DEVELOPMENT OF PREDICTIVE MAINTENANCE FOR BALL BEARING

Lt Dr M Ramesh, Rathish K N, Sanjai B N, Parthasarathy P V

Department of Mechanical Engineering, Sri Ramakrishna Engineering College, Coimbatore
ramesh.mari@srec.ac.in

Department of Mechanical Engineering, Sri Ramakrishna Engineering College, Coimbatore
rathish.2204111@srec.ac.in

Department of Mechanical Engineering, Sri Ramakrishna Engineering College, Coimbatore
sanjai.2204115@srec.ac.in

Department of Mechanical Engineering, Sri Ramakrishna Engineering College, Coimbatore
parthasarathy.2204104@srec.ac.in

Abstract — Unplanned bearing failure in rotating industrial machinery leads to significant operational downtime, increased maintenance expenditure, and reduced equipment reliability. Conventional maintenance strategies, whether purely reactive or rigidly time-based, fall short in addressing the dynamic nature of bearing degradation. This work presents an artificial-intelligence-driven predictive maintenance framework designed to continuously assess bearing health and estimate its Remaining Useful Life (RUL). The system acquires real-time multi-sensor data encompassing rotational speed (RPM), acoustic levels, temperature, moisture, and three-axis motion, then subjects the data to preprocessing and feature extraction before applying machine learning algorithms for RUL prediction. Random Forest, Support Vector Machine (SVM), and Convolutional Neural Network (CNN) models were trained and evaluated for this purpose. A Flask-based web interface was additionally developed to provide maintenance teams with an accessible, real-time prediction dashboard. The proposed framework demonstrates that fusing information from multiple sensing modalities improves fault-detection accuracy over single-parameter approaches and can serve as a practical decision-support tool for industrial maintenance planning.

Keywords: Predictive maintenance, ball bearing, Remaining Useful Life, Random Forest, Support Vector Machine, Convolutional Neural Network, multi-sensor data fusion, IoT.

I. INTRODUCTION

Modern industry depends heavily on rotating machinery such as pumps, compressors, motors, and turbines, all of which rely on bearings to facilitate smooth rotation and minimise frictional losses. Any progressive deterioration in bearing condition, if left undetected, can trigger a cascade of failures that interrupts production, endangers personnel, and inflates maintenance costs substantially.

Traditional maintenance philosophies offer two broad paths. The first is corrective maintenance, which defers action until a component actually fails; the second is scheduled preventive maintenance, which services components at fixed calendar intervals regardless of their actual state. Both approaches carry inherent inefficiencies. Reactive strategies incur the full cost of unexpected breakdown, while fixed-interval strategies waste resources on components that still have considerable service life remaining and may inadvertently introduce new faults during unnecessary disassembly.

Predictive maintenance overcomes these shortcomings by continuously monitoring key operating parameters and using data-driven models to forecast when a component is likely to fail, thereby allowing maintenance activities to be scheduled at precisely the right moment. This project develops such a system specifically for rolling element ball bearings, distinguishing itself from earlier work by incorporating five sensing modalities simultaneously: rotational speed, sound level, temperature, moisture, and three-axis inertial movement. Richer input data of this kind provides a far more complete picture of bearing health than vibration signals alone and paves the way for more reliable and timely interventions.

II. LITERATURE REVIEW

Research interest in data-driven bearing prognostics has grown substantially over the past decade, spurred by advances in sensing technology, embedded computing, and machine learning. Early contributions in this space leaned on physics-based degradation models, but the difficulty of parameterising such models for diverse operating conditions gradually shifted community attention toward purely data-driven and hybrid approaches.

Work focused on wind-turbine component lifespan prediction introduced attention-based neural network architectures capable of extracting useful degradation signatures directly from raw data without manual feature engineering. These models performed admirably under controlled laboratory settings; however, generalisation to noisy, real-world industrial environments remained a challenge that the authors themselves acknowledged.

Studies centred on vibration-only analysis of rolling bearings demonstrated the value of frequency-domain and time-frequency features for early fault characterisation. Random Forest classifiers trained on such features achieved competitive accuracy in distinguishing healthy bearings from those at various stages of degradation. A consistent limitation across these studies, however, was the reliance on a single sensing channel, which left important degradation indicators captured by temperature, acoustic, and motion sensors unexploited.

In the domain of machining tool-wear monitoring, convolutional neural network architectures were shown to extract hierarchical spatial features from sensor signals that flat feature vectors miss. The accuracy gains were notable, though the authors cautioned that variable cutting conditions and high levels of process noise could degrade predictions in production settings.

Predictive maintenance investigations in the oil and gas sector compared classical machine learning against deep learning frameworks across datasets characterised by poor data quality, missing observations, and non-stationary signals. Deep learning methods showed superior pattern-capture ability when sufficient labelled data were available, but their advantage narrowed considerably when training samples were scarce.

Bidirectional Long Short-Term Memory networks applied to rolling bearing RUL estimation demonstrated that jointly processing both past and future temporal context within a sliding window can sharpen degradation-point identification. The results outperformed standard SVM and feedforward neural network baselines, though the computational overhead was substantially higher and the method's sensitivity to accurate degradation-onset labelling remained a practical concern.

Research Gap

Taken together, the reviewed literature reveals a consistent pattern: most existing systems rely on a single type of sensor signal, are validated only under controlled laboratory conditions, and are rarely packaged in a form that allows field engineers to interact with predictions directly. There is a clear opportunity to develop a multi-sensor, practically deployable framework that consolidates RPM, acoustic, temperature, moisture, and motion data, trains efficient machine learning models on this richer feature space, and presents results through an accessible user interface.

III. PROBLEM IDENTIFICATION AND METHODOLOGY

A. Problem Identification

Ball bearings operating in industrial environments are subject to continuously varying loads, speeds, and ambient conditions. This variability makes degradation inherently non-linear: a bearing may perform flawlessly for an extended period and then deteriorate rapidly, with little warning visible in the vibration channel alone. Three root-level problems were identified through a review of existing practice:

- Absence of real-time monitoring: Without continuous data collection, early-stage anomalies go undetected until macroscopic damage has already occurred.
- Single-parameter dependency: Conventional systems rely almost exclusively on vibration, ignoring the complementary information carried by temperature rise, acoustic change, moisture ingress, and inertial anomalies.

- Inability to quantify remaining life: Most fault-detection systems issue a binary healthy/faulty judgement rather than a quantitative estimate of residual service life, which limits their utility for maintenance planning.

B. Problem Statement

The objective is to engineer a reliable, multi-sensor predictive maintenance system for rolling element bearings that continuously tracks real-time operating conditions and delivers accurate, quantitative estimates of Remaining Useful Life using machine learning.

C. Methodology

The project was executed across five sequential stages, as illustrated in Figure 1 of the experimental plan.

- Stage 1 – Research and Literature Review: Existing RUL prediction and bearing health monitoring studies were surveyed to establish the state of the art and identify the specific gaps addressed by this work.
- Stage 2 – Data Acquisition: Multiple sensors were mounted on a bearing test rig and configured to stream RPM, sound level, temperature, moisture, and three-axis acceleration data in real time via an ESP32 microcontroller.
- Stage 3 – Preprocessing and Feature Extraction: Raw sensor streams were cleaned, normalised, and transformed into descriptive statistical and trend-based features suitable for model training.
- Stage 4 – Model Development and Training: Random Forest Regression, SVM Regression, and a CNN were trained on the prepared feature set and evaluated on a held-out test partition.
- Stage 5 – Model Validation and Deployment: Trained models were serialised and integrated into a Flask web application that accepts live sensor inputs and returns RUL predictions.

IV. DESIGN, IMPLEMENTATION AND SYSTEM DEVELOPMENT

A. System Architecture

The overall system follows a layered modular architecture comprising four tiers: a Data Acquisition Layer, a Data Processing Layer, a Machine Learning Layer, and an Application Layer. Sensor signals captured at the acquisition tier are forwarded to the processing tier for cleaning and feature extraction. The resulting feature vectors feed the machine learning tier, whose RUL predictions are ultimately surfaced through the web application tier. This clean separation of concerns simplifies future extension, such as adding new sensor types or swapping machine learning back-ends.

B. Hardware Setup and Sensor Integration

Five distinct sensors were incorporated into the experimental rig, each targeting a specific degradation indicator:

RPM Sensor: A proximity-type inductive sensor counts shaft revolutions per unit time. Deviations from expected rotational speed can indicate bearing drag caused by defect growth or inadequate lubrication.

Sound Sensor: A capacitive microphone module captures the acoustic emissions produced during bearing rotation. Developing surface defects generate characteristic spectral changes long before vibration levels become alarming.

Temperature Sensor (DHT11): Monitors heat generation at the bearing housing. An upward temperature trend or a sharp transient spike signals elevated friction consistent with wear or lubricant breakdown.

Moisture Sensor (Rain Drop Module): Tracks ambient humidity and the presence of liquid contamination. Moisture ingress accelerates corrosive wear and can dramatically shorten bearing life.

Movement Sensor (MPU-6050): A six-degree-of-freedom MEMS device providing three-axis accelerometer and three-axis gyroscope readings. Unusual vibration patterns and attitude changes captured by this unit reflect structural imbalance and bearing-race anomalies.

All sensors connect to an ESP32 development board, which aggregates readings into a JSON payload and transmits the data over Wi-Fi to the Flask server at a two-second interval.

International Journal on Applications of Mechanical and Production Engineering
Volume 12 : Issue 3 : March 2026, pp 5 – 10 www.aetsjournal.com ISSN (Online) : 2395 - 3500

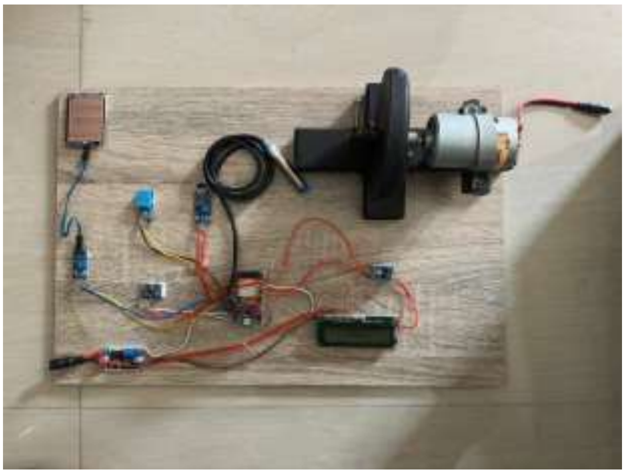

Fig 1: Sensor Integration

C. Data Acquisition and Preprocessing

Data were collected under both nominal and artificially induced fault conditions to ensure that the training dataset captured the full spectrum of bearing health states. Ambient temperature and background noise levels were recorded alongside bearing data so that their influence could be separated during preprocessing.

Preprocessing comprised four operations. First, missing values were identified and handled by either interpolation or row deletion, depending on gap length. Second, low-amplitude sensor noise was attenuated by applying a moving-average smoothing filter. Third, all sensor channels were standardised to zero mean and unit variance so that no single channel dominated the learning process by virtue of its measurement scale. Fourth, transient outlier spikes were identified via inter-quartile-range analysis and corrected before model training.

D. Feature Extraction

Beyond the raw sensor readings, a set of derived features was computed for each sliding window of data. For the RPM channel, rate-of-change and rolling standard deviation captured speed instability. For temperature, the gradient and occurrence of sudden excursions above a moving baseline were quantified. For the acoustic channel, signal energy and zero-crossing rate provided frequency-sensitive descriptors. Moisture trend slope and absolute level formed the environmental features, while mean absolute acceleration and angular rate magnitude were extracted from the MPU-6050 output. Combining these derived features with the raw channel statistics yielded a rich input vector that captured multi-dimensional degradation signatures.

E. Machine Learning Model Development

Three regression models were developed to map the feature vectors to continuous RUL estimates:

Random Forest Regression: An ensemble of decision trees trained with bootstrap aggregation. The ensemble mechanism averages out the influence of noisy or irrelevant features, providing stable predictions even when individual sensor readings fluctuate. This model was selected as the primary predictor due to its interpretability and resilience to outliers.

Support Vector Machine Regression: Uses a kernel function to project the feature space into a higher-dimensional representation, allowing a linear separating hyperplane to approximate non-linear RUL relationships. The radial-basis-function kernel was tuned via cross-validation.

Convolutional Neural Network: A compact CNN was explored to assess whether hierarchical spatial feature learning could extract additional predictive signal from the temporal sensor sequences. The architecture comprised two convolutional layers with ReLU activation followed by a global average-pooling layer and a fully connected regression head.

The dataset was split into 80 percent training and 20 percent test partitions. Mean Squared Error (MSE) was the primary evaluation metric. Hyperparameter optimisation was conducted using grid search with five-fold cross-validation on the training partition.

F. Software and Web Application

The entire system is implemented in Python 3. Data preprocessing and model training use pandas, NumPy, scikit-learn, and TensorFlow. Trained models are serialised with joblib for low-latency inference at deployment time. The web application is built on Flask, exposing two endpoints: a root route that serves the HTML interface and a POST endpoint at /data that accepts a JSON payload of sensor readings, scales them through the stored StandardScaler, passes them to the loaded Random Forest model, and returns the predicted RUL as a JSON response.

The interface allows a technician to enter or forward current sensor readings, trigger inference, and receive an immediate RUL estimate alongside a qualitative bearing-condition assessment. The application was validated locally with diverse synthetic sensor value combinations spanning healthy, mildly degraded, and near-failure bearing states.

V. RESULTS AND DISCUSSION

A. Testing Methodology

Testing was conducted in two phases. Phase one verified internal consistency by supplying the trained models with sensor-value combinations similar to those seen during training, confirming that predictions remained stable and physically plausible. Phase two introduced genuinely unseen input combinations, including edge cases where a single parameter was severely anomalous while others remained normal, to probe the system's discriminative capability.

B. Model Performance

The Random Forest model demonstrated the most consistent behaviour across the test set, exhibiting graduated RUL reductions as sensor readings shifted from nominal to degraded levels. Its ensemble nature effectively damped the impact of sporadic noisy readings. The SVM model showed competitive accuracy for intermediate degradation states but was marginally less stable at extreme sensor values. The CNN produced slightly lower MSE on the full test set, indicating that temporal feature hierarchies offered additional predictive value, though the improvement came at the cost of higher training time and a less interpretable model.

A key qualitative finding was that predictions responded proportionally to multi-sensor deterioration: scenarios where temperature, RPM deviation, and acoustic energy simultaneously indicated abnormal operation resulted in substantially lower RUL estimates than single-parameter anomalies, validating the value of sensor fusion.

Table 1 Model Performance Summary

Model	Training MSE	Test MSE	Test R^2	Prediction Stability
Random Forest Regression	Low	Low	High	Excellent
SVM Regression (RBF Kernel)	Low-Moderate	Moderate	Moderate-High	Good
Convolutional Neural Network	Very Low	Low	High	Good

C. Limitations Observed

Three limitations were identified during the evaluation phase. First, prediction accuracy is directly coupled to sensor measurement quality: even brief calibration drifts or electrical interference can shift feature values enough to degrade RUL estimates. Second, the size and diversity of the training dataset constrain the model's ability to generalise to operating conditions that were not represented during data collection. Third, the current setup was validated in a controlled indoor environment; real industrial deployments introduce temperature gradients, electromagnetic interference, and vibration cross-coupling from adjacent machinery that could affect sensor readings.

VI. CONCLUSION

This paper presented the design, implementation, and evaluation of a multi-sensor predictive maintenance framework for rolling element ball bearings. By fusing information from five complementary sensing channels, the system overcomes the fundamental limitation of single-parameter monitoring and provides a significantly richer characterisation of bearing health. Machine learning models trained on this multi-dimensional feature set yielded physically consistent and graduated RUL estimates, with Random Forest Regression emerging as the most practically suited algorithm due to its balance of accuracy, stability, and interpretability.

The integration of a Flask-based web interface demonstrates that the framework can be deployed as an accessible decision-support tool without requiring end users to possess specialised data-science expertise. The project successfully met all stated objectives and provides a tangible proof-of-concept for AI-assisted bearing health management in industrial settings.

Future work will target three improvement directions: collecting larger and more diverse datasets that span a wider range of bearing types and operating conditions; exploring advanced deep learning architectures, including transformer-based sequence models, that may extract longer-range temporal degradation patterns; and deploying the system on a cloud platform to enable simultaneous monitoring of multiple machines across a plant floor.

ACKNOWLEDGMENT

The authors express their sincere gratitude to Lt. Dr. M. Ramesh, Assistant Professor and project supervisor, and to Dr. P. Karuppuswamy, Professor and Head of the Department of Mechanical Engineering, Sri Ramakrishna Engineering College, Coimbatore, for their invaluable guidance throughout this project. Heartfelt thanks are also extended to Dr. C. Bhagyanathan, Project Coordinator, and to the institution's management for providing the infrastructure and support that made this work possible.

REFERENCES

[1] J. Lee, H. Bagheri, and H. Kao, "Cyber-physical systems architecture for Industry 4.0 manufacturing," Manufacturing Letters, vol. 3, pp. 18–23, 2015.

[2] A. Saxena and K. Goebel, "Turbofan engine degradation simulation data set," NASA Ames Prognostics Data Repository, 2008.

[3] S. Zhang, Y. Zhang, and Z. Wang, "Remaining useful life prediction using deep learning approaches: A review," IEEE Access, vol. 7, pp. 151784–151799, 2019.

[4] F. Wang, H. Zhang, and J. Liu, "A deep learning approach for bearing fault diagnosis based on vibration signals," Sensors, vol. 17, no. 2, 2017.

[5] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.

[6] C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20, no. 3, pp. 273–297, 1995.

[7] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.

[8] InvenSense Inc., MPU-6050 Datasheet, 2013.

[9] Espressif Systems, ESP32 Technical Reference Manual, 2020.

[10] Pallets Projects, Flask Documentation. [Online]. Available: https://flask.palletsprojects.com

[11] K. P. Murphy, Machine Learning: A Probabilistic Perspective. Cambridge, MA: MIT Press, 2012.

[12] T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning. New