2026 Volume: 02 Issue: 11 www.ijrpa.com ISSN 2456-9995 Research Article

International Journal Research Publication Analysis

Page: 01-11

AIRCRAFT ENGINE RISK PREDICTION SYSTEM

***P. Rohan, K. Teja, S. V. Sai Pavan Teja, Bharat Chandra Sahu, Dr. P. Srihari**

Department of CSE (Artificial Intelligence & Data Science) GMRIT – Deemed to be University, Rajam, Andhra Pradesh, India.

Article Received: 09 October 2026
Article Revised: 29 October 2026
Published on: 18 November 2026

***Corresponding Author: P. Rohan**
Department of CSE (Artificial Intelligence & Data Science) GMRIT – Deemed to be University, Rajam, Andhra Pradesh, India.

DOI: https://doi-doi.org/101555/ijrpa.2622

ABSTRACT

Aircraft engines operate under demanding conditions in which changes in vibration, temperature, pressure, lubrication, and engine load can indicate developing mechanical or thermal problems. Unexpected failures can affect operational safety, aircraft availability, and maintenance cost. This paper presents an Aircraft Engine Risk Prediction System that applies machine learning to operational sensor parameters to support early risk assessment and predictive maintenance. The workflow loads engine observations from a CSV dataset, arranges the records by time step, and prepares the data for feature generation. Derived indicators include thermal stress, lubrication risk, mechanical stress, and a degradation score. Rolling statistics are generated for vibration root-mean-square values and sensor temperature over a ten-step window, after which rows containing unavailable rolling values are removed. A Logistic Regression classifier implemented with Scikit-learn is trained using balanced class weights to reduce the effect of class imbalance. The trained model is saved as `model.pkl` and integrated into a Streamlit web application. During prediction, the application accepts current sensor values, recreates the engineered features, and uses the trained model to estimate risk. The interface presents a risk probability and assigns a risk category—SAFE, MODERATE RISK, or HIGH RISK—along with a corresponding monitoring or maintenance recommendation. The system is intended as a decision-support prototype rather than a replacement for certified aircraft-engine inspection or maintenance procedures. The available project material does not include verified final test metrics; therefore, this paper does not fabricate quantitative performance claims. Evaluation using held-out data, class-wise metrics, calibration checks, and validation against operational engineering requirements is necessary before real-world use.

KEYWORDS: Aircraft engine, predictive maintenance, machine learning, Logistic Regression, sensor data, feature engineering, risk classification, Scikit-learn, Streamlit.

1. INTRODUCTION

Aircraft engines are complex systems that operate under high mechanical loads, elevated temperatures, pressure variations, and demanding environmental conditions. Their reliable operation depends on the coordinated performance of mechanical, thermal, fuel, lubrication, and control subsystems. Changes in sensor readings can indicate that an engine is operating outside its usual condition. Monitoring these measurements over time can help maintenance teams recognise developing abnormalities before they become severe.

Traditional maintenance programmes commonly combine scheduled inspection, time-based servicing, and corrective action after a fault is detected. Scheduled maintenance is essential, but a fixed interval may not fully represent the condition of an individual engine. A component may degrade more quickly or more slowly depending on its use, environment, operating history, and maintenance quality. Predictive maintenance attempts to use condition and historical data to estimate risk and support maintenance decisions at a suitable time.

Machine learning provides methods for finding relationships between observed measurements and labelled outcomes. When historical records include sensor values and corresponding health or risk labels, a classification model can learn patterns associated with different engine conditions. The model can then produce a prediction for a new observation. The reliability of such a prediction depends on representative data, correct labels, appropriate preprocessing, suitable validation, and an understanding of the model's limitations.

The proposed Aircraft Engine Risk Prediction System follows this predictive-maintenance idea. The system processes engine sensor observations, creates additional indicators, trains a Logistic Regression classifier, and exposes predictions through a Streamlit interface. It is a prototype for supporting awareness and planning; it is not intended to authorise flight operations or replace approved maintenance procedures.

2. LITERATURE SURVEY / RELATED WORK

The project literature review covers predictive maintenance, remaining useful life (RUL) prediction, fault diagnosis, classification, simulation-based assessment, and machine-learning approaches for aircraft-engine health monitoring. The reviewed work shows a progression from conventional condition monitoring toward temporal deep learning, ensemble models, and hybrid approaches.

Ref.	Focus	Technique / contribution	Year
[1]	Aircraft engine failure prediction	Artificial intelligence for engine-failure prediction	2024
[2]	Remaining useful life	Interpretable ensemble RUL prediction for dynamic maintenance scheduling	2025
[3]	RUL prediction	BiLSTM–Transformer–TCN feature fusion for aircraft-engine RUL	2025
[4]	Predictive maintenance	RUL-based predictive maintenance programme for aircraft engines	2025
[5]	Fault prediction	Adaptive hybrid sampling with BiLSTM for aircraft-engine fault prediction	2025
[6]	Fault diagnosis	1DCNN–BiLSTM with CBAM for aircraft-engine fault diagnosis	2024
[7]	Engine wear assessment	Simulation techniques for assessing aircraft-engine wear	2024
[8]	RUL prediction	Machine-learning-based aircraft-engine RUL prediction	2024
[9]	Fault diagnosis	CNN–BiGRU attention model for aircraft-engine fault diagnosis	2024
[10]	Sensor time-series augmentation	Autoregressive augmentation of multivariate sensor time series for failure prognostics	2024

The reviewed studies collectively indicate that aircraft-engine health monitoring benefits from temporal information, feature extraction, robust handling of imbalance, and models capable of capturing nonlinear relationships. Deep learning approaches can model complex sequences, while interpretable or simpler models remain useful when transparency, computational efficiency, and prototype development are priorities.

3. RESEARCH GAP

The literature includes sophisticated RUL and fault-diagnosis models, but an academic prototype also needs an end-to-end path from raw sensor observations to an understandable risk output. The project addresses this practical gap by integrating data preparation, feature engineering, rolling statistics, classification, model persistence, and a user-facing Streamlit interface in one workflow.

A further challenge is methodological consistency between model training and prediction. Features calculated during training must be recreated in the same way when a user submits new sensor values. The project therefore emphasises a shared feature-generation process, explicit temporal ordering, and validation before operational claims are made.

4. CONTRIBUTIONS

• A prototype workflow for aircraft-engine risk prediction using operational sensor parameters.

• Feature engineering for thermal stress, lubrication risk, mechanical stress, and degradation-related indicators.

• Ten-step rolling mean features for vibration RMS and sensor temperature to introduce recent-history context.

• A Scikit-learn Logistic Regression classifier configured with balanced class weights.

• Persistence of the trained estimator as `model.pkl` and integration with a Streamlit web interface.

• A user-facing output consisting of a model prediction, risk probability, risk category, and associated guidance.

• A documented validation and safety boundary that avoids treating the prototype as a certified aircraft-maintenance system.

5. METHODOLOGY

5.1 System Workflow

The implementation follows an eight-step sequence: data collection, preprocessing, feature engineering, rolling feature generation, model training, model deployment, engine risk prediction, and risk classification with recommendation.

Step	Stage	Main operation	Output
1	Data collection	Load `aircraft_vibration.csv`	Raw observations
2	Preprocessing	Sort by `time_step`, validate and prepare records	Ordered dataset
3	Feature engineering	Calculate stress and degradation indicators	Derived features
4	Rolling features	Ten-step means for vibration RMS and temperature	Contextual features
5	Model training	Fit Logistic Regression with balanced class weights	Trained classifier
6	Deployment	Save/load `model.pkl` in Streamlit	Inference model
7	Prediction	Recreate features and call `predict()` / `predict_proba()`	Class and probability
8	Classification	Map result to category and guidance	User-facing output

5.2 Data Collection and Preprocessing

The project identifies `aircraft_vibration.csv` as the input dataset. The listed sensor observations include RPM, vibration RMS, temperature, oil pressure, load factor, acoustic level, and altitude. The records are sorted by

`time_step` before rolling calculations. Required fields should be checked for presence, numeric compatibility, missing values, duplicate records, and plausible ranges.

The available project material does not establish the dataset dimensions, exact target column, class counts, or complete feature formulas. Those values must be taken from the actual implementation and dataset before a final experimental report is submitted.

5.3 Feature Engineering

Feature group	Purpose	Verification required
Thermal stress	Summarises thermal-condition behaviour	Exact formula, thresholds, units, scaling
Lubrication risk	Represents lubrication-related risk	Oil-pressure mapping and threshold rules
Mechanical stress	Summarises mechanical operating stress	Sensor variables and formula
Degradation score	Combines deterioration-related indicators	Weights, normalisation, label relationship
Rolling vibration mean	Represents recent vibration level	Window alignment and missing-value handling
Rolling temperature mean	Represents recent temperature level	Window alignment and missing-value handling

Feature engineering must avoid leakage. A feature should not use future observations, target labels, or statistics that would not be available at prediction time. The exact equations for the engineered indicators are not supplied in the presentation and therefore are not invented here.

5.4 Rolling Mean Features

A ten-observation rolling mean introduces recent-history context. For a trailing window, the current value and the preceding nine observations form the window once sufficient history exists. Initial rows may therefore contain unavailable rolling values and are removed according to the stated workflow. The implementation should confirm whether the current observation is included and whether a fixed ten-row window corresponds to a meaningful physical time interval.

For a sequence x, a trailing ten-step mean can be expressed conceptually as $x_t = (x + x + \dots + x) / 10$. This equation is explanatory; the actual implementation remains the authoritative definition.

5.5 Logistic Regression

Logistic Regression is a supervised classification algorithm that estimates class probabilities from a weighted combination of input features. For binary classification, the logistic function maps the linear score into a value between zero and one. The project specifies balanced class weights, which increase the contribution of minority-class observations during

training. Balanced weighting can reduce majority-class dominance but does not by itself establish high recall, good calibration, or operational suitability.

$p(y = 1 \mid x) = 1 / (1 + e^{-z})$, where $z = w^T x + b$.

6. SYSTEM IMPLEMENTATION

6.1 Proposed Architecture

The proposed system connects historical sensor data, preprocessing, feature generation, Logistic Regression training, model persistence, and a Streamlit interface. During inference, current sensor values are transformed into the same feature representation used during training and passed to the saved model.

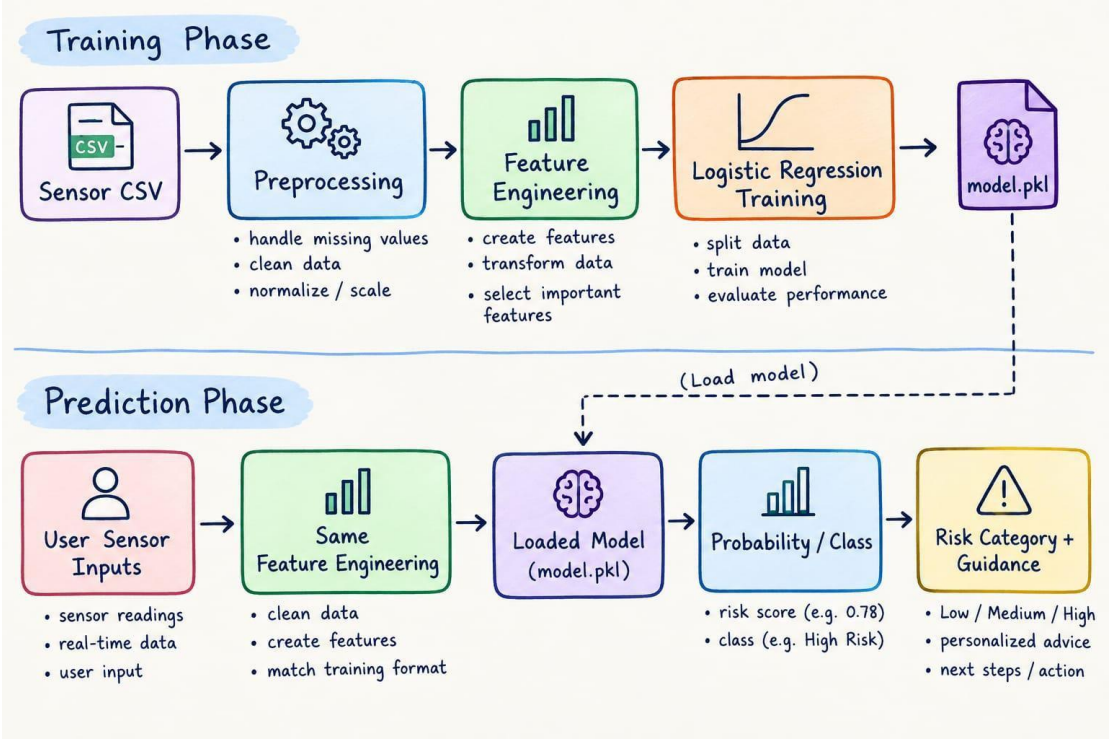

Figure 1: Proposed system architecture.

6.2 Data Processing and Prediction Flow

The training path starts from the CSV dataset and proceeds through validation, temporal ordering, feature generation, rolling calculations, model training, and model persistence. The prediction path begins with user inputs, performs the same feature transformation, loads the trained model, and produces a class and probability that are then mapped to a user-facing risk category.

Figure 2: Data-processing and prediction flow.

6.3 Streamlit Deployment

The trained estimator is saved as `model.pkl` and loaded by the Streamlit application. The application is intended to provide an accessible interface for entering sensor parameters and viewing the model output. Training and inference must use identical feature names, ordering, formulas, and preprocessing rules. A mismatch can produce misleading predictions even when the serialized model loads successfully.

6.4 Risk Classification

Category	General interpretation	Guidance
SAFE	Lowest defined model-risk output	Continue routine monitoring under approved procedures; not proof of zero risk.
MODERATE RISK	Intermediate model-risk output	Review sensor trends and follow authorised inspection guidance.
HIGH RISK	Highest defined model-risk output	Escalate for qualified review under approved maintenance procedures.

The available project material does not specify the numerical probability thresholds separating the displayed categories. Those thresholds must be taken directly from the implemented application rather than inferred or invented.

7. RESULTS AND DISCUSSION

The available project report includes interface outputs demonstrating the intended SAFE, MODERATE RISK, and HIGH RISK presentation. However, the source explicitly states that verified final test-set metrics were not available. Accordingly, the figures below document the interface outputs rather than claiming quantitative model performance.

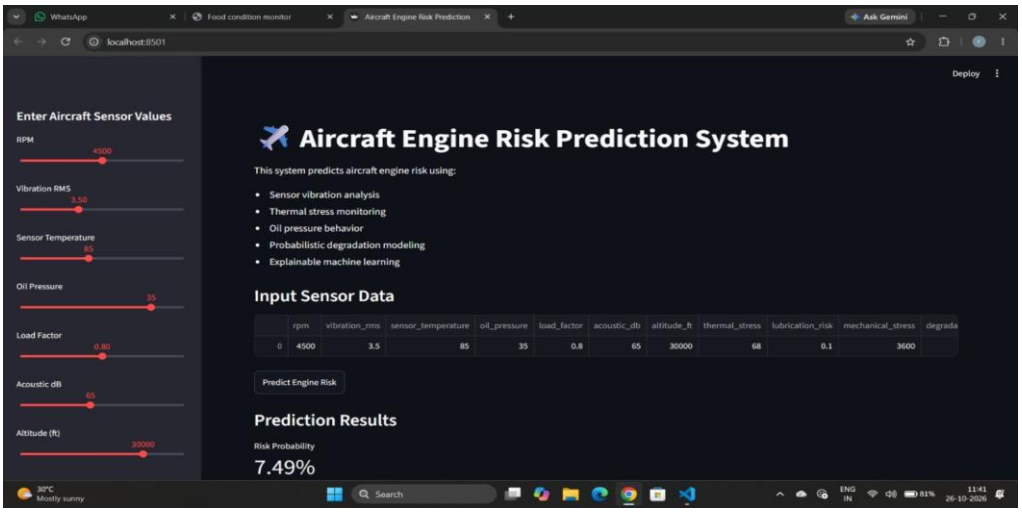

Figure 3: Aircraft Engine Risk Prediction System interface and sensor input values.

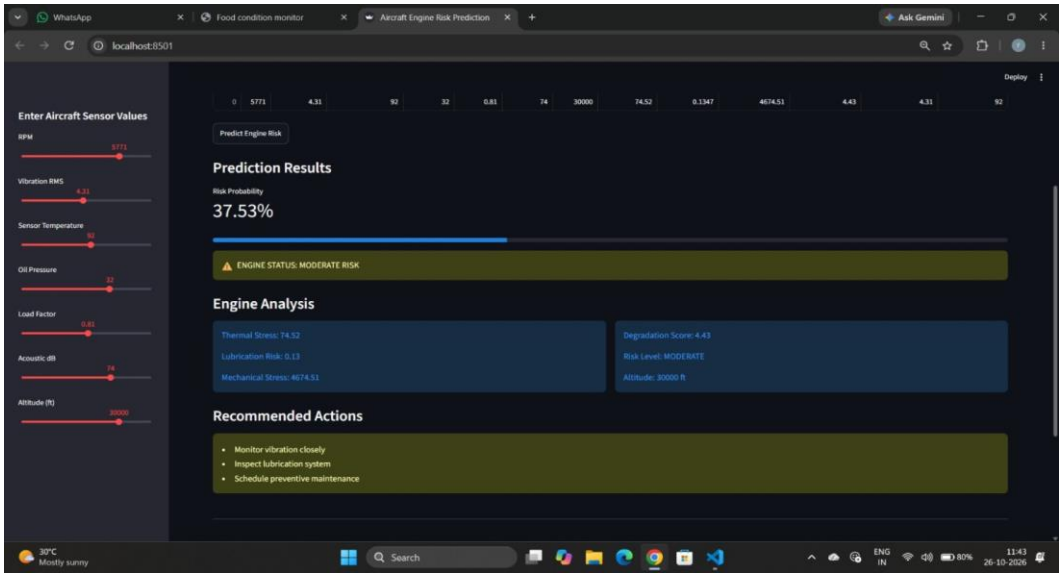

Figure 4: Prediction output showing a moderate engine-risk level.

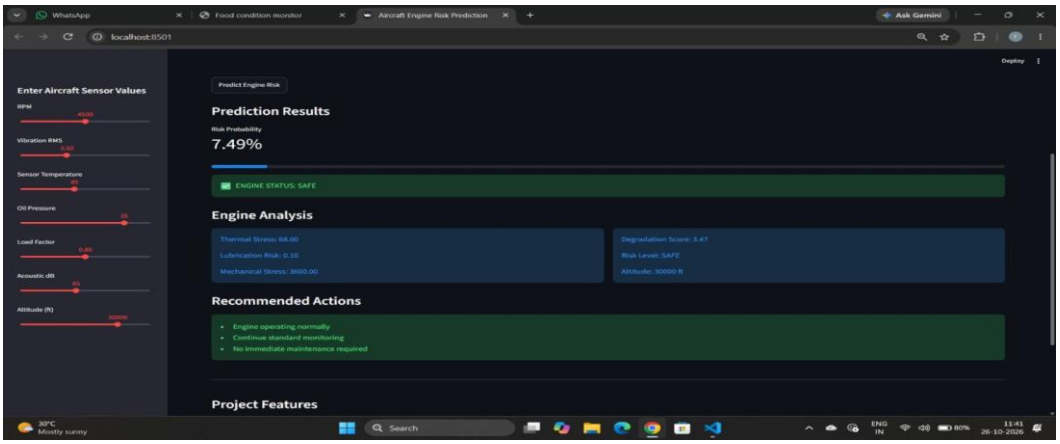

Figure 5: Prediction output showing a safe engine-risk level.

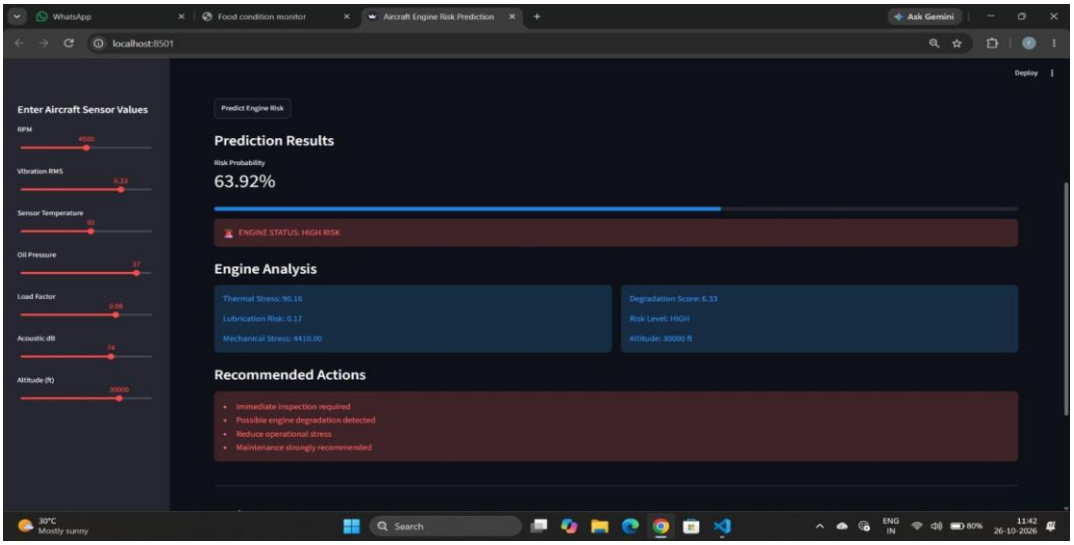

Figure 6: Prediction output showing a high engine-risk level.

7.1 Evaluation Plan

A valid evaluation should separate training and held-out test data without leakage. For temporal observations, a chronological split should be considered. The test set should be evaluated using a confusion matrix, accuracy, precision, recall, and F1-score, with class-wise definitions or averaging methods stated explicitly. If probability outputs are used, calibration and threshold sensitivity should also be examined.

Metric	Purpose	Status in available project material
Accuracy	Overall correct classification rate	Not provided as a verified final result
Precision	Reliability of positive predictions	Not provided as a verified final result
Recall	Detection of actual positive/risky cases	Not provided as a verified final result
F1-score	Balance of precision and recall	Not provided as a verified final result
ROC-AUC	Ranking ability across thresholds	Not provided as a verified final result
Confusion matrix	Counts of true/predicted classes	To be generated from actual test predictions

7.2 Limitations

- The available material does not provide verified final quantitative test metrics.
- The exact dataset size, target-label definition, class distribution, and feature formulas require verification from the implementation.
- The ten-step rolling features require sufficient historical context; the exact inference mechanism for a single current input must be confirmed in the application.
- A Logistic Regression baseline may not capture all complex nonlinear relationships present in engine behaviour.
- The prototype is not a certified aircraft-engine monitoring or maintenance system and must not be used as the sole basis for airworthiness or maintenance decisions.

8. CONCLUSION

This paper presents an Aircraft Engine Risk Prediction System that integrates sensor-data preprocessing, feature engineering, rolling statistics, Logistic Regression, model persistence, and a Streamlit interface into a single predictive-maintenance prototype. The system is designed to transform operational measurements into an interpretable risk output consisting of a predicted class, probability estimate, and associated guidance. The available implementation material demonstrates the intended workflow and interface outputs, but it does not establish verified quantitative performance on a held-out test set. The system should therefore be regarded as an academic decision-support prototype rather than a validated

aviation maintenance solution.

9. FUTURE SCOPE

- Verify the complete dataset specification, target labels, feature formulas, and class distribution from the implementation.
- Run a reproducible held-out evaluation and report class-wise precision, recall, F1-score, confusion matrix, and suitable probability metrics.
- Compare Logistic Regression with appropriate baseline and nonlinear classifiers using the same split and feature definitions.
- Investigate richer temporal models and additional sensor features where justified by the data.
- Improve rolling-feature inference by defining a reliable history-buffer mechanism for live prediction.
- Add model versioning, reproducible environments, input-range validation, and structured experiment tracking.
- Capture and document the actual Streamlit application behaviour and test cases.
- If future operational research is pursued, use representative real-world data, independent validation, engineering review, cybersecurity controls, and relevant aviation compliance processes.

REFERENCES

1. A. Bentaleb, K. Toumlal, and A. Abouchabaka, “Predicting aircraft engine failures using artificial intelligence,” International Journal of Advanced Computer Science and Applications, vol. 15, no. 2, pp. 944–949, 2024.
2. H. Özcan, “Interpretable ensemble remaining useful life prediction enables dynamic maintenance scheduling for aircraft engines,” Scientific Reports, vol. 15, Art. no. 39795, 2025.
3. Y. Yang, X. Su, C. Wang, H. Liu, K. Fu, T. Xie, and Z. Zhang, “BLTTNet: Feature fusion based on BiLSTM-Transformer-TCN for prediction of remaining useful life of aircraft engines,” Scientific Reports, vol. 15, Art. no. 27696, 2025.
4. F. Xue, G. Jin, L. Tan, C. Zhang, and Y. Yu, “Predictive maintenance programs for aircraft engines based on remaining useful life prediction,” Scientific Reports, vol. 15, Art. no. 41756, 2025.
5. J. Hu, X. Jiang, H. Xu, and K. Zhang, “Fault prediction of aircraft engine based on

adaptive hybrid sampling and BiLSTM," Scientific Reports, vol. 15, Art. no. 13726, 2025.

6. J. Wu, L. Kong, S. Kang, H. Zuo, Y. Yang, and Z. Cheng, "Aircraft Engine Fault Diagnosis Model Based on 1DCNN-BiLSTM with CBAM," Sensors, vol. 24, Art. no. 780, 2024.
7. A. Madane, J. Lacaille, F. Forest, H. Azzag, and M. Lebbah, "Assessing Aircraft Engine Wear through Simulation Techniques," 2024.
8. M. Kimollo and X. Liu, "Aircraft Engine Remaining Useful Life Prediction Using Machine Learning," 2024.
9. C. Jin and J. Bai, "Research on Aircraft Engine Fault Diagnosis Based on CNN-BiGRUAttention Model," 2024.
10. D. Baptista de Souza and B. Paes Leao, "Data Augmentation of Multivariate Sensor Time Series Using Autoregressive Models and Application to Failure Prognostics," 2024.