A., Giles, M. B., Scheichl, R., and Teckentrup, A. L. (2011). Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Computing and Visualization in Science*, 14:3–15.
- [2] Cressie, N. (1990). The origins of Kriging. *Mathematical Geology*, 22(3):239–252.
- [3] Cressie, N. (1993). *Statistics for spatial data*. Wiley.
- [4] Cressie, N. and Wikle, C. (2011). *Statistics for spatio-temporal data*. Wiley.

- [5] de Baar, J., Dwight, R., and Bijl, H. (2013). Speeding up kriging through fast estimation of the hyperparameters in the frequency-domain. *Computers & Geosciences*, 54:99 – 106.
- [6] Diez, M., He, W., Campana, E., and Stern, F. (2014). Uncertainty quantification of delft catamaran resistance, sinkage and trim for variable froude number and geometry using metamodels, quadrature and Karhunen-Loeve expansion. *Journal of Marine Science and Technology*, (19):143–169.
- [7] Duvigneau, R., Visonneau, M., and Deng, G. (2003). On the role played by turbulence closures in hull shape optimization at model and full scale. *Journal of Marine Science and Technology*, (8):11–25.
- [8] Forrester, A., Sobester, A., and Keane, A. (2008). *Engineering Design via Surrogate Modelling: A Practical Guide*. Wiley.
- [9] Forrester, A., Sobester, A., and Keane, J. (2007). Multi-fidelity optimization via surrogate modelling. *Proceedings of the Royal Society A*, 463:3251–3269.
- [10] Gandin, L. (1965). *Objective analysis of meteorological fields: Gidrometeorologicheskoe Izdatel'stvo (GIMIZ), Leningrad*. Translated by Israel Program for Scientific Translations, Jerusalem.
- [11] Gerritsma, J., Onnink, R., and Versluis, A. (1981). Geometry, resistance and stability of the Delft Systematic Yacht Hull Series. *7-th HISWA Symposium, 1981, Amsterdam*.
- [12] Haftka, R. (1991). Combining global and local approximations. *AIAA*, 29(9):15231525.

- [13] Hochkirch, K. and Mallol, B. (2013). On the importance of full-scale CFD simulations for ships. *12th International Conference on Computer Applications and Information Technology in the Maritime Industries (COMPIT 2013), Cortona, Italy, April 15th -17th.*
- [14] ITTC (2002). Recommended procedures: Resistance uncertainty analysis, example for resistance test. *International Towing Tank Committee Report 7.5-02-02-02.*
- [15] ITTC (2011). Recommended procedures and guidelines: Resistance test. *International Towing Tank Committee Report 7.5-02-02-01.*
- [16] Kennedy, M. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464.
- [17] Kennedy, M. C. and O’Hagan, A. (2000). Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87:1–13.
- [18] Kerwin, J. (1975). A velocity prediction program for ocean racing yachts. *Report 78-11, Department of Ocean Engineering, Massachusetts Institute of Technology.*
- [19] Keuning, J. and Katgert, M. (2008). A bare hull resistance prediction method derived from the results of the Delft Systematic Yacht Hull Series extended to higher speeds. *International Conference on Innovation in High Performance Sailing Yachts, Lorient, France.*

- [20] Kitanidis, P. K. (1986). Parameter uncertainty in estimation of spatial functions: Bayesian analysis. *Water Resources Research*, 22(4):499–507.
- [21] Kitanidis, P. K. and Lane, R. W. (1985). Maximum likelihood parameter estimation of hydrologic spatial processes by the Gauss-Newton method. *Journal of Hydrology*, 79(1-2):53 – 71.
- [22] Klimke, A. and Wohlmuth, B. (2005). Algorithm 847: spinterp: Piecewise multilinear hierarchical sparse grid interpolation in MATLAB. *ACM Transactions on Mathematical Software*, 31(4).
- [23] Kuya, Y., Takeda, K., Zhang, X., and Forrester, A. I. J. (2011). Multifidelity surrogate modeling of experimental and computational aerodynamic data sets. *AIAA Journal*, 49(2):289–298.
- [24] Mardia, K. V. (1989). Maximum likelihood estimation for spatial models. *Proceedings from the symposium on Spatial Statistics: Past, Present, and Future*, pages 203–253.
- [25] Mardia, K. V. and Marshall, R. J. (1984). Maximum likelihood estimation of models for residual covariance in spatial regression. *Biometrika*, 71(1):135–146.
- [26] Mason, B. H., Haftka, R. T., Johnson, E. R., and Farley, G. L. (1998). Variable complexity design of composite fuselage frames by response surface techniques. *Thin-Walled Structures*, 32(4):235–261.
- [27] Matheron, G. (1963). Principles of Geostatistics. *Economic Geology*, 58:1246–1266.

- [28] Menter, F. (1994). Two-equation eddy-viscosity turbulence models for engineering applications. *AIAA Journal*, 32(8):1598–1605.
- [29] Oden, T., Moser, R., and Ghattas, O. (2010a). Computer predictions with quantified uncertainty. *SIAM News*, 43(9).
- [30] Oden, T., Moser, R., and Ghattas, O. (2010b). Computer predictions with quantified uncertainty. *SIAM News*, 43(10).
- [31] Olivieri, A., Pistani, F., Avanzini, A., Stern, F., and Penna, R. (2001). Towing tank experiments of resistance, sinkage and trim, boundary layer, wake and free surface flow around a naval combatant INSEAN 2340 model. *IIHR Technical Report 421*.
- [32] Queutey, P. and Visonneau, M. (2007). An interface capturing method for free-surface hydrodynamic flows. *Computers & Fluids*, 36:1481–1510.
- [33] Stern, F., Wang, Z., Yang, J., Sadat-Hosseini, H., Mousaviraad, M., Bhushan, S., Diez, M., Yoon, S.-H., Wu, P.-C., Yeon, S., Dogan, T., Kim, D.-H., Volpi, S., Conger, M., Michael, T., Xing, T., Thodal, R., and Grenstedt, J. (2014). Recent progress in CFD for naval architecture and ocean engineering (keynote speaker). *Proceedings of the 11th International Conference on Hydrodynamics, 19–24 October, Singapore*.
- [34] Stern, F., Wilson, R., and Shao, J. (2006). Quantitative V&V of CFD simulations and certification of CFD codes. *International Journal for Numerical Methods in Fluids*, 50:1335–1355.
- [35] Toal, D. (2015). Some considerations regarding the use of multi-fidelity

- Kriging in the construction of surrogate models. *Structural and Multidisciplinary Optimization*, 51(6):1223–1245.
- [36] Viana, F. A., Simpson, T. W., Balabanov, V., and Toropov, V. (2014). Metamodeling in multidisciplinary design optimization: How far have we really come? *AIAA Journal*, 52(4):670–690.
- [37] Visonneau, M., Queutey, P., and Deng, G. (2006). Model and full-scale free-surface viscous flows around fully-appended ships. *European Conference on Computational Fluid Dynamics, ECCOMAS CFD*.
- [38] Volpi, S., Diez, M., Gaul, N., Song, H., Iemma, U., Choi, K., Campana, E., and Stern, F. (2015). Development and validation of a dynamic meta-model based on stochastic radial basis functions and uncertainty quantification. *Structural and Multidisciplinary Optimization*, 51(2):347–368.
- [39] Wikle, C. K. and Berliner, L. M. (2007). A Bayesian tutorial for data assimilation. *Physica D: Nonlinear Phenomena*, 230(1-2):1 – 16.
- [40] Xing, T. and Stern, F. (2010). Factors of safety for Richardson extrapolation. *Journal of Fluids Engineering*, 132:061403–1–061403–13.