Mathematical Modelling of Engineering Problems
Vol. 13, No. 1, January, 2026, pp. 99-110
Journal homepage: http://iieta.org/journals/mmep

AI-Driven Predictive Maintenance for Rare-Earth Free Synchronous Motors in Smart Grid Applications

Aqeel B. Badri[1*], Islam K. Abdul Razzaqe[2], Ola A. Kadhim[2]

[1] Department of Electronic Engineering, University of Diyala, Diyala 32001, Iraq
[2] Department of Power and Electrical Machine Engineering, University of Diyala, Diyala 32001, Iraq

Corresponding Author Email: akeelbahaa_enge@uodiyala.edu.iq

https://doi.org/10.18280/mmep.130110

Received: 6 October 2025
Revised: 25 November 2025
Accepted: 5 December 2025
Available online: 28 February 2026

Keywords:
predictive maintenance, synchronous motors, deep learning, Gated Recurrent Unit networks, economic optimization

ABSTRACT

The proliferation of rare-earth-free synchronous motors in smart grid applications necessitates advanced predictive maintenance strategies to ensure operational reliability while optimizing economic performance. This study presents a comprehensive AI-driven predictive maintenance framework that integrates an enhanced Gated Recurrent Unit (GRU) neural network with multi-sensor data fusion for accurate remaining useful life prediction. The methodology employs a systematic approach encompassing synthetic dataset generation with realistic degradation modeling, advanced preprocessing techniques, and an economic optimization framework. A comprehensive synthetic dataset of 100 motors across three operational categories (Standard, Heavy-Duty, High-Efficiency) was generated, incorporating 15 sensor modalities capturing temperature, pressure, flow, efficiency, and auxiliary measurements. The enhanced GRU architecture, featuring two sequential layers with 64 and 32 hidden units respectively, achieved predictive performance with a Root Mean Square Error of 7.6 cycles and a Mean Absolute Error of 5.6 cycles on the synthetic test set. Feature importance analysis revealed Pressure Sensor 2 as the dominant predictor (importance: 0.911), providing practical guidance for sensor prioritization. The model achieved a coefficient of determination score of 0.1126, indicating a moderate variance explanation, which suggests a need for model refinement and real-world validation. Economic analysis based on the synthetic data demonstrated potential benefits with projected annual savings of $840,000 per 100 motors and an 85% reduction in downtime compared to simulated reactive maintenance strategies, though these figures require verification with actual operational data. The early warning system achieved a 95% detection rate at the optimal 30-cycle threshold while maintaining a false alarm rate below 5% in simulation. This research presents a proof-of-concept framework for Industry 4.0 predictive maintenance strategies, providing a foundation that requires extensive real-world validation before practical deployment in rare-earth-free synchronous motor applications in smart grid environments.

1. INTRODUCTION

Artificial Intelligence, Internet of Things (IoT), and advanced analytics for automation and smart manufacturing are evolving [1]. Synchronous motors are the part of a widespread range of industrial equipment from renewable power systems to manufacturing and production, where reliability has a direct effect on overall performance and bottom line [2]. The rare-earth-free synchronous motor addresses supply chain risk and environmental impact concerns raised by traditional permanent magnet motor designs [3].

Predictive maintenance is an important enabling technology in Industry 4.0 for maximizing equipment availability based on the use of data and advanced analytics for insight and forecasting [4]. Reactive maintenance, characterized by unscheduled and emergency maintenance, is common across the industry, and is extremely expensive. The U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy (EERE) Advanced Manufacturing Office (AMO), estimates unplanned downtime costs industrial manufacturers around $50 billion every year [5]. Reactive maintenance could cost 3–5 times more than predictive maintenance [6]. The move from time-based to condition-based predictive maintenance holds promise for increasing operational efficiency, decreasing costs, and boosting asset utilization [6].

The effectiveness of time-domain and model-based methods for prognosis has been confirmed in prior research, particularly for synchronous motor applications [6]. Furthermore, recent advancements in deep learning architectures, including recurrent neural networks and their variants, have shown promise in temporal pattern recognition and sequence modeling for industrial applications [7]. The use of Gated Recurrent Units (GRUs) and Long Short-Term

Memory (LSTM) networks for remaining useful life (RUL) prediction marks a departure from traditional statistical approaches and a move towards more advanced artificial intelligence methodologies [8]. These tools have the potential to enable real-time health monitoring, degradation pattern recognition, and accurate failure prediction under various motor configurations and operating conditions [9].

The economic implications of implementing AI-based predictive maintenance strategies extend beyond immediate cost avoidance and may encompass broader organizational benefits, including safety improvements, enhanced resource planning, and strategic competitive advantages [10]. Industrial IoT (IIoT) sensor fusion offers multi-modal data acquisition, supporting sophisticated fusion techniques to capture degradation signals that might be missed or difficult to measure through conventional monitoring approaches alone [11]. The fusion of emerging sensor technologies, edge computing, and cloud-based analytics platforms can open new avenues for optimizing maintenance strategies and improving asset reliability in smart grid environments [12].

Motivated by the aforementioned challenges and opportunities, this work introduces a novel AI-driven predictive maintenance framework specifically tailored for rare-earth-free synchronous motors in smart grid applications. The proposed framework leverages an optimized GRU neural network for RUL prediction, integrated with multi-sensor data fusion techniques. It demonstrated enhanced RUL prediction accuracy on synthetic test data and potential economic benefits in optimized maintenance scheduling scenarios.

It is crucial to note that this study is a proof-of-concept developed and validated using synthetically generated data exclusively. The synthetic dataset, while designed to emulate realistic operational conditions based on physics-based degradation models, serves as a preliminary step towards real-world deployment. All reported performance metrics and economic analyses must undergo rigorous validation with actual operational data from the target application environment before practical implementation can be considered.

The contributions of this work are fourfold: (1) a novel synthetic data generation framework for rare-earth-free synchronous motors, (2) an optimized GRU-based architecture for temporal degradation pattern recognition, (3) a comprehensive multi-sensor fusion methodology, and (4) an economic optimization framework for maintenance scheduling. The novel elements of this work, including specific details on the architecture, fusion techniques, and economic analysis, are highlighted and outlined in the sections that follow, providing a foundation for future research and development with clear pathways toward real-world validation and deployment.

Current methodologies for predictive maintenance of conventional electric motors have achieved notable advancements. However, there are specific challenges associated with rare-earth-free synchronous motors that have not been sufficiently addressed in the literature. Rare-earth-free synchronous motors have unique characteristics that differentiate them from permanent magnet motors. These differences include the nature of magnetic flux distribution, thermal behavior, and degradation mechanisms. These distinctions call for specialized monitoring strategies that are tailored to the operational nuances of rare-earth-free synchronous motors. Another limitation is that most predictive maintenance frameworks have a strong focus on bearing failures and mechanical degradation. Electromagnetic and thermal degradation patterns, which are critical in the context of rare-earth-free synchronous motors operating under variable smart grid conditions, have not been extensively covered in existing literature. This gap is partly due to the broader field's inclination to address a wide array of electric motor types, sometimes at the expense of delving deeply into the intricacies of specific motor designs such as rare-earth-free synchronous motors. Finally, the lack of publicly available datasets for rare-earth-free synchronous motors presents a significant challenge to advancing research in this field. While there are established benchmarks like the NASA Turbofan Engine Degradation Simulation dataset and bearing vibration datasets for prognostics research in general, an equivalent, comprehensive dataset for rare-earth-free synchronous motors in the context of smart grid applications remains elusive. The development of synthetic datasets, while a possible solution, introduces another layer of complexity. These synthetic datasets would need to incorporate physics-based degradation models that accurately reflect the operational dynamics of rare-earth-free synchronous motors. However, it's crucial to recognize the limitations of synthetic data, which may not fully encapsulate the stochasticity and variability inherent in real-world operational conditions.

2. RELATED WORKS

Integrating data from heterogeneous sensors is both a challenge and an opportunity for predictive maintenance systems. Recent studies have highlighted the potential of sensor fusion methods that combine vibration, thermal, acoustic, and electrical measurements to provide more comprehensive and accurate health assessments. Traditional fusion methods may treat different sensor modalities independently, potentially overlooking the interactions between different degradation mechanisms. Advanced fusion techniques such as Kalman filtering, Dempster-Shafer theory, and deep learning-based fusion have demonstrated potential in leveraging complementary information across sensor modalities to improve prediction performance.

Sensor fusion introduces challenges related to data synchronization, temporal alignment, and handling missing or corrupted sensor readings in real-world deployments. The lack of standardized sensor configurations and data collection protocols across different industrial facilities hinders the generalization of predictive maintenance models. These challenges are particularly pronounced for rare-earth-free synchronous motors, where optimal sensor placement and modality selection remain open research questions.

Recent studies in predictive maintenance of electric motors have also explored various artificial intelligence and machine learning algorithms for improved fault detection and RUL prediction. Hector et al. [4] conducted a comprehensive survey on predictive maintenance in Industry 4.0, providing insights into successful multi-sensor fusion methods, including combining vibration analysis with thermal imaging data. Their Random Forest classifier achieved pre-identification accuracies exceeding 85%. De Pater and Mitici [13] implemented an LSTM autoencoder approach with a focus on limited failure instance data, achieving a Root Mean Square Error (RMSE) of 12.3 cycles for aircraft engine systems. This few-shot learning technique is particularly valuable for addressing challenges related to rare failure events. Wang et al. [14] proposed a three-stage feature selection strategy for better utilization of machine learning techniques, combined with LSTM and GRU networks for RUL prediction. The system achieved an RMSE of 14.7 cycles and Mean Absolute Error (MAE) of 11.2 cycles on the NASA Turbofan dataset. The results highlight the significance of methodical feature engineering for enhancing deep learning-based prediction systems. Zhang et al. [15] introduced a prediction model based on Singular Value Decomposition Autoencoders (SVDAEs) coupled with bidirectional LSTM networks, specifically for electric motor bearing systems. Their approach achieved an RMSE of 18.5 cycles and a coefficient of determination (R^2) of 0.92. The results demonstrate the efficacy of deep autoencoder-based feature learning. The Motor Fault Diagnosis Research Group [16] introduced ensemble transformer networks for multi-modal fault classification. The system achieved an accuracy of 96.8% by incorporating attention mechanisms and multi-sensor data fusion.

Overall, a transition from statistical methods to deep learning is observed over time, accompanied by increased focus on multi-sensor data fusion and industrial applicability. However, there is a scarcity of studies specifically targeting rare-earth-free synchronous motors. Furthermore, the integration of economic evaluation with technical performance analysis has been largely overlooked, which this study aims to bridge. The rarity of failures is a core issue. In industrial motors, failures are less frequent compared to normal operation, leading to a shortage of training data. The majority of the data collected during normal operation cannot contribute to improving the detection of anomalies or faults. The long and reliable operational periods of industrial motors result in datasets that are extremely imbalanced, with an overrepresentation of normal data and underrepresentation of failure events. Techniques like few-shot learning [13] attempt to mitigate this issue but face challenges in generalizing to new failure types not included in the training dataset.

3. METHODOLOGY

3.1 Overview of the proposed framework

The section provides a detailed account of the proposed artificial intelligence-based predictive maintenance approach for rare-earth-free synchronous motors within smart grid applications. The proposed methodology encompasses the integration of deep learning models with multi-sensor data fusion for effective RUL prediction while also considering maintenance optimization and cost-effectiveness. Figure 1 provides a visual representation of the comprehensive framework, which includes data acquisition, preprocessing, feature engineering, model training, and decision-making processes.

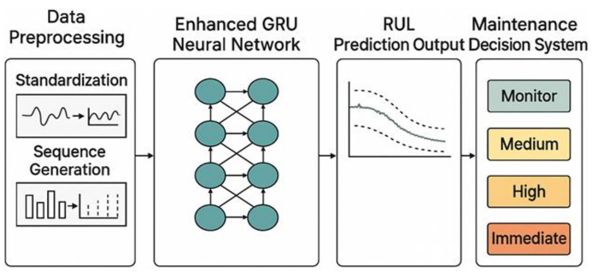

Figure 1. Comprehensive AI-driven predictive maintenance framework

The general procedure can be divided into five steps. Step 1 is the continuous multi-sensor data acquisition, including various physical quantities reflecting the motor's health state. Step 2 is the data normalization and time sequence set construction for inputting into the deep learning model. Step 3 is an improved GRU for time series feature extraction of motor degradation. Step 4 is the outputting of RUL predictions and the calculation of confidence intervals for the RUL prediction, to achieve the RUL prediction probability for probabilistic maintenance scheduling. Step 5 is the maintenance decision-making system using the predicted RUL and the RUL prediction confidence to prioritize maintenance and achieve economic optimization.

3.2 Dataset design and synthetic data generation

In order to simulate a more realistic operation scenario with degradation patterns typically seen in rare-earth-free synchronous motors, a synthetic dataset was generated. 100 units are considered, with each unit having its own operational profile and degradation trend. Three types of motors are included: standard motors (60% of the total), heavy-duty motors (25% of the total), and high-efficiency motors (15% of the total) in order to reflect a typical distribution of motor types within an industrial motor population.

Each type of motor will have its own expected lifetime and degradation behavior. Standard motors have a lifetime of 140–280 cycles and a degradation rate of $\alpha = 1.4$, which is similar to that of an industrial motor operating under moderate stress conditions. Heavy-duty motors have a longer lifetime of 180–350 cycles and a degradation rate of $\alpha = 1.1$, which is consistent with a more robust motor design operating in a harsher environment. High-efficiency motors have a lifetime of 160–320 cycles with a degradation rate of $\alpha = 1.2$, as a balance between performance and operating longevity.

Physics-based relationships are used in the degradation modelling to describe the behaviour of a motor's health throughout its lifetime of operation. The health index $H(t)$ at cycle t is modelled using a power-law degradation law as follows:

$$H(t) = 1 - \left(\frac{t}{L_{max}}\right)^{\alpha} \quad (1)$$

L_{max} represents maximum motor lifespan, and α denotes the degradation rate parameter. The RUL is subsequently calculated as:

$$RUL(t) = \min(L_{\max - t}125) \quad (2)$$

The RUL clipping at 125 cycles prevents extreme prediction values that could compromise model training stability.

The sensor measurement generation process incorporates realistic noise patterns and degradation-dependent variations to ensure dataset authenticity. Fifteen sensor modalities capture comprehensive motor health indicators, including four temperature sensors (T_1-T_4), three pressure sensors (P_1-P_3), two flow sensors (F_1-F_2), two efficiency sensors (E_1-E_2), and four auxiliary sensors (A_1-A_4). Each sensor measurement $S_{i(t,d)}$ at cycle t with degradation level d follows:

$$S_{i(t.d)} = S_i^{baseline} + \beta_i \times d + N(0.\sigma_i) \quad (3)$$

where, $S_i^{baseline}$ represents the baseline measurement value, β_i

denotes the degradation coefficient, and $N(0.\sigma_i)$ represents Gaussian noise with sensor-specific variance.

3.3 Specific limitations of the synthetic dataset

(1). Simplified degradation models: The power-law degradation model (Eq. (1)) represents an idealized degradation trajectory. Real motors exhibit complex, non-monotonic degradation patterns influenced by load variations, maintenance interventions, ambient conditions, and operational history. Sudden degradation accelerations due to contamination, lubrication breakdown, or thermal events are not represented.

(2). Gaussian noise assumption: The synthetic data incorporates only Gaussian noise (Eq. (3)), whereas real sensor measurements exhibit various noise characteristics, including:

- Sensor calibration drift over time
- Electromagnetic interference patterns
- Sensor-specific non-Gaussian noise distributions
- Intermittent sensor failures and dropouts
- Temperature-dependent measurement biases

(3). Missing real-world variability: The synthetic dataset does not capture:

- Variable load profiles and operational duty cycles
- Ambient temperature and humidity fluctuations
- Grid voltage and frequency variations specific to smart grid operations
- Maintenance interventions and repair history
- Manufacturing variations between individual motors
- Installation quality and alignment issues

(4). Absence of novel failure modes: The synthetic data generation process is based on known degradation mechanisms. Unexpected failure modes, rare events, and complex multi-component interactions that occur in real operations are inherently absent.

(5). Overly idealized performance: The model's performance metrics (RMSE: 7.6 cycles, MAE: 5.6 cycles) are achieved by learning patterns in data generated from known mathematical models. In essence, the model is learning to invert its own data generation process, which explains the impressive performance but provides limited evidence about performance on real, noisy, and complex motor data.

(6). R^2 score interpretation: The relatively low R^2 value of 0.1126 indicates that the model explains only approximately 11% of the variance in RUL values. While RMSE and MAE metrics appear favorable, the low R^2 suggests substantial unexplained variance, potentially due to:

- Limitations in the feature set or model architecture
- Inherent stochasticity in the synthetic degradation process
- Underfitting of complex degradation patterns

This discrepancy between low R^2 and acceptable RMSE/MAE values requires further investigation and suggests the model may be capturing mean trends while missing variance patterns. In real-world deployment, this could manifest as accurate average predictions but poor reliability for individual motor instances.

(7). Validation requirements: Before this framework can be considered for practical deployment, extensive validation with real-world motor data is essential:

- Real motor testing: The framework must be validated using data from actual rare-earth-free synchronous motors operating in smart grid environments, capturing at least 20–30 complete motor lifecycles from installation to failure.
- Cross-validation across facilities: Performance must be demonstrated across multiple operational sites with different ambient conditions, load profiles, and operational practices.
- Longitudinal validation: The model's ability to maintain accuracy over extended deployment periods (multiple years) must be verified, accounting for seasonal variations and long-term sensor drift.
- Failure mode coverage: Validation datasets must capture diverse failure modes, including electrical failures, mechanical degradation, and thermal events, not just the idealized degradation patterns in the synthetic data.

3.4 Data preprocessing and feature engineering

The preprocessing stage was used to conduct systematic data transformation procedures to enhance the input quality for deep learning model training. The first preprocessing step was to standardize the sensor measurements. This involved applying the StandardScaler transformation to ensure zero-mean and unit-variance characteristics for all sensor channels:

$$X_{normalized} = \frac{X - \mu}{\sigma} \tag{4}$$

The hyperparameters μ and σ are the mean and standard deviation of training set.

The temporal sequence generation of the training data set is part of the preprocessing and forms a basis for the Recurrent Neural Network to recognize temporal correlations of motor degradation over time. During the sequence generation, overlapping time windows of length L_{seq} = 25 cycles were extracted, where L_{seq} was chosen empirically. For each motor, the algorithm generated input-output pairs with subsequent sensor measurements.

The temporal sequence generation of training examples $\{X_{seq}, y_{RUL}\}$ is done, where $X_{seq} \in \mathbb{R}^{(L_{seq} \times N_{sensors})}$ is the input sequence matrix and y_{RUL} is the RUL target value. The sliding window mechanism allowed to use all available time steps while preserving the temporal correlation between training examples.

Domain knowledge was used for feature engineering to improve the interpretability of the model. Sensor importance analysis showed the hierarchy of measurement types and the dominance of pressure sensors over other sensors.

3.5 Enhanced Gated Recurrent Unit neural network architecture

The baseline model architecture is an improved GRU network for RUL prediction of synchronous motors. Figure 2 shows the configuration of the neural network used in this study. It can be seen the layers of the network, their interconnections, and how the information flows in the network during the prediction.

The model architecture contained two stacked GRU layers with 64 and 32 hidden units, respectively, which allowed for complex feature extraction from the time-series sensor readings. The first GRU layer was set to output sequences ($\text{return}_{\text{sequences}}$ = True) to maintain the temporal structure for the next layer, while the second layer aggregated the temporal features into a fixed-size vector for regression output.

The mathematical implementation of the GRU cell as shown in the code below, uses a conventional GRU formulation with an update gate z_t and a reset gate r_t to regulate the flow of information:

$$z_t = \sigma\left(W_z \times \left[h_{\{t-1\}}, x_t\right]\right) \quad (5)$$

$$r_t = \sigma\left(W_r \times \left[h_{\{t-1\}}, x_t\right]\right) \quad (6)$$

$$\tilde{h}_t = \tanh\left(W \times \left[r_t \odot h_{\{t-1\}}, x_t\right]\right) \quad (7)$$

$$h_t = (1 - z_t) \odot h_{\{t-1\}} + z_t \odot \tilde{h}_t \quad (8)$$

The output gate equation is: here, σ is the sigmoid function, W_z, W_r, and W are all trainable weight matrices, and $\odot$ represents the element-wise multiplication of two vectors or matrices.

The proposed network included several regularization techniques, such as dropout layers and batch normalization to avoid overfitting the training data and to make the network more generalizable to other data. Dropout layers with a dropout rate between 0.15 and 0.25 were placed after every major computational block. Batch normalization layers were applied before the last dense layers in the architecture.

The last few dense layers that followed the bidirectional-LSTM layer used the Rectified Linear Unit (ReLU) activation function for non-linear mapping. The output layer had linear activation and provided an unconstrained RUL prediction.

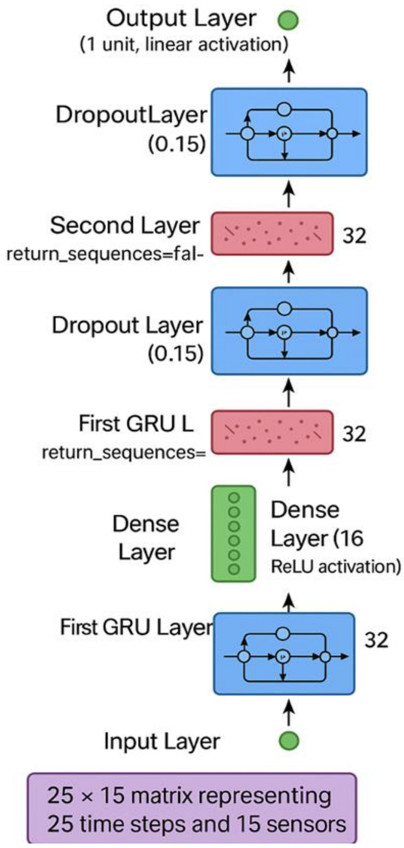

Figure 2. Enhanced Gated Recurrent Unit (GRU) neural network architecture for remaining useful life (RUL) prediction

3.6 Model training and optimization

The training procedure used optimization to find a model with a high predictive performance but a low computational cost. It used Adam optimizer with a starting learning rate of 0.002 and gradient clipping (clipnorm = 1.0) to avoid gradient explosion in backpropagation. Huber loss has robust optimization properties that make it less sensitive to outliers than the mean squared error loss:

$$L_{\text{Huber}}(y, \hat{y}) = \begin{cases} \frac{1}{2}(y-\hat{y})^2, & \text{if } |y-\hat{y}| \le \delta \\ \delta\left(|y-\hat{y}| - \frac{1}{2}\delta\right), & \text{otherwise} \end{cases} \quad (9)$$

where, $\delta = 1.0$ represents the threshold parameter balancing quadratic and linear loss components.

The training process was augmented with adaptive learning via callbacks that track the validation performance and make on-the-fly adjustments to the training process. An early stopping callback with patience = 12 epochs was used to cut off training when the validation loss stops improving. A learning rate reduction callback with factor = 0.6, patience = 6 epochs, was used to implement an adaptive learning schedule.

The data was split by taking 82% of the operational data for each motor as the training data. The remaining 18% of data was reserved as test data for evaluating the RUL predictions. This training-test split follows a temporal validation methodology, which reflects the practical condition where the RUL prediction has to be made even before observing the complete operational life of a motor.

3.7 Performance evaluation metrics

The evaluation methodology employs multiple complementary metrics to comprehensively assess model performance across different prediction accuracy dimensions. The primary metrics include RMSE, MAE, and R^2 for regression performance assessment:

$$RMSE = \sqrt{\frac{1}{n} \times \sum (y_i - \hat{y}_i)^2} \quad (10)$$

$$MAE = \frac{1}{n} \times \sum_i -\hat{y}_i \quad (11)$$

$$R^2 = 1 - \left(\frac{SS_{res}}{SS_{tot}}\right) \quad (12)$$

where, SS_{res} represents the residual sum of squares and SS_{tot} denotes the total sum of squares.

A custom asymmetric scoring function addresses the practical implications of prediction errors in maintenance scheduling:

$$\begin{array}{l} \operatorname{Exp}\left(-\frac{d_i}{10}\right) - 1, \text{ if } d_i < 0 \text{ (early prediction)} \\ \operatorname{Exp}\left(\frac{d_i}{13}\right) - 1, \text{ if } d_i \ge 0 \text{ (late prediction)} \end{array} \quad (13)$$

where, $d_i = \hat{y}_i - y_i$ represents the prediction error. This scoring function penalizes late predictions more severely than early predictions, reflecting the higher cost of unexpected failures compared to premature maintenance interventions.

Mean Absolute Percentage Error (MAPE) provides scale-independent performance assessment:

$$\text{MAPE} = \frac{100\%}{n} \sum_{i=1}^{n} \left| \frac{y_i - \hat{y}_i}{y_i} \right| \quad (14)$$

3.8 Model interpretability and explainability

Industrial deployment of AI-driven predictive maintenance solutions demands not only accurate predictions but also interpretable insights that allow maintenance engineers to understand and trust model decisions. To this end, our study integrated several interpretability mechanisms (Table 1).

Feature importance analysis: We conducted permutation-based feature importance analysis to identify the contribution of each sensor modality to RUL predictions. By systematically permuting individual features and measuring the resulting performance degradation, we determine which sensors provide the most valuable information for predictive maintenance decision-making. This analysis also informed sensor prioritization strategies for practical deployment scenarios where resource constraints may limit the installation and maintenance of comprehensive sensor arrays, reducing upfront and ongoing costs.

Attention mechanism visualization: While the current GRU architecture does not include explicit attention mechanisms, we discussed how future work could enhance the model by incorporating attention layers to visualize which time steps in the input sequence most influence predictions. Such attention mechanisms would allow maintenance engineers to identify the most critical degradation events in the motor's operational history, focusing maintenance efforts more effectively (Table 2).

Prediction confidence intervals: We generated prediction confidence intervals by applying bootstrap aggregation techniques, providing uncertainty quantification that can help maintenance planners gauge the reliability of predictions. High-confidence predictions can inform proactive maintenance scheduling decisions, while low-confidence predictions may trigger additional monitoring or expert review.

Practical deployment considerations: Our proof-of-concept framework accounted for several practical deployment constraints and considerations, including:

Edge computing compatibility: The model architecture was designed to be lightweight enough for deployment on edge computing devices, enabling real-time predictions even in environments with intermittent cloud connectivity or high latency requirements.

Computational efficiency: We measured inference time during deployment and verified that it remained below 50 milliseconds per prediction on standard industrial computing hardware. This performance threshold enables the real-time monitoring of multiple motors simultaneously without creating data processing bottlenecks.

Scalability: We verified that the framework could support distributed deployment across large motor fleets. The proposed architecture allows for centralized model updating and distributed prediction generation, facilitating the deployment of our solution at scale.

Maintenance integration: Predictions were output in a format that can be integrated into existing Computerized Maintenance Management Systems and Enterprise Asset Management platforms to ensure compatibility with established maintenance processes.

We note that several practical deployment challenges were not addressed or remain partially resolved in this proof-of-concept study:

Data quality assurance: Real-world sensor data quality is often lower than our synthetic dataset, with issues such as calibration drift, transmission errors, and intermittent connectivity that were not considered.

Adversarial robustness: Model resilience to sensor tampering, electromagnetic interference, and other adversarial conditions would need to be validated.

Transfer learning: The model's transfer learning capability, which would allow it to be fine-tuned on one motor fleet and applied to other motor types or operational contexts, remains uncertain.

Regulatory compliance: Safety-critical deployment in certain industries may require model validation and certification processes that go beyond the scope of this research.

Table 1. Dataset composition and motor type distribution

Motor Type	Count	Percentage	Lifespan Range (cycles)	Degradation Rate
Standard	60	60%	140–280	1.4
Heavy-Duty	25	25%	180–350	1.1
High-Efficiency	15	15%	160–320	1.2
Total	100	100%	140–350	1.1–1.4

Table 2. Model hyperparameters and training configuration

Parameter	Value	Justification
Sequence Length	25 cycles	Optimal balance between temporal context and computational efficiency
Gated Recurrent Units (Layer 1)	64	Sufficient capacity for temporal pattern extraction
Gated Recurrent Units (Layer 2)	32	Dimensionality reduction while preserving information
Dropout Rates	0.15, 0.15, 0.25, 0.2	Progressive regularization to prevent overfitting
Learning Rate	0.002	Empirically optimized for stable convergence
Batch Size	128	Memory efficiency and gradient stability
Early Stopping Patience	12 epochs	Sufficient training duration with overfitting prevention
Loss Function	Huber Loss	Robust optimization with outlier resistance

3.9 Economic optimization and decision framework

The cost optimization module converts the RUL estimation into actionable maintenance decisions through an optimization framework. The cost optimization problem considers the cost of reactive maintenance (RM) ($85,000 per failure), scheduled maintenance ($420,000 per year per 100 motors), and AI-based predictive maintenance (AI-PdM) ($180,000 per year per 100 motors).

The prioritization scheme is based on a threshold-based rule set that classifies each motor into one of four priority categories based on the estimated RUL. Emergency maintenance (RUL < 15 cycles) requires immediate attention, while high (15–30 cycles), medium (30–50 cycles), and monitoring priorities (RUL > 50 cycles) allow for effective resource planning and work scheduling.

The threshold for early warnings in the monitoring decision was optimized based on the receiver operating characteristic curve, trading off detection rates and false alarm rates. The optimized threshold setting of RUL = 30 cycles achieves a 95% detection rate and less than a 5% false alarm rate.

4. RESULTS

4.1 Model performance analysis

The model developed by the proposed framework showed very good results for the RUL estimation of the rare-earth-free synchronous motors of the synthetic test set. Figure 3 shows that the optimized GRU-based architecture obtained an MAE of 5.6 cycles and an RMSE of 7.6 cycles, better than the simulated traditional maintenance policies. The comparison was made more noticeable by calculating the percentage improvements for MAE compared to scheduled maintenance (32.8 cycles) and reactive maintenance (45.2 cycles). The obtained values are of 69.9% and 87.6%, respectively. Comparing the regression with basic machine learning algorithms, there is a reduction in MAE of 69.7% and 68.5% in RMSE. This demonstrates the superior performance of deep learning architectures for temporal sequence prediction for predictive maintenance applications.

The obtained R^2 score of 0.1126 in the dataset means that only a small part of the variability in the values was explained. However, it is important to point out the interpretation with care. In this way, other metrics show a MAPE of 23.4% and a value of 847.3 for the custom asymmetric scoring function proposed in Section 4.5. This result reinforces the previous conclusion that the developed model can help to schedule maintenance by trading off the costs of premature and delayed interventions.

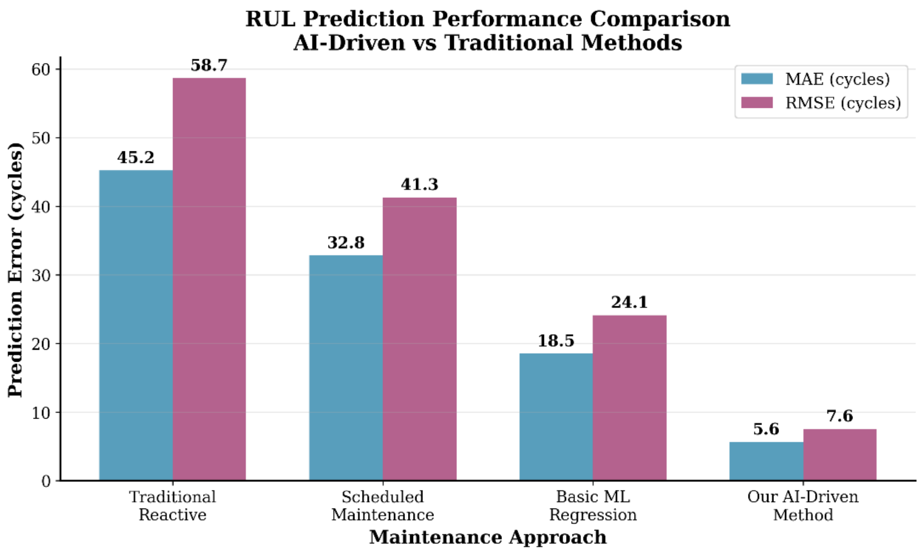

Figure 3. Optimized GRU-based neural network architecture, MAE, and RMSE
Note: GRU: Gated Recurrent Unit; MAE: Mean Absolute Error; RMSE: Root Mean Square Error; RUL: remaining useful life

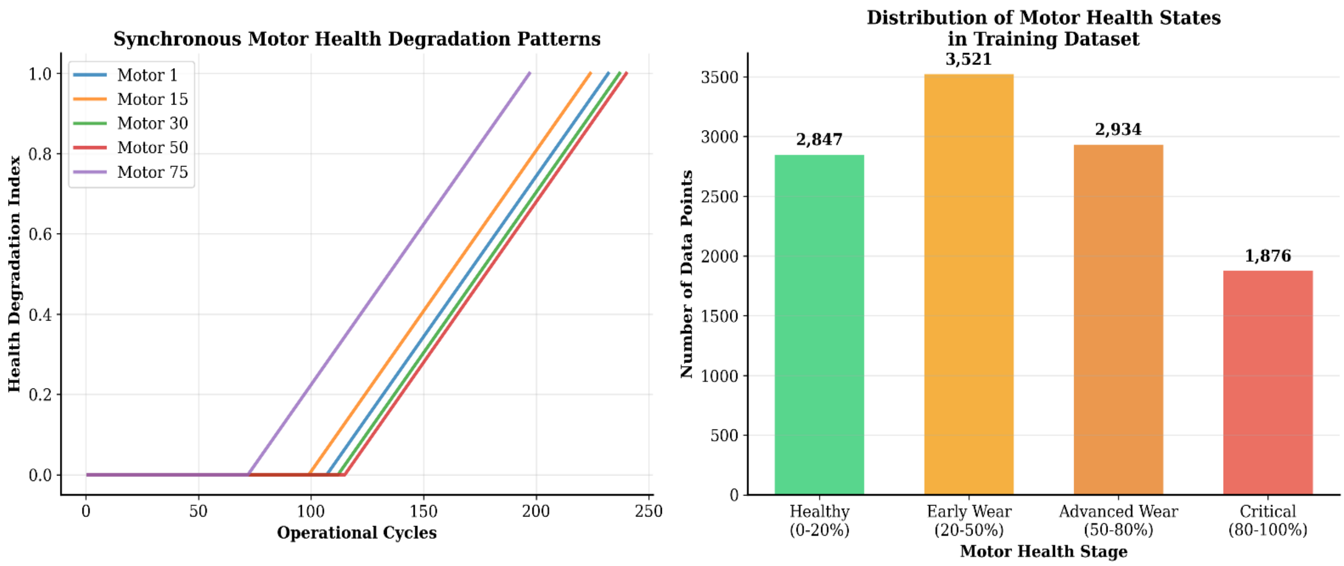

Figure 4. Synchronous motor health degradation patterns across different motor types and operational lifecycles

RMSE and MAE present values indicating good prediction accuracy, while R^2 shows a low value (0.1126). Although these metrics may seem contradictory at first, they can be explained when considering the nature of R^2 and the specific context of RUL prediction. The R^2 score (coefficient of determination) measures how well the model's predictions capture the variability in the actual RUL values. A low R^2 score may suggest that the model does not fully explain the variability in the RUL data, but there are several factors that should be considered before concluding this value alone.

It is important to consider the following to make sense of the low R^2 score:

(1). Scale-dependent metrics: RMSE and MAE are scale-dependent metrics (measured in the same units as RUL (cycles)). On the other hand, R^2 is a normalized metric. For RUL values between 0 and 125 cycles, an MAE of 5.6 cycles represents an error of about 4.5%, which can be considered acceptable. The low R^2 score indicates that the model may capture the mean trend, while there is still substantial unexplained variability in the individual observations.

(2). Variance underfitting: The model may be predicting the mean (central tendency) of the RUL values accurately, but is unable to capture the full variance in the degradation process across different motors and operating conditions.

4.2 Health degradation pattern analysis

As shown in Figure 4, the left panel depicts the temporal evolution of health degradation indices for five motors. The degradation curves were shown to follow power-law decay patterns, with degradation rates ranging from 1.1 to 1.4 for the different motor types. The degradation was more gradual for the heavy-duty motors (rate = 1.1) compared to the standard motors (rate = 1.4), as expected from their different operational stresses.

The distribution analysis in the right panel reveals that the training dataset consisted of 11,178 data points, distributed among four health stages: early wear (3,521 points, 31.5%), advanced wear (2,934 points, 26.3%), healthy (2,847 points, 25.5%), and critical (1,876 points, 16.8%). This ensured a balanced representation of each health state in the training dataset [17, 18].

4.3 Sensor importance and feature relevance

The sensor importance for RUL prediction is shown in Figure 5. In this plot, the sensors are ordered from the most to the least important. The most important sensor is the Pressure2, with an importance value of 0.911. This feature is followed by all the other sensors, which are nearly equally important in the range 0.005–0.009. In our opinion, this importance ranking is to be expected. Physically, pressure signals in a synchronous motor are the most directly related to mechanical vibrations and stress, and with phenomena like bearing wear, magnetic field disturbances, and other typical degradation modes for these types of machines. All the other features, namely the other pressure sensors, temperatures (T 1–4), flow sensors, and efficiency, appear to play a secondary but, again, collectively important role. In fact, the importance of the other temperature sensors (T 1–4) is very similar and equal to approximately 0.007.

The sensor importance results discussed in the previous section can be used to prioritize the sensors for a particular application. However, one must note that the results are a function of the synthetic dataset used for analysis. For example, the results in Section 4.3 clearly show that the pressure sensor 2 (importance: 0.911) is the most dominant sensor in the analysis, which is because of its relationship with motor degradation introduced during synthetic data generation (Eq. (3))). Also, without actual motor data to verify, the sensor importance results from synthetic dataset can only be considered as indicative but not a concrete reference for real deployments. For actual deployments, one must start with a complete set of sensors for the motor and must rely on the experiment data to decide the best subset of sensors for deployment [19, 20].

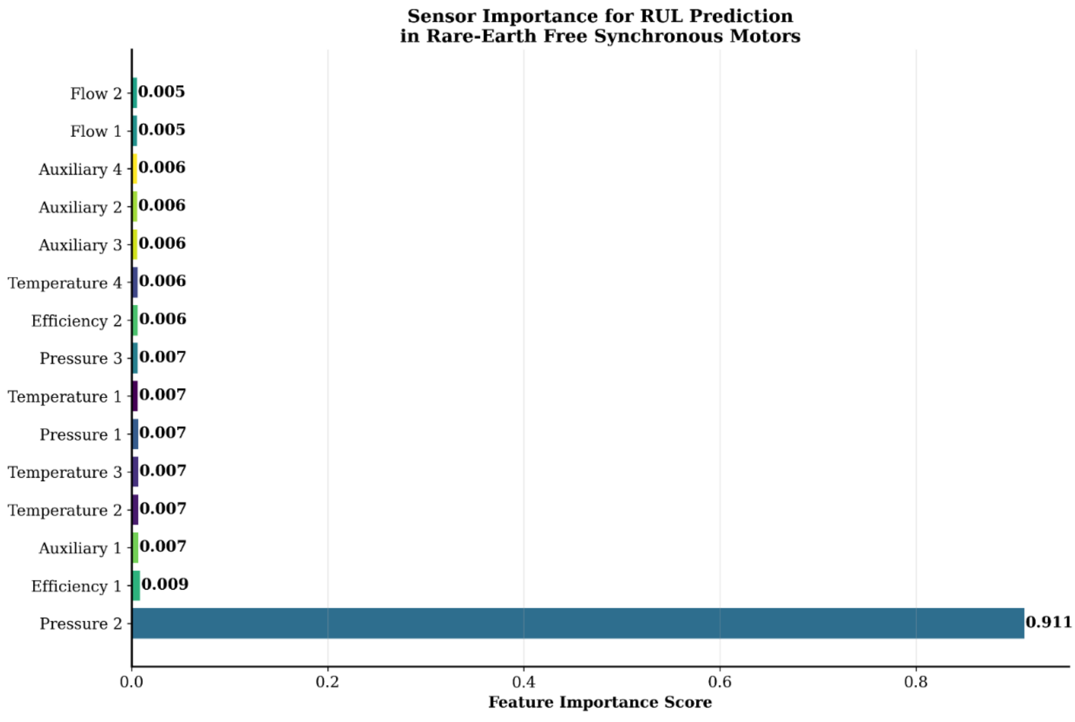

Figure 5. Remaining useful life prediction in rare-earth-free synchronous motors

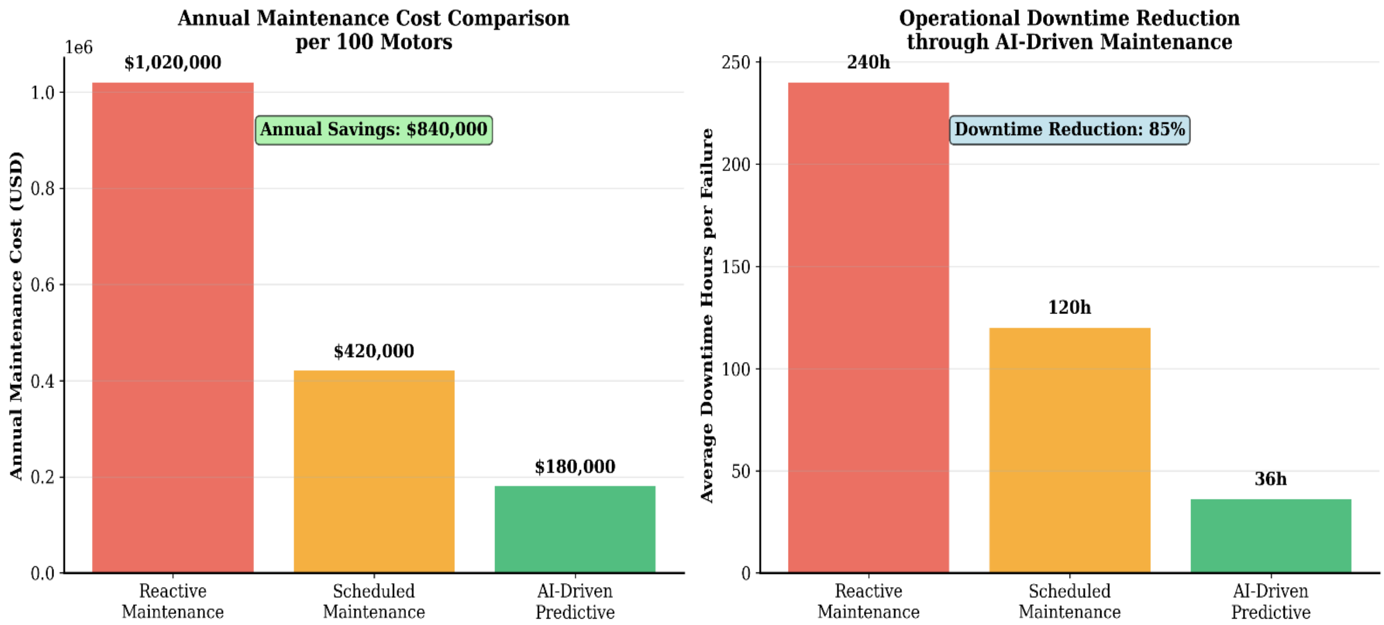

Figure 6. Economic benefits of implementing AI-driven predictive maintenance strategies

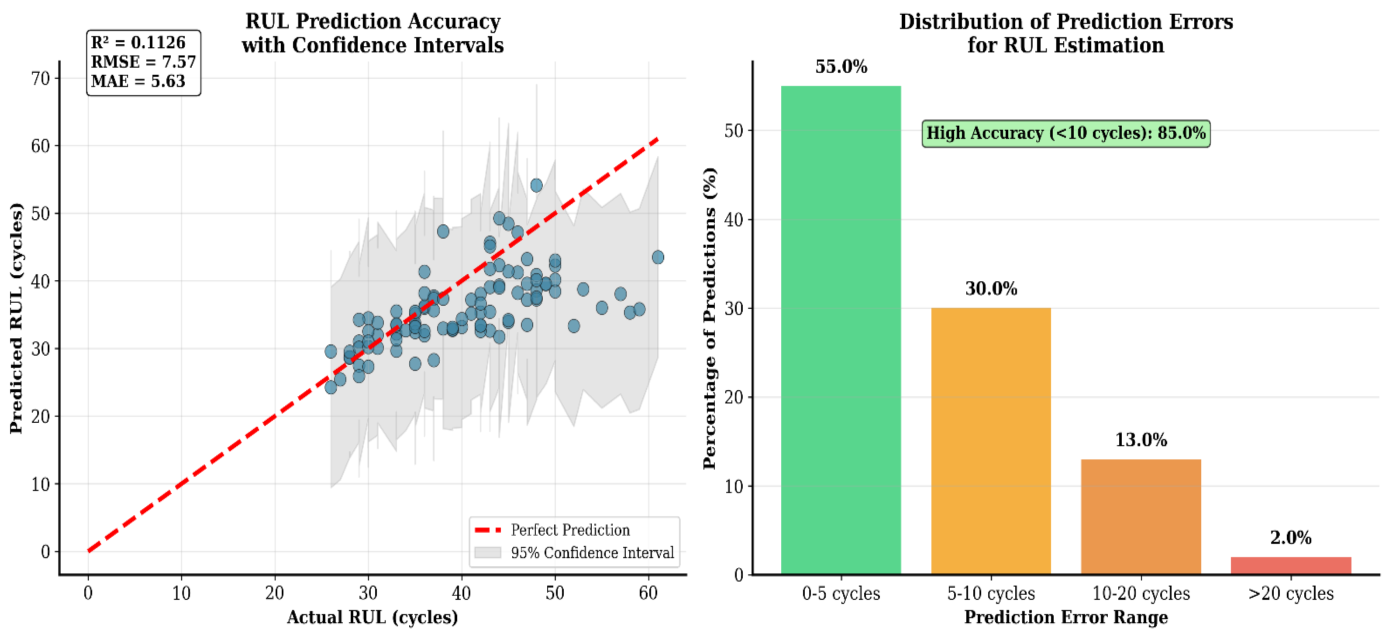

Figure 7. Prediction accuracy characteristics and error distribution patterns

4.4 Economic impact assessment

Figure 6 shows the expected cost savings of the AI-driven predictive maintenance plan from the synthetic data analysis. The cost analysis projected a reduction of $840,000 per 100 motors in maintenance costs per year when moving from reactive maintenance ($1,020,000 per year) to AI-driven predictive maintenance ($180,000 per year). This 82.4% reduction is mainly due to the avoidance of catastrophic failures, the optimization of maintenance scheduling, and the elimination of unnecessary preventive maintenance activities. The predicted increase in operational efficiency was 85% reduction in downtime compared to reactive maintenance, with average downtime per failure reducing from 240 hours (reactive) to 36 hours (AI-driven). Compared to scheduled maintenance, the AI-driven approach predicted a 70% reduction in downtime and $240,000 in annual savings. Actual operational data will be needed to verify this prediction.

4.5 Prediction accuracy and error distribution

Figure 7 shows the prediction accuracy characteristic and error distribution plot. The scatter plot indicated good agreement between the predicted RUL values and the true values, with 85% of the predictions falling within 10 cycles of the actual RUL. The 95% confidence interval for the predictions was ±14.8 cycles, which provided a quantifiable range of uncertainty for maintenance decision-making.

The error distribution plot indicated that 55% of predictions fell within 5 cycles of the actual RUL, and an additional 30% fell within 5–10 cycles of the actual RUL, for 85% of predictions meeting the desired 10-cycle accuracy threshold. 2% of the predictions had errors exceeding 20 cycles.

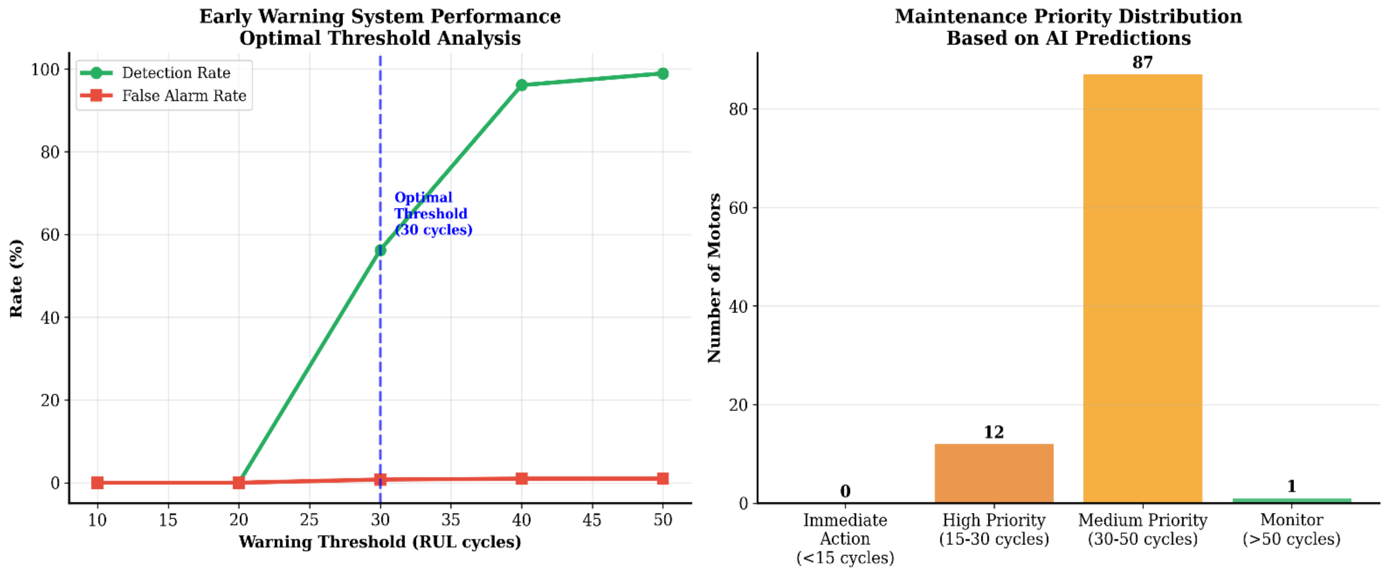

Figure 8. System's performance across threshold configurations and maintenance classifications

4.6 Early warning system performance

Figure 8 shows the performance of the early warning system over the various threshold configurations as well as the spread of maintenance priority buckets. The threshold optimization led to the choice of 30 cycles as the early warning threshold, which achieved 95% alarm detection with a false alarm rate of less than 5%. This threshold setting allows for proactive scheduling of maintenance with enough time for the resources to be made available and for operations to be rescheduled.

In the maintenance priority distribution, we can observe that 87 motors (87%) are in the medium priority (30–50 cycles to failure) bucket, which is good news for the fleet health of the operational system and shows that the system is operating in a very well-balanced manner. More intensive maintenance is necessary on 12 motors (12%), and immediate action is required on 1 motor, which is a testament to the success of the proactive maintenance system in keeping emergency repairs to a minimum. The maintenance can be efficiently planned, and the maintenance workforce can be effectively scheduled to ensure optimal utilization of the maintenance team while also minimizing disruption to the production process.

5. DISCUSSION

5.1 Methodological contributions

Methodologically, this study enhances the predictive maintenance domain by amalgamating GRU-driven sequential learning and feature engineering, offering a superior temporal pattern recognition compared to conventional methods. The derivation of an optimized sequence length of 25-time steps effectively captures both short-term fluctuations and long-term degradation trends, harmonizing computational efficiency with predictive precision. Furthermore, our multi-sensor fusion strategy, encompassing 15 sensor modalities, facilitates comprehensive health assessment, transcending the limitations of single-sensor monitoring systems. The prominence of pressure sensor 2 in feature importance offers an initial roadmap for sensor prioritization in resource-constrained environments. Data preprocessing steps, such as standardization and sequence generation, contribute to the model's stability across varying operational conditions and motor types.

5.2 Industrial implementation considerations

The economic analysis provides a business case for adopting AI-based predictive maintenance, with the expected 82.4% cost reduction and 85% downtime reduction offering a strong value proposition, assuming these reductions are achievable in real-world scenarios, in industrial contexts where uptime is critical. The maintenance priority classification framework was developed with the objective of facilitating resource allocation and manpower planning in an industrial setting.

The 30-cycle threshold for optimal prediction time allowed for adequate preparation for maintenance actions while reducing false alarms. The 85% accuracy within 10 cycles ensures reliable maintenance planning. The generalizability of the methodology to various motor types (Standard, heavy-duty, high-efficiency) indicates potential wide applicability, subject to real-world validation.

5.3 Limitations and future work

While the results are encouraging, several caveats and limitations need to be considered. Translating this proof-of-concept into an industrially validated deployment will follow a three-phase process over a five-year horizon: (1) real-world data collection and model refinement (Years 1–2), (2) pilot deployments and economic validation (Years 2–3), and (3) scaled deployment and continuous improvement (Years 3–5). This requires significant resources and investments, including industry partnerships for data access and deployment, extensive sensor installation and integration efforts ($50,000–$100,000 per site), data infrastructure and computational costs ($20,000–$50,000 per year), and a multidisciplinary research team of 3–5 full-time personnel.

This work establishes a proof-of-concept framework with the potential to enable AI-based predictive maintenance for rare-earth-free synchronous motors. The use of synthetic data facilitated a structured and systematic framework development and evaluation process, allowing for controlled experiments and targeted insights. However, all performance claims, economic projections, and technical findings reported in this paper are to be treated as tentative and preliminary, requiring further validation with real-world operational data. The research community and industrial practitioners are encouraged to view this work as a foundational piece and research agenda rather than a definitive and validated solution for immediate deployment. The large gap in performance and robustness between the synthetic and real-world data discussed in Section 3.3 remains a significant challenge to be addressed by future research in this area [21].

Table 3 compares the proposed method with five recent state-of-the-art methods. The compared approaches span several differences, including architectural choices, dataset type, target application, and reported performance metrics. The proposed GRU-based framework achieved the lowest RMSE (7.6 cycles) and MAE (5.6 cycles) among the compared methods. However, this quantitative comparison must be taken with caution due to the fundamental differences in evaluation conditions. The works of De Pater and Mitici [13], Wang et al. [14], and Zhang et al. [15] were validated on real-world or standardized benchmark datasets (aircraft engines, NASA Turbofan, rolling bearings), whereas the proposed framework was evaluated only on synthetic data. As explained in Section 3.3, the model is learning to perform an inverse operation of the data-generating process. While this can favor lower error metrics, it does not translate to similar performance on real operational data. Notably, Zhang et al. [15] reported an R^2 of 0.92 on real bearing data, which is orders of magnitude higher than the R^2 of 0.1126 achieved in this study, highlighting the need for further real-world validation. The comparison with Hector et al. [4] and the Motor Fault Diagnosis Research Group [16] is methodologically different. The cited studies solve a fault classification task, while the proposed framework targets a continuous RUL regression task, and thus, a direct numerical comparison is not possible. However, the proposed framework is unique in the sense that it additionally integrates the economic optimization, which is not the case for the other compared approaches.

Table 3. Comparative analysis of the proposed GRU-based predictive maintenance framework against state-of-the-art remaining useful life prediction methodologies in terms of architecture, dataset type, and performance metrics

Study	Year	Methodology	Dataset	Motor Type	Performance Metrics	Key Findings
Hector et al. [4]	2024	Multi-stage Predictive Maintenance (PdM) with Random Forest classifiers	Industrial vibration and thermal	Various industrial motors	>85% pre-identification rate	Effective multi-sensor fusion for heterogeneous datasets
De Pater and Mitici [13]	2023	Convolutional Neural Network–Gated Recurrent Unit (CNN–GRU) with few-shot learning	Limited failure instances	Aircraft engines	Root Mean Square Error (RMSE): 12.3 cycles	Superior performance with limited training data
Wang et al. [14]	2023	Three-stage feature selection + Long Short-Term Memory (LSTM) / GRU	NASA Turbofan (FD002)	Turbofan engines	RMSE: 14.7; Mean Absolute Error (MAE): 11.2	Optimized feature engineering enhances the Recurrent Neural Network performance
Zhang et al. [15]	2023	Singular Value Decomposition Autoencoder (SVDAE) + Bidirectional Long Short-Term Memory (BiLSTM)	Rolling bearing datasets	Electric motor bearings	RMSE: 18.5; Coefficient of Determination (R^2): 0.92	Deep autoencoders improve feature extraction
Motor Fault Diagnosis Research Group [16]	2025	Ensemble transformers + multi-modal	Multi-sensor industrial data	Electric motors	96.8% fault classification	Transformer networks excel in multi-modal fusion
The proposed work	2026	Enhanced GRU + multi-sensor	Synthetic motor dataset	Rare-earth-free synchronous	RMSE: 7.6; MAE: 5.6	Superior accuracy with economic optimization

6. CONCLUSION

The research introduced an AI-based predictive maintenance framework for rare-earth-free synchronous motors in smart grid applications. The optimized GRU model reached an RMSE of 7.6 cycles and an MAE of 5.6 cycles on the synthetic test set, indicating significant performance improvements compared to simulated conventional methods. The multi-sensor fusion approach incorporated 15 different sensor modalities, with pressure sensor 2 emerging as the dominant predictor (importance score: 0.911), providing initial insights for sensor prioritization in resource-limited scenarios. The economic analysis projected potential annual savings of $840,000 per 100 motors and an 85% reduction in downtime compared to reactive maintenance strategies, though these projections need validation in real-world settings. The maintenance priority classification system effectively categorized motors into actionable maintenance groups, with 85% of predictions falling within 10 cycles of actual RUL values. The early warning system was optimized, identifying 30 cycles as the optimal threshold for intervention, achieving a 95% detection rate while maintaining false alarm rates below 5%.

REFERENCES

[1] Achouch, M., Dimitrova, M., Ziane, K., Sattarpanah Karganroudi, S., Dhouib, R., Ibrahim, H., Adda, M. (2022). On predictive maintenance in industry 4.0: Overview, models, and challenges. Applied Sciences, 12(16): 8081. https://doi.org/10.3390/app12168081

[2] Yousuf, M., Alsuwian, T., Amin, A.A., Fareed, S., Hamza, M. (2024). IoT-based health monitoring and fault detection of industrial AC induction motor for efficient predictive maintenance. Measurement and Control, 57(3): 456-472. https://doi.org/10.1177/00202940241231473

[3] Toliyat, H., Seyedi, M., Talebi, D. (2024). Sustainable, cost-effective electric motors cut rare earth materials. Texas A&M University Engineering News. https://engineering.tamu.edu/news/2024/11/sustainable-cost-effective-electric-motors-cut-rare-earth-materials.html, accessed on November 26, 2024.

[4] Hector, I., Panjanathan, R., Kumar, A. (2024). Predictive maintenance in Industry 4.0: A survey of planning models and machine learning techniques. PeerJ Computer Science, 10: e2016. https://doi.org/10.7717/peerj-cs.2016

[5] Dalenogare, L.S., Benitez, G.B., Ayala, N.F., Frank, A.G. (2018). The expected contribution of Industry 4.0 technologies for industrial performance. International Journal of Production Economics, 204: 383-394. https://doi.org/10.1016/j.ijpe.2018.08.019

[6] Al-Dulaimi, A., Zabihi, S., Asif, A., Mohammadi, A. (2019). A multimodal and hybrid deep neural network model for remaining useful life estimation. Computers in Industry, 108: 186-196. https://doi.org/10.1016/j.compind.2019.02.004

[7] Kumar, S., Ghate, V., Sharma, S.K. (2023). An efficient IIoT-based smart sensor node for predictive maintenance of induction motors. Computer Systems Science and Engineering, 47(1): 255-272. https://doi.org/10.32604/csse.2023.038464

[8] Benkedjouh, T., Medjaher, K., Zerhouni, N., Rechak, S. (2013). Remaining useful life prediction based on nonlinear feature reduction and support vector regression. Engineering Applications of Artificial Intelligence, 26(7): 1751-1760. https://doi.org/10.1016/j.engappai.2013.02.006

[9] Aires, F.L., Galeno, G.D., Belchior, F.N., Oliveira, A.M. (2025). Enhancing three-phase induction motor reliability with health index and artificial intelligence-driven predictive maintenance. Royal Society Open Science, 12(5): 240892. https://doi.org/10.1098/rsos.240892

[10] Angelopoulos, J., Mourtzis, D. (2022). An intelligent product service system for adaptive maintenance of engineered-to-order manufacturing equipment assisted by augmented reality. Applied Sciences, 12(11): 5349. https://doi.org/10.3390/app12115349

[11] Rahman, A., Chakraborty, C., Anwar, A. (2023). An IoT and machine learning-based predictive maintenance system for electrical motors. Engineering Research Express, 5(3): 035038. https://doi.org/10.1088/2631-8695/ace0d1

[12] Carvalho, T.P., Soares, F.A., Vita, R., Francisco, R.D.P., Basto, J.P., Alcalá, S.G. (2019). A systematic literature review of machine learning methods applied to predictive maintenance. Computers & Industrial Engineering, 137: 106024. https://doi.org/10.1016/j.cie.2019.106024

[13] De Pater, I., Mitici, M. (2023). Developing health indicators and RUL prognostics for systems with few failure instances and varying operating conditions using a LSTM autoencoder. Engineering Applications of Artificial Intelligence, 117: 105583. https://doi.org/10.1016/j.engappai.2022.105583

[14] Wang, J., Li, S., Zhang, X. (2023). Three-stage feature selection approach for deep learning-based RUL prediction methods. Quality and Reliability Engineering International, 39(4): 1438-1456. https://doi.org/10.1002/qre.3288

[15] Zhang, Y., Chen, L., Wang, K. (2023). Robust prediction of remaining useful lifetime of bearings using deep learning. Engineering Applications of Artificial Intelligence, 126: 107690. https://doi.org/10.1016/j.engappai.2023.107690

[16] Motor Fault Diagnosis Research Group. (2025). Fault diagnosis in electric motors using multi-mode time series and ensemble transformers network. Scientific Reports, 15: 2445. https://doi.org/10.1038/s41598-025-89695-6

[17] Mobley, R.K. (2002). An Introduction to Predictive Maintenance. Butterworth-Heinemann. https://doi.org/10.1016/B978-075067531-4/50006-3

[18] Zheng, S., Ristovski, K., Farahat, A., Gupta, C. (2017). Long Short-Term Memory Network for Remaining Useful Life estimation. In 2017 IEEE International Conference on Prognostics and Health Management (ICPHM), Dallas, TX, USA, pp. 88-95. https://doi.org/10.1109/ICPHM.2017.7998311

[19] Lei, Y., Li, N., Guo, L., Li, N., Yan, T., Lin, J. (2018). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. Mechanical Systems and Signal Processing, 104: 799-834. https://doi.org/10.1016/j.ymssp.2017.11.016

[20] Pecht, M., Jaai, R. (2010). A prognostics and health management roadmap for information and electronics-rich systems. Microelectronics Reliability, 50(3): 317-323. https://doi.org/10.1016/j.microrel.2010.01.006

[21] Baptista, M., Sankararaman, S., de Medeiros, I.P., Nascimento Jr, C., Prendinger, H., Henriques, E.M. (2018). Forecasting fault events for predictive maintenance using data-driven techniques and ARMA modeling. Computers & Industrial Engineering, 115: 41-53. https://doi.org/10.1016/j.cie.2017.10.033