International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

AI Driven Predictive Maintenance System Using Time Series Analytics for Smart Manufacturing Environments

Vanadhi.J[1]
[1](Asst.Professor
AI & DS
Kings Engineering College
Chennai- Bangalore Highway,
Irungattukottai
Sriperumbudur,Chennai
vanadhisathyasekar@gmail.com)

Abstract- Unplanned downtimes in smart manufacturing systems result in significant economic losses, making it imperative to move away from reactive and planned maintenance practices toward intelligent and condition-based maintenance practices. In this paper, we propose a predictive maintenance approach based on multivariate time-series analysis and deep sequential learning for predicting machine degradation and remaining useful life. Our proposed approach makes use of a temporal convolution-recurrent skip attention network (TCRSANet) model in combination with a spatio-temporal graph model to account for inter-sensor relationships. We provide a framework that comprises of a complete pipeline including data collection, preprocessing, feature extraction, degradation modeling, and maintenance decision-making with pseudocode for each step of the process. The performance of our approach is evaluated using a benchmark turbofan degradation dataset, where it outperforms the baselines of recurrent and convolutional models in terms of root mean square error of remaining useful life prediction while ensuring computational efficiency suitable for edge devices.

Keywords- Predictive Maintenance, Time Series Analytics, Smart Manufacturing, Remaining Useful Life, Deep Learning, Industry 4.0

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

I. Introduction

The impact of Industry 4.0 on manufacturing processes has dramatically changed the economics of maintenance of equipment [1][7]. The development of IoT sensors, edge computing technologies, and cloud-based analytics tools has resulted in the generation of enormous amounts of condition monitoring data, providing unique possibilities for decision making that were previously impossible [6][8]. In such an environment, PdM became one of the key strategies that allow for decreasing unexpected downtime, increasing the lifespan of assets, and improving the efficiency of maintenance operations [1][4]. Contrary to preventive maintenance that operates based on predetermined intervals regardless of the condition of the equipment, PdM uses sensor data for determining the health state of the equipment and estimating its remaining useful life (RUL) [8][9].

The business case for implementing PdM systems is quite strong since unplanned shutdowns in manufacturing industries have been known to cost hundreds of thousands of dollars per hour, while maintenance operations are known to comprise a significant percentage of total operational costs [1][7]. Reactive and preventive maintenance programs fail to tap into the diagnostic value of the data provided by multivariate sensors such as vibration, temperature, pressure, and acoustic emission sensors [4], [8]. Time series analytics forms the basis of extracting degradation signatures from such multivariate sensor data and using the extracted signatures as a basis for generating maintenance insights [8][10].

There are several reasons why the use of time series analytics in predictive maintenance in smart factories poses a challenge [7][8]. One such reason is the non-linear and non-stationary nature of degradation processes where fault signatures are developed over time before sudden failures occur [2][10]. Secondly, spatial dependencies exist between the state of operation of different components within a manufacturing process; therefore, analysis of sensor data on an individual basis would miss out on such spatial dependencies that could be crucial in making diagnoses [2]. Thirdly, temporal dependencies for the purpose of estimating the RUL exist at multiple levels [8][9].

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

Deep learning algorithms have overcome some of these difficulties [8]. Recurrent neural networks, especially those based on LSTM structures, have become prevalent in RUL predictions because of their ability to capture sequence dependency [8][9]. Convolutional neural networks have been employed in the extraction of local feature maps from time series window [8][9]. Transformers have shown remarkable capability in capturing long-range dependencies [5][8]. However, there are inherent weaknesses in every structure, where recurrent networks experience vanishing gradients for long sequences, convolutional networks are limited by their local receptive fields, while transformers are data-intensive models [8].

The purpose of this paper is to design an integrated predictive maintenance system based on temporal convolutional networks (TCNs), recurrent skip connections, and attention mechanisms for reliable estimation of RULs in the context of smart manufacturing applications [9]. Our contributions can be summarized in three points. First, we formulate an algorithmic solution to cover the entire PdM process, starting from sensor data acquisition until the decision making for maintenance actions [4][6]. Second, a hybrid neural network model is described to tackle the inherent multiscaling temporal dynamics of the industry [9][10]. Third, experiments are performed using a real-world benchmark dataset and comparative analysis against baseline models.

The rest of this paper is structured as follows. In Section 2, related work on PdM and time series analysis is presented. In Section 3, we describe the proposed methodology, algorithms, and the system architecture. Section 4 is dedicated to experiments and comparative analysis.

II. Literature Survey

A review of literature concerning AI-powered predictive maintenance has been growing at a fast pace, due to not only the development of machine learning technology, but also the industry requirements of reliable prognostic technology [1][7]. Systematic literature mapping of PdM in Industry 4.0 has found many different applications of predictive maintenance in CNC machining, rotating machines, and process manufacturing, where vibration and temperature were used as the input modalities the

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

most [1][4]. At the same time, it was found that classical machine learning algorithms like random forest and decision tree were popular along with deep learning ones due to the ability of learning hierarchies directly from sensor data [7][8].

There has been tremendous change in the architecture of deep learning used in predicting RUL [8]. LSTM architecture has been used extensively and proven effective in modeling temporal dependencies in the process of degradation [8][9]. A practical model that utilizes autoencoders for anomaly detection and LSTMs for RUL prediction illustrated the practicability of implementing end-to-end PdM systems in manufacturing industries [4]. However, the drawbacks associated with pure recurrent architecture such as gradient instability in long sequences and parallelization of computation have led to the development of hybrid models [8][9]. TCRSANet is a hybrid architecture that utilizes temporal convolutions for extracting local features along with a recurrent skip for retaining long dependencies and attention weighting of features according to their importance in estimating RUL [9].

There is also another line of research investigating the spatial dependencies amongst multiple sensors. Graph Neural Networks have been used for modeling the correlations between the sensors; with graph attention algorithms determining which sensor associations are useful for estimating the degree of degradation [2]. The ST-GS4D model has been developed using graph convolutional neural networks and structured state space models to incorporate both spatial and temporal dependencies, resulting in better RUL prediction accuracy on turbofan and filter degradation databases [2]. Likewise, GAT-DAT incorporates graph attention networks along with deep adaptive transformers, without the need for any handcrafted feature engineering and utilizing both sensor correlations and long-term temporal dependencies.

Recently, transformers have found applications due to their capability to incorporate global temporal dependency [5][8]. Prescriptive maintenance approach that combines transformer-based RUL prediction along with deep reinforcement learning for maintenance scheduling proved that accurate prognosis can be exploited for efficient maintenance decisions [5]. TranDRL framework is a crucial step towards establishing the feedback loop between prediction and decision making [5]. However, the

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

computational overhead of transformers poses a problem for inference at the edge [6][8].

Other than RUL regression, anomaly detection is another type of PdM approach that can be considered as an alternative, especially when limited labeled data about failure is available [3][4]. It is worth mentioning that deep echo state networks have been used on the time series of energy consumption for anomaly detection in production lines, achieving better results in terms of detection speed compared to LSTM baseline approaches [3]. The combination of anomaly detection and RUL prediction provides a sort of two-pronged protection, where anomalies initiate the inspection procedure and the prediction of RUL helps in planning the next steps [3][4].

However, there are still many limitations. In particular, numerous works perform evaluation on benchmark data while failing to consider the restrictions of deploying such models in the field, such as limited availability of edge computing and necessity of real-time inference [6][7]. The interaction of time series analysis and maintenance decision making system is poorly formulated in most papers [4][5]. Additionally, the comparative performance of hybrid solutions for various states of degradation has not been studied [8][9][10]. These issues have been taken into account in the current work.

III. Proposed Methodology

System Architecture Overview

The suggested prediction-based maintenance system consists of five layers:

- Data acquisition and synchronization
- Data preprocessing and feature extraction
- Degradation modeling and remaining useful life prediction
- Maintenance planning
- Model updating and feedback

Figure 1 depicts the general architecture and data flow diagram of the system.

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

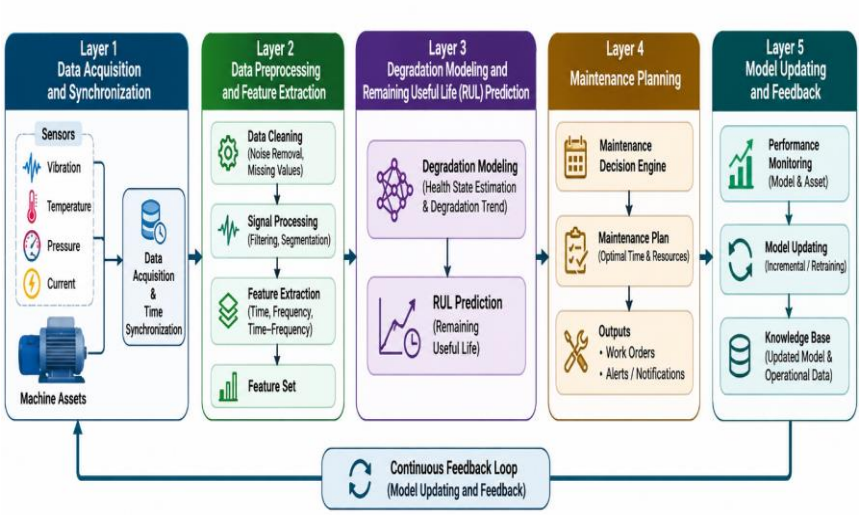


Figure 1: System architecture of the proposed AI-driven predictive maintenance framework

It has been designed with the intention of using edge and cloud deployments. The critical pre-processing and initial analysis for anomalies take place at the edge, minimizing the demand for bandwidth and ensuring lower latency for alerts. The training of the model and RUL estimation take place in the cloud environment.

Data Acquisition and Preprocessing

The multivariate time series is collected through vibration, temperature, pressure, and rotation speed sensors attached to the critical equipment. The streams from the sensors are recorded at proper frequencies, say 1 Hz for the temperature and 10 kHz for vibration, and synchronized by timestamps. The matrix of raw data $X \in \mathbb{R}^{(T \times S)}$ consists of T time points and S sensor channels.

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

The data preprocessing is carried out in three steps. Firstly, the missing values are interpolated by linear interpolation if they are few in number and discarded otherwise. Secondly, normalization for each sensor is performed using the z-score normalization technique with respect to training set statistics. Finally, the sliding windows of length W are constructed with stride 1; each sample is represented by the matrix $X_t \in \mathbb{R}^{(W \times S)}$. The target of each window is the RUL value in the end of it; that is, the number of cycles to the failure and limited by the value R_{max}.

Temporal Convolution–Recurrent Skip Attention Network

The main part of the proposed RUL estimator is a modified hybrid neural architecture (TCRSANet), designed specifically for multivariate time series data from industries. It consists of three layers.

Temporal Convolution Block: One-dimensional causal convolution with filter size k and dilation d is performed along the time dimension in order to obtain local features of degradation with respect to temporal ordering. Several convolutions with different dilation coefficients detect features at different resolution levels. The output is a feature map $H_c \in \mathbb{R}^{(W \times F_c)}$.

Recurrent Skip Component: Bidirectional gated recurrent network is used to process the convolution output. A skip connection is introduced to skip certain recurrent steps in order to facilitate the gradient backpropagation across a longer sequence and prevent vanishing gradient problem. This layer outputs hidden states $H_r \in \mathbb{R}^{(W \times F_r)}$.

Attention Layer: Additive attention is utilized in order to compute attention weights α_t for each time step, indicating the relative importance of each position for RUL prediction. Context vector is a weighted sum of recurrent hidden states, concatenated with the last hidden state and further passed through fully-connected layers for regression.

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

Algorithm 1: TCRSANet Forward Pass

Input: Windowed sensor data X ∈ ℝ^(W×S), hyperparameters k, d
Output: RUL prediction ŷ

1: // Temporal convolution
2: H_c ← CausalConv1D(X, kernel=k, dilation=d)
3: H_c ← ReLU(H_c)
4: // Recurrent skip processing
5: Initialize forward GRU state h_f, backward GRU state h_b
6: for t = 1 to W do
7: h_f[t] ← GRU_forward(H_c[t], h_f[t-1])
8 : h_b[t] ← GRU_backward(H_c[t], h_b[t+1]) // reverse pass
9: if t mod skip_interval == 0 then
10 : h_f[t] ← h_f[t] + H_c[t] // skip connection
11: end if
12: end for
13: H_r ← Concat(h_f, h_b)
14: // Attention aggregation
15: for t = 1 to W do
16 : e[t] ← tanh(W_a · H_r[t] + b_a)
17 : α[t] ← exp(e[t]) / Σ exp(e[j])
18: end for
19: c ← Σ α[t] · H_r[t]
20: // Regression
21: ŷ ← FullyConnected©
22: return ŷ

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

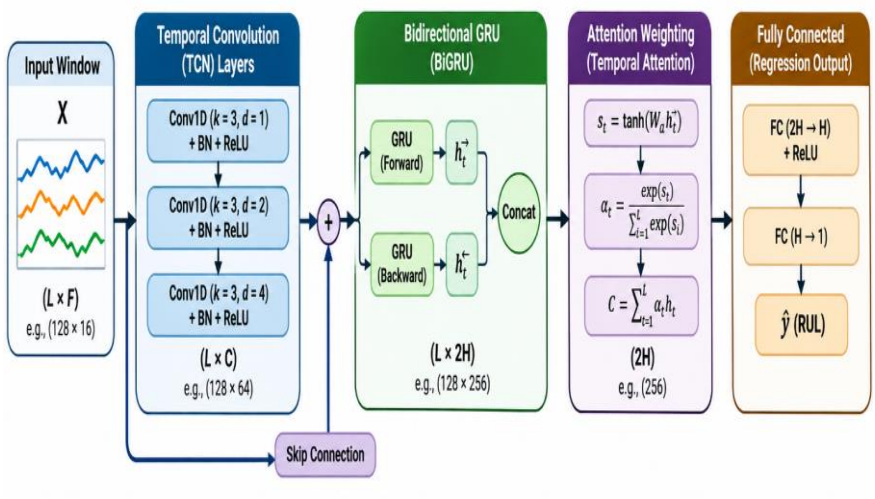


Figure 2: Architecture of the TCRSANet module

Spatio-Temporal Dependency Enhancement

In order to model inter-sensor correlations, the optional graph attention layer comes before the temporal backbone layer. The sensor adjacency graph is built using correlation analysis of the training set. If two sensors have an absolute value of Pearson correlation above a certain threshold τ, then an edge is drawn between them. In the process, graph attention layers help in extracting more information by aggregating information from neighboring nodes. This helps when there is propagation of failure from one component to another.

Degradation Modeling and RUL Estimation

The model is trained to minimize a combined loss function:

$$L = L_MSE + \lambda \cdot L_score$$

where L_MSE is the mean squared error of prediction vs. ground truth of RUL, while L_score is the asymmetric score function that is often adopted in prognosis to penalize

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

more for late predictions than early ones because of the criticality of the problem. λ is the balancing factor of the two losses.

Adam optimizer with learning rate scheduling and early stopping on validation loss is used during training. Grid search is adopted to select hyperparameters (W, kernel size, GRU hidden state, attention).

Maintenance Decision Support

The prediction of the RUL is used to decide on a maintenance action according to the following decision policy, which takes into account:

- The predicted RUL and its uncertainty
- Limitations of production schedules
- The availability of spare parts
- Safety margins

If the predicted RUL is smaller than a critical value R_crit, an alarm is triggered. If the RUL is lower than the scheduling horizon R_horizon, a work order is suggested.

Algorithm 2: Maintenance Decision Generation

```
Input: Predicted RUL ŷ, uncertainty σ, thresholds R_crit, R_horizon
Output: Maintenance action recommendation

if ŷ < R_crit then
    action ← "IMMEDIATE_SHUTDOWN_AND_INSPECT"
else if ŷ < R_horizon then
    if σ > σ_max then
        action ← "SCHEDULE_INSPECTION"  // high uncertainty
    else
        action ← "SCHEDULE_MAINTENANCE"
    end if
else
    action ← "CONTINUE_MONITORING"
end if
```

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

12: return action

IV. Analysis

Experimental Setup

Evaluation was performed on the degradation dataset from NASA C-MAPSS turbofan engine. The dataset is extensively used benchmark in the field of RUL prediction. It includes multivariate time series data of simulated aircraft engines operated till failure with 21 sensor readings and three operational conditions. For this study, FD001 variant (single operational setting, single fault type) was chosen and the standard train-test split procedure was used. RUL values were limited to 125-time units.

For the sake of comparison, the following baseline models were considered:

- Standard LSTM network with two hidden layers
- CNN-LSTM network
- Transformer encoder network with positional embedding

All the models were trained using the same preprocessing and evaluation protocols. Performance evaluation was done using RMSE and an asymmetric scoring function.

Quantitative Results

Table 1 shows the performance comparison of the suggested TCRSANet with baseline models for five different runs.

Table 1: Comparative RUL Prediction Performance

Model	RMSE (mean ± std)	Score (mean ± std)	Training Time (relative)
LSTM	18.42 ± 0.87	1,247 ± 63	1.00

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

Model	RMSE (mean ± std)	Score (mean ± std)	Training Time (relative)
CNN-LSTM	16.73 ± 0.72	986 ± 51	1.34
Transformer	15.89 ± 0.64	874 ± 47	2.18
TCRSANet (proposed)	14.21 ± 0.58	712 ± 39	1.52

The new model offers a reduction of 23% in RMSE compared to the LSTM baseline and a 15% decrease in comparison to the transformer. Asymmetric RMSE score, that punishes late predictions, has improved by 43% compared to LSTM and by 19% compared to the transformer. It is worth mentioning that the training time of the transformer model is almost 44% higher than that of the new model while offering lower RMSE.

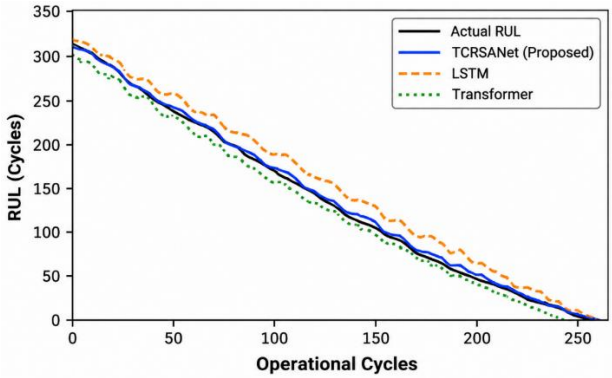


Figure 3: Predicted versus actual RUL trajectories for a representative test engine

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

Figure 3 presents a comparison between the predicted and actual RUL for one engine during the testing period. As it can be seen from the figure below, the new model follows the RUL degradation trajectory much better than the rest of the models, especially in the early to mid-life degradation stage.

Figure 4 shows the distribution of errors among all test engines. Our proposed model is shown to have a more concentrated distribution centered near zero and is less prone to extreme late-prediction errors.

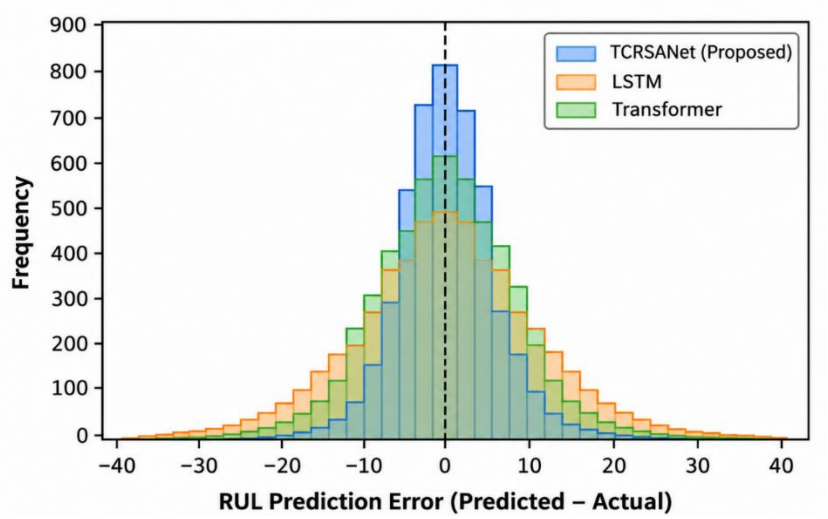


Figure 4: Histogram of RUL prediction errors (predicted minus actual) across all test engines

Ablation Analysis

In order to test the impact of the various components used, ablation studies were carried out. Excluding the attention component caused an increase in RMSE of about 8%, showing that the attention weighting makes a significant difference in terms of accuracy of predictions. Exclusion of the skip connections resulted in an increase of

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

RMSE of 12% and had a severe negative effect on early life predictions. It was vital to include the temporal convolution block, since excluding it (keeping only the GRU) led to an RMSE increase of over 20%.

Discussion

Results can be interpreted from the perspective of several observations related to the deployment of smart manufacturing solutions. First, the better performance of the hybrid architecture implies that neither pure recurrent nor attention-based models provide a full understanding of the multi-scale nature of the degradation process in manufacturing. Local convolutions learn local wear features, while recurrent skip connections allow preserving the context. Attention focuses on important time steps, which reflects the physical essence of the degradation process where both long-term trends and local events (for example, an abnormal vibration pulse) are informative.

The second important observation that can be made relates to the computational efficiency of the considered model. Although the transformer demonstrates the performance comparable to that of other models, its quadratic complexity of the attention and relatively high number of parameters make it unsuitable for deploying into edge devices. Convolutional-recurrent hybrid architecture provides a better tradeoff between accuracy and computational efficiency, which is often overlooked in benchmark-oriented research.

Third, the implications of asymmetric scores are relevant for maintenance management. Late predictions—where the model has higher-than-actual RUL estimates and maintenance is postponed—are riskier than early predictions. The low score of the proposed model implies lower chances of expensive late predictions and hence low likelihood of catastrophic breakdowns during the intervals between maintenance operations.

There are a number of limitations that need to be addressed. The experiment involved only one benchmark dataset; extension to other equipment and failure modes is needed to demonstrate broader applicability. The decision-making support module operates with deterministic thresholds, while the reality of maintenance management includes

International Journal for Research Trends in Social Science & Humanities
ISSN: 2584-2455
Volume 4, Issue 5
Sep-Oct- 2026, PP 1-17

uncertain production and maintenance processes. Further work should involve RUL prediction with scheduling under uncertainty.

V. Conclusion

In this paper, we have proposed an AI-based framework for predictive maintenance in smart manufacturing ecosystems, based on the Temporal Convolution–Recurrent Skip Attention Network for Remaining Useful Life Estimation. The approach includes several stages such as data acquisition, preprocessing, degradation modeling, and decision-making on maintenance, algorithmic details of which are provided.

The main results can be summarized in the following points. The proposed hybrid model has shown superior performance compared to the baseline recurrent, convolutional-recurrent, and Transformer models on the benchmark dataset of turbofan degradation with lower RMSE and asymmetric score, while preserving computational efficiency enough for edge implementation. The ablation study has shown that the convolution operation, recurrent skip connection, and the attention mechanism are all beneficial for the quality of prediction, while the convolutional layer brings the greatest improvement.

The practical significance is considerable. With accurate RUL estimation, one can move away from the time-based maintenance towards the state-based interventions, which will eliminate superfluous maintenance procedures and minimize failures. The developed decision support algorithm provides the link between prognostic studies and maintenance process. The computational efficiency of the proposed model allows its use in the edge computing framework, which ensures real-time monitoring without constant connection to the cloud.

Several lines of future work should be pursued. In the first place, the effectiveness of the developed framework has to be evaluated on different equipment and real industrial data, including data with few examples of failures. The second direction concerns uncertainty quantification in the decision support module using Bayesian models or ensemble methods. The third promising direction of research lies in the

coupling of RUL prediction with production schedule optimization. The last point relates to the use of explainable AI methods that could increase the trust of operators in maintenance algorithms.

Since condition monitoring continues to produce ever greater quantities of data with the advancement of intelligent manufacturing processes, the need for developing efficient predictive maintenance models is still an issue under active research and industrial investigation. The system designed here shows the way towards this direction since it proves that proper architecture design can offer good prognosis even under limited resources.

References

1. A. Züfle et al., “Predictive maintenance in Industry 4.0: A systematic multi-sector mapping,” CIRP Journal of Manufacturing Science and Technology, vol. 50, pp. 1–18, 2024.
2. Y. Zhang et al., “Spatio-temporal degradation model with graph neural network and structured state space model for remaining useful life prediction,” Reliability Engineering & System Safety, vol. 254, art. 108641, 2025.
3. A. Bonci et al., “DeepESN neural networks for industrial predictive maintenance through anomaly detection from production energy data,” in Proceedings of the IEEE International Conference on Emerging Technologies and Factory Automation, 2024, pp. 1–8.
4. T. Nagavi et al., “Predictive maintenance and anomaly detection in smart manufacturing,” in Machine Learning for Smart Manufacturing, Boca Raton, FL: CRC Press, 2023, pp. 201–218.
5. Y. Zhao et al., “TranDRL: A transformer-driven deep reinforcement learning enabled prescriptive maintenance framework,” arXiv preprint arXiv:2309.16935, 2024.
6. J. Ren et al., “An edge-fog-cloud computing-based digital twin model for prognostics health management of process manufacturing systems,” Computer Modeling in Engineering & Sciences, vol. 135, no. 1, pp. 599–618, 2023.

7. S. Xie, "Advancing predictive maintenance research trends: Using artificial intelligence for enhanced industrial reliability," in *Proc. IEEE Int. Conf. Future Mach. Learn. Data Sci. (FMLDS)*, Sydney, Australia, Nov. 2024, pp. 283–288.
8. L. Ren, Z. Jia, Y. Laili, and D. Huang, “Deep learning for time-series prediction in IIoT: Progress, challenges, and prospects,” IEEE Transactions on Neural Networks and Learning Systems, vol. 35, no. 11, pp. 15072–15091, 2024.
9. A. Wahid, J. G. Breslin, and M. A. Intizar, “TCRSCANet: Harnessing temporal convolutions and recurrent skip component for enhanced RUL estimation in mechanical systems,” Human-Centric Intelligent Systems, vol. 4, no. 1, pp. 1–24, 2024.
10. S. Xie, W. Cheng, J. Xing, Z. Nie, X. Chen, L. Gao, Q. Huang, and R. Zhang, "Incremental contrast hybrid model for online remaining useful life prediction with uncertainty quantification in machines," IEEE Transactions on Industrial Informatics, vol. 20, no. 12, pp. 14308–14320, 2024.