FUDMA Journal of Sciences (FJS)
ISSN online: 2616-1370
ISSN print: 2645 - 2944
Vol. 9 No. 12, December (Special Issue), 2025, pp 581 – 589
DOI: https://doi.org/10.33003/fjs-2025-0912-4237

ATTENTION-ENHANCED TEMPORAL MODELING FOR PROGNOSTICS OF TURBOFAN ENGINE REMAINING USEFUL LIFE

*[1,2]Rexcharles E. Donatus, [1]Loveth O. Uhiah and [1]Ajiboye E. Olanrewaju

[1]Department of Aerospace Engineering, Air Force Institute of Technology, Kaduna, Nigeria
[2]Africa Centre of Excellence on Technology Enhanced Learning, National Open University of Nigeria

*Corresponding authors' email: charlly4eyims@yahoo.com

ABSTRACT

Advances in sensor technology and automation are shifting aviation maintenance from fixed schedules to condition-based predictive maintenance (CBPM), which leverages real-time sensor data and machine learning to anticipate failures and optimize interventions. In this study, a deep learning architecture is presented, integrating BiLSTM with a multi-head self-attention module for RUL prediction, and its performance is assessed using the NASA C-MAPSS dataset. The BiLSTM captures bidirectional temporal dependencies in degradation sequences, while the attention mechanism adaptively emphasizes critical cycles and sensor signals. Pre-processing involved piecewise RUL labelling (capped at 125 cycles), cluster-based normalization, rolling statistical features, and sliding-window sequence generation. On FD004, the BiLSTM–attention model achieved an MAE of 9.45, RMSE of 15.52, and PHM score of 3853.21, outperforming the baseline LSTM (MAE 17.80, RMSE 25.14, PHM 4211.21). On FD001, the BiLSTM–attention delivered comparable accuracy, with an MAE of 11.42, RMSE of 15.05, and PHM score of 387.15, matching or exceeding baseline performance (MAE 11.23, RMSE 15.12, PHM 395.55). These findings demonstrate that integrating bidirectional sequence modelling with adaptive attention enhances predictive robustness across varying operating conditions. The proposed approach not only achieves strong generalization but also outperforms state-of-the-art benchmarks in aircraft engine Remaining Useful Life prediction, offering practical benefits for predictive maintenance through improved safety, reduced operational costs, and extended fleet availability.

Keywords: Attention Mechanism, Aviation Maintenance, Bidirectional LSTM, Prognostics, Remaining Useful Life

INTRODUCTION

The aerospace sector, along with other industries reliant on complex machinery, is witnessing a paradigm shift from conventional preventive maintenance (PM) strategies to more sophisticated condition-based predictive maintenance (CBPM) approaches (Ensarioğlu et al., 2023; Fan et al., 2024). Traditional PM practices, which are typically governed by fixed inspection and servicing intervals, often result in redundant maintenance activities and elevated operational costs (Andenyangtso et al., 2024; Asif et al., 2022; Donatus et al., 2025). Conversely, CBPM leverages advances in sensing technologies and data-driven analytics to continuously assess system health in real time, enabling proactive interventions grounded in empirical evidence (Adryan Fitra Azyus et al., 2025). Within this framework, accurate estimation of an asset's Remaining Useful Life (RUL) stands out as a key objective, providing forecasts of operational cycles prior to failure and thereby enhancing safety, reducing downtime, and optimizing maintenance planning (Fan et al., 2024; Wu et al., 2024).

In industrial systems, aviation, and aerospace, predictive maintenance is consequently a key priority, with RUL estimation of turbofan engines central to ensuring system reliability (Elsherif et al., 2025; Ensarioğlu et al., 2023; Ohoriemu & Ogala, 2024). The NASA Commercial Modular Aero-Propulsion System Simulation (C-MAPSS) dataset provides a widely adopted benchmark, simulating engine degradation under varying operating conditions and generating time-series data from nominal operation to failure (Adryan Fitra Azyus et al., 2025; Elsherif et al., 2025). This dataset has facilitated the application of diverse modelling strategies, ranging from survival analysis and ensemble learning (e.g., random forests, XGBoost) to deep learning methods particularly long short-term memory (LSTM) architectures (Adryan & Sastra, 2021; AsifKhan et al., 2022) Despite notable advances, ensuring robust and generalizable predictions across complex operating conditions remains an open challenge (Azyus, 2022; Remadna et al., 2021).

Recent studies have explored advanced architectures such as GRUs, CNNs, autoencoders, and self-attention to improve degradation modelling and temporal feature extraction (Dida et al., 2025; Ensarioğlu et al., 2023). Hybrid models that combine CNNs with recurrent networks, particularly BiLSTMs, consistently deliver stronger performance in complex operating conditions (Y. Liu & Wang, 2021; Muneer et al., 2021). Approaches integrating convolutional autoencoders with GRU or BiLSTM layers further demonstrate the effectiveness of multi-architecture designs for turbofan RUL prediction (Fan et al., 2024; Li et al., 2025).

Despite rapid progress, contemporary RUL approaches exhibit several recurring limitations. Many data-driven models depend heavily on careful feature engineering or dataset-specific pre-processing, which reduces their robustness when transferred to new operating regimes (Elsherif et al., 2025; Kumar et al., 2024). Deep recurrent or convolutional architectures can model temporal dynamics effectively but often lack uncertainty quantification and interpretability, making them difficult to trust for safety-critical decisions (Adryan & Sastra, 2021; Wu et al., 2024). multi-condition datasets such as C-MAPSS FD004, simpler models overfit or fail to generalize, while more complex models suffer from high computational cost and unstable training, especially for long sequences (Azyus, 2022; Kumar et al., 2024; Remadna et al., 2021).

Turbofan RUL prediction is challenging because degradation signals are non-stationary, noisy, and influenced by changing operational settings and multiple concurrent fault modes; this heterogeneity obscures consistent failure signatures and complicates learning (Kumar et al., 2024). Engines operate under variable loads and environmental conditions that change sensor distributions over time, necessitating normalization or domain-aware schemes to avoid spurious correlations (Dida et al., 2025; Pan et al., 2025). Moreover, the scarcity of labelled run-to-failure data for some fault types and the long horizons needed for safe maintenance decisions increase sensitivity to late prediction errors and motivate metrics (e.g., PHM score) that penalize such mistakes (Cohen et al., 2021).

The recent literature on turbofan engine Remaining Useful Life (RUL) prediction highlights a progressive evolution of methods from foundational pre-processing strategies to cutting-edge state-space models. Asif et al (2022) established a strong baseline by demonstrating that deep LSTM models can achieve improved RUL prediction when combined with meticulous pre-processing, including piecewise linear labelling and correlation-based sensor selection. Building on this, Ensarioğlu et al (2023) proposed an engine-specific labelling method using change-point detection alongside handcrafted features, implemented in a hybrid 1D-CNN-LSTM model, which enhanced predictive accuracy by better capturing individual degradation patterns. Fan et al (2024) advanced the field by introducing a Two-Stage Attention-Based Hierarchical Transformer, which systematically applies temporal and sensor-wise attention to capture long-range dependencies more effectively than traditional RNNs. Complementing this, Elsherif et al (2025), developed a hybrid Convolutional Autoencoder and Attention-based LSTM (CAELSTM) framework, where the autoencoder reduces noise and extracts features before attention-driven temporal modelling, yielding highly competitive prognostic performance. Most recently, Li et al (2025) proposed a Bidirectional Mamba model combined with causal discovery, representing a new generation of approaches that improve computational efficiency for long sequences while enhancing interpretability by uncovering causal relationships.

Collectively, these studies illustrate the trajectory of RUL prognostics research, moving from pre-processing and hybrid deep learning toward increasingly efficient, interpretable, and attention-driven architectures.

Recent reviews and experimental studies indicate a gap at the intersection of (a) architectures that both capture long-range temporal context and provide sensor-wise interpretability, and (b) rigorous cross-condition evaluations that include ablation, uncertainty estimation, and computational-cost reporting. In particular, while attention and transformer variants have shown promise, many studies either (i) apply attention without clear ablation vs. strong bidirectional recurrent baselines, or (ii) report accuracy gains without demonstrating robustness across multi-regime subsets such as FD004 or without reporting training stability and compute trade-offs (Z. Liu et al., 2024).

However, despite these advances, existing approaches often struggle to generalize across diverse operating conditions and to effectively capture long-range temporal dependencies in degradation patterns. To address this gap, this study proposes a hybrid framework that combines bidirectional LSTM networks with multi-head attention, designed to enhance temporal modelling while adaptively focusing on the most informative features for robust Remaining Useful Life prediction.

i. We propose a hybrid BiLSTM–multi-head self-attention architecture that combines bidirectional temporal modelling with sensor-wise attention to improve robustness across variable operating conditions.
ii. ii We introduce pre-processing steps (cluster-based normalization, rolling statistics and fixed-length sliding windows) tailored to mitigate distribution shifts among C-MAPSS subsets.
iii. We provide a thorough empirical evaluation on FD001 and FD004 including ablation studies (LSTM vs BiLSTM vs BiLSTM+Attention), reporting MAE, RMSE, PHM score, and computational considerations to demonstrate both accuracy and practical trade-offs.
iv. We analyse model interpretability by visualizing attention weights to identify sensors and time windows most predictive of imminent failure, improving operational trust.
v. We discuss limitations, uncertainty behaviour, and deployment considerations for condition-based maintenance in aircraft operations.

MATERIALS AND METHODS

Dataset and Problem Formulation

This study employs the NASA C-MAPSS dataset, a widely used benchmark for turbofan Remaining Useful Life (RUL) prediction (Graves et al., 2023). Each engine trajectory contains 26 variables, including an engine identifier, cycle count, three operating settings, and 21 sensor measurements, with realistic noise added to replicate operational conditions. RUL is defined as the number of cycles remaining before engine failure, making the task a regression problem. To stabilize model training, a piecewise labelling approach was applied with RUL values capped at 125 cycles.

The dataset consists of four subsets (FD001–FD004). For this work, FD001, which contains one operating condition and a single fault mode, and FD004, which includes six operating conditions and two concurrent fault modes, were selected to balance simple and complex scenarios. This combination allowed a comprehensive evaluation of the model's ability to generalize across different degradation complexities. The characteristics of the four C-MAPSS subsets are summarized in Table 1 to highlight their operating conditions and fault modes.

Table 1: Overview of C-MAPSS Dataset Subsets and Operating Conditions

Subset ID	Number of Train Trajectories	Number of Test Trajectories	Operational Conditions	Fault Modes	Primary Faults
FD001	100	100	One (Sea Level)	One	HPC Degradation
FD002	260	259	Six	One	HPC Degradation
FD003	100	100	One (Sea Level)	Two	HPC, Fan Degradation
FD004	248	249	Six	Two	HPC, Fan Degradation

Data Pre-processing

The raw dataset was reformatted into structured input sequences for model training. Pre-processing involved piecewise RUL labelling capped at 125 cycles, cluster-based normalization to ensure comparability across operating conditions, and the use of rolling statistical features to smooth fluctuations and highlight degradation trends. Temporal dependencies were captured with a sliding-window approach of 100 cycles and stride of one, with the RUL of the last cycle serving as the target. Shorter sequences were zero-padded to maintain consistent dimensions. These steps ensured robust and representative inputs for subsequent learning. The complete pre-processing and modelling pipeline is illustrated in Figure 1.

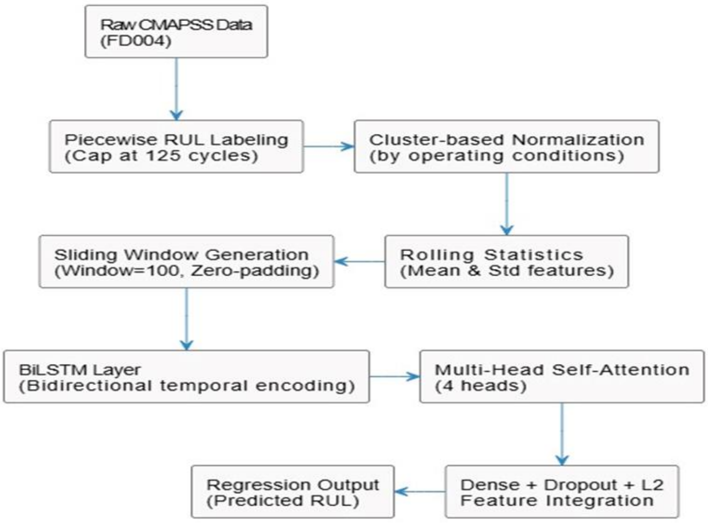

Figure 1: Proposed Data Pre-processing and BiLSTM–Attention Model Pipeline for RUL Prediction

Model Architectures

To evaluate performance fairly, three model categories were implemented. As a traditional benchmark, a Random Forest regressor was trained on rolling statistical features using 100 estimators and a maximum depth of 20 with bootstrap aggregation. For a deep learning baseline, a stacked LSTM was constructed with two recurrent layers of 128 and 64 hidden units, followed by dense layers for regression. Dropout of 0.2 and L2 regularization with a coefficient of 1e-4 were applied to improve generalization. The proposed BiLSTM with multi-head attention extended this baseline by incorporating bidirectional LSTM layers of 128 and 64 units, followed by a self-attention block with four heads to adaptively weight informative time steps and sensor signals. The outputs were passed to dense layers of 64 and 32 units before the final regression output. All deep models were optimized with Adam using an initial learning rate of 1e-3, mean squared error (MSE) loss, dropout, early stopping, and regularization. A learning-rate scheduler was further applied in the BiLSTM–Attention model to enhance convergence. The detailed configurations of the Random Forest, LSTM baseline, and proposed BiLSTM–Attention model are presented in Table 2 for clarity and reproducibility.

Table 2: Summary of Model Architectures

Model	Input Features & Window	Layers / Units	Regularization	Output	Training Settings
Random Forest	Rolling statistical features, no sequence	100 trees, max depth = 20 (bootstrap)	N/A	RUL (regression)	Default scikit-learn, train-test split
LSTM (Baseline)	Sliding windows (length = 100, stride = 1)	LSTM (128) → LSTM (64) → Dense (64, ReLU)	Dropout = 0.2, L2 = 1e-4	Linear	Adam (lr=1e-3), batch = 256, 30 epochs, early stopping
BiLSTM + Attention	Sliding windows (length = 100, stride = 1)	BiLSTM(128) → BiLSTM(64) → Multi-Head Attention (4 heads) → Dense(64, ReLU) → Dense(32, ReLU)	Dropout = 0.2, L2 = 1e-4	Linear	Adam (lr=1e-3), batch = 64–128, 30–50 epochs, early stopping

Training and Evaluation Metrics

Input sequences were generated with a sliding-window length of 100 cycles and training was conducted with early stopping based on validation loss. All experiments were implemented in Python using Google Colab with an NVIDIA T4 GPU to ensure reproducibility and efficient computation.

Model performance was evaluated using four metrics: mean absolute error (MAE), which measures the average prediction deviation; root mean squared error (RMSE), which emphasizes large errors; coefficient of determination (R^2), which represents the proportion of variance explained by the model; and the PHM score, which penalizes late predictions more heavily than early ones, reflecting safety-critical maintenance requirements. Together, these metrics provide a comprehensive measure of the models' predictive accuracy and generalization across datasets.

RESULTS AND DISCUSSION

Quantitative Performance

The proposed BiLSTM–multi-head self-attention framework was evaluated on the NASA C-MAPSS dataset using four performance metrics: MAE, RMSE, R^2, and the PHM score. MAE and RMSE measure absolute prediction accuracy, R^2 reflects explained variance, while the PHM score penalizes late predictions more heavily than early ones to align with safety-critical requirements. The comparative performance of Random Forest, baseline LSTM, and the proposed BiLSTM–Attention model on FD001 and FD004 is summarized in Table 3, highlighting the framework's effectiveness across both simple and complex prognostic settings.

Table 3: Comparative Performance of Baseline Models and Proposed Model FD001 and FD004 Datasets

Dataset	Model	MAE ↓	RMSE ↓	R^2 ↑	PHM Score ↓
FD001	Random Forest	13.856	19.110	0.789	1116.025
FD001	Baseline LSTM	11.233	15.123	0.868	395.553
FD001	BiLSTM + Attention	11.422	15.045	0.869	387.148
FD004	Random Forest	43.947	52.462	-2.689	169096708.757
FD004	Baseline LSTM	17.799	25.136	0.787	4211.214
FD004	BiLSTM + Attention	9.450	15.517	0.770	3853.213

For the simpler FD001 subset, Random Forest achieved reasonable baseline accuracy, but both deep learning models significantly improved results. The BiLSTM–Attention model achieved MAE of 11.42, RMSE of 15.05, and R^2 of 0.869, providing a slight improvement over the baseline LSTM (MAE 11.23, RMSE 15.12, R^2 0.868). In contrast, FD004 presented more complex operating conditions, where Random Forest completely failed to generalize. Here, the BiLSTM–Attention model substantially reduced errors with an MAE of 9.45, RMSE of 15.52, and PHM score of 3853.21, outperforming the baseline LSTM (MAE 17.80, RMSE 25.14, PHM 4211.21).

Figures 2 and 3 provide a visual comparison of model predictions on FD001 and FD004, respectively, showing how the proposed BiLSTM–Attention approach improves alignment with true RUL trajectories compared to baseline models.

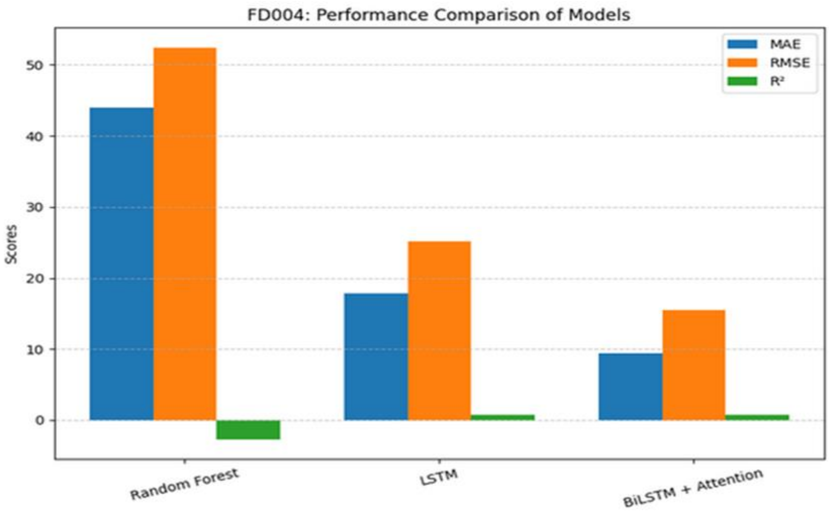

Figures 2: Assessment of Baseline Models Performance and Proposed Model on FD004

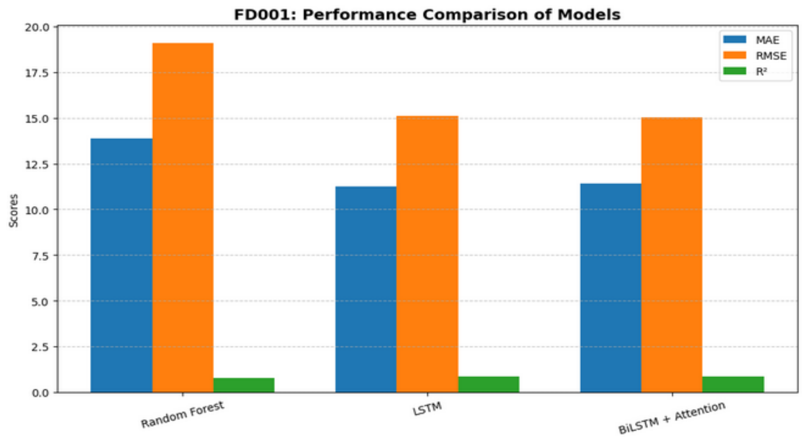

Figures 3: Assessment of Baseline Models Performance and Proposed Model on FD001

These results confirm that while attention brings only marginal gains in simpler scenarios, it becomes critical in complex multi-condition environments, where it enables the model to prioritize informative time steps and mitigate late prediction penalties. Similar observations were reported by (Dida et al., 2025), who highlighted the importance of attention-based sequence modelling in handling multi-regime degradation.

Visualization and Error Analysis

To complement the quantitative metrics, visualization analyses were performed. Scatter plots of predicted versus actual RUL show that the baseline LSTM generally follows the diagonal trend but underestimates at higher RUL values. In contrast, the BiLSTM–Attention model produces tighter clustering around the diagonal, especially in the mid-to-high RUL range, demonstrating better stability and reduced variance. To further assess predictive reliability, Figures 4 and 5 present scatter plots of predicted versus actual RUL values for the baseline LSTM and BiLSTM–Attention models, illustrating the degree of deviation from the ideal diagonal trend.

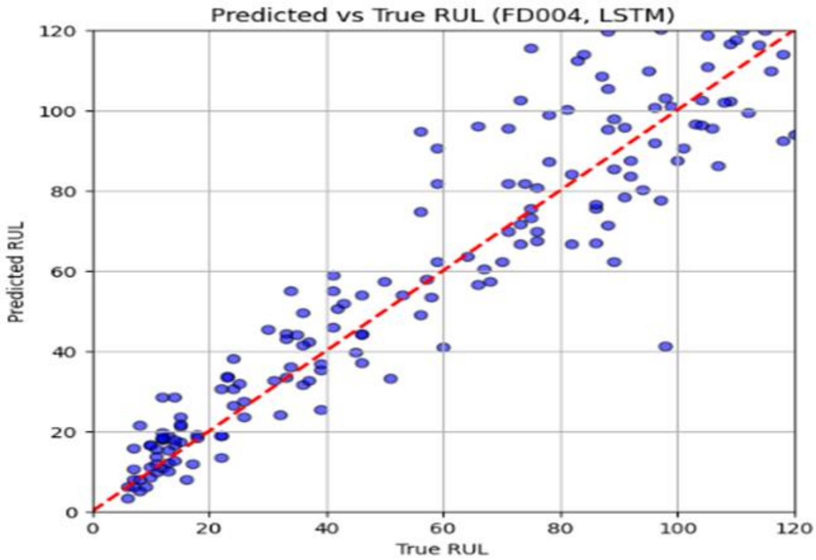

Figure 4: Predicted versus True RUL for the Baseline LSTM Model on the FD004 Dataset

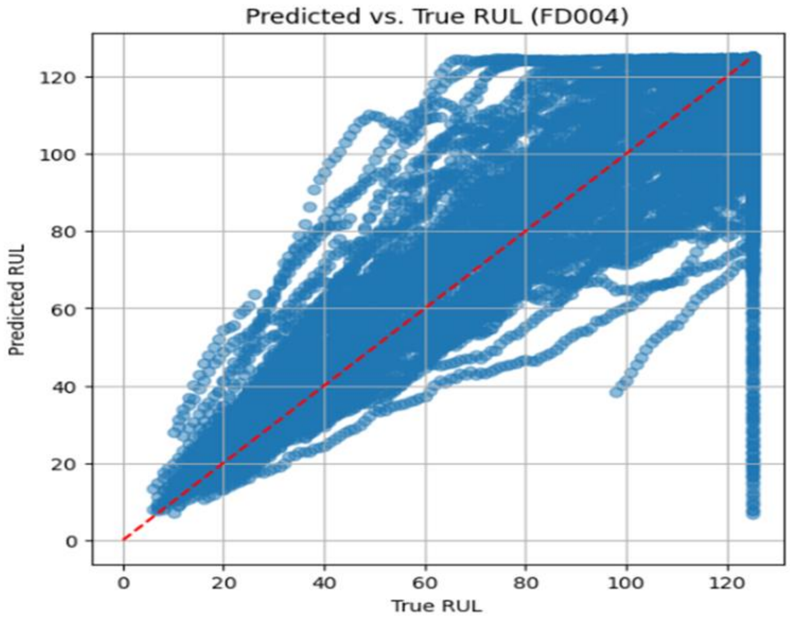

Figure 5: Predicted versus True Remaining Useful Life (RUL) for the BiLSTM–Attention Model on the FD004 Dataset

Training and validation loss curves further highlight model behaviour. The BiLSTM–Attention model converged smoothly with lower final validation error, while the LSTM plateaued earlier with higher residual loss. The convergence behaviour of the models is shown in Figures 6 and 7, where the training and validation loss curves reveal differences in stability and generalization between the LSTM and BiLSTM–Attention frameworks.

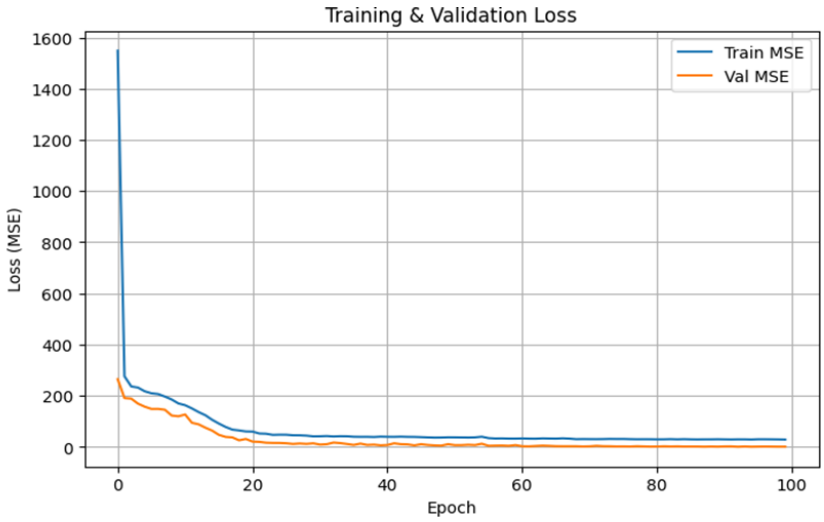

Figure 6: Training and Validation Loss (MSE) of the Proposed Model on the FD004 Dataset

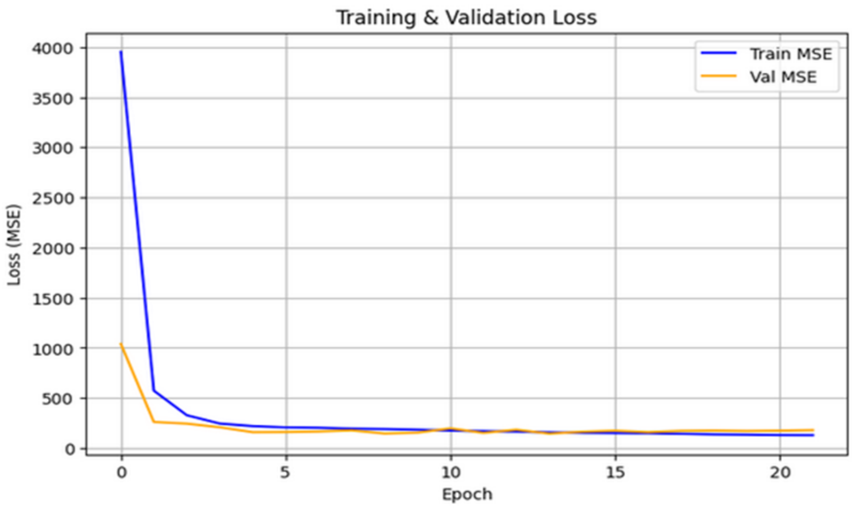

Figure 7: Training and Validation Loss (MSE) of the LSTM Model on the FD004 Dataset

Error distributions confirm this pattern: the BiLSTM–Attention model produced errors tightly centered around zero, indicating stable and unbiased predictions, while the LSTM showed a wider spread with higher variance. Figures 8 and 9 illustrate the error distributions of the baseline LSTM and BiLSTM–Attention models on the FD004 dataset, providing insights into prediction variance and bias.

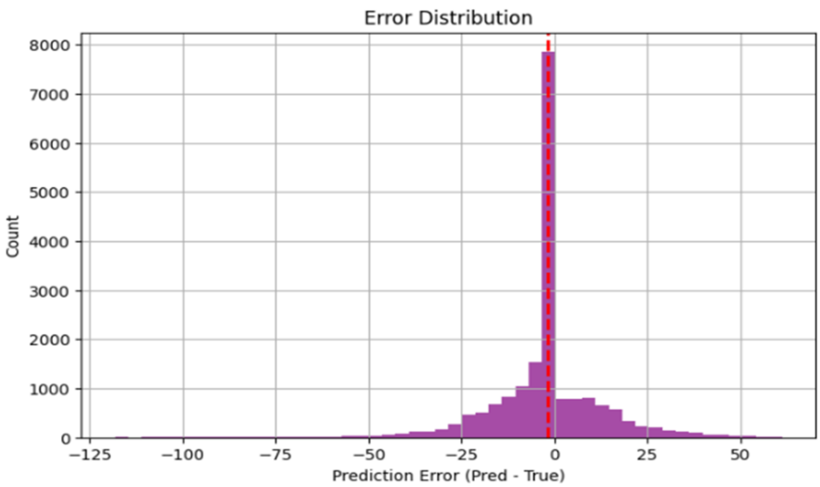

Figure 8: Error Distribution of BiLSTM Predictions (FD004)

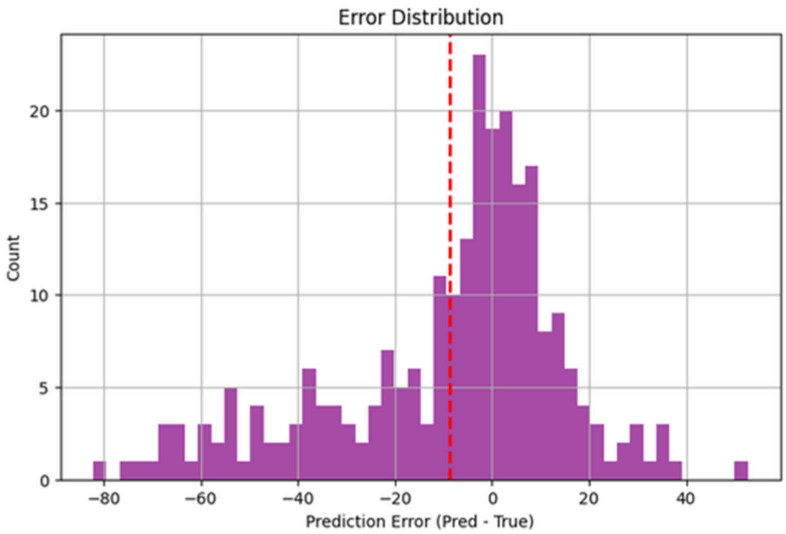

Figure 9: Error Distribution of LSTM Predictions (FD004)

Finally, Figures 10 and 11 illustrate the degradation trajectories of a sample engine from FD004, showing that the BiLSTM–Attention predictions closely follow the true RUL curve while the baseline LSTM fluctuates significantly in mid-life cycles, confirming the superior stability and robustness of the proposed framework.

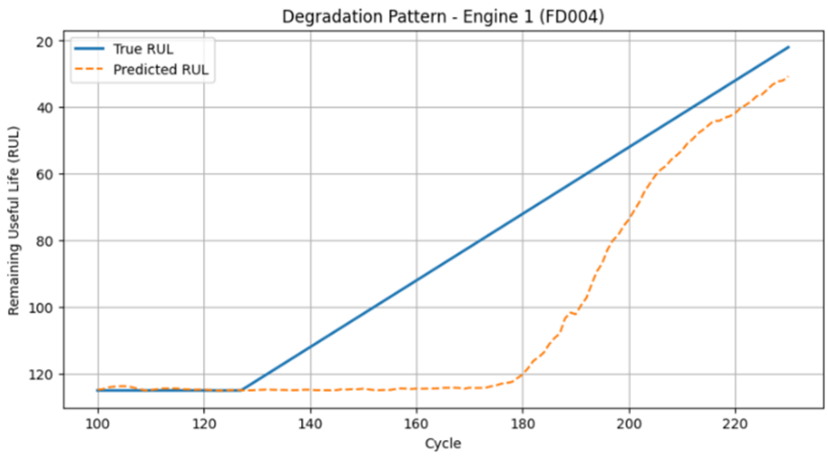

Figure 10: Degradation Pattern of Engine 1 Using the Proposed Model

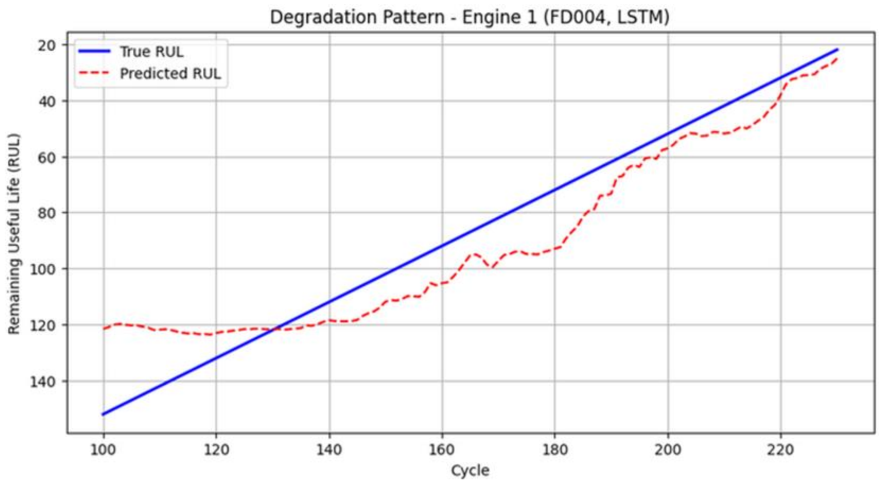

Figure 11: Degradation Pattern of Engine 1 Using the LSTM Baseline

These visualization results complement the quantitative analysis, confirming that attention enhances temporal stability and mitigates late prediction penalties across varying operating conditions.

Discussion of Findings and Implications

The results can be explained by two key design choices. First, the BiLSTM captures bidirectional temporal dependencies, enabling the model to preserve long-term degradation patterns. Second, the attention mechanism highlights the most informative sensors and time steps, reducing noise and late prediction penalties. Such architectural choices are consistent with prior prognostics research emphasizing the complementary strengths of bidirectional recurrent models and attention for long-sequence learning (Dida et al., 2025; Li et al., 2025). This dual-stage design yields a richer and more stable representation of engine health than either recurrent networks or attention alone.

From a practical perspective, the improved predictive accuracy supports reliable condition-based maintenance. By reducing both premature and delayed predictions, the framework lowers maintenance costs, extends engine service life, and enhances aviation safety. These findings align with recent studies showing that hybrid recurrent–attention architectures outperform traditional deep learning approaches for complex prognostic tasks (Dida et al., 2025; Elsherif et al., 2025; Fan et al., 2024).

CONCLUSION

This study introduced a BiLSTM–multi-head attention framework for turbofan RUL prediction and achieved clear improvements over baseline methods. On the C-MAPSS dataset, the model recorded MAE 11.42 / RMSE 15.05 on FD001 and MAE 9.45 / RMSE 15.52 on FD004, outperforming the baseline LSTM and reducing late-prediction penalties. These results demonstrate that integrating bidirectional sequence modelling with adaptive attention improves prediction stability and accuracy, particularly under multi-condition scenarios, offering practical value for condition-based maintenance through earlier fault detection, optimized scheduling, and enhanced operational safety. Nevertheless, the approach is limited by its reliance on a single dataset and the absence of uncertainty quantification, which is essential for real-world prognostics. Future work will extend the framework to multimodal inputs, investigate transformer-based long-sequence models, and incorporate digital-twin simulations and uncertainty-aware prediction methods to support deployment in operational aviation environments.

REFERENCES

Adryan, F. A., & Sastra, K. W. (2021). Predictive Maintenance for Aircraft Engine Using Machine Learning: Trends and Challenges. Avia, 3(1), 37–44. https://doi.org/10.47355/avia.v3i1.45

Adryan Fitra Azyus, Sastra Kusuma Wijaya, & Budhy Kurniawan. (2025). Remaining Useful Life Prediction of Turbofan Engines Using CNN-GRU: A Comparative Study on C-MAPSS Dataset. Journal of Mechanical, Civil and Industrial Engineering, 6(1), 19–27. https://doi.org/10.32996/jmcie.2025.6.1.3

Andenyangtso, N. E., Enyinna, R. D., & Bonet, M. U. (2024). Effective Maintenance Of Aircraft Antiskid Brake System. 6(1), 19–29

Asif, O., Haider, S. A. L. I., Naqvi, S. R., Zaki, J. F. W., Kwak, K. S., & Islam, S. M. R. (2022). A Deep Learning Model for Remaining Useful Life Prediction of Aircraft Turbofan Engine on C-MAPSS Dataset. IEEE Access, 10(August), 95425–95440. https://doi.org/10.1109/ACCESS.2022.3203406

AsifKhan, K. M., M. A., Menouar, H., Eldeeb, A., Abu-Dayya, A., & Salim, F. D. (2022)., Menouar, H., Eldeeb, A., Abu-Dayya, A., & Salim, F. D. (2022). On the Detection of Unauthorized Drones - Techniques and Future Perspectives: A Review. IEEE Sensors Journal, 22(12), 11439–11455. https://doi.org/10.1109/JSEN.2022.3171293

Azyus, A. F. (2022). Determining the Method of Predictive Maintenance for Aircraft Engine Using Machine Learning. Journal of Computer Science and Technology Studies, 4(1), 01–06. https://doi.org/10.32996/jcsts.2022.4.1.1

Cohen, J., Huan, X., & Ni, J. (2021). Fault Prognosis of Turbofan Engines : Eventual Failure Prediction and Remaining Useful Life Estimation. ArXiv Preprint ArXiv, 1–10

Dida, M., Cheriet, A., & Belhadj, M. (2025). Remaining Useful Life Prediction Using Attention-LSTM Neural Network of Aircraft Engines. International Journal of Prognostics and Health Management, 16(2), 1–12. https://doi.org/10.36001/ijphm.2025.v16i2.4274

Donatus, R. E., Ubadike, O. C., Bonet, M. U., & Iyaghiba, S. D. (2025). Automated Aircraft Structural Defect Detection Using Deep Learning and Computer Vision. MEKATRONIKA JOURNAL OF MECHATRONICS AND INTELLIGENT MANUFACTURING, 7(2), 108–123

Elsherif, S. M., Hafiz, B., Makhlouf, M. A., & Farouk, O. (2025). A deep learning-based prognostic approach for predicting turbofan engine degradation and remaining useful life. Scientific Reports, 15(1), 1–31. https://doi.org/10.1038/s41598-025-09155-z

Ensarioğlu, K., İnkaya, T., & Emel, E. (2023). Remaining Useful Life Estimation of Turbofan Engines with Deep Learning Using Change-Point Detection Based Labeling and Feature Engineering. Applied Sciences (Switzerland), 13(21). https://doi.org/10.3390/app132111893

Fan, Z., Li, W., & Chang, K. C. (2024). A Two-Stage Attention-Based Hierarchical Transformer for Turbofan Engine Remaining Useful Life Prediction. Sensors, 24(3). https://doi.org/10.3390/s24030824

Graves, R., Pankaj, P., Kuruvilla, V. J., Johnson, R., & Inoue, M. (2023). Data-Driven Prognostics and Diagnostics of Industrial Machinery — A Turbofan Engine Case Study. PHM Society Asia-Pacific Conference, 4(1), 1–6. https://doi.org/10.36001/phmap.2023.v4i1.3690

Kumar, S., Raj, K. K., Cirrincione, M., Cirrincione, G., Franzitta, V., & Kumar, R. R. (2024). A Comprehensive Review of Remaining Useful Life Estimation Approaches for Rotating Machinery. MDPI Energies, 1–46

Li, M., Zhu, L., Ke, M. L. and, & Ting. (2025). Remaining Useful Life Prediction of Airplane Engine Based on Bidirectional Mamba and Causal Discovery. Sensors, 25(11), 3429. https://doi.org/10.3390/s20164537

Liu, Y., & Wang, X. (2021). Deep & Attention : A Self-Attention based Neural Network for Remaining Useful Lifetime Predictions. 98–105

Liu, Z., Zheng, X., Xue, A., Ge, M., & Jiang, A. (2024). Multi-Head Self-Attention-Based Fully Convolutional Network for RUL Prediction of Turbofan Engines. MDPI ALGORITHMS, 17

Muneer, A., Taib, S. M., Fati, S. M., & Alhussian, H. (2021). Deep-learning based prognosis approach for remaining useful life prediction of turbofan engine. Symmetry, 13(10), 1–21. https://doi.org/10.3390/sym13101861

Ohoriemu, O. B., & Ogala, J. O. (2024). Integrating artificial intelligence and mathematical models for predictive maintenance in industrial systems. Fudma Journal of Sciences, 8(3), 501–505

Pan, Y., Kang, S., & Kong, L. (2025). Remaining useful life prediction methods of equipment components based on deep learning for sustainable manufacturing: a literature review. Artificial Intelligence for Engineering Design, Analysis and Manufacturing, 1–17

Remadna, I., Terrissa, L. S., Ayad, S., & Zerhouni, N. (2021). Rul estimation enhancement using hybrid deep learning methods. International Journal of Prognostics and Health Management, 12(1). https://doi.org/10.36001/IJPHM.2021.V12I1.2378

Wu, F., Wu, Q., Yusong, T., & Xinghua, X. (2024). Remaining Useful Life Prediction Based on Deep Learning: A Survey. SENSORS, 1–30.