aerospace

MDPI

Article

Analysis of Wing Structures via Machine Learning-Based Surrogate Models

Hasan Kiyik [1,*], Metin Orhan Kaya [1] and Peyman Mahouti [2]

[1] Department of Aeronautical and Astronautical Engineering, Faculty of Aeronautics and Astronautics, Istanbul Technical University, Istanbul 34469, Türkiye; kayam@itu.edu.tr
[2] Department of Electronics and Communication Engineering, Yıldız Technical University, Istanbul 34220, Türkiye; pmahouti@yildiz.edu.tr
* Correspondence: kiyik16@itu.edu.tr

Abstract

Accurate structural analysis is essential for the design and optimization of aircraft wings; however, repeated high-fidelity finite element analysis (FEA) becomes computationally expensive when embedded in iterative design loops. This study presents a machine learning-based surrogate modeling framework for the efficient analysis and optimization of metallic commercial wing structures. A detailed Airbus A320-like wing model was developed and analyzed in ANSYS 2023 R1 under modal, static, and eigenvalue buckling conditions. The general dimensions of the Airbus A320 wing were used only as a reference; the resulting model is a conceptual benchmark rather than a one-to-one geometric replica or a validated digital twin of a specific aircraft wing. Using Latin Hypercube Sampling, 340 high-fidelity samples were generated, with 300 samples used for training and validation and 40 retained as an independent holdout set. The proposed Pyramidal Deep Regression Network (PDRN), a deep learning-based surrogate model whose architecture is automatically tuned using Bayesian Optimization, was benchmarked against Artificial Neural Networks (ANNs), Ensemble Learning, Support Vector Regression (SVR), and Gaussian Process Regression (GPR). On the unseen test set, the PDRN achieved the best overall predictive performance, with RMS errors of 0.8% for mass, 3.1% for the first natural frequency, 11.5% for load factor, and 11.4% for safety factor. To evaluate its practical utility, the trained PDRN was embedded into a PSO-based optimization framework for mass minimization under minimum safety factor, load factor, and first-frequency constraints. The surrogate-guided optimum was verified in ANSYS and remained feasible, yielding a mass of 10,485 kg, a first natural frequency of 1.4142 Hz, a load factor of 1.307, and a safety factor of 1.158. Compared with direct ANSYS in-the-loop optimization, the proposed workflow reached a comparable feasible design with substantially fewer high-fidelity evaluations. These results demonstrate that the PDRN provides an accurate and computationally efficient surrogate for rapid wing analysis and constraint-driven structural optimization.

Keywords: surrogate modeling; deep learning; wing structures; buckling analysis; modal analysis

Check for updates

Academic Editor: Sujin Bureerat

Received: 8 March 2026
Revised: 30 March 2026
Accepted: 31 March 2026
Published: 3 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The structural integrity and dynamic performance of aircraft wings are central to the safety, efficiency, and economic viability of commercial aviation. Buckling and modal analyses are foundational in the design and certification of wing structures. Buckling analysis

ensures that wings can withstand compressive and shear loads without catastrophic instability, directly impacting the load-carrying capacity and preventing failure below design limits [1]. Modal analysis characterizes the natural vibration modes, which is essential for avoiding resonance and ensuring aeroelastic stability, as changes in modal properties can signal proximity to buckling or other critical states [2]. High-fidelity simulations and experimental studies confirm that integrating buckling and modal analyses into the design process leads to safer, more efficient, and certifiable wing structures [3–5]. These analyses are not only vital for initial design but also for ongoing optimization and structural health monitoring, further supporting the safety and operational reliability of commercial aircraft.

In the context of commercial aircraft, metallic wing structures remain an important and practically relevant benchmark for structural sizing and optimization studies, owing to their continued industrial use and comparatively well-understood structural behavior [5,6]. Liang and Yin [5] present a nonlinear buckling analysis of a large commercial aircraft's metallic wing structure, focusing on realistic flight cases. Their work underscores the necessity of precise buckling analysis for metallic wings, as these structures are prone to local instability under operational loads. The study demonstrates that targeted reinforcement based on detailed buckling analysis can significantly improve the safety and efficiency of metallic wing designs in commercial aviation. A study by Houston et al. [6] investigates metallic stiffened wing panels, emphasizing the ongoing relevance of metallic structures in commercial aircraft. The work demonstrates that advances in manufacturing (such as welding and extrusion) continue to enhance the performance and cost-effectiveness of metallic wing panels. The authors highlight that accurate buckling analysis is essential for optimizing mass and ensuring structural stability, with experimental results showing significant mass savings and validated sizing methods for metallic wing covers. However, the increasing complexity of modern wing designs, driven by demands for higher efficiency, reduced emissions, and novel configurations (e.g., high-aspect-ratio wings, morphing structures), has rendered traditional high-fidelity finite element analysis (FEA) approaches computationally expensive, especially when embedded within optimization loops [7–9]. Gray and Martins [7] highlight that as wing designs become more flexible and incorporate higher aspect ratios, the need for geometrically nonlinear high-fidelity FEA increases, which significantly raises computational costs—especially when such analyses are embedded within optimization loops for structural sizing and aeroelastic performance. Shizuno et al. [8] note that nonlinear FEA for high-aspect-ratio wings, which are common in modern efficient aircraft, becomes computationally expensive as the number of elements and the complexity of the wing shape increase, motivating the development of reduced-order methods to mitigate this cost. Stanford and Beran [9] discuss the large computational burden associated with gradient-based optimization of nonlinear flapping wing structures using FEA, and present cost reduction techniques to make such analyses more tractable for complex, highly flexible wings. Liang and Yin [5] specifically address the high computational cost and convergence challenges of nonlinear FEA for large and complex wing structures, especially when realistic flight cases and optimization constraints are considered.

Optimization of wing structures for weight, strength, or dynamic performance typically requires thousands of analysis iterations, with each iteration often involving a full-scale finite element analysis (FEA) for buckling and modal response. This process can result in simulation times that are prohibitive or even infeasible for practical engineering workflows, especially as the complexity of the optimization problem and the fidelity of the models increase [5,10,11]. The computational bottleneck becomes even more pronounced when uncertainties, multidisciplinary constraints, or the need for real-time decision support—such as in digital twin frameworks—are considered, as these factors demand

even more iterations and higher-fidelity simulations [12–14]. To address these challenges, researchers have explored strategies like multi-scale FEA, surrogate-based optimization, and advanced algorithms to balance computational cost with optimization accuracy, but the fundamental issue of high computational demand remains a significant concern in the field [10–14]. It should be noted, however, that the present study is not based on geometrically nonlinear FEA. Instead, the surrogate models developed here are trained on linear elastic static, modal, and eigenvalue buckling analyses, which were intentionally selected as a computationally manageable benchmark for methodological validation in the conceptual sizing of a metallic wing structure.

To address the computational challenges of high-fidelity analysis in aerospace structural optimization, the aerospace community has increasingly adopted data-driven surrogate modeling techniques, particularly those based on machine learning (ML) such as Gaussian Process Regression (GPR), neural networks, ensemble learning, and support vector regression (SVR) [15–18]. These surrogate models are trained on a limited set of high-fidelity simulations and can then rapidly and accurately predict structural responses like buckling loads and natural frequencies across a wide design space, dramatically reducing computational cost [15–19]. This efficiency enables not only faster optimization and uncertainty quantification but also supports real-time applications such as structural health monitoring and decision-making in digital twin frameworks [15,20].

Recent research provides strong evidence for the effectiveness of deep neural networks and ensemble learning in multi-fidelity surrogate modeling for aerodynamic and structural optimization. Multi-fidelity deep neural network (DNN) surrogate models have been shown to significantly improve optimization efficiency by blending high- and low-fidelity simulation data, enabling accurate predictions of aerodynamic performance with far fewer expensive high-fidelity simulations [16,21–23]. For example, frameworks using transfer learning and multi-task learning with DNNs can reduce optimization costs by over 90% compared to single-fidelity approaches, while maintaining or even improving prediction accuracy for complex design tasks such as airfoil and wing optimization [24].

While much of the recent literature has explored surrogate-assisted methods in the context of both metallic and composite wings, and across civil and military sectors, the present work focuses on a metallic wing benchmark in a commercial-aircraft setting. This choice was made primarily to provide a structurally meaningful, industrially relevant, and computationally manageable benchmark for methodological validation of the proposed surrogate framework. In this study, the general dimensions of the Airbus A320 wing were used only as a reference to define a realistic structural benchmark. The resulting Airbus A320-like wing model is a conceptual design platform for methodological validation and comparative assessment; it was not developed as a one-to-one replica or as a validated digital twin of a particular industrial wing. Although more advanced nonlinear and composite-wing problems may further amplify the value of surrogate modeling, the present study intentionally adopts a linear elastic metallic-wing benchmark in order to validate the proposed methodology under a controlled and interpretable structural setting. The following sections therefore focus on demonstrating the capability of machine learning-based surrogate models, including deep learning-based approaches, to reduce the computational burden of repeated high-fidelity structural analyses within this benchmark problem.

2. Wing Model and Generation of Data Sets for Surrogate Models

In this section, the details of the studied wing model and the required steps for creating the surrogate model are presented.

2.1. Wing Model

The wing geometry, based on the NACA 23015 airfoil profile, was modeled using CATIA V5. The structural assembly consists of two C section spars, eleven C section ribs and eight L section stringers. Figure 1 illustrates the wing geometry designed in CATIA. The Airbus A320 is selected as the reference platform for this study [25]. Table 1 presents the wingspan, airfoil type, and the locations of the front and rear spars.

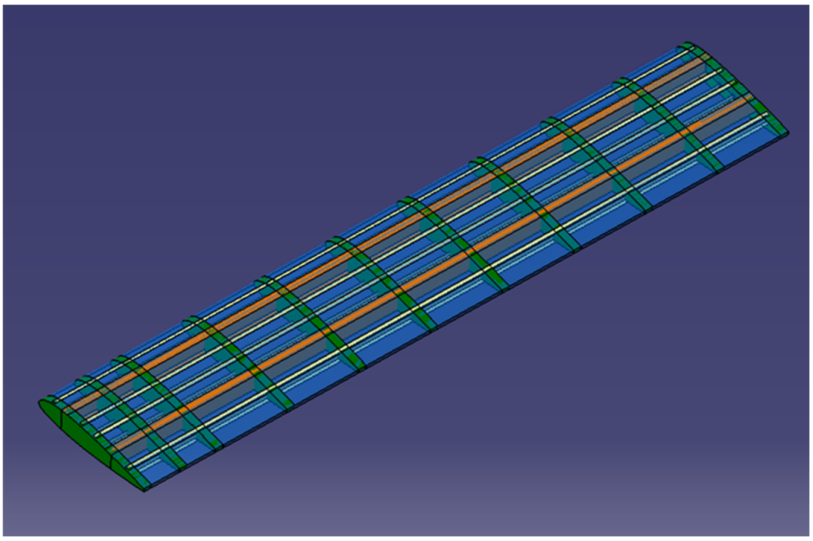

Figure 1. CATIA V5 Model of The Wing.

Table 1. General Specifications for Airbus A320.

Parameter	Value
Wingspan	33.910 [m]
Airfoil	NACA 23015
Position of front spars	20% of the chord [m]
Position of rear spars	65% of the chord [m]

For the preliminary structural design, Aluminum Alloy 7050-T7451 is selected for the wing skin, spars, ribs, and stringers, which are critical for enhancing buckling resistance. This material choice aligns with standard aerospace practices for metallic airframes. The material properties are derived from the MMPDS (Metallic Materials Properties Development and Standardization) handbook [26], specifically adhering to the AMS 4050 specification (plate form) due to its favorable balance of high strength, fracture toughness, and resistance to stress corrosion cracking. The mechanical properties of Aluminum Alloy 7050-T7451 are presented in Table 2. In the present study, Aluminum Alloy 7050-T7451 was modeled as a linear elastic material for all simulations used in dataset generation. Nonlinear elastic–plastic material behavior was not included, because the feasible design space was intentionally restricted to remain within the elastic regime through yield- and buckling-related structural constraints.

To simulate the wing–fuselage attachment, a Remote Displacement boundary condition was applied to the root interface faces of the front and rear spars over a spanwise length of approximately 300 mm. In this formulation, all six degrees of freedom were constrained, i.e., the translations (U_X, U_Y, U_Z) and the rotations (ROT_X, ROT_Y, ROT_Z) were set to zero in order to represent the attachment of the wing root to the fuselage support structure. At the same time, the "Deformable" behavior option was selected to avoid artificial localized stiffening at the constrained region. Unlike a rigid coupling, the deformable formulation

allows the underlying surface nodes to preserve their local deformation compatibility, thereby enabling a more physically realistic transfer of the bending and torsional moments from the wing root into the support region. The ANSYS model and the applied boundary conditions are presented in Figure 2.

Table 2. Mechanical Properties of Aluminum Alloy 7050-T7451.

Parameter	Value
Poisson ratio	0.33
Density	2823 [kg/m^3]
Young's modulus	71,016 [MPa]
Compressive yield strength	434 [MPa]
Tensile yield strength	441 [MPa]
Shear modulus	26,698 [MPa]

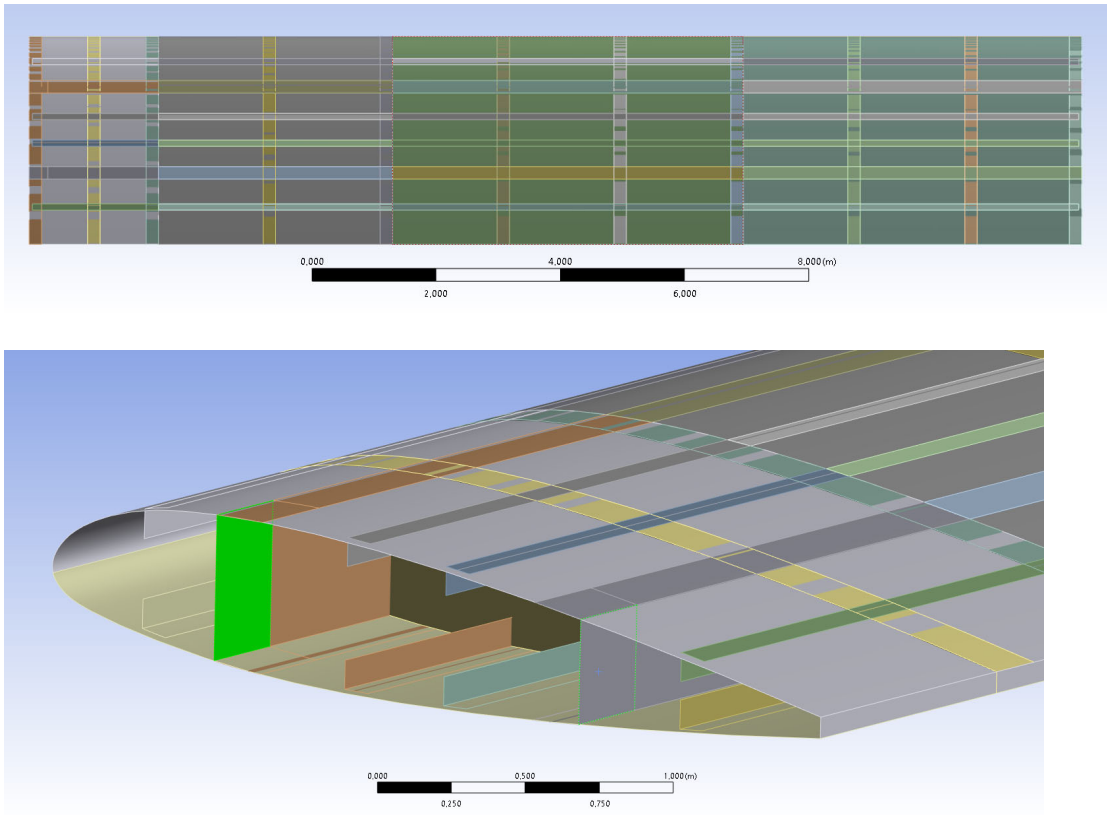

Figure 2. ANSYS Model of The Wing and Boundary Conditions.

The finite element discretization was performed using linear-order elements, with a global element size of 50 mm. To improve resolution at structurally important interfaces, the mesh was locally refined to 25 mm in the component connection regions, resulting in a total of 192,996 nodes and 183,062 elements. This discretization strategy was selected as a practical compromise between numerical resolution and computational feasibility, because the present study required a large number of repeated high-fidelity analyses for dataset generation and optimization comparison. Under the available computational resources, substantially finer global meshes or higher-order element formulations were not practically feasible within the scope of the study. Accordingly, the present mesh should be interpreted as a benchmark-level discretization suitable for global structural response prediction, while more refined meshes would be desirable for detailed local stress assessment and formal

mesh convergence verification. A formal mesh convergence study with progressively finer meshes was not performed in the present work and is acknowledged here as a limitation.

The structural interactions between the ribs, spars, stringers, and skin were defined as bonded contacts using the multipoint constraint (MPC) formulation to effectively link the nodes across mating surfaces. Additionally, to guarantee robust contact detection and accommodate geometric tolerances, a pinball region with a radius of 30 mm was defined. In the present benchmark model, this bonded MPC formulation was adopted to ensure stable and computationally efficient transfer of loads across the structural interfaces during repeated dataset generation analyses. It should be noted, however, that this assumption simplifies the real behavior of riveted or bolted joints and may underestimate local stress concentrations and local connection flexibility at the interface regions. The generated mesh structure of the wing is shown in Figure 3.

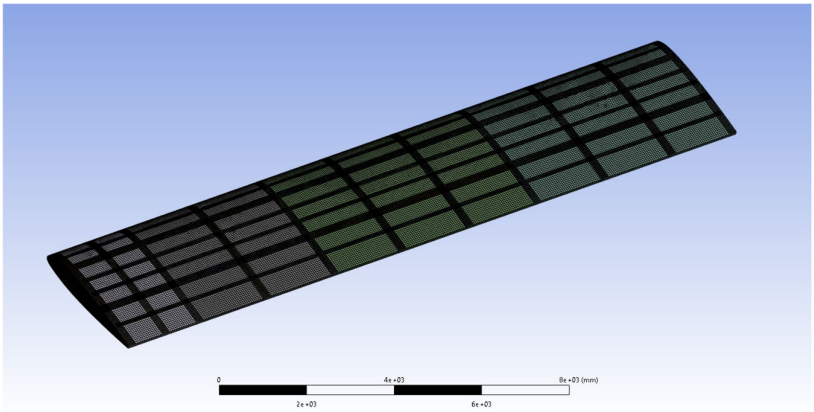

Figure 3. Mesh Preview for the Wing Model.

The structural loading of the wing was modeled to reflect critical operational conditions. A distributed fuel mass of 6400 kg was assigned to the lower wing skins, representing the maximum fuel capacity for a single semi-span configuration. This value was determined based on the Airbus A320 Aircraft Characteristics for Airport and Maintenance Planning data [27], which specifies a total wing fuel capacity of approximately 15,591 to 15,959 L. Assuming a standard jet fuel density of 0.785 kg/L, the assigned mass aligns with the official technical specifications. Furthermore, a volumetric consistency check performed on the 3D CAD model confirmed that the internal wing cavity provides sufficient volume to accommodate the designated fuel mass, ensuring the physical validity of the simulation model.

The aerodynamic pressure loads defined in the three zones are applied normal to the lower wing surface. This directionality is selected because fluid pressure acts perpendicular to the structural boundaries. In accordance with the simplifications inherent in Schrenk's method [28], the calculated spanwise lift distribution is applied as a distributed pressure on the pressure side (lower skin) to simulate the upward bending moment. Specifically, the loading profile is defined as 0.25 bar for the root section (spanning 5851.67 mm), 0.15 bar for the mid-span region (spanning 5651.67 mm), and 0.02 bar for the tip section (spanning 5451.66 mm). The 3000 N aileron load was derived based on the lift increment characteristics described in Schrenk's (NACA TM 948) [28]. For a typical A320 approach speed of 75 m/s and an aileron area of 1.6 m^2, the localized aerodynamic force was analytically calculated as 2976.75 N. This value was subsequently rounded to 3000 N to provide a conservative margin for maneuvering load cases and to simplify the numerical input for the optimization process. To maintain physical fidelity while simplifying the structural complexity of the control surface attachments, this force was applied as a distributed load on the rear spar

surface within Zone 2. The load was specifically positioned at the spanwise midpoint of the wing, with a 200 mm chordwise offset from the rear spar to simulate the resultant torque.

Regarding the mass and load definitions, the modal analysis considers the fuel inventory strictly as a non-structural mass contribution alongside the wing components to account for inertial effects. It should be noted that aerodynamic pressure loads and the aileron force are excluded from the modal analysis, as this simulation focuses solely on the undamped free vibration characteristics of the wing structure under its mass distribution. However, in both the static and eigenvalue buckling simulations (Figure 4), a consistent loading scenario is applied; this includes the spanwise aerodynamic pressure distribution, the 3000 N aileron load, and the full integration of gravitational acceleration (g). Applying the same load set for both static and buckling analyses ensures that the structural stability limits (buckling) are evaluated under the exact operational stress states experienced by the wing. This distinction ensures that while the modal analysis identifies the fundamental dynamic modes, the static and buckling simulations accurately reflect the wing's performance and safety margins under combined aerodynamic and gravitational loading.

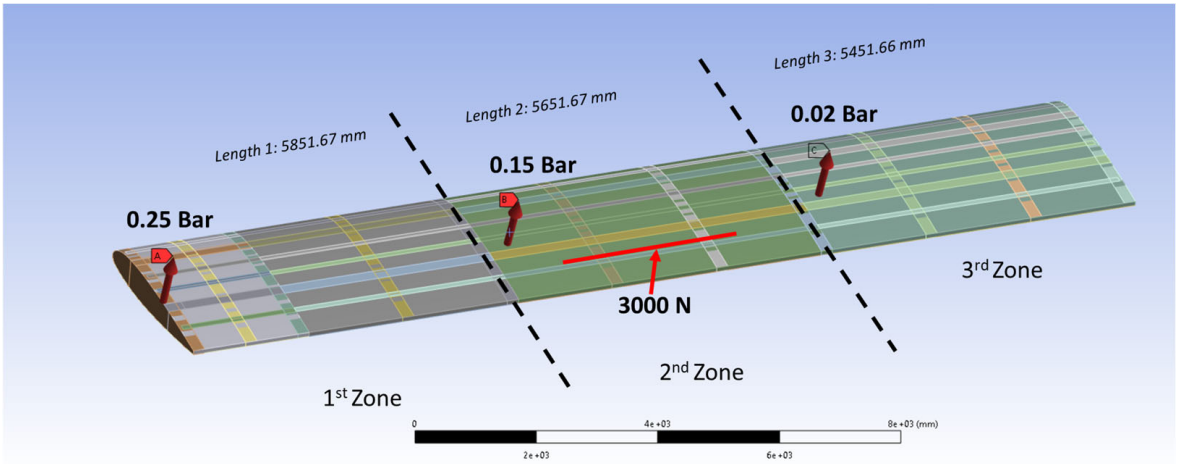

Figure 4. Eigenvalue buckling analysis of the wing and boundary conditions.

Accordingly, the high-fidelity dataset generated in this study represents a linear analysis regime and should be interpreted within that scope. The surrogate models trained on this dataset therefore inherit the assumptions of linear elastic static response and eigenvalue-based buckling prediction.

Structural components within each wing section, specifically the front and rear spars, upper and lower skins, and stringers upper and lower 1 through 4 are assigned distinct thickness parameters. For the rib structure, the eleven ribs positioned along the span are defined with independent thickness parameters, allowing for discrete sizing optimization of each rib.

However, to reduce the dimensionality of the input space and improve computational efficiency, specific parameters were assigned equal values during the dataset generation process. The general distribution of the thickness parameters is illustrated in Figure 5.

To reduce the computational complexity of the optimization process, several structural parameters were coupled based on their spanwise locations. For the wing panels, the thicknesses of the upper and lower skins were maintained identical across all sections. Specifically, independent skin thickness parameters were defined for Section 1 and Section 2, while the skins for Sections 3 and 4 were assigned a unified thickness value (Figure 5). A similar approach was applied to the stringers; the upper and lower stringer thicknesses were equated throughout the wing. While a distinct thickness was assigned to the stringers

in Section 1, a single unified parameter was used for all stringers spanning Sections 2, 3, and 4.

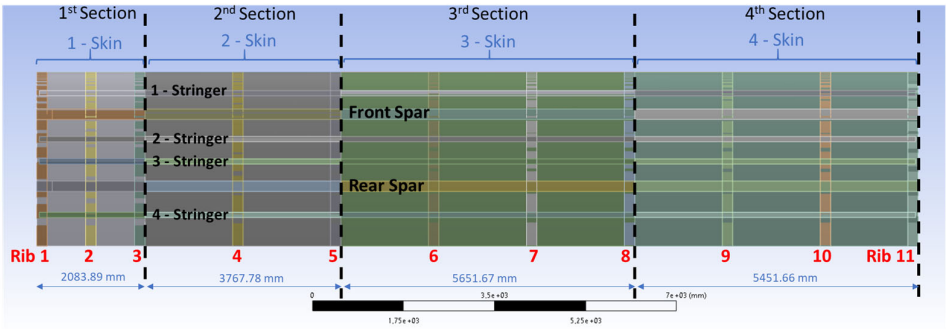

Figure 5. Thickness parameters of the wing.

The wing model incorporates a total of 11 ribs, which were grouped to simplify the design space. Accordingly, Ribs 1–2 and Ribs 3–4 were assigned identical thicknesses within their respective groups, while a single thickness parameter was defined for the remaining ribs from Rib 5 to Rib 11. Regarding the spar structures, the thicknesses of the front and rear spars were set to be equal within each section. A specific spar thickness was defined for Section 1, whereas the spar thickness parameters for Sections 2, 3, and 4 were coupled and assigned a unified value. All optimizable thickness parameters are detailed in Table 3.

Table 3. Design Variable Definitions and Component Grouping Configuration.

Component	Geometric Region	Variable Grouping/Constraint	Parameter and Thickness Range [mm]
Ribs	-	$t_{Rib1} = t_{Rib2}$	P_1: 5–15
	-	$t_{Rib3} = t_{Rib4}$	P_2: 1–10
	-	$t_{Rib5} = t_{Rib6} = \ldots = t_{Rib11}$	P_3: 1–10
Spars	Section 1	$t_{front\ spar} = t_{rear\ spar}$	P_4: 15–25
	Sections 2, 3, 4	$t_{front\ spar} = t_{rear\ spar}$	P_5: 1–10
Skins	Section 1	$t_{uppper\ skin} = t_{lower\ skin}$	P_6: 15–25
	Section 2	$t_{uppper\ skin} = t_{lower\ skin}$	P_7: 5–15
	Sections 3 & 4	$t_{uppper\ skin} = t_{lower\ skin}$	P_8: 1–10
Stringers	Section 1	$t_{stringer1} = t_{stringer2} = \ldots = t_{stringer4}$	P_9: 5–15
	Sections 2, 3, 4	$t_{stringer1} = t_{stringer2} = \ldots = t_{stringer4}$	P_{10}: 1–10

Following the definition of all design parameters, loading conditions, and material properties, the simulation workflow was established within ANSYS Workbench. As illustrated in Figure 6, the project schematic was configured to include the Modal, Static Structural, Eigenvalue Buckling, and Direct Optimization systems.

Based on the initial parameters specified in Table 4, the results obtained from the Modal, Static Structural, and Eigenvalue Buckling analyses are presented in the following sections. The natural frequencies obtained from the modal analysis for the first five modes are presented in Table 5.

The visual inspection of the mode shapes reveals the critical dynamic characteristics of the wing structure. Mode 1 (1.355 Hz) corresponds to the fundamental out-of-plane bending, representing the primary flexible axis under lift loads. Mode 2 (5.597 Hz) indicates the in-plane bending stiffness, while Mode 3 (7.383 Hz) represents the first torsional mode. The frequency separation between the fundamental bending and torsional modes suggests a distinct structural response, which is crucial for aeroelastic stability. Mode 4 (9.444 Hz) demonstrates a higher-order bending behavior coupled with torsional effects.

However, while the first four modes represent the global dynamic characteristics of the wing structure, the fifth mode at 17.814 Hz exhibits localized skin vibration modes rather than global deformation. Therefore, the visual representations are focused on the fundamental global modes (Modes 1–4), which are the primary drivers for structural design. The visual representations of the first four mode shapes are presented in Figure 7.

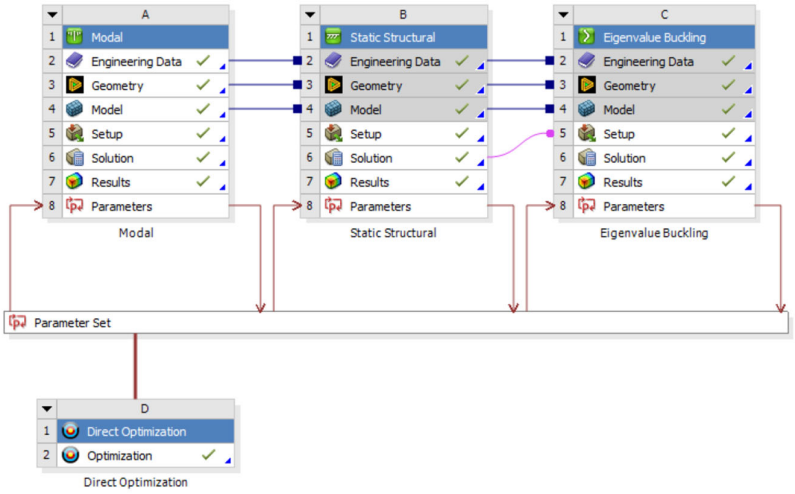

Figure 6. Schematic of ANSYS Workbench.

Table 4. Initial Design Parameters and Thickness Values.

Parameters	Initial Thickness Value [mm]
$P_{1\text{-}i}$	10 mm
$P_{2\text{-}i}$	7 mm
$P_{3\text{-}i}$	5 mm
$P_{4\text{-}i}$	21 mm
$P_{5\text{-}i}$	10 mm
$P_{6\text{-}i}$	14 mm
$P_{7\text{-}i}$	9 mm
$P_{8\text{-}i}$	6 mm
$P_{9\text{-}i}$	8 mm
$P_{10\text{-}i}$	6 mm

Table 5. Modal analysis results.

Order	Frequency, f (Hz)	Modal Shape
1	1.355	1st Out-of-Plane Bending
2	5.597	1st In-Plane Bending
3	7.383	1st Torsional Mode
4	9.444	2nd Out-of-Plane Bending
5	17.814	Local Skin Vibration/Higher Order Coupled Mode

The results of the eigenvalue buckling analysis, including the first four positive critical load multipliers and their respective locations, are summarized in Table 6. The analysis indicates that the fundamental buckling mode (Mode 1) occurs at a load multiplier of 1.1163. This value confirms that the wing structure maintains its elastic stability under design limit loads and remains within the safe operational regime (Load Multiplier > 1.0).

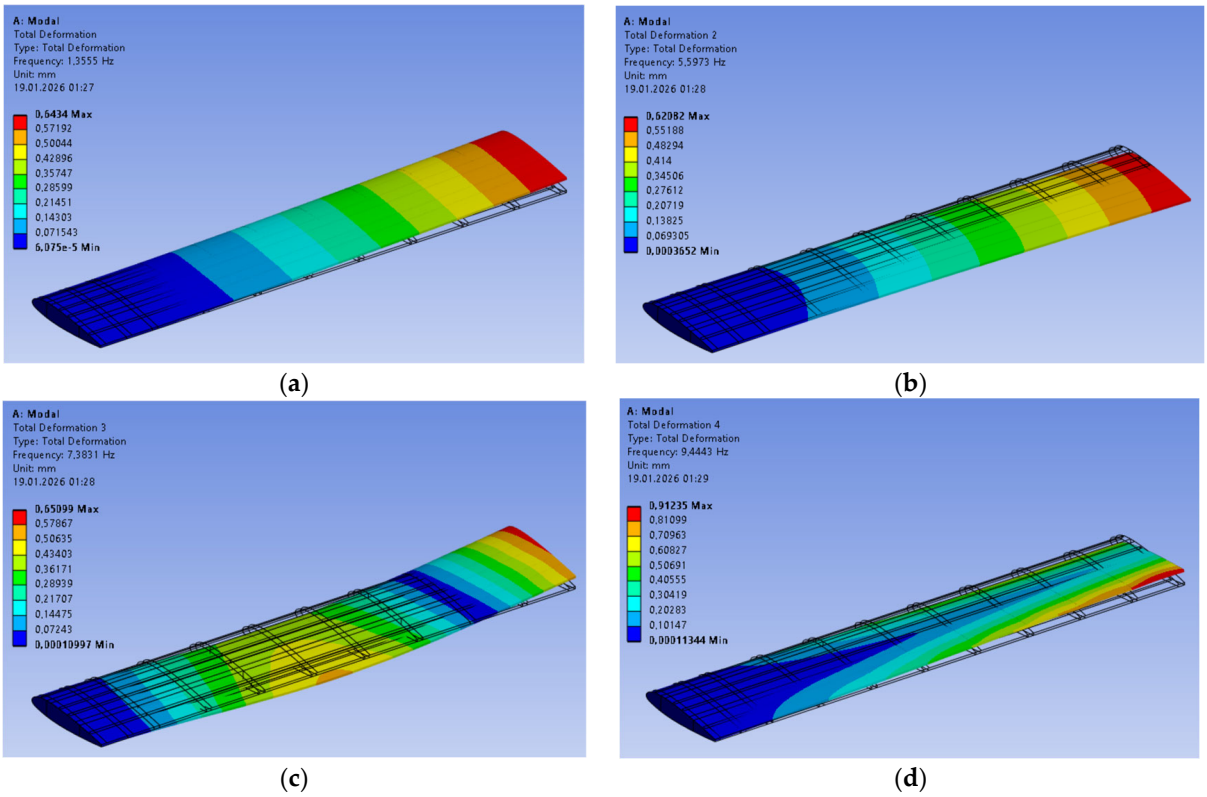

Figure 7. The first four fundamental natural frequency mode shapes of the baseline wing model: (**a**) 1st Out-of-Plane Bending Mode (1.355 Hz), (**b**) 1st In-Plane Bending Mode (5.597 Hz), (**c**) 1st Torsional Mode (7.383 Hz), (**d**) 2nd Out-of-Plane Bending Mode (Coupled with Torsion) (9.444 Hz).

Table 6. Critical buckling modes and load factors.

Order	Load Factor
1	1.1163
2	1.2384
3	1.445
4	1.4851

The distribution of the first four eigenvalues shows a significant "clustering" effect, as seen in Table 6, where the values are remarkably close to one another. This phenomenon is attributed to the repetitive geometric nature of the stringer-stiffened panels in Section 2, where identical bay widths and thicknesses result in nearly simultaneous instability across adjacent sections. The physical characteristic of this critical instability is illustrated in Figure 8. As shown in the figure, the failure mode is identified as local skin buckling. It is observed that while the skin panels between the stringers exhibit visible wave patterns, the primary load-bearing components—specifically the spars, ribs, and stringers—remain structurally stable and retain their original form. This behavior is highly desirable in semi-monocoque aircraft structures, as it demonstrates a safe failure hierarchy where local skin buckling occurs well before the catastrophic global collapse of the primary spars.

According to the initial static structural analysis results, the maximum equivalent (von Mises) stress was calculated as 421.47 MPa. The minimum yield-based safety factor, defined from the linear elastic static analysis as the ratio of the material yield strength to the maximum von Mises stress under the applied loading, was identified as 1.0463, occurring on the rear spar surface proximal to the wing root. These results indicate that the baseline

design, under the defined operational loads, operates very close to the material's yield limits at the root section. These results indicate that the baseline design, under the defined operational loads, operates very close to the material's yield limits at the root section. In the initial configuration of the wing model, the total mass was calculated as 10,853 kg. This total comprises a fuel mass of 6400 kg and a structural mass (wing components) of 4453 kg. For the subsequent optimization process, the fuel mass was treated as a constant non-structural mass, while the structural mass was defined as the primary objective function for reduction through thickness variations.

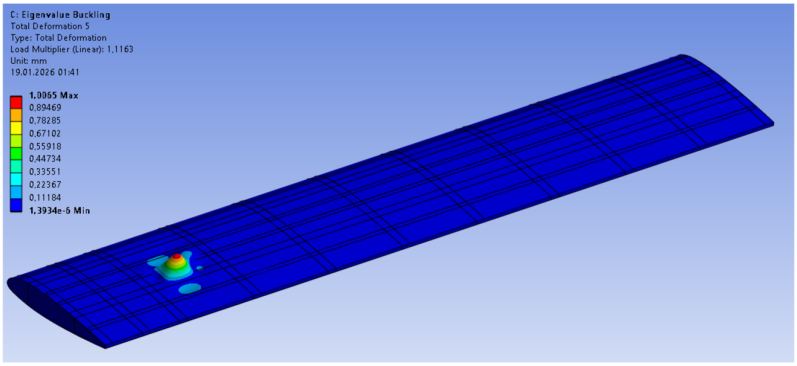

Figure 8. Visual representation of the first critical buckling mode (Mode 1) occurring on the upper skin panel. The instability is characterized by local skin buckling with a critical load multiplier of 1.1163.

2.2. Generation of Data Sets

The development of accurate and reliable machine learning-based surrogate models necessitates a high-quality dataset that effectively maps the relationship between the input design parameters and the structural performance outputs. This section details the methodology employed for defining the design space, selecting an appropriate sampling strategy, and executing the high-fidelity Finite Element Analysis (FEA) simulations.

The design space for this study is defined by ten independent thickness parameters (P_1 to P_{10}), which are strategically categorized into functional groups as detailed in Table 3. To manage the dimensionality of the design space and improve the convergence of the surrogate model, specific structural components were coupled based on their spanwise locations and symmetric roles. This approach involves grouping ribs into distinct sets, equating upper and lower skin thicknesses, and unifying stringer and spar thicknesses across corresponding sections. These ten parameters serve as the primary inputs for the optimization process, where their values are varied to capture the structural response. The dependent output variables tracked are structural mass, natural frequencies, eigenvalue buckling load factors, and minimum yield-based safety factors obtained from the linear elastic static analysis. This structured configuration allows for an efficient exploration of the trade-offs between mass reduction and structural integrity requirements. Traditional grid-based or full-factorial sampling methods quickly become computationally infeasible as the number of variables increases, due to the exponential growth in required sample size (the so-called "curse of dimensionality"). For example, even a coarse grid with four levels per variable in a ten-dimensional space would require 4^{10} = 1,048,576 simulations, which is often prohibitive for engineering applications [29]. Random sampling approaches, such as Monte Carlo, can also be inefficient in high dimensions, frequently resulting in uneven coverage and leaving portions of the design space unexplored [29]. Recent advances propose adaptive and sparse grid sampling methods, which can dramatically reduce the number of required simulations while maintaining accuracy, with some approaches achiev-

ing comparable performance using only a fraction of the samples needed by conventional methods [29,30].

To ensure an efficient and representative sampling of the design space, the Latin Hypercube Sampling (LHS) method was employed. LHS is a stratified sampling technique specifically designed for computer experiments. It operates by dividing the range of each input variable into N equally probable intervals, where N is the desired number of simulation runs. LHS guarantees that exactly one sample is placed within each interval for every input variable, and is widely recognized for its improved space-filling properties in high-dimensional engineering problems [31,32].

The primary advantage of LHS for surrogate modeling lies in its superior space-filling properties. By maximizing the stratification of the samples, LHS ensures uniform coverage of the entire design space, including the boundaries and corners. This is crucial because surrogate models rely on interpolation between known data points; therefore, gaps (voids) in the training data can lead to significant prediction errors. LHS mitigates this risk, enabling the construction of accurate surrogate models with a relatively smaller number of computationally expensive high-fidelity simulations compared to other methods [33–36]. The LHS methodology was implemented using the Design of Experiments module within ANSYS Workbench. Parametric analyses were conducted by automatically varying the ten thickness parameters according to the LHS design and executing the coupled modal and eigenvalue buckling analyses for each design point. Using the LHS procedure, a total of 340 samples were generated and randomly divided into two disjoint subsets: 300 samples for training/validation and 40 samples for independent holdout testing. The 300-sample subset was used for model development, including training and hyperparameter tuning via 5-fold cross-validation. The 40-sample holdout subset was kept completely isolated during this process and was used only after model finalization to assess generalization performance on unseen data. This setup allows the generalization capability of the proposed approach to be evaluated explicitly under limited training data.

2.3. Regression Methods for Surrogate Modeling

The generation of a high-fidelity data set, as detailed in the preceding section, provides the necessary empirical foundation for the development of a computationally efficient surrogate model. The primary objective of this stage is to construct a fast yet accurate predictive model capable of emulating the complex physical responses of the wing structure, thereby circumventing the need for computationally expensive simulations in subsequent design and analysis tasks. This work adopts and evaluates a pyramidal deep learning architecture, referred to here as the Pyramidal Deep Regression Network (PDRN), as the primary surrogate modeling technique for wing structural response prediction. To rigorously validate the performance and establish the efficacy of the proposed PDRN, its predictive capabilities are systematically benchmarked against a curated selection of established regression methods, including both deep learning and non-deep learning approaches.

2.3.1. Counterpart Methods

To establish a comprehensive performance baseline, four distinct and widely recognized regression algorithms were selected as counterpart methods: Ensemble Learning (ENS), Support Vector Regression (SVR), Gaussian Process Regression (GPR), and the conventional Artificial Neural Network (ANN). These techniques were chosen not only for their demonstrated robustness and widespread adoption in engineering surrogate modeling applications but also for their diverse theoretical underpinnings. This selection ensures

that the proposed PDRN is evaluated against a spectrum of methodological philosophies, providing a more holistic and challenging assessment of its capabilities.

- Ensemble Learning

Ensemble learning is a meta-heuristic approach founded on the principle that combining the predictions of multiple individual models, or "base learners," can yield a single, consolidated model with superior predictive performance and robustness. This "wisdom of the crowd" paradigm leverages the diversity among individual models to produce more accurate and reliable outcomes than any single constituent model could achieve in isolation [37,38].

The theoretical strength of ensemble methods is often explained through the bias-variance tradeoff framework. While individual complex models may exhibit low bias on the training data, they often suffer from high variance, leading to poor generalization on unseen data—a condition known as overfitting. By strategically aggregating a diverse set of learners, ensemble techniques can effectively average out the individual model errors, leading to a significant reduction in the variance component of the total prediction error, often without a substantial increase in bias. This results in models with enhanced generalization capabilities, a critical attribute for any surrogate intended for predictive applications.

Given their capacity to achieve state-of-the-art performance, ensemble methods have been successfully applied to construct high-accuracy surrogate models for a wide range of complex engineering problems. Its inclusion as a counterpart method is therefore essential, as it represents a high-performance, non-parametric benchmark that sets a demanding standard for the predictive accuracy of the proposed PDRN.

- Support Vector Regression Machine

Support Vector Regression (SVR) is a powerful regression algorithm derived from the rigorous mathematical framework of statistical learning theory, specifically Vapnik's structural risk minimization principle. In contrast to conventional regression techniques that aim to minimize the empirical error on the training data, SVR seeks to minimize an upper bound on the generalization error, thereby building a model that is inherently designed for robust performance on new, unseen data [39,40].

The operational mechanism of SVR is distinguished by its use of an ϵ-insensitive loss function. This function defines a margin of tolerance, often visualized as a "tube" of radius ϵ, around the regression hyperplane. Data points falling within this tube do not contribute to the regression error, while points outside are penalized. This unique property leads to a sparse solution, where the final regression model is defined only by a subset of the training data—the so-called support vectors—which are the points that lie on or outside the margin boundaries. This reliance on a small subset of data points makes SVR particularly robust to the presence of outliers.

A core strength of SVR is its ability to model highly non-linear relationships through the "kernel trick". By employing kernel functions, such as the polynomial or Radial Basis Function, SVR implicitly maps the input data into a high-dimensional feature space. Within this space, a linear separation or regression can be effectively performed. This allows SVR to capture complex, non-linear phenomena, such as those governing wing structural dynamics, without the need to explicitly compute the transformations into this higher-dimensional space. Due to its robust theoretical foundation and proven efficacy, SVR is widely employed for surrogate modeling in various engineering disciplines. It serves as an ideal counterpart method because its statistical learning approach is fundamentally different from that of neural networks, providing a methodologically distinct benchmark.

- Kriging (Gaussian Process Regression, GPR)

Kriging is a probabilistic surrogate modeling approach that has become a standard reference method in many engineering design and analysis studies [41]. It constructs an approximation of an expensive simulation by assuming that the underlying response varies smoothly and that points that are close to each other in the input space tend to produce similar outputs. The surrogate is built by combining (i) a global trend describing the overall behavior of the response and (ii) a local correction term that captures deviations from this trend through learned correlations in the data.

A defining feature of Kriging is that it is not limited to producing a single predicted value; it also provides an estimate of predictive uncertainty [42,43]. This uncertainty quantification is especially valuable when the available training data are limited or costly, because it indicates where the model is confident and where it is likely extrapolating. As a result, Kriging is frequently used in data-efficient workflows such as design space exploration, sensitivity analysis, and sequential sampling, where new simulations can be strategically selected to reduce uncertainty and improve surrogate fidelity. The predictive capability of Kriging depends on the choice of correlation (kernel) function, which controls how rapidly similarity decays with distance and how smooth or rough the predicted response can be. Its hyperparameters are commonly calibrated from the training data using likelihood-based optimization, allowing the model to adapt its effective length scales and noise level to the problem at hand. When properly tuned, Kriging is well known for achieving high accuracy with relatively small datasets, particularly for smoothly varying structural responses. Kriging is included here as a counterpart method because it offers a rigorous statistical baseline that is fundamentally different from neural network-based regressors. Its data efficiency and built-in uncertainty awareness provide a strong and methodologically distinct benchmark for assessing the predictive performance and practical advantages of the proposed PDRN.

- Artificial Neural Network

The Artificial Neural Network (ANN), particularly the feedforward Multi-Layer Perceptron (MLP), serves as the foundational architecture for modern deep learning [44–46]. Inspired by the structure of biological nervous systems, an ANN is composed of interconnected nodes, or neurons, organized in layers: an input layer, one or more hidden layers, and an output layer. The theoretical power of ANNs stems from the universal approximation theorem, which states that a feedforward network with a single hidden layer containing a finite number of neurons and a non-linear activation function can, in principle, approximate any continuous function to an arbitrary degree of precision. This capability enables ANNs to learn and represent complex, non-linear input-output mappings directly from data.

The learning process in an ANN is typically achieved through an algorithm known as backpropagation. During training, the network's internal parameters—the weights and biases of the connections between neurons—are iteratively adjusted to minimize a loss function that quantifies the discrepancy between the network's predictions and the actual target values. The presence of non-linear activation functions within the neurons is critical, as it allows the network to capture and model the non-linearities inherent in the data. The application of ANNs for surrogate modeling in engineering is a well-established and extensively documented practice. The inclusion of a conventional ANN as a counterpart method is crucial because it represents a direct and essential baseline. It allows for a clear contextualization of the architectural and methodological innovations of the proposed PDRN. Any performance improvements demonstrated by the PDRN can be directly attributed to its advancements over this fundamental neural network paradigm, addressing common challenges such as the manual, often heuristic, process of architecture selection.

2.3.2. Pyramidal Deep Regression Network (PDRN)

The Pyramidal Deep Regression Network (PDRN) employed in this study is a pyramidal deep learning architecture adapted from the framework reported by Koziel et al. [47] and specialized here for regression-based surrogate modeling of wing structural responses. The PDRN provides a parametrically defined network architecture together with an automated hyperparameter optimization process, thereby constituting a systematic approach to surrogate model development for highly nonlinear problems with limited data (Figure 9) [47]. Accordingly, the contribution of the present work is not the proposal of a fundamentally new neural network architecture, but the adaptation, implementation, and validation of the pyramidal deep network concept for a new aerospace structural analysis problem, together with its integration into a surrogate-assisted optimization workflow.

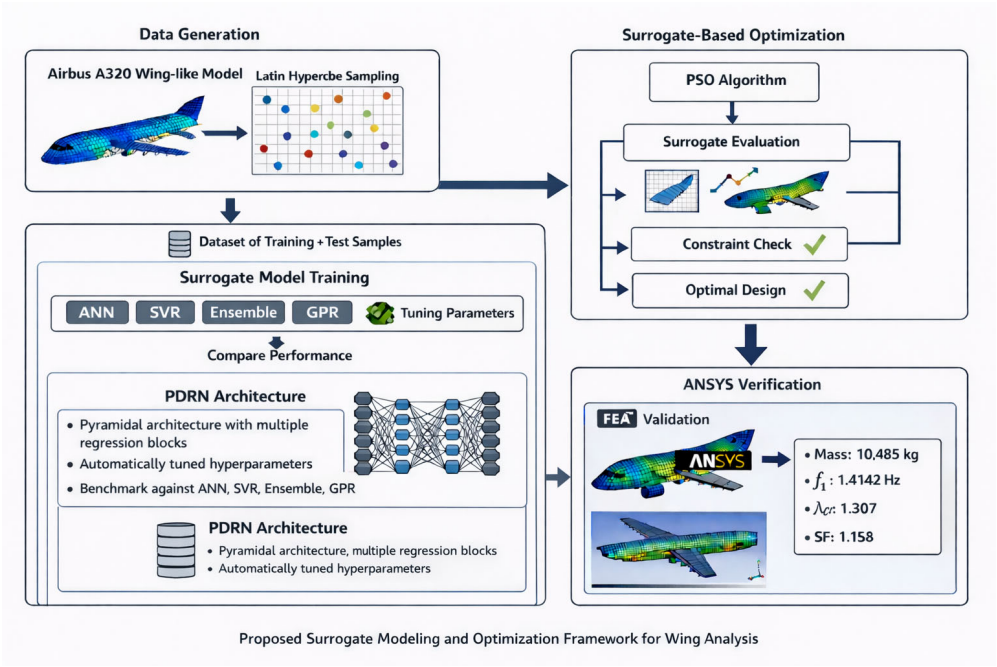

(**a**)

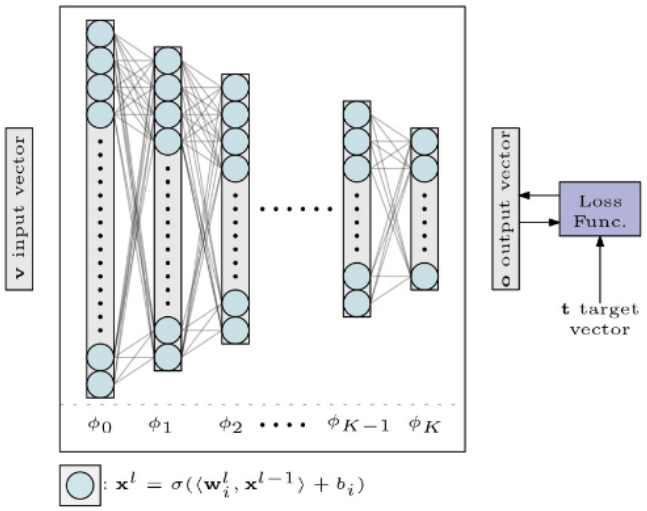

(**b**)

Figure 9. Schematic of the (**a**) Proposed surrogate modeling-based optimization framework, (**b**) PDRN generated by AI tool.

The central strategy of the PDRN is rooted in a two-stage information processing philosophy inspired by the feature-extraction mechanisms of Convolutional Neural Networks. First, the input parameter vector is transformed into a very high-dimensional space. This initial expansion is theoretically motivated by the need to increase the model's degrees of freedom, which facilitates the untangling and effective modeling of the complex, non-linear relationships that are characteristic of the physical responses in wing structures. Following this expansion, the network architecture is structured to systematically distill this high-dimensional information. This is achieved by progressively reducing the dimensionality through successive layers, culminating in the final low-dimensional output space that represents the predicted structural response. This entire process can be conceptualized as a "dimensional evolution" from input to output.

The PDRN topology is defined by a pyramidal, or trapezoidal, shape. The architecture is not determined by manually specifying the neuron count for each layer but is instead governed by a minimal set of three structural hyperparameters:

1. ϕ_K: The number of neurons in the final hidden layer.
2. ds: The depth of the model, defined as the number of hidden layers.
3. sc: A scale ratio that deterministically calculates the number of neurons in all preceding layers relative to ϕ_K. The number of neurons in the l-th layer from the output is given by $r(\phi_K \times sc^{l-1})$, where $r(\cdot)$ is a rounding operator.

Compared with conventional deep networks that use fixed or manually selected layer widths, the pyramidal topology provides a structured allocation of model capacity. The earlier hidden layers are wider, allowing the network to represent complex nonlinear interactions among the input variables, whereas the subsequent layers gradually compress these features toward the final response space. This progressive reduction acts as an architecture-level regularization mechanism and reduces the need for ad hoc specification of each hidden-layer width. In this formulation, ϕ_K, ds, and sc serve as compact structural descriptors of the network, transforming architecture selection into a low-dimensional and reproducible optimization problem.

It should also be emphasized that Bayesian optimization is used here to identify the best PDRN configuration for the present dataset based only on training/validation performance; it is not, by itself, a proof of the universal superiority of the pyramidal topology. Rather, the advantage of the topology is assessed empirically through comparison with counterpart models on the unseen holdout set.

This parsimonious parameterization transforms the otherwise vast and intractable search space of possible network architectures into a well-defined and controllable problem, which is amenable to automated optimization.

The PDRN architecture incorporates specific design choices to address known challenges in training deep neural networks. A critical component is the exclusive use of the Leaky Rectified Linear Unit (LReLU) activation function. This choice is a direct response to the "vanishing gradient" problem, which can stall the training of deep networks that use traditional saturating activation functions like sigmoid or tanh. LReLU maintains a small, non-zero gradient for all negative inputs, as defined by the leakage parameter α in Equation (1):

$$\sigma_\alpha(x) = \begin{cases} x, & x \geq 0 \\ \alpha x, & otherwise \end{cases} \quad (1)$$

This ensures that a gradient signal can consistently propagate backward through all layers of the network, enabling effective learning even in deep architectures. The leakage parameter α constitutes the fourth key hyperparameter of the PDRN model. To further enhance training stability and prevent overfitting, the framework mandates a Z-score

standardization preprocessing step for input data and integrates standard regularization techniques such as Dropout and Batch Normalization layers into the model during training.

A core theoretical strength of the PDRN framework is its automated approach to architecture and hyperparameter determination, which replaces the conventional, often subjective, trial-and-error tuning process. This is accomplished using Bayesian Optimization (BO), a powerful and sample-efficient global optimization strategy.

The BO algorithm iteratively searches for the optimal set of PDRN hyperparameters S = {ϕK, ds, sc, α}. It operates by building a probabilistic surrogate model—specifically, a Gaussian Process (GP)—of the PDRN's performance (e.g., its k-fold cross-validation loss) as a function of these hyperparameters. At each iteration, an acquisition function, such as Expected Improvement (EI), is maximized to select the next hyperparameter configuration to evaluate. This function intelligently balances exploration (sampling in regions where the model is most uncertain) and exploitation (sampling in regions predicted to yield high performance), guiding the search efficiently toward the global optimum. The process continues until a predefined computational budget is exhausted, and the final PDRN architecture selected is the one that yielded the minimum observed average loss. This automated search protocol ensures that the final model is systematically optimized for the specific dataset at hand, enhancing both performance and reproducibility.

Once the optimal architecture is determined via BO, the final PDRN model is retrained using the 300-sample training/validation dataset, while the 40-sample holdout set remains fully excluded from this process and is used only for the final independent performance evaluation. The model's weights and biases are optimized using the Adam optimizer, a stochastic gradient-based algorithm that is highly effective for training deep networks. The optimization objective is the minimization of the Mean Absolute Error (MAE) loss function, defined as $L_{MAE}(o) = E\{|o - t|\}$, where o is the network's prediction and t is the target vector.

3. Results

In this part of the study, a series of regression models, including both machine learning- and deep learning-based approaches, were deployed to construct surrogate models using the training and test datasets prepared in the earlier stages (300 and 40 samples, respectively). The predictive performance of these algorithms was rigorously assessed using the Relative MAE and RMS error. To provide a transparent overview, the outcomes corresponding to each output variable are reported separately in the following table. All computational experiments were carried out in MATLAB version 2025a.

For model development, 300 randomly generated samples were used for training and validation, whereas the remaining 40 samples were reserved as an independent holdout test set. This holdout set was not used in any part of hyperparameter tuning, Bayesian optimization, model selection, or training control. It was employed only after final model fitting to evaluate generalization performance on unseen data. The random partition was generated with the random number generator fixed to state zero to ensure reproducibility.

In order to establish a fair and consistent comparison across the different algorithms, the hyperparameters of all models were tuned using a systematic optimization procedure. This approach ensured that no single algorithm gained an advantage due to hyperparameter settings, thereby allowing for an unbiased performance evaluation. The optimization process was limited to a maximum of 30 epochs, and the robustness of the results was further supported by employing the K-fold cross-validation technique (k-fold = 5). Additionally, to guarantee reproducibility and minimize the influence of stochastic variations, the Random Number Generator (RNG) was fixed to state zero during the entire process. Consequently, all reported results can be reproduced under identical conditions. This

systematic setup ensured that the comparative analysis reflects the intrinsic capabilities of the algorithms rather than artifacts arising from randomness or unequal parameterization. More specifically, for all surrogate models considered in this study, hyperparameter selection was carried out in MATLAB using Bayesian Optimization coupled with 5-fold cross-validation. The Bayesian optimization budget was limited to 30 evaluations, and the random number generator was fixed to RNG state 0 to ensure reproducibility. For the proposed PDRN, each candidate architecture evaluated during Bayesian optimization was trained for 500 epochs. After the optimal hyperparameters were identified, the final PDRN model was retrained on the 300-sample training/validation dataset for 1000 epochs, while the 40-sample holdout set remained excluded and was used only for the final independent performance evaluation.

As it can be observed from Table 7 and Figure 10, across outputs, the results point to a clear deployment strategy. Table 7 summarizes the predictive accuracy of the considered surrogate models using two complementary metrics: MAE, reported in the physical units of each output, and RMS error expressed in percent, which better reflects overall fidelity and is more sensitive to occasional large deviations. Considering both measures is important because a model may achieve a low average error while still producing localized outliers that increase RMS% and can undermine constraint satisfaction in subsequent design studies. For Mass, the models exhibit relatively similar behavior, consistent with an input–output mapping that is close to linear over the sampled domain. As a result, additional model capacity yields limited marginal improvement. GPR and the PDRN provide the smallest RMS errors (0.9% and 0.8%), indicating high reliability in reproducing mass. Although Mass is comparatively easier to learn, keeping its prediction error low remains important because it often directly affects objective values and couples into multiple constraints.

Table 7. Benchmark MAE/RMS [%] results of surrogate models.

Output	MAE/RMS Error [%]				
	ANN	**ENS**	**SVR**	**GPR**	**PDRN**
Mass [Kg]	32.4/1.00	44.3/1.4	37.4/1.1	30.9/0.9	29.8/0.8
1st Frequency [Hz]	0.21/23.1	0.39/44	0.41/46	0.06/6.9	0.02/3.1
Load Factor	0.20/28.1	0.43/51.6	0.15/17.4	0.12/14.6	0.08/11.5
Safety Factor	0.75/60.3	0.51/44.3	0.32/29.4	0.12/16.37	0.07/11.4

A clear performance separation appears for the 1st natural frequency, which reflects coupled stiffness–mass dynamics and stronger nonlinear interactions. The ANN, ensemble learning, and SVR show substantially higher RMS errors (23.1–46%), indicating reduced generalization and larger deviations in regions where the modal response varies more sharply. In contrast, GPR and the PDRN reduce the RMS error to 6.9% and 3.1%, respectively, demonstrating a markedly improved ability to capture the dynamic response, with the PDRN providing the best overall accuracy (MAE 0.02 Hz). Load Factor further highlights differences in robustness. While the baseline learners reduce MAE to varying degrees, their RMS errors remain relatively high (17.4–51.6%), pointing to non-uniform accuracy and the presence of high-error pockets across the design space. GPR improves this behavior (14.6% RMS), and the PDRN achieves the lowest error levels (11.5% RMS), indicating more stable predictions with fewer large deviations, which is particularly important for constraint-driven analyses.

Finally, Safety Factor is the most challenging output, as it is typically governed by peak-response behavior and compounded nonlinearities. This difficulty is reflected in the large RMS errors of the ANN and ensemble learning (60.3% and 44.3%) and the moderate performance of SVR (29.4%). GPR and the PDRN again deliver the strongest results, with

the PDRN achieving the lowest RMS error (11.4%) and MAE (0.07), supporting the conclusion that the proposed model provides the most consistent accuracy on the design-critical outputs. A plausible explanation for this performance gap is the representational structure of the PDRN itself. As described in Section 2.3.2, the model first maps the input variables into a higher-dimensional latent space and then progressively compresses this information through its pyramidal layer sequence. This hierarchical transformation allows the network to model complex interactions among the grouped thickness parameters and to distill the most informative combinations across successive layers. Such a mechanism is especially relevant for the 1st Frequency, which is governed by coupled stiffness–mass dynamics, and for the Safety Factor, which is affected by localized peak-response behavior and compounded nonlinearities. In addition, the use of LReLU supports stable learning of sharp nonlinear trends by maintaining non-zero gradients, while the Bayesian-optimized structural hyperparameters, together with normalization, dropout, and batch normalization, help the model adapt to the present dataset without excessive overfitting. For these reasons, the PDRN appears to provide a more suitable inductive structure for the more nonlinear and design-critical outputs than the baseline ANN and Ensemble Learning models.

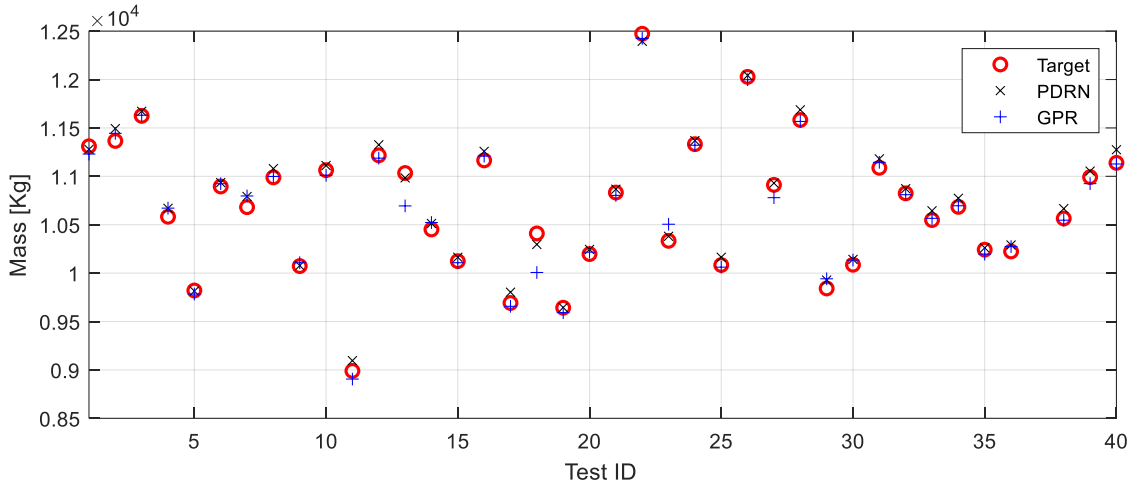

(**a**)

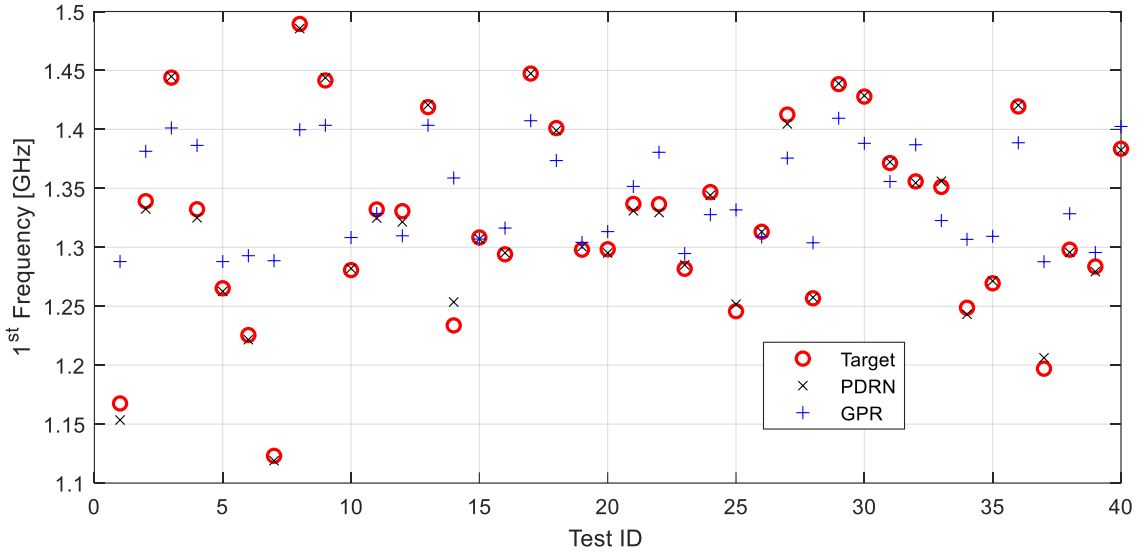

(**b**)

Figure 10. *Cont.*

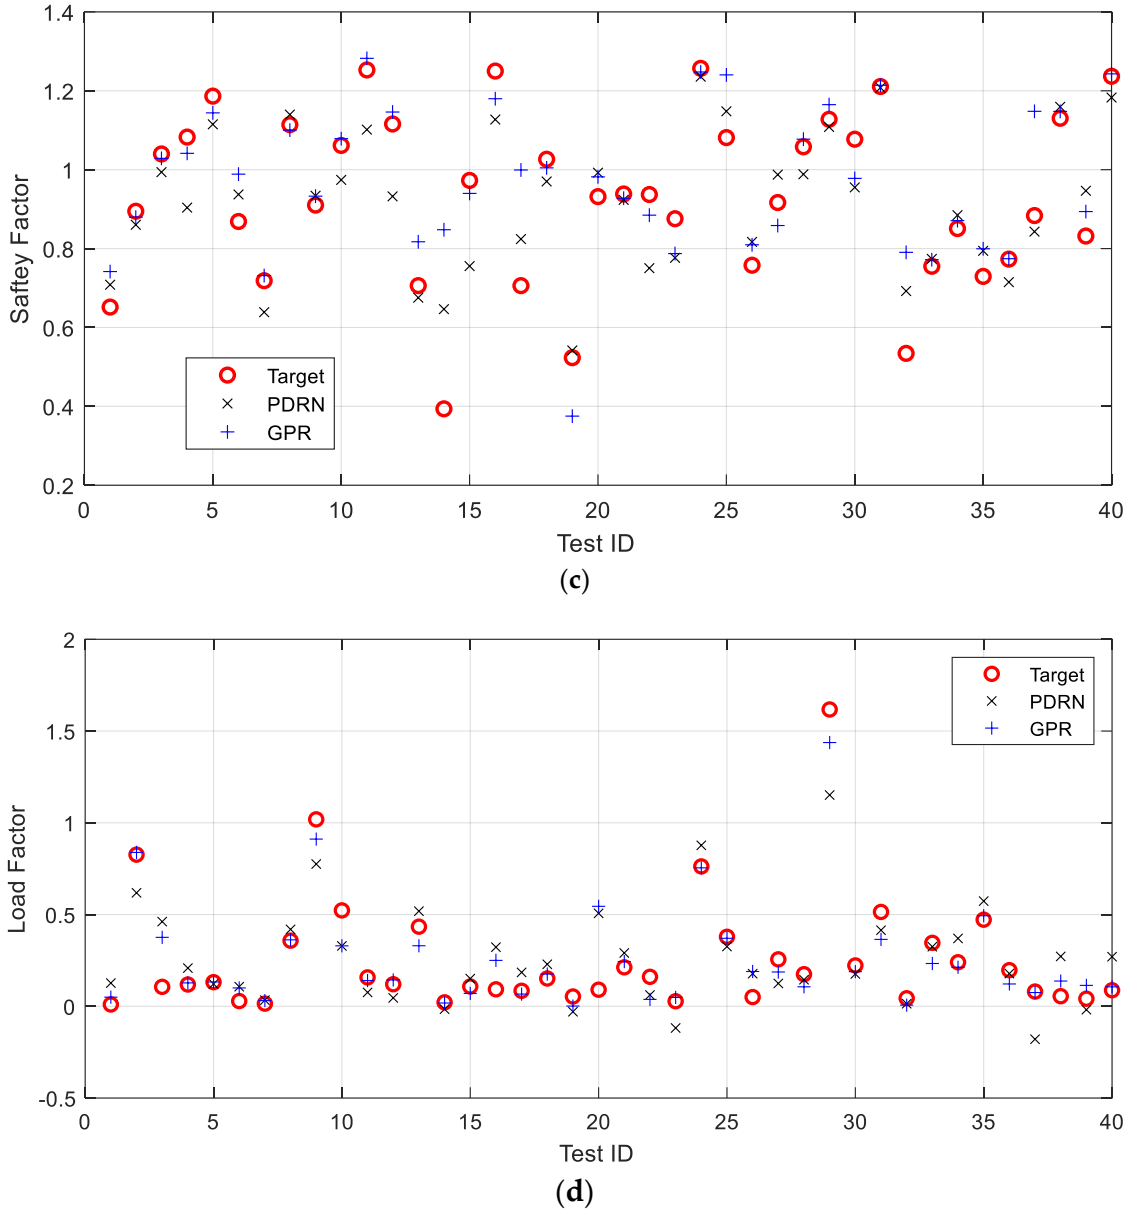

Figure 10. Benchmark results of surrogate models for test data corresponding to (**a**) Mass, (**b**) 1st Frequency, (**c**) Safety Factor, (**d**) Load Factor.

These results naturally motivate the next step of this study: assessing whether the observed predictive improvements translate into measurable benefits in a practical design task. Accordingly, the trained surrogates are integrated into an optimization framework for the wing model, where the aim is to minimize structural mass while satisfying minimum-performance requirements on safety factor, load factor, and the first natural frequency. In this formulation, x denotes the vector of wing design variables defined in Section 2 (with prescribed bounds), and the targets Mass T_{mass}, T_{sf} (targeted safety factor), T_{lf} (Targeted Load factor), and T_f (targeted first frequency) define feasibility thresholds for the respective responses.

To enable a direct and fair comparison with a conventional high-fidelity workflow, the optimization problem is posed using the penalty-based cost function shown in Equations (2)–(6). The objective term is the normalized mass, while large additive penalties are imposed whenever any of the constraints are violated. This construction ensures that feasible designs are primarily ranked by mass, whereas infeasible candidates are strongly

discouraged during the search. In the surrogate-assisted setting, the PDRN is queried at each iteration to rapidly predict the required responses and evaluate the cost, allowing the optimizer to explore the design space at negligible computational expense. The final surrogate optimum is then re-evaluated in ANSYS to verify constraint satisfaction and quantify surrogate-to-simulation discrepancies.

$$Cost(x) = C_0 + C_1 + C_2 + C_3 \tag{2}$$

$$C_0 = \begin{cases} 1000 & if \quad Mass(x) \geq T_{mass} \\ Mass(x) & if \quad SF(x) < T_{mass} \end{cases} \tag{3}$$

$$C_1 = \begin{cases} 1000 & if \quad SF(x) < T_{sf} \\ 0 & if \quad SF(x) \geq T_{sf} \end{cases} \tag{4}$$

$$C_2 = \begin{cases} 1000 & if \quad LF(x) < T_{Lf} \\ 0 & if \quad LF(x) \geq T_{Lf} \end{cases} \tag{5}$$

$$C_3 = \begin{cases} 1000 & if \quad f(x) < T_f \\ 0 & if \quad f(x) \geq T_f \end{cases} \tag{6}$$

where *Tmass*, *Tsf*, *Tlf*, and *Tf* are the target thresholds (*Tmass* = 10,500, *Tsf* = 1.15, *Tlf* = 1.15, *Tf* = 1.2 in this study). Here, *Tsf* denotes the minimum yield-based safety factor derived from the linear elastic static analysis, whereas *Tlf* denotes the minimum eigenvalue buckling load factor associated with structural stability.

In parallel, direct optimization is performed with ANSYS in the loop using the same design variables, bounds, cost definition, and stopping criteria. This paired experimental design isolates the impact of replacing expensive simulations with the PDRN during the iterative search and enables a rigorous comparison in terms of achieved high-fidelity mass, feasibility with respect to T_{mass}, T_{sf}, T_{lf}, and T_f, and total computational cost measured by the number of ANSYS evaluations and overall runtime.

Table 8 reports the final optimization outcomes obtained using the proposed PDRN-assisted optimization and compares them with a direct ANSYS in-the-loop optimization run performed under the same design variables, bounds, cost definition, and stopping criteria. This paired setup isolates the effect of replacing expensive high-fidelity simulations with the learned surrogate during the iterative search and enables a direct comparison in terms of achieved mass, feasibility with respect to the target constraints (minimum Mass, Tsf, Tlf, Tf), and total computational cost measured by the number of ANSYS evaluations (and runtime). Given that high-fidelity FEA is widely recognized as computationally prohibitive when embedded within optimization loops, reducing the number of solver calls is directly tied to practical reductions in turnaround time and design cost.

In the proposed workflow, the computational budget is concentrated into a single offline data generation stage, where the design space is sampled using LHS and each sample is evaluated once via ANSYS to produce the training/test dataset. Specifically, two datasets of 300 training and 40 holdout/test samples are used for training/validation (with k-fold) and for independent generalization checking, respectively. After this one-time cost, the optimization loop can evaluate candidate designs at negligible cost by querying the PDRN instead of repeatedly invoking high-fidelity simulations. This is precisely where the computational advantage emerges: the surrogate allows dense exploration and rapid convergence behavior without paying the runtime penalty of FEA at every iteration, while the final reported optimum is still verified using ANSYS to ensure high-fidelity feasibility and performance (Table 8).

Table 8. Optimal design variables obtained using the proposed PDRN-assisted optimization and their high-fidelity ANSYS verification, compared with the best solution returned by the direct ANSYS in-the-loop optimization (same design bounds, objective, and stopping criteria).

Case	Optimized [x]	1st Frequency [Hz]	Load Factor	Safety Factor	Mass [Kg]
PDRN + PSO	5.01, 5.54, 4.89, 12.13, 4.71, 5.00, 5.48, 24.9, 6.56, 14.99	1.4192	1.247	1.182	10,494
ANSYS Validation		1.4142	1.307	1.158	10,485
ANSYS Direct Opt. @ 290 evaluation	9.99, 5.55, 3.22, 11.53, 4.83, 11.02, 8.31, 24.49, 9.52, 22.94	1.5052	1.309	1.313	11,220
ANSYS Direct Opt. @ 340 evaluation	7.06, 5.88, 3.35, 13.85, 4.94, 8.61, 6.93, 22.91, 7.31, 19.29	1.4992	1.461	1.166	10,976
ANSYS Direct Opt. @ 440 evaluation	8.33, 5.68, 3.15, 11.53, 4.67, 5.32, 5.99, 24.49, 7.70, 16.90	1.4408	1.234	1.207	10,576
ANSYS Direct Opt. @ 540 evaluation	7.16, 3.72, 3.34, 10.54, 4.63, 8.95, 6.93, 24.50, 6.84, 16.83	1.4214	1.223	1.189	10,490
ANSYS Direct Opt. @ 680 evaluation	7.21, 3.52, 3.16, 10.51, 4.67, 5.26, 5.49, 24.4, 6.81, 16.8	1.4098	1.152	1.183	10,354

Table 8 reports the final optimization outcomes obtained using the proposed PDRN-assisted optimization and compares them with direct ANSYS in-the-loop optimization under the same design variables, bounds, objective formulation, and stopping criteria. Because the two workflows allocate computational effort differently, the comparison is presented in terms of both high-fidelity solver usage and runtime breakdown. In the proposed approach, most of the expensive cost is incurred once during the offline generation of the surrogate dataset, after which the optimization proceeds using the trained surrogate at negligible query cost and the final optimum is re-evaluated once in ANSYS for verification. By contrast, in direct optimization, the high-fidelity solver is invoked throughout the iterative search. Therefore, the proposed method should be interpreted as a redistribution of computational cost from online optimization to offline surrogate construction, rather than as a direct one-to-one replacement of the direct optimization workflow. For the direct ANSYS in-the-loop optimization, the comparison was reported at predefined high-fidelity evaluation checkpoints of 290, 340, 440, 540, and 680 ANSYS evaluations under the same design-variable bounds, objective function, and constraint definitions. These checkpoints were used to provide a transparent and directly traceable comparison against the surrogate-assisted workflow in terms of both feasible solution discovery and computational effort.

As seen in Table 8, the direct ANSYS optimization did not satisfy the target mass threshold at 290, 340, or 440 evaluations, whereas the surrogate-guided design was verified in ANSYS as feasible with a mass of 10,485 kg. The first direct ANSYS run that reached a comparable feasible solution occurred at 540 evaluations, yielding a mass of 10,490 kg. Therefore, when measured in terms of high-fidelity solver usage, the proposed workflow

required 341 ANSYS analyses (340 for offline dataset generation and 1 for final validation), compared with 540 analyses for the first comparable feasible direct solution, corresponding to approximately 1.58× fewer high-fidelity evaluations. This comparison should be interpreted in terms of solver usage rather than wall clock time alone; therefore, the runtime breakdown is additionally reported in Table 9. It should be noted that the computational cost of the PSO search on the trained surrogate is negligible compared with the cost of high-fidelity ANSYS analyses, because surrogate inference is effectively instantaneous and the optimization is performed entirely on the trained model. Therefore, the dominant cost of the proposed workflow is associated with offline dataset generation and final high-fidelity verification.

Table 9. Runtime breakdown of the surrogate-assisted and direct ANSYS optimization workflows under the timing assumptions adopted in this study. All high-fidelity finite element simulations were performed in ANSYS 2025 under Windows 10 on the same workstation equipped with an AMD Ryzen 7 3700X (8-core) processor and 32 GB RAM. The reported runtimes represent approximate wall clock estimates based on the average solution time per ANSYS analysis and include the offline dataset generation cost, surrogate training cost, PSO-based optimization cost, and final ANSYS validation for the PDRN-assisted workflow.

Case	Case Detail	Total Run Time [Hours]
PDRN + PSO	(300 + 40) × [30 min] + Training Time with HP optimization [60 min] + Surrogate Assisted optimization with PSO 10K evaluation [less than a minute = 0.001 s × 10K] + ANSYS Validation [30 min]	171.5
ANSYS Direct Opt. @ 290 evaluation	(290) × [30 min]	145
ANSYS Direct Opt. @ 340 evaluation	(340) × [30 min]	170
ANSYS Direct Opt. @ 440 evaluation	(440) × [30 min]	220
ANSYS Direct Opt. @ 540 evaluation	(540) × [30 min]	270
ANSYS Direct Opt. @ 680 evaluation	(680) × [30 min]	340

As reported in Table 9, the total runtime of the proposed surrogate-assisted workflow is approximately 171.5 h, of which nearly 170 h correspond to the offline generation of the 340 high-fidelity ANSYS analyses. By contrast, the total runtime of the direct ANSYS optimization reaches approximately 145, 170, 220, 270, and 340 h at 290, 340, 440, 540, and 680 evaluations, respectively. These results confirm that the wall clock cost of the proposed framework is dominated by the offline dataset generation stage, whereas the additional costs associated with surrogate training and PSO-based search remain negligible compared with repeated high-fidelity solver calls. When compared with the first direct ANSYS solution that achieved a comparable feasible design at 540 evaluations, the surrogate-assisted workflow also provides a lower total runtime (171.5 h versus 270 h), corresponding to an overall runtime reduction of approximately 36.5% under the timing assumptions adopted in Table 9.

Table 8 summarizes the optimal design vector x obtained by the proposed PDRN + PSO framework and compares it against the direct ANSYS in-the-loop optimization at multiple function-evaluation budgets. The optimization is guided by the revised penalty-based objective in (2)–(6), where feasibility is enforced through hard thresholding: solutions are penalized if the mass exceeds T_{mass} = 10,500 kg, or if any of the performance constraints violate T_{sf} = 1.15, T_{lf} = 1.15, T_f = 1.2. Consequently, the cost function promotes designs that are simultaneously lightweight and constraint-satisfying, while strongly discouraging candidates that fail to meet any target.

As seen in Table 8, the PDRN + PSO approach yields a design that satisfies all thresholds with Mass = 10,494 kg, SF = 1.182, LF = 1.247, and f_1 = 1.4192 Hz. Importantly, this optimum is not reported solely from the surrogate prediction: the same design vector is re-evaluated using high-fidelity ANSYS, producing Mass = 10,485 kg, SF = 1.158, LF = 1.307, and f_1 = 1.4142 Hz, which confirms that the surrogate-guided optimum remains feasible under the high-fidelity solver and that the predicted trends are consistent. This validation step is crucial because it demonstrates that the surrogate is accurate enough not only for prediction, but also for decision-making inside an optimization loop.

The comparison with the direct ANSYS optimization highlights the practical computational advantage of the proposed workflow. At 290, 340, and 440 evaluations, the direct solver-driven optimizer returns designs with mass values exceeding 10,500 kg (11,220, 10,976, and 10,576 kg, respectively), meaning that the solutions remain penalized and are therefore not acceptable with respect to the revised objective. Only after a substantially larger computational budget does the direct ANSYS approach reach the feasible region: at 540 evaluations, it attains Mass = 10,490 kg (with SF/LF/f_1 also above the required thresholds), which is comparable to the feasibility and mass level reached by PDRN + PSO. The key point is that the proposed method achieves this feasibility level after generating only 340 high-fidelity samples (300 training + 40 test) to build the surrogate, whereas the direct approach requires ~540 solver calls to obtain a similarly feasible solution—corresponding to roughly a 40% reduction in high-fidelity evaluations for reaching an acceptable design. In practical terms, fewer ANSYS runs directly translate to reduced compute-hours, shorter queue/turnaround time, and lower software/HPC costs, enabling more iterations of design exploration within the same budget.

Finally, Table 8 also illustrates the expected trade-off between surrogate-assisted and fully simulation-driven optimization. With a sufficiently large number of function evaluations, the direct ANSYS optimizer can continue improving the design and, at 680 evaluations, finds a lower-mass solution (10,354 kg) while still satisfying the constraints—at the expense of nearly double the high-fidelity evaluation count compared to the surrogate training cost. This behavior is consistent with the fact that a purely high-fidelity optimizer may more reliably approach deeper minima as the evaluation budget grows. However, the major advantage of the surrogate-assisted strategy is amortization: once the PDRN is trained, it can be immediately reused to solve additional optimization variants (e.g., different weights, modified targets, alternative metaheuristics, or multi-start studies) with negligible incremental cost, requiring only occasional ANSYS re-validation (or targeted new samples if the design space is expanded). Thus, the proposed PDRN-assisted framework offers a strong time- and cost-effective pathway to obtain high-quality feasible designs rapidly, while preserving the option of running expensive solver-based refinement only when marginal gains justify the additional computational expense.

4. Discussion

This study developed and validated a Pyramidal Deep Regression Network (PDRN)-based surrogate modeling framework for the structural analysis and optimization of a

metallic commercial aircraft wing. Using high-fidelity ANSYS simulations of an Airbus A320-like wing model, a dataset was generated through Latin Hypercube Sampling and used to train and test multiple surrogate models. The comparative evaluation showed that the PDRN delivered the best overall predictive performance among the considered methods, outperforming ANN, Ensemble Learning, SVR, and GPR, particularly for the design-critical responses of first natural frequency, load factor, and safety factor. These results indicate that the adapted PDRN architecture is highly effective in capturing the nonlinear relationships between grouped wing thickness parameters and structural responses in this structural surrogate modeling problem.

Beyond predictive accuracy, the results also demonstrate the practical value of the proposed surrogate in structural optimization. When coupled with a PSO-based design framework, the trained PDRN enabled rapid exploration of the design space and yielded a feasible low-mass solution that was subsequently verified through high-fidelity ANSYS analysis. In this respect, the proposed surrogate-assisted workflow reached a comparable feasible design with fewer high-fidelity solver evaluations than direct ANSYS in-the-loop optimization. In particular, the first comparable feasible direct ANSYS solution required 540 high-fidelity evaluations, whereas the proposed workflow required 341 evaluations in total (340 for offline dataset generation and 1 for final validation), corresponding to approximately 1.58× fewer high-fidelity analyses. Although direct high-fidelity optimization may ultimately converge to lighter solutions when significantly larger evaluation budgets are permitted, the primary advantage of the present approach lies in its computational efficiency, flexibility, and reusability. Once trained, the surrogate model can be employed for repeated optimization studies under modified objectives, constraints, or design scenarios with minimal additional computational cost. At the same time, the scope of the present validation should be interpreted carefully. The framework was assessed only for a metallic Airbus A320-like wing model and only within a linear elastic static/modal/eigenvalue buckling analysis regime, which was intentionally selected as a representative and computationally manageable benchmark for demonstrating the core methodology. Therefore, while the results confirm the effectiveness of the proposed PDRN framework for this class of problem, they should not be interpreted as evidence that the present surrogate directly replaces geometrically nonlinear FEA, aeroelastic simulations, or post-buckling analyses. Further studies are required to establish its broader generalizability in such settings. In principle, the methodology can be extended to more advanced aerospace structures, including composite wings, high-aspect-ratio configurations, and morphing systems. However, such applications would require substantially more sophisticated finite element models involving anisotropic material behavior, nonlinear kinematics, aeroelastic coupling, and more complex failure criteria. The present mesh configuration and the use of linear-order elements should likewise be interpreted as a computationally manageable benchmark discretization for repeated large-scale analyses rather than as a formally convergence-certified setup for local stress or stiffness-sensitive buckling verification. In addition, the use of MPC-based bonded contacts in the present model should be interpreted as a conceptual simplification suitable for global structural response prediction; detailed riveted/bolted joint modeling with nonlinear contact would be required for more accurate local stress assessment at structural interfaces. Incorporating these factors was beyond the scope of the present work, as doing so would have obscured the clear validation of the surrogate modeling strategy itself.

Another important direction for future research is the integration of the proposed framework with multi-fidelity and physics-informed learning strategies. Combining low- and high-fidelity simulation data could improve sample efficiency and reduce the cost of data generation, while embedding physical constraints into the learning process could

enhance robustness, consistency, and extrapolation capability. Such developments would help bridge the gap between purely data-driven surrogates and physics-based structural modeling. These directions will be considered in future work to improve the generality, data efficiency, and physical robustness of the proposed framework.

5. Conclusions

This study presented a Pyramidal Deep Regression Network (PDRN)-based surrogate modeling framework for the structural analysis and optimization of a metallic Airbus A320-like wing. Using 340 high-fidelity ANSYS simulations generated through Latin Hypercube Sampling, the proposed framework was trained and evaluated for the prediction of structural mass, first natural frequency, buckling load factor, and minimum safety factor. Among the compared models, the PDRN achieved the best overall predictive performance on the unseen holdout set, demonstrating its effectiveness as a deep learning-based surrogate for this class of structural problems.

When integrated into a PSO-based optimization framework, the trained PDRN enabled rapid design space exploration and produced a feasible low-mass design that was subsequently verified by high-fidelity ANSYS analysis. In addition, the comparison in Table 8 showed that the proposed surrogate-assisted workflow required 341 high-fidelity ANSYS analyses to reach a comparable feasible solution, whereas the first comparable direct ANSYS optimization result required 540 analyses, corresponding to approximately 1.58× fewer high-fidelity evaluations. This result highlights the computational efficiency of the proposed framework for preliminary structural design studies, while the verified ANSYS re-evaluation confirms the practical reliability of the surrogate-guided optimum.

Overall, the present work shows that the adapted PDRN framework is a promising and computationally efficient surrogate modeling tool for constraint-driven wing structural optimization under limited data. Future work will extend the framework toward more advanced structural configurations and toward multi-fidelity and physics-informed learning strategies to improve generality, data efficiency, and physical robustness.

Future work will consider replacing the present conceptual loading definition with more realistic aerodynamic loads obtained from vortex lattice-based tools, such as AVL and Tornado, and mapped to the structural model through established load-transfer procedures such as thin-plate spline-based interpolation [48–51]. The inclusion of multiple maneuver-dependent load cases, together with more refined and convergence-verified discretizations, is also expected to improve the aerodynamic–structural consistency and overall fidelity of the proposed benchmark framework. In addition, future work will investigate modern adaptive differential evolution-based metaheuristics, such as SHADE, L-SHADE, LSHADE-EpSin, and jSO, not only as additional benchmark optimizers within the surrogate-assisted design framework, but also as candidate tools for the hyperparameter optimization of PDRN-type surrogate models [52–56]. In the present study, the computational cost of the optimization stage is already negligible once the surrogate is trained; therefore, the principal benefit of these advanced optimizers is expected to emerge more strongly in improving surrogate model calibration, regression accuracy, and training efficiency rather than in altering the main structural optimization conclusions of the current benchmark problem.

Author Contributions: Conceptualization, H.K. and M.O.K.; methodology, H.K.; software, H.K. and P.M.; validation, H.K.; formal analysis, H.K. and P.M.; investigation, H.K.; resources, H.K.; data curation, H.K.; writing—original draft preparation, H.K.; writing—review and editing, M.O.K.; visualization, H.K.; supervision, M.O.K.; project administration, H.K.; funding acquisition, P.M. All authors have read and agreed to the published version of the manuscript.

Funding: The authors thank scientific research project coordination of the Yildiz Technical University for founding our research under project number FBA-2025-6780.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: The authors would like to thank Yıldız Technical University's Smart RF and Antenna Systems Research Laboratory for providing the computational resources and software licenses necessary for this research. Authors used OpenAI's ChatGPT 5.3 image generation tool to create an initial schematic illustration of the proposed framework shown in Figure 9a.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

FEA	Finite Element Analysis
PDRN	Pyramidal Deep Regression Network
LHS	Latin Hypercube Sampling
ANN	Artificial Neural Network
SVR	Support Vector Regression
ML	Machine Learning
GPR	Gaussian Process Regression
DNN	Deep Neural Network
MMPDS	Metallic Materials Properties Development and Standardization
BC	Boundary Condition
MPC	Multipoint Constraint
LReLU	Leaky Rectified Linear Unit
BO	Bayesian Optimization

References

1. Kumar, B.R. Investigation on buckling response of the aircraft's wing using finite-element method. *Aust. J. Mech. Eng.* **2020**, *18*, S122–S131. [CrossRef]
2. Chaves-Vargas, M.; Dafnis, A.D.; Reimerdes, H.G.; Schröder, K.U. Modal parameter identification of a compression-loaded CFRP stiffened plate and correlation with its buckling behaviour. *Prog. Aerosp. Sci.* **2015**, *78*, 39–49. [CrossRef]
3. Bhutta, A.H. Optimizing Structural Integrity of Fighter Aircraft Wing Stations: A Finite Element Analysis Approach. *Ingenius* **2024**, *32*, 90–100. [CrossRef]
4. Chiariello, A.; Perillo, G.; Linari, M.; Russo, R.; Orlando, S.; Vitale, P.; Belardo, M. Virtual Full Scale Static Test of a Civil Tilt Rotor Composite Wing in Non-Linear Regime. *Aerospace* **2024**, *11*, 278. [CrossRef]
5. Liang, K.Z.; Yin, Z. Investigation on nonlinear buckling performance of the optimized wing structure under the realistic flight cases. *Aerosp. Sci. Technol.* **2023**, *139*, 108416. [CrossRef]
6. Houston, G.; Quinn, D.; Murphy, A.; Bron, F. Wing Panel Design with Novel Skin-Buckling Containment Features. *J. Aircr.* **2016**, *53*, 416–426. [CrossRef]
7. Gray, A.C.; Martins, J.R.R.A. Geometrically Nonlinear High-fidelity Aerostructural Optimization for Highly Flexible Wings. In Proceedings of the AIAA Scitech 2021 Forum, Virtual, 11–15 and 19–21 January 2021. [CrossRef]
8. Shizuno, Y.; Dong, S.; Kuzuno, R.; Okada, T.; Kawashima, S.; Makihara, K.; Otsuka, K. Nonlinear Dynamic Analysis Framework for Slender Structures Using the Modal Rotation Method. *J. Comput. Nonlinear Dyn.* **2025**, *20*, 021002. [CrossRef]
9. Stanford, B.K.; Beran, P.S. Cost reduction techniques for the design of non-linear flapping wing structures. *Int. J. Numer. Methods Eng.* **2011**, *88*, 533–555. [CrossRef]
10. Farah, S.; Hrairi, M.; Abdullah, A.B.; Brahmi, A. A design optimization strategy of an aircraft composite wing-box based on a multi-scale FEA. *FME Trans.* **2025**, *53*, 157–172. [CrossRef]
11. Long, T.; Wu, Y.; Wang, Z.; Tang, Y.; Wu, D.; Yu, Y. Efficient aero-structure coupled wing optimization using decomposition and adaptive metamodeling techniques. *Aerosp. Sci. Technol.* **2019**, *95*, 105496. [CrossRef]
12. Benaouali, A.; Kachel, S. Multidisciplinary design optimization of aircraft wing using commercial software integration. *Aerosp. Sci. Technol.* **2019**, *92*, 391–399. [CrossRef]

13. Phuekpan, K.; Khammee, R.; Panagant, N.; Bureerat, S.; Pholdee, N.; Wansasueb, K. A Comparison of Modern Metaheuristics for Multi-Objective Optimization of Transonic Aeroelasticity in a Tow-Steered Composite Wing. *Aerospace* **2025**, *12*, 101. [CrossRef]
14. Shrivastava, S.; Tilala, H.; Mohite, P.M.; Limaye, M. Weight optimization of a composite wing-panel with flutter stability constraints by ply-drop. *Struct. Multidiscip. Optim.* **2020**, *62*, 2181–2195. [CrossRef]
15. Yondo, R.; Valero, E. A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses. *Prog. Aerosp. Sci.* **2018**, *96*, 23–61. [CrossRef]
16. Zhang, X.; Xie, F.; Ji, T.; Zhu, Z.; Zheng, Y. Multi-fidelity deep neural network surrogate model for aerodynamic shape optimization. *Comput. Methods Appl. Mech. Eng.* **2021**, *373*, 113485. [CrossRef]
17. Tao, J.; Sun, G. Application of deep learning based multi-fidelity surrogate model to robust aerodynamic design optimization. *Aerosp. Sci. Technol.* **2019**, *92*, 464–476. [CrossRef]
18. Wang, H.; Jin, Y.; Sun, C.; Doherty, J. Offline Data-Driven Evolutionary Optimization Using Selective Surrogate Ensembles. *IEEE Trans. Evol. Comput.* **2019**, *23*, 203–216. [CrossRef]
19. Teng, D.; Feng, Y.; Chen, J.; Lu, C. Intelligent vectorial surrogate modeling framework for multi-objective reliability estimation of aerospace engineering structural systems. *Chin. J. Aeronaut.* **2024**, *37*, 337–352. [CrossRef]
20. Mainini, L.; Willcox, K. Surrogate Modeling Approach to Support Real-Time Structural Assessment and Decision Making. *AIAA J.* **2015**, *53*, 1612–1626. [CrossRef]
21. Du, B.; Shen, E.; Wu, J.; Guo, T.; Lu, Z.; Zhou, D. Aerodynamic Prediction and Design Optimization Using Multi-Fidelity Deep Neural Network. *Aerospace* **2025**, *12*, 292. [CrossRef]
22. Wu, M.; He, X.; Sun, X.; Tong, T.; Chen, Z.; Zheng, C. Efficient aerodynamic shape optimization using transfer learning based multi-fidelity deep neural network. *Phys. Fluids* **2024**, *36*, 075126. [CrossRef]
23. Du, X.; Martins, J.R.R.A.; O'Leary-Roseberry, T.; Chaudhuri, A.; Ghattas, O.; Willcox, K. Learning Optimal Aerodynamic Designs through Multi-Fidelity Reduced-Dimensional Neural Networks. In Proceedings of the AIAA SCITECH 2023 Forum, National Harbor, MD, USA, 23–27 January 2023. [CrossRef]
24. Wu, P.; Liu, Z.; Zhou, Z.; Song, C. Multi-fidelity deep neural network surrogate model for aerodynamic shape prediction based on multi-task learning. In Proceedings of the 2024 3rd International Conference on Energy and Power Engineering, Control Engineering (EPECE), Changchun, China, 12–14 July 2024; pp. 137–142. [CrossRef]
25. Aabid, A.; Zakuan, M.A.M.B.M.; Khan, S.A.; Ibrahim, Y.E. Structural analysis of three-dimensional wings using finite element method. *Aerosp. Syst.* **2022**, *5*, 47–63. [CrossRef]
26. Ding, L.; Zhang, X. Sample and computationally efficient stochastic kriging in high dimensions. *Oper. Res.* **2024**, *72*, 660–683. [CrossRef]
27. United States. Department of Defense. *Military Handbook: Metallic Materials and Elements for Aerospace Vehicle Structures*; US Department of Defense: Washington, DC, USA, 1990; Volume 1.
28. Airbus. A320 Aircraft Characteristics: Airport and Maintenance Planning. 2025. Available online: https://www.aircraft.airbus.com/sites/g/files/jlcbta126/files/2025-01/AC_A320_0624.pdf (accessed on 30 March 2026).
29. Schrenk, O. A simple approximation method for obtaining the spanwise lift distribution. *Aeronaut. J.* **1941**, *45*, 331–336. [CrossRef]
30. Li, H.; Sen, S.; Khazanovich, L. A scalable adaptive sampling approach for surrogate modeling of rigid pavements using machine learning. *Results Eng.* **2024**, *23*, 102483. [CrossRef]
31. Garg, M.; Chang, T.H.; Raghavan, K. SF-SFD: Stochastic Optimization of Fourier Coefficients to Generate Space-Filling Designs. In Proceedings of the 2023 Winter Simulation Conference (WSC), San Antonio, TX, USA, 10–13 December 2023; pp. 3636–3646. [CrossRef]
32. Fan, L.; Wang, Z.; Liao, C.; Wang, J. Efficient optimization of plasma surface high harmonic generation by an improved Bayesian strategy. *arXiv* **2024**, arXiv:2410.24051. [CrossRef]
33. Borisut, P.; Nuchitprasittichai, A. Adaptive Latin Hypercube Sampling for a Surrogate-Based Optimization with Artificial Neural Network. *Processes* **2023**, *11*, 3232. [CrossRef]
34. Damblin, G.; Couplet, M.; Iooss, B. Numerical studies of space-filling designs: Optimization of Latin Hypercube Samples and subprojection properties. *J. Simul.* **2013**, *7*, 276–289. [CrossRef]
35. Cioppa, T.M.; Lucas, T.W. Efficient Nearly Orthogonal and Space-Filling Latin Hypercubes. *Technometrics* **2007**, *49*, 45–55. [CrossRef]
36. Afzal, A.; Kim, K.Y.; Seo, J.H. Effects of Latin hypercube sampling on surrogate modeling and optimization. *Int. J. Fluid Mach. Syst.* **2017**, *10*, 240–253. [CrossRef]
37. Wang, Z.; Hong, Y.; Huang, L.; Zheng, M.; Yuan, H.; Zeng, R. A comprehensive review and future research directions of ensemble learning models for predicting building energy consumption. *Energy Build.* **2025**, *335*, 115589. [CrossRef]
38. Zhang, J.; Wang, Z.; Lu, Q.; Zhu, S.; Bai, Z. Physical information-guided ensemble learning model for surface roughness prediction of galvannealed strip. *Eng. Appl. Artif. Intell.* **2026**, *165*, 113411. [CrossRef]

39. Tanveer, M.; Rajani, T.; Rastogi, R.; Shao, Y.H.; Ganaie, M.A. Comprehensive review on twin support vector machines. *Ann. Oper. Res.* **2024**, *339*, 1223–1268. [CrossRef]
40. Roy, A.; Chakraborty, S. Support vector machine in structural reliability analysis: A review. *Reliab. Eng. Syst. Saf.* **2023**, *233*, 109126. [CrossRef]
41. Virdee, T.S.; Kottegoda, N.T. A brief review of kriging and its application to optimal interpolation and observation well selection. *Hydrol. Sci. J.* **1984**, *29*, 367–387. [CrossRef]
42. Giraldo, R.; Leiva, V.; Castro, C. An overview of kriging and cokriging predictors for functional random fields. *Mathematics* **2023**, *11*, 3425. [CrossRef]
43. Amini, S.; Van Nieuwenhuyse, I. A tutorial on kriging-based stochastic simulation optimization. *arXiv* **2025**, arXiv:2502.05216.
44. Abiodun, O.I.; Jantan, A.; Omolara, A.E.; Dada, K.V.; Mohamed, N.A.; Arshad, H. State-of-the-art in artificial neural network applications: A survey. *Heliyon* **2018**, *4*, e00938. [CrossRef]
45. Bacharidis, E.; Seventekidis, P.; Arailopoulos, A. Study on the Generalization of a Data-Driven Methodology for Damage Detection in an Aircraft Wing Using Reduced FE Models. *Appl. Mech.* **2026**, *7*, 9. [CrossRef]
46. Qian, G.; Zhu, J.; Zhu, I.; Shen, P. Optimization of Energy Harvesting Performance of Fully Passive Parallel Flapping Wings Based on the Combination of Taguchi Method and Neural Networks. *Aerosp. Sci. Technol.* **2026**, *174*, 111919. [CrossRef]
47. Koziel, S.; Çalık, N.; Mahouti, P.; Belen, M.A. Accurate Modeling of Antenna Structures by Means of Domain Confinement and Pyramidal Deep Neural Networks. *IEEE Trans. Antennas Propag.* **2022**, *70*, 2174–2188. [CrossRef]
48. Miranda, L.R.; Elliott, R.D.; Baker, W.M. *A Generalized Vortex Lattice Method for Subsonic and Supersonic Flow Applications*; NASA Contractor Report CR-2865; NASA: Washington, DC, USA, 1977.
49. Drela, M.; Youngren, H. *AVL User Primer*; Massachusetts Institute of Technology: Cambridge, MA, USA, 2010. Available online: https://web.mit.edu/drela/Public/web/avl/AVL_User_Primer.pdf (accessed on 30 March 2026).
50. Melin, T. A Vortex Lattice MATLAB Implementation for Linear Aerodynamic Wing Applications. Master's Thesis, Royal Institute of Technology (KTH), Stockholm, Sweden, 2000.
51. Samareh, J.A. Discrete Data Transfer Technique for Fluid–Structure Interaction. In Proceedings of the 18th AIAA Computational Fluid Dynamics Conference, Miami, FL, USA, 25–28 June 2007; AIAA Paper 2007-4309.
52. Tanabe, R.; Fukunaga, A. Success-History Based Parameter Adaptation for Differential Evolution. In Proceedings of the IEEE Congress on Evolutionary Computation (CEC), Cancun, Mexico, 20–23 June 2013; pp. 71–78.
53. Tanabe, R.; Fukunaga, A. Improving the Search Performance of SHADE Using Linear Population Size Reduction. In Proceedings of the IEEE Congress on Evolutionary Computation (CEC), Beijing, China, 6–11 July 2014; pp. 1658–1665. [CrossRef]
54. Awad, N.H.; Ali, M.Z.; Suganthan, P.N.; Reynolds, R.G. An Ensemble Sinusoidal Parameter Adaptation Incorporated with L-SHADE for Solving CEC2014 Benchmark Problems. In Proceedings of the IEEE Congress on Evolutionary Computation (CEC), Vancouver, BC, Canada, 24–29 July 2016; pp. 2958–2965. [CrossRef]
55. Brest, J.; Maučec, M.S.; Bošković, B. Single Objective Real-Parameter Optimization: Algorithm jSO. In Proceedings of the IEEE Congress on Evolutionary Computation (CEC), Donostia-San Sebastián, Spain, 5–8 June 2017; pp. 1311–1318. [CrossRef]
56. Suganthan, P.N. CODES: Codes for Some of My Co-Authored Journal/Conference Papers; GitHub Repository. Available online: https://github.com/P-N-Suganthan/CODES (accessed on 30 March 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.