

A Unified Multi-Fidelity Aero-Structural Design Framework for Novel Aircraft Configurations

Shubham Shubham ^{*}, Ankit Sharma [†], and Atif Riaz [‡]

Faculty of Engineering and Applied Sciences, Cranfield University, Bedford, MK43 0AL, United Kingdom

This study presents a novel framework that integrates multi-disciplinary simulation data, robust multi-objective optimisation (MOO), and interactive visual analytics to balance multiple objectives and cope with design uncertainties. Using the CRM wing and ATI Narrowbody aircraft geometry as test cases, aerodynamic (e.g., lift-to-drag ratio) and structural (e.g., wing mass) performance measures are combined within a Kriging-based surrogate modelling approach. Monte Carlo Simulation is employed to propagate design uncertainties, while NSGA-III guides the exploration of Pareto-optimal solutions using a robust optimisation framework. An HTML-based dashboard, powered by Python servers, allows users to dynamically inspect trade-offs, filter design alternatives, and compare competing objectives in near real-time. Results indicate that increasing the wing aspect ratio and reducing skin root thickness generally enhances aerodynamic efficiency and flight range, although structural weight and associated uncertainties also rise. By visualising such trade-offs interactively, designers gain insight into where higher fidelity analyses or additional data collection may be most beneficial, and which configurations deliver robust performance under probable variations. The resulting tool fosters a more efficient and informed engineering design process by consolidating data, computations, and insights within a single, accessible environment.

I. Introduction

In many engineering fields, design decisions must balance conflicting criteria, handle large-scale data from disparate sources, and mitigate uncertainties that arise from incomplete or noisy information [1]. Visual Analytics (VA) has emerged as a powerful paradigm to address these challenges by coupling interactive data visualization with advanced computational models [2]. In tandem, Multi-Objective Optimization (MOO) methods such as NSGA-II or NSGA-III provide systematic means of exploring trade-offs across potentially high-dimensional problem spaces [3, 4]. Integrating these methods with Uncertainty Quantification (UQ) ensures that design solutions remain robust when real-world variability is introduced, whether stemming from physical properties, operating conditions, or model fidelity [5].

Although these techniques have broad relevance, spanning automotive, civil infrastructure, and energy systems, aircraft design represents a particularly instructive example, where aerodynamics, structures, and propulsion each demand rigorous performance thresholds [6]. A multi-disciplinary, data-driven approach allows engineers to evaluate how aerodynamic efficiency, structural weight, and acoustic properties evolve under varying parameters, while robust optimization ensures resilience to modelling or manufacturing uncertainties [7]. By synthesizing these methods into a streamlined workflow, engineers gain improved insight into key trade-offs and potential design risks, ultimately driving more confident decision-making.

Aerospace systems are inherently constrained by stringent safety, cost, and performance requirements, placing a premium on effective Multi-Disciplinary Optimization (MDO). Small improvements in aircraft efficiency or mass distribution can yield substantial economic and environmental benefits, yet these gains must be pursued without compromising structural integrity or operational reliability. Traditional MDO workflows often struggle to communicate complex trade-offs, such as those between aerodynamic lift-to-drag ratio, structural load capacity, and fuel consumption, to diverse teams of designers, analysts, and decision-makers. This can often lead to information overload and hinder the extraction of meaningful insights.

Coupling Visual Analytics with robust MDO techniques alleviates these communication hurdles and fosters a collaborative environment where domain experts can rapidly iterate through Pareto-optimal solutions [2]. A web-based tool can enable clients and decision-makers to explore and compare design choices interactively, without requiring

^{*}Post-Doctoral Research Fellow, shubham.shubham@cranfield.ac.uk

[†]Post-Doctoral Research Fellow, ankitsharma@cranfield.ac.uk

[‡]Senior Lecturer, School of Aerospace, Transport and Manufacturing, a.riaz@cranfield.ac.uk

expertise in the backend code or the complexities of the optimization process. In particular, interactive dashboards and advanced surrogate modelling reduce reliance on exhaustive simulation runs and expose crucial patterns in vast engineering datasets. As a result, this combined framework supports rapid prototyping and “what-if” testing of new technologies or configurations, which can shorten design cycles by letting engineers immediately evaluate trade-offs [1].

Early MDO frameworks commonly relied on surrogate models and deterministic optimization schemes to accelerate simulation-based analyses [8]. While these efforts successfully reduced computational costs, they often lacked mechanisms to quantify and propagate uncertainties arising from manufacturing tolerances or incomplete physics modelling [9]. More recent studies have introduced probabilistic design approaches and robust optimization techniques [4], demonstrating improvements in decision-making speed and solution stability under parameter variability.

Keim et al. [10] highlight the growing importance of visual analytics in managing the massive and heterogeneous data volumes that are increasingly common in various domains. They emphasize that visual analytics combines visualization, data mining, human factors and statistical techniques. Siddiqui et al. [11] showed a visual data platform called zenvisage for easily finding desired visual patterns from large datasets. A study by Piotrowski et al. [12] addresses the issues of overplotting and data clutter in static displays. They introduced an interactive toolkit combining parallel coordinates and scatter plots, enhanced with machine learning methods to manage multidimensional engineering design data. In another work, Kipouros et al. [13] applied these principles to aerodynamic turbomachinery design, showcasing how parallel coordinates can manage high-dimensional design spaces effectively. This study highlighted the ability of parallel coordinates to identify crucial relationships between design parameters and objectives, facilitating intelligent decision-making and design space exploration. Sacha [14] presented a knowledge generation model for visual analytics, consisting of a computer and a human part. The computer part deals with analysing and visualising existing data generated using different models used within the system. The human part is a reasoning process composed of exploration, verification and knowledge generation loops.

Despite these advancements and multiple visual analytics tools present, there remains a gap in the literature: existing systems rarely integrate MDO, robust multi-objective optimization with uncertainty quantification, and visual analytics into a cohesive, end-to-end decision-making environment. Studies may focus on one aspect, such as advanced surrogate modelling or interactive visualization, without fully capturing the synergy that can arise from coupling robust design and multi-disciplinary data streams under a single web-based platform. A lot of companies like Airbus and Rolls-Royce use similar in-house tools and therefore, it is necessary to create an open-source framework using visual analytics. Furthermore, only a few reported frameworks incorporate second-order uncertainty, in which variations in fidelity or modelling accuracy themselves become part of the optimization process [15].

In light of these challenges, the primary objectives of this study are to integrate data from different disciplines such as aerodynamics and structures, while accounting for multi-objective optimisation employing the NSGA-III algorithm within a robust design paradigm to identify Pareto-optimal solutions that remain feasible under parameter fluctuations. And, during the process, develop an interactive web-based dashboard to visualize trade-offs and uncertainty propagation, enabling domain experts to rapidly evaluate design configurations without leaving the browser environment. The current study applies the framework to only aerospace applications, but it can be applied to any engineering problems.

The abstract is structured as follows. Section II described the methodology detailing the methods used for generating aerodynamic and structural data, workflow for data generation, surrogate modelling, multi-objective optimization, and robust design optimisation and visual analytics framework. Section III described the geometry of the aircraft and wing used in this study. Section IV presents results from the web-based tool and discusses the implications of the new framework for aircraft wing design. Finally, section V discusses the outline of the final manuscript.

II. Methodology

A. Overall Workflow

Initially, simulation data is generated in different disciplines and for different fidelities, depending on the user requirements. This is done using the Multi-Fidelity Multi-Disciplinary Optimisation (MFMDO) framework shown in Figure 1. The generated data is then processed into the visual analytics workflow, shown in Figure 2 that integrates the multi-fidelity multi-disciplinary data, surrogate modelling, and robust multi-objective optimization employing uncertainty quantification within a web-based visual analytics environment. It proceeds in six stages:

- 1) **Aerodynamic and aeroacoustic simulations:** For low-fidelity data, panel methods or modified strip theory for aerodynamics and Hanson or Amiet model for aeroacoustics are employed, providing rapid estimates of key performance metrics. These fast evaluations guide initial parameter screening before higher-fidelity computations,




Fig. 1 System-level overview of the proposed Multi-Fidelity Multi-Disciplinary Optimisation (MFMDO) workflow for generating initial data for Visual Analytics

- for which Navier-Stokes or Lattice Boltzmann method solve is used for aerodynamics and Ffowcs-Williams and Hawkings for aeroacoustics.
- 2) **Structural simulations:** Simplified beam models or finite element methods then provide structural analyses. This can include different types of stress or strain in the structure, object weight, maximum load bearing, etc.
 - 3) **Data Consolidation and Surrogate Modelling:** Results from both disciplines are merged, normalized, and stored in a unified repository. A Kriging surrogate model is trained on this combined dataset to be used for further analysis.
 - 4) **Multi-Objective Optimization:** Design trade-offs are solved (e.g., lift-to-drag vs. structural weight) with the algorithm NSGA-III. Candidate solutions are iteratively generated, evaluated through the surrogate model, and refined to converge on Pareto-optimal sets.
 - 5) **Uncertainty quantification:** Robust design principle is introduced by varying input design parameters such as aspect ratio and skin root thickness of aircraft wings. The optimizer identifies solutions that remain feasible across these perturbations, achieving robustness against real-world variability. In addition, the framework accommodates second-order uncertainty arising from multi-fidelity modelling, beyond the usual aleatory or epistemic uncertainties, thereby accounting for fidelity-based variability in any given design or operating condition.
 - 6) **Web-Based Visualization:** An HTML- and Python-powered front end delivers interactive plots (e.g., 1D or 2D Kriging results, Pareto fronts, parallel coordinates) and seamless user experience to upload multi-fidelity data. Designers can inspect, filter, and compare results, enabling rapid iteration and informed decision-making.




Fig. 2 System-level overview of the proposed Visual Analytics Framework

B. Physics-based models

Only low-fidelity aerodynamic and structural data is utilised in this study to show the fast iterative process of the framework, from generating data to analysing it in the visual analytics dashboard. However, the framework can easily accommodate multi-fidelity data or even experimental data. Open-source software QBlade is used to generate data.

1. Aerodynamic solver: Lifting Line Method

The lifting line method is adopted to compute both the aerodynamic forces and the flow field. Prior studies have confirmed its accuracy and efficiency for low Reynolds number rotor configurations [16–19]. Here, a three-dimensional (3D) lifting line theory is combined with a free vortex wake model, following the nonlinear lifting line formulation described by van Garrel [20] and further detailed in Balduzzi et al. [21].

An incompressible, inviscid, and irrotational fluid domain is assumed, while the wing is represented by a single line of vortices located at the quarter-chord points. In each time step, vortex line elements are shed from the trailing edges and move convectively through the flow (“free wake method”), where their positions evolve under the local velocity field—a superposition of inflow velocity and the induced velocity from all existing wake elements. The circulation on the bound vortices is determined from the airfoil lift and drag characteristics, thus introducing a nonlinear component to the lifting line formulation. Lift and drag forces are then computed from the local angle of attack, α .

For desingularisation, we apply van Garrel’s cut-off approach [22], setting a core radius sized to capture viscous effects using kinematic viscosity ν , a turbulent vortex viscosity coefficient δ_v , and a time offset S_c . The resulting core size r_c is defined as

$$r_c = \left(\frac{5.03\delta_v\nu(t + S_c)}{1 + \varepsilon} \right)^{1/2}, \quad (1)$$

where ε is a small parameter related to the cut-off procedure. This treatment ensures a more realistic representation of vortex diffusion in the free wake.

2. Aerodynamic solver: Vortex Lattice Method

A mid-fidelity aerodynamic database was generated using the OpenVSP parametric geometry environment coupled with its aerodynamic analysis tool VSPAERO [23, 24]. The aircraft/wing geometry is first parameterised in OpenVSP and tessellated into a surface discretisation suitable for VSPAERO. In VLM mode, the lifting surfaces are represented by a lattice of bound vortices with trailing wake elements, and the solver enforces a no-penetration boundary condition at surface control points to determine the circulation distribution [25]. Integrated quantities such as C_L , induced drag, and spanwise loading are then obtained efficiently across sweeps of angle of attack and Mach number; stability and control derivatives can also be computed through small perturbations about the trimmed state [23]. Where required, viscous/profile drag may be incorporated via auxiliary airfoil polar corrections, while retaining the rapid turnaround of the inviscid VLM core. The approach is therefore well-suited for conceptual and early preliminary design, design-space sampling, and surrogate-model training, but it is not intended to resolve strong viscous effects, separation/stall, or shock-dominated transonic phenomena [24, 25].

3. Structural solver

Structural analysis is performed by leveraging the finite element (FE) capabilities of the open-source multi-physics engine Project Chrono [26], a cross-platform C++ library designed around an object-oriented architecture [27]. Specifically, the Chrono::Engine module is employed, which handles the assembly and solution of physical systems governed by Newtonian mechanics and finite element methods. For solving FE problems, the SparseLU solver from the EIGEN C++ template library [28] is used.

Within the QBlade-Chrono coupling interface, Euler-Bernoulli beams are modelled in a co-rotational framework [29], enabling a local coordinate system to move with each deformable beam element. The beam’s rigid-body motion (translation and rotation) is captured in the floating frame, while small elastic deformations are measured relative to this local reference. Geometric stiffness contributions are included when assembling the global tangent stiffness matrix in Project Chrono. By default, user-specified aerodynamic and geometric properties are automatically processed to derive the wing’s structural discretization, which is then formed by interconnecting individual structural entities subject to appropriate constraints. Each cross-section in the model is defined through various parameters such as the Centre of Mass, Shear Center, Elastic Center, Principal Bending Axes, and Structural Pitch [7], ensuring that both mass and

stiffness characteristics are accurately reflected in the final structural configuration. The capabilities of this structural method are demonstrated in previous rotor and aircraft wing studies [15, 30, 31].

C. Surrogate modelling

Ordinary Kriging

The Kriging model assumes that points within proximity to each other within a design domain exhibit similar functional values. Within a d -dimensional hypercube, the correlation amongst points is assumed to follow the relationship given by [32]:

$$\mathbf{R}_{ij} = \exp \left(- \sum_{l=1}^d 10^{\theta(l)} \left\| \mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)} \right\|^{p(l)} \right), \quad (2)$$

where θ characterizes the correlation decay rate and p describes the 'smoothness' of the model across each dimension. These parameters, together with a regularization term, 10_λ , which is incorporated into the correlation matrix diagonal [8], can be calculated through the optimization of the concentrated log-likelihood function [33]:

$$\phi = -\frac{n}{2} \ln(\hat{\sigma}^2) - \frac{1}{2} \ln(|\mathbf{R}|). \quad (3)$$

The variance is defined by:

$$\hat{\sigma}^2 = \frac{1}{n} (\mathbf{y} - \mathbf{1}\hat{\mu})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu}), \quad (4)$$

while the mean of the observations is given by:

$$\hat{\mu} = \frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{y}}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}}, \quad (5)$$

where, $\mathbf{1}$ represents an n -dimensional vector of ones, and n equates to the quantity of data points within the sampling framework which is the foundation of the model architecture. Once the hyperparameters are established, the model's predictor

$$\mathbf{y}(\mathbf{x}^*) = \hat{\mu} + \mathbf{r}^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu}), \quad (6)$$

yields the estimated output at an unsampled location within the design space.

The optimization of hyperparameters constitutes a critical step in the development of Kriging models, significantly influencing the accuracy of the resultant model. However, the computational expense associated with inverting the correlation matrix $\mathbf{R}$, which scales in a cubic manner with the number of data points $O(n^3)$, becomes substantial for an extensive training dataset. To mitigate this computational demand and expedite the optimization process, Toal et al. [34] introduced an adjoint-based approach to the likelihood function, enabling efficient computation of the function's derivatives with respect to each hyperparameter. Utilizing the collection of adjoints provided by Giles [35], the adjoint of the correlation matrix $\mathbf{R}$ is expressed as:

$$\bar{\mathbf{R}} = \frac{1}{2\hat{\sigma}^2} \mathbf{R}^{-T} (\mathbf{y} - \mathbf{1}\hat{\mu})^T (\mathbf{y} - \mathbf{1}\hat{\mu})^T \mathbf{R}^{-T} - \frac{1}{2} \mathbf{R}^{-T}. \quad (7)$$

The derivatives of the log-likelihood function relative to θ are given as:

$$\frac{\partial \phi}{\partial \theta_l} = \ln 10 \sum_{ij} -10^{\theta_l} \left\| \mathbf{x}_{i_l} - \mathbf{x}_{j_l} \right\|^{p_l} \mathbf{R}_{ij} \bar{\mathbf{R}}_{ij}, \quad (8)$$

and relative to p are given as:

$$\frac{\partial \phi}{\partial p_l} = \sum_{ij} -10^{\theta_l} \left\| \mathbf{x}_{i_l} - \mathbf{x}_{j_l} \right\|^{p_l} \ln \left\| \mathbf{x}_{i_l} - \mathbf{x}_{j_l} \right\| \mathbf{R}_{ij} \bar{\mathbf{R}}_{ij}. \quad (9)$$

Should a regularization term be incorporated to account for noise in the training data, its derivative can also be calculated as:

$$\frac{\partial \phi}{\partial \lambda} = 10^\lambda \ln 10 \sum_i \bar{R}_{ii}. \quad (10)$$

To address the optimization challenge of hyperparameters, the hybridized particle swarm optimization algorithm proposed by Toal et al. [36] is implemented. This algorithm, an adaptation of the particle swarm technique originally formulated by Eberhart et al. [37], implements a particle initialization strategy based on an optimal Latin hypercube. A rank-based selection mechanism is employed to select a particle for localized improvement through sequential quadratic programming (SQP) in each generation (iteration).

D. Uncertainty quantification

Monte Carlo Simulation (MCS)

Monte Carlo Simulation (MCS) is a widely used approach for capturing and propagating aleatory uncertainties in engineering models [38]. Because it is non-intrusive, it lends itself to uncertainty quantification, design optimization, and related analytical tasks [39, 40]. This property allows an existing deterministic solver to be reused for stochastic evaluations by repeatedly running it with randomly sampled input parameters. In many cases, MCS proves advantageous over other uncertainty quantification methods, such as generalized Polynomial Chaos (gPC) [41], because it does not require specialized code modifications.

Under MCS, multiple realizations (or samples) of random input parameters are generated according to their respective probability distributions. Each realization defines a deterministic problem instance, which is solved to produce an output dataset [42–44]. Statistical evaluations are then carried out on this ensemble to characterize the system’s response [45]. The unbiased nature of MCS across varying levels of accuracy renders it a good benchmark for evaluating the performance of alternative methodologies. Typically, one focuses on calculating first-order statistics (e.g., mean), though higher-order or distributional features can also be explored. The mean π is defined as

$$\pi = E[f(\mathbf{x})] = \int f(\mathbf{x})p(\mathbf{x})d\mathbf{x}, \quad (11)$$

where $f(\mathbf{x})$ is the output function of interest and $p(\mathbf{x})$ is the probability density function of the random vector $\mathbf{x}$. Directly sampling from $p(\mathbf{x})$ and computing an empirical average

$$\hat{\pi} = \frac{1}{N} \sum_{i=1}^N w_i f(X^{(i)}), \quad (12)$$

enables the convergence of this estimate as $N \rightarrow \infty$. N is independent distributed random samples $\{X^{(i)}; i = 1, \dots, N\}$, each with associated quadrature weights $\{w_i; i = 1, \dots, N\}$. Although alternative numerical quadrature methods exist [46–48], MCS remains popular due to its simplicity and unbiased nature [49], despite the computational burdens that can arise for large sample sizes [42–44].

Additionally, epistemic uncertainty arises when knowledge of certain model parameters is incomplete or imprecise, often as a result of limited data or imperfect modelling assumptions. In this framework, interval analysis is used to address such epistemic sources by bounding uncertain variables within plausible minimum and maximum values, helping to identify how assumptions or data gaps affect system performance outcomes.

E. Multi-Objective Optimisation

This study uses the NSGA-III (Nondominated Sorting Genetic Algorithm III) which is a recent extension of the popular NSGA-II method originally developed for handling multi-objective optimization problems [3]. Unlike its predecessor, NSGA-III is specifically designed to tackle many-objective optimization scenarios where the number of objectives (e.g., four or more) makes conventional crowding-distance mechanisms less effective [4].

NSGA-II introduced the idea of fast nondominated sorting, elitism, and crowding distance to maintain a diverse set of solutions. However, as the dimensionality of the objective space grows, relying solely on crowding distance becomes insufficient for preserving population diversity [3]. NSGA-III addresses this challenge by incorporating a set of predefined reference points (or reference lines) in the objective space to guide the selection of offspring [4]. These reference points act as “anchors” that encourage the search to explore under-represented regions of the Pareto front, ultimately improving both convergence speed and solution distribution in higher-dimensional problems.

During each generation, NSGA-III applies nondominated sorting similar to NSGA-II but leverages the reference points to perform niche-preservation. This mechanism systematically checks how well solutions align or cluster around these reference points, ensuring that the algorithm maintains diverse trade-offs among the objectives [4]. As a result, NSGA-III can more reliably discover a spread of Pareto-optimal solutions even when the dimensionality of the objectives grows beyond three or four. In practical applications, NSGA-III has proven beneficial over NSGA-II for many-objective problems in engineering, logistics, and data analytics, making it a robust choice for advanced multi-objective optimization tasks.

III. Numerical setup

A. Geometry and flow conditions

Two different test cases are demonstrated in this study. First, the Common Research Model (CRM) wing is redesigned (Figure 3). The baseline CRM wing geometry and its aerodynamic data are presented by Vassberg et al. [6]. The portion beyond the crank is rescaled to accommodate variations in aspect ratio, and the overall skin thickness distribution is derived from the baseline root thickness. This simplified strategy strikes a balance between limiting the number of design variables and achieving a satisfactory structural configuration. The chosen input parameter bounds (Table 1) are sampled with 100 points using a Latin Hypercube scheme under uniform distributions. Subsequently, an augmentation step interpolates the initial data to produce 1,000 additional design points. Throughout this study, the true airspeed of the wing is kept at 240 m/s and the angle of attack at 0° .




Fig. 3 Top view of the baseline CRM wing, with the two modifications highlighted in blue and red; adopted from Spinelli et al.

Table 1 Input parameters bounds for the CRM wing

Parameter	Lower Bound	Upper Bound
Wing Aspect Ratio	7	12
Root Skin Thickness [mm]	5	30

The second test case uses the ATI Narrowbody aircraft geometry, shown in Figure 4. Table 2 shows the parameter bounds for the two operational parameters varied, angle of attack and Mach number. Full factorial sampling was employed, resulting in a total of 112 design points being sampled.




Fig. 4 ATI Narrowbody aircraft geometry

Table 2 Input parameters bounds for the ATI Narrowbody aircraft

Parameter	Lower Bound	Upper Bound
Angle of Attack	0	15
Mach number	0.1	0.7

IV. Results

A. CRM wing

Figure 5 presents four surrogate model contour plots, where the left column (a, c) depicts the predicted values for aircraft range and wing mass (objective parameters), while the right column (b, d) shows the corresponding uncertainties in the whole design space comprising aspect ratio (AR) and skin root thickness. The latter shows the design uncertainty that is expected to be present if any particular pair of design points is selected. The surrogates were constructed from the mid-fidelity datasets in this study, and each point in the design space represents a pair of aspect ratio and skin root thickness. To obtain the plots, the data were uploaded to the interactive dashboard in the web-based tool, and the 2D Kriging model was chosen instead of 1D Kriging since two different design parameters (AR and skin root thickness) are present (the web interface allows the user to select between 1D and 2D, depending on the application).

To interpret the range plots, the Breguet range equation is used:

$$\text{Range} = V t_f = V \times \underbrace{\left(\frac{L}{D}\right)}_{\text{aircraft designer}} \times \underbrace{I_{sp}}_{\text{propulsion system designer}} \times \underbrace{\ln\left(\frac{W_i}{W_f}\right)}_{\text{structural designer}}.$$

where W_i and W_f are the initial and final aircraft weights, I_{sp} is overall efficiency, and L/D is the aerodynamic efficiency. In Sub-figure 5 (a), the surrogate indicates that a higher AR and smaller skin thickness yield a longer flight range. Meanwhile, the associated uncertainty Sub-figure 5 (b) reveals comparatively low percentage uncertainty in regions with large aspect ratios; variations in skin thickness do not significantly alter the uncertainty or confidence level in aircraft wing design. So an optimised wing design should have higher AR and low skin thickness, prompting a stronger manufacturing material to be used.

For wing mass (Sub-figures 5 c and d), mass is relatively insensitive to AR but rises with increasing skin root thickness which is expected from a structural point of view. The uncertainty map (d) shows an inverse relationship, where decreasing thickness leads to higher percentage uncertainty, while AR contributes little to uncertainty variations. This shows an inverse relationship between the physical parameter and uncertainty in that parameter, prompting a balanced approach in choosing the skin thickness of the wing.

Figure 6 illustrates the different types of sampling distributions using the quasi-random Halton sequence available for user selection through the web-based interactive dashboard. Halton sequence is selected since it was found to converge




Fig. 5 2D Kriging surrogate model results using initial data provided by the user and % uncertainty obtained for the 2D design space; AR - aspect ratio

faster as compared to other sequences such as random or Sobol [15]. Sub-figure (a) shows a combination of a normal (Gaussian) distribution for aspect ratio and a uniform distribution for skin root thickness. This choice reflects a common approach when certain parameters, like aspect ratio, are expected to exhibit natural variability (especially around a central value), while others, like skin root thickness, are deliberately constrained within a known range. Sub-figure (b) demonstrates the case where both design parameters follow a normal distribution. This may represent a scenario where both parameters are modelled as uncertainties, where the mean value and standard deviation are estimated based on experimental or historical data. The dashboard allows users to toggle between these distribution options and analyse how variations in sampling assumptions influence uncertainty propagation. The quantified uncertainties from these sampling distributions are critical inputs for the robust design optimisation step in our methodology, as described in Figure 2. By selecting different distributions, users can evaluate the effect of sampling uncertainty from each design discipline (e.g. aerodynamics, structures, and propulsion) and incorporate them into the robust design process. The resulting optimisation considers the propagated uncertainties and identifies solutions that remain resilient under varied design and operational conditions.

Figure 7 presents Pareto fronts generated by the interactive dashboard, where each sub-figure pairs two different objectives chosen by the user. For this study, the choices include the 3 physical objectives (wing mass, lift-to-drag ratio, or range) with their corresponding uncertainties. Ideally, for an aircraft, range and L/D need to be high and wing mass




Fig. 6 Different types of sampling distribution for uncertainty quantification using the Halton sequence

low, whereas uncertainty in all parameters needs to be as low as possible. The negative sign for parameters such as L/D or range indicates that the algorithm is set to maximize them. Sub-figure 7 (a) shows the trade-off between maximizing L/D and minimizing wing mass: users can opt for a lightweight design (e.g. ≈ 6000 kg) with a moderate L/D (≈ 17.3) or a heavier design (e.g. ≈ 13500 kg) with a higher L/D (≈ 17.44), depending on material, geometry, or aerodynamic constraints. Similarly, sub-figure 7 (b) contrasts the absolute range with its uncertainty, revealing how longer ranges (e.g. ≈ 8100 nmi) often come with slightly higher uncertainty ($\approx 0.48\%$). The final choice of parameter values may depend on the system design and range requirements of the aircraft.




Fig. 7 Results for multi-objective optimisation showing Pareto fronts and design trade-offs for a robust design

Such analyses help pinpoint balanced compromises by simultaneously weighing performance (e.g., wing mass, L/D or range) against uncertainty or confidence in those predictions. Ultimately, a multi-objective optimization approach is critical in engineering design because real systems, like aircraft wings, are governed by multiple competing criteria

under uncertain conditions. By embedding these capabilities into a web-based tool, designers can rapidly iterate through potential solutions, adjust key parameters, and immediately visualize trade-offs, substantially enhancing design exploration and collaborative decision-making.

B. ATI Narrowbody model

Figure 8 summarises the 2D Kriging prediction for the ATI narrowbody aerodynamic database, with Mach number and angle of attack (AoA) as inputs and the training samples overlaid as markers. The surrogate-predicted C_L exhibits a smooth, predominantly monotonic increase with AoA across the investigated range, with a secondary dependence on Mach; at higher Mach and higher AoA, the model predicts the largest C_L values, consistent with increased aerodynamic loading under compressible flow conditions. The corresponding C_D surface similarly increases with both Mach and AoA, with stronger curvature in Mach than in C_L , indicating that compressibility effects drive a more pronounced drag rise as the flow approaches the upper end of the Mach envelope. The right-hand figures show that the uncertainty increases towards lower AoA values, reaching the highest levels near extreme Mach/AoA combinations. Notably, the C_D uncertainty is larger than the C_L uncertainty, reflecting the higher nonlinearity/sensitivity of drag with respect to Mach and AoA and the greater difficulty of constraining C_D with limited data.




Fig. 8 2D Kriging surrogate model results for the ATI Narrowbody case study

Figure 9 reports the outcome of the robust multi-objective optimisation. In Figure 9 (a), the objective pair is aerodynamic efficiency expressed through the simultaneous minimisation of C_D and maximisation of C_L (shown with

the conventional sign change in the optimisation formulation, hence the negative C_L values on the abscissa). The red markers denote the simulated (evaluated) design points explored during the search, while the blue markers highlight the subset of non-dominated solutions. The resulting Pareto front exhibits a clear monotonic trade-off: designs achieving higher lift (moving left along the C_L axis) incur a progressively higher drag penalty, whereas the lowest-drag solutions correspond to lower achievable lift. This curve therefore, defines the attainable aerodynamic envelope under the imposed constraints and provides the decision-maker with a discrete set of compromise candidates rather than a single “best” design.




Fig. 9 Results for multi-objective optimisation showing Pareto fronts and design trade-offs for a robust design

Figure 9 (b) extends the analysis to robustness by reporting lift uncertainty (in %) as a function of C_L for the same non-dominated set. A systematic trend is observed whereby configurations delivering higher lift are associated with larger predictive/propagated uncertainty, while lower-lift solutions show reduced uncertainty. This behaviour is consistent with robust design principles: as the optimiser pushes towards more aggressive operating points (higher loading), the solutions become increasingly sensitive to modelling discrepancies and input variability, leading to wider uncertainty bands. Together, the two panels show that the proposed framework not only quantifies the classical performance trade-off between C_L and C_D , but also exposes the associated robustness penalty, enabling selection of designs that balance aerodynamic performance against uncertainty (e.g., preferring a slightly lower-lift solution if it yields substantially lower uncertainty and therefore higher confidence in meeting constraints).

The computational gain of the Kriging surrogate is quantified through the Surrogate Speed-up Index (SSI), defined as the ratio between the cost of performing all evaluations with the physics-based solvers and the cost of the same study using the surrogate, including its one-off cost. In the CRM wing example, a single mid-fidelity aero-structural evaluation of wing mass, range and L/D takes $T_{\text{phys}} = 10$ s, whereas a Kriging prediction for the same quantities requires only $T_K = 3 \times 10^{-4}$ s. Training the surrogate once costs $T_{\text{train}} = 30$ s. For an optimisation study using Monte Carlo with $N_{\text{eval}} = 10^4$ samples, the traditional approach would require $T_{\text{phys,total}} = N_{\text{eval}} T_{\text{phys}} = 10^5$ s (about 28 h) of solver time, while the surrogate-based workflow costs only $T_{K,\text{total}} = T_{\text{train}} + N_{\text{eval}} T_K = 30 + 3 = 33$ s. The resulting index is therefore

$$\text{SSI} = \frac{T_{\text{phys,total}}}{T_{K,\text{total}}} \approx 3.0 \times 10^3,$$

i.e. the same information on wing mass, range and L/D is obtained roughly three orders of magnitude faster. This illustrates the key difference between the proposed methodology and a traditional wing-design loop: instead of being limited to a small number of expensive simulations and relying on the designer intuition to extrapolate the physics, the Kriging surrogate allows us to map the entire design space (via Monte Carlo analysis and multi-objective optimisation) with near-interactive turnaround, while still retaining the underlying physics learned from the original mid-/high-fidelity data.

V. Conclusions

This paper demonstrates an integrated framework that combines multi-disciplinary simulation data, robust multi-objective optimisation, and uncertainty quantification. Motivated by the complexity of balancing aerodynamics, structures, and other design considerations under uncertainty, a web-based tool (HTML+Python) was developed, incorporating Kriging surrogate modelling and NSGA-III algorithm for multi-objective MDO problems.

The Pareto fronts and contour plots revealed tangible trade-offs between competing objectives (e.g., maximising lift-to-drag ratio vs. minimising wing mass) and illuminated localised design optima. The approach captured different design uncertainties propagated, demonstrating how design choices in one discipline (e.g., root skin thickness of wing or angle of attack of an aircraft) can increase or decrease uncertainty or confidence in another (e.g., aircraft range or lift coefficient). The surrogate modelling significantly accelerated the design cycle, enabling real-time “what-if” scenarios, rapid exploration and immediate feedback on the feasibility or robustness of a given solution. Although the present work has focused on a simplified wing and aircraft design application, the underlying architecture, particularly the hybrid data and modelling pipeline, lends itself to broader engineering contexts. For instance, the same framework could be applied to wind turbine blade optimisation, renewable energy systems, or structural health monitoring. The dashboard can be made more interactive by incorporating a 3D Kriging surrogate to visualise a surrogate surface or a 3-objective optimisation to visualise a Pareto surface in 3D. By harnessing a web-based environment, domain experts gain immediate access to large simulation datasets and can fine-tune or extend the analysis without specialised coding skills.

For future studies, high-fidelity or experimental data can be tested in addition to existing low-fidelity sources to form a multi-fidelity dataset. This would help validate and refine surrogate models and enhance predictive accuracy. The dashboard can be enriched with 3D surrogate model plots and three-objective MOO algorithms to visualise 3D surfaces, supporting deeper insights into higher-dimensional trade-off spaces. The user’s interactive visualisation features can be extended (e.g., brushing, linking, or VR/AR integrations) to foster advanced scenario comparisons and collaborative real-time exploration. Advanced epistemic or interval-based uncertainty analyses can be incorporated for domains with sparse data or strongly non-linear model behaviour.

Acknowledgments

This project has received funding from Innovate UK under Grant Agreement No 10003388.

References

- [1] Thomas, J. J., *Illuminating the path: [the research and development agenda for visual analytics]*, IEEE Computer Society, 2005.
- [2] Sacha, D., Stoffel, A., Stoffel, F., Kwon, B. C., Ellis, G., and Keim, D. A., “Knowledge generation model for visual analytics,” *IEEE transactions on visualization and computer graphics*, Vol. 20, No. 12, 2014, pp. 1604–1613.
- [3] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., “A fast and elitist multiobjective genetic algorithm: NSGA-II,” *IEEE transactions on evolutionary computation*, Vol. 6, No. 2, 2002, pp. 182–197.
- [4] Deb, K., and Jain, H., “An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints,” *IEEE transactions on evolutionary computation*, Vol. 18, No. 4, 2013, pp. 577–601.
- [5] Forrester, A. I., Söbester, A., and Keane, A. J., “Multi-fidelity optimization via surrogate modelling,” *Proceedings of the royal society a: mathematical, physical and engineering sciences*, Vol. 463, No. 2088, 2007, pp. 3251–3269.
- [6] Vassberg, J., Dehaan, M., Rivers, M., and Wahls, R., “Development of a common research model for applied CFD validation studies,” *26th AIAA applied aerodynamics conference*, 2008, p. 6919.
- [7] Marten, D., “QBlade: a modern tool for the aeroelastic simulation of wind turbines,” 2020.
- [8] Forrester, A. I., Keane, A. J., and Bressloff, N. W., “Design and analysis of “Noisy” computer experiments,” *AIAA journal*, Vol. 44, No. 10, 2006, pp. 2331–2339.
- [9] Martins, J. R., and Lambe, A. B., “Multidisciplinary design optimization: a survey of architectures,” *AIAA journal*, Vol. 51, No. 9, 2013, pp. 2049–2075.
- [10] Keim, D. A., Mansmann, F., Schneidewind, J., Thomas, J., and Ziegler, H., *Visual analytics: Scope and challenges*, Springer, 2008.

- [11] Siddiqui, T., Kim, A., Lee, J., Karahalios, K., and Parameswaran, A., "Effortless data exploration with zenvisage: an expressive and interactive visual analytics system," *arXiv preprint arXiv:1604.03583*, 2016.
- [12] Piotrowski, W., Kipouros, T., and Clarkson, P. J., "Enhanced interactive parallel coordinates using machine learning and uncertainty propagation for engineering design," *2019 15th International Conference on eScience (eScience)*, IEEE, 2019, pp. 339–348.
- [13] Kipouros, T., Inselberg, A., Parks, G., and Savill, A. M., "Parallel coordinates in computational engineering design," *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2013, p. 1750.
- [14] Sacha, D., "Knowledge generation in visual analytics: integrating human and machine intelligence for exploration of Big Data," 2018.
- [15] Shubham, S., Kipouros, T., Dash, S., and Ianakiev, A., "Uncertainty propagation and management of mixed uncertainties for multi-fidelity multi-disciplinary analysis of propeller with different blade sweep," *AIAA SCITECH 2024 Forum*, 2024, p. 1487.
- [16] Shubham, S., Wright, N., Avallone, F., and Ianakiev, A., "Aerodynamic and aeroacoustic investigation of vertical axis wind turbines with different number of blades using mid-fidelity and high-fidelity methods," *AIAA AVIATION 2023 Forum*, 2023, p. 3642.
- [17] Shubham, S., Avallone, F., Brandetti, L., Wright, N., and Ianakiev, A., "Effect of struts and central tower on aerodynamics and aeroacoustics of vertical axis wind turbines using mid-fidelity and high-fidelity methods," *AIAA SciTech Forum*, 2024.
- [18] Thambidurai Arasi, T. R., Shubham, S., and Ianakiev, A., "Effect of Blade Shape on Aerodynamic and Aeroacoustic characteristics of Vertical Axis Wind Turbines using mid-fidelity and high-fidelity methods," *AIAA SCITECH 2024 Forum*, 2024, p. 1488.
- [19] Grande, E., Shubham, S., Avallone, F., Ragni, D., and Casalino, D., "Computational aeroacoustic study of co-rotating rotors in hover," *Aerospace Science and Technology*, Vol. 153, 2024, p. 109381.
- [20] Van Garrel, A., "Development of a wind turbine aerodynamics simulation module," 2003.
- [21] Balduzzi, F., Marten, D., Bianchini, A., Drofelnik, J., Ferrari, L., Campobasso, M. S., Pechlivanoglou, G., Nayeri, C. N., Ferrara, G., and Paschereit, C. O., "Three-dimensional aerodynamic analysis of a Darrieus wind turbine blade using computational fluid dynamics and lifting line theory," *Journal of Engineering for Gas Turbines and Power*, Vol. 140, No. 2, 2018.
- [22] Marten, D., "QBlade Guidelines v0. 95," *Technical University of Berlin, Berlin, Germany*, 2016.
- [23] "OpenVSP Documentation," <https://openvsp.org/docs/>, ????. Accessed: 2025-12-19.
- [24] "OpenVSP Workshop: VSPAERO Ground School," <https://openvsp.org/wiki/doku.php?id=workshop2022>, ????. Accessed: 2025-12-19.
- [25] Katz, J., and Plotkin, A., *Low-Speed Aerodynamics: From Wing Theory to Panel Methods*, 2nd ed., Cambridge University Press, 2001.
- [26] Tasora, A., Serban, R., Mazhar, H., Pazouki, A., Melanz, D., Fleischmann, J., Taylor, M., Sugiyama, H., and Negrut, D., "Chrono: An open source multi-physics dynamics engine," *High Performance Computing in Science and Engineering: Second International Conference, HPCSE 2015, Soláň, Czech Republic, May 25-28, 2015, Revised Selected Papers 2*, Springer, 2016, pp. 19–49.
- [27] Projectchrono.org, "Project Chrono - An Open-Source Physics Engine," URL <https://projectchrono.org/about/>, 2019.
- [28] Tuxfamily, "EIGEN - a C++ template library for linear algebra : matrices, vectors, numerical solvers and related algorithms," URL <http://eigen.tuxfamily.org>, ????
- [29] Recuero, A., and Negrut, D., "Co-rotational formulation in chrono," *ProjectCHRONO, Technical documentation, Università di Parma, Parma, Italy, accessed Feb*, Vol. 19, 2016, p. 2018.
- [30] Shubham, S., Spinelli, A., and Kipouros, T., "Multi-fidelity multi-disciplinary optimisation of propeller design by visual analytics," *Computer Graphics Visual Computing (CGVC)*, 2024.
- [31] Shubham, S., Spinelli, A., and Kipouros, T., "Visual Analytics Framework for Multi-Disciplinary Multi-Objective Optimisation of Aircraft Design," *Available at SSRN 5114731*, ????

- [32] Toal, D. J., "Applications of multi-fidelity multi-output Kriging to engineering design optimization," *Structural and Multidisciplinary Optimization*, Vol. 66, No. 6, 2023, p. 125.
- [33] Jones, D. R., Schonlau, M., and Welch, W. J., "Efficient global optimization of expensive black-box functions," *Journal of Global optimization*, Vol. 13, 1998, pp. 455–492.
- [34] Toal, D. J., and Keane, A. J., "Efficient multipoint aerodynamic design optimization via cokriging," *Journal of Aircraft*, Vol. 48, No. 5, 2011, pp. 1685–1695.
- [35] Giles, M. B., "Collected matrix derivative results for forward and reverse mode algorithmic differentiation," *Advances in Automatic Differentiation*, Springer, 2008, pp. 35–44.
- [36] Toal, D. J., Bressloff, N., Keane, A., and Holden, C., "The development of a hybridized particle swarm for kriging hyperparameter tuning," *Engineering optimization*, Vol. 43, No. 6, 2011, pp. 675–699.
- [37] Eberhart, R., and Kennedy, J., "A new optimizer using particle swarm theory," *MHS'95. Proceedings of the sixth international symposium on micro machine and human science*, Ieee, 1995, pp. 39–43.
- [38] Metropolis, N., and Ulam, S., "The monte carlo method," *Journal of the American statistical association*, Vol. 44, No. 247, 1949, pp. 335–341.
- [39] Helton, J. C., Johnson, J. D., Sallaberry, C. J., and Storlie, C. B., "Survey of sampling-based methods for uncertainty and sensitivity analysis," *Reliability Engineering & System Safety*, Vol. 91, No. 10-11, 2006, pp. 1175–1209.
- [40] Landau, D., and Binder, K., *A guide to Monte Carlo simulations in statistical physics*, Cambridge university press, 2021.
- [41] Xiu, D., and Karniadakis, G. E., "The Wiener–Askey polynomial chaos for stochastic differential equations," *SIAM journal on scientific computing*, Vol. 24, No. 2, 2002, pp. 619–644.
- [42] Liu, J. S., and Liu, J. S., *Monte Carlo strategies in scientific computing*, Vol. 75, Springer, 2001.
- [43] Shonkwiler, R. W., and Mendivil, F., *Explorations in monte carlo methods*, Springer Science & Business Media, 2009.
- [44] Robert, C. P., Casella, G., and Casella, G., *Monte Carlo statistical methods*, Vol. 2, Springer, 1999.
- [45] Cunha Jr, A., Nasser, R., Sampaio, R., Lopes, H., and Breitman, K., "Uncertainty quantification through the Monte Carlo method in a cloud computing setting," *Computer Physics Communications*, Vol. 185, No. 5, 2014, pp. 1355–1363.
- [46] Davis, P. J., and Rabinowitz, P., *Methods of numerical integration*, Courier Corporation, 2007.
- [47] Li, K. S., and Lumb, P., "Reliability analysis by numerical integration and curve fitting," *Structural safety*, Vol. 3, No. 1, 1985, pp. 29–36.
- [48] Lin, Y., "A quadrature-based technique for robust design with computer simulations," Ph.D. thesis, Massachusetts Institute of Technology, 2007.
- [49] Hu, X., Chen, X., Parks, G. T., and Yao, W., "Review of improved Monte Carlo methods in uncertainty-based design optimization for aerospace vehicles," *Progress in Aerospace Sciences*, Vol. 86, 2016, pp. 20–27.



DSPACE

<https://dspace.org/>

A unified multi-fidelity aero-structural design framework for novel aircraft configurations

Shubham, Shubham; Sharma, Ankit; Riaz, Atif

2026-01-12

American Institute of Aeronautics and Astronautics (AIAA)

<https://doi.org/10.2514/6.2026-2651>

Shubham S, Sharma A, Riaz A. (2026) A unified multi-fidelity aero-structural design framework for novel aircraft configurations. In: AIAA SCITECH 2026 Forum, 12-16 January 2026, Orlando, Florida, USA, Article number 2026-2651