This is a preprint of the following article: Sabater, C., Bekemeyer, P. and Görtz, S. "Efficient Bilevel Surrogate Approach for Optimization Under Uncertainty of Shock Control Bumps". AIAA Journal 2020. The published article may differ from this preprint and is available at: https://doi.org/10.2514/1.J059480

An Efficient Bi-Level Surrogate Approach for Optimization under Uncertainty of Shock Control Bumps

Christian Sabater[†] Philipp Bekemeyer[‡] and Stefan Görtz[§]
German Aerospace Center (DLR), Institute of Aerodynamics and Flow Technology, 38108 Braunschweig, Germany

The assessment of uncertainties is essential in aerodynamic shape optimization problems to come up with configurations that are more robust against operational and geometrical uncertainties. However, exploring the stochastic design space significantly increases the computational cost. The aim of this paper is to develop a framework for efficient optimization under uncertainty by means of a bi-level surrogate approach and to apply it to the robust design of a retrofitted shock control bump over an airfoil. The framework combines a surrogate-based optimizer with an efficient surrogate-based approach for uncertainty quantification. The optimizer efficiently finds the global optimum of a given quantile of the quantity of interest through the combination of adaptive sampling and a moving trust region. At each iteration of the optimization, the surrogate-based uncertainty quantification uses an active infill criterion to accurately quantify the quantile requiring a reduced number of samples. Two different quantiles of the drag are chosen for the design of the shock control bump, the 95% to increase the robustness at off-design conditions, and the 50% for a configuration that is preferred for day to day operations. In both cases, the optimum bumps are more robust, compared to the one obtained through classical deterministic optimization.

Nomenclature

ξ	=	uncertainty parameters
C_D	=	drag coefficient
$C_{D,0}$	=	drag coefficient of initial configuration
$C_{D_{95\%}}$	=	95% quantile of the drag coefficient

Presented as Paper AIAA 2019-2214 at the AIAA Scitech 2019 Forum, San Diego, CA, 7-11 January 2019

[†]PhD student, $C^2A^2S^2E$, Lilienthalplatz 7, AIAA Member
[‡]Team Leader, $C^2A^2S^2E$, Lilienthalplatz 7, AIAA Member
[§]Department Head, $C^2A^2S^2E$, Lilienthalplatz 7, AIAA Member

$C_{D_{50\%}}$ = 50% quantile of the drag coefficient

C_L = lift coefficient

C_P = pressure coefficient

h_{bump} = bump maximum height in percentage of the chord

J = objective function

l_{bump} = bump length in percentage of the chord

M_0 = freestream Mach number

QoI = quantity of interest

t_{bump} = bump slope factor

PDF = Probability Density Function

X = design parameters

x_{flap} = flap starting location in percentage of the chord

$x_{h_{bump}}$ = bump maximum height location in percentage of the bump length

x_{start} = bump start location in percentage of the chord

I. Introduction

Traditionally, the use of simulation-based design optimization in aerodynamics has been carried out in a deterministic fashion, neglecting uncertainty [1, 2]. In this case both the design variables such as shape parameters and operating conditions such as Mach number and lift coefficient, are fixed in each iteration of the optimization process. However, the sensitivity of the final solution to small changes in the geometry due to manufacturing tolerances can affect the real performance of the aircraft [3]. In this case, tightening the manufacturing tolerances may not be a feasible solution due to an increase in the production cost. Moreover, fluctuations in the operating conditions such as the Mach number (velocity, air density, temperature) or lift coefficient (change in air-craft weight) cannot be avoided. Even a multi-point approach does not take into account these uncertainties, as it only deals with a discrete (and limited) set of flight conditions. The design point of an aircraft (cruise lift coefficient at a given Mach number) is used by engineers to simplify the design process as in practice, a deterministic optimum can lead towards a serious performance loss when accounting for the uncertainty. As a result, it is necessary to directly assess the effect of these relatively small unknown, aleatory (or irreducible) uncertainties in the design process [4].

Robust optimization techniques can cope with real-world aerodynamic shapes by accounting for these uncertainties. Therefore they prevent a configuration that does not match the specifications for which it was designed when slightly deviating from the design conditions [3]. However, when uncertainties are added to global optimization algorithms, the computational cost may increase by several orders of magnitude [5]. The objective function is changed from a

deterministic quantity of interest QoI, i.e. drag coefficient at a given flight condition, to a statistic of the quantity of interest, that is now a random variable. As a result, the stochastic space needs to be characterized at each iteration of the optimization [6]. This can be unfeasible if the underlying model is expensive to evaluate as it is the case of simulations involving computational fluid dynamics (CFD). In order to deal with this problem, it is necessary to reduce both the number of optimization iterations and the number of evaluations required for the uncertainty quantification (UQ). Surrogate approaches for uncertainty quantification are quite popular, as the sampling can be directly done in an approximation of the QoI at a significantly reduced cost [7]. Usually these approximations consists on Kriging [8, 9], Polynomial Chaos Expansion [10, 11] and stochastic collocation [12].

To reduce the number of optimization iterations in robust design, the use of a gradient-based optimizer together with an adjoint formulation [8, 13] can prove effective in aerodynamic shape optimization while accounting for uncertainties. However, an adjoint formulation is not always available when dealing with black box problems, and a gradient based approach will result in a local optima at best. Another possibility is the use of gradient-free methods such as Covariance Matrix Adaptation Evolutionary Strategies [14], Subplex [9] or surrogate models for robust single objective [11, 15] and multi-objective [16] problems. As in Uncertainty Quantification, the use of surrogate models such as Kriging can also reduce the number of iterations by approximating the design space. Another approach is the use of a global surrogate model containing both design parameters and uncertainties [17]. However, this approach can yield accuracy problems specially when uncertainties are at the same time also design parameters, or when the number of dimensions and design parameters increases.

From an engineering point of view, the interest lies in obtaining an accurate global robust solution at a reduced number of iterations (optimization) and function evaluations (uncertainty quantification). Traditionally, the objective function can be a combination of mean and standard deviation of the QoI [9, 18]. It is also possible to match the probability distribution function [19] or cumulative distribution function [20] to prescribed targets. Another alternative is the optimization of a given quantile of the QoI. This is a flexible approach in engineering problems [21] because it ensures a minimal performance with controlled probability and also does not assume any target distribution. Different quantiles can be chosen by designers as objective function to meet their specific design requirements.

The objective of this paper is the development of an efficient global Optimization Under Uncertainty framework involving expensive black box simulations and the application to the robust design of a shock control bump (SCB). The global optimizer follows a modified surrogate-based optimization (SBO) [22] approach that is designed to carefully balance exploration and exploitation. At each iteration of the optimization, a Surrogate-Based Uncertainty Quantification (SBUQ) approach with an active infill is used to efficiently obtain a given τ quantile of the quantity of interest.

The framework will be applied to the robust design of shock control bumps (SCBs). These are passive flow control devices that improve the performance of transonic wings by altering the flow around near-normal shock waves [23]. The flow is isentropically (gradually) decelerated with respect to the baseline configuration where no bump is present,

and wave drag is reduced. Shock Control bumps were first introduced in 1992 for the mitigation of wave drag [24]. Further studies took place in Europe within the EUROSHOCK II project [25] and in the USA [26] to investigate its full potential. From the 2000s, the focus has been on understanding the flow physics [27] and the realization of optimization studies [28], [29]. A complete overview of Shock Control Bumps is given in [30]. A crucial aspect that is gaining more attention recently is the need for a robust SCB configuration for industrial applications [31], [32].

The robustness of shock control bumps, i.e., its ability to effectively reduce wave drag at different flight conditions, is of primary concern as they are highly sensitive to the shock wave location [30]. At freestream velocities or lift coefficients different from their design point, SCBs suffer from adverse effects as the shock wave is not located in the designed location [23]. For example, when the shock wave moves upstream (due to a decrease in flight speed or lift coefficient), the flow is re-accelerated due to the curvature of the bump, leading to a second supersonic region behind the first shock. On the contrary, if the shock is located downstream of the SCB, no lambda shock structure is generated and the flow is further, leading to an increased shock strength and possibly causing flow separation. As a result, efficient robust optimization techniques can cope with the realistic design of SCBs by accounting for the uncertainties during the design stage.

First, Section II presents the bi-level surrogate framework for optimization under uncertainty. Section III introduces the deterministic and robust SCB design problem, the parametrization of the geometry as well as uncertainties and the chosen numerical model. In Section IV the results of both optimizations are presented, with emphasis in the comparison in the physics between the deterministic and robust optima configurations. Finally, conclusions are presented in Section V.

II. Optimization Under Uncertainty Framework

Global optimization problems require a large number of function evaluations to find the global minimum. This can be unfeasible if the underlying model is expensive to evaluate as it is the case when performing CFD simulations. In the case of optimization under uncertainty, the computational cost is several orders of magnitude more expensive, as a characterization of the stochastic space should be done at each iteration. Next, a framework for the efficient optimization under uncertainty applied to expensive black box problems is presented.

A. Problem Definition

The classical deterministic optimization problem is defined through the optimization of a given QoI $\in \mathbb{R}$ (performance measure, usually the drag coefficient), depending on the design variables $X \in \Omega$ at given operational conditions, A

$$J^*(X^*) = \min_X \text{QoI}\,(X, A) \tag{1}$$

Under the constraints:

$$c_1 : g(X) \leq 0 \tag{2}$$

The robust optimization involves the presence of uncertainties $\xi \in \Theta$ either in operational conditions, geometry, or both. As a result, the QoI becomes a random variable and the problem is shifted from the optimization of the QoI towards the optimization of an statistic of the QoI. In this case, a robust design approach is defined by minimizing a given τ quantile of the QoI, for some $\tau \in (0, 100)$.

$$J^*(X^*, \xi) = \min_X \text{QoI}_{\tau\%}(X, \xi) \tag{3}$$

A high quantile, such as the 95% or 99% can be chosen to minimize extreme events, instead of choosing a more conservative min-max optimization. In addition, by properly accounting for the uncertainties in the design process, confidence intervals can be used instead of conservative margins of safety. An alternative is to consider the median (50% quantile), to focus on the middle value of the realizations.

B. Bi-Level Surrogate-Based for Optimization under Uncertainty

The bi-level surrogate model is developed by combining a surrogate-based optimization framework as outer layer, with a surrogate-based uncertainty quantification framework in the inner layer, as shown in Fig. 1. The outer level (red surface) takes into account the optimization which depends only on the design variables X and has the statistic of the QoI as output (in this case the quantile, $\text{QoI}_{\tau\%}$). The inner level (blue surfaces) is responsible for the uncertainty quantification process of the QoI at a given design point X_i under uncertainties ξ. In aerodynamic shape optimization, the quantity of interest is usually the drag coefficient, and its statistic is a given τ quantile of the drag. The gradient-free framework is limited by the curse of dimensionality and is able to handle up to a moderate number of design parameters (up to 40 to 50) [33], primarily defined through the application of Kriging surrogate models during the Surrogate Based Optimization. Regarding the number of uncertainties, up to 15 to 20 uncertainties can be accurately handled [9]. The framework could also be adapted to handle mixed epistemic/aleatory uncertainties, as in the inner level an statistic of interval approach could take place [8], even though different infill criteria would be required.

This approach focuses on the refinement of the region where the robust solution is located. As a result, the efficiency over other optimization algorithms improved, as shown in Sections IV.A. Another advantage is the parallelization of the DoE stage in the UQ level. As the initial samples are defined beforehand, these can all be computed at the same time. In addition, if the black box solver provides a native parallelization, an additional level of parallelization is possible.

Next, each of the stages that form the bi-level surrogate approach are explained in detail.

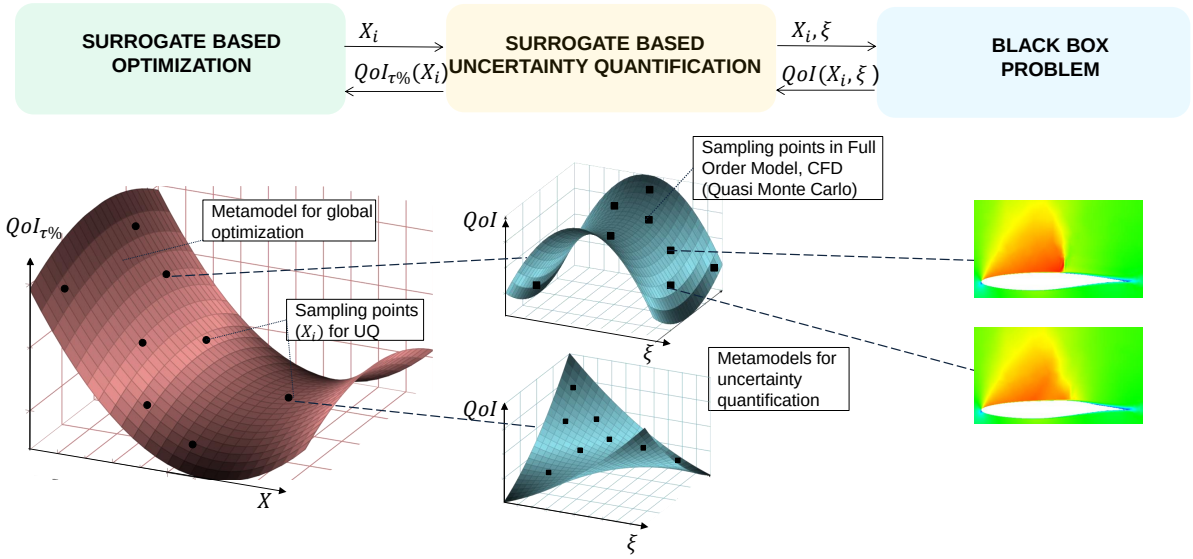

Fig. 1 Bi-level approach for uncertainty quantification; Surrogate of statistics (left); Surrogate of random variable (middle); Full order model evaluation (right)

C. Surrogate-Based Optimization

The outer layer of the bi-level approach is the surrogate-based optimization. The main objective is to make the optimization more efficient by reducing the number of function evaluations of the expensive black box problem [34]. A surrogate model that is cheaply evaluated is built following the parametrization and the quantity of interest. The Surrogate-Modelling for Aero-Data Toolbox (SMARTy) developed by DLR is used for the initial Design of Experiments sampling and for the creation of the Kriging surrogate [35].

The overall framework of the optimization process is shown in Fig. 2. In case of deterministic optimization, the objective function is the QoI, obtained by the black box solver. When dealing with optimization under uncertainty, the objective function is the quantile of the QoI that will be obtained through the uncertainty quantification process later introduced in Fig. 3. The framework has been tailored for the efficient application to aerodynamic shape optimization problems. Each of the steps in the optimization process is explained below.

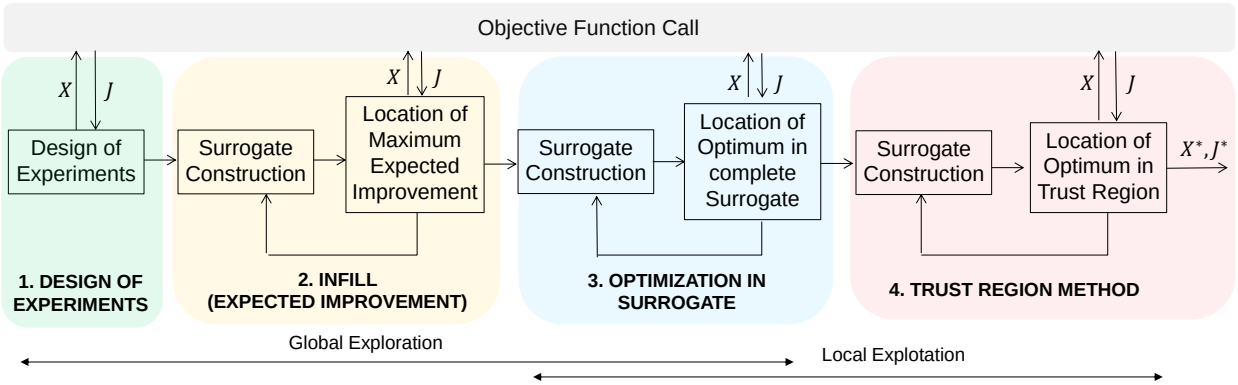

Fig. 2 Surrogate-based optimization framework to efficiently find global minimum of complex functions

1. Design of Experiments

One of the key elements for the construction of surrogate models is the sampling strategy for the evaluation of the objective function. Design of Experiments (DoE) aims to analyze the data efficiently by defining better sampling strategies. Arguably, the most common sampling approach is the use of Latin Hypercube Sampling developed by Mckay [36]. However, Quasi Monte Carlo techniques employing deterministic low discrepancy sequences [37] can prove more effective to obtain an even sampling distribution with a good space filling capability [38]. As a result it can better represent the design space with few sample points.

Sobol Sequences are a quasi-random low discrepancy sequence that use a base of two to successively create fined uniform partitions of the unit interval [39]. In contrast with Latin Hypercube Sampling, Sobol sampling is deterministic. Adding more samples is therefore straightforward by continuing the sequence without degradation in the space filling. As a result, when building surrogate models, it is possible to improve their accuracy by adding additional samples and reusing the existing points. As a rule of thumb, when approximating the objective function with a Gaussian process, the number of samples for an effective initial design of experiments should be about ten times the input dimension [40].

2. Infill through Expected Improvement

Once the initial sampling is evaluated, the surrogate model is constructed. The approximation of the objective function by Kriging, using Bayesian statistics, is among the most used surrogate models in optimization [34, 41]. A detailed explanation of the construction of Kriging surrogates is given by [42]. In particular, Universal Kriging with a Gaussian Kernel (exponent fixed to 2) is used. The hyperparameters of the correlation model are fine-tuned by the maximizing the maximum likelihood [43]. In order to obtain a global solution, this optimization is carried out applying differential evolution [44]. A regularization constant, also called nugget effect, is added to convert the interpolation into regression to better fit noisy data.

The Kriging surrogate provides the normal distribution prediction of the objective function at any given point (mean $\hat{y}$ and standard deviation $\hat{\sigma}(X)$). This enables the calculation of the probability of improvement from the best sampled solution y_{min}. The Expected Improvement (EI) is then defined as the product of the improvement between the potential better solution, $\hat{y}(X)$, and the current best one, y_{min} [22].

$$E[I(X)] = (y_{min} - \hat{y}(X))\,\Phi\left(\frac{y_{min} - \hat{y}(X)}{\hat{\sigma}(X)}\right) + \hat{\sigma}(X)\,\phi\left(\frac{y_{min} - \hat{y}(X)}{\hat{\sigma}(X)}\right) \tag{4}$$

where Φ and ϕ are the cumulative and probability distribution functions of the standard normal distribution, respectively. A large expected improvement is present in regions where a smaller solution than the current best is possible (first term of equation 4) and/or when the model error is large (second term of equation 4). Thus, EI balances exploration with exploitation.

During the refinement phase, the location with maximum EI in the surrogate is obtained by differential evolution and used as infill criteria to update the surrogate model. A population size ten times the number of design parameters is selected. Together with a conservative convergence criteria, the optimizer is able to obtain the global maximum Expected Improvement at relatively low cost. This process is iteratively performed until a predefined convergence criteria is met. The three possibilities are convergence in the objective function value y, in the L2 distance between consecutive locations the design space or in the EI value itself. In addition, this stage can also stop if the maximum number of allocated samples (computational budget) for that stage is met.

The EI criteria can be used in many cases as the best route to find the global minimum. However, in other occasions it may provide a very slow convergence rate [34]. As a result, two further stages are introduced to guarantee convergence.

3. Iterative Optimization in Surrogate

Once the refinement is finished, the acquisition function is changed from the Expected Improvement to the mean of the Kriging surrogate, and determined applying the same differential evolution optimizer as before. Then, the optimum point in the surrogate is validated and computed in the black box solver. It is unlikely that both surrogate and true solution will match, so the surrogate is recomputed with this infill point. The process is repeated until any of the convergence criteria previously explained is met.

4. Trust Region Method

In high dimensional, nonlinear functions (such as aerodynamic analysis which involve shock waves) the convergence of the previous stage is not always guaranteed. The trust region method is a local exploitation method used to find the local minimum in the region of influence of the current sampled minimum [45]. In order to converge towards a global minimum, a good balance between exploration and exploitation is required in the previous stage.

The optimum point in the surrogate is again found with differential evolution, but this time the search space is centered in the current real minimum, y_{min}, (sampled in the black box solver) and bounded to a given trust region, δ. The optimum location is validated in the Full Order Model, and the surrogate is updated. The next trust region δ_m is reduced if the surrogate is not accurate enough or increased otherwise. In case the QoI at the new evaluated point is lower than the current sampled minimum (a new optima is found), the center of the trust region is updated. Then, the process is repeated until convergence. Details of the implementation can be found in [34].

5. Additional Comments

The transition from one infill criteria to another takes place when one of the convergence criteria is met. However, it is recommended to allocate a maximum budget for each of the stages in case convergence can not be achieved. The heuristic follows from allocating 60% of the infill budget to the EI improvement stage, while 20% is used for both the

optimization in the surrogate and the trust region approach. The infill budget should consist of 2 to 4 times the number of DoE samples. Finally, the SBO is not limited to the presented infill criteria. The modularity of the framework makes possible the introduction of additional acquisition functions such as the probability of improvement or entropy-search techniques [46].

D. Surrogate-Based Uncertainty Quantification

The main problem in UQ is the propagation of the uncertainty through the expensive black box model [6]. For the determination of both statistical moments and quantiles of the QoI, the classical stochastic approach of Monte Carlo sampling is too expensive. The solution to this problem is the use of an approximation of the QoI through a surrogate model which is cheaper to evaluate using Monte Carlo [47] than the full order model. The surrogate-based uncertainty quantification (SBUQ) process follows the stages presented in Fig. 3. This framework balances exploration (global representation of the QoI) with exploitation (focus on the value of the QoI closer to the quantile).

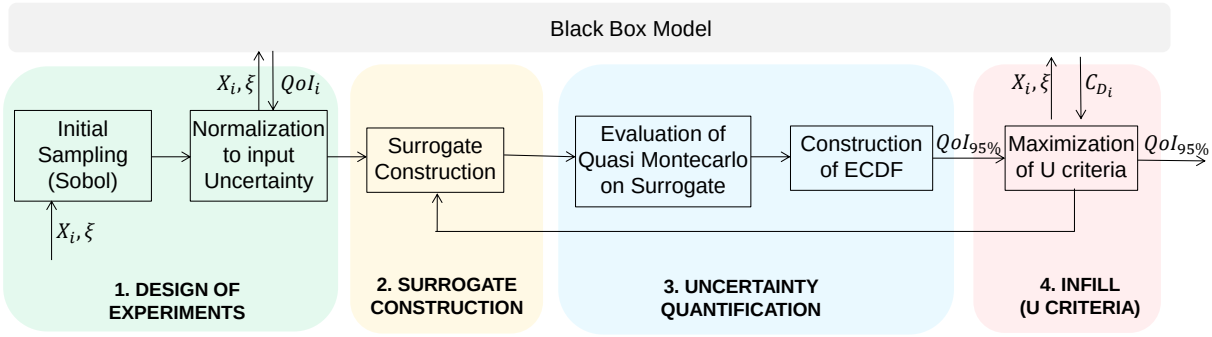

Fig. 3 Surrogate-based uncertainty quantification framework to efficiently characterize the quantile

1. Design of Experiments

Schillings [3] and Maruyama [9] have shown that for UQ, the accuracy of the surrogate can be improved if the initial sampling follows the distribution of the input uncertainties ξ. In the design of experiments, more samples should be placed along the mean, than along the tails of the input PDFs. Following this approach, when performing Monte Carlo on the surrogate with a large number of samples, the accuracy will be larger in the points that are evaluated more often, and lower in those regions that have less probability of being evaluated. As a result, for the problem at hand, an initial sampling based on Sobol sequences is scaled to the input distribution. Samples are computed by calling the black box model in a non-intrusive fashion.

2. Surrogate Construction

Kriging models can not only be used in optimization, but also to approximate the stochastic landscape of the QoI at a given design X_i. This is due to their good adaptation to nonlinear functions and the availability of the error estimator

[48]. Based on the initial sampling, the Kriging model is constructed following the same steps as decribed in section II.C.2.

3. Evaluation of statistics

Under uncertainty, the QoI is a random variable. The evaluation of a given statistic of the QoI in the surrogate is done through the direct integration of 1,000,000 Quasi Monte-Carlo samples. Although the surrogate provides a fast evaluation of the response, the convergence rate of Quasi Monte Carlo is faster than that of classical Monte Carlo. As a result, the sampling is predefined (derandomized algorithm) through Sobol Sequences [38] and the stochastic space is evaluated more efficiently. In this case, the statistic of the QoI that is sought any given τ quantile. It is the value below which the random variable is obtained with a probability of τ%. With a high quantile ($\tau = 95$%), this approach is preferred to classical min-max problems, as these are often too conservative [21].

To obtain the quantile, the Empirical Cumulative Distribution Function (ECDF) [49] is computed from the Quasi Monte Carlo sampling on the Kriging meta-model. This approach has been used in the past to evaluate the quantile of the drag in optimization under uncertainty problems involving airfoil design [21].

$$\mathrm{ECDF}_n(x) = \frac{\text{number of elements in sampling} \leq \mathrm{x}}{n} = \frac{1}{n}\sum_{i=1}^{n} I(x_i \leq x) \tag{5}$$

In equation 5 n is the number of samples evaluated through Quasi Monte Carlo, and I is the indicator function with two possible values: 1 if the event inside the brackets occurs and 0 if not:

$$I(x_i \leq x) = \begin{cases} 1, & \text{if} \quad x_i \leq x \\ 0, & \text{if} \quad x_i > x \end{cases} \tag{6}$$

By using the surrogate to evaluate the statistics, the number of samples is large enough and the ECDF converges to the Cumulative Distribution Function, CDF, according to the Glivenko-Cantelli theorem [50]. By numerically inverting the CDF it is possible to obtain any given τ quantile.

4. Active Infill Criterion

The active infill criterion follows the U-function [51] and equivalent probability of misclassification P_M [52]. At any given point, ξ, where the Kriging mean value is lower than the current quantile using the surrogate, $\hat{y}(\xi) < \hat{q}_\tau$, it is possible to obtain the probability of exceeding this quantile value, $P_M = P(\hat{y}(\xi) > \hat{q}_\tau)$. The opposite, the probability of overestimating the quantile given that the Kriging mean value is higher than the quantile, can also be obtained. Then, the probability of misclassification is defined as:

$$P_M = \min\left\{\Phi\left(\frac{\hat{y}(\xi) - \hat{q}_\tau}{\hat{\sigma}(\xi)}\right), \Phi\left(\frac{\hat{q}_\tau - \hat{y}(\xi)}{\hat{\sigma}(\xi)}\right)\right\} = \Phi\left(\frac{|\hat{q}_\tau - \hat{y}(\xi)|}{\hat{\sigma}(\xi)}\right) \tag{7}$$

Locations where the probability is maximum will have a strong influence on the predicted quantile by the surrogate, and new samples should be added in these regions [53] to increase the predictive accuracy of the quantile. The probability of misclassification is maximum when the fraction $\frac{|\hat{q}_\tau - \hat{y}(\xi)|}{\hat{\sigma}(\xi)}$ is minimum. As a result, it is equivalent to minimizing the so called U function:

$$U(\xi) = \frac{|\hat{q}_\tau - \hat{y}(\xi)|}{\hat{\sigma}(\xi)} \tag{8}$$

Then, the location of the new sample point is found by:

$$\xi^* = \text{argmax}\, P_M(\xi) = \text{argmin}\, U(\xi) \tag{9}$$

A new surrogate model is constructed with this additional sample and a new location is iteratively determined by equation 9 until the required level of accuracy is achieved. The infill criterion balances exploitation with exploration by sampling points closer to the calculated quantile or with a high error. There is a possibility that the proximity to the quantile masks the contribution of the standard deviation, and a multi-objective approach between numerator and denominator of equation 8 may be required. However, this behavior was not observed (as will be shown in Figure 12), and the differential evolution optimizer used to find the optimal infill sample is able to balance both terms equally. The quantile value is calculated with high accuracy at a relatively low cost due to the combination of the surrogate and active infill criterion. This approach is especially useful when dealing with high quantiles, which would otherwise require a large number of function evaluations to achieve the required accuracy.

The implications of equation 7 can be better understood in Figure 4. The thick dotted gray line represents the full model model dependent on the uncertainty ξ, that follows a uniform distribution. By propagating the uncertainty in the full order model, the reference 90% quantile is obtained at $QoI = -0.12$. Based on four samples (black dots), the Gaussian Process surrogate model is constructed (thick orange line), and the predicted quantile is obtained at $QoI = 0.26$. Also, the predicted regression error of the surrogate is shown up to three standard deviations. The probability of misclassification of the quantile is also shown for two locations. On the one hand, for location 1 (square), the surrogate error is maximum, but the probability of misclassification is not. In this case, there a probability of having a sample below the predicted quantile, and consequently, wrongly computing the quantile. This probability is outlined by the lower tail of the distribution in dark yellow. On the other hand, location 2 coincides with the intersection of the quantile with the surrogate. In this case the probability of misclassification is the largest (both upper or lower than the quantile) and a new sample should be added.

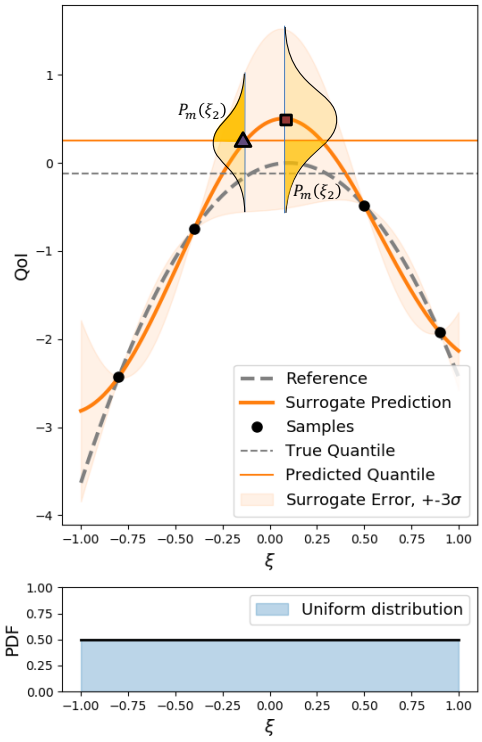

Fig. 4 Probability of misclassification of the 90% quantile for two different locations of the Gaussian Process.

III. Robust Design of Shock Control Bumps - Problem Definition

This section introduces the formulation of the deterministic and robust optimization of a shock control bump over an airfoil, the parametrization of the design parameters and uncertainties, the numerical model and the CFD process chain. The classical deterministic problem setup serves as verification of the optimization process at a single flight condition and will be compared with the robust solution.

Both the deterministic and the robust formulations aim to find the optimum bump shape to be retrofitted to the current supercritical RAE 2822 airfoil for the operation at dash condition: Mach number of 0.734 and constant lift coefficient of 0.789 (angle of attack of 2.77 degrees), which is defined by a higher freestream velocity than the design point. At this condition a strong near-normal shock wave is present over the upper surface and the addition of a shock control bump is expected to considerably reduce drag by weakening the shock intensity. The main objective is then to considerably improve the performance of the airfoil by modifying only the shape in the region near the shock wave.

A. Deterministic Optimization Problem

The deterministic global solution is defined through the classical drag optimization at constant lift coefficient, $C_L = 0.789$, and freestream Mach number, $M_0 = 0.734$:

$$J^*(X^*) = \min_X \frac{C_D\,(X, M_0, C_L)}{C_{D_0}} \tag{10}$$

Under the constraints:

$$c_1: \quad x_{start} + l_{bump} - x_{flap} \leq 0 \tag{11}$$

$$c_2: \quad C_L = C_{L_0} \tag{12}$$

The drag coefficient C_D depends on the design variables X that models the retrofitted bump. The drag is normalized by the initial drag of the baseline airfoil without bump, C_{D_0}. The first constraint guarantees that the end of the bump does not interfere with the flap (last 20% of the chord), following Figure 5. By defining the bump starting location x_{start} and length l_{bump}, it is not possible to fix the end through the bounds of these parameters. The constraint is introduced by a quadratic penalty function within the optimization in the surrogate. This approach allows the construction of a continuous surrogate in the design space (without penalty function), while avoiding the exploration of locations where the penalty is active. The constant lift coefficient constraint is handled internally by the CFD solver as it iteratively changes the angle of attack during the solution.

B. Robust Optimization Problem

The robust solution involves both operational and geometric uncertainties by taking into account the realistic stochastic conditions in flight. As a result, the drag coefficient (objective function) becomes a random variable. Two different cases are considered: a high quantile, $\tau = 95\%$, and the median, $\tau = 50\%$.

$$J^*(X^*, \xi) = \min_X \frac{C_{D_{\tau\%}}(X, \xi)}{C_{D_0}} \tag{13}$$

On the one hand, the optimization of the 95% quantile $C_{D,95\%}$ is equivalent to the minimization of the value below which 95% of the smallest realizations of drag occur due to the uncertainties. On the other hand, by optimizing the 50% quantile, $C_{D,50\%}$, the focus is in the middle value of the realizations, without an interest in extreme events. It is expected that two different configurations will arise. The designer will be responsible to come up with the best trade-off by selecting the proper quantile which is more suitable for his operations.

C. Parametrization of the Design Vector

The bump is defined as a local perturbation to the original airfoil by means of a Hicks-Henne Sinusoidal Function with five design parameters, following the parametrization of [54] and [55].

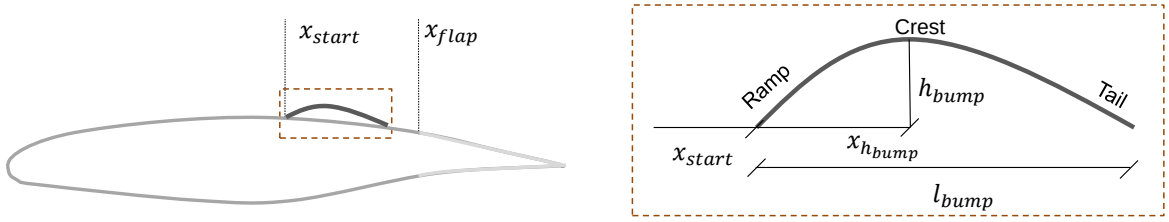

(a) Bump location relative to the airfoil **(b) Zoom in of the bump and design parameters**

Fig. 5 Parametrization of the shock control bump

$$z(x) = h_{bump}\left[\sin\left(\pi\left(\frac{x - x_{start}}{l_{bump}}\right)^{m}\right)\right]^{t_{bump}}, \qquad x_{start} \le x \le x_{start} + l_{bump} \tag{14}$$

where m is used to modify the asymmetry of the bump:

$$m = \frac{\log(0.5)}{\log\left(x_{h_{bump}}\right)} \tag{15}$$

The slope of the curvature of this function is zero at both ends. As a result, no discontinuities are present between the bump and the airfoil. A second advantage with respect to other parametrization such as NURBS curves is that the Hicks-Henne parameters represent the physical geometry of the bump, enabling a straightforward definition of the design space as shown in Figure 5 and Table 1. Note that the parameters are normalized according to the airfoil chord length, c.

Table 1 Design parameters of the shock control bump

Parameter	Description	Lower Bound	Upper Bound
h_{bump}	Bump maximum height	0	0.015
$x_{h_{bump}}$	Bump maximum height location (as a function of the bump length)	0.4	0.85
l_{bump}	Bump length	0.15	0.45
x_{start}	Bump starting location	0.3	0.54
t_{bump}	Bump slope factor	0.2	2

The bounds of the parameters have been chosen according to physical constraints in order to reach a realistic design. For example, a bump height of 0 is equivalent to the baseline configuration of an airfoil with no bump. The upper limit of 0.015 is chosen to account for the correlation of boundary layer thickness and bump height. The bump is centered along the shock wave location of the initial airfoil, around 55% of the chord. The design parameters X are normalized between 0 and 1, equivalent to the lower and upper bounds in Table 1.

D. Parametrization of Uncertainties

Table 2 shows the operational and geometrical uncertainties defined in this case. Mach number and lift coefficient strongly affect the shock wave location, while the SCB height is the most important geometrical parameter influencing the smearing effect of the shock wave [24].

Each uncertainty follows a beta distribution with the mean μ_ξ, according to nominal conditions, and standard deviation σ_ξ according to Table 2. A main advantage of beta distributions for UQ is that they are bounded (truncated PDF) [56] and thus represent a more realistic representation of the uncertainties than unbounded normal distributions. Also, beta distributions are able to adapt to both normal and uniform distributions by modifying their shape parameters α_1 and α_2. In this case the distribution is required to be symmetric ($\alpha_1 = \alpha_2 = 5$) to resemble a truncated normal distribution. Then, the location β_1 and scale β_2 are calculated to have the required mean μ and standard deviation σ.

Table 2 Definition of uncertainties in the robust optimization problem

Parameter	Description	Mean	Standard Deviation	Minimum Value	Maximum Value
$M_{0,\xi}$	Freestream Mach number	0.734	0.0045	0.719	0.749
$C_{L_{0,\xi}}$	Lift coefficient	0.791	0.0045	0.776	0.805
$h_{bump,\xi}$	Bump height	0	0.008	-0.0262	0.262

$$\text{Beta}\,(x, \alpha_1, \alpha_2, \beta_1, \beta_2) = \frac{\gamma\,(\alpha_1 + \alpha_2)\left(\frac{x-\beta_1}{\beta_2}\right)^{(\alpha_1-1)}\left(1-\frac{x-\beta_1}{\beta_2}\right)^{(\alpha_2-1)}}{\gamma(\alpha_1)\gamma(\alpha_2)} \tag{16}$$

E. Numerical Model

In order to obtain the aerodynamic performance of the airfoil, the Reynolds Average Navier Stokes (RANS) equations are solved by using the DLR TAU solver [57] with the Spalart-Allmaras turbulent model. In particular a 3v multigrid cycle, lower/upper symmetric Gauss-Seidel implicit method for time integration in backward Euler solver and a central flux discretization is chosen. The converge density residual is fixed to 1e-8. The unstructured mesh is hybrid and quasi two-dimensional, with tetrahedral and prism elements, as shown in Fig. 6. The number of grid nodes is 29,000, with 380 surface nodes around the airfoil. The domain is decomposed in four regions following the TAU built-in partitioner for the efficient parallelization of the solution. The mesh is validated with experimental data from the RAE2822 [58] as shown in Figure 7. The simulation is run at Mach number 0.792 and angle of attach of 2.31 degrees, following the validation example of [59], and obtaining almost identical results. When an additional wind tunnel correction is added with Mach number of 0.732, the exact location of the shock wave is found. The mesh is suitable to predict the physical transonic flow field characteristics.

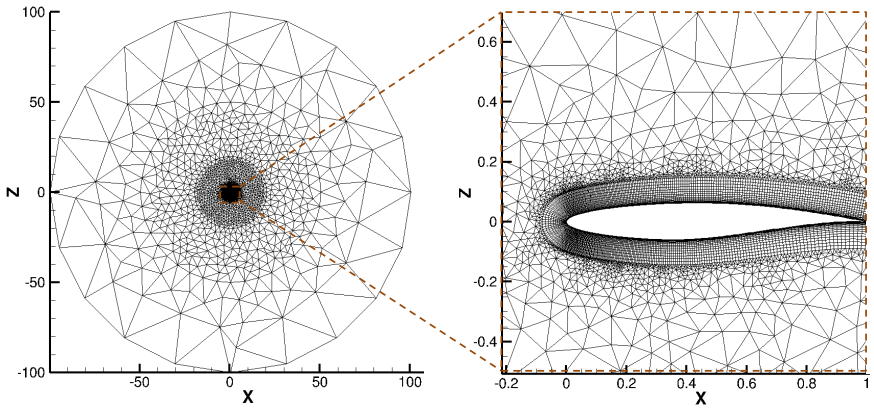

Fig. 6 CFD Grid used in the Optimization Problem

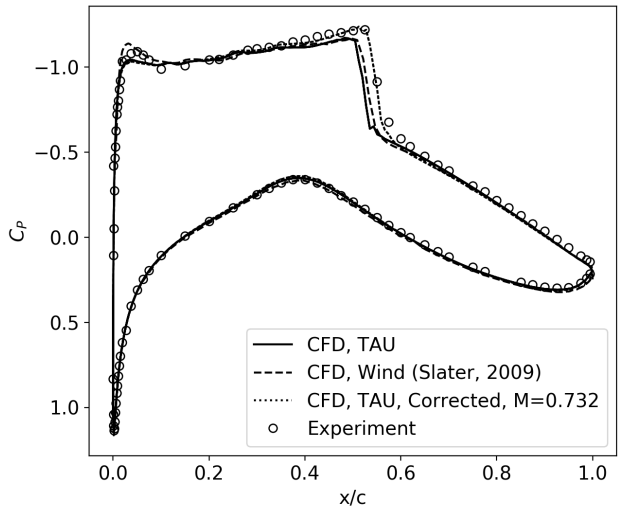

Fig. 7 Validation of numerical grid with experimental data

F. Process Chain

Every time the drag coefficient is required from a given design vector X_i, (for deterministic optimization), or from a given design vector X_i and uncertainties ξ, (for the UQ in robust design), the process chain of Fig. 8 is active. The geometry is changed through mesh deformation tool based on radial basis functions developed by DLR [60]. After the addition of the boundary conditions, the RANS equations are solved. The drag coefficient is obtained by the integration of the pressure and viscous forces along the surface of the airfoil. The process chain relies on the Flow Simulator framework [61] and is able to handle convergence errors, restarts and parallel execution.

IV. Robust Design of Shock Control Bumps - Results

This section presents the results for the deterministic and the robust optimizations of the SCB retrofitted to the RAE2822 using the framework previously described. It also provides the validation of the statistics of the quantity of

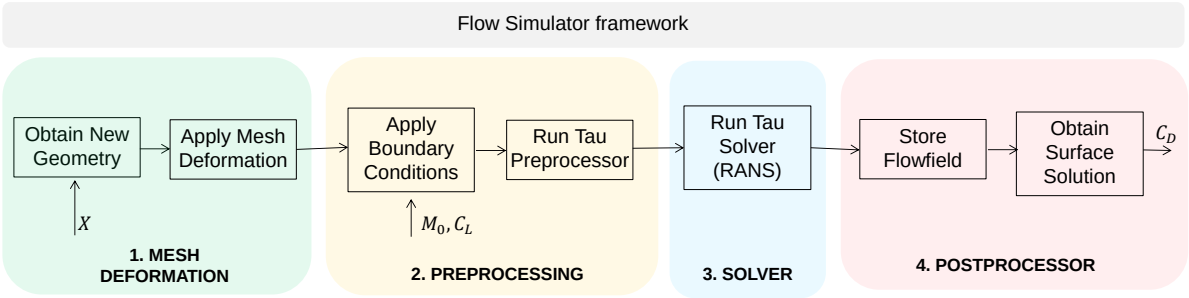

Fig. 8 Framework for the evaluation of the objective function through CFD, from the design vector and operating conditions to the drag coefficient

interest by comparing the surrogate-based UQ with direct integration.

A. Deterministic Optimization

The deterministic optimization problem in equation 10 has been computed by employing three optimizers, with the objective of assessing the global accuracy and the time savings of the surrogate-based optimization approach:

1) A Differential evolution algorithm [44] is used as a reference with an initial population size of 30 individuals (6 times the number of design parameters). It is the same algorithm as the one used for obtaining the minimum value on the surrogate in the SBO approach, but this time the full order model is directly evaluated.
2) Subplex, a gradient-free method of Rowan [62], which effectively decomposes the design space into low dimensional subspaces and searches for the convex hull. It is more efficient than the Simplex method since it scales linearly with the dimensionality and can be applied to noisy functions. It is dependent on the initial starting point of the optimization. Two convergence criteria based on the difference in the objective function and on the difference in the design vector between two consecutive iterations is defined.
3) The SBO method, which has been presented in depth in section II.C. For this case, a total of 50 samples are defined in the initial DoE.

Figure 9 shows the optimization history of the normalized drag coefficient for the three approaches against the optimization iteration number. It is assumed that the optimum found through differential evolution, after 1140 function evaluations, is the global one. The SBO approach is able to find an optimum solution very close to this global optimum at a tenth of a cost. Also, SBO finds a better optimum than the Subplex method with half the iterations. As a result, the SBO method is able to explore the design landscape very efficiently.

The results of all three optimization runs are summarized in Table 3 in terms of optimum geometry, drag reduction and required number of iterations. For all, the optimum bump extends to maximum length about 80% of the chord of the airfoil, which is close to the hinge point of the hypothetical flap defined in the constraint.

As shown by the pressure coefficient distribution along the airfoil surface in Figure 10 a, the bump reduces the wave

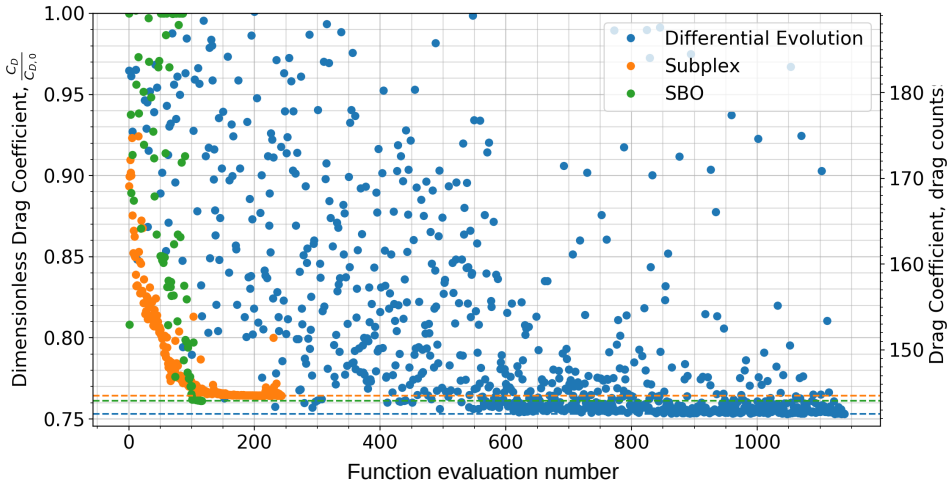

Fig. 9 Convergence history of the drag coefficient for deterministic optimization

Table 3 Deterministic optimization results

Optimizer	Optimum Geometry					Drag Reduction		Function
	h_{bump}	$x_{h_{bump}}$	l_{bump}	x_{start}	t_{bump}	Percentage	Drag counts	Evaluations
Diff. Evolution	0.009	0.644	0.45	0.35	1.282	24.694%	46.74	1140
Subplex	0.008	0.555	0.374	0.426	1.237	23.56%	44.59	244
SBO	0.0077	0.598	0.402	0.398	1.111	23.87%	45.18	116

drag by replacing the near-normal shock wave of the baseline configuration with an isentropic compression wave. The bump influences the pressure only locally. The mitigation of the shock wave is also shown in the pressure coefficient contour of Figure 10 b, in which the lambda shock can be appreciated. The total drag reduction is almost 25%, as the shock control bump is efficient in smearing the strong initial shock wave.

B. A Priori Validation of Uncertainty Quantification

Following Figure 3, for an accurate estimation of a high quantile using the surrogate, it is necessary to determine an appropriate number of samples for the DoE and infill stages. Before the robust optimization, the uncertainty quantification framework is validated on the deterministic optimum shape obtained with SBO under operational and geometrical uncertainties. The reference 95% quantile of the drag, $C_{D_{95\%}}$, is obtained with 10,000 Quasi Monte Carlo realizations of the full order model (CFD). Figure 11 shows the convergence history of the surrogate based uncertainty quantification using different sampling strategies. For each convergence, there is a difference in the proportions between the DoE (exploration) and infill (exploitation) samples. The required accuracy is set to 0.3 drag counts. Approximately, a reduction in one drag count enables an increase in payload of 90kg (1 more passenger) [63], and is in line with the minimum improvement in the performance required for a designer in order to choose one configuration over the other. By setting the requirement to 0.3 drag counts a good trade-off between accuracy and efficiency based on the physics can

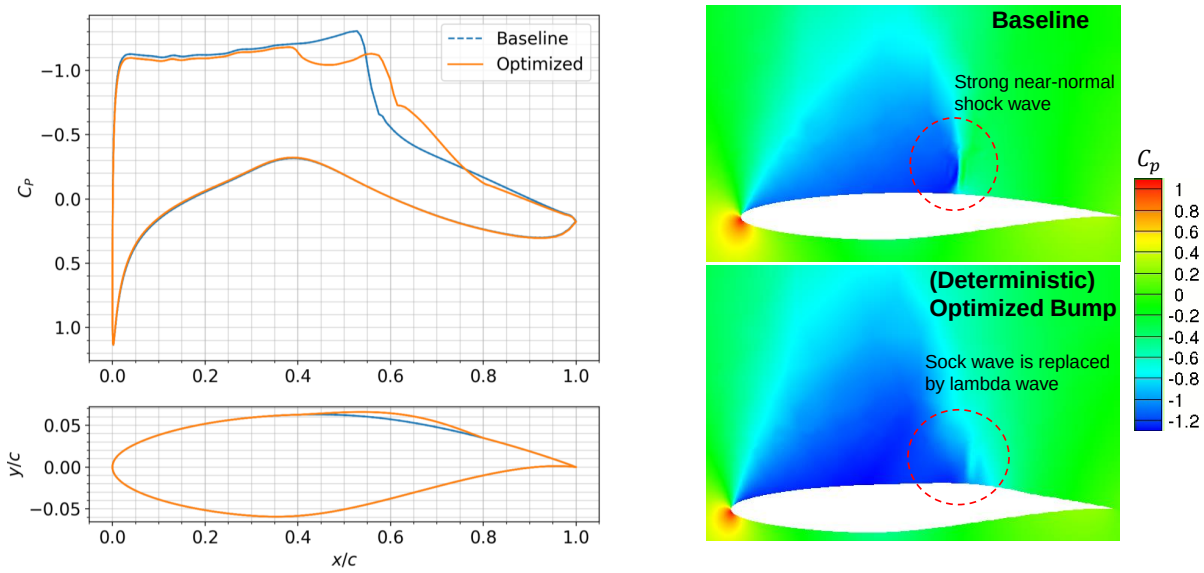

(a) Surface pressure coefficient **(b) Field pressure coefficient**

Fig. 10 Comparison between baseline and deterministic optimum configurations

be established.

As expected, the error in the 95% quantile reduces if the number of samples used for the construction of the surrogate increases. However, the use of training samples following only a global DoE approach (pink line) is not enough to reduce the quantile error to the defined level of accuracy. Using the active infill criterion further reduces the error, but there is clear trade-off between exploration and exploitation of the stochastic space. Without an accurate global representation of the stochastic space provided by the DoE, the initial location of the quantile can not be well estimated.

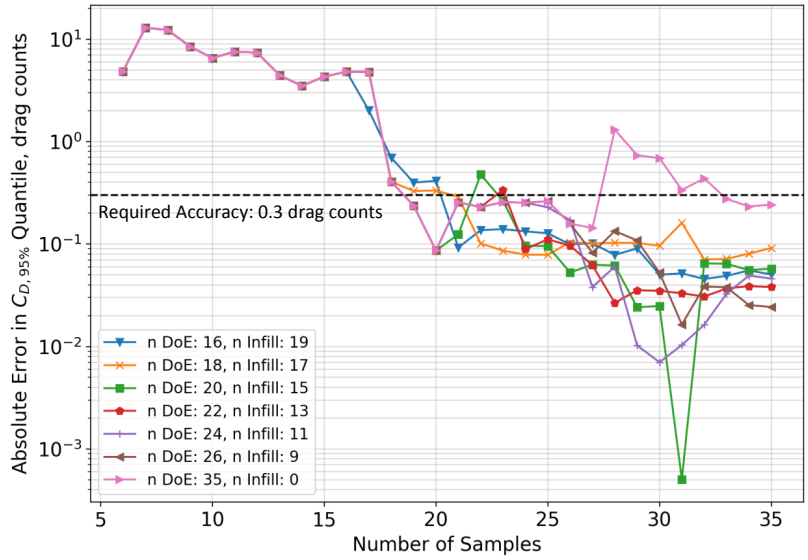

Fig. 11 Convergence error of drag quantile, $C_{D_{95\%}}$, for deterministic optimum configuration under different sampling strategies. Reference solution obtained with 10,000 Quasi Monte Carlo realizations

Based on this study, at each iteration of the optimization in the outer level, a surrogate is constructed in the stochastic space with 24 DoE samples and 5 additional infill samples following the U criterion. Alternatively, a reduced number of samples such as 22 DoE and 3 infill samples could have also been chosen and met our requirements. However, with the focus on accuracy and convergence stability for all UQ surrogate models, it was decided to proceed with a total number of 29 samples. This guarantees a high accuracy in the estimation of the statistic, which is critical specially when looking at high quantiles. The resulting uncertainty space with this sampling strategy is shown in Figure 12. The drag coefficient is represented as a function of the uncertain lift coefficient and Mach number, for the nominal bump height. After drawing Monte Carlo samples based on the surrogate (small black dots), the 95% quantile of the drag is obtained (black line). The quantile divides the 5% of the realizations with highest drag, while keeping below the 95% remaining with lower drag. Note how the infill samples are distributed along this quantile, while the DoE samples are focused on the locations where the realizations are most common.

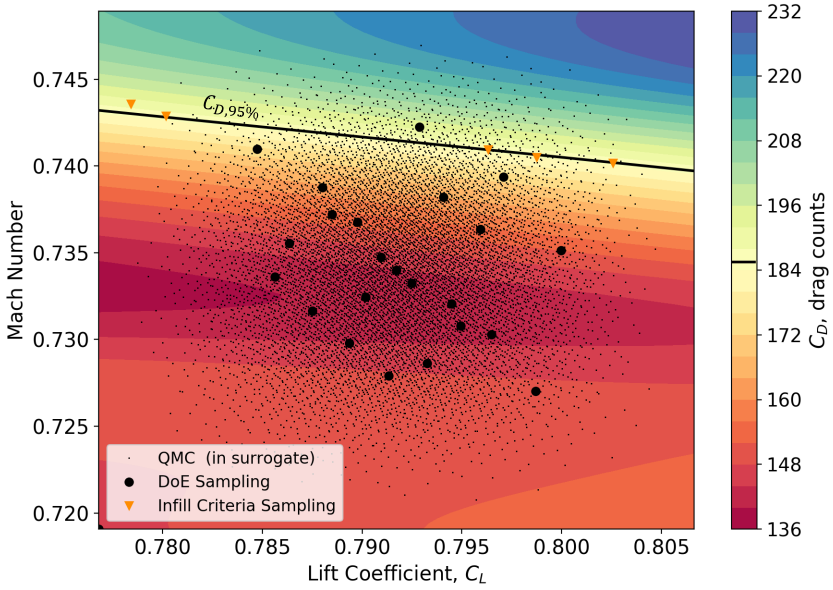

Fig. 12 Drag field as a function of stochastic operating conditions and zero uncertainty in bump height for the deterministic optimum configuration

C. Robust Solutions

Once the capabilities of the framework have been demonstrated in terms of the deterministic optimization and the uncertainty quantification, the presented optimization under uncertainty of the shock control bump is performed through the bi-level surrogate approach. Two different optimum configurations, one in terms of the 50% quantile of the drag coefficient and another in terms of the 95% quantile are found respectively. Also, the optima configuration for a worst-case configuration (100% quantile), is obtained.

Figure 13 shows the convergence history of the robust optimization. The objective function, the 95% quantile of the

drag coefficient, is represented in black squares, while the other colors represent the five normalized design variables at each iteration. These range from one, the upper limit in table 1, to zero, the lower limit. Highlighted markers represent the current optimum configuration until a better one is found. After the initial DoE and EI phases, the optimum solution stabilizes, and is further refined through the trust region approach. The total number of CFD evaluations is 3219 (the stochastic space is analyzed 111 times with 29 CFD evaluations in each iteration for the surrogate-based uncertainty quantification). The parallelization strategy considerably reduces the computational time.

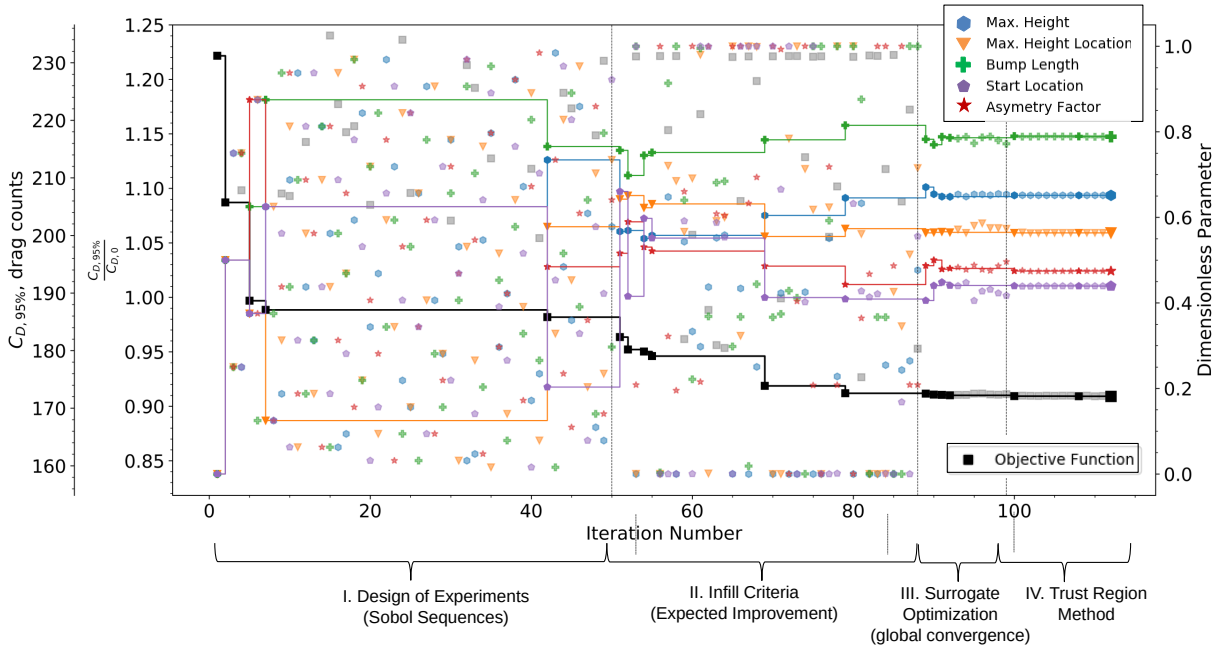

Fig. 13 Convergence history of robust optimization. $C_{D_{95\%}}$ and design parameters vs iteration number

Figure 14 shows the histograms of the two different robust optimum, the deterministic optimum, the worst-case optima and the baseline airfoil without bump. The 50% and 95% quantiles are also shown for each geometry. The optimum robust bump geometry which focuses on the 95% quantile (blue histogram) provides a decrease in the 95% quantile of the drag by 25.5% (58.9 drag counts) with respect to the baseline RAE2822 airfoil where the bump is not present (red histogram). This configuration is able to deal better with the uncertainties under off-design conditions making it also more robust than the bump shape obtained in the deterministic optimization (orange histogram). In particular, it reduces the 95% quantile by 6.7% (12.4 drag counts) compared to its deterministic counterpart due to the increase in bump height. By increasing the bump height, the SCB is more effective alleviating strong shock waves at high velocities and lift coefficients. As a result, its histogram features the shortest tail.

The main drawback of the robust shock control bump with focus in the 95% quantile of drag is the higher mean drag compared to the deterministic optimum configuration, since the high bump is not as effective at alleviating the drag at nominal operating conditions. To deal with this problem, the configuration with the optimum 50% quantile

displaces most of the weight of its PDF to the left. As expected, it is the configuration with the lowest median (2.5 drag counts lower than the deterministic configuration) and mean. This comes at the expense of a longer tail, and worst off-design conditions. This configuration has the worst 95% quantile (excluding the baseline configuration). However, it is an appealing design when the dealing with day to day operational events and off design conditions are not critical for performance. Finally, the optima configuration for worst-case scenario, corresponding to the 100% quantile, is also shown for comparison. This configuration is able to effectively mitigate high drag situations (strongest shock waves). Despite having the shortest tail and lowest variability, it is outperformed by the 95% and 50% robust configurations in terms of average performance and desirability from a practical point of view. As a result, optimizing for worst-case leads to very conservative configurations, and choosing instead a high quantile such as the 95% is a more balanced option.

The deterministic solution presents (as with the optimum 50% quantile configuration), a strong positive skewness. The deterministic shock control bump is very effective around the deterministic design point (Mach number, lift coefficient and shape constant). When geometrical and operational uncertainties are present, the performance is strongly deteriorated from these nominal conditions even at lower lift coefficients or freestream Mach numbers, where the drag is traditionally reduced if no bump is present.

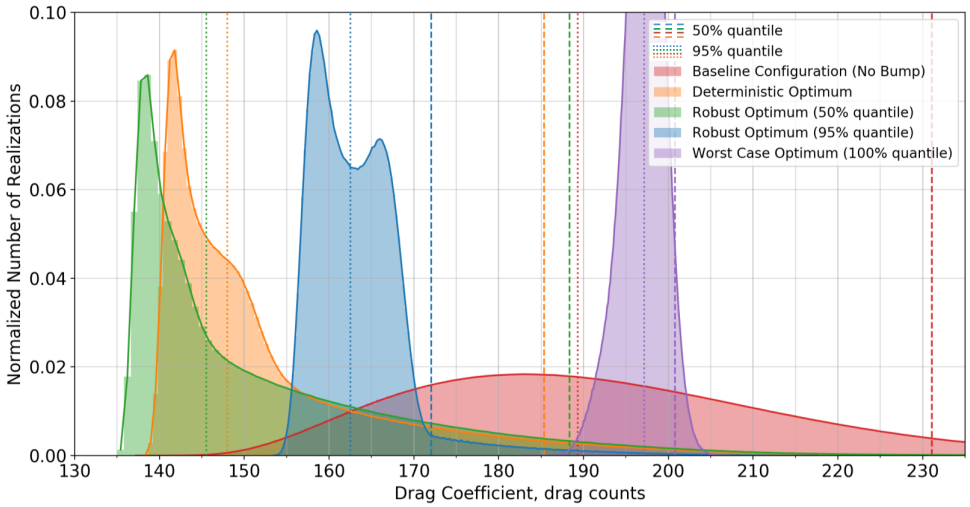

Fig. 14 Histograms of the uncertain drag coefficient for configurations at hand: 95% robust optimum, 50% robust optimum, worst case optima, deterministic optimum and baseline configurations

A complete overview of both robust configurations is shown in Table 4 in terms of design vector and quantiles. For comparison, the characteristics of the baseline and deterministic optimum configurations are also shown. In Figure 15 the two optimum bump geometries are represented. The optimum bump with focus on the 95% quantile is the highest and is located the furthest downstream.

To physically investigate the results of Fig 14, 10,000 stochastic realizations for each of the optimum configurations have been obtained in the full order model following Quasi Monte Carlo. The pressure coefficient of 100 of these

Table 4 Robust Optimization Results

Configuration	Optimum Geometry					Quantile (drag counts)	
	h_{bump}	$x_{h_{bump}}$	l_{bump}	x_{start}	t_{bump}	$C_{D_{95\%}}$	$C_{D_{50\%}}$
Baseline	0	0	0	0	0	231.1	189.3
Deterministic Optimum	0.00772	0.598	0.402	0.398	1.11	185.4	148
95% Optimum	0.00977	0.654	0.386	0.414	0.991	172.1	162.5
50% Optimum	0.00726	0.619	0.432	0.369	1.129	188.5	145.5

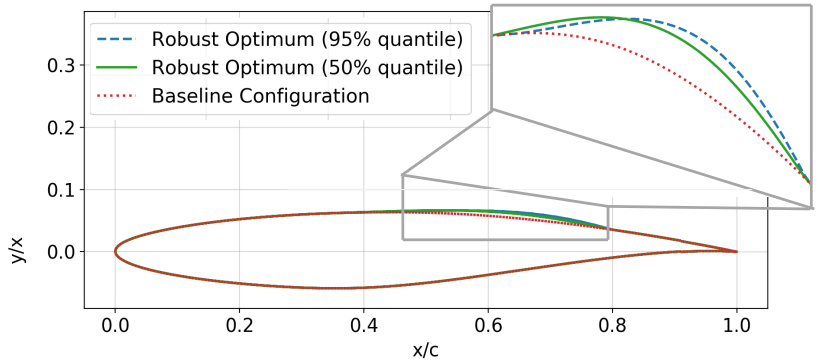

Fig. 15 Geometries of baseline and robust optimum configurations

realizations along the airfoil surface is shown in Fig 16, represented by grey lines. This is done for the robust optimum under the 50% quantile (left) and the robust optimum for the 95% quantile (right). Locations with the largest concentration of grey lines represent the most common scenarios. The thick lines represent those realizations (from the 10,000 initial samples) equivalent to the 5%, 50% and 95% quantile.

On the one hand, it can be deduced from the pressure field that the robust configuration aiming at the 50% quantile has the shock wave in the ideal location during most of the time. This is in line with the positive skewness of its histogram and its low mean drag. However, at off-design conditions, the performance rapidly deteriorates. In its 95% realization a strong normal shock wave is present (black line). As expected, this shock wave is stronger than the one from the robust optimum focused on off-design conditions. This explains its large 95% quantile value and longer tail of the histogram.

On the other hand, the robust bump with focus on the 95% quantile presents a weaker normal shock wave at off-design conditions motivated by the increase in bump height. In addition, the bump is located further downstream and is shorter, increasing its effectivity at higher Mach numbers. However, as previously discussed, this increase in height deteriorates the behaviour around the 50% quantile. In this case, at medium to low Mach numbers and lift coefficients, an initial normal shock wave is followed by an expansion wave as the flow is further accelerated along the bump. This leads to a secondary normal shock wave. As a result, the shock wave location is longitudinally more

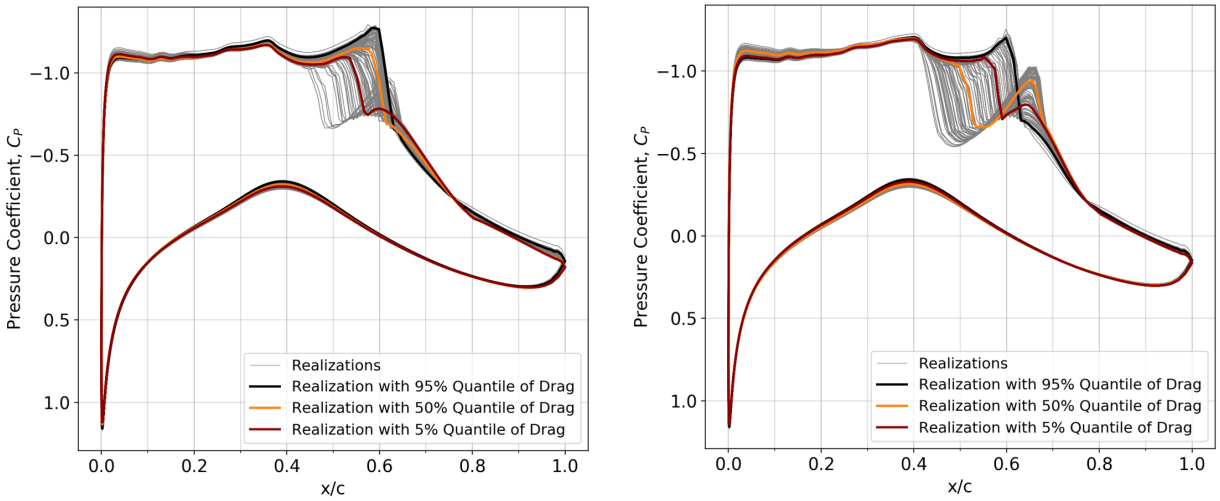

(a) Optimum robust configuration in $C_{D,50\%}$

(b) Optimum robust configuration in $C_{D_{95\%}}$

Fig. 16 Realizations of the pressure coefficient due to uncertainty. Highlighted the realizations corresponding to 5%, 50% and 95% quantile of drag

uniformly distributed with respect to the first configuration. With an increase in Mach number and lift coefficient, the initial shock wave moves forward and the expansion and secondary shock wave decrease until disappearing at around the 95% quantile. This configuration presents a good trade-off between nominal and off design conditions.

The pressure coefficient field for the 50% and 95% realizations, as well as the mean and standard deviation are shown in Figure 17 for the baseline airfoil and two robust optimum configurations. As expected and previously explained, the robust configuration with focus in the 95% quantile is most effective in smearing the shock wave in the 95% realization. The mean pressure field has no physical meaning, but helps in understanding the average behaviour under uncertainties. For example, the optimum configuration with focus in the 50% quantile has a smeared shockwave (softest transition from negative to positive pressure coefficients) during most of the realizations. In opposition, the robust configuration with focus in the 95% quantile shows a secondary shock wave in order to better account for off design conditions. Finally, the baseline configuration always presents a strong shock wave.

The standard deviation pressure field shows that most of the uncertainty occurs around the foot of the SCB, at the location of the main shock wave. The longitudinal displacement of the shock wave due to the uncertainties is largest in the optimum configuration with focus on the 95% quantile. However, this longitudinal displacement makes the quantile more robust w.r.t. off design conditions, specially compared to the baseline configuration. Looking at the baseline airfoil with no bump, the longitudinal displacement of the normal shock wave is the shortest, but with the largest intensity (stronger shock wave).

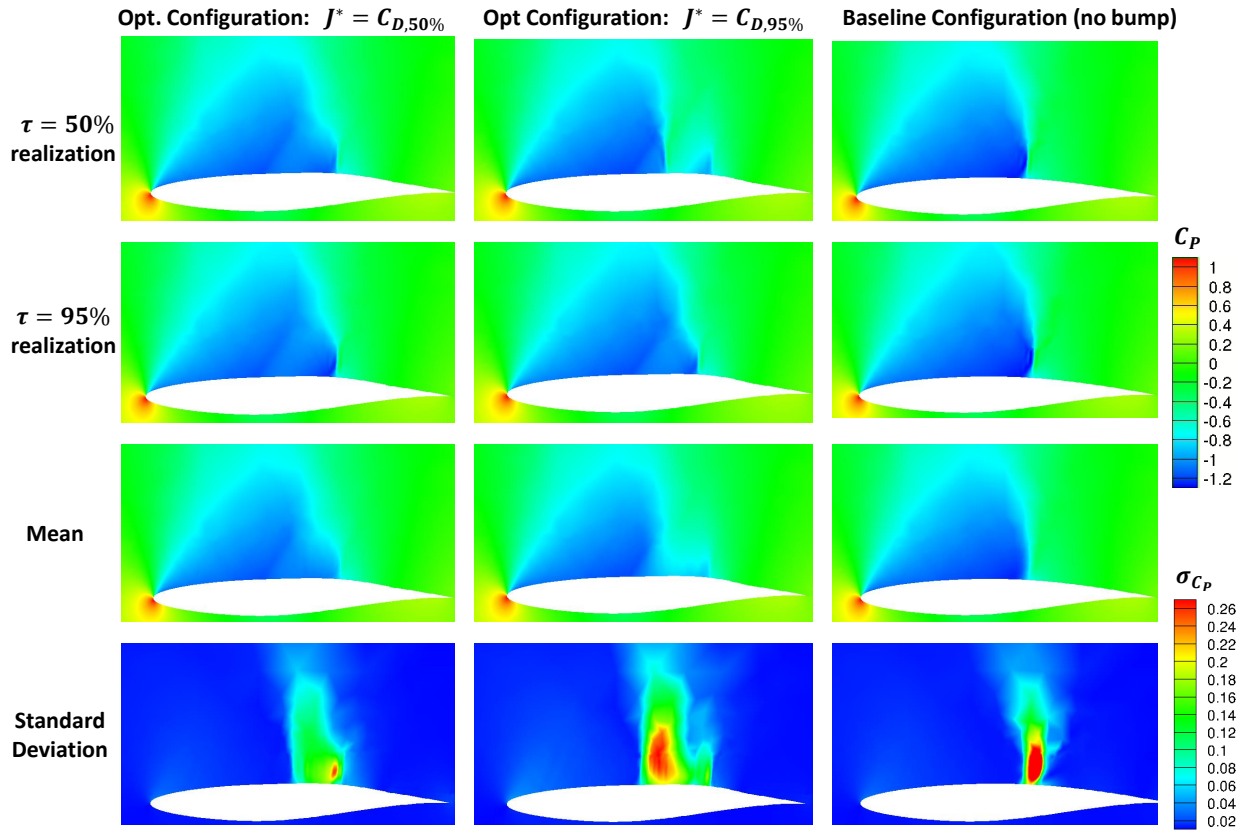

Fig. 17 Pressure coefficient field for robust and baseline configurations. 50% and 95% realizations, mean value and standard deviation value

D. A Posteriori Validation of the Uncertainty Quantification

The optimum results were obtained by the surrogate-based uncertainty quantification framework where only 29 CFD samples were used to estimate the quantile of the QoI in each iteration. After the optimization, a posteriori validation of the optimum statistics is required to assess the capabilities of the model. The reference quantile of each optimum configuration is obtained by computing 10,000 realizations through CFD following Quasi Monte Carlo sampling. Table 5 shows the absolute and relative error of the quantile of the QoI for the optimum configurations between the surrogate and reference. Overall, it has been found that using the proposed approach, the error in the quantile obtained with the surrogate is smaller than the initial requirement of 0.3 drag counts.

As shown in Fig. 18, the reference Empirical Cumulative Distribution Function, obtained with 10,000 samples, is equivalent to the one obtained with the surrogate built with 29 training samples. When no surrogate is present, the ECDF can be obtained from direct integration with equivalent 29 samples. In this case, a large error in the quantiles is present. As a result, the use of the surrogate model together with the infill criteria enables an accurate estimation of the quantile at a low computational cost.

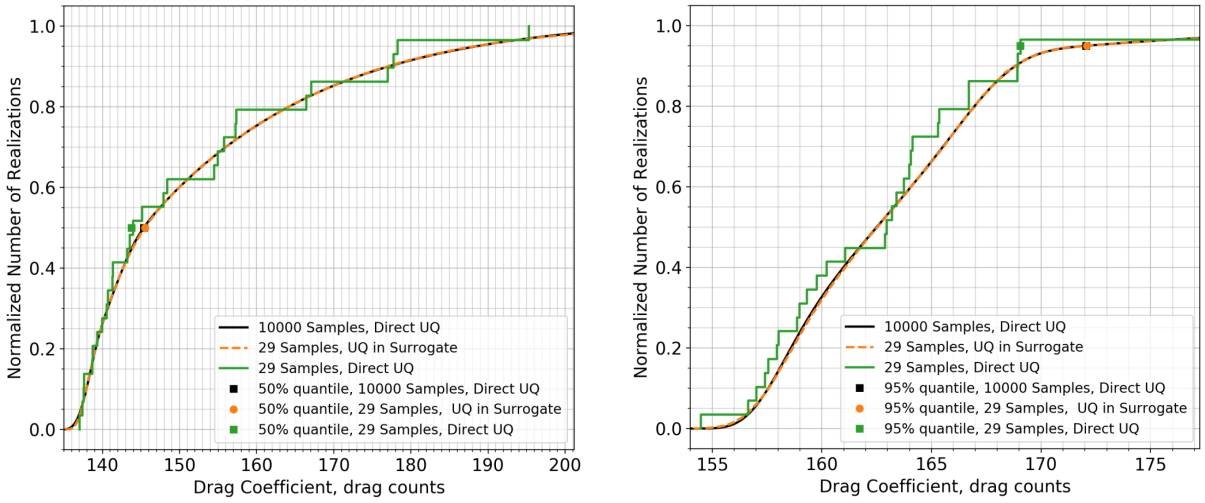

(a) ECDF for robust optimum on $C_{D_{50\%}}$

(b) ECDF for robust optimum on $C_{D_{95\%}}$

Fig. 18 Comparison of Empirical Cumulative Distribution Functions for direct integration with 29 samples (with and without surrogate) and 10,000 samples

Table 5 Error in surrogate-based uncertainty quantification

Configuration	Absolute Error (drag counts)		Relative Error	
	50% Quantile	95% Quantile	50% Quantile	95% Quantile
50% Optimum	0.18	0.097	0.125%	0.05%
95% Optimum	0.04	0.023	0.024%	0.013%

V. Conclusions

A framework for efficient optimization under uncertainty has been presented, validated and applied to the robust design of shock control bumps. The bi-level surrogate approach combining a surrogate-based optimization with surrogate-based uncertainty quantification is able to effectively explore both design and stochastic landscapes. The surrogate-based optimization combines an initial sampling scheme based on Sobol sequences with a Kriging surrogate model, an adaptive sample refinement strategy following the maximum expected improvement criterion, an iterative evaluation of the global minimum of the surrogate with a differential evolution algorithm and a trust region method. Due to the good balance between exploration and exploitation, the surrogate based optimizer efficiently finds the global optimum, while requiring an order of magnitude less function evaluations than traditional global optimization approaches.

The uncertainty quantification, performed at every iteration of the surrogate-based optimization, is based on a Kriging meta-model and an active infill criterion. It is able to reproduce the reference cumulative distribution function and the quantile of interest of the optimum configurations with high accuracy. In this case, a trade-off is present between the required global refinement of the stochastic space and the local refinement in areas close to the quantile. Following this approach, the determination of any given τ quantile of the quantity of interest can be computed with considerably

less function evaluations than with direct integration through Quasi Monte Carlo.

When accounting for operating and geometrical uncertainties, the use of a robust optimization framework with focus on the quantile of the airfoil drag coefficient is useful in the determination of a robust shock control bump. If the interest is the minimization of extreme events, the optimization of the 95% quantile of drag is considered. The resulting configuration has the highest bump height that smears the normal shock wave at high values of lift and Mach, around the 95% quantile. This robust configuration is able to better withstand operating and geometrical uncertainties than the configuration obtained with deterministic optimization. It distributes the longitudinal location of the shock wave more uniformly at an expense of a secondary shock wave at the end of the bump at nominal conditions. The secondary shock wave can be removed if optimizing for the median or 50% quantile. In this case, most common events correspond to a lower value of drag due to the weaker normal shock wave over the bump. However, this configuration is not as effective as the previous in reducing the drag at off-design conditions.

By reducing both the number of iterations in the optimization stage and the number of CFD evaluations in the uncertainty quantification, it is possible to come up with an efficient process chain for robust design within a reduced CPU time. In the future, the framework will be applied to the robust design of high dimensional aerodynamic shape optimization problems involving natural laminar flow wings and three dimensional shock control bumps.

Acknowledgments

This work is funded by the European Commission's H2020 programme, through the UTOPIAE Marie Curie Innovative Training Network, H2020-MSCA-ITN-2016, Grant Agreement number 722734.

References

[1] Skinner, S., and Zare-Behtash, H., "State-of-the-art in aerodynamic shape optimisation methods," *Applied Soft Computing*, Vol. 62, 2018, pp. 933–962. doi:https://doi.org/10.1016/j.asoc.2017.09.030.

[2] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., "Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark," *AIAA Journal*, Vol. 53, No. 4, 2015, pp. 968–985. doi:https://doi.org/10.2514/1.j053318.

[3] Schulz, V., and Schillings, C., "Optimal Aerodynamic Design under Uncertainty," *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, Springer Berlin Heidelberg, 2013, pp. 297–338. doi:https://doi.org/10.1007/978-3-642-36185-2_13.

[4] Beyer, H.-G., and Sendhoff, B., "Robust optimization – A comprehensive survey," *Computer Methods in Applied Mechanics and Engineering*, Vol. 196, No. 33-34, 2007, pp. 3190–3218. doi:https://doi.org/10.1016/j.cma.2007.03.003.

[5] Schuëller, G., and Jensen, H., "Computational methods in optimization considering uncertainties – An overview," *Computer Methods in Applied Mechanics and Engineering*, Vol. 198, No. 1, 2008, pp. 2–13. doi:https://doi.org/10.1016/j.cma.2008.05.004.

[6] Duvigneau, R., "Aerodynamic Shape Optimization with Uncertain Operating Conditions using Metamodels," resreport RR-6143, INRIA, 2007.

[7] Swiler, L., Wojtkiewicz, S., Eldred, M., Giunta, A., and Trucano, T., "Perspectives in Optimization Under Uncertainty: Algorithms and Applications," Vol. 3, 2004. doi:https://doi.org/10.2514/6.2004-4451.

[8] Rumpfkeil, M., "Robust design under mixed aleatory/epistemic uncertainties using gradients and surrogates," *Journal of Uncertainty Analysis and Applications*, Vol. 1, 2013, p. 7. doi:10.1186/2195-5468-1-7.

[9] Maruyama, D., Liu, D., and Görtz, S., "An Efficient Aerodynamic Shape Optimization Framework for Robust Design of Airfoils Using Surrogate Models," *Proceedings of the VII European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS Congress 2016)*, NTUA Greece, 2016. doi:https://doi.org/10.7712/100016.2450.8838.

[10] Dodson, M., and Parks, G. T., "Robust Aerodynamic Design Optimization Using Polynomial Chaos," *Journal of Aircraft*, Vol. 46, No. 2, 2009, pp. 635–646. doi:https://doi.org/10.2514/1.39419.

[11] Zhao, H., Gao, Z., Wang, C., and Yuan, G., "Robust Design of High Speed Natural-Laminar-Flow Airfoil for High Lift," *55th AIAA Aerospace Sciences Meeting*, American Institute of Aeronautics and Astronautics, 2017. doi:https://doi.org/10.2514/6.2017-1414.

[12] Swiler, L. P., and Eldred, M. S., "Efficient algorithms for mixed aleatory-epistemic uncertainty quantification with application to radiation-hardened electronics. Part I, algorithms and benchmark results." 2009. doi:10.2172/972887.

[13] Lockwood, B., Anitescu, M., and Mavriplis, D., "Mixed Aleatory/Epistemic Uncertainty Quantification for Hypersonic Flows via Gradient-Based Optimization and Surrogate Models," 2012. doi:"https://10.2514/6.2012-1254".

[14] Pisaroni, M., Nobile, F., and Leyland, P., "Continuation Multilevel Monte Carlo Evolutionary Algorithm for Robust Aerodynamic Shape Design," *Journal of Aircraft*, Vol. 56, No. 2, 2019, pp. 771–786. doi:https://doi.org/10.2514/1.c035054.

[15] Zhang, X., Wu, Z., and He, W., "An effective approach for robust design optimization of wind turbine airfoils with random aerodynamic variables," *Advances in Mechanical Engineering*, Vol. 11, No. 9, 2019, p. 1687814019879263. doi:https://doi.org/10.1177/1687814019879263, URL `https://doi.org/10.1177/1687814019879263`.

[16] Li, M., "Robust Optimization and Sensitivity Analysis with Multi-Objective Genetic Algorithms: Single- and Multi-Disciplinary Applications," 2007.

[17] Zhang, J., Taflanidis, A., and Medina, J., "Sequential approximate optimization for design under uncertainty problems utilizing Kriging metamodeling in augmented input space," *Computer Methods in Applied Mechanics and Engineering*, Vol. 315, 2017, pp. 369–395. doi:10.1016/j.cma.2016.10.042.

[18] Wang, X., Hirsch, C., Liu, Z., Kang, S., and Lacor, C., "Uncertainty-based robust aerodynamic optimization of rotor blades," *International Journal for Numerical Methods in Engineering*, Vol. 94, No. 2, 2012, pp. 111–127. doi:10.1002/nme.4438, URL `https://doi.org/10.1002/nme.4438`.

[19] Seshadri, P., Constantine, P., Iaccarino, G., and Parks, G., "A density-matching approach for optimization under uncertainty," *Computer Methods in Applied Mechanics and Engineering*, Vol. 305, 2016, pp. 562–578. doi:https://doi.org/10.1016/j.cma.2016.03.006.

[20] Cook, L. W., and Jarrett, J. P., "Horsetail matching: a flexible approach to optimization under uncertainty," *Engineering Optimization*, Vol. 50, No. 4, 2018, pp. 549–567. doi:https://doi.org/10.1080/0305215X.2017.1327581.

[21] Quagliarella, D., Petrone, G., and Iaccarino, G., "Optimization Under Uncertainty Using the Generalized Inverse Distribution Function," *Computational Methods in Applied Sciences*, Springer Netherlands, 2014, pp. 171–190. doi:https://doi.org/10.1007/978-94-017-9054-3_10.

[22] Jones, D. R., Schonlau, M., and Welch, W. J., "Efficient Global Optimization of Expensive Black-Box Functions," *Journal of Global Optimization*, Vol. 13, No. 4, 1998, pp. 455–492. doi:https://doi.org/10.1023/a:1008306431147.

[23] Ogawa, H., Babinsky, H., Pätzold, M., and Lutz, T., "Shock-Wave/Boundary-Layer Interaction Control Using Three-Dimensional Bumps for Transonic Wings," *AIAA Journal*, Vol. 46, No. 6, 2008, pp. 1442–1452. doi:https://doi.org/10.2514/1.32049.

[24] Ashill, P., Fulker, J., and Shires, J., "A Novel Technique for Controlling Shock Strength of Laminar-Flow Airfoil Sections," *Proceedings of the 1st European Forum on Laminar Flow Technology*, Hamburg, 1992, pp. 175–183.

[25] Stanewsky, E., Délery, J., Fulker, J., and de Matteis, P., "Synopsis of the Project EUROSHOCK II," *Drag Reduction by Shock and Boundary Layer Control*, edited by E. Stanewsky, J. Délery, J. Fulker, and P. de Matteis, Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 1–124. doi:https://doi.org/10.1007/978-3-540-45856-2_1.

[26] McGowan, A., "AVST morphing project research summaries in fiscal year 2001," Tech. Rep. nasa tm-2002-2 11769, NASA, Aug. 2002.

[27] Bruce, P. J. K., and Babinsky, H., "Experimental Study into the Flow Physics of Three-Dimensional Shock Control Bumps," *Journal of Aircraft*, Vol. 49, No. 5, 2012, pp. 1222–1233. doi:https://doi.org/10.2514/1.c031341.

[28] Lee, D. S., Periaux, J., Onate, E., Gonzalez, L. F., and Qin, N., "Active Transonic Aerofoil Design Optimization Using Robust Multiobjective Evolutionary Algorithms," *Journal of Aircraft*, Vol. 48, No. 3, 2011, pp. 1084–1094. doi:https://doi.org/10.2514/1.c031237.

[29] Paetzold, M., Lutz, T., Kramer, E., and Wagner, S., "Numerical Optimization of Finite Shock Control Bumps," *44th AIAA Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, 2006. doi:https://doi.org/10.2514/6.2006-1054.

[30] Bruce, P. J. K., and Colliss, S. P., "Review of research into shock control bumps," *Shock Waves*, Vol. 25, No. 5, 2014, pp. 451–471. doi:https://doi.org/10.1007/s00193-014-0533-4.

[31] Jinks, E. R., Bruce, P. J., and Santer, M. J., "Adaptive Shock Control Bumps," *52nd Aerospace Sciences Meeting*, American Institute of Aeronautics and Astronautics, 2014. doi:https://doi.org/10.2514/6.2014-0945.

[32] Nuebler, K., Lutz, T., Kraemer, E., Colliss, S., and Babinsky, H., "Shock Control Bump Robustness Enhancement," *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, American Institute of Aeronautics and Astronautics, 2012. doi:https://doi.org/10.2514/6.2012-46.

[33] Han, Z.-H., Abu-Zurayk, M., Görtz, S., and Ilic, C., "Surrogate-Based Aerodynamic Shape Optimization of a Wing-Body Transport Aircraft Configuration," *AeroStruct: Enable and Learn How to Integrate Flexibility in Design*, edited by R. Heinrich, Springer International Publishing, Cham, 2018, pp. 257–282. doi:https://doi.org/1007/978-3-319-72020-3_16.

[34] Forrester, A. I., and Keane, A. J., "Recent advances in surrogate-based optimization," *Progress in Aerospace Sciences*, Vol. 45, No. 1-3, 2009, pp. 50–79. doi:https://doi.org/10.1016/j.paerosci.2008.11.001.

[35] Han, Z.-H., Görtz, S., and Zimmermann, R., "Improving variable-fidelity surrogate modeling via gradient-enhanced kriging and a generalized hybrid bridge function," *Aerospace Science and Technology*, Vol. 25, No. 1, 2013, pp. 177–189. doi: https://doi.org/10.1016/j.ast.2012.01.006.

[36] McKay, M. D., Beckman, R. J., and Conover, W. J., "Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code," *Technometrics*, Vol. 21, No. 2, 1979, pp. 239–245. doi: https://doi.org/10.1080/00401706.1979.10489755.

[37] Caflisch, R. E., "Monte Carlo and quasi-Monte Carlo methods," *Acta Numerica*, Vol. 7, 1998, p. 1. doi:https://doi.org/10.1017/s0962492900002804.

[38] Kucherenko, S., Albrecht, D., and Saltelli, A., "Exploring multi-dimensional spaces: a Comparison of Latin Hypercube and Quasi Monte Carlo Sampling Techniques," *ArXiv e-prints*, 2015.

[39] Sobol, I., "On the distribution of points in a cube and the approximate evaluation of integrals," *USSR Computational Mathematics and Mathematical Physics*, Vol. 7, No. 4, 1967, pp. 86–112. doi:https://doi.org/10.1016/0041-5553(67)90144-9.

[40] Loeppky, J. L., Sacks, J., and Welch, W. J., "Choosing the Sample Size of a Computer Experiment: A Practical Guide," *Technometrics*, Vol. 51, No. 4, 2009, pp. 366–376. doi:https://doi.org/10.1198/tech.2009.08040.

[41] Han, Z.-H., Abu-Zurayk, M., Görtz, S., and Ilic, C., "Surrogate-Based Aerodynamic Shape Optimization of a Wing-Body Transport Aircraft Configuration," *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, Springer International Publishing, 2018, pp. 257–282. doi:https://doi.org/10.1007/978-3-319-72020-3_16.

[42] Forrester, A. I. J., Sabester, A., and Keane, A. J., *Engineering Design via Surrogate Modelling*, John Wiley & Sons, Ltd, 2008. doi:https://doi.org/10.1002/9780470770801.

[43] Zimmermann, R., "Asymptotic Behavior of the Likelihood Function of Covariance Matrices of Spatial Gaussian Processes," *Journal of Applied Mathematics*, Vol. 2010, 2010, pp. 1–17. doi:https://doi.org/10.1155/2010/494070.

[44] Storn, R., and Price, K., "Differential Evolution. A Simple and Efficient Heuristic for global Optimization over Continuous Spaces," *Journal of Global Optimization*, Vol. 11, No. 4, 1997, pp. 341–359. doi:https://doi.org/10.1023/a:1008202821328.

[45] Alexandrov, N. M., Dennis, J. E., Lewis, R. M., and Torczon, V., "A trust-region framework for managing the use of approximation models in optimization," *Structural Optimization*, Vol. 15, No. 1, 1998, pp. 16–23. doi:https://doi.org/10.1007/bf01197433.

[46] Hennig, P., and Schuler, C., "Entropy Search for Information-Efficient Global Optimization," *Journal of Machine Learning Research*, Vol. 13, 2012, pp. 1809–1837.

[47] Liu, D., and Görtz, S., "Efficient Quantification of Aerodynamic Uncertainty due to Random Geometry Perturbations," *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, Springer International Publishing, 2014, pp. 65–73. doi:https://doi.org/10.1007/978-3-319-03158-3_7.

[48] Kawai, S., and Shimoyama, K., "Kriging-model-based uncertainty quantification in computational fluid dynamics," *32nd AIAA Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics, 2014. doi:https://doi.org/10.2514/6.2014-2737.

[49] Vaart, A. W. v. d., *Asymptotic Statistics*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 1998. doi:https://doi.org/10.1017/CBO9780511802256.

[50] Dudley, R. M., *Uniform Central Limit Theorems*, Cambridge University Press, 1999. doi:https://doi.org/10.1017/cbo9780511665622.

[51] Echard, B., Gayton, N., and Lemaire, M., "AK-MCS: An active learning reliability method combining Kriging and Monte Carlo Simulation," *Structural Safety*, Vol. 33, No. 2, 2011, pp. 145–154. doi:https://doi.org/10.1016/j.strusafe.2011.01.002.

[52] Bect, J., Ginsbourger, D., Li, L., Picheny, V., and Vazquez, E., "Sequential design of computer experiments for the estimation of a probability of failure," *Statistics and Computing*, Vol. 22, No. 3, 2011, pp. 773–793. doi:https://doi.org/10.1007/s11222-011-9241-4.

[53] Schöbi, R., "Surrogate models for uncertainty quantification in the context of imprecise probability modelling," Ph.D. thesis, ETH Zürich, Zürich, Switzerland, 2017.

[54] Mazaheri, K., and Nejati, A., "The Multi-point Optimization of Shock Control Bump with Constant-Lift Constraint Enhanced with Suction and Blowing for a Supercritical Airfoil," *Flow, Turbulence and Combustion*, Vol. 96, No. 3, 2015, pp. 639–666. doi:https://doi.org/10.1007/s10494-015-9671-8.

[55] Tian, Y., Liu, P., and Feng, P., "Shock control bump parametric research on supercritical airfoil," *Science China Technological Sciences*, Vol. 54, No. 11, 2011, pp. 2935–2944. doi:https://doi.org/10.1007/s11431-011-4582-y.

[56] Huyse, L., "Free-form Airfoil Shape Optimization Under Uncertainty Using Maximum Expected Value and Second-order Second-moment Strategies," techreport 2001-211020, NASA, 2001.

[57] Gerhold, T., "Overview of the Hybrid RANS Code TAU," *MEGAFLOW - Numerical Flow Simulation for Aircraft Design*, Springer Berlin Heidelberg, 2015, pp. 81–92. doi:https://doi.org/10.1007/3-540-32382-1_5.

[58] Cook, P., McDonald, M., and Firmin, M., "Aerofoil RAE 2822 - Pressure Distributions, and Boundary Layer and Wake Measurements," Experimental Data Base for Computer Program Assessment," *AGARD Report AR 138*, 1979.

[59] Slater, J. W., "RAE2822 transonic airfoil, NPARC Alliance Verification and Validation Archive," *AGARD Report AR 138*, 2009. URL `http://www.grc.nasa.gov/WWW/wind/valid/raetaf/raetaf.html`.

[60] Gerhold, T., and Neumann, J., "The Parallel Mesh Deformation of the DLR TAU-Code," *Notes on Numerical Fluid Mechanics and Multidisciplinary Design (NNFM)*, Springer Berlin Heidelberg, 2006, pp. 162–169. doi:https://doi.org/10.1007/978-3-540-74460-3_20.

[61] Meinel, M., and Einarsson, G., "The FlowSimulator framework for massively parallel CFD applications," *PARA 2010*, 2010.

[62] Rowan, T. H., "Functional Stability Analysis of Numerical Algorithms," Ph.D. thesis, University of Texas Austin, Austin, TX, USA, 1990. UMI Order No. GAX90-31702.

[63] Hitchens, F., *The Encyclopedia of Aerodynamics*, Andrews UK Limited, 2015. URL `https://books.google.de/books?id=izv0CgAAQBAJ`.