

PHYSICS-INFORMED MACHINE LEARNING FOR FLIGHT DYNAMICS IDENTIFICATION OF THE SAAB 340B USING FLIGHT TEST DATA

Angel Navarro[1], Mushfiqul Alam[1], Alex Elliott[1], & Edward Jinks[2]
[1] Centre for Aeronautics, Power and Propulsion, Faculty of Engineering and Applied Science (FEAS), Cranfield University, United Kingdom
angel.navarrogil@cranfield.ac.uk, mushfiqul.alam@cranfield.ac.uk, alex.j.elliott@cranfield.ac.uk
[2] Siemens Digital Industries Software, Siemens UK, Bristol, United Kingdom
edward.jinks@siemens.com

Abstract

This paper presents a physics-informed machine-learning framework for identifying the short-period longitudinal dynamics of the Cranfield University Saab 340B aircraft using flight-test data. The proposed approach combines sequential measurements of pitch rate, vertical velocity and elevator deflection with static flight-condition and mass-property inputs, including dynamic pressure, aircraft mass and pitch moment of inertia. A hybrid neural-network architecture is adopted, in which Long Short-Term Memory layers process the time-series data, while fully connected layers process the static inputs. Rather than directly predicting the aircraft response, the neural network estimates the coefficients of a reduced-order short-period state-space model. The identified coefficients are then used to simulate the pitch-rate and vertical-velocity responses and to compute the associated short-period damping ratio and natural frequency. Physical consistency is introduced through a loss function based on the residuals of the reduced-order equations of motion, with additional penalty terms used to discourage non-physical modal behaviour. The framework is evaluated using Saab 340B short period flight-test manoeuvres and assessed using correlation factors and root mean square errors. Manual hyperparameter tuning improves the baseline model by 20.2%, while Siemens HEEDS-based optimisation using the SHERPA algorithm achieves a 29.3% improvement relative to the same baseline. The HEEDS-optimised model provides the best overall agreement with the flight-test data, giving improved correlation factors and reduced prediction errors for both pitch rate and vertical velocity. The results demonstrate that physics-informed machine learning can provide a physically interpretable approach to aircraft system identification, while automated optimisation can improve model accuracy and reduce manual tuning effort.

Keywords: physics-informed machine learning, aircraft system identification, flight dynamics, short-period mode, flight-test data, LSTM neural networks, hyperparameter optimisation, Saab 340B.

1 Introduction

Aircraft system identification is the process of deriving mathematical models of aircraft dynamic behaviour from measured response data. It is a central activity in flight dynamics because the resulting models support flight-control development, simulation model validation, handling-qualities assessment, flight-test analysis and safety evaluation [1], [2]. Conventional approaches, including equation-error, output-error, frequency-domain and maximum-likelihood methods, remain effective when the model structure is well defined and the available data provide sufficient excitation of the relevant aircraft modes [1]. Modern aircraft modelling tasks increasingly require models that remain accurate across wider regions of the flight envelope. Variations in airspeed, altitude, aircraft mass, inertia and configuration can alter the identified dynamics, while nonlinear effects may become more significant away from nominal operating conditions [3], [4]. Developing local models at individual flight conditions can therefore become inefficient, while global modelling approaches require sufficient data coverage and careful model-structure selection.

Machine-learning methods offer a flexible data-driven route for modelling complex dynamic systems. Neural networks can learn nonlinear relationships from data and can be structured to process both time histories and operating-condition variables [5]. However, purely data-driven models may lack physical interpretability and can extrapolate poorly outside the training data. For aerospace

https://doi.org/10.71945/icas2026_0686

applications, where stability, reliability and physical consistency are essential, this is a significant limitation [6], [7].

Physics-informed machine learning addresses this limitation by embedding physical constraints within the learning process. In the present work, a neural network is not trained as a black-box response predictor. Instead, it is trained to estimate the coefficients of a reduced-order short-period state-space model. This allows the identified model to retain direct physical meaning through its stability and control derivative terms, while still benefiting from neural-network-based learning.

This paper applies the approach to the short-period longitudinal dynamics of the Cranfield University Saab 340B aircraft operated by the National Flying Laboratory Centre, shown in Figure 1. The main contributions are: first, the development of a physics-informed neural-network framework for Saab 340B short-period system identification; second, the use of a hybrid LSTM and fully connected architecture to combine sequential flight-test data with static flight-condition and mass-property inputs; third, the estimation of physically meaningful reduced-order model coefficients rather than direct black-box response prediction; and fourth, the comparison of manual and Siemens HEEDS-based hyperparameter optimisation [8].



Figure 1. Cranfield University Saab 340B aircraft operated by the National Flying Laboratory Centre [photo credit: Keith Wilson SFB Photographic].

The remainder of this paper is organised as follows. Section 2 describes the Saab 340B flight-test dataset and preprocessing procedure. Section 3 presents the reduced-order short-period model and the physics-informed neural-network formulation. Section 4 outlines the training, validation and hyperparameter optimisation procedures. Section 5 presents the results and discussion, and Section 6 summarises the main conclusions and future work.

2 Flight-Test Data and Preprocessing

2.1 Saab 340B Flight-Test Dataset

The flight-test data used in this study were obtained from the Cranfield University Saab 340B aircraft operated by the National Flying Laboratory Centre. This work focuses on short-period longitudinal manoeuvres, in which elevator excitation produces measurable pitch-rate (q) and vertical-velocity (w) responses. The available cases cover a range of airspeed and altitude conditions, are shown in Figure 2.

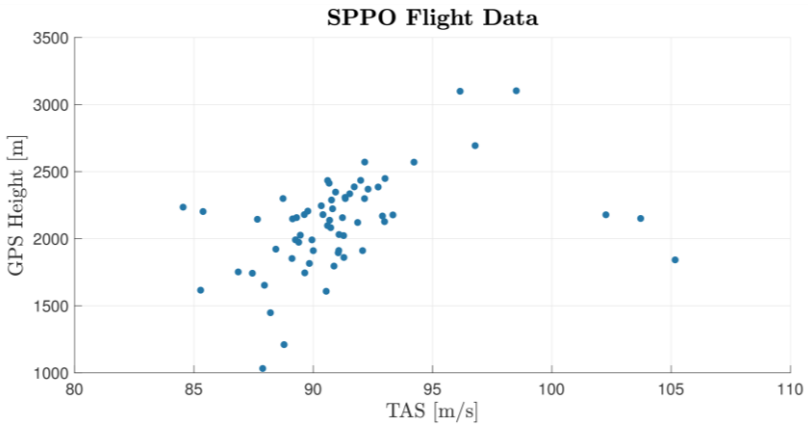
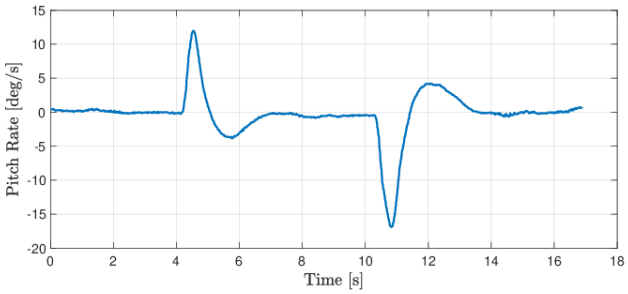


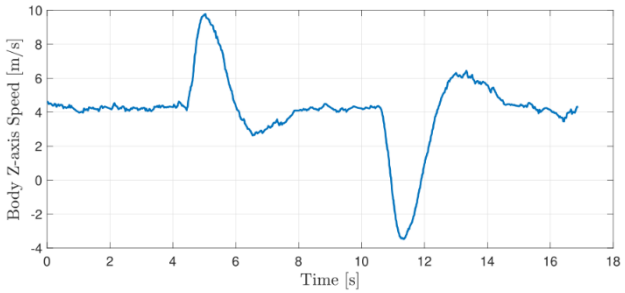
Figure 2. Available Saab 340B short period flight-test points across the test envelope, shown in terms of airspeed and altitude.

https://doi.org/10.71945/icas2026_0686

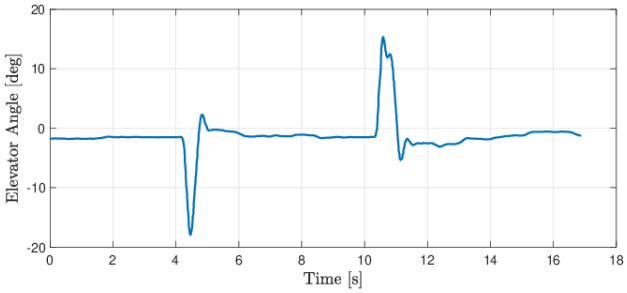
Each flight-test case is represented using two groups of inputs. The sequential inputs are the time histories associated with the short-period response: pitch rate, vertical velocity and elevator deflection. Representative examples of these signals are shown in Figure 3. The static inputs describe the flight condition and aircraft mass properties associated with each manoeuvre. These are dynamic pressure ($\bar{q}$), aircraft mass (m) and pitch moment of inertia (I_{yy}).



(a) Pitch rate (q)



(b) Vertical velocity (w)



(c) Elevator input (δe)

Figure 3. Representative Saab 340B short period flight-test time histories.

Dynamic pressure is used to combine the effect of air density and aircraft speed:

$$\bar{q} = \frac{1}{2}\rho U_e^2 \tag{1}$$

where $\rho(h)$ is the air density at altitude (h), and (U_e) is the forward trimmed aircraft speed. Aircraft mass and pitch moment of inertia are treated as constant during each manoeuvre because of the short duration of the excitation.

The raw flight-test records were pre-processed before neural-network training. The preprocessing sequence consisted of trimming, resampling, filtering and scaling. Each time history was trimmed to retain the pre-manoeuvre trimmed condition and the subsequent short-period response. The signals were then resampled to a uniform time step of $\Delta t = 0.05\ s$. A fourth-order Butterworth low-pass filter with a cut-off frequency of $f_c = 1\ Hz$ was applied to reduce high-frequency measurement noise while retaining the dominant short-period response.

The sequential variables were scaled to the range $[-1,1]$, while the static variables were scaled to the range $[0,1]$. The preprocessing procedure was checked by comparing the processed and raw signals before scaling. The mean correlation factors were 0.9990 for pitch rate (q), 0.9998 for vertical velocity (w) and 0.9890 for elevator deflection (δe), indicating that the preprocessing retained the dominant dynamic content of the manoeuvres.

3 Physics-Informed Identification Method

The longitudinal dynamics of an aircraft can be represented using a state-space model of the form [2]:

$$\dot{x}(t) = Ax(t) + Bu(t), \qquad y(t) = Cx(t) + Du(t) \tag{2}$$

where $x(t)$ is the state vector, $u(t)$ is the input vector, $y(t)$ is the output vector, and A, B, C, and D are the state-space matrices. For the full longitudinal model, the state variables include the body-axis velocity perturbations, pitch rate, and pitch attitude. The corresponding inputs include elevator deflection and throttle setting.

For the short-period approximation, the dominant states are vertical velocity and pitch rate, while elevator deflection is retained as the control input [2]. The reduced-order model used in this work is

https://doi.org/10.71945/icas2026_0686

$$\begin{bmatrix}\dot{w}\\ \dot{q}\end{bmatrix}=\begin{bmatrix}z_w & z_q\\ m_w & m_q\end{bmatrix}\begin{bmatrix}w\\ q\end{bmatrix}+\begin{bmatrix}z_{\delta e}\\ m_{\delta e}\end{bmatrix}\delta e \tag{3}$$

where w is the vertical velocity, q is the pitch rate, δe is the elevator deflection, and z_w, z_q, m_w, m_q, $z_{\delta e}$, and $m_{\delta e}$ are the reduced-order stability and control derivative terms to be identified. The variables w, q, and δe are treated as perturbation quantities relative to the pre-manoeuvre trimmed condition. This model provides a computationally efficient representation of the short-period dynamics while retaining direct physical meaning [2]. The elements in the above matrix can be written as:

$$z_q=\frac{\mathring{Z}_q+mU_e}{m-\mathring{Z}_{\dot{w}}} \qquad z_w=\frac{\mathring{Z}_w}{m-\mathring{Z}_{\dot{w}}}$$

$$m_q=\frac{\mathring{M}_q}{I_{yy}}+\frac{(\mathring{Z}_q+mU_e)\mathring{M}_{\dot{w}}}{I_{yy}(m-\mathring{Z}_{\dot{w}})} \qquad m_w=\frac{\mathring{M}_w}{I_{yy}}+\frac{\mathring{Z}_w\mathring{M}_{\dot{w}}}{I_{yy}(m-\mathring{Z}_{\dot{w}})} \tag{4}$$

$$z_{\delta e}=\frac{\mathring{Z}_{\delta e}}{m-\mathring{Z}_{\dot{w}}} \qquad m_{\delta e}=\frac{\mathring{M}_{\delta e}}{I_{yy}}+\frac{\mathring{M}_{\dot{w}}\mathring{Z}_{\delta e}}{I_{yy}(m-\mathring{Z}_{\dot{w}})}$$

The parameter m denotes the aircraft mass, U_e is the equilibrium forward speed, and I_{yy} is the aircraft pitch moment of inertia about the body y-axis. The quantities $\mathring{Z}_w$, $\mathring{Z}_q$, and $\mathring{Z}_{\dot{w}}$ are the dimensional Z-force derivatives with respect to vertical velocity, pitch rate, and vertical acceleration, respectively. Likewise, $\mathring{M}_w$, $\mathring{M}_q$, and $\mathring{M}_{\dot{w}}$ are the corresponding dimensional pitching-moment derivatives. The terms involving $\mathring{Z}_{\dot{w}}$ and $\mathring{M}_{\dot{w}}$ arise from the unsteady aerodynamic dependence on $\dot{w}$, which is rearranged to obtain the reduced-order short-period model in standard state-space form. The quantity $\mathring{Z}_{\delta e}$ is the dimensional Z-force derivative with respect to the elevator/control input, while $\mathring{M}_{\delta e}$ is the corresponding dimensional pitching-moment derivative. The parameter m denotes the aircraft mass, and I_{yy}, denotes the aircraft pitch moment of inertia about the body y-axis. These expressions can also be found in Appendix B of Ref [2]

The proposed hybrid neural-network architecture is shown in Figure 4. The sequential branch uses Long Short-Term Memory layers to learn time-dependent features from the pitch-rate, vertical-velocity and elevator-deflection histories. The static branch uses fully connected layers to process dynamic pressure, mass and pitch moment of inertia. The outputs of the two branches are concatenated and passed through a final fully connected layer, which outputs the six reduced-order model coefficients.

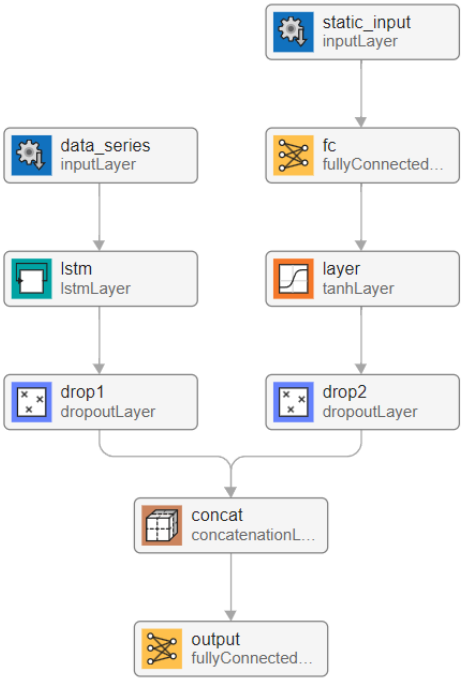


Figure 4. Hybrid physics-informed neural-network architecture. The LSTM branch processes sequential flight-test measurements, while the fully connected branch processes static flight-condition and mass-property inputs. The network output is the reduced-order short-period model coefficient vector [9].

https://doi.org/10.71945/icas2026_0686

Dropout layers are included after the LSTM and fully connected layers to reduce overfitting during training. During each training update, a specified fraction of layer activations is randomly deactivated, forcing the network to avoid reliance on a small subset of neurons and improving its ability to generalise to validation flight conditions. During prediction, all neurons are active and the dropout operation is disabled.

The network estimates the coefficient vector

$$\boldsymbol{\theta}_{NN} = [z_w \quad z_q \quad m_w \quad m_q \quad z_{\delta \mathrm{e}} \quad m_{\delta \mathrm{e}}]^T \tag{5}$$

This formulation differs from a purely black-box response-prediction model because the network output remains physically interpretable. The predicted coefficients define a reduced-order flight-dynamics model, from which both time-domain responses and modal characteristics can be computed.

Physical consistency is imposed through a loss function based on the residuals of the reduced-order equations of motion. For each training sequence, the equation residuals are:

$$\begin{aligned} r_w &= \dot{w} - z_w w - z_q q - z_{\delta e} \delta e \\ r_q &= \dot{q} - m_w w - m_q q - m_{\delta e} \delta e \end{aligned} \tag{6}$$

where r_w and r_q are the residuals of the vertical-velocity and pitch-rate equations, respectively. The residual loss is calculated as the mean-square value of these residuals over the training sequence:

$$L_{\mathrm{residual}} = \frac{1}{N} \sum_{i=1}^{N} \left(r_{w,i}^2 + r_{q,i}^2 \right) \tag{7}$$

where N is the number of time samples in the training sequence. This term encourages the neural network to estimate coefficients that satisfy the reduced-order equations of motion.

A sign penalty is also included to discourage coefficient combinations that are inconsistent with the expected short-period dynamics. In this work, the principal stabilising terms z_w and m_q are constrained to remain negative, since these terms represent the damping contributions of vertical velocity and pitch rate in the reduced-order model. The sign penalty is implemented using a squared hinge function:

$$L_{\mathrm{sign}} = \max(0, z_w)^2 + \max\left(0, m_q\right)^2 \tag{8}$$

This penalty is zero when $z_w < 0$ and $m_q < 0$, and increases when either coefficient takes a non-physical positive value. The elevator-related coefficients are not assigned a universal sign constraint because their sign depends on the adopted elevator-deflection convention.

The modal penalty is used to ensure that the identified reduced-order state matrix produces physically meaningful short-period modal characteristics. For each predicted coefficient set, the short-period natural frequency and damping ratio are calculated from

$$\begin{aligned} \omega_n &= \sqrt{z_w m_q - z_q m_w} \\ \zeta &= -\frac{z_w + m_q}{2\omega_n} \end{aligned} \tag{9}$$

where ω_n is the natural frequency and ζ is the damping ratio. The modal penalty is applied when the predicted coefficients produce a non-positive value of ω_n^2, a negative damping ratio, or modal values outside prescribed physically acceptable bounds. It is written as

$$L_{\mathrm{modal}} = \max(0, -\omega_n^2)^2 + \max(0, -\zeta)^2 + \max(0, \zeta - \zeta_{\max})^2 + \max\left(0, \omega_n - \omega_{n,\max}\right)^2 \tag{10}$$

where $\zeta_{\max}$ and $\omega_{n,\max}$ are upper limits used to prevent unrealistically large damping ratios or natural frequencies. This term therefore penalises unstable or non-physical modal behaviour while allowing the network to learn physically plausible variations across the flight-test envelope.

For a physically consistent model, both residuals should approach zero over the time history. The total loss is defined as

$$L_{\mathrm{total}} = 0.70 L_{\mathrm{residual}} + 0.15 L_{\mathrm{sign}} + 0.15 L_{\mathrm{modal}} \tag{11}$$

where the residual term enforces consistency with the reduced-order equations of motion, while the sign and modal penalty terms discourage unstable or non-physical coefficient combinations.

The residual term is given the largest weight because satisfying the reduced-order equations of motion is the primary physics-informed constraint. The sign and modal penalties are assigned smaller but non-

https://doi.org/10.71945/icas2026_0686

zero weights to discourage unstable or non-physical coefficient combinations without dominating the data-driven learning process. The weighting values were selected from preliminary tuning studies as a compromise between equation-residual minimization and modal stability of the identified models.

4 Training, Validation and Hyperparameter Optimisation

The processed dataset was divided into 41 training cases and 15 validation cases. The validation cases were selected to provide representative coverage of the available flight conditions and were not used to update the network weights. The training and validation distribution is shown in Figure 5.

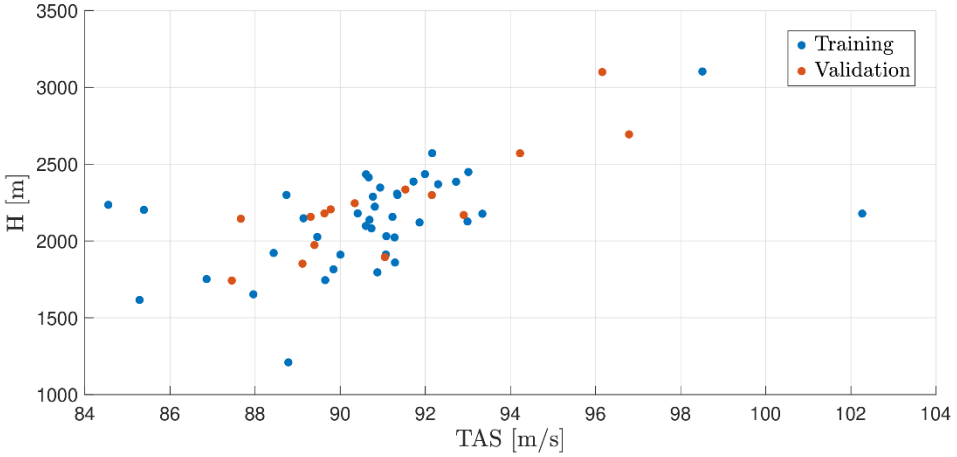


Figure 5. Distribution of training and validation cases across the available Saab 340B short period flight-test envelope.

Training was performed using mini batches extracted from the flight-test time histories. Mini-batching provided a consistent sequence length to the LSTM branch, while overlap between adjacent mini-batches reduced the loss of temporal information at window boundaries. The network weights and biases were updated using the Adam optimiser. The learning rate was reduced during training according to

$$\alpha_i = \frac{\alpha_0}{1 + d.i} \tag{12}$$

where α_0 is the initial learning rate, d is the decay factor and i is the training iteration.

A baseline configuration was first defined to provide a reference against which the tuned models could be compared. Manual hyperparameter tuning was then used as a comparator for the automated optimisation workflow. A one-factor-at-a-time tuning strategy was applied to the main architecture and training parameters. The final manually tuned configuration was evaluated using the same training and validation cases and the same accuracy metrics as the baseline and HEEDS-optimised models.

Siemens HEEDS was integrated with the MATLAB-based neural-network training framework to automate the hyperparameter search using the SHERPA algorithm [10]. The design variables included mini-batch size, mini-batch overlap, number of LSTM layers, number of LSTM hidden units, number of fully connected layers, number of fully connected neurons, initial learning rate, learning-rate decay and dropout coefficient. The optimisation objectives were to minimise the root mean square errors and maximise the correlation factors for pitch rate (q) and vertical velocity (w).

The baseline, manually tuned and best HEEDS hyperparameter configurations are summarised in Table 1.

Table 1. Baseline manually tuned and HEEDS-optimised neural-network configurations.

Hyperparameter	Baseline	Manually tuned	HEEDS-optimized
LSTM layers	1	2	1
Hidden units per LSTM layer	64	128	225
Fully connected layers	1	1	1
Neurons per fully connected layer	50	75	60
Dropout coefficient	0.6	0.6	0.8
Mini-batch size	100	150	150
Mini-batch overlap	50%	25%	10%

https://doi.org/10.71945/icas2026_0686

Initial learning rate	1×10^{-3}	1×10^{-3}	5×10^{-4}
Learning-rate decay	1×10^{-2}	1×10^{-4}	1×10^{-6}
Number of evaluated designs	1	43	130
Approximate computation time	—	19.02 h	4 days, 15 h, 20 min

The identified reduced-order models were validated by simulating the short-period response using the measured elevator input for each manoeuvre. The predicted pitch-rate and vertical-velocity responses were then compared with the measured flight-test responses. The predicted model was evaluated using two criteria correlation factors $(corr)$ and root mean square error $(RMSE)$.

The correlation factor and root mean square error are computed separately for the pitch-rate and vertical-velocity responses. The pitch-rate correlation factor, q_{Corr}, is calculated as

$$q_{\mathrm{Corr}} = \frac{\sum_{i=1}^{n}(q_i - \overline{q})\left(\hat{q}_i - \overline{\hat{q}}\right)}{\sqrt{\sum_{i=1}^{n}(q_i - \overline{q})^2}\sqrt{\sum_{i=1}^{n}\left(\hat{q}_i - \overline{\hat{q}}\right)^2}} \tag{13}$$

where q_i and $\hat{q}_i$ are the measured and predicted pitch-rate responses, respectively, and $\overline{q}$ and $\overline{\hat{q}}$ are their corresponding mean values. Similarly, the vertical-velocity correlation factor, w_{Corr}, is calculated as

$$w_{\mathrm{Corr}} = \frac{\sum_{i=1}^{n}(w_i - \overline{w})\left(\hat{w}_i - \overline{\hat{w}}\right)}{\sqrt{\sum_{i=1}^{n}(w_i - \overline{w})^2}\sqrt{\sum_{i=1}^{n}\left(\hat{w}_i - \overline{\hat{w}}\right)^2}} \tag{14}$$

where w_i and $\hat{w}_i$ are the measured and predicted vertical-velocity responses, respectively. A correlation factor close to unity indicates strong agreement in response shape and phase.

The root mean square errors for pitch rate and vertical velocity are computed as

$$q_{\mathrm{RMSE}} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{q}_i - q_i)^2} \tag{15}$$

$$w_{\mathrm{RMSE}} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\hat{w}_i - w_i)^2} \tag{16}$$

where n is the number of time samples. The correlation factors therefore assess the similarity in response shape, while the RMSE values quantify the magnitude of the prediction error.

5 Results and Discussion

The accuracy of the baseline, manually tuned, and HEEDS-optimised models is compared in Figure 6 and Table 2.

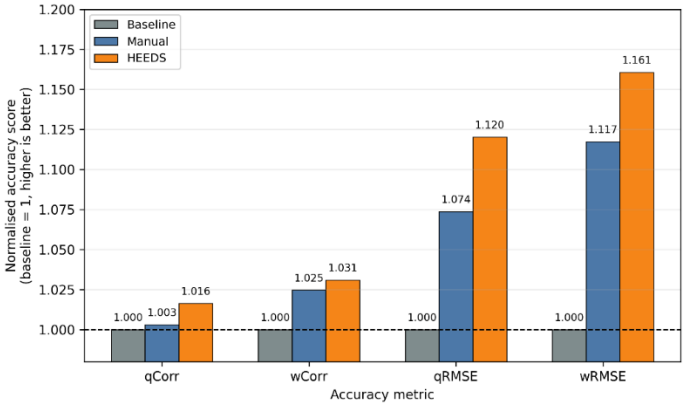


Figure 6. Accuracy comparison between the baseline, manually tuned and HEEDS-optimised models. Correlation terms are normalised by the baseline value, while RMSE terms are plotted as the ratio of baseline RMSE to model RMSE, so that higher values indicate improved performance for all metrics.

https://doi.org/10.71945/icas2026_0686

Table 2. Accuracy comparison between the baseline, manually tuned and HEEDS-optimised models.

Model	q_{Corr}	w_{Corr}	q_{RMSE}	w_{RMSE}	Combined improvement
Baseline	0.924 ± 0.038	0.937 ± 0.054	0.127 ± 0.040	0.159 ± 0.070	—
Manually tuned	0.927 ± 0.048	0.960 ± 0.026	0.118 ± 0.042	0.142 ± 0.056	20.2%
HEEDS-optimised	0.940 ± 0.052	0.966 ± 0.030	0.113 ± 0.0410	0.137 ± 0.057	29.3%

The baseline model captured the dominant short-period response, achieving mean correlation factors above 0.9 for both pitch rate and vertical velocity. Manual tuning improved the combined accuracy metric by 20.2% relative to the baseline. The improvement was most significant for the vertical-velocity response, where the mean correlation factor increased and the RMSE decreased.

The HEEDS-optimised model provided the best overall performance. The best design was obtained at design number 98 out of 130 evaluations and achieved a 29.3% combined improvement relative to the baseline. The optimisation first exceeded the manually tuned improvement of 20.2% at design number 46, where a 21.3% improvement was achieved. This indicates that the automated search was able to match and then surpass the manually tuned configuration within the explored design space. The final HEEDS design also outperformed the manually tuned model, giving higher mean correlation factors and lower mean RMSE values for both pitch rate and vertical velocity. The evolution of the HEEDS performance metric is shown in Figure 7, illustrating the progressive improvement achieved by SHERPA as additional designs were evaluated.

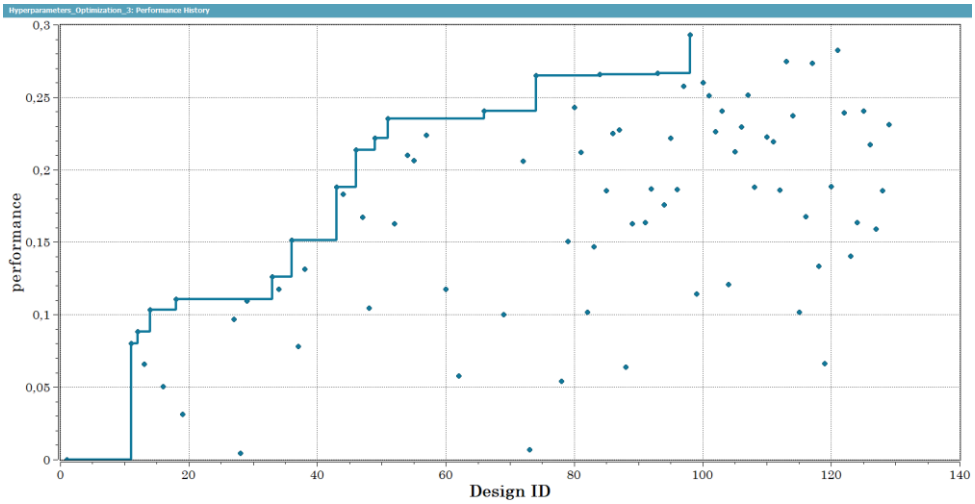


Figure 7. HEEDS performance history during automated hyperparameter optimisation, showing the improvement in combined prediction accuracy with design iteration.

Representative time-domain comparisons are shown in Figure 8. In many cases, both the manually tuned and HEEDS-optimised models capture the dominant short-period response. However, the HEEDS-optimised model provides improved agreement in selected cases, particularly around peak response regions where the manually tuned model underpredicts or overpredicts the response amplitude. This supports the quantitative improvement observed in the correlation-factor and RMSE metrics.

https://doi.org/10.71945/icas2026_0686

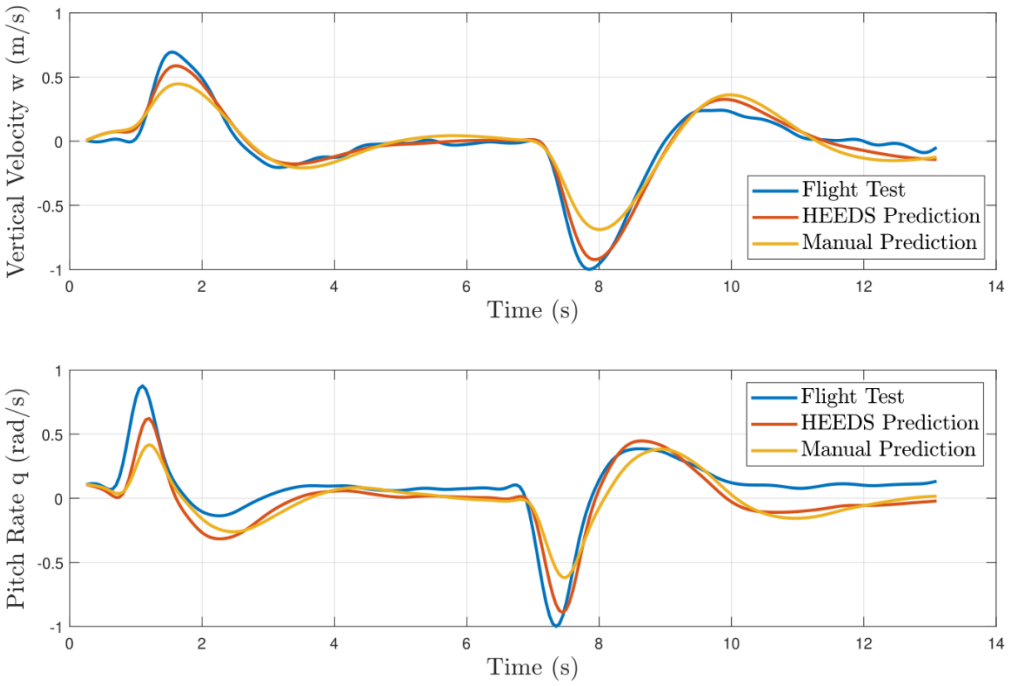
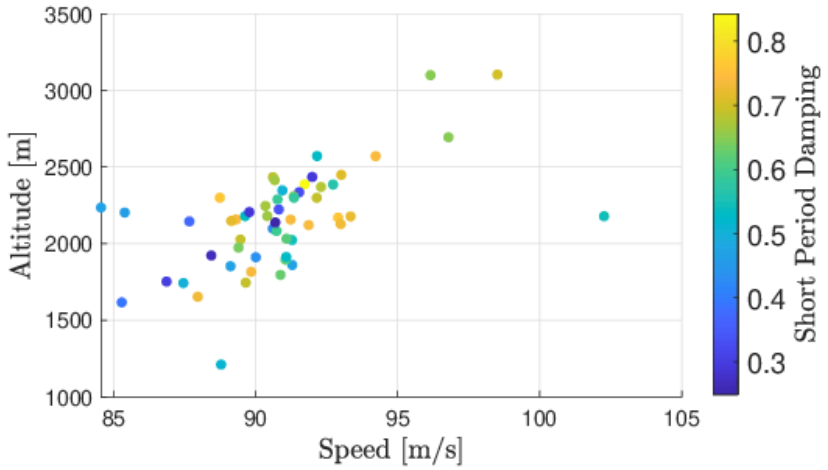


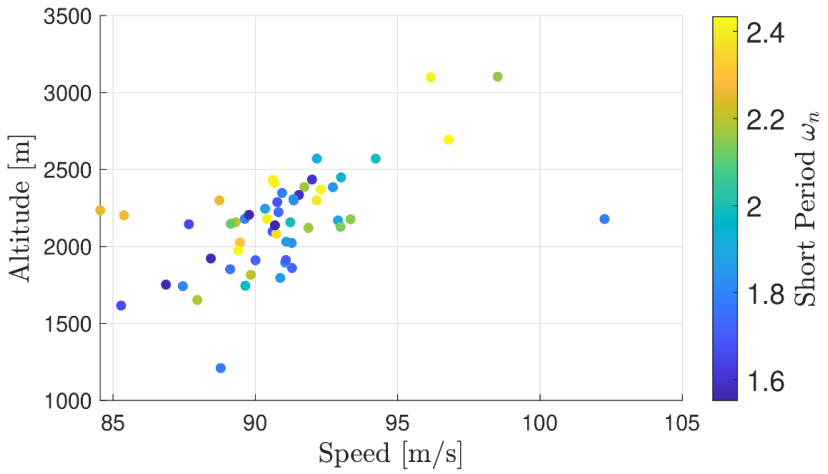
Figure 8. Example time-domain response comparison showing improved HEEDS prediction relative to the manually tuned model for selected Saab 340B short-period manoeuvres.

The estimated reduced-order coefficients were used to compute short-period damping ratio and natural frequency across the evaluated flight conditions. The resulting modal distributions are shown in Figure 9. These modal characteristics provide a physically interpretable output from the machine-learning framework, rather than only a time-domain response prediction.



(a) Damping ratio, ζ (-).

https://doi.org/10.71945/icas2026_0686



(b) Natural frequency, ω (rad/s)

Figure 9. HEEDS-predicted short-period modal characteristics across the available flight envelope.

The predicted modal characteristics remained physically meaningful across the evaluated cases. This indicates that the physics-informed loss formulation successfully constrained the network to produce stable and interpretable reduced-order models. The modal maps also highlight a limitation of the dataset: most flight-test points are concentrated within a relatively narrow region of the flight envelope. Predictions near sparsely populated regions should therefore be interpreted cautiously until additional data are available.

Overall, the results show that physics-informed machine learning can identify reduced-order short-period dynamics from Saab 340B flight-test data while retaining physical interpretability. Manual tuning improves the baseline model but requires significant user involvement. HEEDS optimisation provides the best overall accuracy and a more systematic tuning workflow, although at the cost of increased machine computation time.

6 Conclusion

This paper presented a physics-informed machine-learning framework for identifying the short-period longitudinal dynamics of the Cranfield University Saab 340B aircraft from flight-test data. The proposed architecture combines LSTM layers for sequential flight-test measurements with fully connected layers for static flight-condition and mass-property inputs. The network estimates the coefficients of a reduced-order short-period state-space model, allowing pitch-rate and vertical-velocity responses, damping ratio and natural frequency to be obtained from a physically interpretable model.

The preprocessing procedure retained the dominant dynamic content of the flight-test data, with correlation factors above 0.98 between raw and processed signals. The baseline physics-informed model captured the main short-period response, achieving mean correlation factors above 0.9 for both pitch rate and vertical velocity. Manual hyperparameter tuning improved the baseline model by 20.2%, while Siemens HEEDS-based optimisation achieved the best overall performance, with a 29.3% improvement relative to the baseline.

The results demonstrate that physics-informed machine learning can support aircraft flight-dynamics system identification while retaining physical interpretability through the estimated reduced-order coefficients. HEEDS-based optimisation of hyperparameters using SHERPA algorithm provided a practical route for improving model accuracy and reducing manual tuning effort, although at the cost of greater machine computation time. Future work should include additional flight-test data, extension to phugoid and lateral-directional dynamics, and direct comparison with classical aircraft system-identification methods.

7 Acknowledgements

The authors gratefully acknowledge the support of Cranfield University and the National Flying Laboratory Centre for providing access to the Saab 340B flight-test data used in this study. The authors also thank the flight-test, aircraft operations, and data-support teams involved in the acquisition and preparation of the experimental datasets. The contribution of Siemens Digital Industries Software is also acknowledged, particularly in relation to the use of HEEDS for automated hyperparameter

optimisation. The views expressed in this paper are those of the authors and do not necessarily represent those of the supporting organisations.

8 Contact Author Email Address

Mailto: mushfiqul.alam@cranfield.ac.uk

9 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 proceedings.

References

[1] E. A. Morelli and J. A. Grauer, “Advances in Aircraft System Identification at NASA Langley Research Center,” *J. Aircr.*, vol. 60, no. 5, pp. 1354–1370, May 2023, doi: 10.2514/1.c037274.

[2] M. V Cook, *Flight Dynamics Principles*, 2nd ed. Elsevier Ltd., 2007.

[3] C. C. de Visser and D. M. Pool, “Stalls and Splines: Current Trends in Flight Testing and Aerodynamic Model Identification,” *J. Aircr.*, vol. 60, no. 5, pp. 1480–1502, May 2023, doi: 10.2514/1.c037283.

[4] E. A. Morelli, “Global nonlinear parametric modelling with application to F-16 aerodynamics,” in *Proceedings of the 1998 American Control Conference. ACC (IEEE Cat. No.98CH36207)*, 1998, pp. 997–1001 vol.2. doi: 10.1109/ACC.1998.703559.

[5] M. Millidere, M. Alam, S. Place, S. Mirjalili, and J. Whidborne, “Aerodynamic Modelling of Saab 340B Development Using Binary Particle Swarm Optimization,” in *AIAA SciTech Forum and Exposition, 2024*, 2024. doi: 10.2514/6.2024-1495.

[6] Y. Yu, H. Yao, and Y. Liu, “Aircraft dynamics simulation using a novel physics-based learning method,” *Aerosp. Sci. Technol.*, vol. 87, pp. 254–264, May 2019, doi: 10.1016/j.ast.2019.02.021.

[7] T. Stachiw, A. Crain, and J. Ricciardi, “A physics-based neural network for flight dynamics modelling and simulation,” *Adv. Model. Simul. Eng. Sci.*, vol. 9, no. 1, p. 13, Dec. 2022, doi: 10.1186/s40323-022-00227-7.

[8] Siemens Digital Industries Software (2025), “Simcenter HEEDS (Version 2504)” , Plano TX:Siemens

[9] A. Navarro, “Physics Informed Machine Learning for Flight Dynamics Identification of Saab 340 from Flight Test Data,”, Cranfield University, Bedford, 2025.

[10] R. C. Technology, “SHERPA – An Efficient and Robust Optimization/Search Algorithm,” May 2023.

https://doi.org/10.71945/icas2026_0686