

Multi-fidelity Surrogate Assisted Design Optimisation of an Airfoil under Uncertainty using Far-Field Drag Approximation^{*}

Elisa Morales^{1,4}[0000-0002-0267-5322], Péter Zénó Korondi^{2,3}[0000-0002-9935-1059],
Domenico Quagliarella¹[0000-0003-3828-4937], Renato Tognaccini⁴[0000-0003-4528-5540],
Mariapia Marchi²[0000-0002-7107-9073], Lucia Parussini³[0000-0003-1790-9731], and
Carlo Poloni^{2,3}[0000-0001-8077-0563]

¹ Italian Aerospace Research Centre, Capua, Italy
{e.moralestirado, d.quagliarella}@cira.it

² ESTECO, Trieste, Italy
{korondi, marchi, poloni}@esteco.com

³ Department of Engineering and Architecture, University of Trieste
lparussini@units.it

⁴ Università degli Studi di Napoli “Federico II”, Naples, Italy
renato.tognaccini@unina.it

Abstract. Uncertainty-based optimisation techniques provide optimal airfoil designs that are less vulnerable to the presence of uncertainty in the operational conditions (i.e., Mach number, angle-of-attack, etc.) at which an airfoil is functioning. These uncertainty-based techniques typically require numerous function evaluations to accurately calculate the statistical measure of the quantity of interest. To render the computational burden down, the design optimisation of the airfoil is performed by a multi-fidelity surrogate-based technique. The high-fidelity aerodynamic performance is calculated with a compressible RANS solver using a fine grid. At the low-fidelity level a coarser grid is used. To obtain accurate drag predictions despite the lower grid resolution the so-called far-field drag approximation is employed.

Keywords: robust design optimisation · Gaussian processes · conditional value-at-risk · far-field drag · multi-fidelity learning

1 Introduction

The aerodynamics performance of real-world applications is inherently uncertain due to manufacturing errors, uncertain environmental conditions [18, 19] and other physical phenomena like icing [4]. Therefore, uncertainties must be accounted already during the aerodynamic design of airfoils.

Uncertainty-based optimisation techniques provide optimal airfoil designs that are less vulnerable to the presence of uncertainty in the operational conditions (i.e., Mach

^{*} This research has been developed with the partial support of the H2020 MCSA ITN UTOPIAE grant agreement number 722734.

number, angle-of-attack, etc.) at which the airfoil is functioning. In such optimisation techniques, the Quantity of Interest (QoI) is a statistical measure instead of a deterministic value. The accurate calculation of a statistical measure requires numerous function evaluations, which increases the computational demand significantly. In particular, the Computational Fluid Dynamic (CFD) calculations burden the computational budget as the QoIs (a statistical measure of lift, drag, or moment coefficients) are computed by solving the Reynolds-averaged Navier-Stokes (RANS) equations numerically.

In order to reduce the computational time, a multi-fidelity surrogate assisted method is adopted here. Surrogates are data-based mathematical models constructed using only a few expensive function evaluations. With their help, the aerodynamic performance of an airfoil can be predicted at a low computational cost. The accuracy of a surrogate highly depends on the size of its training data. Therefore, when the function evaluations are truly expensive, for example, a CFD simulation with very fine grid, the training data can be complemented by function evaluations of lower fidelity. The information coming from various fidelity levels can be fused together with multi-fidelity Gaussian process regression (MF-GPR). This technique was introduced by [9].

In this work, the drag coefficient (c_d) of the MH 114 [2] propeller airfoil is minimised by a multi-fidelity surrogate assisted optimisation technique. The open-source fluid-dynamic solver SU2 [5] is used for calculating the c_d . SU2 solves the compressible RANS equations numerically and calculates c_d by integrating the stress over the body surface with the so-called near-field method.

The drag coefficient c_d can have different levels of fidelity by using diverse grid refinements. A calculation with a fine grid provides a high-fidelity c_d prediction. However, fine meshes are very demanding from a computational point of view. Coarse grids are computationally cheap, but they introduce a higher proportion of spurious drag. This numerically introduced drag stems from the truncation error of the used numerical methods and the artificial dissipation of solving the RANS equations with a coarse grid. The artificial dissipation is added in the numerical schemes to boost the convergence of the flow and to stabilise the scheme. Hence, the prediction of the near-field c_d deteriorates.

Nevertheless, there are far-field methods for the estimation of the drag force that allow the c_d prediction with a level of accuracy similar to a fine grid by identifying the spurious drag sources. A review of all these methods is given in [6]. In this work, the formulation described in [13] has been implemented.

The prediction of the drag coefficient using the far-field method will be used for the low-fidelity level on a coarse grid. This procedure will allow a better estimation of the drag coefficient with respect to the near-field value computed on the same grid, thus resulting in an increased accuracy while preserving the computational cost. In addition, the near-field c_d estimation obtained with a fine grid will be used for the high-fidelity level.

This paper is organised as follows. Section 2 gives a brief overview of MF-GPR. The aerodynamic computational chain is detailed in Section 3. The far-field approach for the drag estimation is explained in Section 4. Section 5 introduces a deterministic airfoil design problem. The airfoil design problem under uncertainty is given in Section 6. Specifically, the drag coefficient of a propeller blade airfoil will be minimised under geometrical and environmental constraints. The uncertainty will be introduced

on the angle-of-attack modelled by a four-parameter beta distribution. We propose a multi-fidelity surrogate assisted optimisation pipeline in Section 7. The results of the deterministic and probabilistic optimisations are discussed in Section 8. The interpretation of the results is concluded in Section 9.

2 Multi-fidelity Gaussian Process Regression

The drag coefficients obtained with the far-field approximation and SU2 are fused together into a single surrogate by multi-fidelity Gaussian process regression (MF-GPR). The recursive formulation is adapted here as proposed by [11]:

$$\tilde{f}_1(\mathbf{x}) = \mathbf{h}_1^T(\mathbf{x})\boldsymbol{\beta}_1 + \tilde{\delta}_1(\mathbf{x}), \quad (1a)$$

$$\tilde{f}_2(\mathbf{x}) = \rho(\mathbf{x})\tilde{f}_1(\mathbf{x}) + \mathbf{h}_2^T(\mathbf{x})\boldsymbol{\beta}_2 + \tilde{\delta}_2(\mathbf{x}), \quad (1b)$$

$$\rho(\mathbf{x}) = \mathbf{g}^T(\mathbf{x})\boldsymbol{\beta}_\rho, \quad (1c)$$

where indices 1 and 2 denote the low and high-fidelity levels, respectively. The mean trend of the fidelity level is formulated as a least-squares regression $\mathbf{h}_i(\mathbf{x})\boldsymbol{\beta}_i$ with the vector of regression functions $\mathbf{h}_i(\mathbf{x})$ and the vector of regression coefficients $\boldsymbol{\beta}_i$. The local variations of the model are modelled as zero-mean Gaussian distributions with σ_i^2 variance and incorporated into $\tilde{\delta}_i(\mathbf{x}) \sim \mathcal{N}(0, \sigma_i^2)$. This recursive formulation, first, trains a standard GPR surrogate using the low-fidelity samples calculated by the far-field method. Then, the posterior of the low-fidelity GPR is combined together with the high-fidelity observations of SU2 by training an additional GPR. This recursive formulation avoids the need to construct a large covariance matrix containing the low- and high-fidelity designs as in [9]. Even if the training cost of the surrogate is negligible compared to the aerodynamic design evaluation, the reduced covariance size is advantageous as the model is frequently re-trained during the optimisation process.

3 Aerodynamic computational chain

When aerodynamic shape design problems are faced, it is crucial to have a self-operating aerodynamic computational chain. It takes as input the design variables given by the optimiser and generates the candidate (in this case a single component airfoil) to be evaluated, builds the computational mesh, and runs the computational fluid dynamic flow solver. Finally, the obtained performance of the candidate is provided to the optimiser. In the design problem studied here, the performance of the candidate is the airfoil drag coefficient. Moreover, for the low-fidelity runs, instead of providing directly the c_d given by the CFD solver, the drag coefficient is calculated with a far-field approach and later is provided to the optimiser. This pipeline is shown in Fig. 1, and explained in the following sub-sections. Note that the far-field formula is described in depth in Section 4.

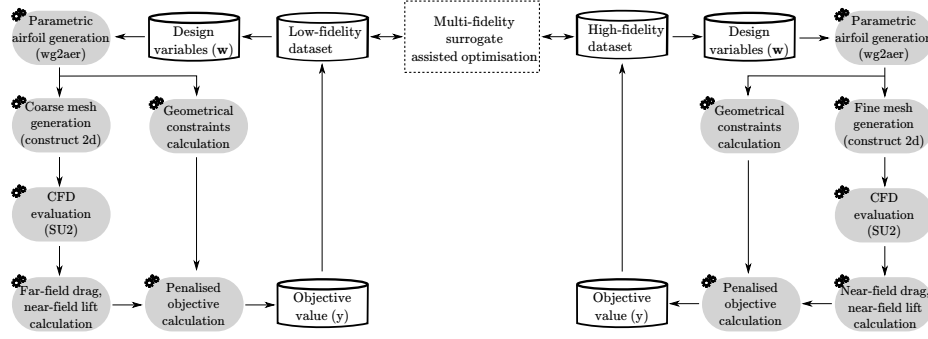


Fig. 1. Aerodynamic computational chain

Airfoil generation

The candidate airfoils are generated using *wg2aer*¹. The input for the program is a set of values of the design variables, and it modifies a specified starting airfoil (MH 114) accordingly to some modification functions. To generate the baseline airfoil, the design variables are set equal to 0. Specifically, the airfoil is parametrized as a linear combination of the initial geometry $(x_0(s), y_0(s))$, and the applied modification functions $y_i(s)$. Thus, the airfoil is described as:

$$y(s) = k \left(y_0(s) + \sum_{i=1}^n w_i y_i(s) \right), \quad (2)$$

$$x(s) = x_0(s)$$

the airfoil shape is controlled by the design parameters w_i and k , the scale factor to fulfil the maximum thickness criterion.




Fig. 2. Modified airfoils example. Baseline airfoil (—)

Specifically, ten design parameters are considered. This will lead to eight optimisation design variables. The first and second design parameters describe a thickness mode

¹ Program developed at the Italian Aerospace Research Centre (CIRA)

so that they have the same value but opposite sign ($w_2 = -w_1$). In addition, the third and fourth design parameters represent a camber mode, hence both are equal ($w_4 = w_3$). The range for the design parameters are:

$$\begin{aligned} w_1, w_2 &\in [-2, 2], & w_3, w_4 &\in [-2, 2], \\ w_5, w_6, w_7, w_8 &\in [-1, 1], & w_9, w_{10} &\in [-0.2, 0.2] \end{aligned} \quad (3)$$

Regarding the modification functions, the first four are polynomials affecting the whole airfoil, while the rest are Hicks-Henne bump functions that have the location of the bump at different position of the airfoil chord. In Fig. 2, an example of airfoil modifications is provided.

Grid generation

Construct2D is an open-source grid generator designed to create 2D grids for CFD computations on airfoils [1]. The grids are generated in Plot3D format; however, the source code has been changed to provide also the grid in SU2 format. Given the coordinates of the modified airfoil, a C-type grid is generated. The number of cells of the mesh will depend on the level of fidelity that has to be run. The possible mesh sizes are provided in Table 1.

Table 1. Mesh size parameters for low- and high-fidelity simulations. (N_b : number of cells on the body surface, N_w : number of cells in the wake, N_j : number of cells in far-field direction, N_{total} : total number of cells)

	N_b	N_w	N_j	N_{total}
Low fidelity (LF)	96	48	48	16384
High fidelity (HF)	512	256	256	262144

The far-field is located at 500 airfoil chords. The possible computational meshes are depicted in Fig. 3




Fig. 3. Possible grids depending on the fidelity level.

CFD evaluation

The CFD solver used for the aerodynamic shape design optimisation problem is the open-source fluid-dynamic solver SU2 [5]. Particularly, the compressible Reynolds-averaged Navier-Stokes (RANS) equations are solved. The turbulence model used is the Spalart-Allmaras (SA) [17]. Furthermore, for the spatial integration, JST central scheme with artificial dissipation coupled with an implicit Euler method for the pseudo-time stepping is used.

4 Far-field Drag Coefficient Calculation

The far-field method implemented in this work was introduced in [13]. It allows the decomposition of the drag force in three components: wave, viscous, and spurious drag. Specifically, the method is based on entropy variations. The entropy drag is expressed as volume integral, which allows the decomposition of the drag into the above components. Hence, a proper selection of each region is needed. The entropy drag is defined as:

$$D_{\Delta s} = D_w + D_v + D_{sp} \quad (4)$$

where D_w , D_v , and D_{sp} are the wave, viscous, and spurious contributions, respectively.

D_{sp} is the drag source related to the entropy introduced by the truncation error and the artificial dissipation of the numerical schemes used by the Computational Fluid Dynamics flow solver. Hence, by identifying the D_{sp} contribution and subtracting it from Equation (4), a prediction of the drag coefficient, with an accuracy close to fine grids, is obtained on a coarser mesh. This will imply a considerable advantage while facing aerodynamic optimisation problems since the use of coarser grids allows a significant reduction of the required computational time. This advantage for optimisation has already been shown in [7, 12].

Furthermore, a test to verify this advantage for the design problem of the propeller blade airfoil, here studied, has been carried out. In particular, a viscous flow with working conditions $M_\infty = 0.2$, $Re_\infty = 4.97 \times 10^6$ and $c_l = 1.0$ is performed on the MH 114 airfoil. The compressible RANS equations are solved using the SU2 flow solver with the SA turbulence model [17]. Five C-type grids of an increasing number of cells are studied. The grid size is obtained by the square root of the ratio between the number of cells of the finest grid and the grid under evaluation ($h = \sqrt{N_{h=1}/N_i}$). The number of cells on the body surface (N_b), on the wake (N_w), in the far-field direction (N_j), and the total number of cells (N_{total}), as well as, the near-field value of the drag coefficient ($c_{d_{nf}}$), and the far-field value (c_{d_v}) are given in Table 2. Note that in this test case the only drag contribution is the viscous (D_v), and the drag values are expressed in drag counts ($1dc = 10^{-4}$).

Table 2. Mesh sizes and computed drag coefficients. Viscous test at $M_\infty = 0.2$, $Re_\infty = 4.97 \times 10^6$ and $c_l = 1.0$.

	N_b	N_w	N_j	N_{total}	$c_{d_{nf}}$ [dc]	c_{d_v} [dc]
$h = 10.7$	96	48	48	9216	165.7	115.0
$h = 8$	128	64	64	16384	137.1	107.8
$h = 4$	256	128	128	65536	120.7	106.5
$h = 2$	512	256	256	262144	113.9	104.2
$h = 1$	1024	512	512	1048576	110.0	102.6

In Fig. 4, the comparison between near-field and far-field drag coefficients versus the grid size is given. In addition, the pressure coefficient distribution on the body surface (c_p) at the different grid refinements is also plotted.

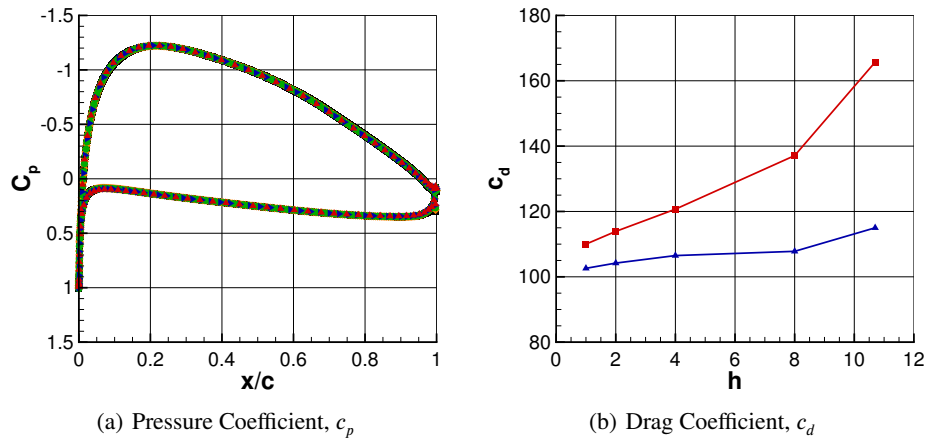


Fig. 4. MH 114 test at $M_\infty = 0.2$, $Re_\infty = 4.97 \times 10^6$ and $c_l = 1.0$. *Left:* Pressure coefficient distribution on the body surface at $h = 1$ (■), $h = 2$ (◆), $h = 4$ (●), $h = 8$ (►), and $h = 10.7$ (▲). *Right:* Near-field (—■—) and far-field (—►—) drag coefficients versus mesh size.

In Fig. 4 (a) it can be observed that the differences between the c_p on the body surface are barely visible; hence a good local accuracy of the solution is demonstrated also when the coarsest grid is used. However, Fig. 4 (b) shows how the near-field value of the drag coefficient converges as the grid is refined ($h \rightarrow 0$). The variation on c_d between the coarsest and finest mesh sizes is given by the spurious drag source introduced by the numerical method and the artificial dissipation. Contrarily, using the far-field analysis of the drag force, the spurious drag contribution is removed. Hence, a better estimation of the c_d is found. Thus, the drag coefficient value for the lower fidelity of the surrogate model is improved in accuracy while keeping the same computational time.

5 Deterministic Design Optimisation Problem

The shape optimisation design problem studies the minimisation of the drag coefficient (c_d) of a propeller blade airfoil subjected to geometric and aerodynamic constraints. The baseline design is the Martin Hepperle MH 114 airfoil for a propeller. The flow conditions are $M_\infty = 0.2$, $Re = 4.97 \times 10^6$, and $\alpha = 2^\circ$. The lift coefficient of the airfoil is required to be at least one ($c_l \geq 1$). Geometrical constraints are imposed for obtaining realistic shapes. The percentage thickness with respect to the airfoil chord ($t_\%$) is fixed to the value of the baseline. The Leading Edge Radius (LER) and the Trailing Edge Angle (TEA) are constrained by minimum values not to fall below their baseline values with more than 10%.

In mathematical terms, the deterministic optimisation example reads:

$$\left\{ \begin{array}{l} \min_{\mathbf{w}} c_d(\mathbf{w}) \\ \text{subject to:} \\ c_l(\mathbf{w}) \geq 1.0 \\ t_\%(\mathbf{w}) = 13.05 \\ \text{LER}(\mathbf{w}) \geq 0.011377 \\ \text{TEA}(\mathbf{w}) \geq 6.0^\circ \end{array} \right. \quad (5)$$

A penalty approach will be used to handle the constrained problem.

$$\min_{\mathbf{w} \in W \subseteq \mathbb{R}^n} c_d(\mathbf{w}) + p_{c_l} \max(0, 1 - c_l(\mathbf{w})) + p_{\text{LER}} \max(0, 0.011377 - \text{LER}(\mathbf{w})) + p_{\text{TEA}} \max(0, 6 - \text{TEA}(\mathbf{w})), \quad (6)$$

where the $p_{c_l} = 1000$, $p_{\text{LER}} = 100000$ and $p_{\text{TEA}} = 100$. The equality constraint of the thickness is imposed by scaling the modified airfoil shape to the given value.

6 Probabilistic Design Optimisation Problem

In order to improve the performance of the airfoil under uncertainty, a robust optimisation must be performed. Here, the angle-of-attack is the only parameter considered uncertain, thus representing uncertainty in the environmental conditions. Particularly, the uncertainty has been modelled with a four-parameter beta distribution. The variability range is $\alpha_{ref} \pm 0.25$. The statistical measure chosen to solve the robust design problem is the Conditional Value-at-Risk (CVaR) measure [3, 16], or super-quantile, at a confidence level γ equal to 0.95. This type of risk measures was introduced in the financial sector, but they have been advantageously applied to aerodynamic design optimisation problems [14, 15].

$$\left\{ \begin{array}{l} \min_{\mathbf{w}} \text{CVaR}^{0.95}(c_d(\mathbf{w}, \mathbf{u})) \\ \text{subject to:} \\ \text{CVaR}_{loss}^{0.95}(c_l(\mathbf{w}, \mathbf{u})) \geq 1.0 \\ t_\%(\mathbf{w}) = 13.05 \\ \text{LER}(\mathbf{w}) \geq 0.011377 \\ \text{TEA}(\mathbf{w}) \geq 6.0^\circ \end{array} \right. \quad (7)$$

Therefore, the robust optimisation problem reads:

$$\begin{aligned} \min_{\mathbf{w} \in W \subseteq \mathbb{R}^n} & \text{CVaR}^{0.95}(c_d(\mathbf{w}, \mathbf{u})) + p_{c_l} \max(0, 1 - \text{CVaR}_{loss}^{0.95}(c_l(\mathbf{w}, \mathbf{u}))) + \\ & + p_{\text{LER}} \max(0, 0.011377 - \text{LER}(\mathbf{w})) + p_{\text{TEA}} \max(0, 6 - \text{TEA}(\mathbf{w})), \end{aligned} \quad (8)$$

where $\text{CVaR}_{loss}^{0.95}(c_l) = -\text{CVaR}^{0.95}(-c_l)$ is the loss Conditional Value-at-Risk. The random perturbations of the angle-of-attack impacts only the aerodynamic force requirements (c_d and c_l). The geometric constraints can be evaluated for each design configuration (\mathbf{w}) deterministically. The same penalty parameters have been used as in the deterministic case.

7 Optimisation pipeline

The multi-fidelity surrogate assisted design optimisation strategy of expensive problems [10] has been tailored to reduce the computational cost of the airfoil optimisation problem here studied. Computationally more economic far-field drag predictions are used to populate the training dataset as the computational budget affords only a handful of high-fidelity RANS simulations. The optimisation pipeline is presented in Fig. 5.

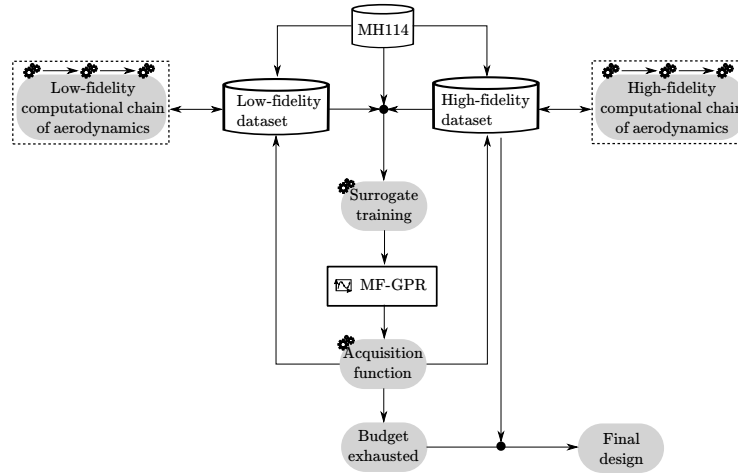


Fig. 5. Optimisation pipeline with multi-fidelity surrogate

Our constrained expected improvement formulation is:

$$\begin{aligned} \text{cEI} &= \mathbb{E} \left[\max \left(0, F_{obj}^* - F_{obj}(\tilde{c}_d, \tilde{c}_l) \right) \right] \mathbb{P} [\tilde{c}_l \geq 1], \\ &= \left(\left(F_{obj}^* - F_{obj}(\hat{c}_d, \hat{c}_l) \right) \Phi \left(\frac{F_{obj}^* - F_{obj}(\hat{c}_d, \hat{c}_l)}{\hat{\sigma}_{c_d}^2} \right) + \right. \\ &\quad \left. + \hat{\sigma}_{c_d}^2 \phi \left(\frac{F_{obj}^* - F_{obj}(\hat{c}_d, \hat{c}_l)}{\hat{\sigma}_{c_d}^2} \right) \right) \Phi \left(\frac{\hat{c}_l - 1}{\hat{\sigma}_{c_l}^2} \right), \end{aligned} \quad (9)$$

where $\hat{\sigma}_{c_d}^2$ and $\hat{\sigma}_{c_l}^2$ are the standard deviations of the drag and lift coefficient, respectively. F_{obj} is the penalised objective given by Eq. (6) and Eq. (8) for the deterministic and robust optimisation studies respectively. The best evaluated objective value is given by F_{obj}^* . The Φ and ϕ symbols denote the cumulative distribution function and probability density function of the standard normal distribution respectively. \mathbb{E} is the expected value and \mathbb{P} is the probability operator. Our optimisation method updates the multi-fidelity surrogate in every iteration with one additional design evaluation. The design of which predicted performance maximises the constrained expected improvement function given in Eq. (9) is evaluated.

$$\mathbf{w}_{new} = \max_{\mathbf{w} \in W \subseteq \mathbb{R}^n} \text{cEI} \quad (10)$$

After finding the most promising design candidate \mathbf{w}_{new} with Eq. 10, the algorithm chooses the fidelity level of the aerodynamic solver based on the Scaled Expected Variance Reduction (SEVR) measure [10]:

$$l = \max_{LF, HF} \text{SEVR}_l, \quad (11)$$

where the SEVR is defined as:

$$\text{SEVR}_{LF} = \frac{\rho(\mathbf{w}_{new}) \hat{\sigma}_{c_d, LF}^2(\mathbf{w}_{new})}{c_{LF}}, \quad (12)$$

$$\text{SEVR}_{HF} = \frac{\rho(\mathbf{w}_{new}) \hat{\sigma}_{c_d, LF}^2(\mathbf{w}_{new}) + \hat{\sigma}_{c_d, \delta_{HF}}^2(\mathbf{w}_{new})}{c_{HF}}, \quad (13)$$

where $c_{LF} = 1$ and $c_{HF} = 10$ are the costs of the low- and high-fidelity simulations respectively.

To calculate the CVaR statistical measures, the computational chains of low- and high-fidelity has to be complemented by an additional step for computing the lift and drag coefficients CVaR risk measures. Due to the heavy computational demand, the risk measure is calculated with a surrogate-based uncertainty quantification approach. At the low fidelity level, for each design configuration, five LF samples are used for constructing a local GPR model, while at the high-fidelity level, five LF samples and three HF samples are used for constructing a local MF-GPR model. These models can then be used to draw a statistically significant number of samples to calculate the risk

measure. The number of HF samples is set to the minimum required number of samples necessary for training the local probabilistic model. We arbitrarily decided to increase the number of LF samples by 20% w.r.t. the HF approximation.

The constructed local probabilistic models of the baseline configuration are presented in Fig. 6. Furthermore, in Fig. 7, the convergence of the risk measure value of the aerodynamic force coefficients in relation to the number of virtual samples is depicted. Based on the CVaR convergence, in the present work 100000 virtual samples of the local probabilistic models are generated to calculate the CVaR values for both aerodynamic coefficients (c_l and c_d).

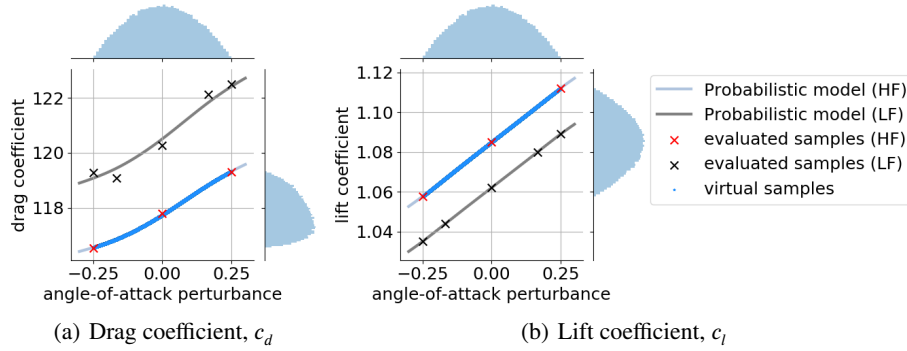


Fig. 6. Local probabilistic models of the aerodynamic force coefficients of the baseline configuration.

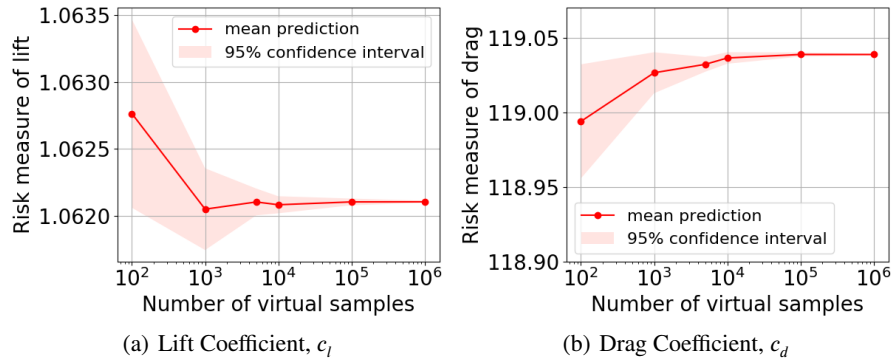


Fig. 7. Convergence of the risk measure value of the aerodynamic force coefficients.

The computational costs of the fidelity levels are set according to their true computational time required for the probabilistic optimisation. The cost of running a low-fidelity CFD evaluation is considered as 1. According to the computational time, the cost for a high-fidelity is 16 times greater than the low-fidelity runs. Considering that to build the low-fidelity probabilistic model five LF samples are needed, the total cost of the LF model is 5. On the other hand, to construct the high-fidelity probabilistic model three HF and five LF evaluations are required, thus the total cost is 53. Therefore, a 1 to 10 cost-ratio is used in this study. Note that the computational costs of training the surrogates and calculating the acquisition function are considered negligible in comparison to the CFD evaluations.

Finally, the computational chain of the aerodynamic forces with the probabilistic model is shown in Fig. 8.

8 Results

In this section, the main obtained results are shown. A two-step optimisation was carried out. Firstly, a deterministic optimisation, which results are explained in Sub-section 8.1, and, later, a probabilistic one. The results of the latter are given in Sub-section 8.2.

8.1 Deterministic optimisation

The deterministic design optimisation was solved with multi-fidelity and single-fidelity surrogate-based techniques, and, also, with population-based technique. A resume of all the optimisation results is presented in Table 3.

Table 3. Comparison of multi-fidelity, single-fidelity surrogate-based and population-based optimisation results. (The cost of a single high- and low-fidelity simulations are 10 and 1 respectively.)

	c_d [dc]	c_l	TEA[°]	LER	Samples [LF,HF]	Cost
baseline (MH114)	117.8	1.09	6.60	0.014324	[0, 1]	10
best design MF-GPR	112.0	1.00	6.03	0.013228	[120, 18]	300
best design GPR	117.6	1.08	7.76	0.011973	[0, 30]	300
best design CMA-ES ($h = 10.7$)	117.1	1.02	6.17	0.018025	[1800, 0]	1800
best design CMA-ES ($h = 8$)	115.6	1.02	6.34	0.020900	[1800, 0]	2700

Regarding the surrogate approach, the same computational budget was used for comparing the single- and multi-fidelity surrogates. In both cases, 10 high-fidelity samples were used for constructing the initial surrogate. In the case of MF-GPR, the surrogate was complemented with 100 low-fidelity samples. Therefore, to keep the same computational budget, only 20 additional samples were generated to complement the high-fidelity surrogate. In Table 3, it can be observed that the optimisation that uses MF-GPR is able to find a better solution. This is due to the fact that the computational budget was severely limited, hence this did not allow us to have enough high-fidelity samples. The

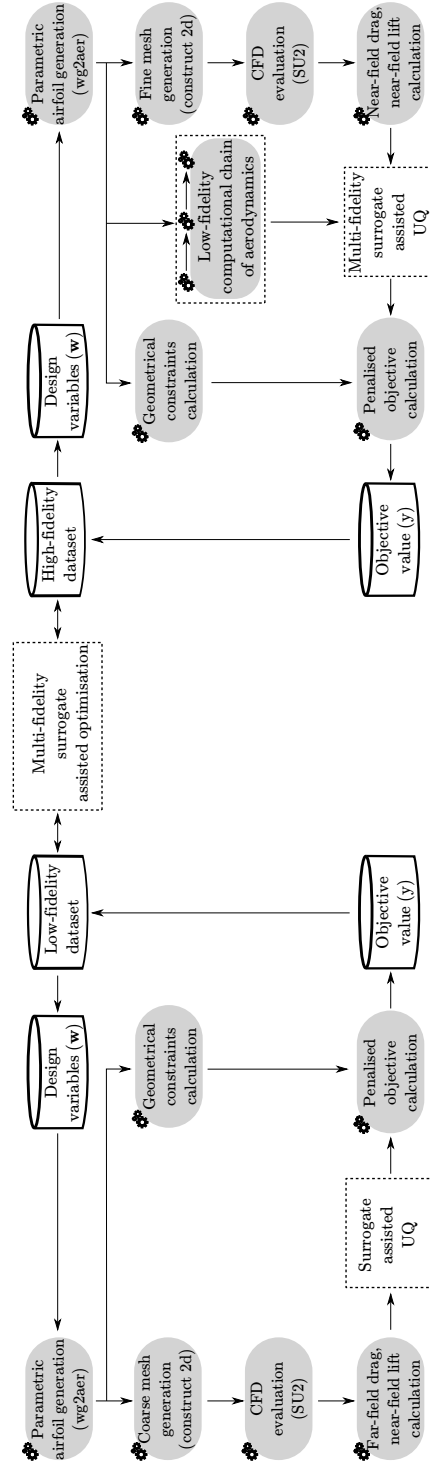


Fig. 8. Computational chain of the aerodynamic forces with the probabilistic model.

lack of HF samples prohibits the construction of an accurate GPR model. However, by introducing low-fidelity information obtained from computationally cheaper samples, the MF-GPR could provide a much better approximation of the performance landscape. Clearly, this handful of new samples were not enough to find a sufficiently good design.

From an aerodynamic point of view, the optimal airfoil of the MF-GPR approach has a lower drag coefficient since its camber line is lower than the camber lines of the other two airfoils, the MH114 and the optimal airfoil obtained at the optimisation using GPR approach. This can be observed in Fig. 9. Note that, for visualization purposes, the axes in Fig. 9 and Fig. 10 are not dependent. Specifically, by decreasing the camber and keeping the free stream angle-of-attack (AoA), the effective AoA of the airfoil actually perceives decreases. Thus, a lower lift coefficient is obtained. This implies a reduction of the lift-induced drag coefficient, hence of the total drag.




Fig. 9. Baseline and deterministic optimal airfoil comparison. MH114 (—), MF-GPR optimal airfoil (—), and GPR optimal airfoil (—). The *dashed* lines are the camber of each airfoil. Axes are not dependent.

In addition, the presented design optimisation approach is compared with a popular population-based algorithm, namely CMA-ES [8]. The optimisation was performed using only low-fidelity CFD evaluations. The evolutionary algorithm was not able to find a similar optimal design to the one given by the MF-GPR approach. Contrarily, the airfoil was barely optimised. Therefore, it was decided to increase the mesh size to $h = 8$ (instead of $h = 10.7$) and redo the optimisation. In this case, the algorithm found a similarly best design configuration to the presented method (MF-GPR). The cost of a CFD evaluation on $h = 8$ grid size is 1.5 times the cost of a low-fidelity one. In addition, to perform the population-based optimisation, 1800 evaluations were needed. This implies that the computational cost is significantly increased with respect to the multi-fidelity approach presented in this paper. Therefore, the presented optimisation method presents an advantage with respect to single-fidelity surrogate-based optimisation. By adding low-fidelity samples, the performance landscape is better approximated. This results in a better allocation of computational resources. The multi-fidelity surrogate-based approach can find better airfoil designs compared to classical population-based optimisation and single-fidelity techniques as the HF evaluations are performed only for promising design candidates.




Fig. 10. Baseline and deterministic optimal airfoil comparison. MH114 (—), MF-GPR optimal airfoil (—), and CMAES optimal airfoil using $h = 8$ grid size (—). The *dashed* lines are the camber of each airfoil. Axes are not dependent.

Furthermore, Fig. 10 shows the comparison of the optimal airfoils obtained with CMA-ES and with MF-GPR approaches with respect to the baseline. In addition, the camber line of each airfoil is also plotted. In this case, both optimal airfoils have a similar camber line. However, the obtained optimal design with CMA-ES has the maximum airfoil thickness placed at a forward position ($x/c = 0.227$) with respect to the MF-GPR best design ($x/c = 0.297$). This implies that the effective angle-of-attack is greater than the one perceived by the MF-GPR optimal airfoil. Hence, a higher c_l and, consequently, c_d are found.

Table 4. Comparison of prediction error of multi- and single-fidelity surrogate models. (Prediction error is defined as the arithmetic mean value of the relative error of the high-fidelity predictions during the course of optimisation.)

	\hat{c}_d	\hat{c}_l	\hat{F}_{obj}	$F_{obj}(\hat{c}_d, \hat{c}_l, \text{TEA, LER})$	HF iterations
MF-GPR surrogate	2.04 %	0.71 %	34.51 %	5.43 %	8
GPR surrogate	3.11 %	5.65 %	17.14 %	11.53 %	20

To evaluate the quality of the produced surrogate models, the mean prediction error is calculated. The results are summarised in Table 4. Particularly, the prediction error was calculated for each HF iteration by using the prediction and true values calculated at the new infill design point. It can be seen that MF-GPR predicted the aerodynamic forces of the new designs significantly better. The poor design configuration found by the GPR-based optimisation is also due to the penalised approach which was employed here. The lift coefficient is not well predicted which drives the algorithm to waste computations on designs which are unfeasible and have high objective values. Moreover, in Table 4, it is also shown that the objective cannot be accurately predicted by a surrogate directly. However, by independently predicting the aerodynamic forces and calculating the objective afterwards based on these predictions, it was more accurately predicted.

8.2 Probabilistic optimisation

For improving the performance of the airfoil under uncertainty, a probabilistic optimisation was carried out. Specifically, only uncertainty on the angle-of-attack was considered. Considering that the advantages of using MF-GPR were already shown in the deterministic optimisation, the probabilistic one was only made using the proposed optimisation approach.

The predicted lift and drag distributions of the optimal designs and the baseline airfoil are shown in Fig. 11. By perturbing the angle-of-attack, the obtained optimal design for the deterministic problem violates the constraint imposed for lift coefficient, as shown in Fig. 11 (a). The MF-GPR algorithm is able to take uncertainty into account during the optimisation; hence, it can find an optimal design which respects the imposed lift constraint. Particularly, for the probabilistic design optimisation, it was decided that $\text{CVaR}_{\text{loss}}(\tilde{c}_l)$ must be greater or equal to one. Figure 11 (a) shows how the probabilistic optimum design fulfils the imposed c_l constraint. Consequently, this design has a higher drag.

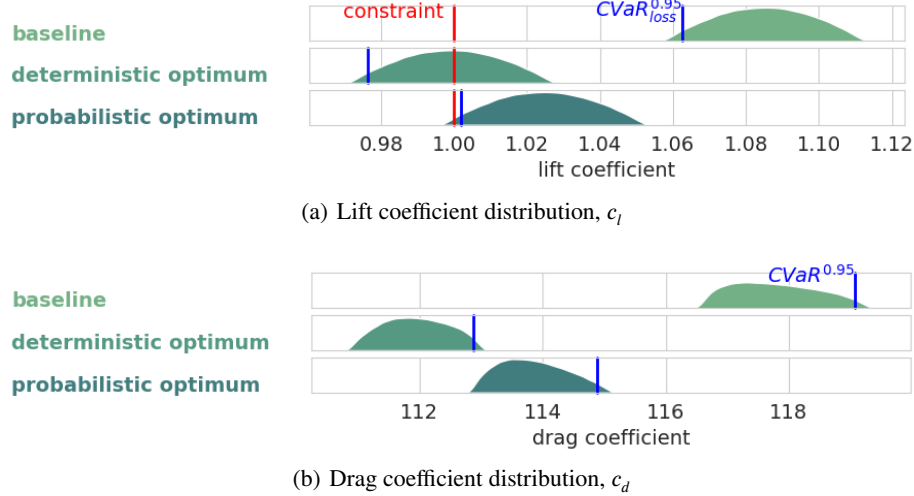


Fig. 11. Prediction of the distributions for the baseline and optimal designs.

In Table 5, a comparison of the baseline airfoil with the optimal designs is given in terms of lift, drag, and shape characteristics. Besides, in order to calculate the CVaR, the MF-GPR technique requires the computation of three high-fidelity samples. Hence, the computational budget is triple. However, this is still limited as the size of the budget is equal to the cost of only 90 high-fidelity simulations.

Table 5. Comparison of the requirements considering environmental uncertainty.

	Objective	CVaR(\tilde{c}_d) [dc]	CVaR _{loss} (\tilde{c}_l)	TEA	LER	Samples [LF,HF]	Cost
baseline (MH114)	119.07	119.07	1.063	6.60	0.014324	[5, 3]	35
deterministic opt.	136.55	112.87	0.976	6.03	0.013228	[5, 3]	35
probabilistic opt.	114.89	114.89	1.002	6.38	0.016937	[370, 51]	880

The obtained optimal airfoils are compared in Fig. 12. It is appreciated that both the deterministic and the probabilistic optimisation resulted in a smaller camber line curvature airfoil than the baseline. The MH114 airfoil generates a significant higher lift coefficient than the required constraint value, hence the optimisation tends to reduce the camber curvature, so that the lift reduces and so does the drag. Moreover, by comparing the deterministic and probabilistic optimal designs, it can be seen that the probabilistic optimum has a stronger S-shaped lower side. This increases drag and lift coefficient; hence, resulting in a feasible airfoil design.




Fig. 12. Baseline, deterministic optimum, and robust optimum airfoil comparison. MH114 (—), deterministic optimum (—), and probabilistic optimum (—). The *dashed* lines are the camber of each airfoil. Axes are not dependent.

Finally, in Fig. 13, a comparison of the pressure coefficient distribution and friction coefficient on the body surface of the baseline, deterministic and probabilistic optimal designs is presented. Analysing both optimal designs, it can be observed that the deterministic optimum presents a smoother expansion rate on the upper surface of the airfoil. Specifically, the maximum is reached at 30% of the chord, whereas the probabilistic optimum has the peak at 10% of the chord. Comparing the pressure coefficient distribution of the optimal airfoils, the contribution of pressure to drag coefficient is clearly higher for the probabilistic optimum. Besides, for the contribution of friction on drag coefficient, a similar conclusion can be applied. The maximum value of skin friction coefficient is higher for the probabilistic optimum on the suction side of the airfoil and, also, to a minor extent on the pressure side. Thus, the friction drag is higher for the probabilistic solution, mainly due to the peak of maximum c_f .




Fig. 13. Pressure coefficient (*left*) and friction coefficient (*right*) on the body surface comparison. MH114 (—), deterministic optimum (—), and probabilistic optimum (—).

9 Conclusion

In this present work, a complete optimisation workflow is presented for expensive aerospace applications under uncertainty. The workflow is employed to find optimal airfoil designs that produce minimal drag and respect both aerodynamic force and geometrical constraints. The prediction of aerodynamic forces is an expensive process as the Reynolds Averaged Navier-Stokes partial differential equations have to be solved numerically. To decrease the computational effort, a multi-fidelity surrogate-based technique was used. The drag and lift predictions obtained with two significantly different size meshes were fused together with a hierarchical Gaussian process regression technique. To increase the correlation of the high- and low-fidelity drag predictions, the spurious drag was compensated by performing the far-field drag prediction at low-fidelity level. Our approach was compared against a classical single-fidelity surrogate-based and an evolutionary algorithm. The results showed that classical methods could struggle to find significantly improved designs due to the limited computational budget. The highest potential of our multi-fidelity surrogate-based approach relies on solving problems under uncertainty where the required numerous probabilistic samples can be efficiently obtained by introducing a multi-fidelity probabilistic model.

References

1. Construct2d. <https://sourceforge.net/projects/construct2d/>
2. Martin hepperle mh 114 for a propeller for ultralight. <http://airfoiltools.com/airfoil/details?airfoil=mh114-il>, accessed: 2020-06-28
3. Acerbi, C., Tasche, D.: Expected shortfall: A natural coherent alternative to value at risk. *Economic Notes* **31**(2), 379–388 (2002). <https://doi.org/10.1111/1468-0300.00091>, <http://dx.doi.org/10.1111/1468-0300.00091>

4. Arizmendi, B., Bellosta, T., del Val, A.I., Gori, G., Prazeres, M.O., Reis, J.: On real-time management of on-board ice protection systems by means of machine learning. In: AIAA Aviation 2019 Forum. p. 3464 (2019)
5. Economon, T.D., Palacios, F., Copeland, S.R., Lukaczyk, T.W., Alonso, J.J.: SU2: An open-source suite for multiphysics simulation and design. AIAA Journal **54**(3), 828–846 (2016). <https://doi.org/10.2514/1.J053813>, <https://doi.org/10.2514/1.J053813>
6. Fan, Y., Li, W.: Review of far-field drag decomposition methods for aircraft design. Journal of Aircraft **56**(1), 11–21 (2019)
7. Gariepy, M., Trepanier, J.Y., Petro, E., Malouin, B., Audet, C., LeDigabel, S., Tribes, C.: Direct search airfoil optimization using far-field drag decomposition results. In: 53rd AIAA Aerospace Sciences Meeting. AIAA SciTech Forum, American Institute of Aeronautics and Astronautics (Jan 2015). <https://doi.org/10.2514/6.2015-1720>, <https://doi.org/10.2514/6.2015-1720>
8. Hansen, N., Müller, S.D., Koumoutsakos, P.: Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es). Evolutionary computation **11**(1), 1–18 (2003)
9. Kennedy, M.C., O'Hagan, A.: Predicting the output from a complex computer code when fast approximations are available. Biometrika **87**(1), 1–13 (2000)
10. Korondi, P.Z., Marchi, M., Parussini, L., Poloni, C.: Multi-fidelity design optimisation strategy under uncertainty with limited computational budget. Optimization and Engineering (2020)
11. Le Gratiet, L., Garnier, J.: Recursive co-kriging model for design of computer experiments with multiple levels of fidelity. International Journal for Uncertainty Quantification **4**(5) (2014)
12. Morales Tirado, E., Quagliarella, D., Tognaccini, R.: Airfoil optimization using far-field analysis of the drag force. In: AIAA Scitech 2019 Forum. p. 0972 (2019)
13. Paparone, L., Tognaccini, R.: Computational fluid dynamics-based drag prediction and decomposition. AIAA Journal **41**(9), 1647–1657 (2003)
14. Quagliarella, D.: Uncertainty Management for Robust Industrial Design in Aeronautics: Findings and Best Practice Collected During UMRIDA, a Collaborative Research Project (2013–2016) Funded by the European Union, chap. Value-at-Risk and Conditional Value-at-Risk in Optimization Under Uncertainty, pp. 541–565. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-319-77767-2_34, https://doi.org/10.1007/978-3-319-77767-2_34
15. Quagliarella, D., Tirado, E.M., Bornaccioni, A.: Risk measures applied to robust aerodynamic shape design optimization. In: Flexible Engineering Toward Green Aircraft, pp. 153–168. Springer (2020)
16. Rockafellar, R.T., Uryasev, S.: Conditional value-at-risk for general loss distributions. Journal of Banking and Finance **26**, 1443–1471 (2002)
17. Spalart, P., Allmaras, S.: A one-equation turbulence model for aerodynamic flows. In: 30th aerospace sciences meeting and exhibit. p. 439 (1992)
18. Wang, X., Hirsch, C., Liu, Z., Kang, S., Lacor, C.: Uncertainty-based robust aerodynamic optimization of rotor blades. International journal for numerical methods in engineering **94**(2), 111–127 (2013)
19. Zang, T.A.: Needs and opportunities for uncertainty-based multidisciplinary design methods for aerospace vehicles. National Aeronautics and Space Administration, Langley Research Center (2002)