

Explainable remaining useful life prediction of rolling bearings using open-source vibration data

Kai Wang *

School of Automotive Engineering, Henan College of Industry & Information Technology, Jiaozuo, China

Article Info

Article History:

Received 14 July 2026

Accepted 17 Sep 2026

Keywords:

Remaining useful life (RUL) prediction;
Bearing prognostics;
Vibration data analysis;
Condition monitoring;
Explainable AI;
Rolling bearings

Abstract

Remaining useful life (RUL) prediction is more useful for maintenance decision-making when the basis of a warning can also be inspected. This study develops an explainable RUL prediction framework for rolling bearings using the open-source NASA IMS 1st_test run-to-failure dataset. Each one-second vibration record is divided into 1024-sample windows, from which RMS, Peak-to-Peak, Kurtosis, and Crest Factor are extracted and aggregated as a four-feature observation. Ten consecutive observations form the input to a lightweight gated recurrent unit (GRU) regressor. To avoid imposing an arbitrary degradation-onset boundary on the prediction target, normalized RUL is defined directly from elapsed life and decreases continuously from 1 to 0. Using the first 70% of the run for training and the final 30% for chronological testing, five random initializations give an MAE of 0.329 ± 0.084 and an RMSE of 0.351 ± 0.098 for the GRU model, compared with MAE/RMSE values of $0.348 \pm 0.128/0.393 \pm 0.149$ for the 1D-CNN and $0.363 \pm 0.011/0.386 \pm 0.013$ for the LSTM. Gradient sensitivity, Integrated Gradients, SHAP, and permutation analysis consistently identify Kurtosis as the most influential late-life feature, while a feature-deletion test produces the largest average MAE increase when Kurtosis is removed. The study is intentionally limited to one fixed-speed, fixed-load IMS run and should therefore be interpreted as controlled benchmark validation rather than vehicle-level validation.

© 2026 MIM Research Group. All rights reserved.

1. Introduction

Background and Problem Formulation

Rolling bearings in vehicle drivetrains and industrial rotating machinery operate under cyclic loading, friction, and speed variations. Progressive damage such as pitting or spalling can increase vibration and, if undetected, lead to unplanned downtime or secondary damage. Remaining useful life (RUL) prediction seeks to estimate the useful operating time remaining before a defined end-of-life condition and is therefore an important element of condition-based and predictive maintenance [1-3]. Vibration measurements are widely used for this purpose because they contain both slowly evolving energy changes and impulsive signatures associated with bearing degradation.

A numerical RUL estimate alone, however, provides limited information about why a warning has been issued. For maintenance personnel, an inspectable link between the prediction and degradation-related vibration indicators can improve traceability and support engineering review. This work therefore considers RUL prediction together with feature-level explanation, while deliberately limiting its experimental claims to a controlled open-source bearing benchmark.

*Corresponding author: wangkai95956@outlook.com

DOI: <http://dx.doi.org/10.17515/resm2026-1869vk0714rs>

Res. Eng. Struct. Mat. Vol. x Iss. x (xxxx) xx-xx

1.1 Related Works and Literature Gap

With the rapid development of data acquisition and intelligent computing, data-driven vibration analysis has become a major direction in RUL research. Hybrid architectures combining convolutional neural networks (CNNs) with recurrent models such as LSTM, BLSTM, and GRU are widely used because convolution can capture local patterns while recurrent units' model temporal dependence [4,5]. To improve generalization under limited labelled data or operating-condition shifts, transfer-learning and semi-supervised strategies have been explored [6-8]. Graph-based models provide another route to learning coupled degradation representations [9]. Open benchmark data such as NASA IMS remain important for controlled method comparison and reproducibility [10-12].

Recent RUL research increasingly combines prediction with uncertainty quantification, operational decision-making, and physical knowledge. Deep learning-stochastic ensembles have been linked to predictive-maintenance scheduling and dynamic mission-abort policies [13], while RUL-aware physics-informed Bayesian digital twins connect probabilistic prognostics with safety-constrained control [14]. Deep ensembles have also been used for dynamic burn-in optimization with precise lithium-ion battery RUL prediction [15]. In parallel, explainable RUL frameworks increasingly incorporate degradation physics and attribution mechanisms: Kim et al. developed a physics-informed explainable RUL framework [16], and Zhang and Wang combined rolling-bearing RUL uncertainty prediction with SHAP-based interpretation [17]. These developments show that explain ability is no longer limited to post-hoc visualization. Against this background, the present work addresses a narrower and deliberately transparent question: whether a compact set of physically interpretable time-domain vibration indicators can support temporal RUL prediction while providing a directly inspectable feature-level explanation on an open bearing benchmark.

Related studies also address interpretable intelligent fault diagnosis in rotating machinery [18], explainable predictive maintenance [19], general condition-monitoring RUL modelling [20], predictive maintenance for multi-rotor systems [21], explainable Industry 4.0 fault detection [22], and explainable aero-engine RUL estimation [23]. Work on tool-wear RUL prediction [24] and degradation-start detection [25] further illustrates the breadth of prognostics and interpretability research. Together with the recent studies above, this literature motivates transparent prognostic modelling while leaving room for a compact, feature-level analysis on an open rolling-bearing run-to-failure benchmark.

1.2 Research Objectives and Primary Contributions

The study is intentionally confined to the NASA IMS 1st_test run-to-failure bearing dataset and does not claim a universally superior neural architecture or direct validation for vehicle duty cycles. Its contribution is a compact and reproducible workflow that links physically interpretable vibration indicators, temporal RUL regression, and feature-attribution checks. The central question is whether a small set of established degradation indicators can provide useful RUL information while retaining an inspectable basis for late-life warnings.

Data and features. The analysis uses the IMS 1st_test records at their native acquisition resolution; no artificial up sampling is applied. Bearing 3 is selected because an inner-race defect was reported at the end of this run, and Channel 5 is used as the primary vibration measurement. Each one-second record contains 20,480 samples. It is divided into twenty non-overlapping 1024-sample windows, and four time-domain indicators—RMS, Peak-to-Peak, Kurtosis, and Crest Factor—are computed for each window and averaged to form one chronological feature observation.

Prediction model and baselines. Ten consecutive four-feature observations are organized as a temporal input sequence. A lightweight GRU regressor is used as the proposed sequence predictor and is compared with 1D-CNN and LSTM reference models under the same chronological 70%/30% split. All three models use the same feature sequences, training portion, RUL targets, optimization budget, and evaluation metrics. Results are repeated across five random initializations to expose sensitivity to model initialization.

Explain ability. The numerical RUL estimate is complemented by feature-level attribution. Gradient sensitivity is used as the primary local explanation and is cross-checked with Integrated Gradients, SHAP, permutation sensitivity, and a feature-deletion test. Explanations are evaluated across early-middle-, and late-life stages and across five random initializations. The attribution scores are interpreted as predictive sensitivity of the trained model rather than as causal proof of a physical damage mechanism.

1.3 Chapter Arrangement

Section 2 describes the IMS data, signal preprocessing, RUL target and sequence construction, prediction models, attribution methods, and chronological evaluation protocol. Section 3 reports predictive performance, a data-driven degradation-onset check, feature-deletion results, and attribution robustness. Section 4 discusses engineering interpretation and generalization limitations, and Section 5 concludes the study.

2. Materials and Methods

2.1 Data Source and Study Scope

RUL model development requires observations spanning healthy operation through late-life degradation. The NASA IMS bearing run-to-failure dataset provides a controlled public benchmark with complete vibration histories and is therefore used for reproducible method evaluation. The present study does not treat IMS as a surrogate for vehicle duty cycles. Instead, it uses the dataset to examine the prediction-and-explanation workflow under fixed operating conditions; transfer to variable-speed or vehicle environments remains outside the validated scope of this work.

In IMS Test 1, four Rexnord ZA-2115 double-row roller bearings are mounted on a common shaft and operated at 2000 RPM under a 6000 lb radial load. Each vibration file contains 20480 samples acquired over approximately one second, and 2156 chronological files are available from the run. Bearing 3 developed an inner-race defect by the end of Test 1; therefore, Channel 5, one of the two vibration channels associated with Bearing 3, is used as the primary signal in this study. The first 70% of the chronological observations are reserved for model training and the final 30% are held out for testing. The experimental configuration and data-acquisition procedure are summarized in Fig. 1.




Fig. 1. Schematic diagram of the experimental test rig and data acquisition system

2.2 Signal Preprocessing and Degradation Feature Engineering

The 2156 IMS files are ordered by acquisition time. No temporal resampling or artificial up sampling is performed. For each file, Channel 5 is divided into twenty non-overlapping windows of 1024 samples. Four time-domain indicators are calculated in every window and then averaged across the twenty windows to obtain one four-dimensional feature observation for that acquisition

time. This preserves the chronological degradation trend while keeping the sequence model compact. Four features are computed for each 1024-sample window:

- Root Mean Square (RMS) measures the overall vibration-energy level in the analyzed window. Progressive wear, looseness, or increasing dynamic interaction can raise the RMS level, although RMS is also sensitive to operating-condition changes.
- Peak-to-Peak amplitude is the difference between the maximum and minimum signal values within the window. It is sensitive to extreme excursions and therefore provides complementary information on intermittent impacts and transient vibration events.
- Kurtosis quantifies the impulsiveness and non-Gaussianity of the amplitude distribution. High kurtosis is commonly associated with isolated high-amplitude events and can therefore respond to localized contact or surface damage.
- Crest Factor (CF) is the ratio between the absolute peak amplitude and RMS. It indicates how large the strongest peak is relative to the overall vibration-energy level and therefore complements the scale-sensitive RMS indicator.
- The four indicators jointly represent vibration energy (RMS), amplitude range (Peak-to-Peak), impulsiveness (Kurtosis), and peak-to-energy ratio (Crest Factor). Their chronological evolution is shown in Fig. 2. A data-driven degradation-onset check is superimposed for descriptive analysis; it is not used to construct the RUL target.




Fig. 2. Normalized degradation-feature trends and data-driven change point for IMS Test 1, Bearing 3, Channel 5

2.3 RUL Target and Sequence Construction

To remove dependence on an arbitrarily selected first-prediction-time boundary, the revised RUL target is defined directly from elapsed life rather than by imposing a healthy plateau followed by a manually selected decay stage. Let t denote the chronological observation index and T the final observation index. The normalized RUL target is

$$y_t = \frac{T - t}{T}, 0 \leq t \leq T \quad (1)$$

This definition gives $y_t=1$ at the beginning of the observed run and $y_t=0$ at its end. The former fixed 70% FPT label rule is therefore no longer required. A separate data-driven change-point analysis is used only to characterize when the vibration health trajectory changes most strongly; it does not alter the RUL target.

At each observation t , the four-feature vector is $x_t = [RMS_t, P2P_t, Kurtosis_t, CF_t]^T$. To include recent degradation history, ten consecutive feature vectors are concatenated as $X_t = [X_{t-9}, \dots, X_t] \in \mathbb{R}^{10 \times 4}$.

Feature-wise Min-Max scaling is fitted using only the first 70% training portion and then applied unchanged to the held-out 30% test portion. This prevents future test information from entering the preprocessing stage.

2.4 Sequence-Based RUL Prediction Model

The proposed sequence predictor is a lightweight gated recurrent unit (GRU) regression network. Each normalized 10×4 input sequence is passed to a single GRU layer with 32 hidden units. The final hidden state is mapped through a fully connected layer with 16 ReLU units and a single sigmoid output neuron that produces the normalized RUL estimate $\hat{y}_t \in [0,1]$.

$$h_t = GRU(X_t; \theta_g) \quad (2)$$

$$z_t = ReLU(W_1 h_t + b_1) \quad (3)$$

$$\hat{y}_t = \sigma(W_2 z_t + b_2) \quad (4)$$

The model is trained by minimizing mean squared error, $\mathcal{L}_{RUL} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2$. Adam optimization is used with an initial learning rate of 1×10^{-3} , a batch size of 128, and 150 training epochs. To quantify initialization sensitivity, all reported predictive results are repeated over five random initializations. For comparison, the LSTM baseline contains one 32-unit LSTM layer followed by a sigmoid regression output. The 1D-CNN baseline uses two temporal convolution layers with 16 and 32 channels (kernel size 3), ReLU activations, adaptive average pooling, and a sigmoid regression output. The baselines use the same four-feature sequences, optimizer, training/test split, and RUL target.

2.5 Feature Attribution and Robustness Validation

The primary local explanation uses input-gradient sensitivity. For feature j , sensitivity over a sequence of length L is defined as $A_j = \frac{1}{L} \sum_{\tau=1}^L \left| \frac{\partial \hat{y}_t}{\partial x_{\tau,j}} \right|$. The four sensitivities are normalized as $\omega_j = \frac{A_j}{\sum_{k=1}^4 A_k}$, so that $\sum_{j=1}^4 \omega_j = 1$. These weights quantify model-level predictive sensitivity and are not interpreted as causal contributions to physical bearing damage. Gradient sensitivity is cross-checked using three complementary methods: Integrated Gradients with a zero baseline in the scaled feature space, SHAP expected-gradient attribution using background sequences from the training portion, and permutation sensitivity obtained by permuting one feature trajectory across sequences while retaining the others. A deletion check additionally replaces each feature by its training-set mean and records the resulting change in test MAE. Attribution results are aggregated over early life ($RUL > 0.66$), middle life ($0.33 < RUL \leq 0.66$), and late life ($RUL \leq 0.33$), and over five random model initializations. The complete prediction and feature-attribution framework is summarized in Fig. 3.




Fig. 3. Architecture of the proposed explainable remaining useful life prediction framework

2.6 Experimental Protocol and Evaluation Metrics

To avoid future-information leakage, the 2156 chronological observations are split without random shuffling: the first 70% form the training set and the final 30% form the held-out test set. Min-Max scaling is fitted only on the training portion. The GRU, 1D-CNN, and LSTM models are trained under the same split and repeated across five random initializations. The test portion is not used for parameter fitting.

Prediction performance is evaluated using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) on normalized RUL. MAE reports the average absolute prediction error, whereas RMSE penalizes larger errors more strongly. Results are reported as mean $\pm$ standard deviation across five initializations.

$$MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \quad (5)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2} \quad (6)$$

For the degradation-onset robustness check, a 31-observation rolling mean is applied to the four feature trajectories, followed by robust scaling and one-dimensional principal-component fusion. A two-segment least-squares change point is then selected by minimizing within-segment squared error over candidate locations between 40% and 90% of observed life. This change point is reported descriptively and is not used to create the RUL labels.

3. Results and Discussion

3.1 Model Predictive Performance Baseline Comparison

All predictive results use the held-out final 30% of IMS Test 1. Across five random initializations, the proposed GRU model obtains an MAE of 0.329 ± 0.084 and an RMSE of 0.351 ± 0.098 . The 1D-CNN baseline obtains 0.348 ± 0.128 MAE and 0.393 ± 0.149 RMSE, while the LSTM obtains 0.363 ± 0.011 MAE and 0.386 ± 0.013 RMSE. Thus, the GRU gives the lowest mean error under the common protocol, although the standard deviations also show that initialization affects the GRU and 1D-CNN more strongly than the LSTM.

Table 1. RUL prediction performance over five random initializations

Model	Mean Absolute Error (MAE)	Root Mean Squared Error (RMSE)
1D CNN	0.348 ± 0.128	0.393 ± 0.149
LSTM	0.363 ± 0.011	0.386 ± 0.013
Proposed GRU explainable model	0.329 ± 0.084	0.351 ± 0.098

The comparison is intended to establish that the compact sequence model is competitive on the controlled benchmark rather than to claim universal architectural superiority. The difference between the GRU and the two baselines should be interpreted together with the single-run experimental scope and the observed variation across random initializations.

Late-life prediction remains the most demanding region because the test set contains RUL values below those observed during much of training. RMSE is therefore reported together with MAE so that occasional large late-life errors are not hidden by average absolute error. No conclusion is drawn from these results regarding performance under variable speed, load, temperature, or vehicle duty cycles.

3.2 Degradation-Onset and Feature-Deletion Checks

The former manuscript imposed a 70% first-prediction-time boundary on the RUL target. In the revised protocol this manual boundary is removed: normalized RUL decreases continuously with elapsed life, while degradation onset is examined independently. The data-driven two-segment change-point analysis identifies the strongest transition at approximately 85.2% of the observed life (file timestamp 2003.11.22.17.16.56). The difference between this data-driven location and the former 70% assumption confirms that a fixed FPT should not be treated as a known physical degradation onset.

Feature deletion provides an additional check on whether the attribution ranking corresponds to predictive dependence. Averaged over five GRU initializations, replacing Kurtosis by its training-set mean increases test MAE by 0.0249, compared with increases of 0.0045 for Peak-to-Peak and 0.0015 for RMS; replacing Crest Factor changes MAE by -0.0019 on average. Kurtosis therefore produces the largest degradation in predictive performance when removed, consistent with its top ranking in the attribution analyses.

The deletion result should not be interpreted causally: a feature can be important to the trained predictor because it is correlated with the degradation trajectory. Nevertheless, agreement between feature deletion and multiple attribution methods provides a more robust basis for interpreting model dependence than gradient magnitude alone.

Because the RUL label is now independent of a manually specified FPT, model accuracy is no longer conditioned on choosing 60%, 70%, or 80% as a label breakpoint. The change-point result is retained only as an independent description of when the measured vibration indicators exhibit their strongest distributional transition.

3.3 Explainable Attribution Analysis

Feature attribution is evaluated using Gradient sensitivity, Integrated Gradients, SHAP, and permutation sensitivity. Across the early-, middle-, and late-life stages, all four methods rank Kurtosis as the most influential feature on average. Inter-method Spearman rank agreement is 0.77 in early life and 0.90 in both middle and late life. Across the five random initializations, mean pairwise rank correlations range from 0.64 to 1.00 depending on the method and life stage, indicating moderate-to-high explanation stability rather than a single-seed result.

In late life ($RUL \leq 0.33$), the mean normalized Kurtosis importance is 0.762 for Gradient sensitivity, 0.723 for Integrated Gradients, 0.731 for SHAP, and 0.536 for permutation sensitivity. Peak-to-Peak is generally the second most influential feature, especially under permutation sensitivity. Figure 4 summarizes the late-life attributions, and Table 2 reports the corresponding values.




Fig. 4. Late-life feature attribution/sensitivity comparison across four explanation methods; error bars show variation across five GRU initializations

The attribution pattern is physically plausible but should be interpreted as predictive association. Kurtosis measures impulsiveness and departure from a Gaussian amplitude distribution; its dominant late-life ranking is consistent with increasingly pronounced impact-like events as localized surface damage develops. Peak-to-Peak captures extreme amplitude excursions and therefore complements Kurtosis when intermittent impacts become stronger. RMS represents overall vibration energy and is useful for monitoring broad energy growth, while Crest Factor expresses the peak amplitude relative to RMS and can distinguish isolated peaks from a general rise in vibration level.

The result differs from the original single-gradient interpretation in which RMS appeared dominant. The revised analysis averages explanations over multiple observations and random initializations and cross-checks them with independent attribution mechanisms. The consistent ranking of Kurtosis, together with the feature-deletion result, indicates that the trained model relies strongly on impulsive information in the selected IMS channel. This does not establish that Kurtosis is a universal bearing-health indicator or a causal damage mechanism.

For predictive-maintenance use, these explanations provide context for a low-RUL estimate: an impulse-dominated warning suggests a different signal pattern from one driven primarily by broadband energy growth. Such context can support engineering review, but the attribution does not identify a specific fault location by itself. Frequency-domain or envelope features would be required for more direct differentiation of inner-race, outer-race, and rolling-element fault mechanisms.

Table 2. Mean late-life normalized feature attribution/sensitivity over five GRU initializations

Method	RMS	Peak-to-Peak	Kurtosis	Crest Factor
Gradient	0.043	0.126	0.762	0.069
Integrated Gradients	0.083	0.134	0.723	0.060
SHAP	0.065	0.123	0.731	0.081
Permutation	0.069	0.207	0.536	0.187

4. Discussion and Engineering Generalization Limitations

The revised results support two conclusions within the controlled IMS benchmark. First, compact time-domain features retain sufficient information for a sequence model to follow the observed life trajectory, although initialization variability is non-negligible. Second, explanation robustness improves when gradient sensitivity is considered together with Integrated Gradients, SHAP, permutation sensitivity, and feature deletion rather than being treated as a stand-alone causal measure.

Generalization beyond this experiment has not been established. IMS Test 1 operates at a constant 2000 RPM and a constant 6000 lb radial load, whereas automotive bearings experience speed and torque transients, start-stop operation, temperature variation, lubrication changes, structural transmission paths, and road-induced disturbances. Accordingly, the title and claims of the revised manuscript refer to rolling-bearing benchmark validation rather than validated vehicle-bearing prognosis. Additional IMS runs and datasets with variable speed/load conditions are required before deployment-oriented conclusions can be made.

The explanation is also limited by the selected feature space. RMS, Peak-to-Peak, Kurtosis, and Crest Factor distinguish broad energy and impulsive behaviour but do not directly identify the physical fault location. Attribution values therefore indicate which measured statistics the predictor relies on, not why a particular physical defect formed. More complete diagnostic interpretation would require frequency-domain, envelope, or order-related information together with validation across naturally occurring faults and multiple operating conditions.

5. Conclusions

This study develops a compact explainable RUL workflow for the open-source NASA IMS rolling-bearing benchmark. Vibration records from Bearing 3, Channel 5 are represented using RMS, Peak-to-Peak, Kurtosis, and Crest Factor, arranged into ten-step temporal sequences, and mapped to normalized RUL using a lightweight GRU regressor. The arbitrary 70% FPT label rule used in the previous manuscript has been removed; normalized RUL is defined directly from elapsed life, while a separate data-driven change-point analysis characterizes degradation onset without controlling the target labels.

Under a chronological 70%/30% train-test split and five random initializations, the GRU achieves an MAE of 0.329 ± 0.084 and an RMSE of 0.351 ± 0.098 . The corresponding 1D-CNN values are 0.348 ± 0.128 and 0.393 ± 0.149 , and the LSTM values are 0.363 ± 0.011 and 0.386 ± 0.013 . These results indicate competitive predictive performance on the selected run but do not establish architectural superiority beyond this benchmark.

Explainability is evaluated using four complementary approaches rather than gradient sensitivity alone. Gradient, Integrated Gradients, SHAP, and permutation sensitivity all rank Kurtosis as the most influential feature on average in late life, and the deletion test produces the largest average MAE increase when Kurtosis is removed. The attribution pattern is physically consistent with stronger impulsive activity during advanced degradation, but it is interpreted as model-level predictive dependence rather than physical causality.

The present evidence is limited to a single fixed-speed, fixed-load IMS run. The framework should therefore be regarded as a transparent benchmark-oriented RUL workflow, not a field-validated vehicle prognostic system. Validation on additional run-to-failure bearings and under varying speed, load, temperature, lubrication, and vehicle duty cycles is required before deployment. Future work may also add frequency-domain and envelope indicators to connect RUL explanations more directly to specific bearing fault mechanisms.

Acknowledgments

This report is dedicated to our supervising professors and laboratory colleagues for their guidance, support, and interesting discussions. We would also like to thank the university's computer resources and the open vibration data sets provided by the PHM community.

Funding

Henan Province Science and Technology Research Project Project Name: Dynamic Characteristics and Accuracy Control of Electric Spindle Rotation under Time Varying Load Excitation (Project Number: 262102220050)

References

- [1] Liu X, Zhang Z, Li Z, Wang J, Zhu Y, Ma H. Advancements in bearing health monitoring and remaining useful life prediction: techniques, challenges, and future directions. *Measurement Science and Technology*. 2025;36(3):032003. <https://doi.org/10.1088/1361-6501/ad824a>
- [2] Li W, Chen J, Chen S, Li P, Zhang B, Wang M, et al. A comprehensive review of artificial intelligence-based algorithms for predicting the remaining useful life of equipment. *Sensors*. 2025;25(14):4481. <https://doi.org/10.3390/s25144481>
- [3] Kamat P, Kumar S, Sugandhi R. Vibration-based anomaly pattern mining for remaining useful life (RUL) prediction in bearings. *Journal of the Brazilian Society of Mechanical Sciences and Engineering*. 2024;46(5):290. <https://doi.org/10.1007/s40430-024-04872-4>
- [4] Yang L, Jiang Y, Zeng K, Peng T. Rolling bearing remaining useful life prediction based on CNN-VAE-MBiLSTM. *Sensors*. 2024;24(10):2992. <https://doi.org/10.3390/s24102992>
- [5] Bharatheedasan K, Maity T, Kumaraswamidhas LA, Durairaj M. An intelligent of fault diagnosis and predicting remaining useful life of rolling bearings based on convolutional neural network with bidirectional LSTM. *Sādhana*. 2023;48(3):142. <https://doi.org/10.1007/s12046-023-02179-8>
- [6] Sandgren J. Transfer Learning for Predictive Maintenance in Workshop Environments: Vibration Based Analysis of Mechanical Systems. Master's thesis, Uppsala University, 2025.

- [7] Kim JW, Park KS. Remaining Useful Life Prediction of Bearings via Semi-Supervised Transfer Learning Based on an Anti-Self-Healing Health Indicator. *Sensors*. 2025;25(12):3662. <https://doi.org/10.3390/s25123662>
- [8] Park K, Kim JW. Remaining Useful Life Prediction of Bearing Via Semi-Supervised Transfer Learning Based on Anti-Self-Healing Health Indicator. *SSRN Electronic Journal*. 2025. <https://doi.org/10.2139/ssrn.5218342>
- [9] Jin Z, Wu Z, Zhao J, He D, Zhuang Y. Graph Convolutional Neural Network Algorithms for Bearing Remaining Useful Life Prediction: A Review. *Journal of Failure Analysis and Prevention*. 2025;25(3):1040–1056. <https://doi.org/10.1007/s11668-025-02111-2>
- [10] Sternharz G. Operational Modal Analysis and Prediction of Remaining Useful Life for Rotating Machinery. PhD thesis, Brunel University London, 2023.
- [11] Su H, Lee J. Machine Learning Approaches for Diagnostics and Prognostics of Industrial Systems Using Open Source Data from PHM Data Challenges: A Review. *International Journal of Prognostics and Health Management*. 2024;15(2). <https://doi.org/10.36001/ijphm.2024.v15i2.3993>
- [12] Sayyad S, Kumar S, Bongale A, Kamat P, Patil S, Kotecha K. Data-Driven Remaining Useful Life Estimation for Milling Process: Sensors, Algorithms, Datasets, and Future Directions. *IEEE Access*. 2021;9:110255–110286. <https://doi.org/10.1109/ACCESS.2021.3101284>
- [13] Faizanbasha A, Rizwan U. Deep learning-stochastic ensemble for RUL prediction and predictive maintenance with dynamic mission abort policies. *Reliability Engineering & System Safety*. 2025;259:110919. <https://doi.org/10.1016/j.res.2025.110919>
- [14] Faizanbasha A, Rizwan U, Hussainy ST, Almakayeel N, Asif FB. Remaining useful life aware physics-informed Bayesian digital twin for safety-constrained control in robot-integrated battery manufacturing. *Robotics and Computer-Integrated Manufacturing*. 2026;101:103318. <https://doi.org/10.1016/j.rcim.2025.103318>
- [15] Faizanbasha A, Rizwan U, Hussainy ST, Almakayeel N, Asif FB. Deep learning ensemble for dynamic burn-in optimization with precise Remaining Useful Life prediction for lithium-ion batteries. *Electrochimica Acta*. 2026;548:148001. <https://doi.org/10.1016/j.electacta.2025.148001>
- [16] Kim M, Yoo S, Son S, Chang SY, Oh KY. Physics-informed deep learning framework for explainable remaining useful life prediction. *Engineering Applications of Artificial Intelligence*. 2025;143:110072. <https://doi.org/10.1016/j.engappai.2025.110072>
- [17] Zhang T, Wang H. Explainable remaining useful life uncertainty prediction method for rolling bearing. *Engineering Applications of Artificial Intelligence*. 2025;161:112114. <https://doi.org/10.1016/j.engappai.2025.112114>
- [18] Chen G, Yuan J, Zhang Y, Zhu H, Huang R, Wang F, et al. Enhancing Reliability Through Interpretability: A Comprehensive Survey of Interpretable Intelligent Fault Diagnosis in Rotating Machinery. *IEEE Access*. 2024;12:103348–103379. <https://doi.org/10.1109/ACCESS.2024.3411488>
- [19] Pashami S, Nowaczyk S, Fan Y, et al. Explainable Predictive Maintenance. *arXiv preprint arXiv:2306.05120*. 2023. <https://doi.org/10.48550/arXiv.2306.05120>
- [20] Surucu O. A General Condition Monitoring Model for Remaining Useful Life Estimation Using Degradation Signals. Doctoral dissertation, University of Guelph, 2022.
- [21] Ozkat EC, Bektas O, Nielsen MJ, la Cour-Harbo A. A data-driven predictive maintenance model to estimate RUL in a multi-rotor UAS. *International Journal of Micro Air Vehicles*. 2023;15:1–14. <https://doi.org/10.1177/17568293231175641>
- [22] Asutkar S, Tallur S. An explainable unsupervised learning framework for scalable machine fault detection in Industry 4.0. *Measurement Science and Technology*. 2023;34(10):105123. <https://doi.org/10.1088/1361-6501/ace21c>
- [23] Protopapadakis G, Apostolidis A, Kalfas AI. Explainable and Interpretable AI-Assisted Remaining Useful Life Estimation for Aeroengines. *Proceedings of ASME Turbo Expo 2022*. 2022;Vol. 2:GT2022-80777, V002T05A002. <https://doi.org/10.1115/GT2022-80777>
- [24] Zhang X. Deep Learning Driven Tool Wear Identification and Remaining Useful Life Prediction. PhD thesis, Coventry University, 2020.
- [25] Schaad Concha CI. A Framework for Degradation Start Detection and RUL Prognosis of Physical Assets. Thesis, Universidad de Chile, 2020.