

AN AUTOMATED PARAMETRIC FRAMEWORK FOR EARLY-STAGE STRUCTURAL DESIGN AND OPTIMISATION OF COMPOSITE AIRFRAMES

Nikolaos Ziakos[1] & Andrea Cini[1]

[1]Department of Aerospace Engineering, Universidad Carlos III de Madrid, Avd. de la Universidad, 30 Leganés, 28911, Madrid, Spain

Abstract

During aircraft conceptual design, it is necessary to explore and evaluate multiple alternatives rapidly, yet current methods, relying on empirical formulas and legacy data, offer limited accuracy for novel configurations and composite structures. An automated parametric framework for the early-stage design of composite airframes, combining rapid model generation tools with physics-based composite modelling and data-efficient surrogate-based optimisation methods, is presented in this paper. Python modules coupled with HyperMesh TCL scripting generate analysis-ready GFEMs of wingbox and fuselage-section structures, automatically assigning materials, properties, loads, and boundary conditions. The framework is demonstrated on a wingbox benchmark using a nested layout–sizing formulation, where component sizing is handled by OptiStruct in the inner loop together with a staged multi-run strategy reducing solver-induced variability. A Gaussian-process surrogate is used to predict the optimised wingbox mass for different layout variables, while a Bayesian Optimisation loop guides the selection of new layouts under a limited evaluation budget. Results show that the suggested workflow identifies a competitive low-mass layout, in agreement with the outcome of a global-search baseline algorithm, while requiring approximately 80 % fewer expensive evaluations.

Keywords: Parametric modelling, Global Finite Element Model, Composite Structures, Structural Optimisation, Surrogate-Based Optimisation

1. Introduction

The development of lightweight and structurally efficient airframes remains a central objective of aircraft design, due to the direct effect of structural mass on fuel consumption, emissions, and operating cost. Composite materials provide significant opportunities for structural integration and weight reduction, but they also increase modelling and optimisation complexity because of the large number of geometric, layout, and laminate-related design parameters. During conceptual and preliminary design, aircraft structures are still commonly sized using handbook methods, statistical relationships, and simplified weight-estimation models [1, 2]. These approaches are computationally efficient and suitable for rapid trade-off studies, but their accuracy is limited for unconventional configurations and composite structures, where legacy data are often not available.

Finite Element (FE)-based structural analysis and optimisation can improve early estimates of mass, stresses, strains, and stability eigenvalues. However, their use in early-stage design is limited by two practical bottlenecks. First, structural models must be generated and modified rapidly as geometric and layout parameters change. Second, each candidate design requires a constrained sizing optimisation, making exhaustive design-space exploration computationally expensive. This motivates the development of workflows combining automated parametric GFEM generation with data-efficient optimisation strategies.

Several parametric tools have been developed to support rapid aircraft geometry definition and preliminary design. Geometry-level frameworks, such as Vehicle Sketch Pad (VSP) [3, 4], Engineering Sketch Pad (ESP) [5] and FAST-OAD [6, 7], provide rapid CAD-based geometry generation and enable conceptual trade-off studies and weight estimation. Other approaches extend parametric modelling toward structural model generation. Sensmeier and Samareh [8] developed OptWing, which can automatically generate wingbox layouts for early-stage weight estimation, while Tang et al. [9] demonstrated a knowledge-based tool for wing FE model generation. Tools, such as *ELWIS* [10], generate detailed wing Global Finite Element Models (GFEMs) for early weight estimation and sizing integration using CPACS.

Efforts toward broader airframe model generation have also been reported. Mason and Quinlan [11] presented the Conceptual Design Shop (CDS) for rapid airframe structural modelling. Hwang et al. [12] developed a geometry and structural modelling framework for high-fidelity aircraft conceptual design optimisation. Parametric fuselage modelling has also been investigated, for example by Schwinn et al. [13] for preliminary sizing and crashworthiness applications. These works demonstrate the growing importance of automated structural model generation. Despite this progress, tools combining automated parametric generation of composite GFEMs with data-efficient coupled structural layout–sizing optimisation remain limited.

Aircraft structural layout optimisation introduces an additional difficulty because layout variables, such as the number of ribs or stringers, are integers, whereas component sizing variables are continuous or can sometimes be treated as such for composite laminates. Solving the full coupled problem directly can be computationally demanding, particularly when each evaluation involves high-fidelity FE analysis. A practical strategy is to use nested or multilevel optimisation formulations, where the outer loop explores the structural layout and the inner loop performs component sizing using a gradient-based optimiser [14, 15]. Although this approach separates integer layout variables from sizing variables, each outer-loop evaluation still requires a computationally expensive inner-loop constrained sizing optimisation.

Surrogate models, which can be used to approximate the computationally expensive evaluation with a cheaper data-driven model, have seen rapid growth in the literature [16, 17]. Gaussian Process Regression (GPR) models, also known as kriging, are particularly attractive for low-data engineering problems as they combine flexible nonlinear approximation with predictive uncertainty quantification [18, 19]. In this context, GPR models are commonly used for surrogate-based Bayesian Optimisation (BO), which provides an efficient strategy for reducing the number of expensive evaluations required during design exploration [20, 21]. The Efficient Global Optimisation (EGO) framework, introduced by Jones et al. [22], established Expected Improvement (EI) as a standard acquisition criterion for efficient BO, balancing the exploitation of promising regions with the exploration of uncertain ones, while naturally accommodating constraints and derivative information [20, 23].

Surrogate-based optimisation has been successfully integrated into various aerospace studies. Priem et al. [24] demonstrated that BO can be embedded in an industrial aircraft MDO framework to reduce the number of high-fidelity evaluations. Using the nested optimisation approach, Stanford et al. [15, 25] employed surrogate models for wingbox layout–sizing optimisation studies. Similarly, the AMIEGO and AMIEGO-MIMOS frameworks [26, 27] targeted large-scale problems and multiple infill criteria. More broadly, parallel surrogate-assisted global optimisation has been identified as a key strategy for expensive engineering problems where multiple high-fidelity evaluations can be run concurrently [28].

Although nested layout–sizing optimisation frameworks have been demonstrated, their application to detailed composite airframe FE models remains limited. In addition, when the inner-loop sizing problem is solved by a gradient-based optimiser, solver-induced variability and convergence to different local minima can affect the consistency of the mass values returned to the outer loop. This variability can degrade surrogate training and infill selection if the observed response is treated as fully deterministic. Therefore, the surrogate-based infill strategy can benefit from an explicit treatment of observation noise, as demonstrated in [29].

An automated parametric framework for composite airframe GFEM generation and surrogate-assisted structural layout exploration is presented below. Python modules coupled with HyperMesh TCL

scripting are used to generate wingbox and fuselage-segment structural models, assign materials and properties, create mesh connectivity, and prepare FE analysis or sizing optimisation input files. The optimisation capability is demonstrated on the Simple Transonic Wing (STW) benchmark [30]. A nested layout–sizing formulation is adopted, where the outer loop varies the number of ribs, the number of stringers per skin panel, and the spar-cap width, while the inner loop performs composite sizing optimisation in OptiStruct [31]. A GP surrogate trained on a limited number of optimised layouts is subsequently used for BO. The resulting design is compared with a nested global-search reference to assess both solution quality and computational efficiency.

The remainder of this paper is organized as follows. Section 2. presents the parametric airframe modelling framework, including the wingbox and fuselage-section generation modules while section 3. describes the STW benchmark, the composite sizing setup, and the nested layout–sizing formulation. Section 4. introduces the sampling strategy, GP surrogate model, and BO procedure, followed by Section 5. presenting the STW results, including surrogate accuracy, variable trends, BO convergence, and comparison with the global-search reference. Section 6. summarises the main findings and outlines future work.

2. Parametric Airframe modelling Framework

The suggested nested optimisation methodology relies on an automated procedure to generate the required GFEM models and prepare the corresponding composite sizing optimisation setup without manual intervention. The automated airframe modelling framework is implemented through in-house Python modules defining the geometry, structural layout, meshing parameters, materials, loads, boundary conditions and optimisation parameters, and then interfacing with Altair HyperMesh through automatically generated TCL scripts. Design parameter modifications are thus quickly applied to the GFEM model, avoiding repeated manual preprocessing.

The objective of the framework is to support early-stage design exploration, where multiple structural layouts must be generated and evaluated with limited computational and modelling time. The current implementation includes two main modules. The first module generates wingbox models and is directly used in the optimisation study presented in the following sections. The second module generates representative fuselage-segment models and is included to demonstrate the extensibility of the developed framework toward broader airframe structures.

2.1 Wingbox Generation

The wingbox generation module creates 2D GFEM models of transport-aircraft wingbox structures from a set of geometric and internal-layout parameters. The wing loft definition is based on planform quantities such as aspect ratio, surface area, sweep, dihedral, taper ratios, airfoil sections, and twist distribution. The internal structure includes ribs, front and rear spar webs and caps, and T-shaped stringers (Figure 2). These components are created as two-dimensional surfaces and subsequently meshed using shell elements.

A wide range of planform and structural-layout parameters is available, allowing the same module to generate different wing configurations without modifying the modelling script. Table 1 lists the parameters currently available for the wingbox definition. The external geometric parameters define the Outer Mold Line (OML) and planform shape, while the internal configuration parameters determine the structural layout of the stiffening components. For the subsequent optimisation, only a subset of these parameters is varied, namely the number of ribs, the number of stringers per skin panel, and the spar-cap width, although more parameters can be included in the future.

Figure 1 shows examples of wing planforms and wingbox geometries generated by the module. This is useful in early-stage structural design, where both global planform changes and internal structural-layout changes may be relevant.

After the geometry is created, the module generates a two-dimensional shell mesh in HyperMesh according to the defined mesh density. The meshing procedure is designed to preserve nodal connectivity across component interfaces, ensuring that adjacent shell components are consistently connected through shared nodes. A hybrid structured meshing strategy is adopted, consisting mainly of CQUAD4 elements, allowing CTRIA3 elements only in transition regions where a fully structured quadrilateral mesh is not practical.

Table 1 – Available parameters for the parametric creation of the wing model.

Wing Geometry Definition	**Internal Wing Configuration**
Aspect ratio (AR)	Rib spacing and positioning
Surface area (S)	Stringer cross-section dimensions
Sweep angle and dihedral	Stringer spacing/number
Kink span percentage (if applicable)	Spar-cap cross-section dimensions
Taper ratios (TR)	Spar locations (front/rear as % chord)
Airfoil sections and twist distribution	Stringer orientation and run-out

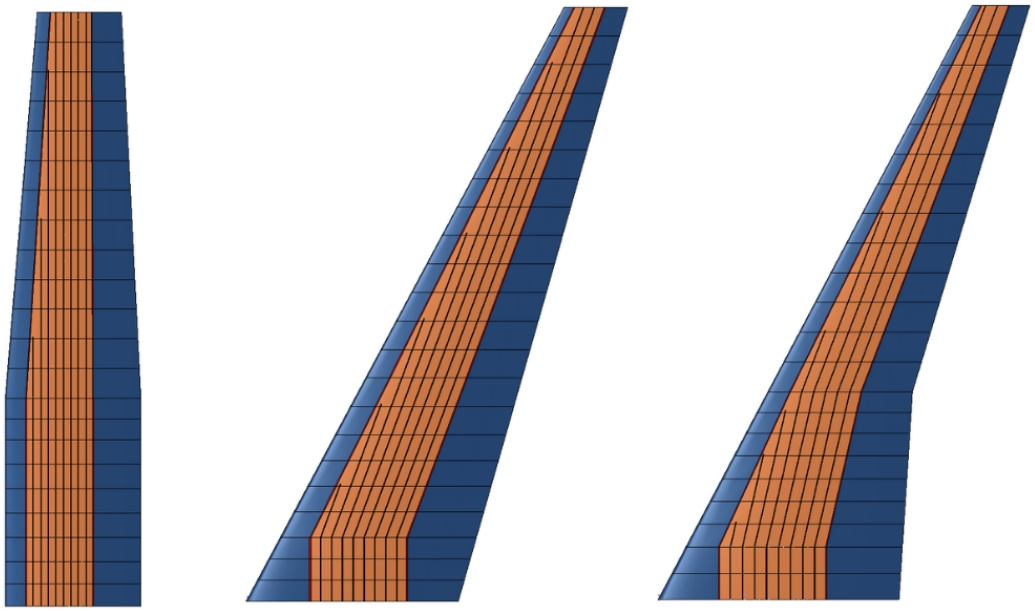


Figure 1 – Examples of wing planforms generated by the parametric tool.

An element-quality check is performed within the workflow. Regions containing low-quality elements, typically located near stringer run-outs, local transitions, or intersections between structural components, are identified and locally remeshed using HyperMesh cleanup and refinement capabilities.
Figure 2 shows a representative generated wingbox mesh. At the interfaces between skin panels and spar or stringer caps, superimposed shell elements with different properties and offsets are used. This allows overlapping structural components to be modelled while preserving the intended load path and avoiding the need for separate contact or tie definitions.

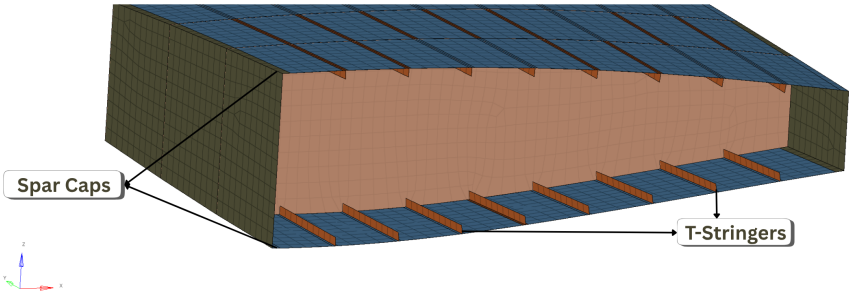


Figure 2 – Typical wingbox GFEM mesh generated by the parametric tool.

2.2 Fuselage Segment Generation Module

A second Python module is developed to generate representative fuselage-section models using the same general approach. The purpose of this module is to extend the parametric framework beyond isolated wingbox models providing the basis for future coupled wing–fuselage structural studies. The external fuselage geometry consists of a constant cylindrical section followed by an aft tapered section. The model represents a typical transport-aircraft fuselage section and includes structural components such as frames, stringers, passenger and cargo floors, floor stiffeners, stanchions, and reinforced cutouts (Figure 3).
Compared with the wingbox module, the fuselage model requires a larger number of geometric feature definitions because of the presence of cutouts, floor structures, local reinforcements, and section transitions. The outside fuselage geometry must be specified for both the constant and tapered sections. Cutout positions and dimensions, including those of passenger and cargo doors, emergency exits, and windows, are also defined parametrically. The model includes omega stringers, Z-frames, cutout reinforcements using intercostals and sills and floor-support structures.

Table 2 – Available parameters for the parametric creation of the fuselage model.

Fuselage Geometry and Cutouts	**Internal Structural Layout**
OML characteristics of fuselage sections	Frame spacing and cross-section
Passenger and cargo door positions and dimensions	Stringer spacing and cross-section
Emergency exit door positions and dimensions	Main and cargo floor cross-section
Window number and cut-out positioning	Main and cargo floor supports cross-section
Main and cargo floor vertical position	Intercoastal and sill dimensions for cut-out reinforcement

The passenger floor is modelled using longitudinal stiffeners and transverse beams, which are connected to the frames through vertical stanchions. The cargo floor is represented using longitudinal and transverse support beams, also connected to the surrounding frame structure. This level of detail is intended to provide a representative GFEM fuselage model suitable for preliminary structural studies, rather than a detailed local model including every connection or bracket. Figure 3 shows a representative generated fuselage-section.
Table 2 lists the parameters required for the fuselage model definition. Some of those parameters define the global geometry and cutout layout, while others control structural features such as frame spacing, stringer spacing, floor position, and reinforcement dimensions.

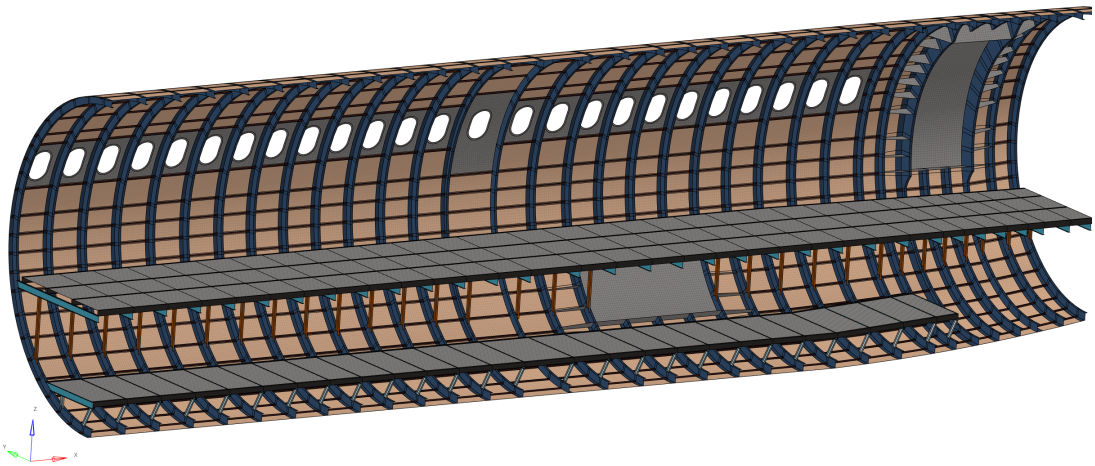


Figure 3 – Fuselage section model generated by the parametric tool.

3. Benchmark Wingbox Model and Nested Optimisation Problem

The remainder of the paper focuses on the wingbox module, which is coupled to the nested sizing and surrogate-assisted optimisation workflow to demonstrate the procedure using the STW benchmark [30].

3.1 Simple Transonic Wing Benchmark

The STW benchmark is based on the high-gross-weight variant of the Boeing 717 wing. The STW wingbox includes an unswept inboard region extending to the wing–fuselage intersection (wing root), located at a semispan of $y = 1.5$ m attached to an outboard swept section. Outboard of the root, the wingbox extends from 15% to 65% of the local wing chord. Table 3 lists the wing geometric parameters and aircraft mass breakdown used in the current model.

Table 3 – Aircraft and wing parameters of the STW benchmark [30].

Quantity	Description	Value
	Baseline wing geometry	
b	Semispan	14 m
C_{root}	Root chord	5 m
C_{tip}	Tip chord	1.5 m
S	Planform area (single wing)	45.5 m^2
MAC	Mean aerodynamic chord	3.56 m
	Aircraft masses	
M_{payload}	Payload mass	14.5×10^3 kg
M_{frame}	Operating empty mass, excluding wing	25×10^3 kg
$M_{\text{fuel,res}}$	Reserve fuel mass	2×10^3 kg

Two manoeuvre limit load cases are considered for sizing: a $+2.5g$ pull-up manoeuvre and a $-1g$ pull-down manoeuvre at Mach number $M = 0.458$ and sea-level conditions. Ultimate load cases are obtained by applying a safety factor of 1.5 to the limit load factors, resulting in $+3.75g$ and $-1.5g$ conditions. The required lift for each manoeuvre is calculated using the Landing Gross Mass (LGM), assuming an initial total wing mass of 6 tonnes. Aerodynamic loads are computed using the Athena Vortex Lattice (AVL) code for the aerodynamic surface, and applied as pressure distribution on the skin elements. Since the structural model only includes the wingbox, aerodynamic loads acting on the leading and training edges are introduced as concentrated forces on dedicated load-transfer nodes connected to the spar nodes through RBE3 connections (Figure 4). Inertial loads associated with each manoeuvre are also included. For each load case, both linear static and linear buckling

analyses are performed. The external wing geometry and aerodynamic pressure distributions are kept fixed throughout the optimisation, corresponding to a rigid-wing loading assumption.

3.2 Nested Layout–Sizing Formulation

The structural optimisation problem is formulated using a nested optimisation approach. The outer loop includes the structural layout variables, while the inner loop performs gradient-based component sizing for each fixed layout. The outer-loop optimisation can then be solved using a gradient-free method, such as BO, avoiding issues with integer variables.
Three outer-loop layout variables are considered:

$$\mathbf{x}_L = \left[N_{\text{ribs}}, N_{\text{str}}, w_{\text{cap}}\right]^T, \tag{1}$$

where N_{ribs} is the number of ribs, N_{str} is the number of stringers per skin panel, taken to be the same for the upper and lower skins, and w_{cap} is the spar-cap width. The selected variable bounds are:

$$22 \leq N_{\text{ribs}} \leq 30, \qquad 6 \leq N_{\text{str}} \leq 14, \qquad 20\ \text{mm} \leq w_{\text{cap}} \leq 50\ \text{mm}, \tag{2}$$

where N_{ribs} and N_{str} are integer-valued variables whereas w_{cap} is continuous.
For each candidate layout $\mathbf{x}_L$, the parametric framework automatically regenerates the wingbox GFEM and prepares the corresponding OptiStruct sizing optimisation. The inner-loop problem then computes the minimum feasible wingbox mass by optimising the composite ply-thickness variables and metallic rib thicknesses under strength, stability, and manufacturing constraints. The nested response evaluated by the outer loop can therefore be written as:

$$y(\mathbf{x}_L) = \min_{\mathbf{x}_s} M_{\text{wingbox}}(\mathbf{x}_L, \mathbf{x}_s), \tag{3}$$

subject to the sizing constraints described in Sec. 3.4, where $\mathbf{x}_s$ denotes the inner-loop sizing design-variable vector.

3.3 Generated Wingbox Structural Model

For each layout configuration, the STW wingbox structural model is automatically generated, including lower and upper skins, front and rear spar webs and caps, ribs, and stringers as distinct structural components. Ribs are modelled using Aluminium 7010-T651, while all the other components are modelled using unidirectional Carbon Fibre Reinforced Polymer (CFRP) laminates. Table 4 lists the model material properties. The `PSHELL` property card is used for the metallic ribs, while the ply-based composite modelling approach available in HyperMesh is adopted through the `PCOMPP` card for laminate components.

Table 4 – Mechanical properties of materials used for the STW model.

Aluminium 7010-T651

Property	Units	Value
Young's modulus	E [GPa]	71.7
Poisson's ratio	ν [–]	0.33
Density	ρ [kg/m^3]	2810
Yield strength	σ_y [MPa]	490

CFRP unidirectional ply [30]

Property	Units	Value
Young's modulus, fibre direction	E_{11} [GPa]	117.7
Young's modulus, transverse direction	E_{22} [GPa]	9.7
Shear modulus, 1–2	G_{12} [GPa]	4.8
Shear modulus, 1–3	G_{13} [GPa]	4.8
Poisson's ratio	ν_{12} [–]	0.35
Density	ρ [kg/m^3]	1550

Boundary conditions simulate a semi-wing connected to the fuselage. Symmetry conditions with respect to the x–z plane are imposed at the aircraft centreline ($u_y = \theta_x = \theta_z = 0$). At the wing root, the model is constrained along the chordwise, y, and vertical, z, directions, with $u_x = u_z = 0$, as shown in Figure 4a. Boundary conditions are applied to the external perimeter of the ribs at the corresponding locations. The aerodynamic OML is based on the RAE2822 airfoil with zero spanwise twist. Non-structural masses, including one tonne of fuel, are represented using concentrated mass elements distributed across the integrated tank volume up to 80% of the semispan and connected to the primary structure through RBE3 elements, (Figure 4b).

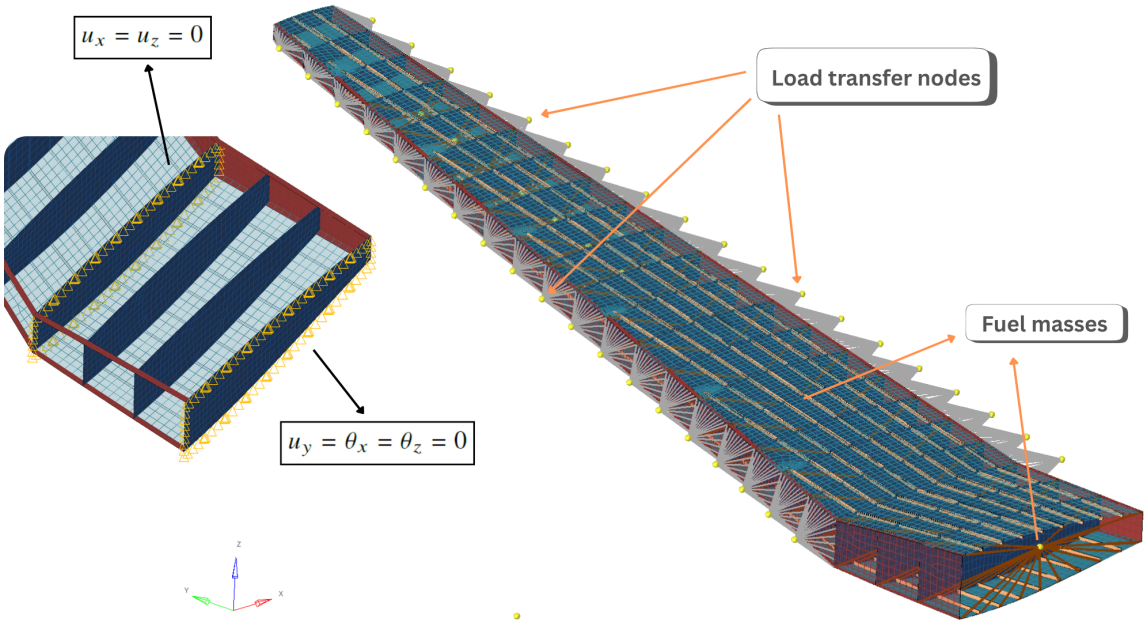


Figure 4 – STW wingbox structural model: a) boundary conditions, b) mesh, concentrated masses, and load-transfer nodes.

3.4 Composite Sizing Optimisation Setup

The wingbox is divided into nine structural components: ribs, upper and lower skins, upper and lower stringers, front and rear spar webs, and front and rear spar caps. In addition, the wing is divided into three spanwise sizing zones, shown in Figure 5. Each zone is characterised by uniform thickness per component. This provides a piecewise spanwise thickness distribution while keeping the sizing problem computationally manageable.

For each composite component and sizing zone, individual unidirectional ply shapes, with a manufacturing ply thickness of 0.17 mm, are defined for the orientations $\theta \in \{0°, \pm 45°, 90°\}$. The laminates are symmetric and balanced by construction. A 10% minimum ply-percentage rule is enforced for each orientation using the `DCOMP` card in OptiStruct [31]. Rib thicknesses are defined for each spanwise sizing zone.

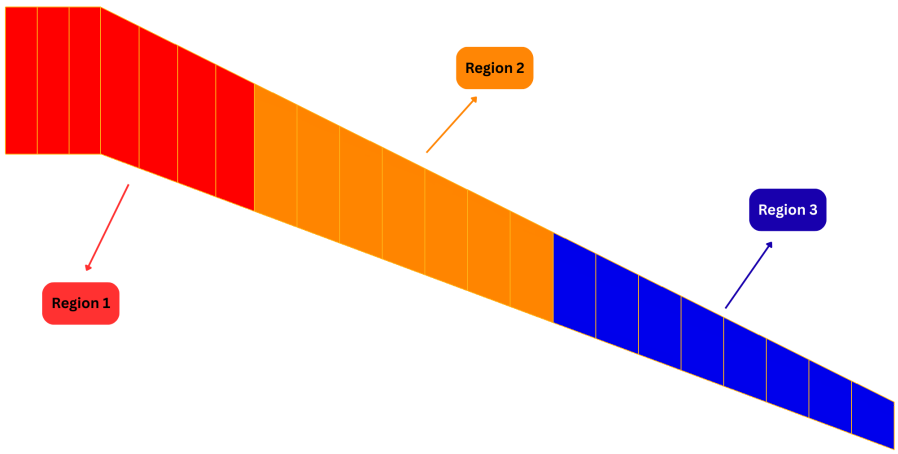


Figure 5 – Spanwise sizing zones used in the wingbox sizing optimisation.

3.4.1 Sizing design variables

Each laminate is modelled using smeared properties, therefore ignoring stacking-sequence effects, and parameterised by the total thickness assigned to each ply orientation (ply-bundle thickness). Since a manufacturing ply thickness of 0.17 mm is specified, the ply-bundle variables are treated

as discrete. OptiStruct employs a gradient-based approach, where a continuous relaxation of the thickness response is used, while enforcing manufacturable ply-bundle thicknesses corresponding to integer multiples of the specified ply thickness. This approach works better the smaller the discrete interval of the variables is [31]. The sizing and stacking-sequence optimisation steps are separated, following the recommended composite optimisation workflow in OptiStruct [31]. Stacking-sequence optimisation is outside the scope of the present work.
For each composite laminate, design variables are assigned to the total number of the 0°, $\pm 45^\circ$, and 90° ply orientations. The same design variable is used for the $+45^\circ$ and -45° ply groups to enforce laminate balance constraints. Since symmetric laminates are used, the variables refer to half of the laminate. In addition, one thickness design variable is assigned to the ribs in each spanwise sizing zone. The total number of sizing variables is therefore

$$n_{\text{ply},\theta}\, n_{\text{comp}}\, n_{\text{zones}} + n_{\text{th,rib}} = 75, \tag{4}$$

where $n_{\text{ply},\theta}$ is the number of ply-orientation variables per laminate, n_{comp} is the number of composite components, n_{zones} is the number of spanwise sizing zones, and $n_{\text{th,rib}}$ is the number of rib-thickness variables.

3.4.2 Structural integrity constraints

For composite components, an ultimate-load strain limit of $\pm 3500\ \mu\varepsilon$ is imposed on the laminate in-plane major and minor principal strains, $(\varepsilon_1, \varepsilon_2)$, evaluated for each element at each ply integration points. This is consistent with a damage-tolerance Barely Visible Damage (BVD) no-growth capability approach [32]. For the aluminium ribs, a modified non-yielding criterion is enforced at limit load using a Von Mises allowable stress of

$$\sigma_{\text{Al}}^{\text{allow}} = 0.7\,\sigma_y = 343\ \text{MPa}. \tag{5}$$

The reduced allowable is used to preliminarily account for fatigue and damage-tolerance considerations. The stress constraint is enforced on both top and bottom surfaces of each shell element.
Structural stability is enforced through linear buckling constraints. Buckling analyses are performed at limit load, and the first buckling eigenvalue is required to satisfy

$$\lambda_1 \geq \lambda_{\text{req}} = 1.5, \tag{6}$$

so that no buckling mode occurs before ultimate load. An allowable constraint violation of 1.5% is used during the OptiStruct sizing optimisation.

3.4.3 Inner-loop sizing problem

For a fixed layout $\mathbf{x}_L$, the objective of the sizing optimisation is to minimise the wingbox structural mass while satisfying the structural integrity and manufacturing constraints. The sizing problem is solved using OptiStruct, where a constraint screening method is used to retain only the active constraints, and the design sensitivities are computed for the retained responses using the direct method. A convergence tolerance of 0.2% is used for the objective. Let $\mathbf{x}_s \in \mathbb{R}^{75}$ denote the sizing design variable vector, including the composite ply-thickness variables and the metallic rib-thickness variables. The inner-loop sizing problem is written as

$$\begin{aligned}
\min_{\mathbf{x}_s} \quad & M_{\text{wingbox}}(\mathbf{x}_L, \mathbf{x}_s) \\
\text{subject to} \quad & |\varepsilon_1(\mathbf{x}_L, \mathbf{x}_s)|,\ |\varepsilon_2(\mathbf{x}_L, \mathbf{x}_s)| \leq 3500\ \mu\varepsilon, && \text{for all composite elements under ultimate loads,} \\
& \sigma_{\text{eq}}(\mathbf{x}_L, \mathbf{x}_s) \leq 343\ \text{MPa}, && \text{for all aluminium rib elements under limit loads,} \\
& \lambda_1(\mathbf{x}_L, \mathbf{x}_s) \geq 1.5, && \text{under limit loads,} \\
& 1 \leq n_{\text{ply},\theta} \leq 40, \quad n_{\text{ply},\theta} \in \mathbb{Z}, && \text{for all components, zones, and } \theta \in \{0^\circ, \pm 45^\circ, 90^\circ\}, \\
& 1\ \text{mm} \leq t_{\text{rib}} \leq 10\ \text{mm}, && \text{for all rib groups,} \\
& p_{j,\theta}(\mathbf{x}_L, \mathbf{x}_s) \geq 0.10, && \text{for all laminates } j \text{ and orientations } \theta.
\end{aligned} \tag{7}$$

where M_{wingbox} is the wingbox structural mass, σ_{eq} is the Von Mises equivalent stress in the aluminium ribs, λ_1 is the first buckling eigenvalue, and $p_{j,\theta}$ is the ply percentage of orientation θ in laminate j. The integer variables $n_{\text{ply},\theta}$ denote the equivalent number of plies associated with each orientation in the half of the symmetric laminate. In OptiStruct, these are implemented as discrete thickness variables based on the manufacturing ply thickness.

3.5 Multi-Run Sizing Strategy

Gradient-based sizing optimisation may converge to different local minima depending on the initial design, optimisation path, and active constraint set. As a result, small changes in layout parameters, or even repeated evaluations of the same layout, may return slightly different optimised masses. This introduces effective solver-induced variability in the objective values observed by the outer-loop optimiser. To reduce this effect, a staged multi-run strategy is applied for each fixed layout configuration, as shown in Figure 6.

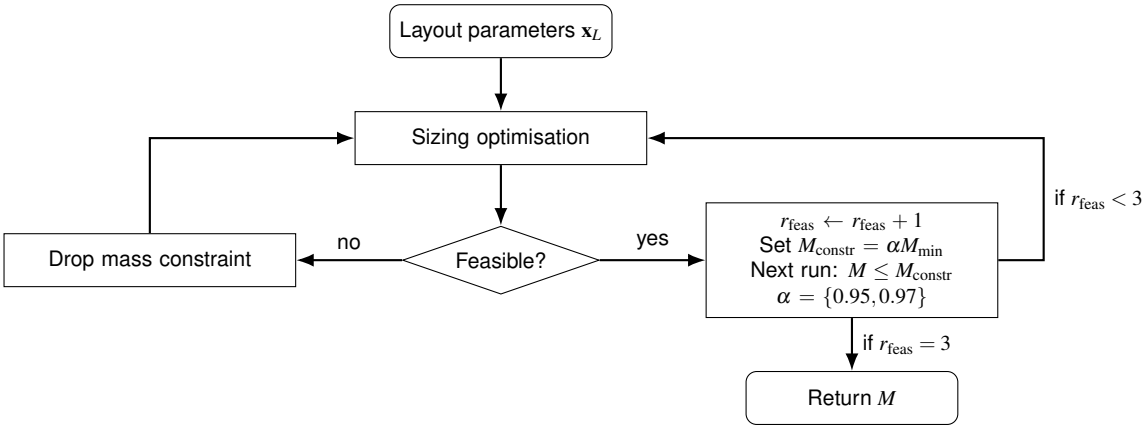


Figure 6 – Flowchart of the staged multi-run strategy used for the inner-loop sizing optimisation.

The first sizing optimisation run is performed without an additional mass constraint. If a feasible solution is obtained, an upper bound on the mass is imposed in the next run and the sizing optimisation is performed again starting from the previous feasible design. The mass bound is defined by multiplying the current minimum mass by a tightening factor α. A value of $\alpha = 0.95$ is used after the first feasible run to encourage a meaningful improvement, while $\alpha = 0.97$ is used in the following feasible run to reduce the likelihood of repeated infeasibility. If a run becomes infeasible, the mass constraint is removed in the next run so that the optimiser can recover feasibility.

The procedure is terminated after three feasible sizing runs. This limit was selected as a compromise between reducing local-minimum sensitivity and keeping each nested evaluation computationally affordable. Additional runs increase the cost approximately linearly, while the corresponding mass improvements tend to diminish. The multi-run procedure reduced the returned wingbox mass by as much as $300\ \text{kg}$ between the first and final feasible sizing runs for some layouts, while improving the consistency of the mass values used for surrogate training and outer-loop optimisation.

4. Surrogate-Assisted Layout Optimisation

The nested optimisation problem described above requires one complete composite sizing optimisation for every candidate layout. Since each evaluation involves automated GFEM generation, FE analysis, sensitivity calculation, and constrained sizing optimisation, exhaustive exploration of the layout design space is computationally expensive. A surrogate-assisted optimisation strategy is therefore adopted to approximate the mapping between the layout variables and the optimised wingbox mass, and guide the selection of new layouts under a limited evaluation budget.

4.1 Sampling Strategy

An initial dataset is generated using Latin Hypercube Sampling (LHS), which provides a structured sampling plan by distributing samples across the admissible range of each variable [20].To improve the space-filling quality of the initial dataset, the LHS plan is optimised using the Morris–Mitchell

maximin criterion [33]. This criterion seeks to maximise the minimum pairwise distance between sample points. An evolutionary algorithm is used to optimise the sampling plan, while preserving the LHS structure and ensuring that the integer variables retain admissible discrete values. For each sampled layout, the parametric framework automatically generates the wingbox GFEM and sets up the OptiStruct sizing optimisation. The optimised structural mass calculated by the sizing procedure is then used as the observed response for surrogate training.

4.2 Gaussian-Process Surrogate Model

The computationally expensive sizing optimisation is approximated using a GPR model, also referred to as Kriging. GPR is widely used in surrogate modelling as it provides a flexible non-parametric approximation together with predictive uncertainty quantification [34, 17]. Let $f(\mathbf{x}_L)$ represent the unknown latent function mapping a layout configuration to the optimised wingbox mass. The observed response y, obtained from the inner-loop sizing optimisation is modelled as

$$y = f(\mathbf{x}_L) + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_n^2), \tag{8}$$

where ε is an additive observation-noise term. In the current problem, this term is used to account not only for numerical noise, but also for the epistemic variability introduced by local convergence effects in gradient-based sizing optimisation.

The GP surrogate provides a predictive distribution for any candidate layout, characterised by a posterior mean $\mu(\mathbf{x}_L)$ and standard deviation $s(\mathbf{x}_L)$. The posterior mean is used as the surrogate prediction of the optimised wingbox mass, while the posterior standard deviation quantifies the prediction uncertainty. This uncertainty estimate is later used by the BO procedure to guide the selection of new expensive evaluations.

A Matérn-$3/2$ covariance function with Automatic Relevance Determination (ARD) is implemented. The Matérn family is commonly used in engineering surrogate modelling because it allows the assumed smoothness of the response to be controlled [34, 17]. The Matérn-$3/2$ kernel provides a compromise between local flexibility and smoothness, while ARD assigns a separate lengthscale to each layout variable. This allows the surrogate to represent different sensitivities of the optimised mass with respect to rib number, stringer number, and spar-cap width. Before training, input variables are normalised to the unit interval using their lower and upper bounds.

The GP hyperparameter vector, θ_{GP}, is defined by:

$$\theta_{\mathrm{GP}} = \left[\sigma_f^2, \ell_{N_{\mathrm{ribs}}}, \ell_{N_{\mathrm{str}}}, \ell_{w_{\mathrm{cap}}}, \sigma_n^2\right]^T, \tag{9}$$

where σ_f^2 is the signal variance, $\ell_{N_{\mathrm{ribs}}}$, $\ell_{N_{\mathrm{str}}}$, and $\ell_{w_{\mathrm{cap}}}$ are the ARD lengthscales associated with the three layout variables, and σ_n^2 is the observation-noise variance. Given a training dataset of size N, the hyperparameters are selected by maximising the logarithmic marginal likelihood, or equivalently by minimising the Negative Log Marginal Likelihood (NLML), of the GP model. The NLML is given by [34]

$$NLML(\theta_{\mathrm{GP}}) = -\log p\left(\mathbf{y} \mid \mathbf{X}_L, \theta_{\mathrm{GP}}\right), \tag{10}$$

where $\mathbf{X}_L$ contains the evaluated layout configurations, $\mathbf{y}$ contains the corresponding optimised structural masses, and $p(\mathbf{y} \mid \mathbf{X}_L, \theta_{\mathrm{GP}})$ is the marginal likelihood of the observations under the GP model.

During hyperparameter optimisation, a lower bound of $\sigma_n \geq 10\,\mathrm{kg}$ is imposed on the observation-noise standard deviation. This value was selected based on repeated inner-loop sizing runs for identical or very similar layout configurations, showing a minimum mass variation on this order caused by the varying convergence of the gradient-based sizing optimiser. The lower bound is small compared to the typical optimised wingbox mass, but large enough to prevent the surrogate from treating solver-induced variability as a deterministic local response variation. It therefore acts as a conservative regularisation choice, allowing the GP to capture the dominant layout-to-mass trend while accounting for the effective variability of the nested sizing response. It is worth noticing that the optimised noise hyperparameter converged to this lower bound, so the fitted surrogate effectively uses a noise standard deviation of $10\,\mathrm{kg}$.

Predictive accuracy is assessed using the Leave-One-Out (LOO) cross-validation approach. For GP models, LOO predictions can be efficiently computed from the inverse of the training covariance matrix following the standard GP formulation [34]. LOO is suitable for the investigated small-data setting because each training point is successively predicted using a model conditioned on the remaining points. The resulting LOO predictions are evaluated using the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE) and coefficient of determination R^2.

4.3 Bayesian Optimisation

The trained GP surrogate is embedded into a BO loop to select new layout configurations in a data-efficient manner. BO is a sequential strategy for expensive black-box optimisation where a probabilistic surrogate model is used to decide where the expensive function should be subsequently evaluated [22, 35, 23]. At each iteration, the current surrogate posterior is used to define an acquisition function $\alpha(\mathbf{x}_L)$, which assigns a utility value to each candidate layout. The next layout can be generically written as

$$\mathbf{x}_{L,n+1} = \arg\max_{\mathbf{x}_L \in \mathscr{X}_L} \alpha_n(\mathbf{x}_L), \tag{11}$$

where $\mathscr{X}_L$ indicates the admissible mixed-variable layout design space.

Since the nested sizing optimisation response is affected by solver-induced variability, a noise-aware improvement-based acquisition function is adopted. In particular, the Logarithmic Noisy Expected Improvement acquisition function, LogNEI, is used. Noisy expected improvement accounts for observation noise by defining improvement with respect to the posterior distribution of the latent objective rather than only the best observed noisy value [29]. The logarithmic formulation improves the numerical conditioning of improvement-based acquisition functions during acquisition optimisation [36]. The BO loop is implemented in BoTorch [37], using GPyTorch for GP model training and posterior prediction [38]. More specifically, the qLogNEI acquisition function available in BoTorch, which uses Monte-Carlo sampling of the GP posterior to select multiple candidate layouts in batch, allowing the corresponding expensive sizing optimisations to be evaluated in parallel, is implemented.

The BO process is initialised from the LHS dataset. New candidate layouts proposed by the acquisition function are evaluated using the same automated GFEM generation and OptiStruct sizing workflow as the initial samples. Once the new optimised mass values are obtained, they are added to the dataset and the GP surrogate is trained again. Algorithm 1 summarises the overall procedure. A fixed budget of 30 expensive layout evaluations is used, consisting of 24 initial LHS evaluations followed by 6 BO infill evaluations. This split was selected to provide broad initial coverage of the layout design space while still allowing adaptive refinement through BO. The six infill evaluations are selected using qLogNEI and evaluated through the parallel automated workflow.

Algorithm 1 Bayesian Optimisation loop

1: Fit GP surrogate on initial dataset $\mathscr{D}_0 = \{(\mathbf{x}, y)\}^N$
2: **for** $n = 0, 1, 2, \ldots, N_{\text{infill}} - 1$ **do**
3: Select new candidate(s) by maximising the acquisition function
4: $\mathbf{x}_{n+1} \leftarrow \arg\max_{\mathbf{x} \in \mathscr{X}} \alpha(\mathbf{x}; \mathscr{D}_n)$
5: Evaluate the expensive nested sizing model to obtain y_{n+1}
6: Update the dataset: $\mathscr{D}_{n+1} = \mathscr{D}_n \cup \{(\mathbf{x}_{n+1}, y_{n+1})\}$
7: Retrain the GP surrogate on $\mathscr{D}_{n+1}$
8: **end for**

5. Results and Discussion

This section presents the application of the suggested framework to the nested structural layout–sizing optimisation of the STW wingbox. The surrogate-based BO results are discussed and compared against a global-search baseline algorithm.

5.1 Initial Surrogate Accuracy and Variable Trends

The initial GP surrogate is trained using the 24 layouts generated by the optimised LHS plan. Table 5 summarises the LOO error metrics and fitted hyperparameters of the initial 24-point Matérn-$3/2$ GP surrogate.

Table 5 – LOO error metrics and fitted hyperparameters for the initial 24-point Matérn-$3/2$ GP surrogate.

Dataset	Outputscale	$\ell_{N_{\text{ribs}}}$	$\ell_{N_{\text{str}}}$	$\ell_{w_{\text{cap}}}$	RMSE [kg]	MAE [kg]	R^2
Initial 24-point dataset	1.2776	0.4495	0.9798	0.3181	31.02	24.77	0.439

The surrogate gives an RMSE of 31.02 kg and an MAE of 24.77 kg. These errors correspond to approximately 1–2% of the typical optimised wingbox mass. Therefore the model is considered to have a reasonably good accuracy, given the limited number of evaluations. The moderate R^2 value reflects the relatively small variation of the optimised mass over the investigated layout space and the effective epistemic variability introduced by the inner-loop sizing optimisation.

The fitted ARD lengthscales provide additional information about the response represented by the surrogate. The outputscale controls the overall amplitude of the latent variation represented by the GP, while the ARD lengthscales describe how rapidly the fitted response changes along each input direction. Since the inputs are normalised to the unit interval, the lengthscales can be compared directly across the three layout variables. The shortest lengthscale is obtained for the spar-cap width, $\ell_{w_{\text{cap}}} = 0.3181$, indicating the strongest local variation of the predicted mass with respect to this variable. The rib-number lengthscale is also relatively short, $\ell_{N_{\text{ribs}}} = 0.4495$, although the corresponding response is less influencing. The stringer-number lengthscale is the largest, $\ell_{N_{\text{str}}} = 0.9798$, indicating a smoother dependence over the investigated range.

Figure 7 shows the LOO parity plot for the initial 24-point surrogate. The model predictions match the general trend, however the scatter around the ideal $y = x$ line indicates that the initial model does not fully capture all the response local variations. In particular, some high-mass designs are under-predicted and some low-mass designs are over-predicted, suggesting a smoothing tendency.

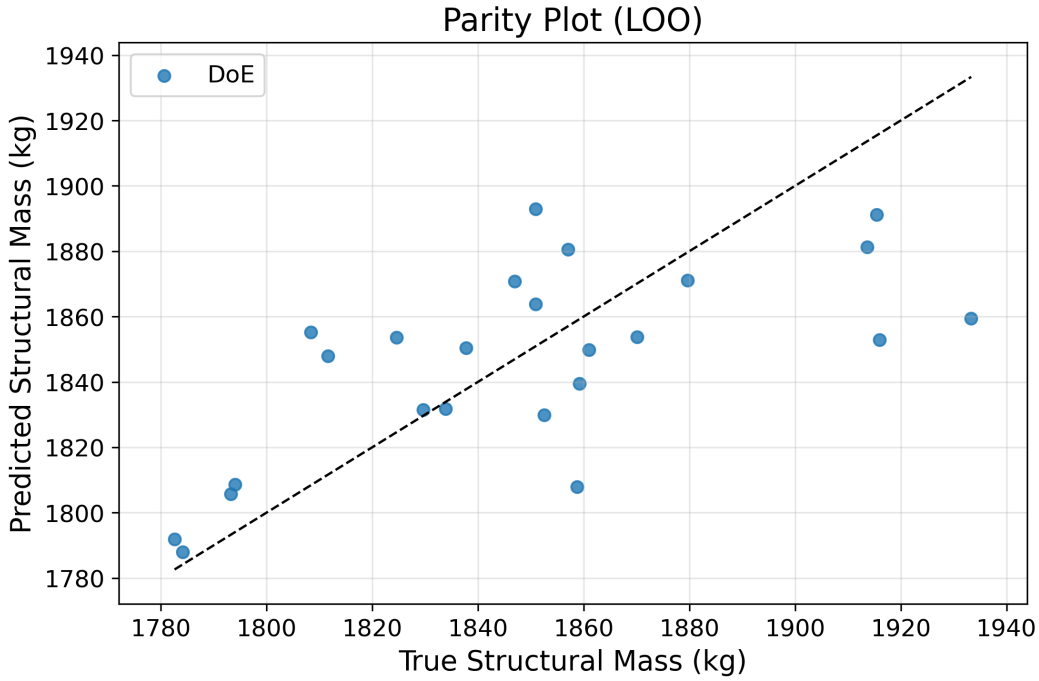


Figure 7 – Leave-one-out parity plot of the initial GP surrogate trained on the 24-point optimised LHS dataset.

To interpret the fitted surrogate response, one-dimensional partial dependence plots are generated for the three layout variables. Partial dependence plots show the average variation of the surrogate-

predicted mass with respect to one selected variable while averaging over the sampled values of the remaining variables [39]. Figure 8 shows the plots correspond to the initial 24-point surrogate, as this model provides broad coverage of the design space before BO infill points become concentrated near promising regions.

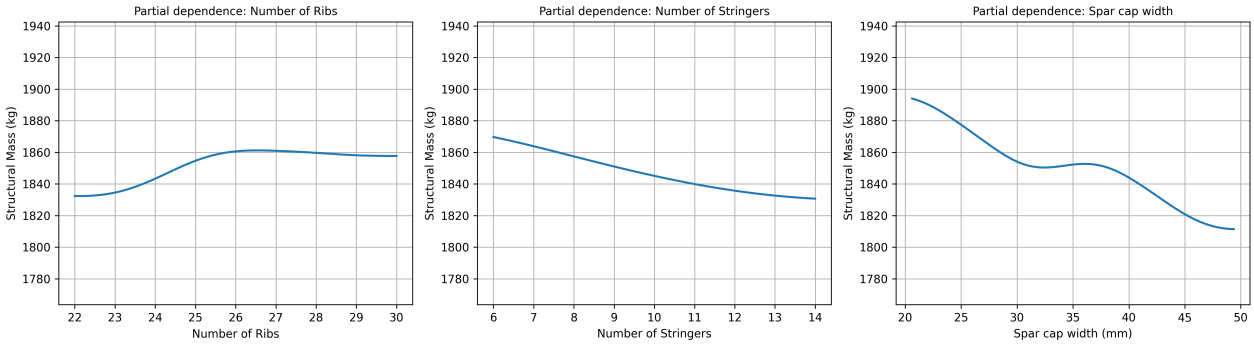


Figure 8 – One-dimensional partial dependence plots of the initial 24-point GP surrogate for the three layout variables.

The trends in Figure 8 indicate that spar-cap width has the strongest influence on the optimised wingbox mass. Increasing w_{cap} generally reduces the predicted structural mass, although the model response is not perfectly monotonic and shows a mild plateau at intermediate values within the variable range. This trend is physically consistent, as wider spar caps increase the wingbox bending stiffness and consequent load carrying capability reducing the required thickness of other components. The number of stringers also shows a mostly monotonic beneficial trend, with larger N_{str} reducing the predicted mass. This is consistent with improved skin-panel support and buckling resistance. In contrast, the effect of rib number is weaker and non-monotonic: the predicted mass increases from 22 to 26 ribs and then remains approximately flat or slightly decreases. This suggests that, within the investigated bounds, adding ribs does not provide a mass benefit comparable to increasing the spar-cap width or the number of stringers. It is worth noticing that these trends reflect the surrogate prediction, not necessarily the true underlying physical response.

5.2 Bayesian Optimisation Results

Starting from the an initial dataset of 24 samples, 6 additional layouts were selected using the qLogNEI acquisition function for BO. These infill points were evaluated using the same automated GFEM generation and inner-loop OptiStruct sizing optimisation as the initial samples. After adding the infill points, the GP surrogate was trained again on the complete 30-point dataset.

Table 6 compares the LOO error metrics of the initial and updated GP surrogates. The RMSE decreases from $31.02\ \text{kg}$ to $28.48\ \text{kg}$, the MAE decreases from $24.77\ \text{kg}$ to $21.07\ \text{kg}$, and R^2 increases from 0.439 to 0.630 as a consequence of the infill points. Although the dataset remains small, this improvement indicates that the BO evaluations effectively refine the surrogate in the relevant low-mass region.

Table 6 – LOO error metrics of the GP surrogate before and after BO infill.

Dataset	RMSE [kg]	MAE [kg]	R^2
Initial 24-point dataset	31.02	24.77	0.439
Final 30-point dataset	28.48	21.07	0.630

Figure 9 shows the LOO parity plot for the updated 30-point surrogate. The infill points, shown by different markers from the initial LHS points, are concentrated in the lower-mass region of the design space, as expected from an improvement-based BO procedure.

The BO study identifies the best layout $(N_{\text{ribs}}, N_{\text{str}}, w_{\text{cap}}) = (23, 13, 50\ \text{mm})$, with an optimised structural mass of $1737.35\ \text{kg}$. The best configuration lies at the upper bound of the spar-cap width and uses a high number of stringers, consistently with the trends inferred from the partial dependence plots.

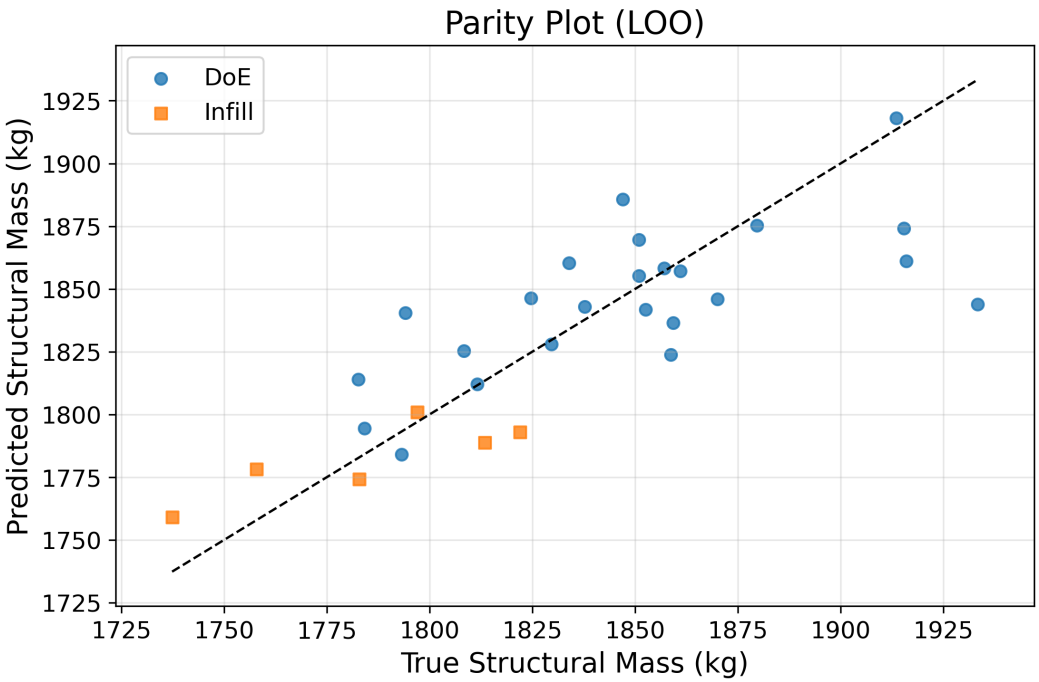


Figure 9 – Leave-one-out parity plot of the updated GP surrogate after adding the 6 BO infill evaluations.

Compared with the best design in the initial LHS dataset, which had a structural mass of $1782.61\ \mathrm{kg}$, the BO infill stage reduces the optimised wingbox mass by $45.26\ \mathrm{kg}$, corresponding to approximately a 2.54% reduction.
Overall, the BO infill stage not only improves the best observed design, but also confirms the main structural trend suggested by the surrogate, i.e. efficient layouts in the investigated design space favour large spar-cap widths and a relatively high number of stringers, while the effect of rib number is comparatively weaker.

5.3 Result comparison with MIDACO baseline

To assess the quality of the BO solution, the optimisation problem is also solved using MIDACO (Mixed Integer Distributed Ant Colony Optimisation, [40, 41]. MIDACO is a gradient-free evolutionary hybrid optimiser for continuous, discrete, and mixed-integer constrained optimisation problems. MIDACO treats the design responses as black-box and is designed for nonconvex, discontinuous, or otherwise challenging optimisation problems, while offering parallelisation capabilities [40, 41]. In this validation campaign, MIDACO optimises only the outer-loop layout variables, while each proposed layout is evaluated using the same OptiStruct sizing optimisation used in the BO study. This makes the comparison direct at the layout-optimisation level, since both methods use the same inner-loop evaluator.
Table 7 lists the best layouts and computational cost of the BO and nested MIDACO studies. The nested MIDACO study evaluated 152 layout configurations and identified the best layout $(N_{\mathrm{ribs}}, N_{\mathrm{str}}, w_{\mathrm{cap}}) = (23, 13, 48.61\ \mathrm{mm})$, with a structural mass of $1732.24\ \mathrm{kg}$.

Table 7 – Comparison between the BO result and the nested MIDACO reference.

Method	N_{ribs}	N_{str}	w_{cap} [mm]	Mass [kg]	Evaluations	Solver CPU [h]	Wall time [h]
BO, 24+6	23	13	50.00	1737.35	30	518.52	56.59
Nested MIDACO	23	13	48.61	1732.24	152	3176.08	350.40

The MIDACO resulting design has a mass that is slightly lighter than the BO result, showing a difference of $5.11\ \mathrm{kg}$, corresponding to approximately 0.29% of the BO mass. This difference is small compared to the observed variability of the inner-loop sizing procedure and should not be interpreted

as evidence of a clearly superior layout. A more important observation is that both methods recover the same discrete layout region, $(N_{\text{ribs}}, N_{\text{str}}) = (23, 13)$, with spar-cap widths close to the upper bound. The BO framework reaches this configuration using 30 expensive layout evaluations compared to the 152 required by MIDACO. That corresponds to an 80.3% reduction in the number of expensive nested evaluations. In terms of wall-clock time, the BO study required 56.59 h, compared with 350.40 h for MIDACO. All studies were performed on the same workstation, equipped with 64 CPU cores and 300 GB of RAM. This comparison demonstrates that the surrogate-assisted framework identifies the same optimum layout region saving significant computational cost.

5.4 Discussion

The optimised layouts indicate that the most influential design parameters in the investigated space are the spar-cap width and the number of stringers, while the effect of the number of ribs appears weaker and less systematic. Although both stringers and ribs improve support against local panel buckling, the results suggest that increasing the number of stringers is a more weight-efficient strategy. This is partially ascribed to the larger mass penalty associated with the metallic ribs. Configurations with a greater number of stringers generally featured reduced skin and stringer thicknesses, indicating that the additional longitudinal support improved panel stability and allowed stiffened panels to carry loads more efficiently. Likewise, wider spar caps were associated with lower overall mass and reduced thickness demand in the main bending-carrying regions, suggesting that they improve bending stiffness. By contrast, increasing the number of ribs did not produce a comparable mass benefit through thickness reductions elsewhere in the structure.

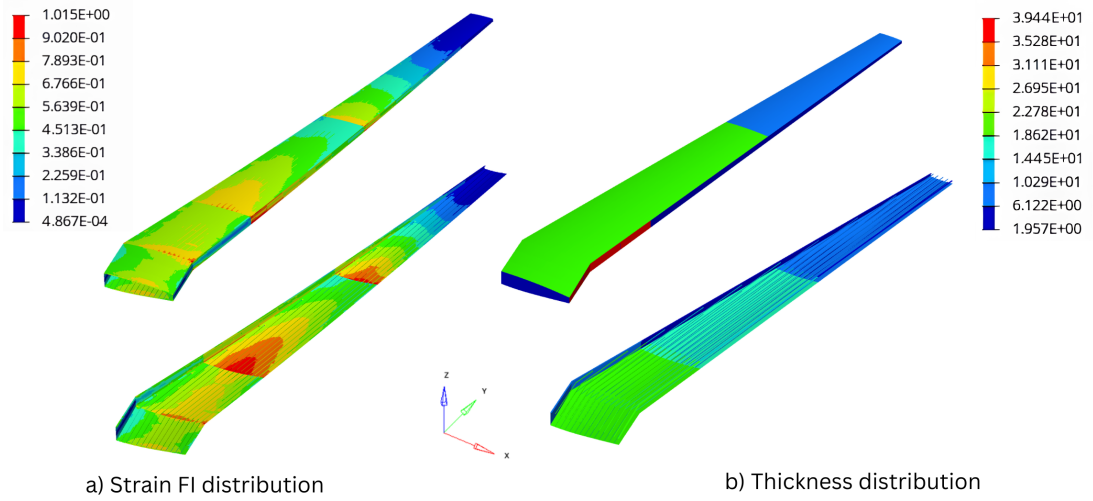


Figure 10 – Optimised thickness and strain-based FI distributions for the best design

Figure 10a) shows the distribution of the principal strain-based Failure Indices (FI) of the critical ply in each laminate for the best design found by BO. Localised peaks are observed near the kink region, as well as in the boundaries between adjacent optimisation regions, where abrupt thickness changes introduce local stiffness discontinuities resulting in strain concentrations. Values exceeding $FI = 1$ can occur because the optimisation was performed with a constraint tolerance of 1.5%. Figure 10b) depicts corresponding thickness distribution of the composite laminates for the same design. Higher thicknesses are observed near the wing root, where higher bending moments are experienced. The first optimisation region also includes the kink, where local stress concentrations increase the required material, specifically on the more loaded rear spar. Finally, the ribs are sized to relatively low values (4.29-1.96 mm) compared to composite components, consistent with the secondary role of rib number observed in the layout-level trends.

Several approach limitations should be pointed out. First, the layout design space is restricted to three variables, and only three spanwise sizing zones are considered. Extending the framework to larger layout spaces or finer sizing parameterisations may provide additional mass reduction, but would require more scalable surrogate strategies. Second, aerodynamic loads are kept fixed during

the optimisation, so aeroelastic load redistribution is not included. Third, the laminate representation is based on smeared ply-thickness variables and does not include stacking-sequence optimisation. Finally, although the multi-run sizing procedure reduces the tendency of finding poor local minima, the inner-loop response is not perfectly deterministic, limiting the validity of conclusions which can be drawn from small mass differences between individual optimised designs.

6. Conclusions

This paper presented an automated parametric framework for composite airframe structural modelling and surrogate-assisted layout optimisation. The framework uses Python modules coupled with HyperMesh TCL scripting to generate optimisation-ready GFEMs with limited manual intervention. The model-generation capability was demonstrated for both wingbox and fuselage structures.

The optimisation workflow was applied to the STW benchmark using a nested layout–sizing formulation. In the outer loop, three layout variables were considered: number of ribs, number of stringers per skin panel, and spar-cap width. For each fixed layout, the inner loop performed composite sizing optimisation in OptiStruct using 75 sizing variables, including ply-thickness variables for the composite components and thickness variables for the metallic ribs. Strength, buckling, and manufacturing constraints were imposed, and a staged multi-run sizing strategy was implemented to reduce local-minimum sensitivity and improve the consistency of the mass values returned to the outer loop.

A GP surrogate was trained on 24 optimised layouts selected by an optimised LHS plan and was refined using 6 BO infill evaluations. The updated surrogate improved the leave-one-out metrics, reducing the RMSE from $31.02\ \mathrm{kg}$ to $28.48\ \mathrm{kg}$ and increasing R^2 from 0.439 to 0.630. Partial dependence plots indicated that the spar-cap width and stringer number have the strongest influence on the optimised wingbox mass within the investigated design space, while the effect of rib number is weaker and non-monotonic.

The best layout identified from BO is $(N_{\mathrm{ribs}}, N_{\mathrm{str}}, w_{\mathrm{cap}}) = (23, 13, 50\ \mathrm{mm})$, with an optimised structural mass of $1737.35\ \mathrm{kg}$. This corresponds to a reduction of $45.26\ \mathrm{kg}$, or approximately 2.54%, relative to the best design in the initial LHS dataset. The MIDACO global-search algorithm, used as validation reference, identified a very similar layout, $(23, 13, 48.61\ \mathrm{mm})$, with a mass of $1732.24\ \mathrm{kg}$. The small mass difference between the two solutions is within the variability expected from the inner-loop sizing procedure, while the agreement in the recovered layout provides confidence in the surrogate-assisted optimisation result. The BO workflow reached this layout using 30 expensive evaluations, compared with the 152 required by the nested MIDACO optimisation, corresponding to an 80.3% reduction.

Results demonstrate that the suggested framework can support early-stage composite airframe design by combining automated parametric GFEM generation, physics-based sizing optimisation, and data-efficient surrogate-assisted layout exploration. Future work will focus on extending the framework toward coupled wing–fuselage modelling, including the wing–fuselage intersection, and on expanding the design space to include additional layout and sizing variables. Further developments will also investigate surrogate formulations for direct constraint prediction and higher-dimensional optimisation, reducing reliance on repeated inner-loop sizing optimisations.

Acknowledgments

This research was carried out within the Spanish national project *TIFON (Tecnologías Inteligentes para la Fabricación, el diseño y las Operaciones en entornos iNdustriales)*, coordinated by *Universidad Politécnica de Madrid (UPM)* funded by the Spanish *Agencia Estatal de Investigación (AEI)* under the “Programa Transmisiones”, grant No. PLEC2023-010251.

7. Contact Author Email Address

Corresponding author email: nziakos@pa.uc3m.es

8. Copyright Statement

The authors confirm that they, and/or their company or organisation, 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] Daniel P. Raymer. *Aircraft Design: A Conceptual Approach*. AIAA Education Series. American Institute of Aeronautics and Astronautics, Reston, VA, USA, 5 edition, 2012.

[2] Odeh Dababneh and Timoleon Kipouros. A review of aircraft wing mass estimation methods. *Aerospace Science and Technology*, 72:256–266, 2018.

[3] Andrew S. Hahn. Vehicle sketch pad: A parametric geometry modeler for conceptual aircraft design. In *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, Orlando, FL, USA, 2010. AIAA.

[4] William J. Fredericks, Kevin R. Antcliff, Guillermo Costa, Nachiket Deshpande, Mark D. Moore, Edric A. San Miguel, and Alison N. Snyder. Aircraft conceptual design using vehicle sketch pad. In *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, Orlando, FL, USA, 2010. AIAA.

[5] John Joe, Viraj Gandhi, John Dannenhoffer, and Hamid Dalir. Rapid generation of parametric aircraft structural models. In *AIAA Scitech 2019 Forum*. American Institute of Aeronautics and Astronautics, January 2019.

[6] Christophe David, Scott Delbecq, Sébastien Defoort, Peter Schmollgruber, Emmanuel Bénard, and Valérie Pommier-Budinger. From FAST to FAST-OAD: An open source framework for rapid overall aircraft design. In *10th EASN International Conference on Innovation in Aviation & Space*, volume 1024 of *IOP Conference Series: Materials Science and Engineering*, page 012062. IOP Publishing, 2021.

[7] Wojciech Grendysa, Joël Jézégou, and Emmanuel Benard. Parametric model of aircraft external geometry as a component of fast-oad aircraft preliminary design system. *Journal of Physics: Conference Series*, 2526(1):012030, June 2023.

[8] Mark Sensmeier and Jamshid Samareh. Automatic aircraft structural topology generation for multidisciplinary optimization and weight estimation. In *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*. American Institute of Aeronautics and Astronautics, April 2005.

[9] Jiapeng Tang, Ping Xi, Baoyuan Zhang, and Bifu Hu. A finite element parametric modeling technique of aircraft wing structures. *Chinese Journal of Aeronautics*, 26(5):1202–1210, October 2013.

[10] Felix Dorbath, Björn Nagel, and Volker Gollnick. A knowledge based approach for automated modelling of extended wing structures in preliminary aircraft design. In *Proceedings of the Deutscher Luft- und Raumfahrtkongress (DLRK) 2011*, Bremen, Germany, 2011.

[11] Brian H. Mason and Jesse Quinlan. Conceptual design shop: A tool for rapid airframe structural modeling. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. American Institute of Aeronautics and Astronautics, January 2018.

[12] John Hwang, Gaetan Kenway, and Joaquim R. R. A. Martins. Geometry and structural modeling for high-fidelity aircraft conceptual design optimization. In *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. American Institute of Aeronautics and Astronautics, June 2014.

[13] D. B. Schwinn, D. Kohlgrüber, J. Scherer, and M. H. Siemann. A parametric aircraft fuselage model for preliminary sizing and crashworthiness applications. *CEAS Aeronautical Journal*, 7(3):357–372, May 2016.

[14] Lars U. Hansen and Peter Horst. Multilevel optimization in aircraft structural design evaluation. *Computers & Structures*, 86:104–118, 2008.

[15] Bret K. Stanford, Christine V. Jutte, and Christian A. Coker. Aeroelastic sizing and layout design of a wingbox through nested optimization. *AIAA Journal*, 57(2):848–857, 2019.

[16] André Ferreira Costa Vieira, Marcos Rogério Tavares Filho, João Paulo Eguea, and Marcelo Leite Ribeiro. Optimization of structures and composite materials: A brief review. *Eng*, 5(4):3192–3211, 2024.

[17] Robert B. Gramacy. *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. CRC Press, Boca Raton, FL, 2020.

[18] Ricardo M. Paiva, André R. D. Carvalho, Curran Crawford, and Afzal Suleman. Comparison of surrogate models in a multidisciplinary optimization framework for wing design. *AIAA Journal*, 48(5):995–1006, 2012.

[19] Timothy W. Simpson, Timothy M. Mauery, John J. Korte, and Farrokh Mistree. Kriging models for global approximation in simulation-based multidisciplinary design optimization. *AIAA Journal*, 39(12):2233–2241, 2001.

[20] Alexander I. J. Forrester, Andras Sobester, and Andy J. Keane. *Engineering Design via Surrogate Modelling: A Practical Guide*. John Wiley & Sons, Chichester, UK, 2008.

[21] Nestor V. Queipo, Raphael T. Haftka, Wei Shyy, Tushar Goel, Rajkumar Vaidyanathan, and P. Kevin Tucker. Surrogate-based analysis and optimization. *Progress in Aerospace Sciences*, 41(1):1–28, 2005.

[22] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.

[23] Rémi Lam, Matthias Poloczek, Peter I. Frazier, and Karen E. Willcox. Advances in bayesian optimization with applications in aerospace engineering. In *20th AIAA Non-Deterministic Approaches Conference (AIAA SciTech Forum)*, 2018.

[24] Remy Priem, Hugo Gagnon, Ian Chittick, Stéphane Dufresne, Youssef Diouane, and Nathalie Bartoli. An efficient application of bayesian optimization to an industrial mdo framework for aircraft design. In *AIAA AVIATION 2020 Forum*, 2020.

[25] Bret K. Stanford and Satadru Roy. Sizing and topology design of an aeroelastic wingbox under uncertainty. In *International Forum on Aeroelasticity and Structural Dynamics (IFASD 2019)*, Savannah, GA, United States, June 2019. Paper IFASD-2019-058. NASA NTRS Document ID: 20200002634. Report No.: NF1676L-31616.

[26] Satadru Roy, Kenneth T. Moore, John T. Hwang, Justin S. Gray, William A. Crossley, and Joaquim R. R. A. Martins. A mixed integer efficient global optimization algorithm for the simultaneous aircraft allocation-mission-design problem. In *58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Grapevine, Texas, USA, January 2017. American Institute of Aeronautics and Astronautics (AIAA). AIAA SciTech Forum. AIAA Paper 2017-1305.

[27] Satadru Roy, William A. Crossley, Bret K. Stanford, Kenneth T. Moore, and Justin S. Gray. A mixed integer efficient global optimization algorithm with multiple infill strategy – applied to a wing topology optimization problem. In *AIAA Scitech 2019 Forum*. American Institute of Aeronautics and Astronautics (AIAA), 2019. AIAA Paper 2019-2356.

[28] Raphael T. Haftka, Diane Villanueva, and Anirban Chaudhuri. Parallel surrogate-assisted global optimization with expensive functions – a survey. *Structural and Multidisciplinary Optimization*, 54(1):3–13, 2016.

[29] Victor Picheny, Tobias Wagner, and David Ginsbourger. A benchmark of kriging-based infill criteria for noisy optimization. *Structural and Multidisciplinary Optimization*, 48(3):607–626, 2013.

[30] Alasdair C. Gray and Joaquim R. Martins. A proposed benchmark model for practical aeroelastic optimization of aircraft wings. In *AIAA SCITECH 2024 Forum*. American Institute of Aeronautics and Astronautics, January 2024.

[31] Altair Engineering, Inc. *Altair OptiStruct User Guide*. Altair Engineering, Inc., 2025. Altair HyperWorks / OptiStruct 2025.1 online documentation.

[32] Andrea Cini, A. Popov, and S. Ratchev. Minimum weight structural optimisation procedure for the preliminary sizing of a composite wing of a compound helicopter. In *6th Aircraft Structural Design Conference*, Bristol, UK, October 2018.

[33] Max D. Morris and Toby J. Mitchell. Exploratory designs for computational experiments. *Journal of Statistical Planning and Inference*, 43(3):381–402, 1995.

[34] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.

[35] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

[36] Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected improvements to expected improvement for bayesian optimization. In *Advances in Neural Information Processing Systems (NeurIPS 2023)*, 2023. 37th Conference on Neural Information Processing Systems.

[37] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. Botorch: A framework for efficient monte-carlo bayesian optimization. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, 2020. arXiv:1910.06403.

[38] Jacob R. Gardner, Geoff Pleiss, David Bindel, Kilian Q. Weinberger, and Andrew Gordon Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

[39] Christoph Molnar. *Interpretable Machine Learning*. 3 edition, 2025.

[40] Martin Schlueter, Stefan Erb, Matthias Gerdts, Stephen Kemble, and Jan-J. Rückmann. MIDACO on MINLP space applications. *Advances in Space Research*, 51(7):1116–1131, 2013.

[41] MIDACO Solver. *MIDACO Solver User Manual*, version 6.0 edition, April 2018. Official user manual.