Multi-objective design optimisation of an airfoil with geometrical uncertainties leveraging multi-fidelity Gaussian process regression *

Péter Zénó Korondi[1,2][0000−0002−9935−1059], Mariapia Marchi[2][0000−0002−7107−9073], Lucia Parussini[1][0000−0003−1790−9731], Domenico Quagliarella[3][0000−0003−3828−4937], and Carlo Poloni[1,2][0000−0001−8077−0563]

[1] Department of Engineering and Architecture, University of Trieste, Piazzale Europa, 1, 34127 Trieste, Italy.
peterzeno.korondi@phd.units.it
[2] ESTECO S.p.A, Building B, 99 Padriciano, Area Science Park, Trieste, 34149, Italy.
[3] Italian Aerospace Research Centre, Capua, Italy

Abstract. This paper presents the multi-objective optimisation of the MH114 high-lift airfoil. We seek the set of Pareto optimal solutions that maximise the airfoil lift and minimise the drag. The lift and drag forces are considered uncertain due to geometrical uncertainties. The uncertainty quantification of the probabilistic aerodynamic force values requires a large number of samples. However, the prediction of the aerodynamic forces is expensive due to the numerical solution of the Navier-Stokes equations. Therefore, a multi-fidelity surrogate assisted approach is employed to combine expensive RANS simulations with cheap potential flow calculations. The multi-fidelity surrogate-based approach allows us to economically optimise the aerodynamic design of the airfoil under uncertainty.

Keywords: multi-objective optimisation under uncertainty · airfoil design · surrogate assisted optimisation· Gaussian process regression

1 Introduction

Shape optimisation of an airfoil is one of the most fundamental problems in aerodynamic design optimisation. The purpose of an airfoil is to generate a pressure difference in a flow so that a force is generated. The force component perpendicular to the flow direction is called lift, its magnitude and sense (in respect of the defined force reference frame) depend on the shape of the airfoil and on the flow conditions. Together with the lift, the presence of the airfoil in the flow inevitably generates a force component parallel to the flow direction, called drag. Most engineering applications exploit the lift, while the drag is

* This work was partially supported by the H2020-MSCA-ITN-2016 UTOPIAE, grant agreement 722734.

an inevitable loss. Therefore, the shape optimisation of an airfoil aims to find the optimal lift-to-drag ratio such that some additional application-dependent requirements are also satisfied.

In reality, solutions that are optimise all objectives simultaneously (as suggested by the reviewer) are typically non-existing. A single-objective problem might have a single global optimal solution. However, formulating a real problem as single-objective implies a decision making on the preferences of various requirements. By formulating our problem as a multi-objective optimisation, the preferences of various requirements are decided after a set of Pareto optimal solutions are found [3]. Consequently, in the context of multi-objective airfoil optimisation, we are searching for a set of airfoil designs which are Pareto optimal to our multi-objective problem formulation. This set is called Pareto front. When the objectives are conflicting, the Pareto set contains more then one solutions that cannot be improved in any of the objectives without degrading at least one of the others.

To obtain the Pareto front various algorithms exist. However, they commonly require many performance evaluations of the underlying problem. This is troublesome for the aerodynamic shape optimisation of an airfoil, as accurate computational fluid dynamics (CFD) calculations are typically expensive in computational time [15].

This issue can be tackled by employing surrogate models [13, 14]. Expensive calculations are performed for only a handful of designs. Then a statistical model is built to approximate the aerodynamics of airfoil designs which have not been evaluated by the expensive CFD code. The accuracy of the statistical model highly depends on the number of available CFD evaluations. Consequently, sparsely sampled design landscapes are hard to approximate accurately with standard surrogate techniques. In such a case, aerodynamic calculations of lower fidelity can be used to provide sufficient information for building an accurate statistical model. The information from lower fidelity calculations can be fused together with high fidelity data by using the multi-fidelity Gaussian process regression (MF-GPR) [9, 11].

There is also another issue to take into account. Often the actual design, or operation point, and its performance are slightly different from the optimisation solutions because of manufacturing, wear off and other operational deformations, like icing and surface pollution. In practice, our design problem is affected by various uncertainty sources which affect the actual performance. This issue can be addressed with uncertainty quantification (UQ) techniques and optimisation under uncertainty methods. UQ can be used to estimate statistical measures of the design performance that can be in turn used as reliable or robust objectives of the optimisation under uncertainty problem (see e.g. the reviews [2, 19]).

In our recent work, we proposed a multi-fidelity optimisation workflow for optimising expensive multi-objective problems under uncertainty [10]. In this work, we investigate the performance and applicability of our proposed workflow for the aerodynamic shape optimisation of an airfoil.

The optimisation problem is presented in Section 2. In Section 3 the used aerodynamic solvers are examined. In Section 4 the MF-GPR method is introduced and Section 5 discusses on the uncertainty treatment. The proposed optimisation workflow is briefly presented in Section 6. Results are analysed in Section 7. and our conclusions are expressed in Section 8.

2 Design optimisation problem of airfoil

Conceptually, aerodynamic design optimisation can be approached in two ways: inverse and direct [20]. In inverse design optimisation a desired pressure distribution is targeted and the optimisation algorithm seeks to find the geometrical shape which produces the targeted pressure distribution or lift force. Therefore, inverse methods are commonly applied in later design phases when target values are known. A direct method, as its name suggests, directly optimises the objective without any predefined target value. In aerodynamic shape optimisation a maximal lift-to-drag ratio is typically desired.

The geometrical shape of the airfoil is defined by superposing modal shape functions on the baseline geometry of MH114. Eight modal shape functions are considered. They are shown in Figure 1 and Table 1. The first two modes modify the thickness and the camber line of the airfoil. The remaining six modes introduce local shape modifications of the upper and lower side of the airfoil at the leading edge (LE), mid span and trailing edge (TE) respectively. The design variables to be optimised are the scaling parameters (x_i) of the modal shape functions. Additionally, the shape of the airfoil is considered to have some uncertainties due to the manufacturing process. Therefore, each shape mode is superposed on the design shape with an uncertain scaling factor (u_i).

Table 1: Design and uncertain variables

Mode	Function type	Design variable	Uncertain variable	Physical interpretation
mode 1	Polynomial	x_1	u_1	thickness
mode 2	Polynomial	x_2	u_2	camber
mode 3	Hicks-Henne	x_3	u_3	upper LE
mode 4	Hicks-Henne	x_4	u_4	lower LE
mode 5	Hicks-Henne	x_5	u_5	upper middle
mode 6	Hicks-Henne	x_6	u_6	lower middle
mode 7	Hicks-Henne	x_7	u_7	upper TE
mode 8	Hicks-Henne	x_8	u_8	lower TE

In this work, the MH114 airfoil is optimised for a high-lift propeller. The set of geometries are sought which are Pareto optimal for maximal lift (L) and minimal drag (D). As the lift and drag forces are uncertain due the geometrical uncertainties, a reliability based multi-objective optimisation problem is considered here:

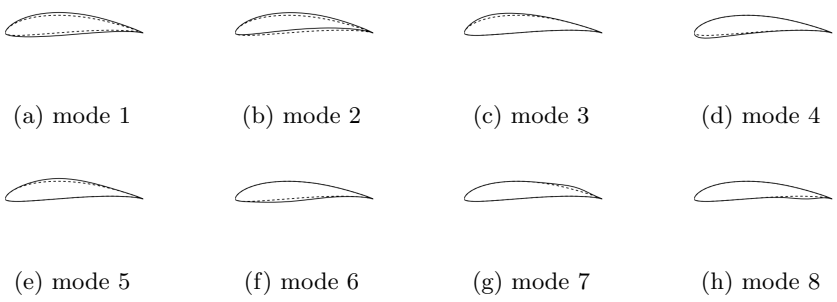

(a) mode 1 (b) mode 2 (c) mode 3 (d) mode 4

(e) mode 5 (f) mode 6 (g) mode 7 (h) mode 8

Fig. 1: Design modes

$$\min_x \quad \mathcal{S}_{95}\left(-C_l(x,\tilde{u})\right) \tag{1a}$$

$$\min_x \quad \mathcal{S}_{95}\left(\frac{C_d(x,\tilde{u})}{C_{d0}}\right) \tag{1b}$$

where $C_{d0} = 0.01$. The S_{95} denotes the 95th superpercentile which is a reliability measure defined in Section 5. The lift coefficient (C_l) and drag coefficient (C_d) are:

$$C_l = \frac{L(x,\tilde{u})}{1/2\rho U^2} \tag{2a}$$

$$C_d = \frac{D(x,\tilde{u})}{1/2\rho U^2} \tag{2b}$$

ρ is the density and U is the free-stream velocity of the air. Compared to the traditional C_l and C_d coefficient definitions, the chord length is omitted from the above definitions as the chord length is considered as unit throughout this work.

3 Solvers

To calculate the aerodynamic forces of the airfoil two solvers are considered: XFOIL [5] and SU2 [6].

The former is an airfoil analysis tool based on potential flow equations (panel methods). For viscous problems, a two-equation integral boundary layer formulation is coupled with the inviscid flow solution [4]. The transition criteria is calculated by the e^N envelope method. XFOIL has a fairly rapid calculation time and provides sufficient accuracy for most engineering applications.

The SU2 software provides a solver for the compressible Reynolds-averaged Navier-Stokes equation. The RANS equation is closed by Menter's Shear Stress

Transport turbulence model [12] which efficiently blends the k-ω turbulence model of the near wall region to the k-ϵ model of the region away from any wall. Various studies have been carried out to compare the results of these two solvers, like e.g. [1, 22]. Both solvers are suitable to accurately predict the aerodynamic forces of an airfoil. For the sake of this work we will consider SU2 as a higher fidelity solver as it implements a more general form of the Navier-Stokes equation.

The aerodynamic evaluations with XFOIL and SU2 are performed with the framework software described in [16,17]. The modal shape function superposition is performed with the *wg2aer*[4]. The modified airfoil geometry is stored in a *Selig* format which can be directly processed by XFOIL. For the CFD evaluation, the modified airfoil and its surrounding domain are discretised with the open-source *Gmsh* software which generates the mesh in *.su2* format. Finally, SU2 performs the aerodynamic analysis of the airfoil and provides the high-fidelity drag and lift predictions.

The lift and drag coefficients of the MH114 airfoil are plotted in Figures 2a and 2b. The calculations are carried out at $Re = 5 \cdot 10^6$ and $M = 0.218$ with standard sea-level conditions. SU2 considers the domain around the airfoil fully turbulent. Therefore, XFOIL was also forced to operate in fully turbulent conditions by setting the transition point location at the beginning of lower and upper airfoil sides (XTRLO and XTRUP set to 0.01). We can see that the two solvers produce similar polar trends; however, there are some deviations in the actual values. This makes the two solvers appropriate candidates for a multi-fidelity optimisation.

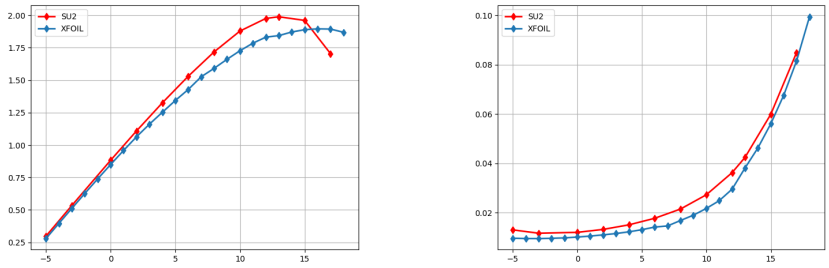

(a) Lift coefficient against angle-of-attack (b) Drag coefficient against angle-of-attack

Fig. 2: Comparison of lift and drag coefficient curves of MH114 calculated with XFOIL and SU2.

[4] Software developed by the Italian Aerospace Research Centre (CIRA)

4 Multi-fidelity Gaussian Process Regression

The multi-fidelity Gaussian process regression (MF-GPR) is briefly discussed in this section. This technique tailors the well-known Gaussian process regression (GPR)[5] to fuse information from various fidelity sources into a single surrogate [8, 9]. We employ the recursive formulation proposed by [11]:

$$\begin{aligned} \tilde{f}_{\mathrm{LF}}(\mathbf{x}) &= \mathbf{h}_{\mathrm{LF}}^{T}(\mathbf{x})\boldsymbol{\beta}_{\mathrm{LF}} + \tilde{\delta}_{\mathrm{LF}}(\mathbf{x}), \\ \tilde{f}_{\mathrm{HF}}(\mathbf{x}) &= \rho(\mathbf{x})\tilde{f}_{\mathrm{LF}}(\mathbf{x}) + \mathbf{h}_{\mathrm{HF}}^{T}(\mathbf{x})\boldsymbol{\beta}_{\mathrm{HF}} + \tilde{\delta}_{\mathrm{HF}}(\mathbf{x}), \\ \rho(\mathbf{x}) &= \mathbf{g}^{T}(\mathbf{x})\boldsymbol{\beta}_{\rho}, \end{aligned} \tag{3}$$

where a least squares regression $\mathbf{h}_i^T(\mathbf{x})\boldsymbol{\beta}_i$ with $i = \mathrm{HF}, \mathrm{LF}$ formulates the mean trend of the fidelity level. Correspondingly, $\mathbf{h}_i(\mathbf{x})$ is the vector of regression functions and $\boldsymbol{\beta}_i$ is the vector of regression coefficients. The local variations of the model are incorporated into $\tilde{\delta}_i(\mathbf{x}) \sim \mathcal{N}(0, \sigma_i^2)$ and modelled as zero mean Gaussian distributions with σ_i^2 variance.

This MF-GPR formulation is hierarchical. The low-fidelity level is modelled by a GPR. The high-fidelity model builds an additional GPR using the posterior distribution of the low-fidelity level. A GPR problem is solved at each level without the need to construct a covariance matrix which contains the observations of all fidelity levels as in [9]. The surrogate is frequently updated during the optimisation; hence, the smaller size of the covariance matrix can result in a significant computational speed-up.

5 Uncertainty treatment

When simulating a flow around an airfoil only a limited number of phenomena are modelled. Therefore, the aerodynamic performance of a real airfoil might deviate from the numerical results. This motivates the construction of probabilistic models which are appended to the design optimisation workflow to predict the variations of the aerodynamic performance. Uncertainty modelling techniques are grouped as: deterministic, probabilistic and possibilistic, according to [2].

In this work, only geometrical uncertainties are considered. It is assumed that their nature is probabilistic and they can be described by a Gaussian distribution. The geometrical uncertainties are propagated through the aerodynamic solver which will result in a probabilistic aerodynamic performance. The comparison of two probability distributions is not a trivial task. The possible realisations of a distribution are multiple [21]. Therefore, only certain properties of a distribution are compared. In this work, the 95th superpercentile[6] of the C_l and C_d distributions are used for evaluating the aerodynamic performance of the airfoil.

[5] Gaussian process regression is also called Wiener–Kolmogorov prediction and Kriging.

[6] The k-th superpercentile of a probabilistic variable is the expected value of all possible realisations which are not smaller than the k-th percentile value. This measure is commonly called as superquantile or conditional Value-at-Risk as well.

This risk measure is employed to ensure reliability. The advantages of the superpercentile measure over other risk measures for engineering applications are discussed in [18, 21]. The 95th superpercentile is calculated with the following equation:

$$\mathcal{S} = \overline{q}_\zeta = \mathbb{E}\left[\tilde{y} | \tilde{y} \geq q_\zeta(\tilde{y})\right] = \frac{1}{1-\zeta/100} \int_{\zeta/100}^{1} q_\tau(\tilde{y}) d\tau, \tag{4}$$

where q_ζ is the ζ-th percentile and $\zeta = 95$.

Analytical propagation of the uncertainty is not possible due the complex aerodynamics solvers. Therefore, the superpercentile values of the C_l and C_d distributions are calculated using empirical values obtained by sampling. To obtain a sufficient number of samples, surrogate assisted uncertainty quantification is performed as in [10]. The probabilistic space is considered to be independent from the design space and for each design a local GPR is built on 15 probabilistic samples. The calculation of the superpercentile is then done with 5000 virtual samples of the local probabilistic space. These values seemed to be adequate after performing a few trial and error checks.

6 Multi-objective optimisation framework for airfoil optimisation under uncertainty

This work employs the optimisation workflow proposed in [10]. The workflow embodies a multi-fidelity Bayesian optimisation for multi-objective problems and in this work we employed it for aerodynamic design optimisation of airfoils using XFOIL and SU2.

The workflow can be divided into three major components: design of experiments (DoE), multi-fidelity surrogate construction and acquisition function. To initialise our optimisation workflow a DoE technique is employed to obtain a dataset for surrogate construction. At each fidelity level the design space is sampled by uniform Latin Hypercube Sampling (LHS).

After obtaining the 95th superpercentiles of lift and drag, two independent multi-fidelity surrogate are trained using the recursive formulation defined by Eq. (3).

The two MF-GPR models are used by the acquisition function to determine which design configuration should be evaluated in the next iteration and which solver should perform the aerodynamic calculation. Since it is a multi-objective problem, the decision on the next design location is made by maximising the hypervolume improvement of the lower confidence bound of the drag and lift coefficients following the suggestion of [7]. For the selected design configuration the superpercentile values of the drag and lift coefficients are calculated by evaluating the corresponding probabilistic sample with XFOIL or SU2. The selection on the solver is based on Scaled Expected Variance Reduction (SEVR) values of the fidelity levels [10].

With the new superpercentile values the surrogate model of the lift and drag can be retrained. The surrogate is updated with new designs until the

computational budget is exhausted. At the end, the set of Pareto optimal designs are presented to decision makers. The complete optimisation workflow is depicted in Figure 3.

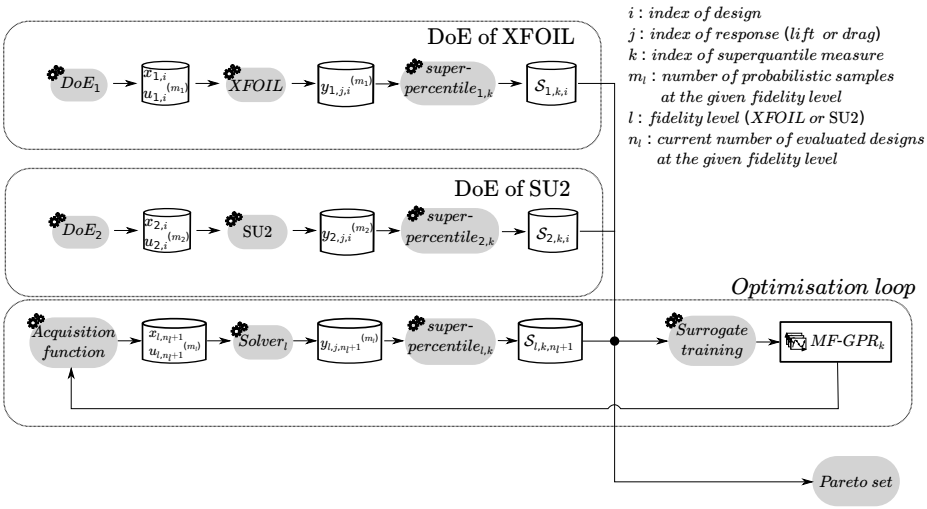

Fig. 3: Multi-objective probabilistic optimisation workflow for aerodynamic shape optimisation of an airfoil with MF-GPR

7 Results

A brief summary of the solved optimisation problem is presented in Table 2. The problem is bi-objective and has 8 design and 8 uncertain variables. The computational budget is set to 136500 units.

Table 2: Summary of optimisation problem

Number of objectives	2
Number of constraints	0
Number of design variables	8
Number of uncertain variables	8
Computational budget	136500

The number of evaluated LF and HF samples and their cost are presented in Table 3. We assigned 300 cost units for running a single evaluation of aerodynamics forces with SU2 and 1 unit for evaluating the design with XFOIL.

Here, we determined the cost of the fidelity levels based on the actual running times of the simulations on the used machine. The optimisation stopped when no further high-fidelity samples could be added to the surrogate training set. After the 435th HF simulation only 4245 units remained in the budget which is not sufficient for generating the required 15 samples to build the probabilistic model. Therefore, only the 96.8 % of the budget was used.

Table 3: Number of LF and HF samples and their costs

	LF DoE	LF total	HF DoE	HF total
Aerodynamic solver	XFOIL	XFOIL	SU2	SU2
Evaluation cost	1	1	300	300
Total number of samples	450	1755	225	435
Budget spent	450 (0.3 %)	1755 (1.2%)	67500 (51.7 %)	130500 (95.6%)

In Figure 4, we can see that the algorithm alternates the fidelity levels. The alternation stems from the fact that, in regions where the expected improvement is high due to large uncertainties, the algorithm will evaluate the new design with the low-fidelity solver. Following this step, high expected improvement values in the region are the results of promising performance prediction with low-level of uncertainty. Therefore, the region can be sampled by high-fidelity simulation without risking the waste of budget.

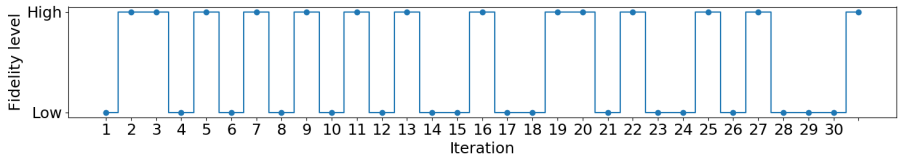

Fig. 4: History of fidelity selection

To investigate the advantage of MF-GPR over single-fidelity GPR, the relative prediction errors of the surrogates are calculated in every iteration when a HF sample is generated. The classical GPR model is built in every iteration using only the HF samples of the actual iteration. Overall, the MF-GPR provides a better prediction; however, in some iterations GPR can temporarily be the best predictor as shown in Figure 5. At each iteration, the prediction error is calculated based on a single sample. When the newly evaluated design lies in a region which can be accurately predicted by a GPR model using only HF samples, it is possible that the single-fidelity GPR model provides a slightly better prediction. However, in the majority of the iterations MF-GPR outperforms the GPR model. In Table 4 we show the mean prediction error of the single- and multi-fidelity

surrogates. Both surrogate models can provide a relatively accurate prediction of the objectives in overall. However, the single-fidelity surrogate model shows significantly bigger prediction errors when the newly evaluated design lies in a region which was not explored by a sufficient number of observations.

Table 4: Mean prediction error

	GPR	MF-GPR
Objective 1 (lift)	0.77 %	0.31 %
Objective 2 (drag)	1.51 %	0.99 %

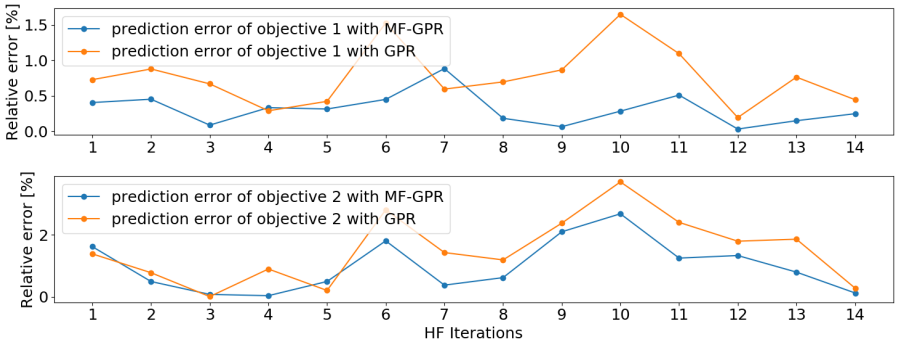

Fig. 5: Comparing relative prediction errors of MF-GPR and GPR (objective 1: lift, objective 2: drag)

Throughout the iterations, the correlation between the predictions at low and high-fidelity levels are high and steady as shown in Figure 6. The high correlation is expected at the beginning of the optimisation as the high number of LF samples in the training data results in a MF-GPR model which predicts performances close to the LF observations. This correlation does not deteriorate by updating the model with high-fidelity samples which suggests that the initial MF-GPR dominated by the LF samples provides a good approximation.

The obtained Pareto optimal solutions are depicted in Figure 7 (red circles, Pareto front HF). The initial HF Pareto front obtained after the DoE (dash-dotted grey line) was significantly improved. The MF-GPR models of the objectives can provide accurate predictions; hence, most of the design locations suggested by the acquisition function are Pareto optimal. In the same figure, the Pareto front of the LF samples is shown. It could seem that the LF Pareto optimal solutions dominate the HF front. However, as Figures 2a and 2b also suggest, the drag and lift force are actually under-predicted by using XFOIL.

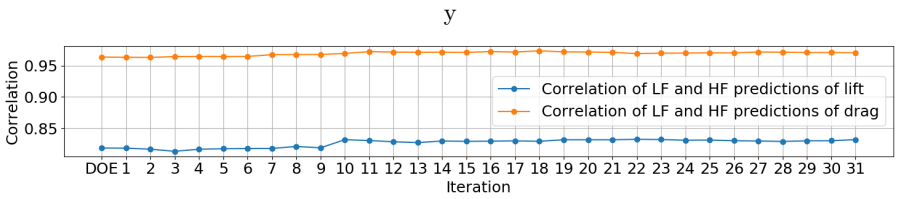

Fig. 6: Correlation history of objectives

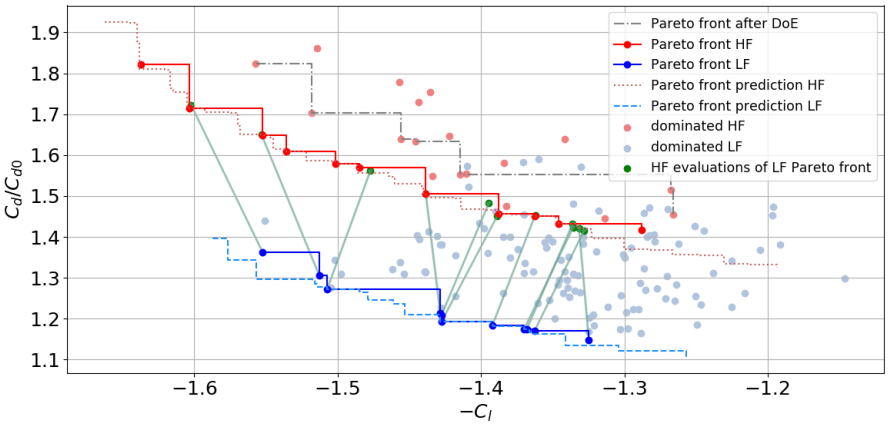

Fig. 7: Comparison of high and low-fidelity Pareto fronts

Indeed, by re-evaluating the Pareto optimal LF optimal designs with the HF solver, the green circles (HF evaluations of LF Pareto front) are found. Thus, we can conclude that the gap between the LF and HF front is due to the approximation error of the LF evaluations. Therefore, the introduction of the HF samples into the surrogate model constructions is beneficial for obtaining an accurate Pareto front.

The list of design variables and objectives of the non-dominated designs are shown in Table 5. We can see that many design variables reach the boundary of the design variable limits. The design variable of the thickness mode set to -1 for every non-dominated design. This is expected as thin airfoils produce significantly less drag.

Depending on some further criteria on the propeller blade, the decision maker can choose the preferred airfoil (among the HF Pareto optimal solutions) for further analysis. For example, Table 6, lists three possible designs corresponding to the minimum drag, highest efficiency and maximum lift design. To compare the predicted probability distributions of their aerodynamic forces, the baseline and the most efficient design are shown in Figure 8b. Violin plots are drawn based on the 5000 virtual samples mentioned in Section 5. By optimising the

Table 5: List of design variables and objectives of the non-dominated designs

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	$\mathcal{S}_{95}\left(-C_l\right)$	$\mathcal{S}_{95}\left(C_d/C_{d0}\right)$
-1.000	1.000	1.000	1.000	1.000	-1.000	0.000	-1.000	-1.637	1.822
-1.000	1.000	-0.084	1.000	-0.119	-1.000	-0.526	-1.000	-1.603	1.715
-1.000	1.000	-1.000	1.000	-1.000	-0.113	-1.000	-1.000	-1.552	1.649
-1.000	0.663	-0.927	0.654	-1.000	-1.000	-0.595	-1.000	-1.536	1.608
-1.000	0.611	-0.458	-1.000	-0.659	-0.608	-0.345	-1.000	-1.501	1.578
-1.000	0.917	0.272	-1.000	-1.000	-1.000	-1.000	-0.199	-1.485	1.569
-1.000	-0.128	-1.000	0.350	-1.000	-1.000	-0.616	-0.921	-1.439	1.505
-1.000	-0.215	1.000	-1.000	-1.000	-1.000	-1.000	-1.000	-1.388	1.457
-1.000	-0.717	-1.000	-1.000	-1.000	-1.000	0.020	-1.000	-1.363	1.450
-1.000	-1.000	-0.379	1.000	-1.000	-1.000	-1.000	-1.000	-1.346	1.433
-1.000	-1.000	1.000	-1.000	0.000	1.000	-1.000	-1.000	-1.288	1.418

Table 6: Optimal designs of various criteria

	$\mathcal{S}_{95}\left(C_l\right)$	$\mathcal{S}_{95}\left(C_d\right)$	η
Baseline design	1.4222	0.01646	86.4
Minimum drag design	1.2879	0.01418	90.8
Most efficient design	1.5360	0.01608	95.5
Maximum lift design	1.6371	0.01822	89.9

95-th superpercentile of the aerodynamic forces, designs are preferred of which aerodynamic performance does not deteriorate due to the geometrical uncertainties. To be more specific, the superpercentile optimisation aims at shifting the probability distribution towards more reliable performance values w.r.t. the distribution of deterministic designs. However, in the case of maximum lift and most efficient designs, we also found more robust solutions (narrower spread of the distribution) with better performance (higher lift and smaller drag distributions) as side effect.

The optimisation problem defined with Eqs. 1a-1b aims to optimise exclusively the drag and lift forces, other aerodynamic and structural parameters of the airfoil are neglected. This can result in airfoil designs which are sub-optimal when structural requirements and other aerodynamic parameters are considered. For example, the most efficient design in Figures 9c has a very thin trailing edge which is undesired from a structural point of view. Nevertheless, we can see that the pressure distribution in Figures 9b, 9d and 9f is well approximated with XFOIL. This explains the effectiveness of the used multi-fidelity approach.

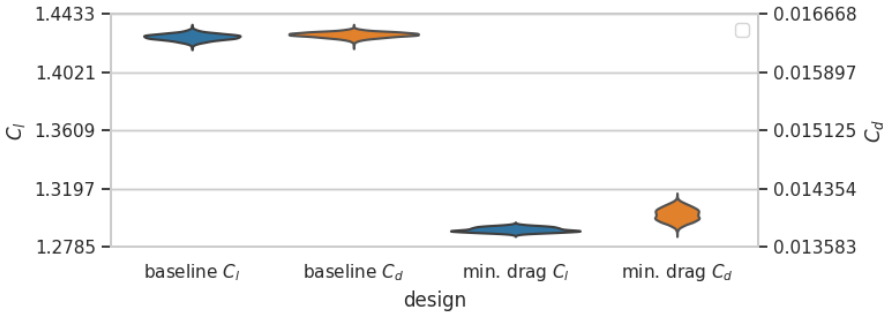

(a) Comparison of the baseline and minimum drag design with uncertainty prediction

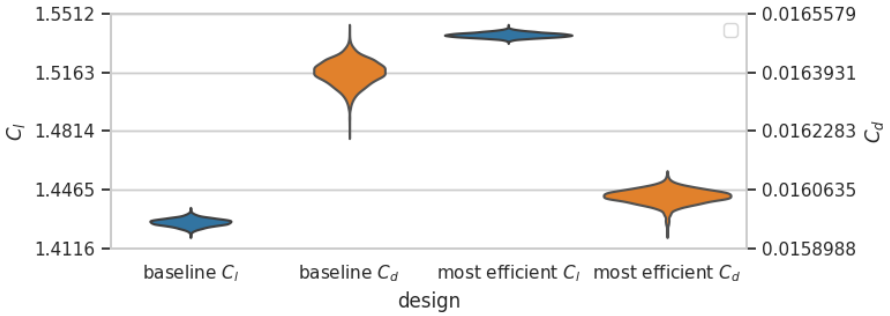

(b) Comparison of the baseline and most efficient design with uncertainty prediction

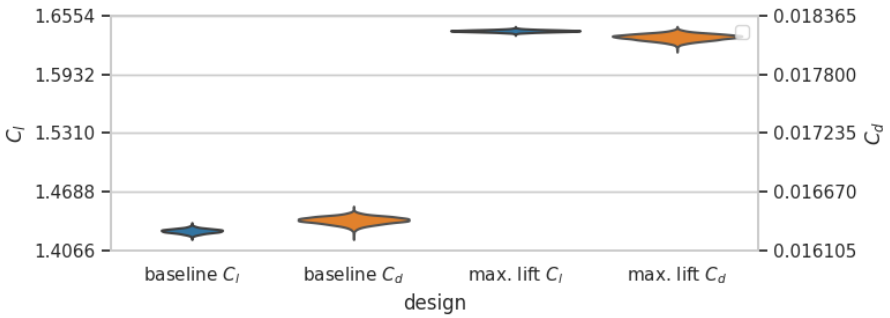

(c) Comparison of the baseline and maximum lift design with uncertainty prediction

Fig. 8: Comparison of the baseline and Pareto optimal designs with uncertainty prediction

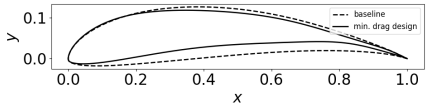

(a) Minimum drag design

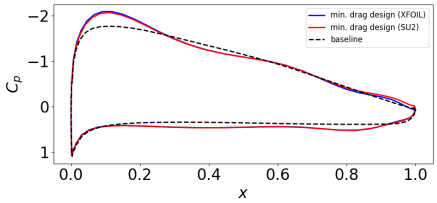

(b) C_p distribution of minimum drag design

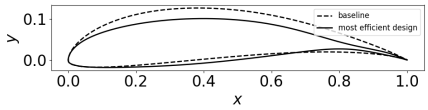

(c) Most efficient design

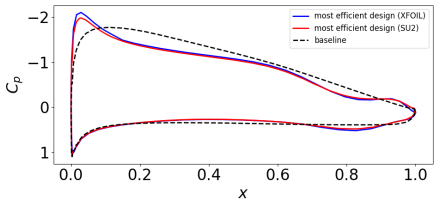

(d) C_p distribution of most efficient design

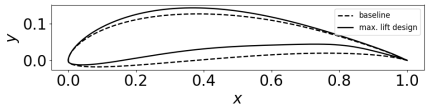

(e) Maximum lift design

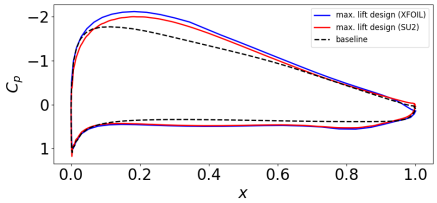

(f) C_p distribution of maximum lift design

Fig. 9: Comparison of Pareto optimal airfoil designs

8 Conclusion

The optimisation of an aerodynamic shape is computationally expensive - the more so, when uncertainties are taken into account. This problem can be tackled by employing multi-fidelity surrogate assisted optimisation. On the one hand, the required number of design evaluations is reduced by using statistical models which helps us to evaluate only promising design candidates. On the other hand, the required number of expensive high-fidelity design evaluations is reduced by employing a MF-GPR which can complement the information obtained from high-fidelity evaluations with low-fidelity information.

In this study, we performed an aerodynamic shape optimisation under uncertainty combining information from XFOIL and the RANS solver of SU2. The multi-fidelity surrogate assisted optimisation provided an accurate Pareto front approximation with only a limited number of high-fidelity RANS simulations. The optimal solutions found by our approach display higher reliability than the baseline solution.

References

1. Barrett, R., Ning, A.: Comparison of airfoil precomputational analysis methods for optimization of wind turbine blades. IEEE Transactions on Sustainable Energy **7**(3), 1081–1088 (2016)
2. Beyer, H.G., Sendhoff, B.: Robust optimization–a comprehensive survey. Computer methods in applied mechanics and engineering **196**(33-34), 3190–3218 (2007)
3. Deb, K.: Multi-objective optimization. In: Search methodologies, pp. 403–449. Springer (2014)
4. Drela, M.: Xfoil 6.9 user primer, xfoil_doc. txt. last updated 11/30/2001
5. Drela, M.: Xfoil: An analysis and design system for low reynolds number airfoils. In: Low Reynolds number aerodynamics, pp. 1–12. Springer (1989)
6. 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)
7. Emmerich, M.T., Giannakoglou, K.C., Naujoks, B.: Single-and multiobjective evolutionary optimization assisted by gaussian random field metamodels. IEEE Transactions on Evolutionary Computation **10**(4), 421–439 (2006)
8. Forrester, A.I., Sóbester, A., Keane, A.J.: Multi-fidelity optimization via surrogate modelling. Proceedings of the royal society a: mathematical, physical and engineering sciences **463**(2088), 3251–3269 (2007)
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. Menter, F.R., Kuntz, M., Langtry, R.: Ten years of industrial experience with the sst turbulence model. Turbulence, heat and mass transfer **4**(1), 625–632 (2003)
13. Poloni, C.: Hybrid ga for multi objective aerodynamic shape optimisation. Genetic Algorithms in Engineering and Computer Science (1995)
14. Poloni, C., Giurgevich, A., Onesti, L., Pediroda, V.: Hybridization of a multi-objective genetic algorithm, a neural network and a classical optimizer for a complex design problem in fluid dynamics. Computer Methods in Applied Mechanics and Engineering **186**(2-4), 403–420 (2000)
15. Poloni, C., Pediroda, V.: Ga coupled with computationally expensive simulations: tools to improve efficiency. Genetic Algorithms and Evolution Strategies in Engineering and Computer Science. Recent Advances and Industrial Applications pp. 267–288 (1997)
16. Quagliarella, D., Diez, M.: An open-source aerodynamic framework for benchmarking multi-fidelity methods. In: AIAA Aviation 2020 Forum. p. 3179 (2020)
17. Quagliarella, D., Serani, A., Diez, M., Pisaroni, M., Leyland, P., Montagliani, L., Iemma, U., Gaul, N.J., Shin, J., Wunsch, D., et al.: Benchmarking uncertainty quantification methods using the NACA 2412 airfoil with geometrical and operational uncertainties. In: AIAA Aviation 2019 Forum. p. 3555 (2019)
18. 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)

19. Schuëller, G.I., Jensen, H.A.: Computational methods in optimization considering uncertainties–an overview. Computer Methods in Applied Mechanics and Engineering **198**(1), 2–13 (2008)
20. Song, W., Keane, A.: A study of shape parameterisation methods for airfoil optimisation. In: 10th AIAA/ISSMO multidisciplinary analysis and optimization conference. p. 4482 (2004)
21. Tyrrell Rockafellar, R., Royset, J.O.: Engineering decisions under risk averseness. ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering **1**(2), 04015003 (2015)
22. Vaithiyanathasamy, R.: Cfd analysis of 2d and 3d airfoils using open source solver su2. University of Twente, Internship report (2017)