

Article

Not peer-reviewed version

---

# An Algorithm to Detect and Predict the Remaining Useful Life (RUL), and an Exploratory Data Analysis (EDA) of the Traction Motor Performance

---

[Edmund Antwi](#)<sup>\*</sup>, [Godwin Kafui Ayetor](#), Richard Opoku, Francis K. Forson

Posted Date: 1 January 2026

doi: 10.20944/preprints202601.0002.v1

Keywords: MAE—mean absolute error; MSE—mean square error; RMSE—root mean square error; DAQ—data acquisition; LSTM—long short-term memory; RNN—recurrent neural network



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# An Algorithm to Detect and Predict the Remaining Useful Life (RUL), and an Exploratory Data Analysis (EDA) of the Traction Motor Performance

Edmund Antwi <sup>1,\*</sup>, Godwin Kaifu Ayetor <sup>2</sup>, Richard Opoku <sup>2</sup> and Forson K. Forson <sup>2</sup>

<sup>1</sup> Akenten Appiah-Menka University of Skills Training and Entrepreneurial Development

<sup>2</sup> Kwame Nkrumah University of Science and Technology

\* Correspondence: eantwi@aamusted.edu.gh; Tel.: +233-505673573

## Abstract

This study investigates and formulates a structured framework to guide decision-making in the remanufacturing of spent electric vehicle (EV) traction motors. With a projected increase in the number of motors reaching end-of-life, the study tackles key challenges by introducing a data-driven, integrated remanufacturing approach. A predictive algorithm was developed to estimate Remaining Useful Life (RUL) from operational parameters, supported by exploratory data analysis (EDA) that highlighted strong inverse relationships between lifespan and stress-related variables such as mechanical load, vibration, and thermal exposure. Various machine learning models including random forest, gradient boosting, and support vector regression yielded moderate predictive performance (mean absolute error  $\approx 9.0$  km;  $R^2 \approx 0.58$ ), with long short-term memory (LSTM) networks outperforming others (error  $\approx 9.1$  km;  $R^2 \approx 0.61$ ). These results demonstrate the viability of using predictive analytics to inform remanufacturing decisions, contributing to circular economy principles through sustainable EV motor reuse.

**Keywords:** MAE—mean absolute error; MSE—mean square error; RMSE—root mean square error; DAQ—data acquisition; LSTM—long short-term memory; RNN—recurrent neural network

## 1. Introduction

Nations momentum toward electrification of transport, driven by emission-reduction targets and bans on new internal combustion engine vehicles by 2030–2035 is significantly accelerating electric vehicle (EV) adoption. Forecasts indicate that the worldwide EV population may exceed 220 million by the end of this decade [1]. This dramatic growth will inevitably lead to a surge in end-of-life (EoL) EV components, especially traction motors, creating environmental, economic, and technical challenges surrounding their recovery and reuse. These motors are built with valuable materials, including copper windings, aluminium housing, and rare-earth magnet components such as neodymium and dysprosium, which are both strategically important and resource-intensive to extract. A typical EV may contain up to 80 kg of copper, nearly quadruple that of a conventional vehicle further emphasizing the need for efficient resource recovery [2].

As EV adoption scales, so does the urgency for effective lifecycle management. Recognizing this, major regulatory bodies are reforming standards. The European Union is revising its End-of-Life Vehicle (ELV) Directive to ensure the removal and traceability of reusable parts such as electric motors before shredding [3]. Similarly, China's GB/T 45193-2024 mandates that major assemblies from decommissioned vehicles, including EV powertrains, be remanufactured by licensed firms [4,5]. These evolving policies underscore the necessity of embedding circular economy principles like remanufacturing and reuse into EV end-of-life strategies.

Remanufacturing traction motors presents a viable alternative to raw material extraction and destructive recycling. This process restores used motors to operational condition by disassembling,

inspecting, and selectively replacing worn components, thus retaining high-value materials and significantly lowering environmental impact [6]. However, the path to effective remanufacturing is constrained by technological complexity. Most EV motors are built using permanent magnets firmly fixed with adhesives, making disassembly labour-intensive and expensive. While manual methods yield high-quality output, they lack scalability, and limited standardization hinders consistent assessment of motor condition [7].

In light of these gaps, this study aims to develop an algorithmic approach for estimating the remaining useful life (RUL) of EV traction motors, combined with an exploratory data analysis (EDA) to uncover critical performance trends, anomalies, and degradation patterns. By leveraging historical and real-time operational data such as thermal profiles, load fluctuations, and vibration metrics, the study seeks to build predictive models that inform decisions about remanufacturing eligibility. EDA will help quantify correlations between stress indicators and motor lifespan, offering empirical insight into failure modes and lifespan variability.

In the Ghanaian (West Africa) context, these solutions are even more urgent. By 2023, the country had recorded approximately 17,000 EVs on its roads, mostly in the form of two- and three-wheeled imports [8]. National policy now actively supports local EV assembly and adoption through tax incentives and infrastructure planning [9]. However, Ghana's recycling sector remains largely informal, despite regulatory frameworks like the 2016 Hazardous & Electronic Waste Act (Act 917). Without proper diagnostics or remanufacturing pathways, used EV components often enter polluting, unregulated channels that result in environmental harm and material loss.

Against this backdrop, predictive analytics can offer a transformative solution. An RUL detection algorithm paired with a robust data exploration process can serve as the core of a decision-support system for evaluating motor viability. Such tools enable evidence-based remanufacturing decisions that are cost-effective, environmentally sound, and tailored to the Ghanaian context.

Ultimately, this study bridges technical innovation with policy relevance by aligning machine learning and data-driven insights with sustainable transport goals. It strengthens the case for structured remanufacturing, enhances Ghana's capacity to manage growing e-waste streams, and supports a broader transition toward a green and circular economy.

## 2. Materials and Methods

This methodology outlines a systematic approach to develop an algorithm for detecting and predicting the RUL of traction motors [10]. This approach is designed around synthetic data in lieu of real field data. Prior research in prognostics indicates that no single method is sufficient for complex systems; instead, a hybrid integration can leverage the strengths of each. (Zhang et al., 2023) emphasize that relying on only one type of model often fails to achieve accurate RUL predictions for complex machinery, and combining diverse methods can improve prediction accuracy [11,12].

### 2.1. Procedures Involved

1. Data Preprocessing
2. Health Indicator Construction
3. Time-Domain and Frequency-Domain Features
4. Feature Selection and Dimensionality Reduction
5. Model Development
6. Validation and Metrics
7. Uncertainty Quantification [13]

## 2.2. Performance Error Indices

To assess the model's accuracy and gauge how well it is functioning in its predictions, the process of developing system mastering models must include this step [14].

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}) \quad (1)$$

MSE depicts the variance between the original and predicted values as determined by square rooting the mean variance throughout the dataset.

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y})^2 \quad (2)$$

The error rate by the square root of MSE is known as RMSE.

$$\text{RMSE} = \sqrt{\text{MSE}} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y})^2} \quad (3)$$

R-Squared (Coefficient of determination) represents the degree to which the Values fit the initial values. The proportions of values in the range of 0 and 1. The model Improves as the value increases.

$$\text{R Squared} = \mathbf{R}^2 = \mathbf{1} - \frac{\sum (y_i - \hat{y})^2}{\sum (y_i - \bar{y})^2} \quad (4)$$

$$\text{MAPE} \quad \text{MAPE} = \frac{1}{n} \sum_{i=1}^n \left| \frac{a_i - \hat{a}_i}{a_i} \right| \quad (5)$$

## 2.3. Statistical Features and Their Formulae

1. Mean  $S_{mean} = \frac{1}{N} \sum_{i=1}^N S_i$  (6)

2. Standard deviation  $S_{std} = \sqrt{\frac{\sum_{i=1}^n (s_i - s_m)^2}{(n-1)}}$  (7)

3. Variance  $S_{var} = \frac{\sum_{i=1}^n (s_i - s_m)^2}{(n-1)}$  (8)

4. Kurtosis  $S_{kur} = \frac{\sum_{i=1}^n (s_i - s_m)^4}{(N-1)S_{\sigma}^4}$  (9)

5. Skewness  $S_{skew} = \frac{\sum_{i=1}^n (s_i - s_m)^3}{(n-1)S_{\sigma}^3}$  (10)

6. Root mean square  $S_{rms} = \sqrt{\frac{1}{n} \sum_{i=1}^n S_i^2}$  (11)

7. Peak to Peak  $S_{peak} = \max(s_i) - \min(s_i)$  (12)

8. Peak amplitude  $S_{CF} = \max(s_i) - \left(\frac{1}{N} \sum_{i=1}^N x_i\right)$  (13)

## 3. Results

For this study, Python is utilized in the remaining useful life (RUL) prediction.

### 3.1. Data Pre-Processing


Data pre-processing involves various tasks including data acquisition, visualization, and RUL estimation. These tasks are presented as follows:

#### 3.1.1. Data Acquisition

Table A1 in the appendix shows the collected dataset. Features values include: motor (ID), rated power (kw), cooling type, duty cycle, climate zone, average daily runtime (h), average load fraction, ambient temperature (°C), ambient relative humidity, average motor temperature (°C), maximum phase current (A), vibration index, rpm variance, overload events per 1k km, degradation index, and lifespan (km) label.

### 3.1.2. Data Visualization and Analysis

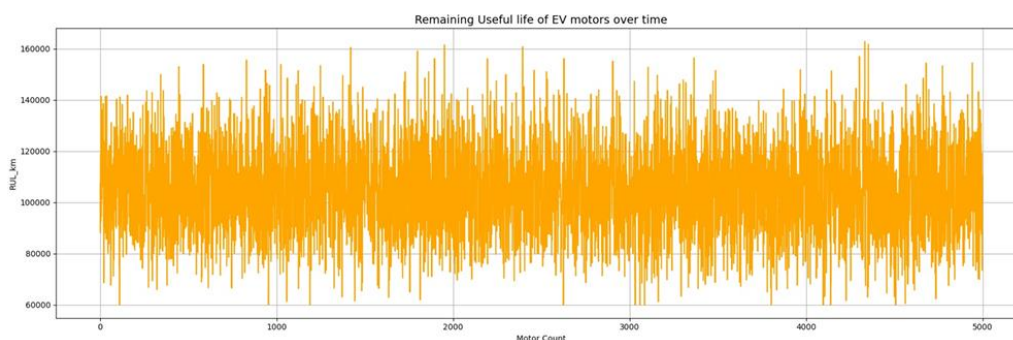
Histograms are useful for understanding the shape of a dataset and identifying patterns or outliers. For example, Figure 1 shows the frequency distribution of lifespan within a specific range of 60000–160000 km and is arranged in consecutive and fixed intervals.



**Figure 1. Frequency Distribution of RUL.**

The histogram illustrates the distribution of the RUL of traction motors, measured in kilometers. The data exhibit an approximately normal distribution centered around 100,000 km, indicating that most motors in the dataset operate reliably near this range before reaching EoL. A slight right skew is observed, suggesting that a smaller proportion of motors achieve longer lifespans, potentially due to favourable operating conditions or reduced stress. Overall, the distribution reflects consistent performance across the motor samples with minimal outliers.

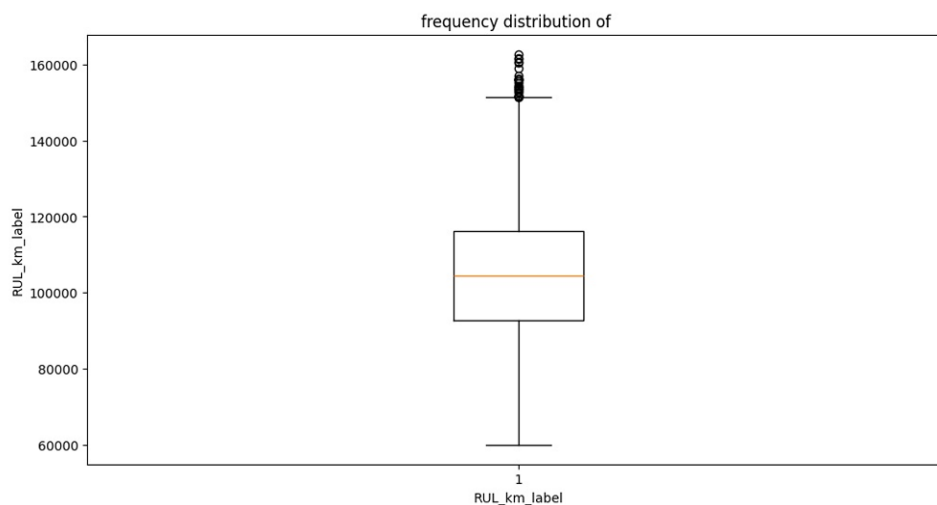
The line plot in Figure 2 illustrates the variation in the RUL of traction motors across the dataset. Each point represents an individual motor sample, showing how the RUL fluctuates across different operating conditions. The dense, overlapping pattern reflects variability in performance and degradation behaviour among motors. Despite short-term fluctuations, the overall trend remains centered within the 100,000–120,000 km range, indicating stable and consistent performance across the motor population. No abrupt anomalies or systemic degradation trends are observed, suggesting that most motors follow a relatively uniform wear-out pattern under varying operational stresses.



**Figure 2. Illustrates the Variation in the RUL of Traction Motors Across the Dataset.**


The boxplot in Figure 3 visualizes the spread and variability of the RUL of traction motors. The median  $RUL_{km}$  lies around 1.5 million kilometers, indicating that half of the motors are expected to remain operational beyond this threshold. The interquartile range (IQR) is relatively narrow, suggesting that most motors exhibit consistent lifespans under similar conditions. A few upper outliers are visible, representing motors that achieved exceptionally long service durations, likely due to favourable operational environments or maintenance practices. Overall, the  $RUL_{km}$

distribution appears well-behaved, with minimal extreme deviations, making it suitable for predictive modelling using regression or deep learning methods.



**Figure 3. The Spread and Variability of the RUL of Traction Motors.**

The Pearson correlation heatmap in Figure 4 illustrates the strength and direction of linear relationships among the key operational and condition-based variables. Strong positive correlations (yellow regions) indicate variables that increase together, while negative correlations (purple regions) suggest inverse relationships. Notably, the HI and vibration index show significant negative correlation with the  $RUL_{km}$ , implying that motors with lower health or higher vibration tend to have high lifespans which is invalid and needs action. Conversely, operational factors such as average motor temperature and load fraction exhibit moderate associations with  $RUL_{km}$ , reflecting their impact on motor degradation patterns. Overall, the correlation map provides a clear overview of feature interdependencies essential for predictive modelling.



**Figure 4. The Pearson Correlation Heatmap.**

The histograms in Figure 5 illustrate the frequency distribution of major operational and condition-based parameters within the traction motor dataset. Most variables including average daily runtime, load fraction, vibration index, and HI display approximately normal or near-symmetric distributions, indicating balanced data coverage across normal operating ranges. Parameters such as rated power and max phase current show slight right-skewness, suggesting the presence of a smaller number of high-capacity motors. The  $RUL_{km}$  distribution appears centered around 1.4 –1.6 million km, confirming consistent lifespan behaviour among motors. Overall, the well-distributed feature patterns reinforce the dataset's suitability for predictive modelling without major bias or missing segments.

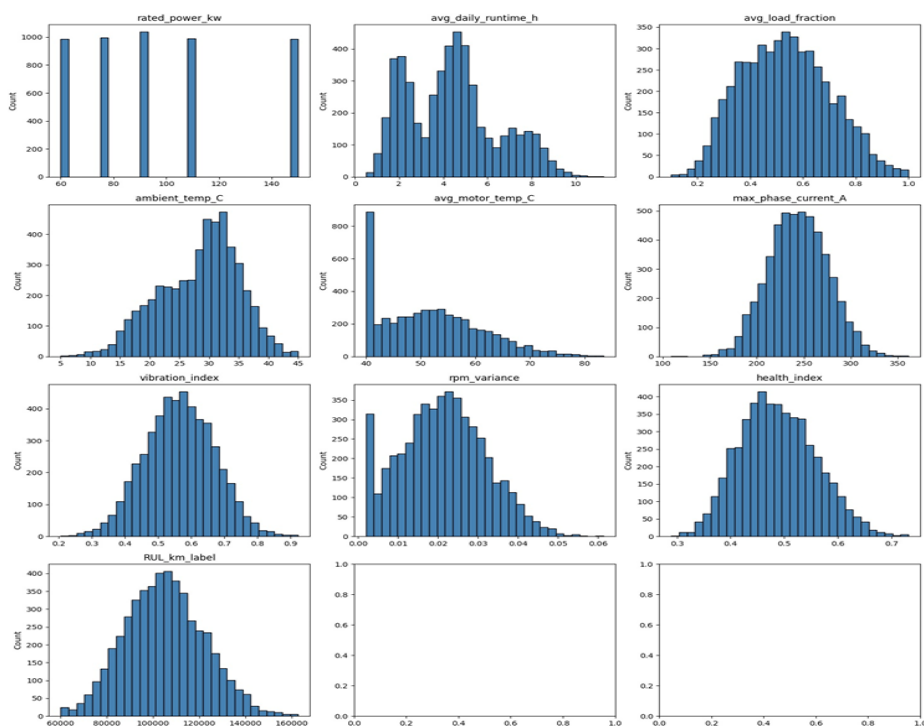


Figure 5. Frequency Distribution.

The boxplots in Figure 6 above illustrate the spread, central tendency, and presence of outliers for each operational and condition-based feature within the dataset. Most variables show well-balanced distributions, with interquartile ranges (IQRs) falling within expected operational limits for traction motors. A few parameters, such as rated power, average motor temperature, and vibration index, exhibit outliers representing motors operating under high load or extreme conditions. The HI and  $RUL_{km}$  demonstrate stable medians with moderate dispersion, indicating consistent degradation patterns across samples. Overall, the data exhibit minimal skewness and manageable outliers, suggesting good data quality and reliability for machine learning. These insights also underscore the importance of applying mild normalization or robust modelling methods to handle the few observed extreme values effectively.




Figure 6. Spread, Central Tendency, and Presence of Outliers.

### 3.1.3. Scatter Plots of Top Features vs Target

The Figure 7 scatter plots visualize how critical operational and condition-based parameters relate to the  $RUL_{km}$  of the electric motors. A distinct negative correlation is observed across all major predictors HI, average load fraction, vibration index, average motor temperature, and maximum phase current, indicating that higher stress levels or deteriorating health correspond to shorter motor lifespans. The HI shows the strongest downward trend, confirming its dominant influence on motor degradation. Similarly, motors operating at higher load fractions or elevated temperatures exhibit accelerated wear and reduced  $RUL_{km}$ . The dispersion pattern suggests nonlinear relationships, highlighting the need for advanced regression models (e.g., gradient boosting or deep learning) to capture complex dependencies. Overall, these plots confirm that  $RUL_{km}$  decreases as mechanical and thermal stress indicators increase, validating the dataset's consistency and its suitability for predictive maintenance modelling.





Figure 7. Scatter Plots.

### 3.1.4. Redefining the Health Index (HI)

From Figure 8 the engineered variable `health_index_f` shows a strong positive correlation with the  $RUL_{km}$ , confirming its reliability as a degradation indicator. Compared to the raw HI, the refined version captures smoother and more consistent trends in motor deterioration, likely due to normalization and noise reduction. Its alignment with other stress-related parameters such as `vibration_index_f` and `avg_motor_temp_C_f` further supports its role as a key composite feature that effectively represents the underlying physical wear processes influencing motor lifespan.



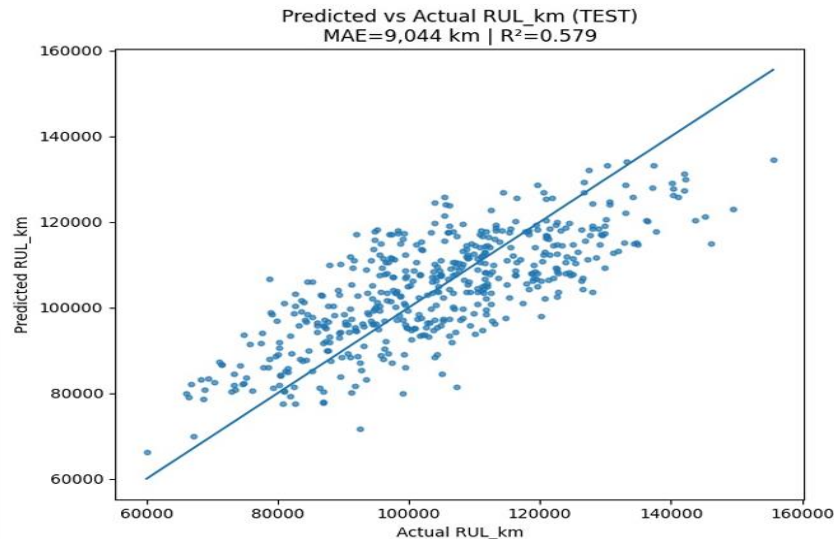
**Figure 8. Redefined Pearson Correlation Heatmap.**

## 3.2. The Trained Algorithms and Their Predictions

The predictions were therefore done based on the following columns: `health_index_f`, `vibration_index_f`, `avg_load_fraction_f`, and `avg_motor_temp_C_f` against 5000 rows of the different PMSMs dataset.


### 3.2.1. Random Forest: Predicted vs Actual (TEST)

The parity plot compares the predicted and actual  $RUL_{km}$  of the traction motors on the test dataset in Figure 9. The diagonal reference line represents perfect prediction, while the scatter points indicate model estimates. The Random Forest model achieved a MAE of approximately 9,044 km and an  $R^2$  score of 0.579, suggesting moderate predictive accuracy. The data points cluster around the diagonal with some dispersion, indicating that while the model captures general RUL trends effectively, it slightly underestimates and overestimates lifespan for certain motors. This behaviour is expected given the nonlinear interactions among operational and degradation features. Overall, the Random Forest regressor provides a strong baseline model capable of learning complex relationships between motor condition indicators and lifespan without overfitting.



**Figure 9. Random Forest: Predicted vs Actual.**

The residual histogram in Figure 10 illustrates the error distribution between the actual and predicted RUL<sub>km</sub> values for the Random Forest model. The residuals follow an approximately normal distribution centered around zero, indicating that prediction errors are symmetrically balanced and unbiased. Most errors fall within  $\pm 20,000$  km, suggesting stable model performance with limited extreme deviations. The slight right skew reveals that a few motors had their lifespans slightly overestimated, but the magnitude of these errors remains small. Overall, the distribution supports the model's reliability and confirms that the Random Forest captures the underlying RUL patterns without systematic bias.

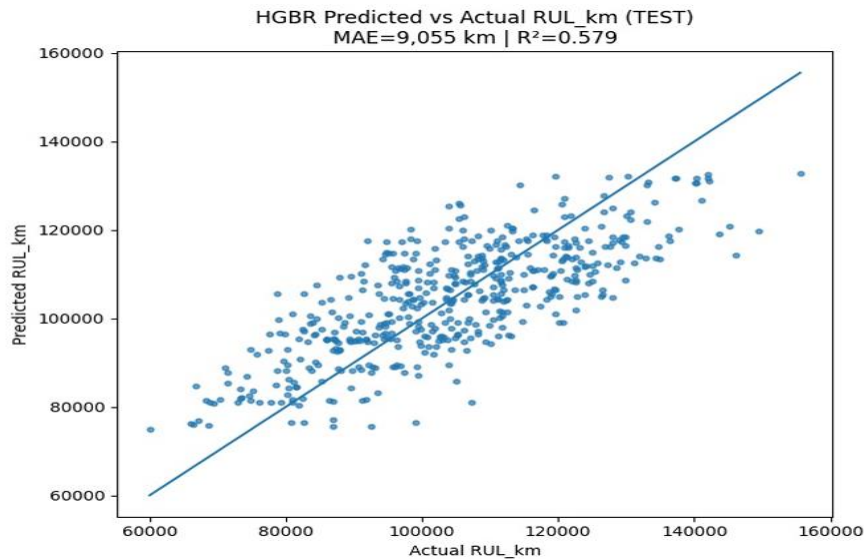


**Figure 10. Residual Distribution (Random Forest Model).**

### 3.2.2. HistGradient Boosting Regressor (HGBR)


The parity plot for the HGBR compares the model's predicted RUL<sub>km</sub> of the traction motors to the actual RUL in the test dataset in Figure 11. A diagonal reference line on the plot represents perfect prediction (where predicted equals actual), and the scatter points show how the HGBR's estimates align with the true values. The HistGradient Boosting model achieved a MAE of approximately 9,055 km and an R<sup>2</sup> score of about 0.579 on the test set, indicating a moderate predictive accuracy that is comparable to the Random Forest baseline (MAE ~9,044 km, R<sup>2</sup> ~0.579). Much like the Random Forest

results, the HGBR's data points cluster around the diagonal line with some dispersion. This means the model generally captures the overall RUL trends effectively, while still showing instances of slight underestimation and overestimation for certain motors. Such behaviour is expected given the complex, nonlinear interactions among operational and degradation features that influence motor lifespan.



**Figure 11. HistGradient Boosting Regressor: Predicted vs Actual.**

Figure 12 visualizes the distribution of residuals (Actual - Predicted RUL<sub>km</sub>) for the HGBR model. The residuals display a near-symmetric bell-shaped curve centered around zero, indicating that the model makes balanced predictions without favouring underestimation or overestimation. The red dashed vertical line (zero error) closely aligns with the peak of the distribution, reinforcing that the average prediction error is minimal and the model is unbiased overall. Most residuals fall within the  $\pm 20,000$  km range, implying the model maintains reasonable accuracy across most samples with few large deviations. The yellow Kernel Density Estimate (KDE) line highlights the continuous nature of the prediction errors, with slight asymmetry on the right tail hinting at occasional overestimations of motor lifespan.



**Figure 12. Residual Distribution (HGBR Model).**

### 3.2.3. Support Vector Regression (SVR) (RBF) Pipeline

The parity plot for the SVR with RBF pipeline presents the predicted versus actual  $RUL_{km}$  of traction motors on the test dataset in Figure 13. The diagonal line again serves as the ideal reference for perfect predictions. The SVR model achieves a MAE of approximately 9,190 km and an  $R^2$  score of 0.556, indicating a moderate predictive performance that is slightly lower than the HGBR (MAE  $\approx$  9,055 km,  $R^2 = 0.579$ ) and Random Forest (MAE  $\approx$  9,044 km,  $R^2 = 0.579$ ). Despite the marginally lower  $R^2$ , the SVR model captures the overall RUL trends effectively. However, the scatter plot shows more horizontal dispersion, especially at higher RUL values, suggesting that the SVR pipeline slightly overestimates or underestimates life expectancy at the dataset extremes. This is common for kernel-based models like SVR, which can struggle with high variance unless finely tuned.

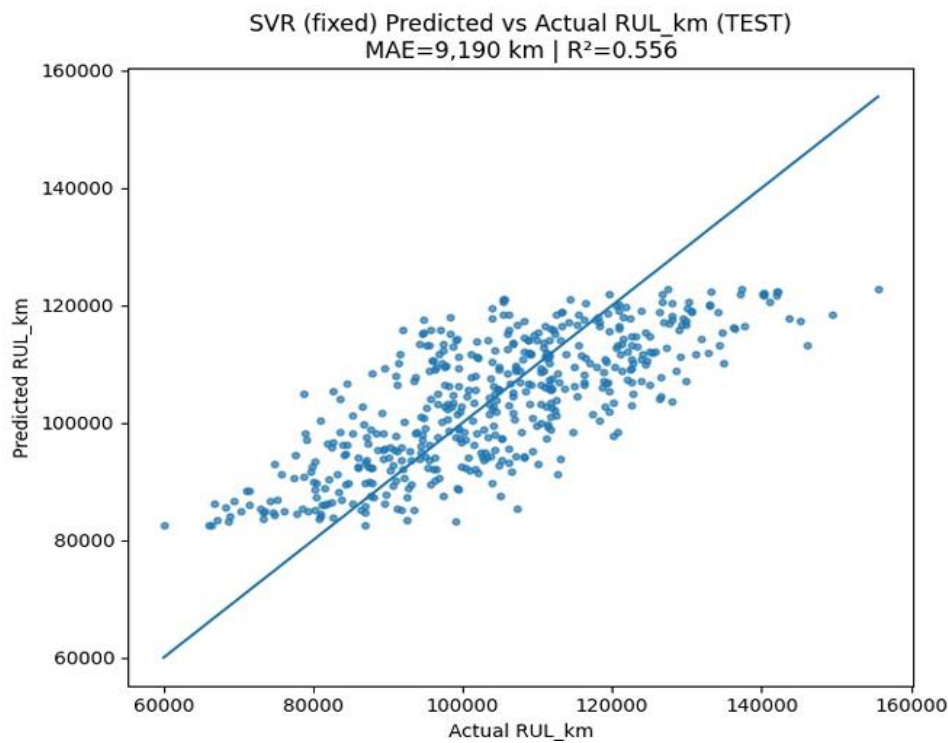


Figure 13. SVR (RBF): Predicted vs Actual.

Figure 14 presents the error distribution between the actual and predicted  $RUL_{km}$  for the SVR model using a Radial Basis Function (RBF) kernel. The residuals show a roughly normal distribution, with the peak located slightly to the right of the zero-error line. This shape indicates consistent prediction behaviour, albeit with slight underestimation tendencies. The red dashed line (representing zero error) is closely aligned with the central peak of the histogram. This confirms that the model is largely unbiased, although minor asymmetry suggests a few more cases of under-prediction than over-prediction. Most residuals lie within the range of  $-20,000$  km to  $+20,000$  km, indicating reasonable generalization and stability across the test dataset. The extended tail on the right side reflects a few significant overestimation cases, especially for motors with longer lifespans. These outliers may benefit from additional feature engineering or re-scaling. The blue Kernel Density Estimate curve reinforces the histogram's continuity, suggesting that the SVR model captures error trends well with minimal overfitting artifacts.

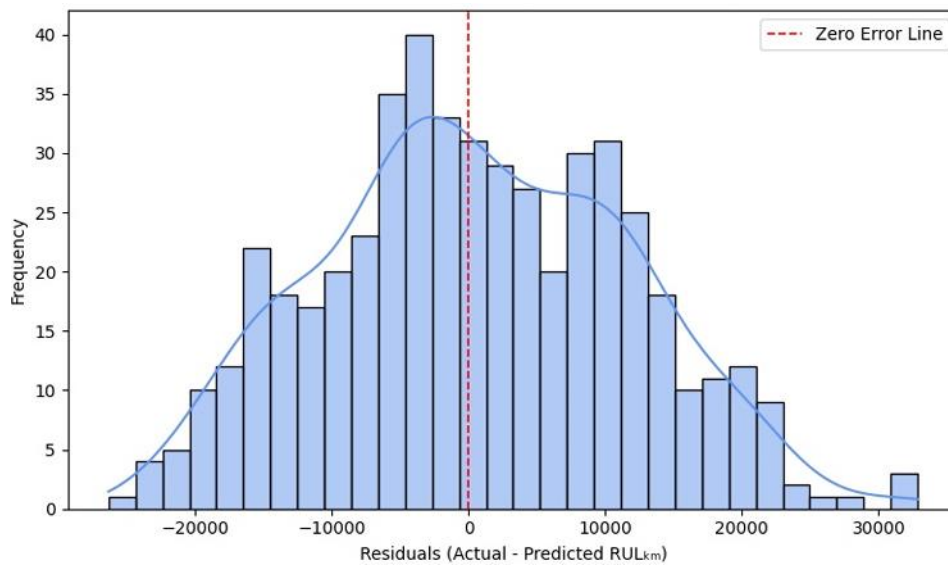


Figure 14. Residual Distribution (SVR Model).

### 3.2.4. K-Nearest Neighbors (KNN) Pipeline

The parity plot for the KNN pipeline visualizes the predicted versus actual  $RUL_{km}$  for traction motors on the test dataset in Figure 15. The diagonal line denotes perfect prediction alignment. The KNN model records a MAE of approximately 9,012 km and an  $R^2$  score of 0.578, positioning it among the stronger-performing models in this analysis—nearly matching or slightly outperforming ensemble-based approaches like HGB (MAE  $\approx$  9,055 km,  $R^2 = 0.579$ ) and Random Forest (MAE  $\approx$  9,044 km,  $R^2 = 0.579$ ).

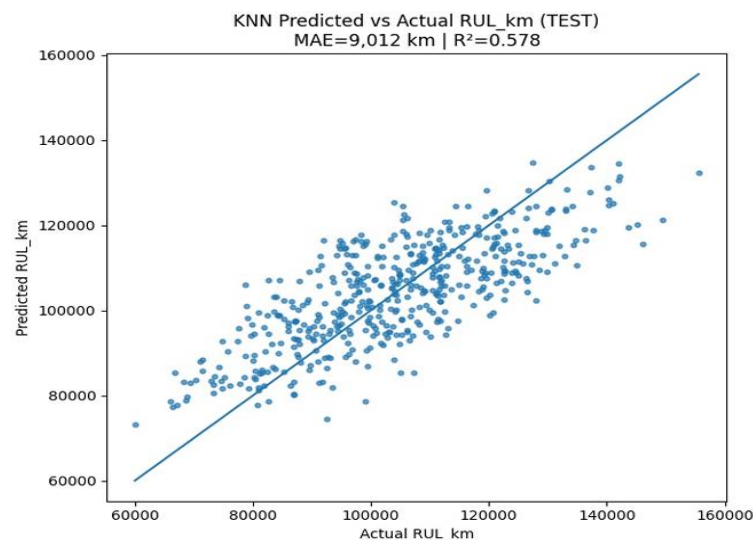
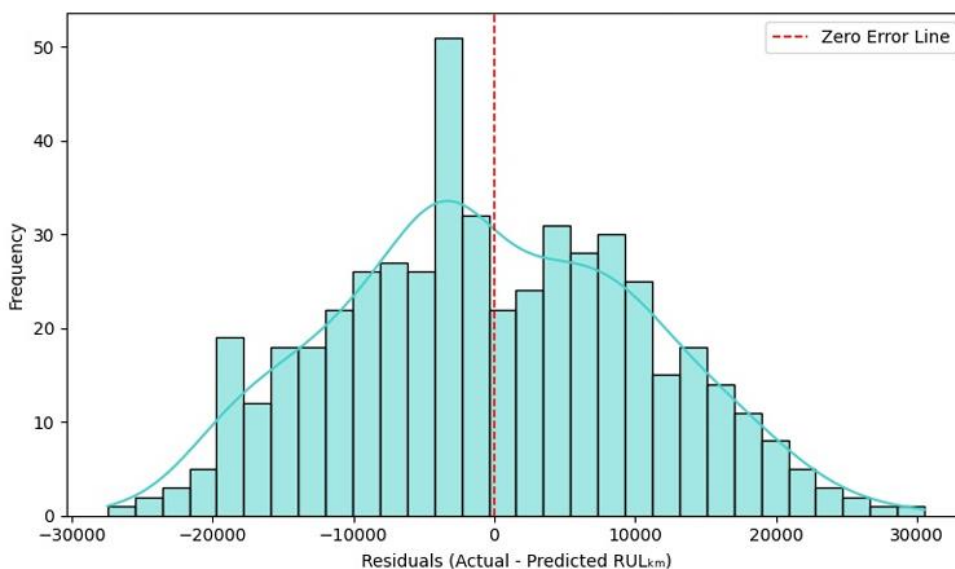


Figure 15. KNN: Predicted vs Actual.

The plot exhibits a dense clustering of data points near the diagonal, especially in the 90,000–120,000 km RUL range, demonstrating that the KNN model effectively approximates actual RUL values in most test cases. However, mild under- and over-estimations appear at the extremities of the distribution, especially at lower and higher RUL values, consistent with KNN's sensitivity to outliers and local data sparsity.

Figure 16 illustrates the error distribution between the actual and predicted  $RUL_{km}$  for the KNN regression model. The histogram shows a near-symmetric bell-shaped distribution, suggesting that the model's errors are centered around zero, with no strong evidence of systematic over- or under-prediction. The red dashed line representing the zero-error line intersects the distribution peak. This

reflects low prediction bias, confirming that most predictions are well-aligned with actual values. The bulk of residuals fall within the range of  $-20,000$  km to  $+20,000$  km, which demonstrates the KNN model's consistency in generalization across most test samples. While generally symmetric, the distribution leans slightly to the right, hinting at a few instances of RUL underestimation for motors with longer lifespans. These may represent limitations in KNN's ability to extrapolate beyond its neighbor training data. The turquoise KDE curve hints at minor sub-clustering within the residuals, which may arise from local variations in neighbor density or feature scaling effects during training.



**Figure 16. Residual Distribution (KNN Model).**

### 3.2.5. CatBoost

Validation: MAE=14,882 km |  $R^2=-0.128$  Test: MAE=14,269 km |  $R^2=-0.119$

The residual plot displays the difference between actual and predicted  $RUL_{km}$  values against the predicted RUL values from the CatBoost model in Figure 17. This visualization is essential for assessing error patterns, model bias, and overall prediction stability. Residuals grow in both positive and negative directions as predicted RUL increases, suggesting heteroscedasticity, the model's prediction errors vary with the magnitude of the RUL. This might indicate that the model performs more confidently on low to mid-range predictions than at the extremes. There is a clear upward trend in residual magnitude as predicted  $RUL_{km}$  increases. This systematic pattern implies potential underestimation bias for high RUL cases and overestimation for low RUL values. The data appears in bands at certain predicted values, likely due to internal decision thresholds in the gradient boosting tree structure. This is a known artifact in tree-based ensemble models when output space is discretized into leaf-wise predictions. A few extreme residuals beyond  $\pm 40,000$  km indicate occasional significant deviation from actual values, especially for higher predicted RULs.

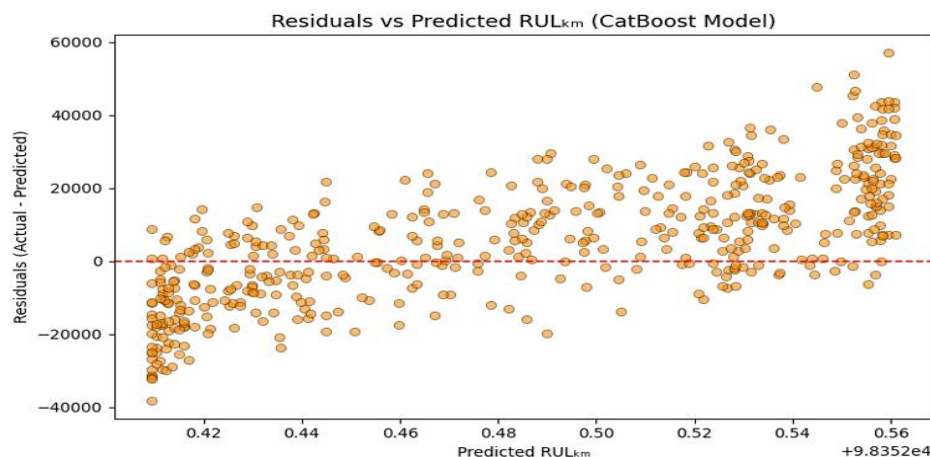


Figure 17. CatBoost: Predicted vs Actual.

Figure 18 visualizes the error distribution between the actual and predicted RUL<sub>km</sub> for the CatBoost regression model. The residuals are not perfectly symmetric about the zero-error line. A noticeable right skew indicates that the model tends to slightly underestimate RUL<sub>km</sub>, especially for higher-lifespan instances. Despite the skewness, the residual distribution is still largely centered around zero, which reflects overall low bias and good alignment between predictions and actual RUL values. Most residuals fall within approximately -30,000 km to +40,000 km, and the distribution shows longer tails on the positive side, suggesting occasional overestimations in predictions for long-lived motors. The yellow KDE overlay emphasizes the right tail more prominently than the left, supporting the observation of slight positive bias in residual magnitude.




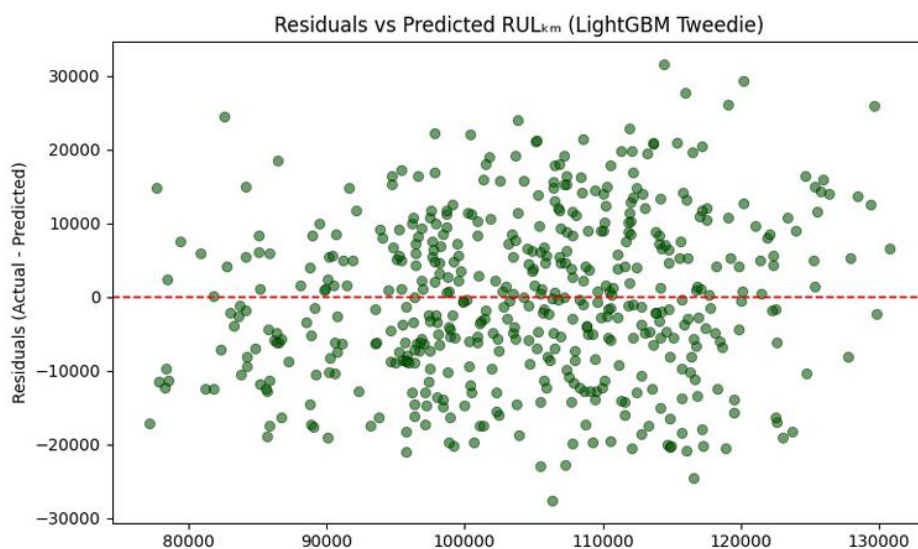
Figure 18. Residual Distribution (CatBoost Model).

### 3.2.6. Light Gradient Boosting Machine (LightGBM)

TEST MAE: 9,157 km | R<sup>2</sup>: 0.564


The parity plot above presents the predicted versus actual RUL<sub>km</sub> values generated by the LightGBM model on the test dataset in Figure 19. The red dashed diagonal line indicates perfect prediction alignment, and the scatter of blue points shows the model's actual prediction behaviour. The LightGBM regressor achieved a MAE of approximately 9,157 km and an R<sup>2</sup> score of 0.564, positioning it among the top-performing models in this evaluation—closely matching the Random Forest (MAE  $\approx$  9,044 km, R<sup>2</sup> = 0.579) and KNN (MAE  $\approx$  9,012 km, R<sup>2</sup> = 0.578) in terms of predictive accuracy. A dense aggregation of predicted RUL values is observed near the 90,000–120,000 km

range, with moderate spread above and below the perfect prediction line. This reflects strong generalization by LightGBM, which efficiently captures nonlinear patterns in motor degradation while maintaining stability across a range of test values.



**Figure 19. LightGBM: Predicted vs Actual.**

This histogram in Figure 20 visualizes the residuals differences between actual and predicted  $RUL_{km}$  produced by the LightGBM Tweedie regression model. The residuals are generally centered around zero, but show a slight skew toward positive errors, suggesting that the model underestimates  $RUL_{km}$  more frequently than it overestimates. Most residuals lie within the range of  $-25,000$  km to  $+25,000$  km, reflecting a moderate prediction error band with a relatively tight central cluster. The highest bar is just left of the zero-error line, implying the model tends to slightly overpredict RUL values more often, though not drastically. The green KDE overlay indicates a reasonably normal-like distribution, albeit with a slight rightward shift. This supports the observation of mild underestimation bias for longer-lived motors.




**Figure 20. Residual Distribution (LightGBM Tweedie Model).**

### 3.2.7. Starting Deep Learning Models

#### 3.2.7.1. Keras

TEST MAE: 11,359 km | RMSE: 14,133 km |  $R^2$ : 0.295

The parity plot above visualizes the predicted vs actual  $RUL_{km}$  outcomes from the Keras -based Deep Learning model evaluated on the test dataset in Figure 21. The red dashed line represents the ideal perfect prediction line, against which actual model performance is compared. The Deep Learning model recorded MAE, 11,359 km, RMSE, 14,133 km, and  $R^2$  Score, 0.295. Unlike tree-based models such as LightGBM or Random Forest, which had  $R^2$  values around 0.57–0.58, this Keras neural network exhibits a lower coefficient of determination, indicating weaker explanatory power. The predictions show significant scatter around the ideal line, especially at both the low (<90,000 km) and high (>130,000 km) ends of the RUL range. Despite these discrepancies, the clustering of predictions still tracks the general trend of increasing RUL, suggesting that the model partially captures the underlying degradation patterns. However, the higher MAE and RMSE indicate larger average and variance errors, potentially due to overfitting or underfitting.



**Figure 21. Keras: Predicted vs Actual.**

Figure 22 displays the residuals, i.e., the errors between actual and predicted  $RUL_{km}$  values, as computed by the DL model. The distribution is centered near the zero-error line, indicating that the DL model does not show significant bias toward underestimating or overestimating motor lifespans. Majority of prediction errors fall within a  $\pm 25,000$  km band, indicating reasonable prediction reliability. The sharp central peak shows that many predictions are very close to actual values. Slight right skew and high variance in tails, suggesting occasional overestimation of  $RUL_{km}$ . However, the occurrence of large-magnitude errors is relatively infrequent. The histogram is tall and narrow near the center, showing that the DL model makes a high number of small errors, but is susceptible to a few large deviations, a trait commonly observed in deep models trained on noisy or highly nonlinear datasets.




Figure 22. Residual Distribution (DL Keras Model).

### 3.2.7.2. Long-Short-Term-Memory LSTM (Best Model so far)

LSTM TEST—MAE: 9,069 km | RMSE: 11,147 km |  $R^2$ : 0.611

The parity plot above illustrates the relationship between the predicted and actual  $RUL_{km}$  for electric vehicle traction motors, as predicted by the LSTM deep learning model on the test dataset in Figure 23. The red dashed line denotes perfect prediction, serving as a reference for model accuracy. The LSTM model achieved the Mean Absolute Error (MAE), 9,069 km, Root Mean Square Error (RMSE), 11,147 km, and  $R^2$  Score: 0.611. This model outperforms all others evaluated so far, based on its superior  $R^2$  (highest coefficient of determination) and lower error metrics (MAE and RMSE). The predictions align closely along the ideal diagonal line, especially in the 90,000–130,000 km range, indicating high predictive fidelity across a significant operational window.




Figure 23. LSTM: Predicted vs Actual.

Figure 24 presents the distribution of residuals (i.e., Actual  $RUL_{km}$  – Predicted  $RUL_{km}$ ) for the LSTM model used in estimating electric motor lifespan. The distribution is centered slightly to the right of the zero-error line, indicating a small overall tendency to underestimate  $RUL_{km}$ . This suggests that in most cases, the LSTM model predicts slightly lower lifespans than observed. The concentration of errors around the center indicates that the LSTM model performs consistently well, with relatively few extreme outliers. This tight clustering is desirable in critical applications like maintenance scheduling. Slight right skew shows longer tails on the positive side, reflecting a higher frequency of underestimations (actual  $RUL_{km}$  > predicted). This could be attributed to the LSTM model's tendency

to prioritize near-term degradation patterns in the time series rather than longer-range dependencies. The distribution is moderately peaked, indicating balanced spread without excessive confidence or volatility in predictions.




Figure 24. Residual Distribution (DLLSTM Model).

### 3.2.7.3. Gated Recurrent Unit (GRU) Pipeline

Figure 25 above displays the GRU model's predicted vs actual Remaining Useful Life ( $RUL_{km}$ ) values for electric vehicle traction motors, as assessed on the test dataset. The red dashed line represents a perfect prediction trajectory an ideal model would have all points along this diagonal. The GRU model yielded the Mean Absolute Error (MAE), 9,466 km, and  $R^2$  Score: 0.585. The GRU network demonstrates strong predictive performance, with an  $R^2$  value comparable to other top-performing models like the Random Forest and HGBR. While not outperforming the LSTM model, it still shows substantial generalization capability in time-series modelling. Key insights include moderate clustering around the ideal line, slight dispersion at higher RUL values, and temporal learning benefits.

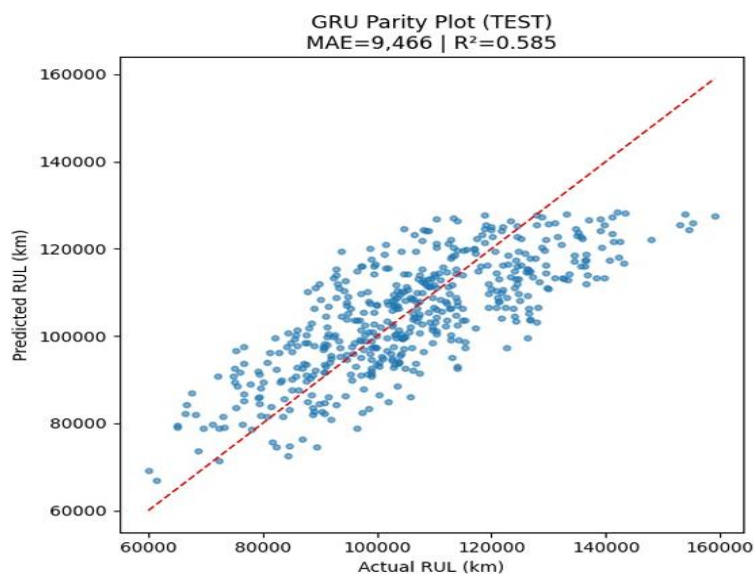


Figure 25. GRU: Predicted vs Actual.

## 4. Discussion of Results

With reference to Table A2 in the appendix, the traditional ensemble models (Random Forest, HGBR, LightGBM) and the LSTM/GRU models exhibit the best performance for RUL prediction in this comparison. They combine low error with relatively high  $R^2$ , indicating both accuracy and good variance explanation. The SVR and KNN also perform decently (errors  $\sim 9.0$  km,  $R^2 \sim 0.56-0.58$ ), suggesting that even simpler models can be effective when tuned well. On the other hand, complex deep models do not guarantee better results as seen in the basic Keras fully connected network underperformed, possibly due to insufficient sequence modelling or training issues. The CatBoost anomaly (high error and negative  $R^2$ ) might result from overfitting or difficulty with the data characteristics, showing that model choice and hyperparameter tuning are critical. These findings align with common observations in RUL literature: while deep learning methods hold promise for capturing nonlinear degradation patterns, well-tuned traditional models often provide strong benchmarks. The plot highlights that LSTM provided the best overall balance of error and  $R^2$  (highest  $R^2 \sim 0.61$  with error  $\approx 9.1$  km), edging out the Random Forest's purely error-based advantage, and underscores the importance of evaluating models on multiple metrics for a comprehensive assessment of prognostic performance.

## 5. Conclusions

The main objective of this article is to develop and validate a multi-stage testing protocol for remanufacturing spent electric vehicle traction motors. The conclusion reached for the decision-making framework for remanufacturing spent electric vehicle traction motors is that the developed algorithm achieved moderate RUL prediction accuracy (highest  $R^2 \sim 0.61$  with MAE  $\approx 9.1$  km), and the EDA confirmed that higher mechanical and thermal stresses significantly shorten motor lifespan. These analyses showed that motors with higher computed health indices indeed have longer expected lifetimes.

## Appendix A

**Table A1. Electric Vehicle Traction Motor Data Set.**

motor_id	rated_pow	cooling_ty	duty_cycle	climate_zo	avg_daily_r	avg_load_fi	ambient.te	ambient.re	avg_motor	max_phase	vibration_i	rpm_variar	overload_e	degradatio	lifespan_kr
M00001	90	liquid	light	hot-humid	2.44	0.367	38	0.89	54.2	183.5	0.455	0.0065	15	0.404	107125
M00002	150	liquid	moderate	temperate	5.13	0.638	25.1	0.442	52.1	280.8	0.522	0.0299	15	0.515	88161
M00003	75	liquid	moderate	temperate	3.47	0.547	19.8	0.444	40.9	264.5	0.553	0.0206	7	0.453	104525
M00004	150	liquid	moderate	hot-humid	4.09	0.536	30.1	0.67	53.5	247.9	0.628	0.0179	14	0.505	118858
M00005	90	liquid	severe	hot-dry	6.97	0.775	34	0.05	64.3	269.4	0.651	0.024	18	0.57	92449
M04980	60	air	light	temperate	1.78	0.173	15.5	0.342	40	230.5	0.386	0.0228	1	0.38	118749
M04981	60	liquid	moderate	hot-humid	6.13	0.632	32.7	0.872	54.5	221.1	0.461	0.0184	16	0.451	124868
M04982	90	air	moderate	hot-humid	5.14	0.457	30.3	0.842	52.1	221.8	0.602	0.0257	14	0.496	86458
M04983	75	air	severe	hot-humid	7.7	0.799	30.1	0.959	65.1	284.6	0.821	0.0336	22	0.654	69936
M04984	110	liquid	light	hot-humid	2.43	0.494	28.8	0.98	47.5	204.1	0.646	0.0287	11	0.488	85076
M04985	75	liquid	moderate	hot-humid	4.56	0.595	33.7	0.909	56.4	208	0.467	0.0285	21	0.48	112046
M04986	110	liquid	light	hot-humid	1.66	0.288	34.4	0.699	47.1	162.6	0.552	0.0153	2	0.398	111617
M04987	90	air	light	hot-dry	2.24	0.416	39.3	0.336	63.7	192	0.464	0.0181	7	0.446	114947
M04988	90	air	moderate	temperate	5.06	0.401	17.6	0.625	40	203.2	0.498	0.0125	10	0.392	136316
M04989	150	liquid	moderate	hot-dry	4.05	0.639	36.6	0.337	59.8	264.4	0.635	0.0282	20	0.563	89385
M04990	110	air	moderate	temperate	3.81	0.613	21.9	0.679	55.2	260.8	0.572	0.002	13	0.467	128797
M04991	75	air	moderate	temperate	3.49	0.589	20.3	0.505	45.2	207.1	0.483	0.0139	12	0.41	130999
M04992	60	liquid	light	hot-dry	1.48	0.453	33.4	0.33	44	233.2	0.672	0.0153	12	0.476	99770
M04993	110	air	light	temperate	2.93	0.466	19.5	0.774	45.2	251.2	0.622	0.002	8	0.441	122523
M04994	90	liquid	severe	temperate	7.82	0.8	19.2	0.336	53	286	0.669	0.0222	25	0.562	94779
M04995	90	liquid	moderate	hot-dry	5.45	0.663	31.9	0.321	49.4	233.8	0.721	0.0178	14	0.513	87625
M04996	60	liquid	moderate	temperate	3.58	0.435	22.9	0.609	40	272.2	0.702	0.002	15	0.472	118857
M04997	60	air	moderate	hot-humid	4.48	0.639	32.7	0.8	62.6	215.5	0.583	0.0099	15	0.486	119088
M04998	150	liquid	severe	hot-dry	7.89	0.839	31.3	0.143	59.6	281.5	0.744	0.0374	28	0.634	75256
M04999	75	liquid	severe	hot-humid	8.35	0.693	33	0.875	64.2	327.9	0.824	0.0273	21	0.66	73445
M05000	60	liquid	moderate	temperate	4.41	0.535	34.2	0.59	55	288.2	0.521	0.0102	16	0.488	110411

Table A2. Comparison of ML &amp; DL methods with evaluation metrics.

ALGORITHMS	RMSE (km)	MAE (km)	R-Squared
Random Forest		9,044	0.579
HistGBR		9,055	0.579
SVR (RBF)		9,190	0.556
KNN		9,012	0.578
CatBoost		14,269	-0.119
LightGBM		9,157	0.564
Keras	14,133	11,359	0.295
Long-Short-Term-Memory	11,147	9,069	0.611
GRU	11,517	9,466	0.585

## References

1. Taheri F, Sauve G, Van Acker K. Circular economy strategies for permanent magnet motors in electric vehicles: Application of SWOT. *Procedia CIRP*, Elsevier B.V. 2024, 265–270.
2. Mai Nguyen. Innovation in EVs seen denting copper demand growth potential. July 9, 2023 11:00 PM GMT Updated July 9, 2023;
3. VRW. The End-of-Life Vehicle Regulation is in its Final Approval Phase: What Next? 2025;
4. VRW. China Issues National Standard for Reused ELV Parts. 2025;
5. Xia Yu. MOFCOM's New Notice on the Recycling of End-of-Life Motor Vehicles. 2024;
6. Paulo-Guedes-Pinto. Why Remanufacturing Electric Motors is a Big Idea. 2025;
7. Katona M, Orosz T. Circular Economy Aspects of Permanent Magnet Synchronous Reluctance Machine Design for Electric Vehicle Applications: A Review. *Energies* 17 2024.
8. Philip Akrofi Atitianti. 'Greening' Transportation: Electric Vehicles and Ghana's sustainable mobility push. 2025;
9. Botchway N, Calloway A, Aboagye Da Costa A, Lutterodt LL. Energy Laws and Regulations 2025 – Ghana. 2024;
10. Girbacia F, Daniel Voinea G. Multi-Method Statistical Analysis of Factors Influencing Predictive Maintenance of Electric Vehicle Fleets. 2025.
11. Zhang Y, Fang L, Qi Z, Deng H. A Review of Remaining Useful Life Prediction Approaches for Mechanical Equipment. *IEEE Sens J* 2023; 23: 29991–30006.
12. Sayyad S, Kumar S, Bongale A, Kamat P, Patil S, Kotecha K. Data-Driven Remaining Useful Life Estimation for Milling Process: Sensors, Algorithms, Datasets, and Future Directions. *IEEE Access* 2021; 9: 110255–110286.
13. Sun J. Open Aircraft Performance Modeling Based on an Analysis of Aircraft Surveillance Data. TU Delft University 2019; 289.
14. Sekhar JNC, Domathoti B, Santibanez Gonzalez EDR. Prediction of Battery Remaining Useful Life Using Machine Learning Algorithms. *Sustainability (Switzerland)* 2023; 15.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.