

AI-assisted design of holistically sustainable aerostructures

Angelos Filippatos[1*], Thanasis Kyriazis[1], Spiros Pantelakis[1], Dionysios Markatos [1]

[1]Department of Mechanical Engineering and Aeronautics, University of Patras, 26504 Patras, Greece;
*angelos.filippatos@upatras.gr

Abstract

This contribution presents an AI-assisted design framework for sustainability-driven aero-structural optimization using data-driven design space modeling. A feedforward neural network is trained as a surrogate model to map geometric and material design parameters to structural performance indicators, sustainability-related metrics, and a composite index representing overall sustainability performance. In parallel, an inverse neural network is developed to generate feasible design configurations that satisfy prescribed sustainability requirements, including target sustainability index (SI) levels. Five representative design cases are formulated based on sustainability-oriented and structural performance requirements commonly encountered in the aerospace industry. The capability of the coupled forward–inverse modeling framework to generate valid design suggestions is then evaluated. The proposed framework enables rapid exploration of design alternatives and improves the interpretation of sustainability trade-offs in early-stage aircraft structural design.

Keywords: AI-assisted design, Machine Learning, Holistic Sustainability, Design for Sustainability, Aerostructures Design

1. Introduction

The global aviation sector is entering a decisive decade, driven by ambitious climate targets, accelerating technological change, and increasing pressure to ensure economic and industrial resilience. In Europe, this transition is strongly shaped by long-term decarbonization objectives, including climate neutrality by 2050, and by coordinated research frameworks that promote disruptive innovation in aircraft technologies [1]. Aircraft structures represent a critical leverage point in this transition, as structural design decisions taken at the conceptual and preliminary stages largely determine mass, life-cycle environmental impact, manufacturing cost, and circularity potential.

Beyond performance and safety, emphasis should be placed on economic viability, environmental impact and resource efficiency [2]. This design strategy is promoted by policy initiatives such as the European Green Deal and the Circular Economy Action Plan, which require manufacturers to reduce emissions and environmental impact of the aerostructure both during its usage and throughout design production and end-of life phase [3]. Those goals lead the engineers to rethink how materials are selected, having lightweight and circularity as the main pillars for the achievement of these goals. For

example, replacing conventional steels with aluminum can reduce component mass by 35–45 %, yielding up to 39 % lower life-cycle CO_2 emissions [4], [5]. Moreover, the recovery of metallic and polymeric components, without compromising the strength and durability, can reduce CO_2 emissions by 40% [6]. Therefore, focusing on lightweight designs by employing recyclable and advanced materials are the key points towards sustainable and eco-friendly design configurations.

Conventional aerostructures design methodologies remain predominantly performance-driven, focusing on strength, stiffness, and certification compliance. Sustainability-related aspects—such as CO_2 footprint, material criticality, recyclability, and life-cycle cost—are often addressed sequentially or at later design stages, when design flexibility is limited. This sequential approach hampers the ability to identify truly optimal trade-offs and leads to sub-optimal solutions from a system-level sustainability perspective. Recent European research efforts advocate a paradigm shift from eco-efficiency towards holistic sustainability-driven design, where structural performance, environmental impact, economic viability, and circularity are addressed simultaneously [7], [8] (Figure 1). However, embedding such holistic perspective significantly increases design complexity, especially for multi-material aerostructures, where geometry, material selection, manufacturing routes, and end-of-life strategies are strongly coupled. Traditional simulation-based optimization techniques, while accurate, are computationally expensive and ill-suited for rapid design space exploration in early design phases. Artificial Intelligence (AI), and in particular Machine Learning (ML), provides a powerful enabler to address this challenge.

Machine Learning (ML) enable systems to continuously learn by analyzing large volumes of data identifying relationships between inputs and outputs. In the engineering sector, ML has been used to optimize performance metrics such as stiffness, stress distribution, and crashworthiness of components [9], [10], [11], but those approaches focus more on the performance aspect of the components and not on the sustainability and economic considerations. Supervised models—including Artificial Neural Networks (ANNs), Random Forests (RF), and Support Vector Machines (SVMs)—have successfully predicted material properties, optimized geometries, and approximated Finite Element Analysis outputs [12], [13], [14], although approaches that emphasize simultaneously on structural performance, CO_2 emissions, life-cycle costs, and circularity for aviation components remain limited.

This paper presents an AI-assisted framework for the holistic sustainable design of aerostructures, aligned with European research priorities and addressing current challenges in sustainable aviation design and multidisciplinary optimization.



Figure 1 - Core pillars of aviation structures sustainability.

2. Methodology

2.1 Holistic Sustainability Framework

Sustainability is introduced as a core design driver through a composite Sustainability Index (SI) (Figure 2), aggregating structural performance, environmental impact, economic viability, and circularity performance into a single quantitative measure. SI integration is inspired by these studies [7], [8], [15], where CO_2 emissions, life-cycle costs, and circularity are directly accounted into the conceptual design of multi-material components while structural performance is maintained. Weighting factors (Wi) allow alignment with stakeholder priorities and regulatory targets. By embedding the Sustainability Index directly into the dataset, sustainability becomes an integral design requirement rather than a post-processing metric.

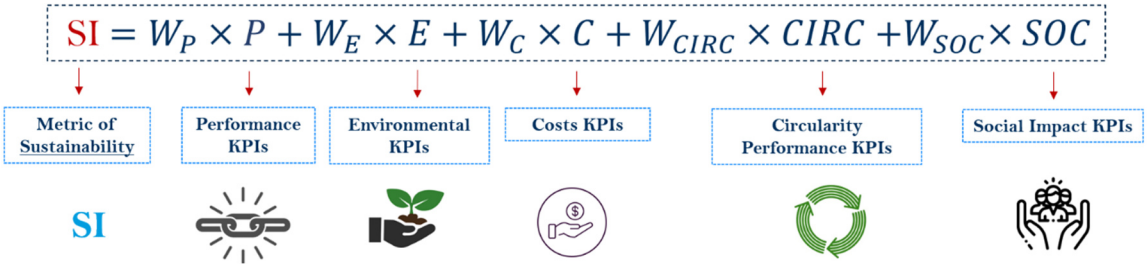


Figure 2 - Sustainability index formulation.

2.2 Case Study

The present study focuses on a Hat-Stiffened Panel, a structural configuration widely employed in primary aircraft structures where high stiffness-to-weight ratios are required, such as fuselage and wing skin panels. The geometry considered in this workflow is illustrated in Figure 3 [16]. The upper part corresponds to the stiffener or stringer, while the lower component represents the skin. The dimensions shown in the figure define the reference geometry used as the basis for generating the remaining design variants.

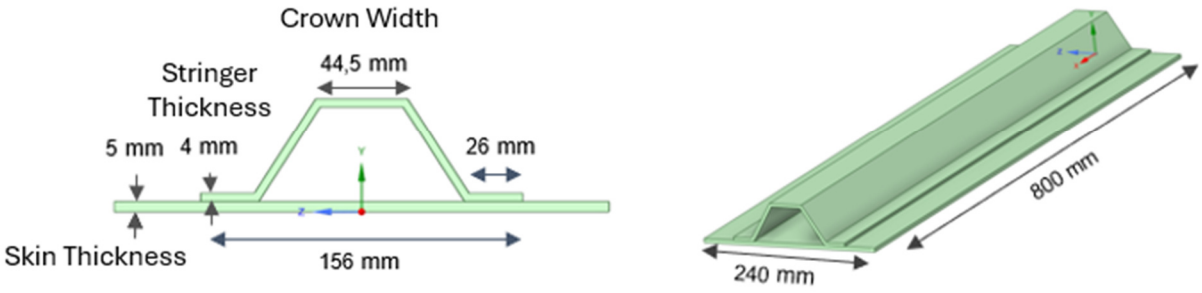


Figure 3 - Geometry and reference dimensions of hat-stiffened panel.

Table 1 presents all the possible material configurations investigated in this study. For the composite configurations, the Carbon Fiber-Reinforced Plastic (CFRP) laminate follows the stacking sequence $[0/45/-45/0/90/0]_2S$, resulting in a total of 24 prepreg carbon-fiber plies.

Table 1 - Material configurations on skin and stringer components.

Component Configurations		
No	**Skin**	**Stringer**

1	Aluminum 2024 T3	Aluminum 2024 T3
2	CFRP	Aluminum 2024 T3
3	Aluminum 2024 T3	CFRP
4	CFRP	CFRP
5	17-4PH Stainless Steel	CFRP

A Design of Experiments (DOE) is conducted where the varying parameters include the material selection for each component and three key geometric variables: skin thickness, stringer thickness and crown width. Finite Element Analyses (FEA) simulations are performed for design variants within the parameter ranges presented in Table 2. Those ranges were selected to ensure the Safety Factor of each design remains above 1.5.

The target responses extracted from the FEA simulations are the maximum total deformation, maximum principal stress, first-mode eigenfrequency, and the individual masses of the skin and stringer components. The first three responses are used to evaluate structural performance score, while the component masses combined with the corresponding material properties, are used to determine the cost, environmental, circularity and social impact metrics of the configuration. Together, these outputs are necessary and sufficient for the calculation of the Sustainability Index (SI) for each design variant. An equal-weight stakeholder policy is adopted for the SI evaluation.

Table 2 - Constraints variables and targets of DOE.

Constraints	Variables	Parameters
Minimum safety factor: >1.5	• Thickness of skin: 3-6mm • Thickness of stringer: 3-6mm • Crown width: 35-65 mm	• Maximum total deformation • Maximum principal stress • Mode 1 eigenfrequency • Skin Mass • Stringer Mass

Based on the FEA simulation results, separate Response Surface Models (RSM) are fit for each material configuration. These surrogate models are subsequently employed to predict the mechanical responses throughout the design space via factorial exploration within the selected parameter ranges. Moreover, the predicted responses are used for the calculation of the SI, leading to the generation of a dataset containing 4487 design samples. Detailed information about the generation of the dataset and the methodology process can be extracted from a previous work [8]. This dataset serves as the input for the Machine Learning models used for the subsequent design suggestion framework.

2.3 Machine Learning

Two Artificial Neural Networks (ANNs), the inverse and forward model, are trained using the generated dataset. Those models operate within a validation framework pipeline, that enables the generation of feasible design configurations satisfying the user-defined requirements.

2.3.1 Inverse Model

The inverse model's architecture along with its hyperparameters are illustrated in Figure 4. The model receives six standardized input parameters corresponding to the target design requirements: maximum total deformation, maximum principal stress, first-mode eigenfrequency, skin mass, stringer mass, and Sustainability Index (SI). The objective of the model is to predict the material configuration

and geometric dimensions that produce a hat-stiffened panel satisfying the specified structural and sustainability requirements.

The network follows a multi-task learning architecture composed of a shared feature extraction trunk and three task-specific output heads. The shared trunk consists of two fully connected layers with 64 and 128 neurons, respectively. Both layers employ the Rectified Linear Unit (ReLU) activation function, followed by batch normalization to improve convergence and stability. To deal with overfitting, a dropout rate of 30% is applied after each dense layer, by randomly deactivating a portion of the neurons during training. This technique encourages model to learn independent and meaningful representations of the design space.

The output of the shared trunk neurons is subsequently distributed to three independent output branches. The first output predicts the skin material while the second the stringer material. Both branches are formulated as multi-class classification tasks and consist of a hidden Dense layer with 64 neurons and ReLU activation, followed by a SoftMax output layer. They both follow a 10% neuron dropout policy to enhance generalization process during training.

The third output performs the regression task, associated with the geometric dimensions of the component's configuration. This branch consists of two dense layers with 64 neurons each and uses Sigmoid activation function at the output layer to produce normalized continuous outputs.

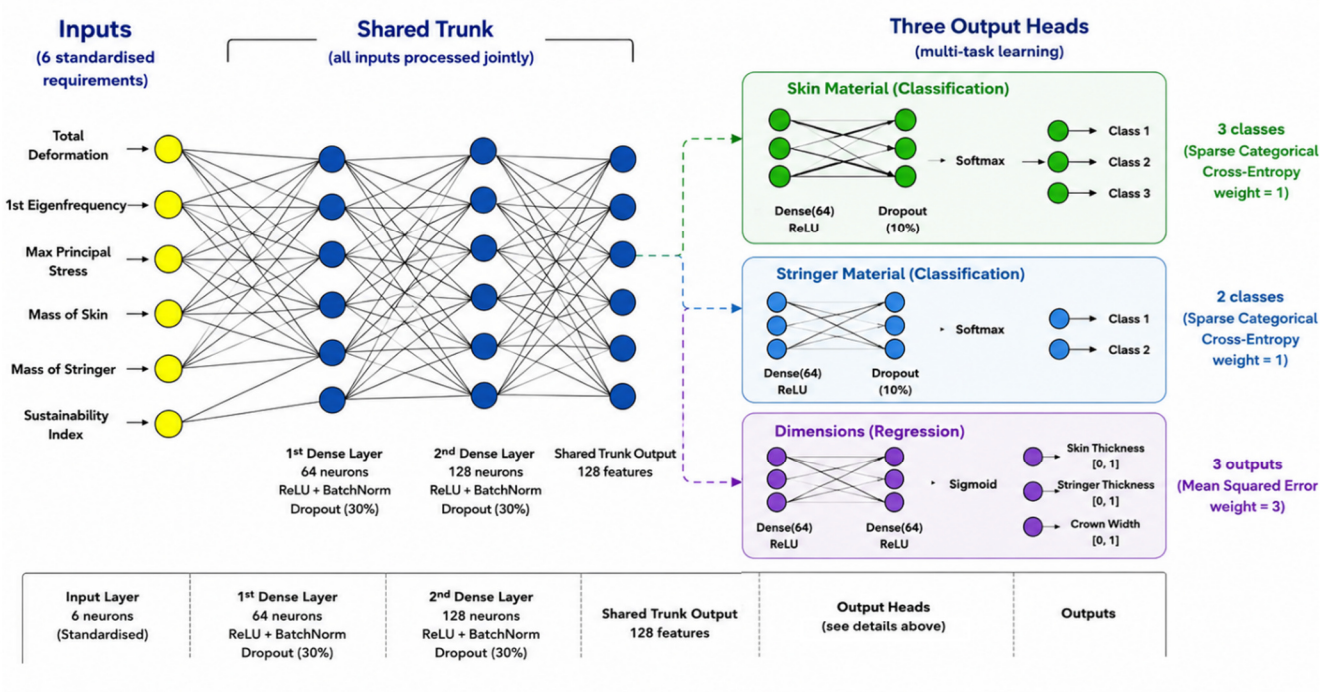


Figure 4 – Inverse model architecture and hyperparameters.

Different loss functions are employed, depending on the nature of each task. Sparse categorical cross entropy is used for the material classification tasks, while the Mean Squared Error (MSE) is adopted for the geometric regression branch. Since the dimensional task was found to be more challenging, a weighting factor of 3 was assigned to the regression loss, whereas for the classification branches were assigned a weight factor of 1.

To further improve generalization capability and reduce overfitting, L2 regularization with a coefficient of $\lambda = 10^{-3}$ is applied to all Dense layers. A 90/10 train-validation split is adopted, and validation loss is continuously monitored throughout the training process. Finally, early stopping with a patience of 20 epochs is implemented to prevent unnecessary training once convergence is achieved.

The final architecture and hyperparameter configuration of the model were determined through an iterative trial and error process, aimed at maximizing model's predicting performance, with emphasis placed on improving the accuracy of the weaker targets.

2.3.2 Forward Model

A forward model is trained to predict the mechanical responses associated with the design configurations generated by the inverse model. It serves as a validation mechanism, evaluating the capability of the inverse model to produce feasible design alternatives that satisfy the user-defined requirements.

The outputs of the inverse model constitute the input features of the forward model. The preprocessing pipeline applied to these inputs is illustrated in Figure 5. The categorical variables are transformed using one-hot encoding, with the first category removed to avoid redundancy and reduce dimensionality. Consequently, the three skin material classes are represented by two encoded variables, while the two stringer material classes are represented by a single encoded variable.

The numerical features are normalized to interval $[0,1]$ using Min-Max scaling. This normalization strategy was selected, since the geometric parameters are physically bounded within predefined ranges (Table 2), making Min-Max scaling more suitable than standard normalization approaches. Following preprocessing, the forward model receives a total of six input features.

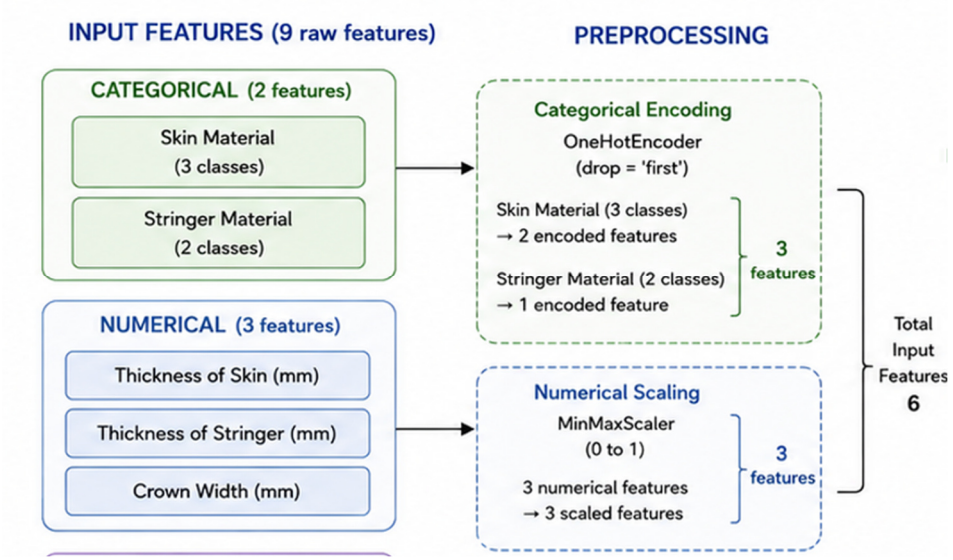


Figure 5 - Pre-processing of forward model's input features.

The architecture of the forward model is presented in Figure 5. The inputs neurons are fully connected to two hidden Dense layers containing 128 neurons each. Both hidden layers employ the Rectified Linear Unit (ReLU) activation function to capture the nonlinear relationships between design parameters and structural responses. The last hidden layer is connected to the 5 output neurons each one corresponding to one of the targets. The output layer contains five neurons corresponding to the predicted target variables: maximum total deformation, first-mode eigenfrequency, maximum principal stress, skin mass, and stringer mass. From those 5

Prior to training, the target variables are standardized using z-score normalization due to the substantial differences in magnitude among the predicted responses. A constant learning rate is adopted, to ensure smoother convergence and the maximum epoch number is set to 5000. However, the training procedure terminates if the validation performance has not improved for 50 consecutive repetitions. Additionally, 15% of the dataset is reserved as a validation subset and is excluded from

the weight update process during training.
The final model configuration and hyperparameter values were determined through an extensive grid-search procedure. The hyperparameters investigated included the number and size of hidden layers, learning rate, activation functions, and L2 regularization strength. After testing all the possible combinations, the configuration, yielding the best predictive performance, was selected as the final architecture of the forward model.

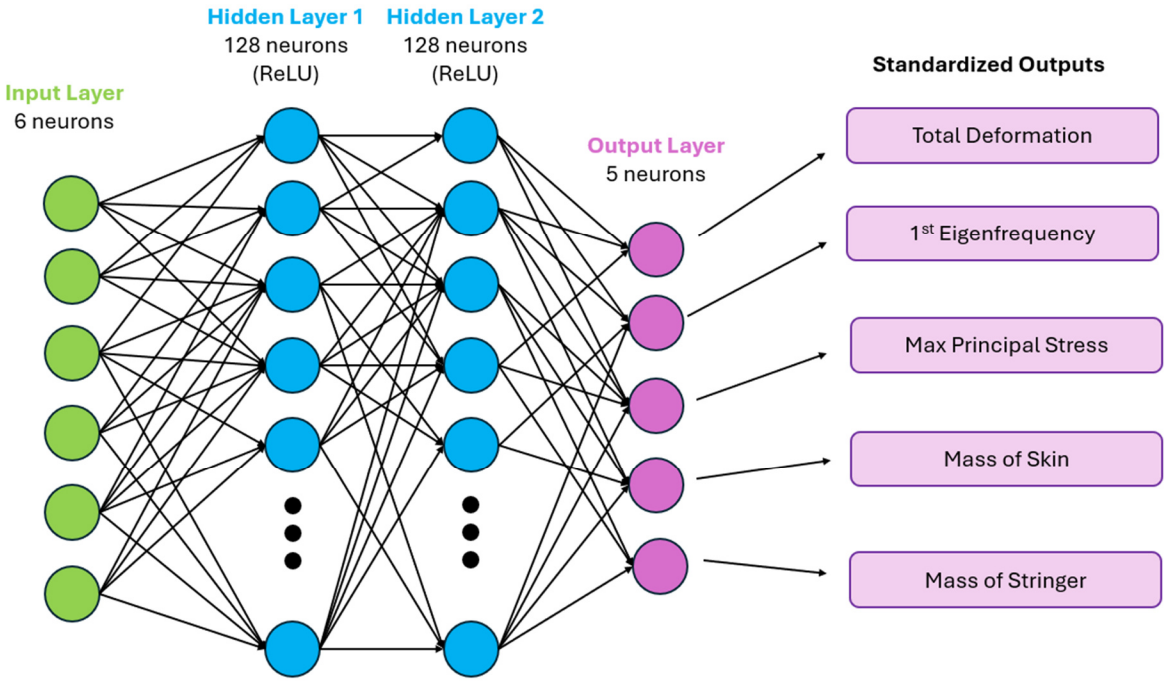


Figure 6 - Forward model's architecture.

2.3.3 Inverse-Forward Validation Framework

The objective of the inverse-forward validation framework is the generation of multiple feasible design configurations that satisfy user-defined structural and sustainability requirements. A methodology pipeline diagram is presented in Figure 7.
Initially, the user specifies lower and upper bounds for the six target responses. These include maximum total deformation, maximum principal stress, first-mode eigenfrequency, skin mass, stringer mass, and Sustainability Index (SI). For each requirement, five equal spaced points are generated within the specified interval, to uniformly explore the desired response space. To avoid requesting predictions exactly at the interval boundaries, a 10% inset is applied to the user-defined limits.
The cartesian combinations of all sampled response points is subsequently constructed and provided as input to the inverse model. Since 5 points are generated for each of the six response parameters, a total of 5^6= 15625 response combinations are evaluated. The inverse model predicts the corresponding material selections and geometric dimensions for each response combination.
Subsequently, the forward model acts as a validation mechanism and tests whether the inverse model's design suggestions fall into the given requirements. A tolerance margin of 5% is introduced during this validation stage to account for the predictive uncertainty of the forward model.
After the validation process, only the configurations satisfying all imposed requirements are retained as acceptable solutions. The ratio between valid and total generated configurations is additionally computed, expressing the feasibility degree of the requested design limits. A large ratio can potentially indicate a set of constraints that are generally easy to satisfy and vice versa, given of course that the

performance of the two models are within acceptable ranges.
Finally, the feasible configurations are grouped according to their material combinations and for each material family, the mean and that standard deviation values of the geometric variables are calculated. The material configuration together with the mean geometric dimensions constitute the final design recommendation of the framework. The corresponding standard deviation values serve as confidence indicators, expressing the allowable variation range around the proposed dimensions, while still complying with the user-defined design requirements.

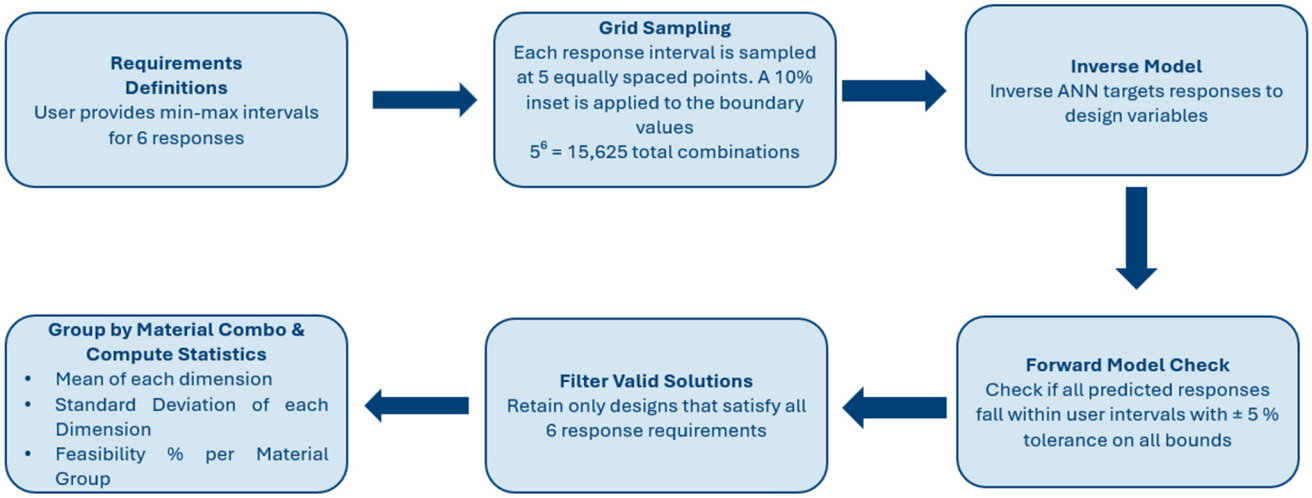


Figure 7 - Inverse forward validation framework.

2.4 Design Scenarios

To evaluate the effectiveness of the inverse-forward validation framework in generating designs based on user-defined requirements, five representative design cases were developed. Table 3 summarizes these five design cases along with their corresponding requirement profiles. The selected cases were devised to emphasize sustainability-oriented design considerations during the preliminary design phase, while simultaneously satisfying structural performance characteristics commonly encountered in aerospace engineering applications.
To ensure that the requested response ranges correspond to regions of the design space sufficiently represented within the training dataset, each case was compared against the statistical quartiles of the dataset distribution. For response variables not specified in “Requirements” column of each case are requested within the interquartile range between the first quartile (Q1) and the third quartile (Q3) of the dataset. Table 4 summarizes the first and the third quartile value of each target along with their minimum and maximum values observed on the dataset.

Table 3 - Design cases to test model’s effectiveness.

Design Cases	Requirements
1st Case – Balanced Mid-Range	$Q1 \leq All\ Targets \leq Q3$
2nd Case – Ultra Lightweight and high SI	$Skin\ and\ Stringer\ Mass \leq Q1$ and $SI \geq Q3$
3rd Case – High SI	$SI \geq Q3$
4th Case – Low SI	$SI \leq Q1$
5th Case – Low Deformation and High Stress	$Deformation \leq Q1$ and $Stress \geq Q3$

Table 4 - Dataset quartile reference.

Targets	Q1	Q3	min	max
Deformation (mm)	2.27	3.85	0.79	17.9
Frequency (Hz)	148.48	222.84	100.59	267.05
Stress (MPa)	516.2	659.7	358.3	980
Skin Mass (kg)	1.393	3.044	0.901	9.056
Stringer Mass (kg)	1.011	1.570	0.706	2.868
SI	0.427	0.516	0.328	0.902

3. Results & Discussion

3.1 Inverse model training

The inverse model demonstrated strong predictive performance for both the categorical and numerical output tasks, as is presented in Table 4.

Table 5 - Losses for categorical and numerical predictions of the inverse model on the validation dataset.

Averaged Scaled MAE for Numerical Outputs	0.028
Averaged Scaled MSE for Numerical Outputs	0.001
Skin Material Prediction Loss	$5x10^{-6}$
Stringer Material Prediction Loss	$3x10^{-6}$

The categorical predictions achieved nearly perfect accuracy on both training and validation sets, from the early training epochs and remained stable throughout the training process (Figure 8). This behavior is also confirmed by the confusion matrices shown in Figure 9, where all validations samples were classified correctly with no one misclassified sample across Aluminum, CFRP and Steel classes.
Thus, material classification is a discrete, well-separated problem within the feature space. More specifically, the different material combinations produce distinct structural response and mass distributions, making the six input features effective fingerprints of the employed materials

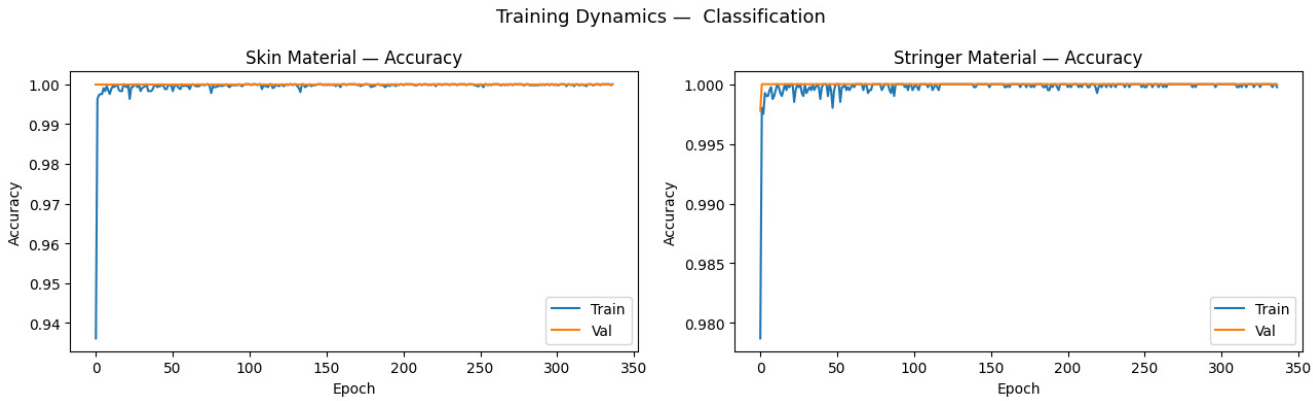


Figure 8 - Training dynamics of classification training and validation loss across training epochs.

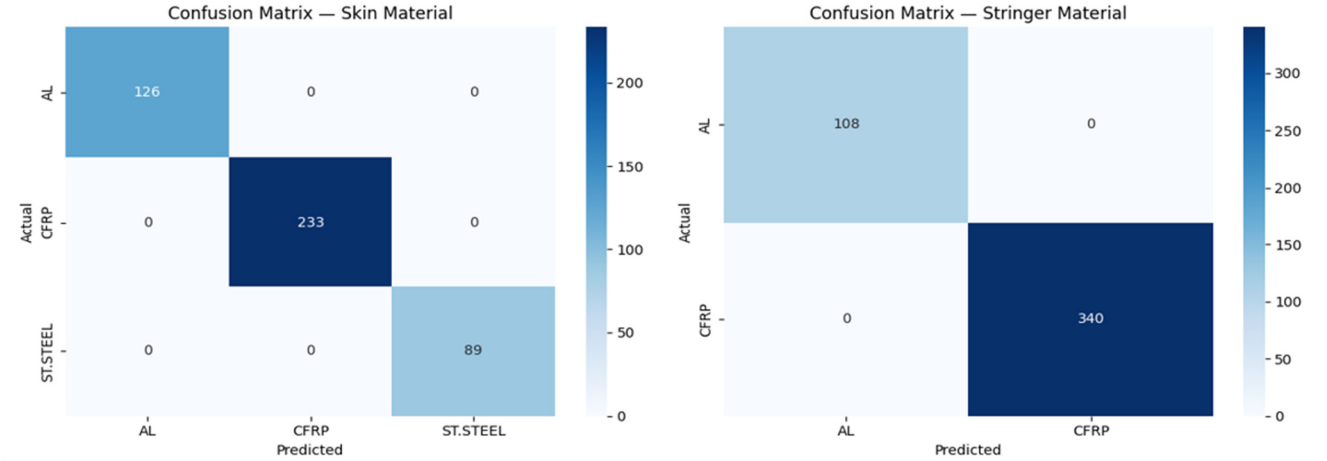


Figure 9 - Confusion matrix of categorical variables prediction.

Figure 10 illustrates the evolution of the training and validation losses associated with the numerical outputs, namely skin thickness, stringer thickness and crown width. The regression loss decreases consistently throughout the training procedure, indicating stable convergence of the inverse mapping problem. The loss decreases rapidly on the initial epochs, followed by a gradual convergence on the remaining epochs, indicating effective learning of the non-linear relationships between the mechanical responses and the geometric design variables.

The training loss remains slightly higher than the validation loss, as a result of the regularization techniques applied during training particularly neuron dropout. This is not inverse overfitting, but rather an expected consequence, since deactivating a random percentage of neurons on every forward pass adds noise and inflates the training loss. On the other hand, during validation, dropout is disabled and all neurons participate on the forward propagation, allowing the model to demonstrate its full predictive capability.

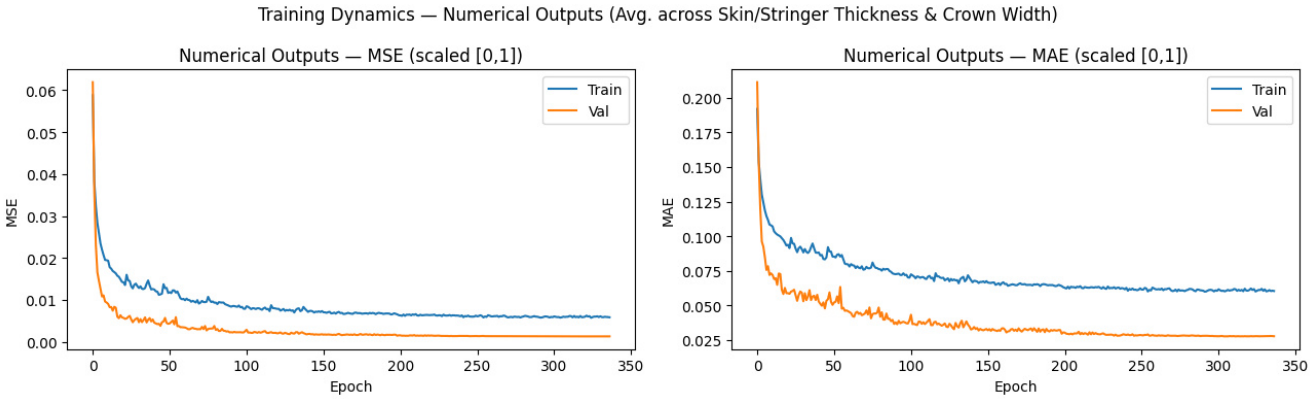


Figure 10 – Numerical Outputs MSE and MAE of training and validation error across training epochs.

Figure 11 illustrates the residual plots for each geometric parameter predicted by the inverse model. The residuals associated with skin and stringer thickness are bounded within $\pm 0.3\ mm$, which is an acceptable accuracy given the $3.0 - 6.0\ mm$ design range of these parameters. Additionally, the residuals remain centered around zero indicating the absence of strong systematic bias. However, large deviations are observed near the parameter boundaries, suggesting reduced prediction confidence in regions where fewer training samples are available.

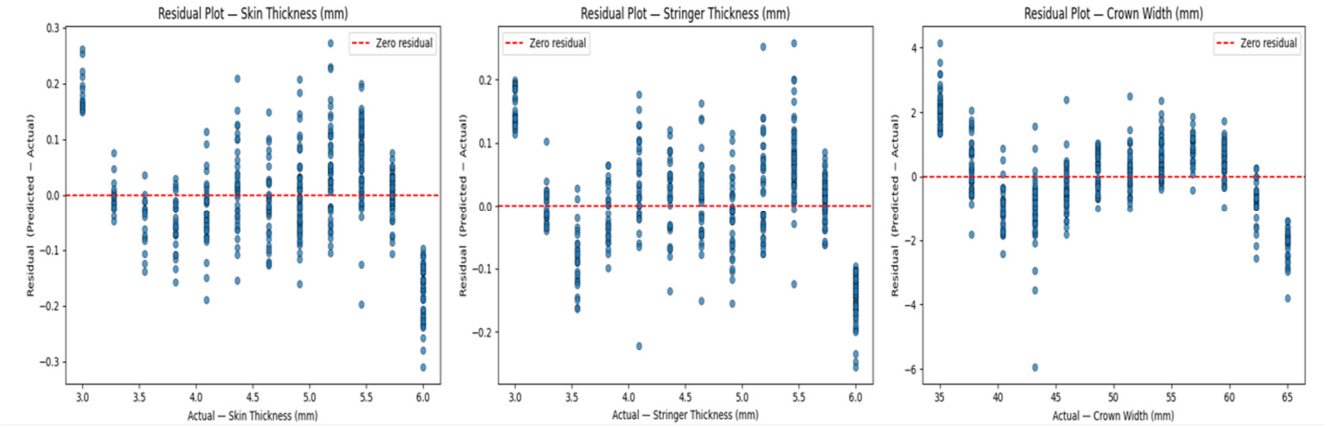


Figure 11 - Residual plots of skin thickness, stringer thickness and crown width.

In contrast, the crown width residuals present higher scatter values which can be approximately equal to $6\,mm$, indicating increased ambiguity on the prediction of the crown width variable. Unlike the material classification task, the prediction of the exact thickness and crown width is a continuous regression problem, where many different design configurations can produce similar mechanical responses. Consequently, the problem becomes ill-posed, since the model must resolve ambiguity that genuinely exists in data, which no size of training data can fully extinct. As a result, there is an upper limit to inverse model's achievable accuracy in the numerical prediction task, on contrast to the material classification problem.

3.2 Forward Model Training

An extensive grid search across influential ANN model hyperparameters was conducted to determine the best performing forward model configuration. The candidate models were ranked according to their prediction performance on the most challenging target variable, the prediction of total deformation. Since deformation exhibited the weakest predictive behavior among all targets, it was selected as the primary criterion for model selection.

The optimal architecture corresponded to an ANN consisting of 128 neurons each, an L2 regularization coefficient of 10^{-3} and a learning rate of $2x10^{-3}$. The performance metrics are summarized in Table 5.

Table 6 - Performance metrics of the final forward model.

Response Target	R^2 coefficient	RMSE	Mean Value	Standard Deviation	Unit
Total Deformation	0.9843	0.148	3.177	1.252	mm
1st Eigenfrequency	0.995	0.95	183.28	41.8	Hz
Max Principal Stress	0.9957	7.27	597	115	MPa
Skin Mass	0.9997	0.043	3.063	2.524	kg
Stringer Mass	0.9997	0.011	1.443	0.582	kg

Due to the large difference in the magnitude of the predicted target variables (e.g. deformation in the order of $10^{-4}m$, stress at $10^{7}Pa$, mass values around 1kg), the raw RMSE values are not directly comparable across targets. To enable visualization of the learning behavior during training, the RMSE of each target was normalized by the standard deviation of the corresponding response variable. Subsequently, the averaged normalized RMSE across all targets was computed and used for the learning curves presented in Figure 12.

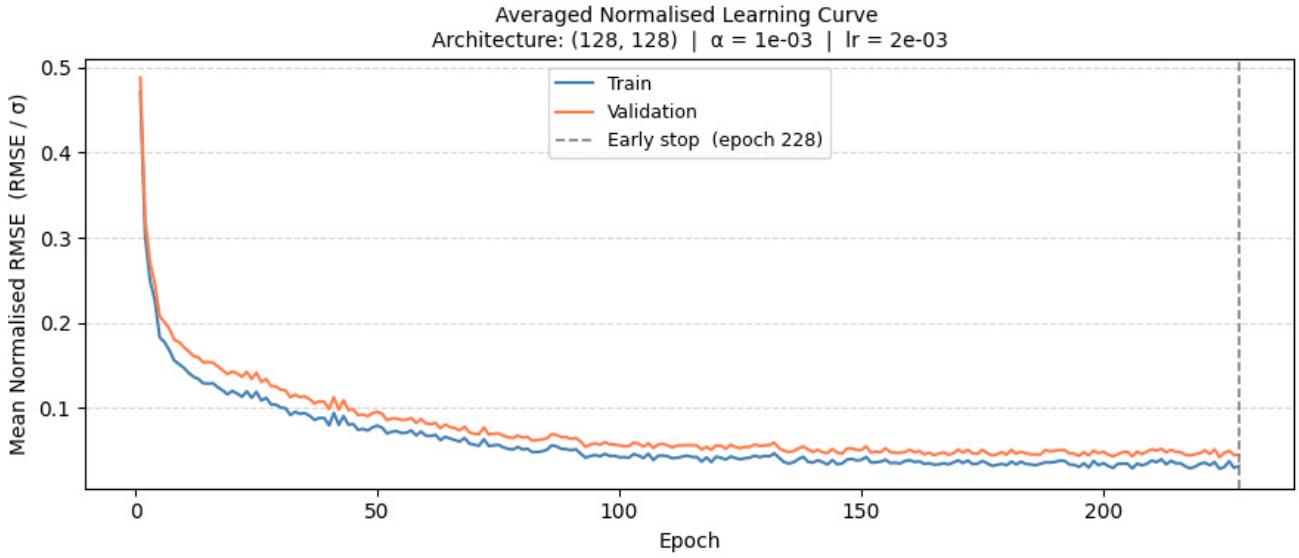


Figure 12 - Averaged normalized learning curves of the forward model.

Inspecting Figure 12, it is evident that both of the loss curves present their steepest decline withing the first 50 epochs, indicating the quick adaptation of the forward model during the initial learning stage. During the rest of the epochs the losses continue to decrease at a slower yet stable rate, progressively improving the accuracy of the network. Finally, the loss metrics present a plateau between the 210th and 230th epoch, forcing the training process to stop without reaching the maximum number of 5000 training epochs. Moreover, the validation and training loss curves present a close agreement throughout the entire training process, indicating that the forward model is able to generalize adequately without overfitting.

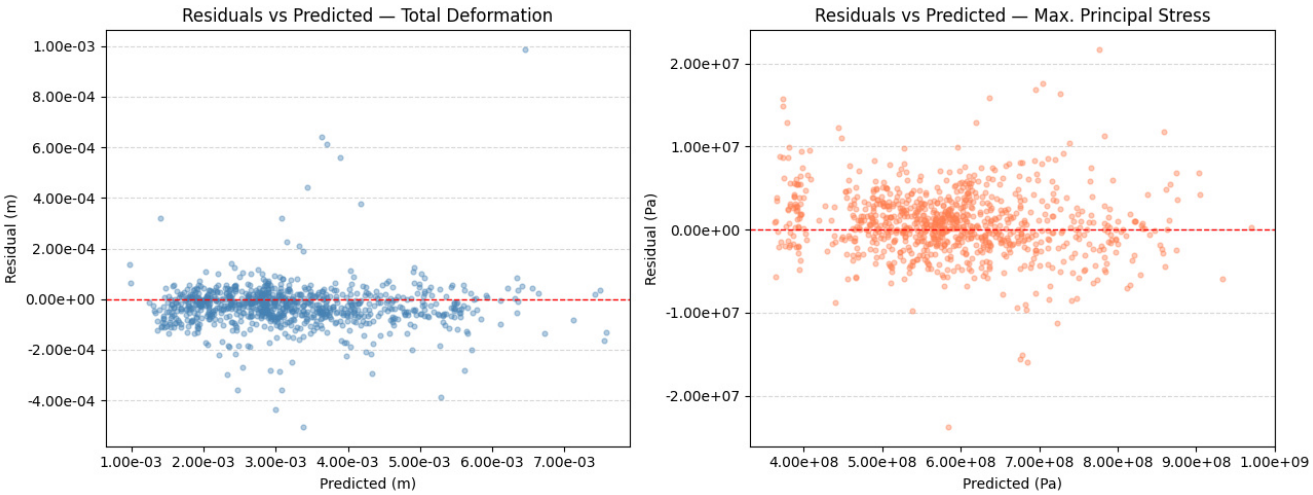


Figure 13 - Residual plots for deformation and stress prediction.

Following training, the forward model was evaluated on the validation dataset, consisting of samples not previously seen during the training phase. Figure 13 presents the residual distributions associated with deformation and stress prediction, which constitute the most demanding response targets of the forward model.

Observing the plots, the majority of the prediction errors remain bounded within $\pm 0.2\,mm$ for deformation and $\pm 10\,MPa$ for stress. Considering the statistical characteristics summarized in Table 5, these deviations indicate that the forward model achieves a satisfactory level of predictive accuracy

for both response variables. Furthermore, the residuals remain generally centered around zero, suggesting the absence of strong systematic prediction bias.

3.3 Scenario Results

Design configurations satisfying the target requirements specified in Table 3 were generated using the inverse–forward model collaboration framework. Table 7 summarizes the suggested skin and stringer material combinations for each design case, together with the predicted geometric dimensions ($mean\ value\ \pm standard\ deviation$). The mean values represent the final geometric recommendations of the framework, while the corresponding standard deviations act as confidence indicators, expressing the allowable variation range around the suggested dimensions that still satisfies the imposed requirements. Additionally, the percentage of feasible configurations over the total number of generated candidates is reported, providing an indication of how restrictive each requirement set is within the explored design space.

Figure 14 illustrates the geometric design recommendations generated by the inverse–forward validation framework for each investigated design scenario. The predicted skin thickness, stringer thickness, and crown width values are presented for the feasible configurations identified by the framework. Due to the different magnitude ranges of the thickness and width parameters, the thickness-related variables are grouped in a single plot, while the crown width predictions are presented separately to improve visualization clarity. The markers represent the mean values of the recommended geometric parameters for each design case, while the corresponding error bars indicate the uncertainty range defined by the standard deviation around the mean value.

Table 7 – Suggested configurations for each design case from the inverse and forward model collaboration.

Case Name	Skin Material	Stringer Material	Skin Thickness	Stringer Thickness	Crown Width	Feasibility Percentage (%)
Balanced Mid-Range	AL	CFRP	5.48 ± 0.24	5.29 ± 0.24	53.24 ± 4.23	1.60
High SI	AL	CFRP	5.52 ± 0.27	3.40 ± 0.38	43.63 ± 6.71	11.01
High SI	CFRP	CFRP	4.39 ± 0.53	4.20 ± 0.74	56.47 ± 7.21	1.08
Low SI	CFRP	CFRP	5.94 ± 0.08	5.43 ± 0.42	63.13 ± 5.10	73.82
Low Deformation and High Stress	AL	CFRP	5.36 ± 0.15	4.32 ± 0.36	61.11 ± 2.46	0.31
Low Deformation and High Stress	CFRP	CFRP	3.98 ± 0.03	4.25 ± 0.06	61.88 ± 2.58	0.04

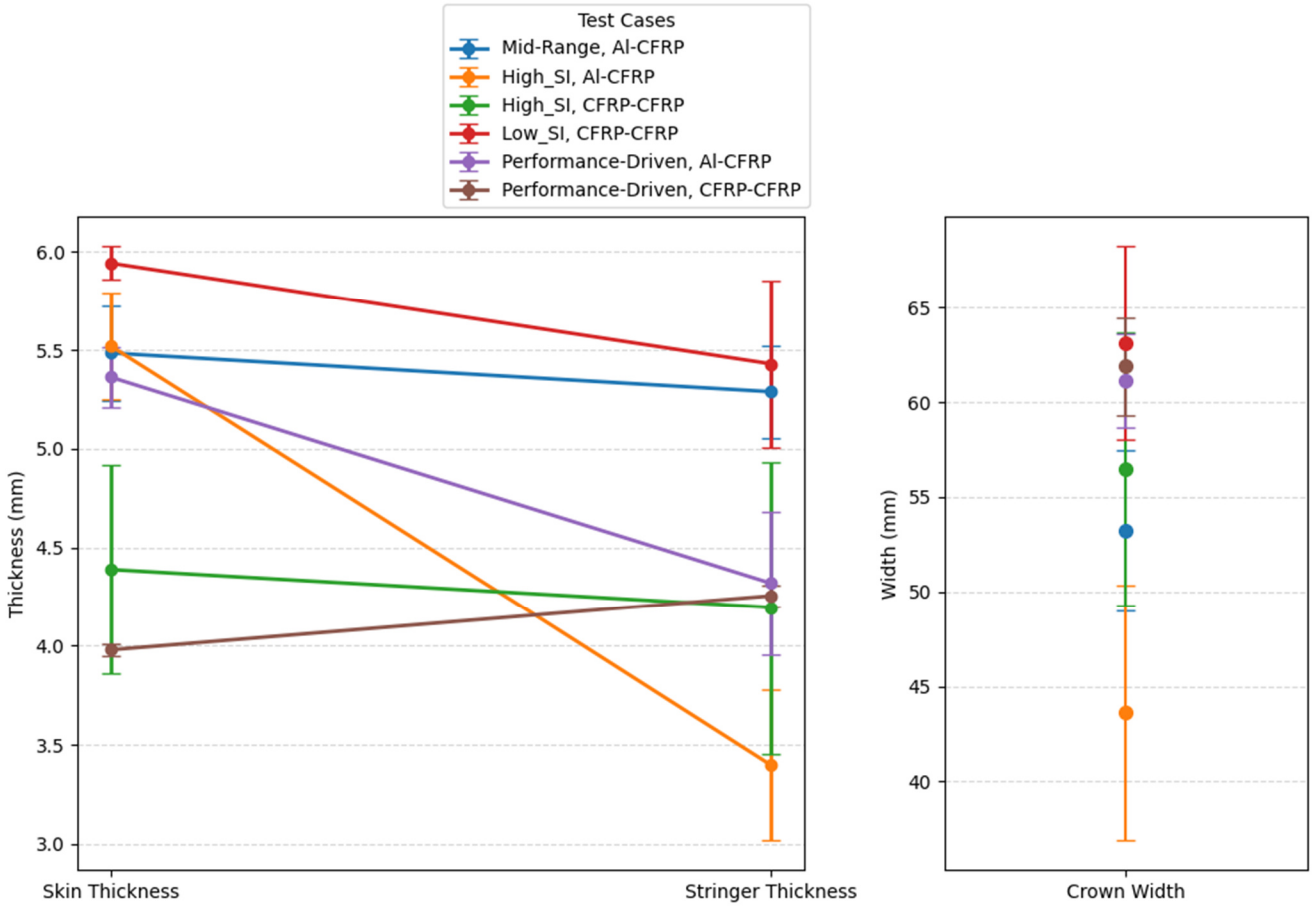


Figure 14 - Suggestion intervals of the numerical design variables provided by the framework for each design scenario.

The obtained results demonstrate that the inverse–forward framework can identify meaningful relationships between structural performance, sustainability objectives, material selection, and geometric characteristics. In particular, the framework consistently associates CFRP configurations with structurally demanding cases, while aluminum-containing combinations are preferred in cases emphasizing sustainability and balanced performance. This behavior indicates that the models successfully captured the underlying trends embedded within the dataset during training, rather than merely memorizing isolated design samples.

The feasibility percentages further reveal the sensitivity of the design space to the imposed requirements. Cases involving contradictory or highly restrictive objectives produced very low feasibility percentages, indicating that only a limited number of configurations within the dataset satisfy such combinations of requirements.

More specifically, the “Balanced Mid-Range” case yielded feasible solutions exclusively based on the AL + CFRP material combination, suggesting that the framework recognizes this hybrid configuration as an effective compromise between structural efficiency and sustainability-oriented objectives. Similarly, the “High SI” case favored lightweight hybrid or fully CFRP configurations with thinner stringers, indicating that the model associates reduced structural mass with improved sustainability performance. On the other hand, the “Low SI” case produced exclusively CFRP + CFRP configurations with large thicknesses and wide crown dimensions, highlighting that the framework identifies fully composite structures as structurally efficient but less favorable from a sustainability perspective due to higher environmental and economic impacts.

Finally, the “Low Deformation and High Stress” case resulted in extremely low feasibility percentages, confirming the difficulty of simultaneously satisfying highly demanding structural requirements.

Nevertheless, the framework consistently proposed configurations with wide crown dimensions and relatively thick structural components, demonstrating that the models successfully learned the relationship between panel stiffness, stress distribution, and geometric reinforcement strategies. The very small standard deviations observed in some of the extreme cases additionally suggest that the feasible design space collapses toward nearly deterministic solutions under highly restrictive performance constraints.

4. Conclusions

In the present study, Finite Element Analysis (FEA) simulations were conducted to investigate the structural response of a hat-stiffened panel under different material configurations and geometric variations. The resulting simulations were compiled into an initial database, which was subsequently expanded through the use of Response Surface Models (RSMs) fitted on the FEM-generated data. Through factorial exploration of the design space, an artificial dataset containing 4,487 design samples was generated.

Based on this dataset, two Artificial Neural Network (ANN) models were developed and trained. The first model, referred to as the inverse model, predicts feasible hat-stiffened panel configurations given target mechanical responses and a sustainability score. The second model, namely the forward model, predicts the structural responses corresponding to a specified configuration. Following training, both models were integrated into an inverse–forward validation framework, where the inverse model generates candidate design solutions and the forward model validates them against the user-defined requirements. The framework ultimately returns feasible material configurations together with geometric recommendations expressed through mean values and standard deviations, where the latter quantify the allowable variation range of the proposed dimensions while still satisfying the imposed design objectives.

To evaluate the proposed methodology, five representative design cases were formulated based on sustainability-oriented considerations and structural performance requirements commonly encountered in the aerospace industry. The framework successfully generated feasible design suggestions for four out of the five cases, demonstrating its ability to capture meaningful relationships between structural performance, sustainability objectives, material selection, and geometric characteristics embedded within the dataset. The only infeasible case involved highly contradictory requirements and as a result the framework was not able to produce designs complying with those restrictions.

Overall, the proposed inverse–forward collaboration framework enables real-time generation and validation of design configurations based on user-defined requirements. The methodology contributes toward targeted exploration of the design space, significantly reducing the need for exhaustive finite element simulations and accelerating the preliminary design process for sustainable aerostructures.

Despite the promising results, several limitations remain. The inverse design problem is inherently ill-posed, meaning that multiple distinct design configurations may satisfy similar performance requirements. This ambiguity complicates the training process of the inverse model and limits its achievable predictive accuracy. Consequently, some generated configurations may not perfectly satisfy highly restrictive user-defined requirements, potentially leading to the exclusion of feasible solutions in demanding design scenarios.

Future work should focus on expanding the dataset to cover wider regions of the design space, thereby improving the ability of the models to learn the complex relationships between geometry, material selection, structural performance, and sustainability metrics. A larger and more diverse dataset is expected to improve the robustness and generalization capability of the framework, particularly for

highly demanding design cases. Additionally, comparative studies involving alternative supervised learning techniques, such as Support Vector Machines (SVMs), Decision Trees, and ensemble learning methods, could provide valuable benchmarking against the proposed ANN-based framework. Finally, the proposed methodology could be extended to additional engineering sectors, including automotive engineering, energy systems, machine design, and civil engineering applications, where rapid identification of feasible configurations could significantly reduce computational cost and design development time.

5. Contact Author Email Address

Mail to: angelos.filippatos@upatras.gr

6. Copyright Statement

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

References

[1] M. Klöwer, M. R. Allen, D. S. Lee, S. R. Proud, L. Gallagher, and A. Skowron, "Quantifying aviation's contribution to global warming," *Environmental Research Letters*, vol. 16, no. 10, Oct. 2021, doi: 10.1088/1748-9326/ac286e.

[2] A. Mayyas, A. Qattawi, M. Omar, and D. Shan, "Design for sustainability in automotive industry: A comprehensive review," *Renewable and Sustainable Energy Reviews*, vol. 16, no. 4, pp. 1845–1862, May 2012, doi: 10.1016/J.RSER.2012.01.012.

[3] European Commision, "Communication from the Commision to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions," in *A new Circular Economy Action Plan For a cleaner and more competitive Europe*, Brussels: Commision, Mar. 2020.

[4] A. Candela, G. Sandrini, M. Gadola, D. Chindamo, and P. Magri, "Lightweighting in the automotive industry as a measure for energy efficiency: Review of the main materials and methods," *Heliyon*, vol. 10, no. 8, p. e29728, Apr. 2024, doi: 10.1016/J.HELIYON.2024.E29728.

[5] G. Copani, P. Shafinejad, T. Hipke, R. Haase, and T. Paizs, "New metals remanufacturing business models in automotive industry," *Procedia CIRP*, vol. 112, pp. 436–441, Jan. 2022, doi: 10.1016/J.PROCIR.2022.09.033.

[6] S. Bobba, P. Tecchio, F. Ardente, F. Mathieux, F. M. dos Santos, and F. Pekar, "Analysing the contribution of automotive remanufacturing to the circularity of materials," *Procedia CIRP*, vol. 90, pp. 67–72, 2020, doi: 10.1016/J.PROCIR.2020.02.052.

[7] A. Filippatos *et al.*, "Sustainability-Driven Design of Aircraft Composite Components," *Aerospace 2024, Vol. 11, Page 86*, vol. 11, no. 1, p. 86, Jan. 2024, doi: 10.3390/AEROSPACE11010086.

[8] A. Filippatos, D. Markatos, A. Theochari, and S. Pantelakis, "Integrating Sustainability in Aircraft Component Design: Towards a Transition from Eco-Driven to Sustainability-Driven Design," *Aerospace 2025, Vol. 12, Page 140*, vol. 12, no. 2, p. 140, Feb. 2025, doi: 10.3390/AEROSPACE12020140.

[9] R. S. Kadadevaramath and M. Sangli, "Application of data science in automotive design optimization for faster convergence," *IOP Conf. Ser. Mater. Sci. Eng.*, vol. 1189, no. 1, p. 012031, Oct. 2021, doi: 10.1088/1757-899X/1189/1/012031.

[10] F. Bodendorf and J. Franke, "A machine learning approach to estimate product costs in the early product design phase: a use case from the automotive industry," *Procedia CIRP*, vol. 100, pp. 643–648, Jan. 2021, doi: 10.1016/J.PROCIR.2021.05.137.

[11] C. Li and I. Y. Kim, "Multi-material topology optimization for automotive design problems," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 232, no. 14, pp. 1950–1969,

Dec. 2018, doi: 10.1177/0954407017737901.

[12] P. V. Badarinath, M. Chierichetti, and F. D. Kakhki, "A Machine Learning Approach as a Surrogate for a Finite Element Analysis: Status of Research and Application to One Dimensional Systems," *Sensors 2021, Vol. 21, Page 1654*, vol. 21, no. 5, p. 1654, Feb. 2021, doi: 10.3390/S21051654.

[13] J. G. Hoffer, B. C. Geiger, P. Ofner, and R. Kern, "Mesh-Free Surrogate Models for Structural Mechanic FEM Simulation: A Comparative Study of Approaches," *Applied Sciences 2021, Vol. 11, Page 9411*, vol. 11, no. 20, p. 9411, Oct. 2021, doi: 10.3390/APP11209411.

[14] L. Liang, M. Liu, C. Martin, and W. Sun, "A deep learning approach to estimate stress distribution: a fast and accurate surrogate of finite-element analysis," *J. R. Soc. Interface*, vol. 15, no. 138, Jan. 2018, doi: 10.1098/RSIF.2017.0844/35772.

[15] S. Malefaki, D. Markatos, A. Filippatos, and S. Pantelakis, "A Comparative Analysis of Multi-Criteria Decision-Making Methods and Normalization Techniques in Holistic Sustainability Assessment for Engineering Applications," *Aerospace 2025, Vol. 12, Page 100*, vol. 12, no. 2, p. 100, Jan. 2025, doi: 10.3390/AEROSPACE12020100.

[16] E. Barkanov, "Optimal weight design of laminated composite panels with different stiffeners under buckling loads," in *Proceedings of the 27th International Congress of the Aeronautical Sciences*, Nice, France, Sep. 2010.