International Forum on Aeroelasticity and Structural Dynamics
IFASD 2026
16-19 June 2026, Goettingen (GERMANY)

A MULTI-FIDELITY NEURAL NETWORK FRAMEWORK FOR ENHANCING STRIP THEORY AERODYNAMICS

Vítor B. Santos[1], Flávio L. Cardoso-Ribeiro[1], Álvaro A. G. Quesada[2], Pedro J. González[2], Flávio J. Silvestre[2], Wolf R. Krüger[3], Henry Späth[4], Zamira Daw[4]

[1]Aeronautics Institute of Technology (ITA)
Praça Marechal Eduardo Gomes 50, 12228-900 São José dos Campos, Brazil
vitor.santos@ga.ita.br
flaviocr@ita.br

[2]Technical University of Berlin
Marchstraße 12, 10587 Berlin, Germany
garcia.quesada@campus.tu-berlin.de
p.gonzalez.ramirez@tu-berlin.de
flavio.silvestre@tu-berlin.de

[3]German Aerospace Center (DLR)
Bunsenstraße 10, 37073 Goettingen, Germany
wolf.krueger@dlr.de

[4]Institute of Aircraft Systems, University of Stuttgart
Pfaffenwaldring 27, 70569 Stuttgart, Germany
henry.spaeth@ils.uni-stuttgart.de
zamira.daw@ils.uni-stuttgart.de

Keywords: Neural networks, surrogate model, transfer learning, aerodynamics, strip theory

Abstract: Accurate yet computationally efficient aerodynamic modeling remains a critical bottleneck in flight dynamics simulations of flexible aircraft. While low-fidelity analytical approaches like strip theory offer rapid inference speeds, they often exhibit significant discrepancies when compared to empirical wind-tunnel data, as observed during the testing of the TU-Flex flying demonstrator. To address this limitation, this paper proposes a multi-fidelity Scientific Machine Learning framework designed to enhance baseline strip theory predictions without the prohibitive computational costs of high-fidelity solvers. The methodology employs a sequential dual-network architecture. The first neural network employs transfer learning to effectively combine the aerodynamic trends captured by Computational Fluid Dynamics simulations with the magnitudes observed in experimental wind-tunnel measurements. This data mixing serves as a high-fidelity baseline to parameterize the physics-informed loss function of a second surrogate network. By embedding the governing equations of generalized aerodynamic forces into its training process, this second model maps operational flight conditions to quasi-steady modal forces, successfully recovering continuous spanwise lift and pitching moment distributions. Results demonstrate that the proposed framework accurately matches target high-fidelity load profiles while requiring minimal computational effort. This surrogate model establishes a pathway for future real-time aeroelastic evaluation and system identification within the TU-Flex flight simulation environment.

1 INTRODUCTION

The estimation of the aerodynamic forces is often the computational bottleneck in aircraft simulation frameworks, especially when structural flexibility is taken into account, and real-time solutions must consider the coupling between flight dynamics and aeroelastic modes [1]. During preliminary design, low-fidelity methods are preferred for both structural dynamics and aerodynamics, allowing design constraints and general flight characteristics to be evaluated in a relatively short time frame and with reduced computational costs, despite compromising accuracy. However, even such low-fidelity approaches can be cumbersome when embedded in flight dynamics solvers for flexible aircraft.

The increasing availability of multi-fidelity data — ranging from low-fidelity analytical methods to wind-tunnel measurements and Computational Fluid Dynamics (CFD) — has provided fertile ground for the development of data-driven aerodynamic models. In this context, we turn our attention to the strip theory aerodynamic model proposed in [2] for the TU-Flex flying demonstrator. While that simple aerodynamic model performs well in the flight dynamics simulation framework developed for the aircraft, the accuracy of the method has been challenged by wind-tunnel observations considering the same flight conditions, as depicted in Figure 1. Therefore, there is a need for a more accurate aerodynamic model while keeping the computational costs at reasonable levels, which means that higher-fidelity methods are certainly prohibitive. This is where physics-inspired, data-driven surrogate models can play an interesting role, and it is exactly what we intend to investigate in this work.

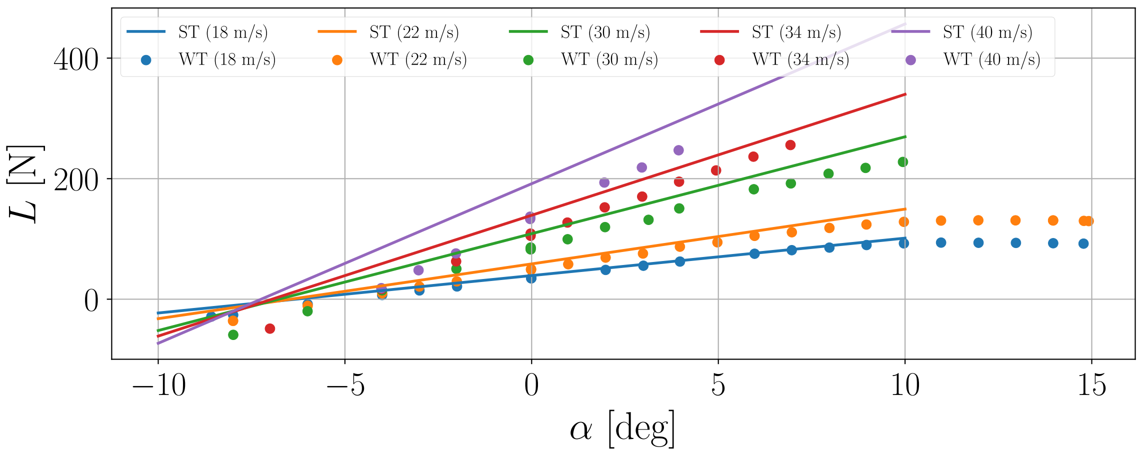

Figure 1: Discrepancies observed in lift from wind-tunnel (WT) and strip theory (ST) data.

We propose a framework structured on Scientific Machine Learning (SciML) concepts aimed at improving one of the aerodynamic models of the TU-Flex aircraft. By feeding CFD and experimental information into a baseline feedforward neural network, we improve the accuracy of the predictions made by a second neural network intended to serve as a surrogate for strip theory aerodynamics. The core idea behind our model is to embed the equation governing the generalized aerodynamic forces term into the loss function to be minimized during the training of the surrogate model.

Preliminary findings show that the proposed method can predict the quasi-steady generalized aerodynamic forces, allowing lift and pitching moment distributions to be accurately retrieved at a low computational cost in terms of inference time compared to high-fidelity aerodynamic models. The results also open up possibilities for tackling more complex tasks, such as unsteady aerodynamics, as long as experimental or numerical data become available.

We begin this document by introducing the TU-Flex flying demonstrator in Section 2, followed in Section 3 by a theoretical background on the use of strip theory for unsteady and incompressible aerodynamics, as well as the main concepts of this emerging area of Scientific Machine Learning. A comprehensive description of our methodology is presented in Section 4, followed by our main findings showcased in Section 5. Finally, in Section 6, we conclude our work and outline directions for future studies.

2 TU-FLEX FLYING DEMONSTRATOR

The lack of experimental setups for variable flexibility wings poses challenges for advances in the literature on how different levels of flexibility influence the flight dynamics, aeroelasticity, and control characteristics of a High Aspect Ratio (HAR) aircraft. In this context, the Chair of Flight Dynamics, Flight Control, and Aeroelasticity (FMRA: *Flugmechanik, -regelung und Aeroelastizität*) at the *Technische Universität Berlin* and the Institute of Aeroelasticity of the German Aerospace Center (DLR) have partnered to develop a flight demonstrator called TU-Flex [3].

The prototype shown in Figure 2, which has a typical cargo aircraft configuration, serves as a convenient testbed for wings of variable flexibility due to its wing module attachment system.

Figure 2: TU-Flex (FW configuration) flying demonstrator rollout. Credits: FMRA, 2025

In this study, we consider the Flexible Wing (FW) configuration of the TU-Flex, for which in-flight wingtip deflections of up to 10% of the half-span are expected. The main properties of the FW are summarized in Table 1.

Table 1: Geometric properties of the TU-Flex's flexible wing.

Span	3.6 m
Root chord	0.3 m
Taper ratio	0.5
Aspect ratio	16
Leading edge sweep angle	20 deg

The TU-Flex aircraft serves as a comprehensive experimental testbed designed to advance the characterization, modeling, and control of flexible and very flexible aircraft. In a ground-testing capacity, it provides a framework to establish new procedures for defining fundamental elastic and aeroelastic properties through static, ground-vibration, and wind-tunnel experiments. During flight operations, the system will gather empirical data on the complex, coupled motion between rigid-body flight dynamics and structural aeroelasticity in response to control inputs, directly supporting the implementation of novel system identification techniques. Furthermore, the platform acts as an adaptable developmental tool to evaluate diverse sensor configurations for real-time structural shape reconstruction, assess alternative control surface arrangements, and deploy the advanced structural regulators and control architectures necessary to ensure the safe stabilization and navigation of highly flexible flight vehicles. A comprehensive explanation of the TU-Flex project can be found in [4–7].

3 THEORETICAL BACKGROUND

In this section, we present some fundamental concepts behind the methods to be used throughout this work. In Subsection 3.1, we describe the unsteady and incompressible aerodynamics coupled with strip theory under a few referenced assumptions. Subsection 3.2 concerns the Scientific Machine Learning ideas and the concepts important for the development of our approach.

3.1 Unsteady and Incompressible Aerodynamics via Strip Theory

During aircraft preliminary design, the use of low- to medium-fidelity aerodynamic models is usually preferred in flight dynamics simulation frameworks due to their computational performance, allowing for efficient integration with control laws and optimization routines.

Following the methodology presented in [8, 9], the lifting surface is discretized into n spanwise stations with uniform properties, as is conventional in strip theory. Consequently, the local downwash and the aerodynamic loads are fully described by two-dimensional aerodynamic theory [10]. Moreover, a modification of the usual strip theory proposed in [11] was employed, in which available data for the lift curve slope ($c_{\ell,\alpha}$) and the aerodynamic center position along the reference chord of each station (x_{ac}) were used instead of those considered under the flat plate assumption.

For clarity, we adopt the following notation throughout this work: lowercase bold italic symbols are used for vectors and one-dimensional arrays (e.g., $\boldsymbol{a}$); uppercase bold symbols for matrices (e.g., $\boldsymbol{A}$); and non-bold italic symbols for scalars (e.g., a). Moreover, time derivatives are denoted by $\dot{\bullet}$, and $\odot$ represent Hadamard (i.e., element-wise) operations.

According to Theodorsen's theory [12] for an airfoil oscillating in pitch and plunge, the instantaneous lift $\boldsymbol{\ell} \in \mathbb{R}^n$ and pitching moment $\boldsymbol{m} \in \mathbb{R}^n$ per unit span can be split into circulatory ($\bullet^{\mathrm{c}}$) and non-circulatory ($\bullet^{\mathrm{nc}}$) terms as follows:

$$\boldsymbol{\ell}^{\mathrm{nc}} = \pi\rho_\infty \frac{\boldsymbol{c}}{4} \odot \left[\ddot{\boldsymbol{h}} + V_\infty \dot{\boldsymbol{\alpha}} - (\boldsymbol{x}_{c/2} - \boldsymbol{x}_{\mathrm{ea}}) \odot \ddot{\boldsymbol{\alpha}}\right], \tag{1a}$$

$$\boldsymbol{\ell}^{\mathrm{c}} = \frac{1}{2}\rho_\infty V_\infty \boldsymbol{c} \odot \boldsymbol{c}_{l_\alpha} \odot \left(\boldsymbol{w}_{3c/4} + \boldsymbol{\lambda}_1 + \boldsymbol{\lambda}_2\right), \tag{1b}$$

$$\boldsymbol{m}^{\mathrm{nc}} = \pi\rho_\infty \frac{\boldsymbol{c}^{\odot 2}}{4} \odot \left\{V_\infty \dot{\boldsymbol{h}} + (\boldsymbol{x}_{c/2} - \boldsymbol{x}_{\mathrm{ea}}) \odot \ddot{\boldsymbol{h}} + V_\infty^2 \boldsymbol{\alpha} - \left[\frac{\boldsymbol{c}^{\odot 2}}{32} + (\boldsymbol{x}_{c/2} - \boldsymbol{x}_{\mathrm{ea}})^{\odot 2}\right] \odot \ddot{\boldsymbol{\alpha}}\right\}, \tag{1c}$$

$$\boldsymbol{m}^{\mathrm{c}} = -\pi\rho_\infty V_\infty \frac{\boldsymbol{c}^{\odot 2}}{4} \odot \boldsymbol{w}_{3c/4} + (\boldsymbol{x}_{\mathrm{ac}} - \boldsymbol{x}_{\mathrm{ea}}) \odot \boldsymbol{\ell}^{\mathrm{c}}, \tag{1d}$$

where $\rho_\infty \in \mathbb{R}$ is the air density, $V_\infty \in \mathbb{R}$ is the unperturbed airspeed, $\boldsymbol{c} \in \mathbb{R}^n$ is the reference chord length (assumed to be at the middle of the spanwise stations), $\boldsymbol{c}_{\ell,\alpha} \in \mathbb{R}^n$ is the lift curve slope, $\boldsymbol{x}_{\mathrm{c}/2} \in \mathbb{R}^n$ is the distance from the mid-chord point, $\boldsymbol{x}_{\mathrm{ea}} \in \mathbb{R}^n$ is the distance from the elastic axis point, and $\boldsymbol{x}_{\mathrm{ac}} \in \mathbb{R}^n$ is the distance from the aerodynamic center point, all measured along the reference chords. The vertical and angular displacements of the load reference axis (elastic axis) in the undeformed body-attached reference frame are represented by $\boldsymbol{h} \in \mathbb{R}^n$ and $\boldsymbol{\alpha} \in \mathbb{R}^n$, respectively.

The downwash induced at the three-quarter chord, $\boldsymbol{w}_{3\mathrm{c}/4}$, is given by:

$$\boldsymbol{w}_{3c/4} = \dot{\boldsymbol{h}} + V_\infty \boldsymbol{\alpha} - (\boldsymbol{x}_{3c/4} - \boldsymbol{x}_{\mathrm{ea}}) \odot \dot{\boldsymbol{\alpha}}. \tag{2}$$

The aerodynamic lag states $\boldsymbol{\lambda}_1 \in \mathbb{R}^n$ and $\boldsymbol{\lambda}_2 \in \mathbb{R}^n$ in Equation 1b can be evaluated considering Jones' approximation [13] for Wagner's indicial function, such that Duhamel's convolution integral has an analytical solution [14], yielding

$$\dot{\boldsymbol{\lambda}}_1 = -0.082 V_\infty \boldsymbol{c}^{\odot -1} \odot \boldsymbol{\lambda}_1 - 0.165 \dot{\boldsymbol{w}}_{3c/4}, \tag{3a}$$

$$\dot{\boldsymbol{\lambda}}_2 = -0.64 V_\infty \boldsymbol{c}^{\odot -1} \odot \boldsymbol{\lambda}_2 - 0.335 \dot{\boldsymbol{w}}_{3c/4}. \tag{3b}$$

For an elastic wing, the solutions of Equations 1a-1d must be transferred to the undeformed-body reference frame, multiplied by the strip widths $\Delta \boldsymbol{y} \in \mathbb{R}^n$, and summed up in order to obtain the total instantaneous lift $L \in \mathbb{R}$ and the pitching moment around the aerodynamic center $M_y \in \mathbb{R}$, i.e.,

$$L = (\boldsymbol{\ell}^{\mathrm{nc}} + \boldsymbol{\ell}^{\mathrm{c}}) \cdot \Delta \boldsymbol{y}, \tag{4a}$$

$$M_y = (\boldsymbol{m}^{\mathrm{nc}} + \boldsymbol{m}^{\mathrm{c}}) \cdot \Delta \boldsymbol{y}. \tag{4b}$$

The modified strip theory formulation presented in this section provides a computationally efficient, time-domain representation of the unsteady aerodynamic loads acting on a lifting surface. By expressing the aerodynamic lag through state-space differential equations, this physics-based model is well known for its seamless integration into flight dynamics simulations and control frameworks. Having established this mathematical baseline for aerodynamics, the following subsection introduces the Scientific Machine Learning concepts that will be coupled with these governing equations to construct our proposed methodology.

3.2 Scientific Machine Learning

Recent advancements in deep learning have demonstrated unprecedented success in purely data-driven tasks such as computer vision and natural language processing. However, classical "black-box" machine learning models are notoriously data-hungry, struggle to extrapolate beyond their training distributions, and frequently yield predictions that violate fundamental physical principles, such as the conservation of mass or energy.

Scientific Machine Learning has emerged as a transformative paradigm designed to bridge the gap between scientific computing and machine learning. As Figure 3 clearly demonstrates, interest in the field has steadily accelerated since 2017, when foundational publications [15–18] were released. In recent years, however, a spike in interest has been observed, showcasing a transition from a niche concept to a mainstream research necessity. Rather than relying on blind data-fitting, SciML forces deep neural networks to adhere strictly to well-known physical laws, ensuring physical consistency. This directly tackles the critical drawbacks of both parent fields. Injecting domain knowledge restricts the network's hypothesis space and eliminates the dependency on massive datasets. In return, the machine learning architecture bypasses the computational bottlenecks that hinder traditional numerical solvers when dealing with high-dimensional, noisy data.

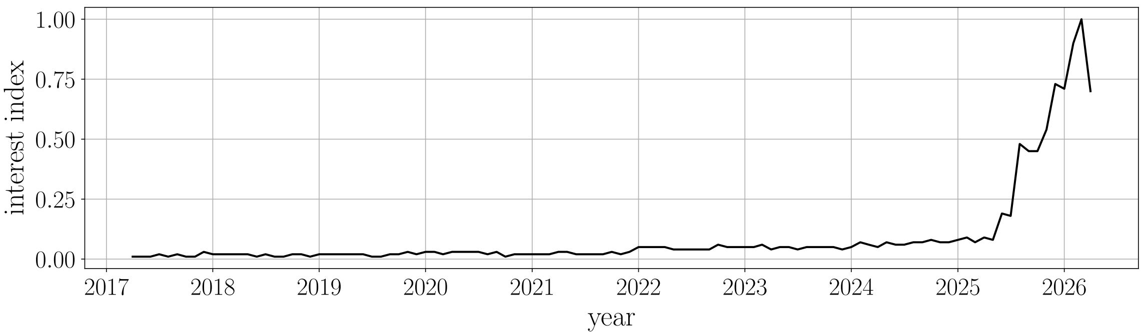

Figure 3: Google search of the term “scientific machine learning”.

Notably, Physics-Informed Neural Networks (PINNs) [15] have fundamentally altered how we approach complex systems by embedding governing equations directly into the neural network's loss function. This allows the network to act as a mesh-free solver capable of tackling both forward and inverse problems. Parallel to PINNs, the development of operator learning — exemplified by DeepONets [19] and Fourier Neural Operators (FNOs) [20] — has shifted the focus from learning solutions in finite-dimensional spaces to learning mappings between infinite-dimensional function spaces. This allows trained models to evaluate solutions orders of magnitude faster than traditional numerical solvers, demonstrating generalization across varying grid resolutions. Furthermore, techniques like the Sparse Identification of Nonlinear Dynamics (SINDy) [21], centered on differentiable programming and symbolic regression, have enabled the automated discovery of governing equations directly from observational data.

Despite its promising potential to revolutionize fields ranging from materials science [22] to climate modeling [23], SciML is not without its challenges. Consequently, researchers are currently grappling with several theoretical and practical hurdles [24–26], which include spectral bias, optimization difficulties in multi-objective loss landscapes, and the pressing need for rigorous uncertainty quantification.

4 METHODOLOGY

We propose a neural network framework based on multi-fidelity data for obtaining a surrogate model to predict the generalized aerodynamic loads (due to lift and pitching moment) acting on the flexible wing of the TU-Flex prototype at a given flight condition in terms of airspeed V_∞ and angle of attack α.

The idea behind our methodology is to combine high-fidelity Computational Fluid Dynamics (Section 4.2) and wind-tunnel data (Section 4.1) to train a neural network using a transfer learning approach. The predictions of this first neural network are then integrated into a second neural network that has a physics-inspired loss function based on quasi-steady strip theory aerodynamics (Section 4.3), in a process described thoroughly in Section 4.4. The training of this second neural network produces a light-weight surrogate model of improved accuracy when compared to the strip theory method employed originally, which can be used in the flight dynamics simulation framework of the TU-Flex.

4.1 Wind-Tunnel Data

A half-span model of the flexible wing, depicted in Figure 4, was built and tested during a wind-tunnel campaign conducted at the DLR facilities in Goettingen, Germany. A detailed discussion of the wind-tunnel testing for the TU-Flex wing can be found in [27]. The data used in this study were collected using a balance attached to the root of the half-wing, which recorded load measurements under several steady flow conditions. Figure 5 illustrates a sketch of the balance and the sign convention for the measured forces and moments.

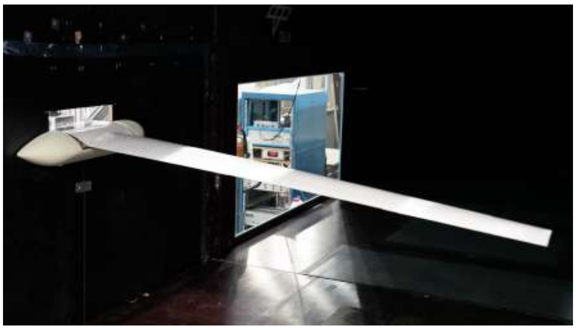

Figure 4: TU-Flex's flexible wing setup at the wind-tunnel. Credits: FMRA, 2024.

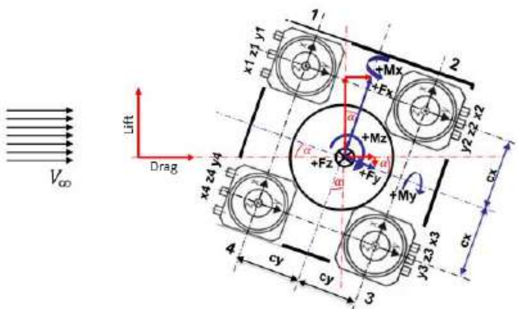

Figure 5: Sketch of the balance used for load data acquisition. Credits: FMRA, 2024.

The steady flow test campaign comprises data acquired during runs in which different combinations of airspeed and angle of attack were imposed on the wing prototype, as listed in Table 2.

Table 2: Nominal steady flow conditions covered during the experimental campaign.

V_∞ [m/s] \ α [deg]	-9	-8	-7	-6	-4	-3	-2	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
18	•	•		•	•	•	•	•		•	•	•		•	•	•	•	•	•	•	•	•	•
22		•		•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
30		•		•	•		•	•	•	•	•	•		•	•	•	•	•					
34			•		•		•	•	•	•	•	•	•	•	•								
40						•	•	•		•	•	•											

For each condition, a certain amount of time was given for the flow to stabilize prior to beginning the data acquisition. The data were then collected for 10 seconds at a frequency of 1000 Hz. This procedure was repeated for every condition listed in Table 2.

The measured quantities are the instantaneous resultant forces and moments in the three axes acting on the root of the wing. Those forces and moments were directly obtained from the balance measurements and were subsequently converted to the aerodynamic axes and corrected to disregard the weight of the structure and to account for balance drifting. The mean and standard deviation for the measured lift and pitching moment are shown in Figure 6.

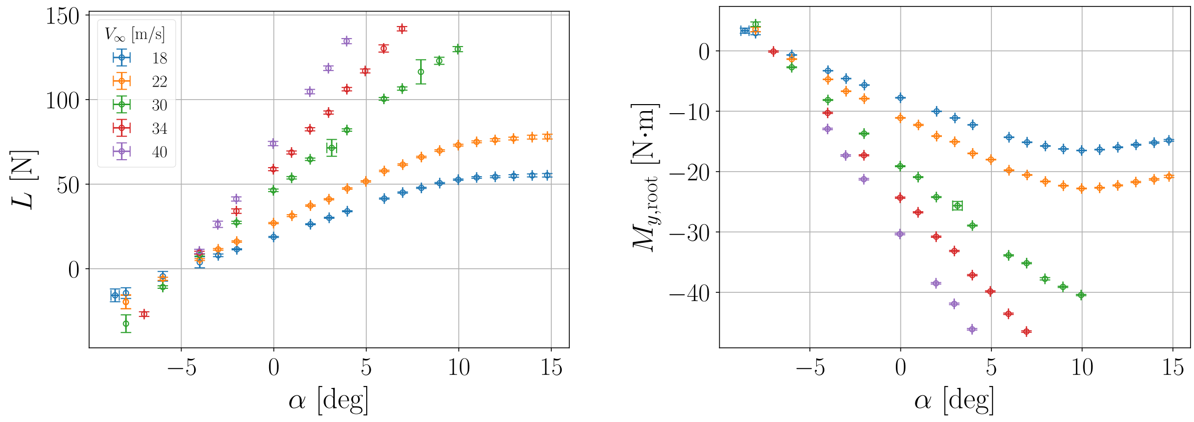

(a)

(b)

Figure 6: Wind-tunnel data – mean and standard deviation – for lift (a) and pitching moment around the root midchord (b).

The resulting mean and standard deviation profiles for both lift and root pitching moment capture the physical behavior of the TU-Flex flexible wing under aerodynamic steady loads, including noticeable non-linearities at higher dynamic pressures. This experimental foundation serves mainly as a benchmark for validating the numerical frameworks presented in subsequent sections.

4.2 Computational Fluid Dynamics Data

Computational Fluid Dynamics (CFD) simulations provided the overall trends for the lift and pitching moment curves, allowing us to evaluate flow regimes that would be impractical or unsafe to test in a wind-tunnel due to the risk of structural damage. The computational mesh was based on the original wing geometry with a minor modification: the semi-circular trailing edge was replaced with a blunt profile to improve aerodynamic modeling and facilitate grid generation.

As introduced in Section 2, the CFD simulations shown in this section were also performed under the responsibility of DLR as part of their collaboration on the TU-Flex project. The computational mesh is composed of a structured surface grid across the upper, lower, and trailing edge surfaces, incorporating boundary-layer cells to ensure a wall resolution of $y^+ \lesssim 1$. The aerodynamic analysis considered steady RANS (Reynolds-Averaged Navier-Stokes) computations with a second-order central spatial discretization coupled with the Menter k-ω SST turbulence model. Furthermore, standard atmospheric conditions were applied, and the flow was assumed to be fully turbulent over the entire wing.

Figure 7 shows the lift and pitching moment around the root midchord of the wing obtained through the simulations.

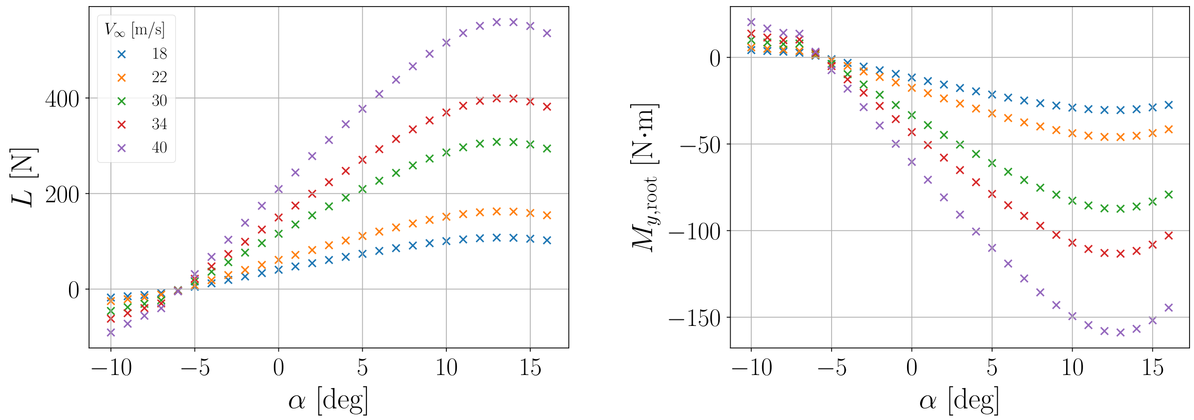

(a) (b)

Figure 7: CFD data for lift (a) and pitching moment around the root midchord (b).

The CFD data provide a reasonable representation of the lift and pitching moment behavior across the considered angles of attack, capturing general trends even in the stall regions where experimental wind-tunnel data could not be obtained. However, notable differences in magnitude exist between the numerical and experimental datasets.

A primary reason for this discrepancy is the assumption of a rigid wing in the CFD simulations, which neglects the structural flexibility of the physical model. In reality, aerodynamic loads cause the physical wing to bend and twist. By omitting this flexibility, the rigid CFD model tends to overestimate the analyzed quantities, especially at higher angles of attack. Moreover, steady RANS computations often struggle to accurately predict the onset of stall, massively separated flows, and unsteady aerodynamic phenomena at high angles of attack. Another reason to consider is that even with corrections, residual wind-tunnel wall interference, blockage effects, or mounting sting/balance interference can cause deviations from the conditions simulated in the CFD environment.

4.3 Strip theory Data

The FW configuration allows for wingtip displacements in the range where geometrically linear structural theory holds; therefore, a modal superposition approach is used to estimate the elastic displacements $\boldsymbol{v} \in \mathbb{R}^{2n} \mid \boldsymbol{v} := [\boldsymbol{h} \quad \boldsymbol{\alpha}]^\top$ in terms of vertical $\boldsymbol{h} \in \mathbb{R}^n$ and angular $\boldsymbol{\alpha} \in \mathbb{R}^n$ displacements of the load reference axis (elastic axis) in the undeformed body-attached reference frame, such that

$$\boldsymbol{v} = \boldsymbol{\Phi}\, \boldsymbol{\eta}, \tag{5}$$

where, $\boldsymbol{\eta} \in \mathbb{C}^m$ is the vector of modal amplitudes, and $\boldsymbol{\Phi} \in \mathbb{C}^{2n\times m} \mid \boldsymbol{\Phi} := [\boldsymbol{\Psi}^\top \quad -\boldsymbol{\Gamma}^\top]^\top$ is the modal matrix (reduced to m modes) of the free and undamped structure, comprised of $\boldsymbol{\Psi} \in \mathbb{C}^{n\times m}$ out-of-plane bending modes and $\boldsymbol{\Gamma} \in \mathbb{C}^{n\times m}$ torsional modes. The negative sign for $\boldsymbol{\Gamma}$ ensures the convention of positive $\boldsymbol{\alpha}$ for pitching up. Note that we have chosen the usual angle of attack notation for the angular displacements, as they, in fact, represent the local geometric angles of attack at a given time instant.

In this work, we followed one of the aerodynamic models proposed in [2] for the TU-Flex prototype. It assumes a quasi-steady approach in which incremental aerodynamic loads due to the elastic deformations of the structure are calculated, while temporal memory effects on those are neglected. The steady aerodynamics come straightforwardly from the equations 1a-1d, since all time-dependent terms vanish, yielding the following simple expressions:

$$\boldsymbol{\ell} = \frac{1}{2}\,\rho_\infty\, V_\infty^2\, \boldsymbol{c} \odot \boldsymbol{C}_{\ell\alpha} \odot \boldsymbol{\alpha}, \tag{6a}$$

$$\boldsymbol{m} = (\boldsymbol{x}_{\mathrm{ac}} - \boldsymbol{x}_{\mathrm{ea}}) \odot \boldsymbol{\ell}. \tag{6b}$$

The generalized forces $\boldsymbol{Q}_{\eta,k} \in \mathbb{R}^n$ for a given mode k, as demonstrated in [8], are calculated as follows:

$$\boldsymbol{Q}_{\eta,k} = \int_{\forall} \boldsymbol{f} \cdot \frac{\partial \boldsymbol{v}}{\partial \eta_k}\, d\forall, \tag{7}$$

where $\boldsymbol{f}$ is the distributed load vector per unit volume $\forall$. The previous equation can be conveniently written as:

$$\boldsymbol{Q}_\eta = \boldsymbol{\Phi}^\top \boldsymbol{f}. \tag{8}$$

After some manipulation and assuming that only lift and pitching moment around the root midchord are the relevant loads, the generalized modal forces can be approximated as follows:

$$\boldsymbol{Q}_\eta \approx -\left[\boldsymbol{\ell} \odot \cos(\boldsymbol{\alpha}) \odot \Delta \boldsymbol{y}\right] \boldsymbol{\Psi} - \left[\boldsymbol{m} \odot \Delta \boldsymbol{y} + \boldsymbol{\ell} \odot (x_{\mathrm{bal}} - \boldsymbol{x}_{\mathrm{ea}}) \odot \Delta \boldsymbol{y}\right] \boldsymbol{\Gamma}. \tag{9}$$

Equation 9 considers the distance from the elastic axis of each strip to the balance position at the root x_{bal} and was solved accounting for the first four out-of-plane bending modes in $\boldsymbol{\Psi}$ and four torsional modes in $\boldsymbol{\Gamma}$, along with the same flow conditions as those in the CFD and experimental data. A dataset was then built consisting of each combination of α and V_∞, along with the respective $\boldsymbol{\eta}$ as inputs and $\boldsymbol{Q}_\eta$ as the output.

4.4 Neural Network Framework

The proposed framework consists of two distinct neural networks designed to effectively fuse the available multi-fidelity datasets. The first network, designated as Model I, is a baseline feedforward architecture that undergoes a two-step training process. In the initial step, the network is pre-trained to fit the general aerodynamic trends derived from the CFD data described in Section 4.2. In the second step, a transfer learning technique is applied to adapt this pre-existing knowledge to the experimental domain. By freezing the weights and biases of the network's first hidden layer, the model retains the generalized lower-layer features learned from the simulation data. The remaining upper layers are then fine-tuned using the wind-tunnel dataset detailed in Section 4.1.

The rationale for this sequential approach is to first capture the broader physical behavior of the lift and pitching moment across various flow velocities, subsequently adjusting the network to match the real magnitudes observed during the experimental campaign. This methodology leverages the benefits of transfer learning by achieving faster convergence and higher generalization accuracy when working with constrained target datasets [28]. A visual schematic of this first neural network architecture is shown in Figure 8.

The second neural network, designated as Model II and shown in Figure 8, shares a similar feedforward architecture with Model I. However, instead of mapping inputs purely to experimental data outputs, Model II is formulated to minimize a hybrid loss function containing a physics-inspired regularization term. Model II maps the input flight conditions and generalized coordinates to a predicted set of generalized modal forces, $\hat{\boldsymbol{Q}}_\eta$.

We assume that the generalized modal forces due to aerodynamic effects are governed by Equation 8. Rather than computing the lift and pitching moment using the analytical lower-fidelity representations of Equations 6a and 6b, we substitute them with the high-fidelity distributions estimated by Model I. After that, we compare them with the predictions of the generalized forces from Model II. Through this configuration, the distributed lift and pitching moment can be estimated from the generalized modal forces via:

$$\begin{bmatrix} \hat{\boldsymbol{\ell}} \odot \Delta \boldsymbol{y} \\ \hat{\boldsymbol{m}} \odot \Delta \boldsymbol{y} \end{bmatrix} = \left(\boldsymbol{\Phi}^\top\right)^{-1} \hat{\boldsymbol{Q}}_\eta, \tag{10}$$

where $\left(\boldsymbol{\Phi}^\top\right)^{-1}$ is the Moore-Penrose pseudoinverse matrix because the reduced modal matrix is non-square. From Equation 10, we integrate the spanwise distributions to estimate the total lift and pitching moment at the root midchord of the wing:

$$\left.\begin{bmatrix} \hat{L} \\ \hat{M}_{y,\text{root}} \end{bmatrix}\right|_{\text{model II}} = \begin{bmatrix} \sum \hat{\boldsymbol{\ell}} \odot \Delta \boldsymbol{y} \odot \cos(\boldsymbol{\alpha}) \\ \sum \left[\hat{\boldsymbol{m}} \odot \Delta \boldsymbol{y} + \hat{\boldsymbol{\ell}} \odot \Delta \boldsymbol{y} \odot (x_{\text{bal}} - \boldsymbol{x}_{\text{ea}})\right] \end{bmatrix} \tag{11}$$

The loss function at a given data point k to be minimized during the training of Model II is composed of a data-driven term and a physics-based residual term:

$$\mathscr{L}_k = \left.\sum_{j=1}^{m} \left(\hat{\boldsymbol{Q}}_{\eta,j} - \boldsymbol{Q}_{\eta,j}\right)^2\right|_k + \left.\left(\left.\begin{bmatrix} \hat{L} \\ \hat{M}_{y,\text{root}} \end{bmatrix}\right|_{\text{model II}} - \left.\begin{bmatrix} \hat{L} \\ \hat{M}_{y,\text{root}} \end{bmatrix}\right|_{\text{model I}}\right)^2\right|_k, \tag{12}$$

where the total loss $\mathscr{L}$ for a given batch of size n_{batch} is defined as:

$$\mathscr{L} = \frac{1}{n_{\text{batch}}} \sum_{k=1}^{n_{\text{batch}}} \left(\mathscr{L}_{k,\text{data}} + \mathscr{L}_{k,\text{physics}}\right). \tag{13}$$

The final objective of Model II is to predict a set of generalized modal forces, $\hat{\boldsymbol{Q}}_\eta$, that simultaneously satisfy the high-fidelity integrated lift and pitching moment targets generated by Model I while maintaining consistency with the baseline experimental generalized forces calculated via the strip theory method described in Subsection 4.3.

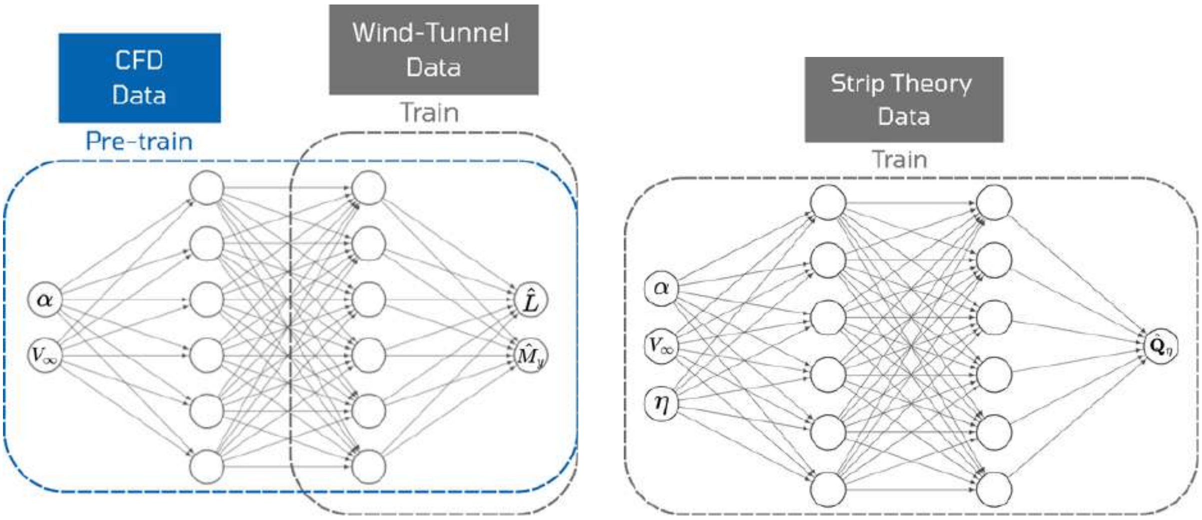

(a) Model I

(b) Model II

Figure 8: Architecture layouts for the proposed neural networks.

During the training phase of both networks, 80% of the available data points was allocated for model optimization, with the remaining 20% reserved for validation tracking. The explicit hyperparameters selected for each architecture are detailed in Table 3. These hyperparameters were chosen from a preliminary parametric sweep to ensure stable loss convergence and to prevent overfitting across both data regimes.

Table 3: Hyperparameters considered for the neural network models, including pre-training.

Hyperparameter	**Model I**	**Model II**
Hidden layers	2	2
Neurons per layer	128	128
Batch size	16 (pre), 16	128
Optimizer	Adam	Adam
Epochs	3×10^4 (pre), 3×10^4	3×10^5
Learning rate	1×10^{-4}	1×10^{-5}
Activation function	$\tanh$	$\tanh$

5 RESULTS

The methodology proposed in Section 4.4 was implemented within a Python framework, and the resulting loss curves for each neural network model are presented in Figure 9. Model I exhibits a two-stage convergence curve characteristic of its training sequence: the initial phase reflects pre-training on the CFD data, followed by a visible transition when the secondary phase begins using the wind-tunnel dataset.

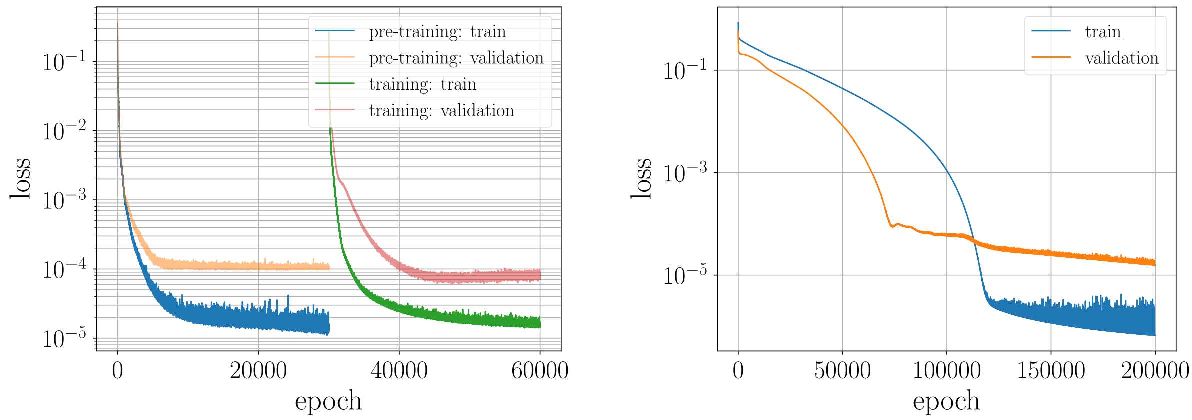

(a) Model I

(b) Model II

Figure 9: Loss history of the training step of the proposed models.

The core objective of Model I is to combine the aerodynamic trends of the numerical simulations with the physical magnitudes observed in the wind-tunnel. Consequently, when extrapolating beyond the experimental data points of Figure 6, the model is expected to capture and mirror the physical behavior established by the CFD data shown in Figure 7. This predicted multi-fidelity alignment is confirmed by the results illustrated in Figure 10.

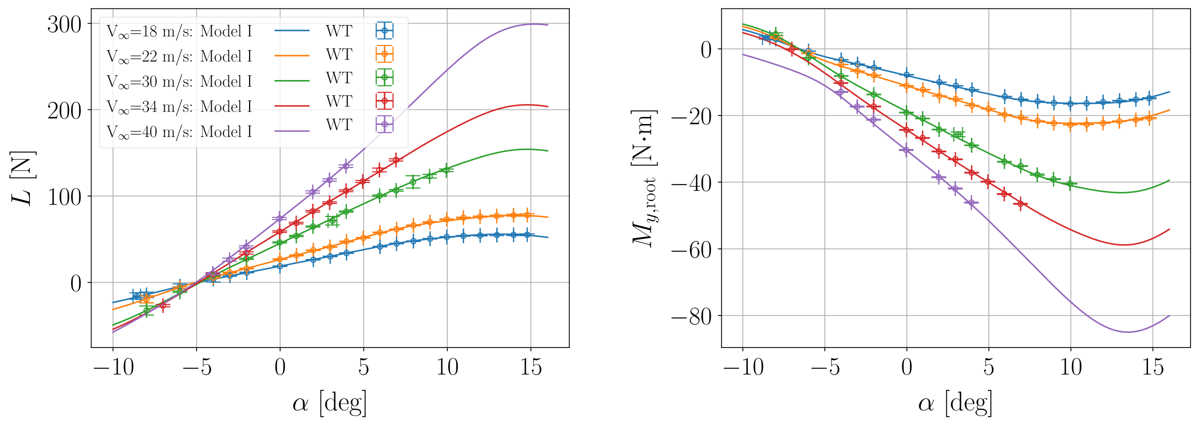

(a)

(b)

Figure 10: Model I predictions for lift (a) and pitching moment around the root midchord (b).

Model I acts as a high-fidelity baseline to parameterize the physics-informed loss function of Model II. Once trained, Model II predicts the generalized aerodynamic modal forces, allowing the continuous lift and pitching moment distributions to be recovered via Equation 10 along the discretized spanwise stations. This recovered behavior is illustrated in Figure 11, which displays the predicted lift distribution across the non-dimensionalized span at a representative angle of attack of $5°$.

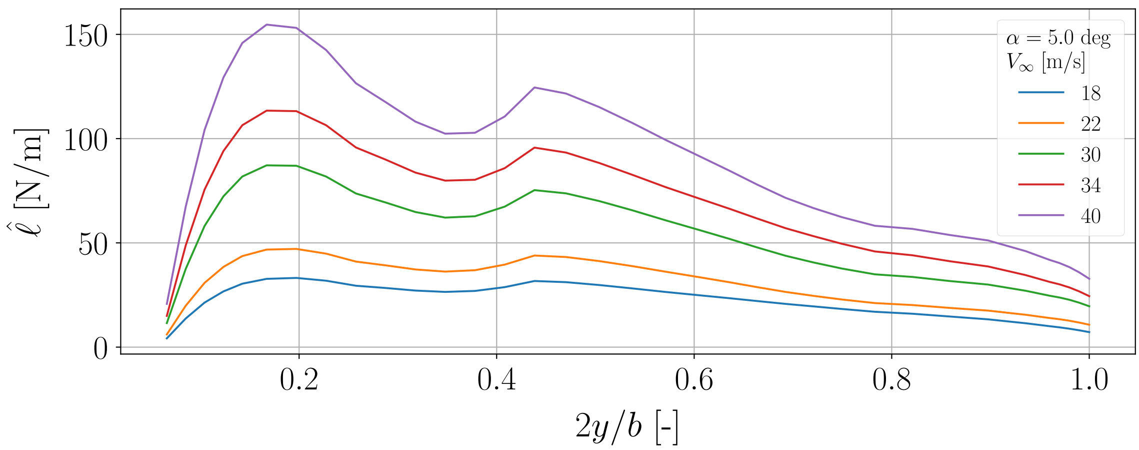

Figure 11: Lift distribution predicted by Model II along the non-dimensionalized span at an angle of attack of 5°.

By integrating these predicted spanwise distributions ($\hat{\boldsymbol{\ell}}$ and $\hat{\boldsymbol{m}}$), the total lift and root pitching moment are computed using Equation 11. The resulting integrated solutions are depicted in Figure 12.

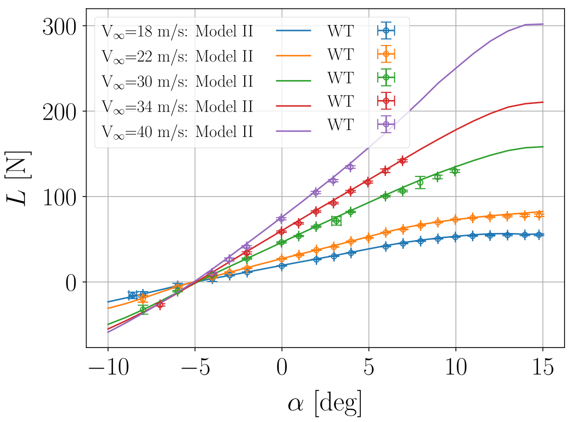

(a)

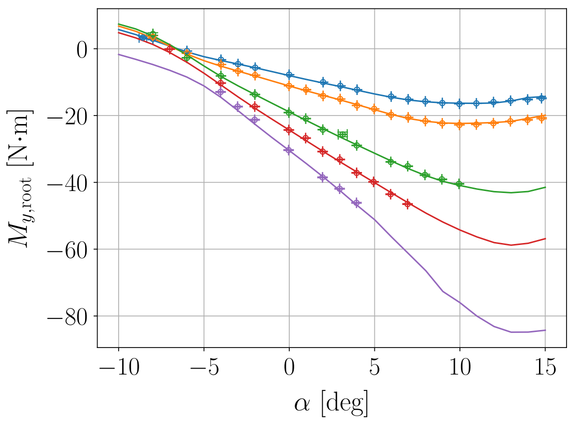

(b)

Figure 12: Model II predictions for lift (a) and pitching moment around the root midchord (b).

The proposed framework successfully predicts generalized aerodynamic modal forces that align with the baseline strip theory method while simultaneously integrating the higher-fidelity total lift and pitching moment constraints provided by Model I. Nevertheless, certain limitations of the current methodology should be highlighted. First, the resolution of the results is limited by the discretization considered for the strip theory model, and the prediction accuracy cannot be guaranteed at higher angles of attack. Even though the CFD data provide the general behavior of the aerodynamic forces, the experimental data were unavailable at higher angles of attack across several of the evaluated airspeeds; thus, the predictions in those regions could not be validated. An additional limitation is that the Model II spanwise lift predictions near the wing root deviate from what is usually expected, likely due to boundary effects. Finally, regarding computational efficiency, further testing is required to ensure the framework's inference speed is sufficient for real-time flight simulations. Addressing this performance constraint is the immediate next step of our research, which aims to implement Model II as an aerodynamic surrogate model embedded within the TU-Flex flight dynamics simulation code.

6 CONCLUSION

In this work, a multi-fidelity Scientific Machine Learning (SciML) framework was successfully developed and implemented to enhance the accuracy of baseline strip theory aerodynamic predictions for the TU-Flex flying demonstrator. By pairing a sequential transfer learning approach with physics-informed regularization, the proposed framework effectively leverages the trends captured by CFD simulations alongside data collected during a wind-tunnel testing campaign.

The results demonstrate that the pre-trained neural network layout (Model I) robustly bridges the discrepancy in magnitude between the numerical model and experimental observations. This multi-fidelity data mixing serves as a higher-fidelity parameter for the physics-informed loss landscape of the subsequent architecture (Model II). Model II successfully acts as a continuous operator, mapping operational flight conditions to quasi-steady aerodynamic modal forces.

Despite the predictive performance of the network under nominal regimes, key limitations of the methodology have been identified. Because wind-tunnel data were restricted across selected airspeeds, model confidence cannot be guaranteed near stall regions. Furthermore, small localized deviations from conventional aerodynamic distributions observed near the wing root highlight a need for better spatial training arrays or modified boundary condition enforcement within the hybrid loss function.

Future investigations will focus on executing thorough computational benchmarking to verify the execution speed of Model II. The next primary phase of this study will be embedding this trained neural network configuration as an efficient aerodynamic surrogate model within the TU-Flex flight dynamics simulation code. Furthermore, while the current architecture effectively captures quasi-steady loads, future iterations of this research aim to extend the methodology to unsteady aerodynamic phenomena. Transitioning to these complex dynamic regimes remains challenging, as it relies heavily on the availability of highly time-resolved experimental or numerical datasets, which are difficult and resource-intensive to acquire. Solving these data acquisition and computational challenges promises to pave the way for real-time system identification, flight dynamics, and aeroelastic control architectures across highly flexible aircraft designs.

ACKNOWLEDGMENTS

This study was financed in part by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES) - finance code 001; by Finep and Embraer S.A. under the research project “Advanced Studies in Flight Physics and Control” - contract 01.22.0552.00; and by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) due to a resolution of the German Federal Parliament within the scope of the LuFo VI-1 project “FlexFuture” - grant number 20E2120A.

7 REFERENCES

[1] Santos, V., Vieira, B., Cardoso Ribeiro, F., et al. (2024). Surrogate modeling of highly flexible structures and aerodynamics using neural networks. In *IFASD 2024 - International Forum on Aeroelasticity and Structural Dynamics*. The Hague, The Netherlands, p. 20.

[2] Quesada, Á., González, P., Barbosa, G., et al. (2024). Influence of nonlinear aerodynamic effects on high aspect ratio aircraft model. In *IFASD 2024 - International Forum on Aeroelasticity and Structural Dynamics*. The Hague, The Netherlands, p. 26.

[3] González, P., Stavorinus, G., Silvestre, F., et al. (2023). TU-Flex: A very-flexible flying demonstrator with a generic transport aircraft configuration. In *AIAA SCITECH 2023 Forum*. doi:10.2514/6.2023-1312.

[4] Gonzáles, P. J., Stavorinus, G., Sahin, H., et al. (2022). A structural preliminary design of a flexible flying demonstrator. In *ICAS 2022 - Congress of the International Council of the Aeronautical Sciences*. ISBN 978-171387116-3, p. 18.

[5] González, P. J., Barbosa, G. C., Quesada, A. A., et al. (2026). Ground experimental tests and flight test design for a flexible aircraft flying demonstrator. In *IFASD 2026 - International Forum on Aeroelasticity and Structural Dynamics,*. Goettingen, Germany.

[6] Barbosa, G. C., Quesada, A. A., González, P. J., et al. (2024). Aeroservoelastic analysis and load alleviation control for very flexible aircraft: TU-Flex study. In *IFASD 2024 - International Forum on Aeroelasticity and Structural Dynamics,*. The Hague, The Netherlands.

[7] Stavorinus, G., González, P. J., Barbosa, G. C., et al. (2026). Aeroelastic model update for a flexible wing using experimental data. In *IFASD 2026 - International Forum on Aeroelasticity and Structural Dynamics,*. Goettingen, Germany.

[8] Silvestre, F. and Paglione, P. (2008). Dynamics and control of a flexible aircraft. In *AIAA Atmospheric Flight Mechanics Conference and Exhibit*. p. 18. doi:10.2514/6.2008-6876.

[9] Silvestre, F. J. and Luckner, R. (2015). Experimental validation of a flight simulation model for slightly flexible aircraft. *AIAA Journal*, 53(12), 3620–3636. doi:10.2514/1.J054023.

[10] Dowell, E., Clark, R., Cox, D., et al. (2004). *A Modern Course in Aeroelasticity*. Solid Mechanics and its Applications. Kluwer Academic Publishers. ISBN 9781402020391.

[11] Yates, E. C. (1966). Modified-strip-analysis method for predicting wing flutter at subsonic to hypersonic speeds. *Journal of Aircraft*, 3(1), 25–29. doi:10.2514/3.43702.

[12] Theodorsen, T. (1934). General theory of aerodynamic instability and the mechanism of flutter. *NACA Report 496*, 413–433.

[13] Jones, R. T. (1938). Operational treatment of the nonuniform-lift theory in airplane dynamics.

[14] Bismarck-Nasr, M. N. (1999). *Structural Dynamics in Aeronautical Engineering*. AIAA Education Series. American Institute of Aeronautics and Astronautics, Incorporated. ISBN 9781600861178.

[15] Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2017). Physics informed deep learning (part I): data-driven solutions of nonlinear partial differential equations. *CoRR*, abs/1711.10561.

[16] Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2017). Physics informed deep learning (part II): data-driven discovery of nonlinear partial differential equations. *CoRR*, abs/1711.10566.

[17] Chen, T. Q., Rubanova, Y., Bettencourt, J., et al. (2018). Neural ordinary differential equations. *CoRR*, abs/1806.07366.

[18] Karpatne, A., Watkins, W., Read, J. S., et al. (2017). Physics-guided neural networks (PGNN): an application in lake temperature modeling. *CoRR*, abs/1710.11431.

[19] Lu, L., Jin, P., and Karniadakis, G. E. (2019). Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *CoRR*.

[20] Li, Z., Kovachki, N. B., Azizzadenesheli, K., et al. (2020). Fourier neural operator for parametric partial differential equations. *CoRR*.

[21] Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15), 3932–3937. doi:10.1073/pnas.1517384113.

[22] Fransen, M., Fürst, A., Tunuguntla, D., et al. (2026). Towards scientific machine learning for granular material simulations: Challenges and opportunities. *Archives of Computational Methods in Engineering*, 33(1), 789–821. doi:10.1007/s11831-025-10322-8.

[23] Pani, T. T., Joshi, P. D., Dandekar, R., et al. (2025). CLIMASIM — climate simulation with scientific machine learning. In *Second Workshop on XAI4Science: From Understanding Model Behavior to Discovering New Scientific Knowledge*. p. 28.

[24] Xu, C., Liu, D., Nassereldine, A., et al. (2025). FP64 is all you need: Rethinking failure modes in physics-informed neural networks. *arXiv*, abs/2505.10949.

[25] Naser, M. (2026). Fundamental flaws of physics-informed neural networks and explainability methods in engineering systems. *Computers & Industrial Engineering*, 212, 111704. ISSN 0360-8352. doi:https://doi.org/10.1016/j.cie.2025.111704.

[26] Rathore, P., Lei, W., Frangella, Z., et al. (2024). Challenges in training PINNs: A loss landscape perspective. In *41st International Conference on Machine Learning*. Vienna, Austria, p. 33.

[27] González, P., Barbosa, G., Quesada, Á., et al. (2024). Wind tunnel testing and modal validation of TU-Flex's high aspect-ratio wings. In *IFASD 2024 - International Forum on Aeroelasticity and Structural Dynamics*. The Hague, The Netherlands, p. 25.

[28] Zhuang, F., Qi, Z., Duan, K., et al. (2019). A comprehensive survey on transfer learning. *CoRR*, abs/1911.02685.

COPYRIGHT STATEMENT