

VALIDATING SURROGATE MODELS AND INCORPORATING UNCERTAINTY QUANTIFICATION IN MULTI-ELEMENT AIRFOIL DESIGN OPTIMISATION

Davide Di Pasquale¹, Carles Bolart¹

¹ Centre for Aeronautics, Applied Aerodynamics Group, Cranfield University, Cranfield, Bedfordshire, MK430A, UK

Abstract

This work deals with the aerodynamics optimisation of a generic two-dimensional three element high-lift configuration. Specifically, it focuses on the development and validation of surrogate models, as well as their integration with optimisation algorithms, uncertainty quantification method, and other computational design tools in order to define and develop a better design methodology for high-lift systems. Special emphasis is put into the process itself to make it fast and highly automated yet keeping accuracy uncompromised. Although the high-lift system is applied only during take-off and landing in the low speed phase of the flight the cost efficiency of the airplane is greatly influenced by it. The ultimate goal of an aircraft high-lift system design team is to define the simplest configuration which, for prescribed constraints, will meet the take-off, climb, and landing requirements usually expressed in terms of maximum L/D and/or maximum C_L . The ability of the calculation method to accurately predict changes in objective function value when gaps, overlaps and element deflections are varied is therefore critical. Despite advances in computer capacity, the enormous computational cost of running complex engineering simulations makes it impractical to rely exclusively on simulation for the purpose of design optimisation. To cut down the cost, surrogate models, also known as metamodels, are constructed from and then used in place of the actual simulation models. A detailed analysis of the integrated design system, the methods as well as the optimisation results are provided.

Keywords: CFD, Engineering Design, High-Lift, Optimisation, Surrogate Models.

1. Introduction

1.1 Context

During the last decades, aerodynamic engineering practices have become more computer-based and less experiment-based, trying to reduce the highly resource-consuming wind tunnel experiments and flight tests. For an aircraft industry, in order to be competitive in today's global market, where aggressive weight targets, shortened development time scale and reduced costs are the primary objectives and constraints, a different approach for the design process is necessary to be introduced. Multi-objective and multi-disciplinary design optimisation tools are becoming a necessity for the development of innovative and more efficient advanced aerodynamic configurations [1]. Computational Fluid Dynamics (CFD) is widely used in vehicle aerodynamics development today, but typically used to study one vehicle configuration at a time. In order to be used for aerodynamic exploration and optimisation the CFD simulation process has to be able to study a large set of design alternatives within the short period of time typically available in the overall aerodynamics development process. Flaig et al. [2] explained that the unaltered, clean wing is based on optimum cruise conditions, since the larger part of the flight consists of cruising towards the destination. Therefore, the aircraft high lift system designer is usually given a wing designed for cruise conditions. The maximum chord of the slat and/or flap(s) is usually dictated by the size of the wing box determined for structural and fuel capacity considerations. Therefore, very little leverage exists on the shape of these elements but more on the spacing with respect to each other (gap and overlap). The highly complex flow-field that develops around wing airfoils especially in high lift configuration, as described by Smith [3], and the computational cost associated to accurately simulate and extract the crucial flow metrics represents a challenging task for the design

optimisation process. The ability of the calculation method to accurately predict changes in objective function value when gaps, overlaps and element deflections are varied is therefore critical. More recently high-lift aerodynamics became a common research field for experimental [4] and numerical studies. The capabilities of numerical methods related to the aerodynamic performance of high-lift devices are reviewed by Rumsey [5]. Most current methods for transport aircraft high lift system design rely on use of wind tunnel testing in conjunction with Computational Fluid Dynamics (CFD) analysis. These approaches do not allow varying the positions of the various elements, such as slats and flaps, in a systematic fashion. The traditional manual approach would find it difficult, if not impossible, to satisfy future design requirements. In addition, aerodynamic simulations are subject to uncertainties in aircraft geometries due to various unpredictable perturbations, e.g. manufacturing tolerances, vibration, operational wear etc. These uncertainties have the potential to dramatically lower aerodynamic performance. Therefore, enabling uncertainty quantification is crucial for the robustness of aircraft designs. Uncertainty quantification (UQ) is a fast-growing research field in the CFD community. In the context of aerodynamic design, most UQ papers focus on the treatment of uncertainties in flow conditions, e.g. Mach number or angle of attack, and only very little attention has been paid to uncertainties in the geometry, as also stated by Liu et al. [6].

1.2 Metamodels for high-lift systems

An important factor concerning the integrated solution process is the time constraint. Most optimisation methods applied directly to the simulation will be slow. A widely accepted method to circumvent this problem is to replace CPU time expensive computer models by CPU inexpensive mathematical functions, called metamodels or surrogates. A review of the state-of-the-art constructing surrogate models and their use in optimisation strategies is to be found in references [7, 8, 9]. They are educated guesses as to what an engineering function might look like, based on a few points in space where it is possible affording to measure the function values. The basic idea is for the surrogate to act as a curve fit to the available data so that the results may be predicted without recourse to the use of the primary source, the computationally intensive simulation codes. Surrogate models are often referred to as Metamodels, because they can be understood as models of the models. See Figure 1.




Figure 1 – Real system, model and metamodel (right to left).

Metamodels are most valuable in fields such as aerodynamics, where models are particularly complex and expensive. However, this field is also specially challenging for metamodels to spread through, due to the intrinsic uncertainty of its computational models. But still, the potential of metamodels is of great interest for the aerodynamics community, especially for costly applications.

In this regard, the possibility of using metamodels for high-lift systems is of enormous interest for the aircraft industry, due to the significant impact these systems have on the sizing, economics and safety of most transport aircraft configurations, as explained by van Dam [10].

High-lift systems are devices able to increase a wing's lift coefficient momentarily and beyond that of nominal cruise flight, which is needed to reduce speed for take-off and landing manoeuvres. As pointed out by Rudolph [11], there are economical limits to the length of runways; safety limits to the take-off and landing speeds; and technical limits for the tires. The need for high-lift systems did not become apparent until 1950s though, when commercial aircraft became faster thanks to the introduction of jet engines. Needless to say, this requirement still holds in modern times and need to comply with the regulations indicated in the FAR/CS-25 documents [12, 13]. A broad range of different high-lift systems have been developed over the years, although the most widely used in civil aircraft is the multi-element

wing, which means deployable elements at the leading edges and trailing of the wing, such as the one in Figure 4. The aerodynamics performance of a multi-element wing is strongly dependent on the interactions between the different elements. Compared with a single element airfoil, additional complexity can be identified in the flow-field that develops around such configuration, as illustrated in Figure 2.

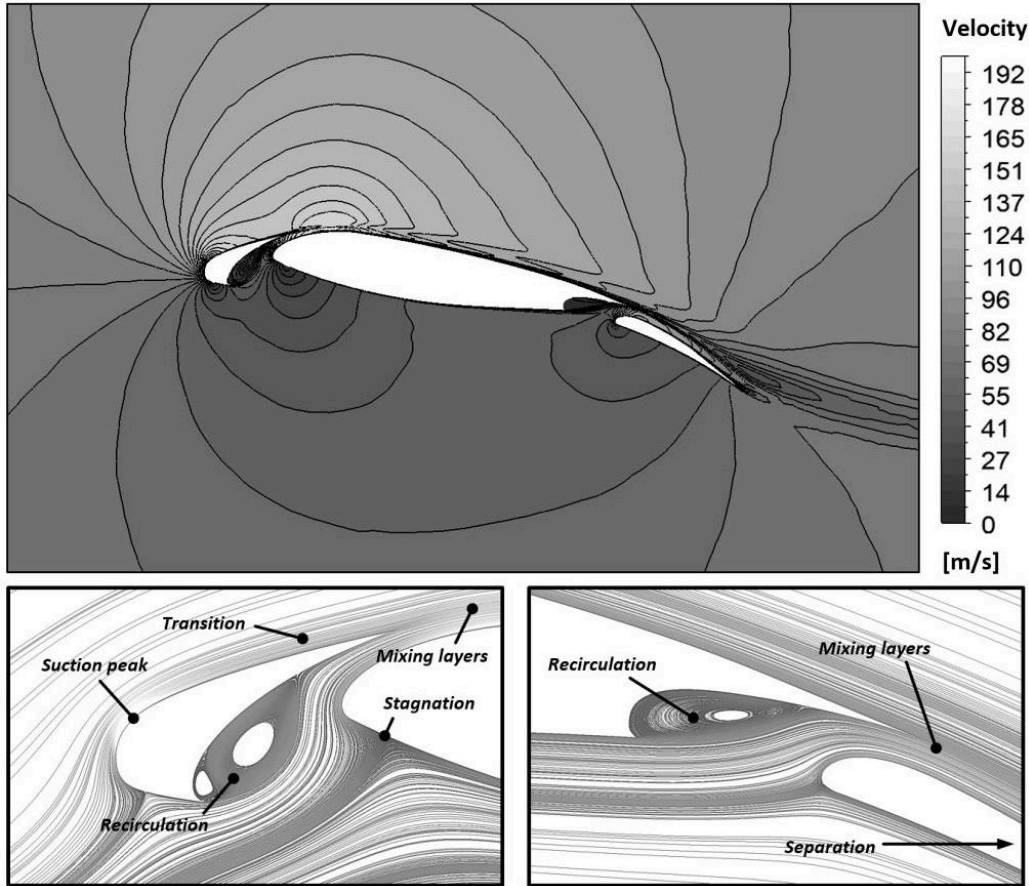


Figure 2 – Top: velocity contours. Bottom: Streamlines and flow features.

In particular recirculation areas develop in the cove region of slat and main element, together with the mixing of the shear layers of the different elements. By changing the geometric spacing (slots) between the elements, the wakes of the corresponding boundary layers can merge, leading to a "confluent boundary layer". This deteriorates the stall characteristics of the airfoil. Therefore, the gap size must be properly balanced.

2. Methodology

As explained by Khondge and Sovani [14], an ideal aerodynamics optimisation simulation process has four major attributes: large number of design points, high automation, speed and high fidelity, even if obviously these constitute two conflicting requirements and the compromise best suited to the application targeted would drive the choices set.

This project builds upon the insights provided by previous work carried out by Lopez [15], yet trying to address those attributes pointed out by Khondge and Sovani with the purpose of producing a solution as compact and industry-worthy as possible:

- **Large number of design points.** More than 200 geometric configurations are tested for both the optimisation and uncertainty quantification objectives.
- **High-automation.** Given that the majority of an engineer's effort is spent in managing inter-software connections and data transfer in the case of an optimisation process made of disjointed building

blocks, thus slowing down the entire optimisation process. In this work the entire optimisation process is performed in a single commercial software platform known as ANSYS Workbench [16], so that the process can be fast and highly automated relying on its seamless integration and automation capabilities.

- **Speed.** Each set of simulations can be run in less than 24h using moderate computing power. Overall, the time scale of the project suits the industrial aerodynamic development process.
- **Accuracy.** The baseline model is validated against experimental data, and finely tuned hybrid mesh ensures well-converged solutions for the vast majority of design points.

Figure 3 depicts the project workflow diagram, the methodology designed for the task at hand. It has three main parts: Baseline Simulation, Metamodel-Based Optimisation (MBO) and Metamodel-Based Uncertainty Quantification (MBUQ). The content of the document closely follows the diagram in Figure 3, as hinted by all section and subsection titles.

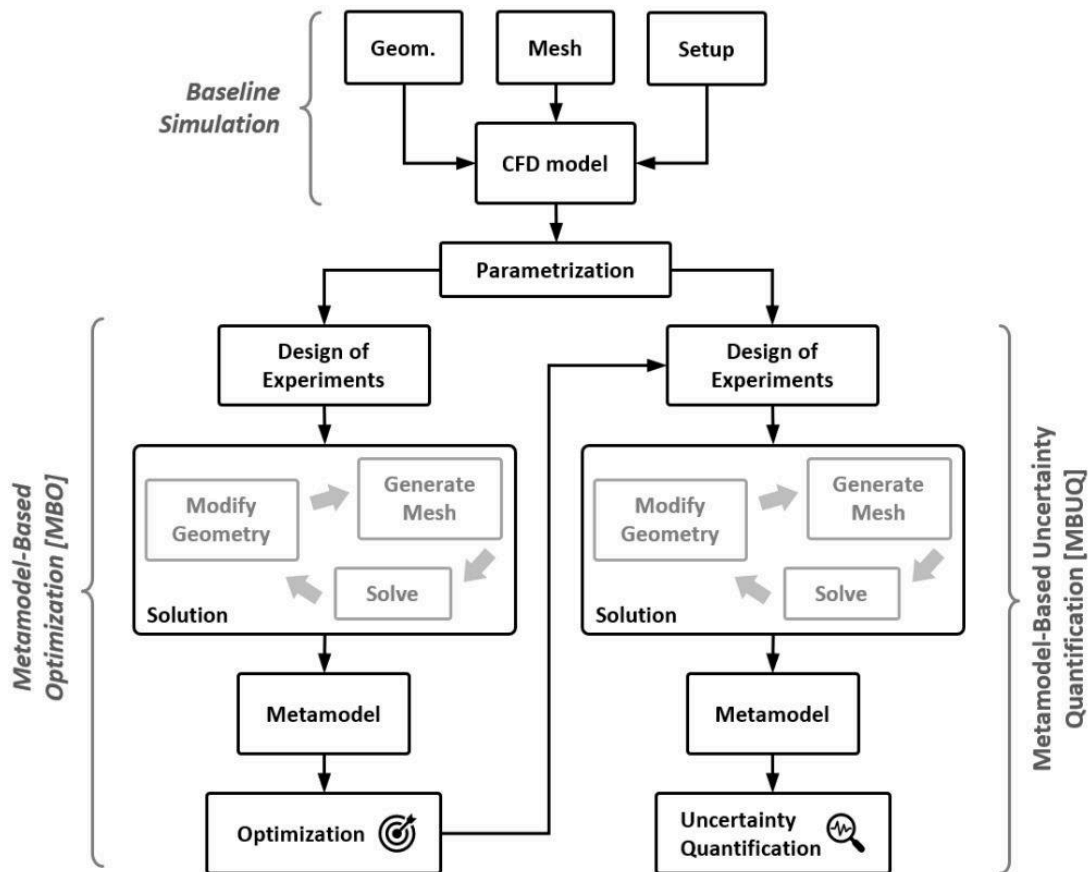


Figure 3 – Project workflow diagram.

Further detail on the practical execution of such methodology, including specific software settings, can be found in the original thesis report [17], from which this paper is built upon.

3. Baseline Simulation

The Baseline is the first of the three major workloads of the project, and the most critical by far. In short, it is a parametric CFD model, which later on needs to be automatically modified and simulated for MBO and MBUQ purposes, without human intervention and for a large number of cases.

3.1 Test Case Geometry

The test case for this project is the Garteur 2D wing section at take-off configuration (also referred to

as *datum*), which is a local sweep-normalized airfoil obtained from the Airbus A310 supercritical wing at its 59% spanwise station, see Figure 4. The reference chord is $c = 1\text{m}$.




Figure 4 – Test case: Garteur 2D wing section at take-off configuration.

The Garteur High-Lift Action Group [18] run extensive experimental tests on this case during the period 1985-1989, as a response to the realization that 2D calculation methods failed at reliably predicting the aerodynamic forces for such complex high-lift flow fields. Because of the large amount of experimental data collected along the program, the case study has been used as a validation reference for CFD codes throughout the years, and there is no different reason for the task at hand.

For this project, the flow conditions are fixed and corresponding to one of those test runs: Mach number $M = 0.22$, Reynolds number $Re = 16.8 \times 10^6$ and angle of attack $\alpha = 12.22$ deg.

3.2 Mesh Generation and Simulation Settings

The geometry of the virtual 2D fluid domain extends $15c$ (15 times the cord) to the back and $10c$ to the front, following a c-shaped inlet boundary. Several sub-regions of the fluid domain are defined too (Figure 5 top left), with the objective of providing denser mesh for the airfoil's surroundings and wake. The mesh needs to be precise for each case yet adaptable to any. This requirement, together with the intricate geometry, make a hybrid mesh the most suitable choice. The structured part is bound to the thin immediate neighbourhoods of the airfoil surface, meant to resolve the viscous mechanisms of the boundary layer (Figure 5 bottom). In ANSYS terminology, this mesh zone is referred to as inflation layer. The rest of the domain is left to the unstructured mesh, which fills the domain in a very isotropic and adaptive way, from the edge of the inflation layer to the outer domain boundaries.

Each of the mesh regions has different options and parameters to be set, and there is no exact science on how to proceed. But two particular aspects require special attention:

- **Wake resolution.** It can be controlled through the sizing of the unstructured mesh behind the airfoil, with the help of the sub-regions defined. The prime importance of the wake refinement is greatly emphasized by Moitra [19] and Murayama and Yamamoto [20].
- **Boundary layer resolution.** It is directly related to the structured inflation layer, and thus can be controlled by its parameters. It happens to be tricky to get right but, two main indicators can come handy:
 - **Dimensionless wall distance** at the surface should be $y^+ \approx 1$.
 - **Turbulent viscosity** contours should peak below the middle of the inflation layer, which in turn should be designed to fully cover the boundary layer with at least 15 cells.

As suited to the project time-scale and scope, the CFD approach taken is Reynolds Averaged Navier Stokes (RANS) with K-omega SST turbulence model, which is a two-equation eddy viscosity model, generally well suited in adverse pressure gradients and separated flows.

Regarding the simulation settings, all ANSYS Fluent setup options and parameters (fluid models, boundary conditions, solution methods, etc.) can again be consulted in the original thesis report [17].




Figure 5 – Geometry and mesh of the virtual fluid domain.

3.3 Solution and Validation

The right combination of mesh and simulation settings proved to produce well-converged, robust and accurate results. These three aspects are discussed next:

- **Robustness** refers to the fact that small changes in mesh density or solver settings do not have a significant impact on the results. It is a minimum requirement to ensure convergence and accuracy of future not assessable geometric configurations.

- **Convergence** indicates that the averaged flow is steadily developed, especially within the boundary layer. It can be checked by lift coefficient (C_L) and drag coefficient (C_D) curves against number of iterations, which proved to reach steady values before 1000. However, the baseline model is configured to 2000 iterations as a safety margin for future configurations that might need more time to converge.
- **Accuracy** can be checked by validation, comparing the pressure coefficient (C_P) distributions of the baseline simulation against experimental data, as seen in Figure 6.




Figure 6 – C_P distribution validation.

For additional proof-check, validation was also successfully performed for a higher angle of attack for which experimental data was also available. To sum up, the baseline model proved to be able to match each suction peak and recovery profile for each of the elements and at both angles, resulting in less than 1% error in the prediction of C_L (refer to the original thesis report [17] for more details).

4. Metamodel-Based Optimisation

4.1 Parameterisation

For simple parameterisation, the main element is locked in place and used as a fixed reference, so that the position and orientation of the other two can be expressed relative to it. Three parameters are then required for each of the two moving elements: vertical position, horizontal position and rotation; which adds up to six parameters or degrees of freedom for the whole airfoil configuration.

In the real system, however, the deployment mechanism would physically link the movement of each element. In Fact, one of the most important class of constraints is represented by kinematics used to deploy the high-lift devices. This aspect has an important influence on limiting the relative positions of slat and flap in respect of the main element. Although, in the current study these constraints are not taking into account, the design space is defined, for both slat and flap, keeping these limitations in mind. However, it is not the purpose of this project to exactly reproduce the behaviour of the real system but to simulate the design process of it.

There are many ways in which these parameters could be defined. It is in the interest of the engineer, though, to make them as physically meaningful as possible while feasible to implement. With this in mind, the parameterisation adopted is the cartesian approach shown in Figure 7.

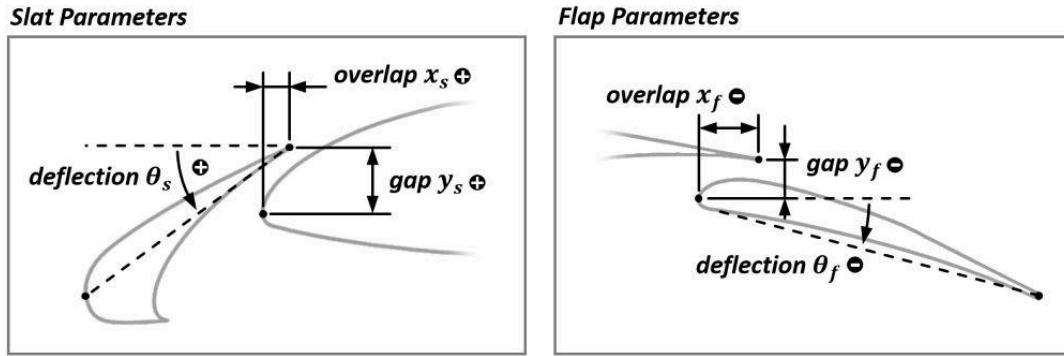


Figure 7 – Parameterisation of the airfoil geometry.

In the software, however, the parameter values are not expressed in their absolute magnitudes shown in the figure, but rather relative to the datum configuration.

4.2 Design of Experiments

Design of Experiments (DOE) refers to the strategy to generate a sample of design points (airfoil configurations) to be tested in order to efficiently map the model response that is, to understand the effect of each input parameter (geometric parameters discussed above) into the output parameters (C_L and C_D coefficients).

First thing to do is to bound the design space domain by restricting each input parameter to a range of values. Table 1 provides the design space limits established, based on visual approval, kinematic constraints in mind and trial and error.

Once this is done, the main concern becomes which design points should be tested and just as with any sufficiently complex problem, the only feasible and most effective way to tackle it is through statistical methods, even if it deterministic in principle.

For the task at hand, the Latin Hypercube Sampling (LHS) method is used, which is an enhanced version of the random approach for a more even distribution. This is the only task, however, that could not be directly and efficiently implemented in ANSYS, since even though the platform features DOE and LHS capabilities, it has no way to anticipate or handle possible geometric interferences.

Table 1 – MBO design space by means of parameter bounds.

Parameters	Symb.	Units	Datum	Min *	Max *
P1 Slat overlap	x_s	m	0.02	-0.05	0.05
P2 Slat gap	y_s	m	0.05	-0.02	0.02
P3 Slat angle	θ_s	deg	35.0	-10.0	10.0
P4 Flap overlap	x_f	m	-0.03	-0.05	0.10
P5 Flap gap	y_f	m	-0.05	-0.01	0.01
P6 Flap angle	θ_f	deg	-15.5	-20.0	10.0

*relative to Datum

To address that, a MATLAB script is used. In essence, the code generates an LHS sample using the built-in *lhsdesign* function, then checks whether each design point is geometrically acceptable, filters out the wrong ones and displays the resulting sample. In addition to the simple geometric interference check, the code implements minimum spacing requirements to filter-out configurations with minute gaps between elements, which would likely cause flow blockage in reality and mesh errors in the simulation. The script also enables the visualization of the design space, which is not possible within ANSYS either. Using the input parameter ranges specified in Table 1 and imposing a minimum spacing of $0.01c$, the resulting sample is displayed in Figure 8. It consists of 220 approved design points filtered from an initial raw sample of 400, which means 45% of design points were either intersecting or leaving too small distances between elements. 220 design points seem high enough as a sample size, well above 30 times the number of input parameters. The MATLAB generated list of design points can be easily

imported into ANSYS through a .csv file.




Figure 8 – MBO design space visualization.

For further visualization, some parameter values are plotted in pairs in Figure 9. The first plot shows how the combined slat gap and slat overlap parameters effectively avoid the interfering region. As a consequence of that, the second plot, just as it would happen with any plot involving one of those parameters, inevitably appears biased. The third plot shows the expected even distribution of two unbounded and unrelated parameters.

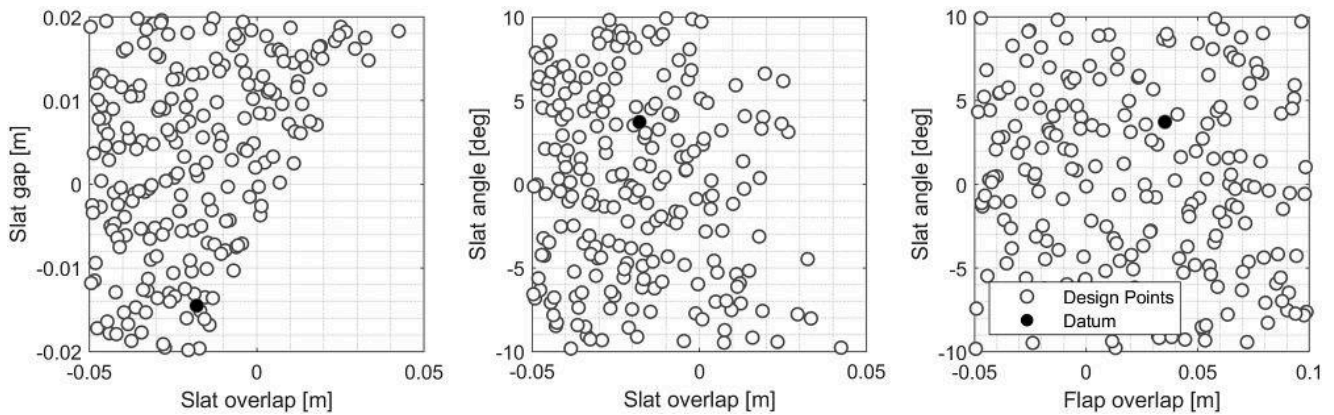


Figure 9 – MBO input parameters distributions.

4.3 Solving

Following the diagram in Figure 3, once the parameterisation is set and the DOE sorted out, ANSYS Workbench will automatically solve all design points. At the push of the *Update Project* button, the program will start to repeatedly modify the baseline geometry, mesh, solve and save the output C_L and C_D values for each of the 220 design points, each consisting of 2000 iterations over 400000 cells. Ideally for the engineer, this step would require little to no effort. In reality, though, it turned out to be a bit more complicated. In any case, using twenty Intel® Core i-7700 machines, the 220 design points could be solved in less than 12 hours.

The baseline model proved to be robust enough, since 100% of design points were successfully solved, and more than 95% achieved convergence. Not a single simulation reported mesh errors or diverging results. The non-converged 7 out of 220 design points, presented highly oscillating convergence plots, a pseudo vortex shedding phenomenon. In reality, those configurations would likely present fully detached flow.

Figure 10 display the raw output parameters (C_L , C_D and Efficiency $E_{ff} = C_L/C_D$) plotted in pairs. These results are very significant, and reveal the fine optimisation achieved by the engineers back in late 1970s, certainly without the computational capabilities or integrated design exploration tools available today. The datum configuration is clearly pareto-optimal.




Figure 10 – MBO output parameters distributions.

Pareto-optimality refers to a quality of a design point resulting from a multi-objective optimisation by which no single objective can be improved for that point without worsening other objectives. For instance, for the datum configuration C_L cannot be further increased without reducing C_D . In other words: for the given C_L , the C_D could not be lower. The Pareto-front is the limit defined by a set of Pareto-optimal design points. Pareto-fronts are clearly distinguishable in Figure 10 as invisible walls limiting the performance of the airfoil. It is also noticeable that the cloud of points seems to be segregated into two groups, probably due to a change in the flow structure.

4.4 Metamodelling

Once the output parameters are computed for all design points, the generation of the metamodel is immediate and could potentially be automated too. The term metamodel, however, does not appear anywhere in the ANSYS environment. Instead, response surface is used.

Four different response surface types are available in ANSYS to select from: Full 2nd Order Polynomial, Non-Parametric Regression, Kriging and Neural Network. For the sake of testing, all response surface types were tried out individually and compared. But there is yet another option named Genetic Aggregation (GA), which automates the process of selecting and configuring the response surface best suited to the data. The GA algorithm has provided the best Root Mean Square (RMS) results (see Table 2). The RMS error is a quadratic scoring rule, which measures the average magnitude of the error, which is a global error measure and provides an understanding of the model accuracy over the entire design domain. It is negatively oriented score: lower values are better.

Table 2 – Goodness of fit for different response surface types.

Response Surfaces	RMS* Error (Best = 0)	
	CL	CD
Genetic Agregation [Kriging]	8.4E-08	3.1E-09
Standard Response Surface	1.9E-01	6.2E-03
Kriging	1.1E-07	3.1E-09
Non-parametric Regression	2.0E-02	1.1E-03
Neural Network, 10 cells	1.1E-01	4.1E-03

*Root Mean Square

The most suitable response surface type selected through the GA algorithm was the Kriging, also known as Gaussian Regression Process. It is a widely used metamodelling algorithm that goes beyond standard regression by combining polynomials with additional local deviation models to perform multi-dimensional interpolation across the design space.

4.4.1 Surface Plots and Sensitivity Analysis

Using the Kriging response surface type selected through the GA algorithm, Figure 11 shows the resulting surface plots generated for both C_L and C_D . It is important to notice here the distinction between

response surface and surface plot. On the one hand, the response surface is the Kriging metamodel, a high-dimensional mathematical object with as many dimensions as input and output parameters. On the other hand, the surface plot is the three-dimensional surface-like graph used to visualize the general effect of specific input parameters to a specific output. So, the metamodel is not solely the surface-like plot displayed in Figure 11, just a planar section cannot possibly describe a three-dimensional object.




Figure 11 – MBO Kriging metamodel surface plots.

Besides their graphical limits, however, surface plots are extremely useful to visualize and understand different aspects of the response surfaces. From Figure 11, for example, some conclusions can be intuitively drawn at first sight; large negative flap angles at moderate overlap values produce high lift; large negative flap angles at high overlap values, however, result in an extremely inefficient configuration that likely features separated flow; the lowest drag is achieved at low flap angles, but this also leads to low lift. Another useful visual tool for understanding response surfaces is the local

sensitivity chart, shown in Figure 12. It allows to see the influence of each input parameter to the outputs. Some conclusions can be drawn from it at first sight: P6, the flap angle, is by far the most influential one; P4, the flap overlap, can have a strong effect on lift without affecting considerably the drag; P5, the flap gap, has little influence on objective function of interest.




Figure 12 – MBO local sensitivity chart.

It is not a coincidence that, for the surface plots in Figure 11, the input parameters are chosen to be the ones having the highest impact on the outputs. In fact, there is a strong relation between surface plots and sensitivity charts, explained next.

It might be striking, a priori, the fact that the surface plots do not go through all design points while the quality metrics indicate perfect fit (recall Table 2). In fact, if any other combination of input parameters were chosen for the surface plots, the mismatch would be much more staggering.

But there is no error whatsoever; it is again a consequence of using low-order approximations to reveal certain aspects of a higher-order function. Looking at Figure 12, two design points with similar flap angle and flap overlap will likely result in similar lift outputs because the other parameters have little influence. Bringing the thought of experiment to the extreme, if a surface plot for a given set of parameters would perfectly pass through all design points, it would mean that all other parameters would have absolutely no influence on the results. In summary, for a given input parameter or pair of input parameters, the closer the design points are to their surface plots the higher their local sensitivity.

4.4.2 Sample Size

214 design points (datum + 220 solved - 7 unconverged) were considered for the metamodelling construction. It is clear, though, that the sample size must have an effect on the response surface and associated plots. In an attempt to examine this effect, some randomly selected samples of 50, 100 and 150 design points were extracted from the original set. These samples were then imported into different ANSYS files and used as if they were the only design points known to generate Kriging response surfaces through the GA. For comparison, Figure 13 shows the contour plots of such surface plots.




Figure 13 – Kriging metamodel contour plots for increasing sample size.

The results obtained were not as clear as anticipated, but rather obscure. Expectedly, as the sample size would increase, the subsequent metamodellers would converge to a single form. But this effect is not

significantly observed in the figure, although all samples manage to roughly capture the shape and the approximate location of maximum lift.

Reconsidering this analysis, however, there were two aspects that made it dubious and unfair: first, the extracted samples were selected randomly from the whole set of design points, missing the benefit of LHS for improved spreading across the design space; and second, each of the extracted samples was not built as an extension of the previous but randomly selected from scratch.

But in any case, for improved metamodel accuracy and better optimisation, it definitely would not be a good practice to simply increase the number of points. Successive refinements around the optimal is a much better approach, as suggested for future work in the conclusion.

4.5 Optimisation

Recalling the diagram in Figure 3, this section culminates the MBO. Recapping, the ultimate goal behind the generation of metamodels is the possibility to estimate the aerodynamic performance of a new configuration without the need to run CFD again. Once the initial effort to simulate all design points is done, the metamodel provides a cheap and fast test bed for all sorts of analysis to be carried out. Some simple analysis, namely surface plots and local sensitivities, were already addressed in the previous section.

Before proceeding to the execution of the optimisation, however, some geometric bounds are needed to prevent the evaluation of geometrically interfering cases. This can be done by defining the inequality $P2 \geq 0.65P1 - 0.006$ [m] using the parameter relationship feature. This linear expression relates slat overlap and slat gap variables to produce the desired optimisation bound, as visualized in Figure 14.




Figure 14 – Optimisation linear bound for the slat.

Optimisation is the task of finding the best solution to a problem according to some criteria. Computer algorithms tackle the problem by iteratively and intelligently trying out different solutions, getting better and better until no significant improvement can be achieved.

The optimisation algorithm used in this case is the Multi-Objective Genetic Algorithm (MOGA), a variant of a popular and standardized algorithm NSGA-II, built-in into ANSYS DesignXplorer [21]. Multi-objective indicates the ability to optimise more than one criterion simultaneously. This is also often referred to as pareto-optimisation. On the other side, genetic algorithm indicates that the code works by a series of rules and mechanisms inspired by the processes of natural selection.

By default, as automatically configured by the program depending on the dataset, the algorithm runs 6000 initial metamodel evaluations and then iterates in cycles of 1200 evaluations until 2% convergence criteria is met or a maximum number of 20 iterations is reached. For a thorough exploration of the design space, several optimisations were carried out using different criteria: two single-objective (maximum C_L and minimum C_D) and one multi-objective (minimum C_D , maximum E and bounded C_L , required to be at least as good as the datum value).

For all MOGA runs, optimal solutions were converged rather fast and before reaching the maximum number of iterations, proving default settings were effective enough. Once the optimal values were found, or rather, estimated, they were validated by means of the CFD model. All the data collected during these analyses is compacted in Table 3, where *Estimated* refers to the optimal values computed by MOGA and *Calculated* refers to the validated CFD simulations. As a general observation, all optimal values were too optimistic, overestimating C_L and underestimating C_D to some degree. Moreover, looking at the relative errors, the max C_L has the largest error in C_L of all, while the min C_D has the largest error in C_D . More alarming, however, is Table 4, which compares these results to the maximum and minimum values among all design points tested previously. The calculated solutions were not only

worse than the estimated but also worse than the best achieved before. So, MOGA seemed to fail in finding a better solution within the design space.

Table 3 – Optimisation results and validation.

Parameters	Units	Datum	Single Objectives		Tradeoff* Optimization			
			Max CL	Min CD	Cand. 1	Cand. 2	Cand. 3	
P1	Slat overlap	mm	0.00	-28.81	-29.89	-16.36	-15.91	-15.89
P2	Slat gap	mm	0.00	2.77	-19.65	-16.22	-15.74	-16.01
P3	Slat angle	deg	0.00	-7.48	2.68	2.25	4.27	2.70
P4	Flap overlap	mm	0.00	30.42	40.44	4.60	5.81	5.82
P5	Flap gap	mm	0.00	-0.03	-8.51	1.45	1.97	1.44
P6	Flap angle	deg	0.00	-15.62	9.97	1.59	1.82	1.55
CL	Estimated	-		3.965	2.420	3.010	2.993	3.017
	Calculated	-	2.978	3.716	2.380	2.886	2.865	2.891
	Relative error	%		-6.3%	-1.6%	-4.1%	-4.3%	-4.2%
CD	Estimated	-		0.086	0.035	0.047	0.047	0.047
	Calculated	-	0.055	0.087	0.040	0.052	0.051	0.052
	Relative error	%		0.9%	13.9%	9.9%	9.1%	9.7%
Eff.	Estimated	-		46.02	69.73	63.72	63.60	63.72
	Calculated	-	53.82	42.73	60.21	55.60	55.79	55.63
	Relative error	%		-7.2%	-13.7%	-12.7%	-12.3%	-12.7%

*Bounded CL, Min CD, Max Eff.

Table 4 – Minimum and maximum force coefficients.

Objectives	DOE	Optimizer	
		Estimated	Calculated
Max CL	3.850	3.965	3.716
Min CD	0.039	0.035	0.040

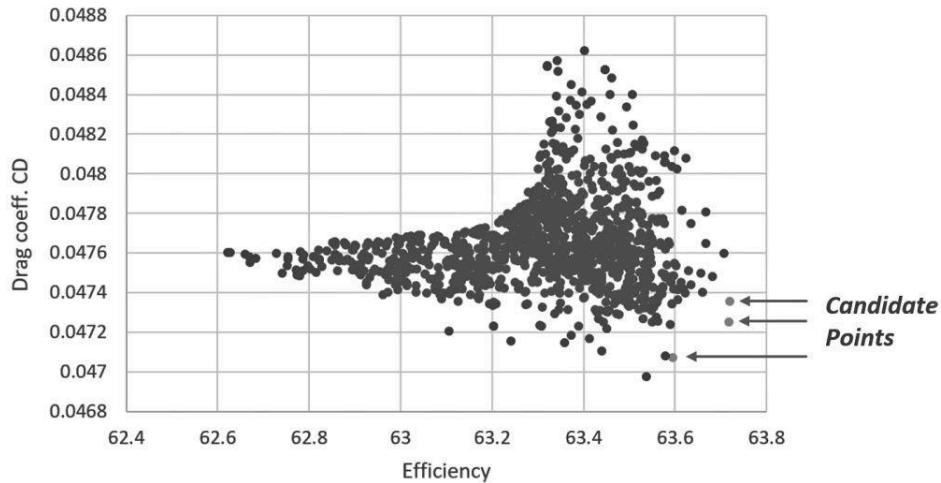


Figure 15 – Multi-objective optimisation evaluations.

But tracing back, the fault came actually from the metamodel itself being too optimistic, which is probably due to the low design space resolution around these maximum performance points. Although disappointing it may sound, without having any advanced metamodeling and optimisation tools, simple inspection would have been enough to hand-pick a better solution for the single objective problems. This discussion will be reopened in the conclusion for the recommendations for further work. For the multi-objective approach, however, it is not straightforward to compare against all design points because the optimised solutions are actually traded-off based on several criteria. Returning to the concept of pareto-optimality, Figure 15 shows the pareto-front encountered by MOGA during the multi-

element airfoil optimisation. The three candidate points indicated in the figure are the ones listed in Table 3. For this case, comparing to the datum, the algorithm is actually able to find points that, once validated, provide lower C_D and higher E values. However, all candidates fail to keep the C_L above the datum as required. Notwithstanding the not-quite-right estimations, the general trend is actually good, and promising insights can be extracted from this analysis. From Table 3, C_L errors for metamodel to CFD are around 5% for all cases, and as commented in section 3.3, the error for CFD to experimental was less than 1%. Assuming these errors remain similar across the design space, the present analysis could be predicting C_L values surely below 10% error compared to experiments, for any configuration and within seconds.




Figure 16 – Maximum C_L , minimum C_D and tradeoff comparison.

Other positive observations that verify the methodology and confirm the overall good trends can be drawn from the comparisons in Figure 16.

- Starting with the first picture, the resulting flap configurations for the three different optimisation objectives is coherent. Being the flap angle the most influential parameter: high lift is achieved by large negative deflections; low drag is achieved by small deflections; and the trade-off just happens to be in between, as expected.
- Moving to the second comparison, the trade-off solution, which tried to simulate the typical take-off requirements, closely matches the actual datum take-off flap configuration.
- Even more positively striking is the third comparison, in which the max C_L candidate, a requirement typically adopted for landing, totally matches the actual landing flap configuration.

For all comparisons, the slat is not really matching, which is also plausible due to the low influence its input parameters have on the output results, as already shown in Figure 12.

5. Metamodel-Based Uncertainty Quantification

Revisiting again Figure 3, this chapter deals with the discussion of the third and last work package. As shown in the diagram, the MBUQ follows a very similar path to the MBO: it uses the parameterised baseline model to generate a DOE, solve it, build a metamodel upon the results and perform some sort of post-processing. In contrast to the MBO, however, the ultimate goal of the MBUQ is not to find an optimal configuration but to study the propagation of uncertainties through the model.

As a matter of fact, deterministically optimised solutions, such as the calculated by MBO, do not offer any guarantee to perform well in reality because real conditions are not fixed and determined but always subject to uncertainties. These uncertainties can potentially affect the performance dramatically, making all effort put into the design process pointless. For a good design, then, more important than optimality is robustness to ensure these unknowns are under control. This is the reason of being of the MBUQ analysis. In principle, the MBUQ could have been worked on without the MBO, since it uses the same parameterised geometry for its baseline model. However, the MBUQ needs a design point to analyse

the robustness upon. For this purpose, an optimal provided by MBO is taken. For the MBUQ task, ANSYS provides the Six Sigma Analysis (SSA) package. SSA works as follows: first, it generates a metamodel based on input-output relations of a statistically-defined set of design points; second, it evaluates the metamodel thousands of times to calculate the probability of certain output values based on the probability of the inputs. In essence, the MBUQ can be simply understood as an MBO process for which only the DOE and the post-processing are different -being the parameterisation, solving and metamodeling steps the exact same.

5.1 Design of Experiments

In contrast to the MBO, the MBUQ is not meant to explore a large design space for an optimal to be found, but to explore the close surroundings of a design point to evaluate its robustness. Hence, the design space for the MBUQ is much smaller.

Another key difference, yet related, is the way in which the design space is defined. Contrary to the parameter bounds needed for the DOE in MBO, SSA requires the DOE in MBUQ to be determined by probability distributions, although under the hood ANSYS takes parameter ranges corresponding to $\pm 3\sigma$, i.e. 3 times the standard deviation from the mean.

For simplicity and testing purposes, normal distributions are considered. For the means, the Candidate 1 input parameters from the trade-off MBO are used. Even though the optimised solutions still have room for improvement, they serve the purpose of analysis. Regarding the standard deviations, $\sigma = 0.001\text{m}$ and $\sigma = 0.1^\circ$ are taken for the longitudinal and angular parameters, respectively as representative of small misalignments. The resulting design space for MBUQ is shown in Table 5.

Table 5 – MBUQ design space by means of parameter normal distributions.

<i>Parameters</i>	<i>Symb.</i>	<i>Units</i>	<i>Mean</i>	<i>Sigma</i>
<i>P1 Slat overlap</i>	x_s	<i>mm</i>	-16.36	1.0
<i>P2 Slat gap</i>	y_s	<i>mm</i>	-16.22	1.0
<i>P3 Slat angle</i>	θ_s	<i>deg</i>	2.25	0.1
<i>P4 Flap overlap</i>	x_f	<i>mm</i>	4.60	1.0
<i>P5 Flap gap</i>	y_f	<i>mm</i>	1.45	1.0
<i>P6 Flap angle</i>	θ_f	<i>deg</i>	1.59	0.0

Because the candidate design point is far from interfering geometrically and the standard deviations are small, there is no threat of possibly intersecting geometry. Hence, the DOE can be safely generated within ANSYS, by means of the Optimal-Space Filling algorithm, equivalent to MATLAB's *lhsdesign* function.

200 design points are generated using such method. In order to check the validity, however, the plotting section of the MATLAB script is used again. In a sense, for the MBUQ analysis, the script is needed to confirm it is not needed indeed. Figure 17 displays the sample of design points for the MBUQ.




Figure 17 – MBUQ design space visualization.

5.2 Solving and Metamodelling

Using the same parameterised baseline model as the MBO, and with the new DOE, the MBUQ sample is automatically solved by means of the same process described in section 3.3. As expected, all 200 design points successfully achieved converged solutions. Once the results are collected, and before going any further, some insights can be drawn from the raw outputs, plotted in pairs in Figure 18.

First thing to notice is that the candidate is indeed not the most optimal according to the minimum C_D and maximum E_{ff} criteria, as expected due to the low resolution of the MBO metamodel and consequent coarse optimisation. However, it is not totally fair to judge it solely based on that because the bounded C_L also played an important role in the trade-off optimisation. Looking closely at Figure 18, a second thing to notice is a slightly appreciable pareto-front. Bearing that in mind, the candidate is actually very close to the true pareto-optimality.

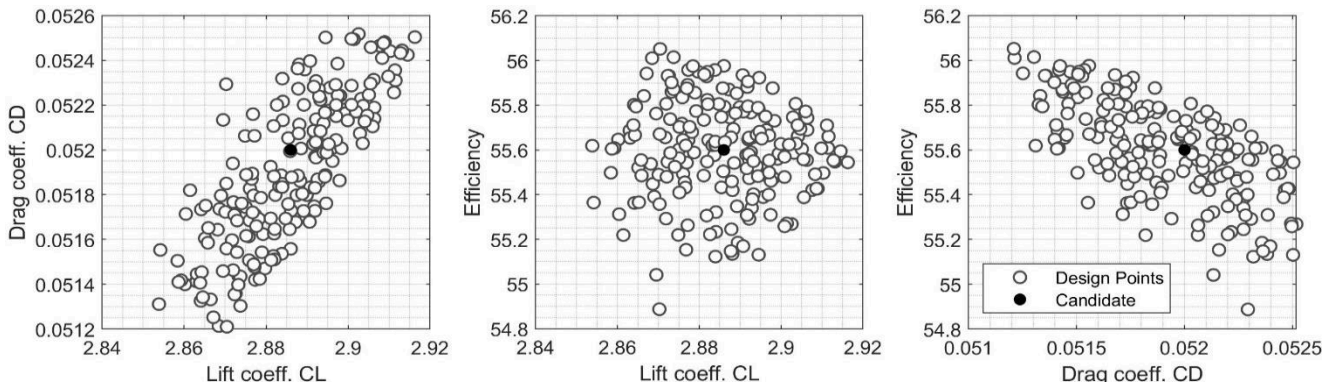


Figure 18 – MBUQ output parameters distributions.

Upon the output results, a new metamodel is immediately generated by means of the GA. The local sensitivity chart for the MBUQ metamodel is displayed in Figure 19. As seen, within the small design space under consideration, the relative influence of each input parameter is different compared to the MBO local sensitivities previously shown in Figure 12.




Figure 19 – MBUQ local sensitivity chart.

Here, the flap overlap loses importance while the slat parameters gain more. The flap angle still dominates, but the contributions are more evenly spread.

As a result of this, no combination of input parameters results in surface plots for MBUQ as nicely fitted as the ones in Figure 11 for MBO. Yet, the goodness of fit for the response surface or metamodel is sufficiently good: $RMS_{CL} = 9.8 \times 10^{-4}$; $RMS_{CD} = 4.3 \times 10^{-5}$.

5.3 Uncertainty Quantification

Taking advantage of the MBUQ metamodels, SSA algorithms can perform, similarly to what the optimisation algorithm did for the MBO metamodels, thousands of evaluations on the model, in a few seconds. 10000 evaluations were performed across the design space using LHS sampling technique.




Figure 20 – MBUQ output probability distributions.

Unlike optimisation algorithms, no successive iterations are needed neither there is a convergence criterion to meet. Instead, Six Sigma Analysis (SSA) uses the input-output relations from metamodels in conjunction with the input probability distributions to quantify the propagation of uncertainties to the outputs.

The resulting output probability distributions are shown in Figure 20. As observed, each of the outputs take the approximate shape of a slightly skewed normal distribution, inherited from the inputs.

Shape parameters for the approximate normal curves, i.e. means, standard deviations and skewness-are listed in Table 5. Most importantly, though Table 5 also gathers the output variation ranges for given probability margins, which are the ultimate objectives of the uncertainty quantification analysis. Absolute values for six-sigma ranges are calculated by default by SSA, which account for 99.99966% of probability. To compute the values for 95% of probability, the two-sigma rule was used.

In the end, no matter how many surface plots, sensitivity charts, scatter plots, probability histograms or complicated mathematics are needed throughout the uncertainty quantification analysis, but the truly valuable result is the ability to make statements like: for the given conditions, for more than 99.99% of cases C_L is guaranteed not to vary more than 1.63% from its nominal optimal value.

Table 5 – SSA numerical results.

		<i>CL</i>	<i>CD</i>	<i>Eff.</i>
Normal Distributions	<i>Mean</i>	2.886	0.0519	55.611
	<i>Standard Deviation</i>	0.008	0.0002	0.117
	<i>Skewness</i>	0.037	0.0208	-0.368
Six Sigma (99.99...%)	<i>Variation</i>	± 0.047	± 0.0011	± 0.701
	<i>Rel. variation</i>	± 1.63%	± 2.12%	± 1.26%
Two Sigma (95.44%)	<i>Variation</i>	± 0.016	± 0.0004	± 0.234
	<i>Rel. variation</i>	± 0.54%	± 0.71%	± 0.42%

6. Conclusion

The context and the necessity behind the project has been exposed. The methodology according to some desired qualities for industry-readiness, namely large design space, high-automation, speed and uncompromised accuracy has been outlined. The process of construction and validation of a metamodel and its use for optimisation purposes has been explained.

An integrated metamodel-based design method featuring optimisation and uncertainty quantification has been developed and tested upon the geometric configuration of a multi-element airfoil, and has demonstrated to be seamless and robust. And being the optimisation process de-coupled from the RANS execution reduce the risk of CFD calculations failure during the process. Several metamodel types were generated upon the simulation results and compared. The Kriging metamodel, tailored by means of the genetic aggregation algorithm, proved to be the best fitted. Surface plots and sensitivity charts helped to understand the behaviour of the model. Several optimisation studies were performed upon the metamodel, using single and multiple objectives. The optimised configurations were verified using CFD. Overall, the metamodel-based optimisation process proved to be sensible and presented good trends, but some fixable shortcomings were also identified. The model is not guaranteed to be accurate over the whole design space, especially if the problem tackled is highly non-linear. Besides, the number of samples required for constructing an accurate surrogate model increases rapidly with the number of design variables. In general, and despite the few setbacks encountered, the method proved to be promising and well fitted for the industrial design time-scale.

6.1 Future Work

Further work on the same topic should focus on the improvement of the metamodel-based optimisation process. A shortcoming of the optimisation results of this project was the low resolution of the metamodel around the optimal, consequence of the fact that the metamodel was generated all at once and evenly spread across the design space and then used as a black box in the optimisation process. Indeed, it was generated in order to be as accurate as possible in all input domain and minimizing average prediction error for the model, and clearly the solution of optimisation problems using a prebuild surrogate model, essentially depends on the training sample. A more efficient approach which is under development, is to iteratively refine the metamodel around the regions of interest, because during an optimisation process there is no need to have good accuracy everywhere. It is needed best precision near optimum points to locate them and in other places it is necessary a model to be accurate enough just to indicate in which regions the optimiser should dig for optimum (or in which regions it is not known the function behaviour) [22]. Future work should focus on an adaptive DoE process where new samples are sequentially inserted to improve the prediction accuracy, providing the possibility to call the source CFD code from the optimisation process. In this way for the same amount of design points, the model would provide a much better optimised solution. As a positive side effect, it could also provide insights on the sample resolution required for reliable metamodeling.

7. Contact Author Email Address

mailto: davide.dipasquale@cranfield.ac.uk

8. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third-party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

- [1] Slotnick J., Khodadoust A., Alonso J., Darmofal, D., Gropp W., Lurie E., and Mavriplis D., CFD vision 2030 study: a path to revolutionary computational aerosciences. Technical report, NASA Langley Research Center, 2013.
- [2] Flaig A, and Hilbig R. High-lift design for large civil aircraft, in *AGARD-CP-515*, 1993.
- [3] Smith A. High-Lift Aerodynamics. *37th Wright Brothers Lecture*, Vol. 12, No. 6, 1975.
- [4] Reckzeh D, Hansen H. High Reynolds-number wind tunnel testing for the design of Airbus high-lift wings, *New Results in Numerical and Experimental Fluid Mechanics V*, edited by H. Rath, C. Holze, H. J. Heinemann, R. Henke, and H. Honlinger, Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Springer, 2006.
- [5] Rumsey L and Ying S. Prediction of high lift: review of present CFD capability. *Progress in Aerospace Sciences*, No. 38, pp. 145-180, 2002.
- [6] Liu D, Litvinenko A, Schillings C, and Schulz V. Quantification of airfoil geometry induced aerodynamic uncertainties. *Society for Industrial and Applied Mathematics*, 2016.
- [7] Cenaero, I,L, Complements on Surrogate Based Optimization for Engineering Design, RTO- EN-AVT-167 5 – 167.
- [8] Forrester J, and Keane A,J. Recent advances in surrogate-based optimization, *Progress in Aerospace Sciences*, Volume 45, Issues 1-3, pp. 50-79, January-April, 2009.
- [9] Queipo N.V, Haftka R.T, Shyy W, Goel T, Vaidyanathan R, Tucker, P.K. Surrogate-based analysis and optimization, *Progress in Aerospace Sciences*, 2005 Volume 41, Issue 1, pp. 1-28.
- [10] Van Dam C. The Aerodynamic Design of Multi-element High-Lift Systems for Transport Airplanes. *Progress in Aerospace Sciences*, Vol. 38, No. 2, pp. 101-144, 2002.
- [11] Rudolph P. High-Lift Systems on Commercial Subsonic Airliners. NASA Contractor Report 4746, No. Sept. 1996.
- [12] FAA Federal Aviation Regulations (FAR-25), 2013.
- [13] EASA Certification Specifications (CS-25), 2013.
- [14] Khondge A. and Sovani S. An Accurate, Extensive, and Rapid Method for Aerodynamics Optimization: The 50:50:50 Method. *SAE Technical Paper Series*, Vol. 1, 2012.
- [15] Lopez M. Validating Surrogate Models and Incorporating Uncertainty Quantification in Aerospace Engineering Design Optimization. MSc thesis report, Cranfield University, 2018.
- [16] ANSYS WORKBENCH Software, www.ansys.com.
- [17] Bolart C. *Validating Metamodels and incorporating Uncertainty Quantification in Multi-element Airfoil Optimisation*. MSc thesis report, Cranfield University, 2018.
- [18] Dargel G and Schinieder H. *Garteur AD AG08 Final Report*. Garteur High Lift Action Group, Tech. Rep., 1989.
- [19] Moitra A. Issues in 2-D High-Lift CFD Analysis. *AIAA 21st Applied Aerodynamics Conference*, No. June, 2003.
- [20] Murayama M, Yamamoto K, and Kobayashi K. Validation of Flows on High-Lift Configurations by Structured and Unstructured Mesh Method. *43rd AIAA Aerospace Sciences Meeting and Exhibit*, Nevada, 2005.
- [21] ANSYS Inc. ANSYS 19.1 DesignXplorer User Guide, 2018.
- [22] Di Pasquale D, Zhu F, Cross M, Savill M, Kipouros T, Integrated system to perform surrogate based aerodynamic optimisation for high-lift airfoil, *EngOpt 2016 - 5th International Conference on Engineering Optimization*, Iguassu Falls, Brazil, 19 - 23 June 2016.

2021-09-10

Validating surrogate models and incorporating uncertainty quantification in multi-element airfoil design optimisation

Di Pasquale, Davide

International Council of the Aeronautical Sciences

Di Pasquale D, Bolart C. (2021) Validating surrogate models and incorporating uncertainty quantification in multi-element airfoil design optimisation. In: ICAS 2021: 32nd Congress of the International Council of the Aeronautical Sciences, 6-10 September 2021, Shanghai, China

<https://dspace.lib.cranfield.ac.uk/handle/1826/17364>

Downloaded from Cranfield Library Services E-Repository