Computers and Fluids 123 (2015) 185–201

Contents lists available at ScienceDirect

Computers and Fluids

journal homepage: www.elsevier.com/locate/compfluid

Uncertainty quantification for a sailing yacht hull, using multi-fidelity kriging

Jouke de Baar [a,*], Stephen Roberts [a], Richard Dwight [b], Benoit Mallol [c]

[a] Australian National University, Australia
[b] Delft University of Technology, The Netherlands
[c] Numeca, Brussels, Belgium

ARTICLE INFO

Article history:
Received 26 May 2015
Revised 5 August 2015
Accepted 5 October 2015
Available online 22 October 2015

Keywords:
Uncertainty quantification
Multi-fidelity
Kriging
RANS
Free-surface

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.

1. Introduction

Initial error analysis for marine applications focused on grid convergence and time discretisation errors [33,34,40]. 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 quantification (UQ) [6,33]. 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 [9,17]. 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 [9,16].

An extensive overview of multi-fidelity methods is given in [36], we provide a brief discussion. In an early application, Haftka [12] presents a multiplicative multi-fidelity analysis of a clamped beam. [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. Haftka [23] uses 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, Mardia and Marshall [35] present 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

* Corresponding author. Tel.: +61 421330253.
E-mail address: j.h.s.debaar@gmail.com (J. de Baar).

http://dx.doi.org/10.1016/j.compfluid.2015.10.004
0045-7930/

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:

$$\boldsymbol{X} \sim \mathcal{N}(\mu, P), \tag{1}$$

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

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

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 $\boldsymbol{X}$. Now, by applying Bayes' rule:

$$p(\boldsymbol{x}|\boldsymbol{y}) = \frac{p(\boldsymbol{y}|\boldsymbol{x})\ p(\boldsymbol{x})}{p(\boldsymbol{y})}, \tag{3}$$

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

$$\mathrm{E}(\boldsymbol{X}|\boldsymbol{y}) = \mu + K(\boldsymbol{y} - \mu), \tag{4}$$

and:

$$\mathrm{cov}(\boldsymbol{X}|\boldsymbol{y}) = (I - KH)P, \tag{5}$$

with the gain matrix:

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

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:

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

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:

$$R = \epsilon^2 I, \tag{8}$$

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 $\boldsymbol{y}$ instead, using a maximum likelihood estimate [20,21,24,25]. This is equivalent to minimising the log-likelihood:

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

with:

$$A = (R + HPH^T), \tag{10}$$

where P depends on θ, as given in (7), and where $|.|$ 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, Fig. 1 shows the process:

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

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 [9,17]. Here we assume that we have a set of high-fidelity CFD simulations $\boldsymbol{y}_{\mathrm{HF}}$, which are expensive to evaluate and provide an accurate output, as well as a set of low-fidelity simulations $\boldsymbol{y}_{\mathrm{LF}}$, 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 [9,17], let us assume that the high-fidelity process $\boldsymbol{X}_{\mathrm{HF}}$ is related to the low-fidelity process $\boldsymbol{X}_{\mathrm{LF}}$ as:

$$\boldsymbol{X}_{\mathrm{HF}} = \rho\ \boldsymbol{X}_{\mathrm{LF}} + \boldsymbol{X}_{\mathrm{D}}, \tag{12}$$

where ρ is a regression parameter. We now have a new process $\boldsymbol{X}_{\mathrm{D}}$ that represents the difference between the low-fidelity and the high-fidelity output. This difference process $\boldsymbol{X}_{\mathrm{D}}$ 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:

$$y_{\mathrm{LF}} = y_{\mathrm{exact}} + Ch_{\mathrm{LF}}^p + \mathrm{H.O.T.}$$
$$y_{\mathrm{HF}} = y_{\mathrm{exact}} + Ch_{\mathrm{HF}}^p + \mathrm{H.O.T.},$$

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

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

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 $\boldsymbol{x}_{\mathrm{D}}$ 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 $\boldsymbol{X}_{\mathrm{D}}$ 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:

$$\boldsymbol{X}_{\mathrm{MF}} = \begin{pmatrix} \boldsymbol{X}_{\mathrm{LF}} \\ \boldsymbol{X}_{\mathrm{D}} \end{pmatrix}$$

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}.