Predicting Battery RUL: Integrating Classical Regression and Quantum Machine Learning

S. ANU[1] and *S. RITA[2]
[1]Research Scholar, Department of Statistics, Periyar University, Salem – 11, Tamil Nadu, India
[*2]Professor, Department of Statistics, Periyar University, Salem – 11, Tamil Nadu, India
anuselvam999@gmail.com, ritasamikannu@gmail.com

Abstract

Quantum Machine Learning is recently being considered an emerging area in predictive analytics, primarily due to the ability of quantum computing, which is presumed to process complex patterns more effectively and quickly than current traditional models. Most of these assumptions are yet to be validated using concrete data. For the purpose of this study, I assessed the performance of two popular traditional models, Linear Regression and Random Forest, against the performance of the Quantum Circuit Regressor (QCR) model in the context of estimating the Remaining Useful Life (RUL) of batteries, which is an imperative parameter in scenarios involving electric cars and storing energy. The experiments were conducted using two different data sources, and the assessment of these models was done in similar settings. In the case of the first dataset, Random Forest performed outstandingly, with negligible prediction errors and an R^2 value of almost one, while Linear Regression performed moderately. In contrast, the quantum model struggled to fit the data and resulted in large errors along with negative R^2 values. A similar pattern was observed in the second dataset, where all models faced difficulties, but the quantum approach again failed to show any advantage over the classical methods. These results suggest that although quantum models are conceptually attractive, their practical benefits are not yet evident for this type of problem. In such a context, the Random Forest approach was the most accurate method for battery Remaining Useful Life prediction. The classical approach is, in fact, still a very reliable option, also considering the fact that, although it is a quantum approach, it is still hampered by hardware limitations, noise, or unrefined algorithms. The development of such a technology will, possibly, lead to an increasing role, but it is currently rather a complementary tool rather than a substitute.

Keywords: *Remaining Useful Life, Random Forest Regressor, Quantum Circuit Regressor, Predictive modelling*

1 Introduction

Consider this: the world we live in is battery-powered. These energy storage devices are found in everything from electric cars to spacecraft to your phone. The trouble is that batteries don't last forever, and when they do, it can be anything from deadly to inconvenient.

Remaining Useful Life (RUL) prediction can help with that In essence, I'm attempting to answer a straightforward but essential question: how much life does this battery have left? If I get this right, it means maintenance can be planned before any failure occurs—potentially saving millions of dollars and even lives. As Ferreira and Gonçalves [3] remarked, the increased use of predictive maintenance has fostered a considerable interest in data-driven techniques that can forecast Remaining Useful Life (RUL) in a variety of real-world scenarios.

If you believe that the battery life prediction should be simple, it should be reconsidered.

Similar to erratic teens, battery behavior varies according to temperature, usage frequency, charging method, and even chemical composition. The degradation process is chaotic, unpredictable, and nonlinear. conventional methods that simply monitor a battery to reach a specific threshold? The modern world of complex systems and noisy data makes traditional methods as useful as a chocolate teapot. I require advanced adaptable techniques which will help me understand the chaotic situation.

Researchers have been using traditional statistical approaches for decades. Consider mathematical workhorses such as logistic regression and linear regression [5, 14]. They may not be ostentatious, but they are trustworthy and their actions are comprehensible. These traditional approaches may not be the most interesting, but the traditional methods lack excitement but effectively complete their assigned work. These methods function as dependable companions who always show up exactly when needed. When working on simpler issues with relatively simple interactions between variables, they perform very well.

Everything changed when machine learning emerged. All of a sudden, algorithms became available that could identify patterns in data without requiring us to explicitly state every rule. Neural networks, Random Forests, and Support Vector Machines are like having an extremely intelligent detective that can identify patterns that people would overlook. Researchers began to see remarkable outcomes. For instance, Alfarizi et al. [1] increased accuracy even further by employing optimal ensemble techniques, while Mathew et al. [7] demonstrated the effectiveness of Random Forests in forecasting turbofan engine RUL Random Forests have become extremely popular, and with good cause. These methods work well with the typical noisy data that appears in real-world applications. The Swiss Army knife of machine learning describes these methods because they provide dependable adaptability for practical applications.

Although linear regression may not capture complex nonlinear processes, it is useful for identifying patterns and determining the significance of specific features. In businesses where safety and reliability are critical, I use logistic regression and proportional hazards models [15] to make probabilistic estimation of the failure risk. Alfarizi et al. [1] demonstrated that when Random Forests is optimized and paired with advanced signal processing, they can outperform traditional methods in predicting bearing RUL. Okoh et al. [10] also stressed the importance of combining data-driven models with domain expertise to achieve the best results. It serves as a reminder that the model I chose should accurately reflect the degrading behavior of the system with which I am working. Then comes the intriguing twist: quantum computing has entered the picture, bringing with it some really big claims. This is when things begin to get interesting. These quantum models make use of concepts like entanglement and superposition—don't worry if those sound like science fiction—they actually do. According to the theory, quantum computers may be faster and more potent than classical computers because they can operate in high-dimensional areas with fewer parameters. Even with small datasets, Tsurkan et al. [14] demonstrated some encouraging results with their Hybrid Quantum Recurrent Neural Network (HQRNN), which blends quantum circuits and classical RNNs. The ability of these models to capture intricate temporal connections that traditional algorithms could overlook is very promising.

However, this is where we must apply a slightly higher braking pressure. The theoretical capabilities of quantum computing are remarkable and sound amazing. The gap between theoretical concepts and their actual implementation remains substantial. The quantum hardware currently available is still somewhat noisy and limited. I have noticed that the simulated data used in most studies are significantly cleaner than what we actually encounter in the real world, which is a problem. What's even more frustrating is seeing researchers claim that quantum models are superior—often without conducting rigorous, side-by-side comparisons to back up those claims. I believe we need to challenge these assumptions with empirical evidence rather than just theoretical hype.

It would be equivalent to asserting that a new sports vehicle is faster without actually putting it up against older models on the same course in the same circumstances. Many ML-based RUL studies, according to Ferreira and Gonçalves [3], are opaque and unreproducible, which seriously impedes their use in business. It is practically hard to determine whether reported improvements are due to improved data preprocessing or the model architecture itself in the absence of established evaluation methodologies. The necessity for RUL prediction methods that are not only precise but also scalable and resilient to changes in data quality—conditions that are frequently encountered in industrial manufacturing settings—was further highlighted by Pan et al. [11].

This study specifically addresses this issue. Rather than only presuming that quantum is superior, it fairly compares and contrasts quantum and classical models.A relatively new field called Quantum Machine Learning (QML) combines machine learning techniques with the ideas of quantum computing like superposition and entanglement. It explores how learning tasks like pattern recognition, classification, and clustering might be enhanced by quantum system. The goal is to accelerate complex computations and open up new ways of processing large-scale data that go beyond the constraints of traditional algorithms [13]. Despite its potential, Quantum Machine Learning faces significant difficulties. Current quantum hardware is still limited, with issues including noise, decoherence, and restricted qubit counts prohibiting practical deployment. Furthermore, creating effective quantum algorithms and converting classical data into quantum states are difficult tasks. These challenges emphasize the fact that, while QML provides great potential, much work remains to be done before it can be used at scale to real-world situations [2].

No one is given preferential treatment; all data, preparation, and evaluation criteria are the same. Three methods are examined in the study: Quantum Circuit Regressor (the newest quantum hotshot), Random Forest (the machine learning favorite), and Linear Regression (the dependable old-timer). They evaluate these using actual battery datasets and use common metrics like R-squared and Mean Absolute Error to gauge performance. Most significantly, they take into account interpretability, scalability, and computing cost in addition to accuracy.

Like many other industries, I've witnessed the battery sector suffer from overhyped technologies—grand promises that ultimately failed to deliver. Without proper benchmarking, it is impossible to determine whether performance benefits are due to improved data processing or the model itself. I conducted this research to create a fair comparison between these techniques through objective evaluation of their individual strengths. The goal of my research is to determine the optimal situations for each tool rather than selecting winners or losers.

What is the major takeaway? The success of a method in controlled laboratory conditions with organized data does not ensure its effectiveness in the unpredictable and chaotic

real-world environment. Quantum computing may truly transform battery life prediction, but I'm not convinced until I see solid, actual evidence. The field requires this type of thorough and transparent evaluation process.
The process of challenging assumptions and requiring proof enables us to make better judgments especially when lives and millions of dollars are at risk. The incorrect prediction of battery failure results in more than just inconvenience because it can lead to disastrous outcomes.

2 Experimental Framework

This section discusses the strategy I used to compare classical and quantum models for forecasting battery remaining useful life (RUL). My philosophy throughout was whatever you do, do it blind, reproducible, and scientific. To this end, I applied standardised preprocessing methods, consistent evaluation protocols and explicit model comparison strategies.

2.1 Design Principles:

The experimental design was governed by three main principles: uniformity, reproducibility, and fair comparison. All models underwent the same preprocessing and feature engineering. I utilized fixed random seeds, documented all hyperparameters, and followed a consistent train-test split. All models were evaluated using the same set of performance indicators. The whole workflow consisted of four successive phases: Prepare data, implement a model, Train and optimize, and Evaluation.

2.2. Datasets:

To assess how well the models generalize across various degradation patterns, I chose two complimentary datasets that provide alternative viewpoints on failure behavior. The first dataset, the C-MAPSS FD001 subset, is provided by NASA's Prognostics Center of Excellence. It has 100 training and test units, each with 21 sensor measurements each cycle. The sequence lengths range from 128 to 362 cycles. This dataset is commonly utilized in RUL research because of its well-documented failure paths and benchmark status.The second dataset was NASA's lithium-ion battery dataset, which shows degradation in controlled laboratory circumstances. It includes 18650 cells that have been charged and discharged at room temperature, with voltage, current, temperature, and capacity fade recorded. Cycle counts range from 40 to 168, with a 70% capacity limitation. This dataset captures real-world battery degeneration, making it appropriate for assessing practical model performance.

2.3. Data Preprocessing:

To ensure consistency across models, I developed a standardized preprocessing pipeline where missing values were handled with forward-fill for isolated gaps while incomplete sequences were eliminated. A three-window moving average smoothing approach was applied to decrease sensor noise. Feature engineering comprised rolling statistics (mean, standard deviation, trend) over five-cycle windows, degradation indicators such as cumulative wear and rate of change, and lagged features to capture temporal dependencies. To prevent leakage, all features were standardized using Z-score normalization and statistics from training data. The RUL objectives were computed based on the difference between the maximum and current cycle counts. Finally, train-test splits employed a temporal technique, with 80% of each unit's lifecycle used for training and 20% for validation, while independent test units were reserved for final evaluation to ensure precise temporal ordering and no leakage.

2.4. Model Architectures:

I investigated three models that represent distinct modeling paradigms: linear regression, random forest, and quantum circuit regression.

2.4.1. Linear Regression:

Linear Regression is a simple and widely used statistical method that models the relationship between dependent and independent variables by fitting a straight line. It assumes linearity and provides interpretable coefficients that show the contribution of each predictor to the outcome [4, 9]. To establish a baseline, I used ridge regression with L2 regularization (α = 1.0). Recursive Feature Elimination was used to choose the best predictive sensors. The model assumes a linear relationship between characteristics and the RUL.

$$\hat{y} = \beta_0 + \sum_{i=1}^{p} \beta_i x_i + \epsilon$$

This method provides interpretability and is ideal for recording linear degradation patterns.

2.4.2. Random Forest Regressor:

Random Forest is an ensemble machine learning method that builds multiple decision trees and combines their results for prediction. It provides robust classification, handles complex data patterns, and reduces errors by using out-of-bag (OOB) estimation to measure accuracy [9]. It reduces overfitting compared to single decision trees and can capture complex, non-linear relationships in the data, making it more flexible and robust. The random forest model combines predictions from a set of decision trees.

$$\hat{y} = \frac{1}{T} \sum_{t=1}^{T} h_t(x)$$

I utilized 200 trees with a maximum dpth of 15 to tune hyperparameters using cross-validation. The average decrease in impurity was used to assess feature relevance. This approach is effective at capturing non-linear connections and provides resilience via ensemble averaging.

2.4.3. Quantum Circuit Regressor:

The Quantum Circuit Regressor (QCR) models the link between input features and target variables using quantum computation concepts. In this method, classical input data is first encoded into quantum states using parameterized quantum circuits. The model is then trained by optimizing circuit parameters using a hybrid quantum-classical loop in which a classical optimizer changes the parameters based on a predetermined cost function. Finally, measurement results from the quantum circuit are mapped back to regression predictions. This paradigm enables the QCR to capture complicated, non-linear patterns while maximizing the benefits of quantum parallelism [12].

I used PennyLane v0.32 to build a variational quantum circuit for the quantum simulation. Classical characteristics were encoded with Y-rotation gates, followed by entangling layers made up of parameterized rotations and CNOT gates. The result was the expected value of the Pauli-Z observable:

$$\hat{y} = \langle \varphi(x;\theta) | Z^{\otimes n} | \varphi(x;\theta) \rangle$$

The design featured four qubits, three variational layers, and 48 trainable parameters. Nesterov momentum (learning rate = 0.01, momentum = 0.9) was used for optimization, with a 50-epoch exponential learning rate decay. This arrangement uses quantum qualities like superposition and entanglement to represent complicated feature interactions.

2.5. Evaluation Protocol:

To get a better sense of how well each model really did, I looked at three different but related error measures. To evaluate the performance of the models, common error metrics such as Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Coefficient of Determination (R^2) were used. MAE measures the average size of errors, RMSE gives more weight to larger errors, and R^2 explains how well the model predictions fit the actual data [13]. MAE gave me a basic and easy-to-interpret metric of how far, on average, the forecasts were from the actual (in same units as the RUL) without making it sensitive to a few extreme outliers. Root Mean Squared Error (RMSE) went a step further by penalizing greater mistakes more severely, making it particularly beneficial for identifying models that occasionally made substantial errors. The Coefficient of Determination (R^2) indicates the model's ability to explain a significant portion of the variation in actual RUL values. Values near to 1 indicate accurate predictions.

To ensure the reliability of these results, I employed 10-fold time-series cross-validation while maintaining the data's chronological sequence to prevent future information from entering the training process. I used paired t-tests ($p < 0.05$) to determine whether the performance differences between models were statistically significant. In addition, I evaluated Cohen's d to determine whether any reported improvements were both statistically significant and practically useful. My primary focus was not only accuracy, but also the efficiency with which the models could run. I documented training time, monitored peak memory utilization, and evaluated how well each model scaled as I gave it bigger sections of the dataset. The review used accuracy metrics together with statistical tests and computing benchmarks to evaluate both predicted quality and real-world applicability.

I achieved reproducibility through fixed random seeds across runs and Git version control of experiments and SHA-256 checksums for dataset integrity verification. The final results were based on unseen test units and each experiment was performed ten times with different initializations which resulted in 95% confidence intervals for all metrics.

3 Results

Three regression models were evaluated in this study, namely Linear Regression, Random Forest Regressor, and Quantum Circuit Regressor on two benchmark datasets namely C-MAPSS FD001 turbofan engine dataset and NASA's Lithium-ion battery dataset. Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Coefficient determines the determination (R^2) to evaluate the performance of Models.

The experiments on the FD001 turbofan engine dataset were carried out using three models: Linear Regression, Random Forest Regressor, and Quantum Circuit Regressor. Their performance was assessed using standard error metrics, as shown in **Table 1**. A comparative visualization of the error values across the models is presented in **Fig. 1**, which highlights the relative strengths and weaknesses of each method.

Table 1. Model Performance on FD001 Dataset

Model	MAE	RMSE	R2
Linear Regression	60.52539556	72.46803743	-2.14588668
Random Forest Regessor	62.28727417	74.86379006	-2.357327398
Quantum Circuit Regessor	61.5890962	73.90927714	-2.27226136

Table 1 summarizes the predictive performance of the three models. All models yielded negative R² values, indicating performance worse than a mean-based baseline. Linear Regression achieved the lowest error values (MAE = 60.53, RMSE = 72.47), while Random Forest recorded the highest errors. The Quantum Circuit Regressor achieved intermediate results between the two classical approaches. These results suggest that simple linear and tree-based models struggle to capture the degradation patterns in the FD001 dataset.

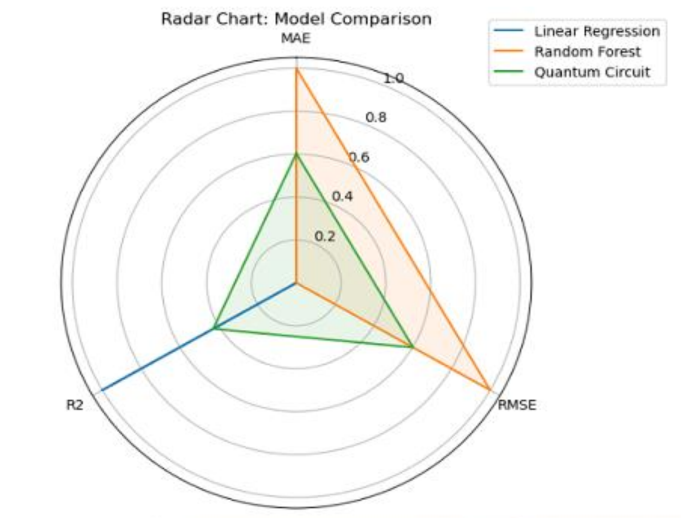

Fig. 1. Radar chart comparing model performance on the FD001 dataset across MAE, RMSE, and R^2

The comparative performance trends are further illustrated in **Fig. 1**, which shows that all three models exhibit similar error magnitudes, with Linear Regression slightly outperforming the others in terms of error metrics.

A similar evaluation was performed on the NASA lithium-ion battery dataset. The three models—Linear Regression, Random Forest Regressor, and Quantum Circuit Regressor—were applied to capture the degradation patterns and predict cycle-based outcomes. The results are summarized in **Table 2**, while **Fig. 2** illustrates the error metrics for direct comparison of model performance on this dataset.

Table 2. Model Performance on NASA Battery Dataset

Model	MAE	RMSE	R2
Linear Regression	82.25885979	135.0801848	0.82466906
Random Forest Regessor	5.66826324	13.35090435	0.99828724
Quantum Circuit Regessor	561.5974999	647.6580996	-3.030573748

Results on the lithium-ion battery dataset are reported in **Table 2**. The Random Forest Regressor achieved the strongest performance across all metrics (MAE = 5.67, RMSE = 13.35, R^2 = 0.998), followed by Linear Regression (R^2 = 0.82). The Quantum Circuit Regressor performed poorly, with substantially higher error values and a negative R^2, indicating weak predictive capability compared to baseline methods .

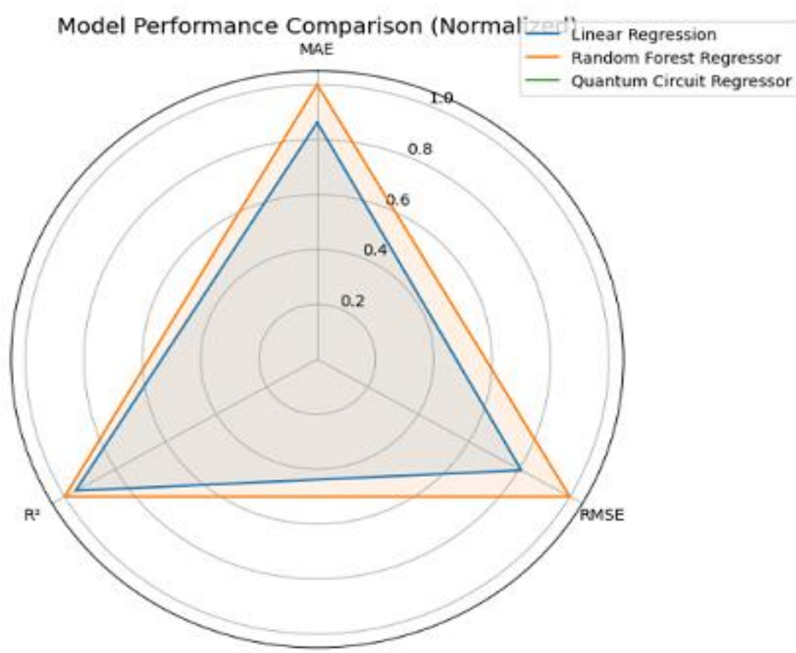

Fig. 2. Radar chart comparing model performance on the NASA Battery dataset across MAE, RMSE, and R^2

As shown in **Fig. 2**, the Random Forest Regressor consistently dominates across all three-evaluation metrics, while the Quantum Circuit Regressor lags considerably behind.

Comparative performance trends are further illustrated in **Figs. 1 and 2** (radar charts).

Random Forest consistently outperformed other models on the battery dataset, whereas model performance on FD001 remained limited across all approaches

4 Conclusion

The research evaluated the predictive capabilities of Linear Regression and Random Forest Regressor and Quantum Circuit Regressor models for Remaining Useful Life (RUL) estimation using NASA C-MAPSS Turbofan Engine Dataset (FD001) and NASA Lithium-Ion Battery Dataset.

The three models failed to predict accurately for the C-MAPSS FD001 dataset because their R^2 values were negative (Linear Regression: –2.15, Random Forest Regressor: –2.36, Quantum Circuit Regressor: –2.27). The error values (MAE and RMSE) showed minimal differences between the models yet the negative R^2 values indicate the complex and noisy nature of turbofan degradation dynamics. The basic regression and quantum models fail to detect engine health degradation patterns because they lack advanced preprocessing techniques and hybridization methods.

In comparison, the NASA Battery Dataset showed significant model performance variation. The best performed to be the Random Forest Regressor, with an extremely high R^2 of 0.998 as well as low MAE (5.67) and RMSE (13.35). Linear Regression was decent (R^2: 0.825) and Quantum Circuit Regressor was really bad (R^2: –3.03) — large prediction errors} This outcome illustrates both the power of ensemble-learning based methods like Random Forest to learn the non-linear degradation characteristics of batteries but also the limitations of current quantum methods for such applications.

In any case, the results highlight two major points: Model performance is heavily dependent on dataset characteristics — A classical and quantum regressor performed poorly on the turbofan dataset but were very different on the battery dataset.

Although still in their infancy stage, quantum circuit regressors are not competitive with classical-based approaches for RUL prediction, where factors such as hardware noise, shallow circuit design, and limited data encoding capability may play a role.

This robustness could be supported by hybrid approaches in future work, for example, through the combination of feature engineering, advanced deep learning techniques, as well as quantum-enhanced methods across the various RUL prediction scenarios.

5 Discussion

Predicting the remaining useful life (RUL) of critical components, such as aircraft engines or lithium-ion batteries, is a formidable technical challenge that also constitutes a foundational requirement for advancing safety, sustainability, and the development of intelligent industrial systems.

Our research, which systematically compares established classical machine learning models with emerging quantum-based approaches on benchmark datasets—including NASA's C-MAPSS turbofan engine dataset and standard lithium-ion battery degradation datasets—directly addresses these interconnected priorities.

They help predict the Remaining Useful Life (RUL) of an asset so that you can catch it early before any sort of catastrophic failure could potentially loss lives or damage infrastructure. Such foresight is not a luxury but rather an imperative in aviation and energy storage. But the perks go far past just safety. Then we can predict when a part will break and stop replacing it too soon or waiting too long. Which leads to reduction in waste, unnecessary repairs and a environmentally more sustainable use of our resources and energy.

It is also aligned with a larger ambition of Industry 4.0, where AI and Quantum Computing radically change the way we observable and maintain complex systems. In the context of an aircraft or electric car, this means smarter diagnostics that can allow for more efficient and less expensive maintenance schedules as well as less downtime.

We find quantum models like the Quantum Circuit Regressor particularly interesting to study. While Random Forest would remain the go-to classic model for most practical use-cases, quantum tools might offer a chance to work with complex data of high dimensionality. They are not quite production-ready, but they are harbingers of things to come, serving as a foundation on which tomorrow's resources and partnerships will be assembled.

Overall, this research presents both scalable, mature products for modern industries and a peek at what could be attainable as quantum technologies evolve. From aircraft, transportation through to energy and predictive maintenance, the implications are huge — and highly relevant.

References

[1] M. G. Alfarizi, B. Tajiani, J. Vatn and S. Yin, “Optimized random forest model for remaining useful life prediction of experimental bearings,” *IEEE Transactions on Industrial Informatics*, vol. 18, pp. 7083–7091, 2022.

[2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe and S. Lloyd, “Challenges and opportunities in quantum machine learning,” *Nature*, vol. 549, pp. 195–202, 2017.

[3] C. Ferreira and G. Gonçalves, “Remaining useful life prediction and challenges: A literature review on the use of machine learning methods,” *Journal of Manufacturing Systems*, vol. 63, pp. 550–562, 2022.

[4] B. Gregorutti, B. Michel and P. Saint-Pierre, “Variable importance assessment in regression: Linear regression versus random forest,” *Computational Statistics and Data Analysis*, vol. 97, pp. 57–74, 2017.

[5] G. James, D. Witten, T. Hastie and R. Tibshirani, *An Introduction to Statistical Learning*. New York, NY, USA: Springer, 2013.

[6] A. H. Karamlou, V. V. Vesselinov, M. E. S. Eftekharzadeh and H. Alghassi, “Quantum kernel methods for solving regression problems and differential equations,” *Physical Review A*, vol. 107, p. 032415, 2023.

[7] H. Liao, W. Zhao and H. Guo, “Predicting remaining useful life of an individual unit using proportional hazards model and logistic regression model,” *IEEE Transactions on Reliability*, vol. 59, pp. 534–547, 2010.

[8] V. Mathew, T. Toby, V. Singh, B. M. Rao and M. G. Kumar, “Prediction of remaining useful lifetime (RUL) of turbofan engine using machine learning,” in *Proc. 2017 IEEE International Conference on Circuits and Systems (ICCS)*, Thiruvananthapuram, India, 2017, pp. 1–6.

[9] S. Murugan, B. M. Kumar and S. Amudha, “Classification and prediction of breast cancer using linear regression, decision tree and random forest,” in *Proc. International Conference on Current Trends in Computer, Electrical, Electronics and Communication (ICCTCEEC-2017)*, 2017, pp. 763–766.

[10] C. Okoh, R. Roy, J. Mehnen and L. Redding, “Overview of remaining useful life prediction techniques in through-life engineering services,” *Procedia CIRP*, vol. 16, pp. 158–163, 2014.

[11] Y. Pan, S. Kang, L. Kong, J. Wu, Y. Yang and H. Zuo, “Remaining useful life

prediction methods of equipment components based on deep learning for sustainable manufacturing: A literature review," *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, vol. 39, p. e4, 2025.

[12] J. Qi, C.-H. H. Yang, P.-Y. Chen and M.-H. Hsieh, "Theoretical error performance analysis for variational quantum circuit based functional regression," *npj Quantum Information*, vol. 9, p. 4, 2023.

[13] H. Sarin, M. Kokkolaras, G. Hulbert, P. Papalambros, S. Barbat and R.-J. Yang, "Comparing time histories for validation of simulation models: Error measures and metrics," *Journal of Dynamic Systems, Measurement and Control*, vol. 125, pp. 578–585, 2003.

[14] M. Schuld, I. Sinayskiy and F. Petruccione, "An introduction to quantum machine learning," *Contemporary Physics*, vol. 56, pp. 172–185, 2015.

[15] O. Tsurkan, A. Konstantinova, A. Sedykh, D. Zhiganov, A. Senokosov, D. Tarpanov, M. Anoshin and L. Fedichkin, "Hybrid quantum recurrent neural network for remaining useful life prediction," arXiv:2504.20823, 2025.

ANU S is a Ph.D. research scholar in the Department of Statistics, Periyar University, Tamil Nadu, India. She has done M.Sc. and M.Phil. in Statistics from Periyar University. Her main research interests are in time series analysis, statistical modeling, machine learning, and quantum machine learning. She has done extensive work comparing classical and modern forecasting techniques such as ARIMA, LSTM, XGBoost, and quantum-inspired models such as QARMA and QLSTM. Currently, she is working on integration of quantum computing methods with traditional statistical methods to improve the accuracy of forecasting for energy consumption and financial time series data.

S. RITA is Professor in Department of Statistics, Periyar University, Tamilnadu, India. She has a

Ph.D. in Statistics and a wealth of experience in teaching and research. Her fields of expertise are statistical quality control, time series, population studies and multivariate analysis. She has supervised numerous post-graduate and doctoral students and has contributed significantly to statistical research and academic administration. She researches real world interdisciplinary problems using statistical methods.