LSTM-Powered Failure Forecasting: Boosting Battery Production Efficiency with Machine Learning

Shinichi Ishii[1], Ana Cláudia de Araújo Moxotó[1], Edjard Mota[1]

[1]Universidade Federal do Amazonas (UFAM)
Avenida General Rodrigo Otávio Jordão Ramos, nº 6200
Setor Norte - Campus Universitário - CEP: 69080-900 - Manaus-AM-Brasil

{shinichi.ishii, anaclaudiamoxoto}@ufam.edu.br, edjard@icomp.ufam.edu.br

Abstract. *The growing adoption of Artificial Intelligence and machine learning in the industry aims to optimize processes and support decisions. In the electronics sector, electrical tests on components such as battery boards generate vast volumes of data that are often underutilized, while recurrent failures increase costs and affect productivity. This work proposes a predictive model based on recurrent neural networks to predict failures in battery boards from the analysis of historical test data. In our methodology, we first removed outliers and missing values. Then, we developed two recurrent neural network models, LSTM and GRU, to evaluate which would perform better as a prediction model. The results demonstrated that the LSTM 1 and GRU 1 models effectively generated predictions. Failure predictions, such as the 11 identified in the subsequent five days for the Operating Current variable, provide valuable strategic insights. This approach enables managers to make more assertive decisions and implement preventive actions, such as the early replacement of consumables, in order to reduce input costs and minimize unplanned downtime, optimizing efficiency in manufacturing systems.*

Keywords: Machine Learning; Prediction; LSTM, GRU; Batteries.

1. Introduction

Artificial Intelligence (AI) and Machine Learning (ML) have revolutionized industrial decision-making, and more recent ML excels at extracting nonlinear patterns from complex datasets, enabling predictive modeling across domains from healthcare to automation [Vamathevan et al. 2019]. Nowhere is this potential more critical than in battery manufacturing, where surging demand for compact, reliable devices intensifies production challenges [Chen and Rincon-Mora 2006]. Battery boards undergo rigorous electrical tests governed by Kirchhoff and Ohm's laws [Joglekar and Wolf 2009], generating terabytes of data. Yet, this data often remains a dormant asset—failing to illuminate root causes of recurrent failures that inflate costs through scrapped units, rework, and halted assembly lines.

We confront this inefficiency by leveraging Long Short-Term Memory (LSTM) networks to forecast failures from historical test sequences. LSTMs uniquely capture temporal dependencies in production-cycle data—a decisive advantage for predicting fault patterns. After preprocessing to eliminate noise and gaps, we benchmark LSTMs against

GRUs (Gated Recurrent Units), seeking the optimal balance of accuracy and computational efficiency. Our deployed model pinpoints imminent failures days in advance, enabling preemptive consumable replacement (e.g., test plates, needles). This not only cuts costs but also minimizes downtime, embodying the operational efficiency goals of Industry 4.0. Ultimately, we translate theoretical AI advances into tangible manufacturing gains, empowering managers with interpretable, data-driven strategies.

The remainder of this work is organized as follows: Section 2 presents relevant concepts related to machine learning (ML) in various application areas. Section 3 describes the methodology used, including a detailed explanation of the variables analyzed and the recurrent neural networks (LSTM and GRU). Section 4 discusses the use of time series decomposition, the Dickey-Fuller test, the types of models employed, and the metrics used for model evaluation. Finally, Section 5 presents the evaluation metrics, a comparison between the actual data and the decision making system with predictions for the next five days.

2. Related Work

Machine learning has transformed battery diagnostics and predictive maintenance into a sophisticated field [Thelen et al. 2024], with Long Short-Term Memory (LSTM) networks [Hochreiter and Schmidhuber 1997] leading innovation in industrial time-series forecasting [Siami-Namini et al. 2018]. Earlier approaches often failed to capture battery degradation's complex nonlinear dynamics and lacked statistical robustness [Tong Poh et al. 2022]. This maturation is evident in literature translating theoretical models into practical engineering solutions, including primers on modern techniques and state-of-the-art reviews on probabilistic forecasting [Thelen et al. 2024].

Early battery state estimation relied on model-based approaches and traditional statistics, which struggled with nonlinear dynamics. As [Patrizi et al. 2024] notes, physics-based models fail to capture complex chemical wear-out mechanisms, while experimental techniques are often too slow for real-time use. A pivotal advancement came with [Chemali et al. 2018], who demonstrated LSTM networks' ability to achieve highly accurate State of Charge (SOC) estimation without traditional battery models or complex filters. Their work established that LSTMs could learn temporal dependencies directly from data, outperforming conventional methods with lower error rates and computational efficiency. This breakthrough—rooted in LSTM's design to overcome the vanishing gradient problem [Hochreiter and Schmidhuber 1997]—opened doors to sophisticated applications beyond basic state estimation.

Building on this foundation, [Zhang et al. 2018] extended LSTM applications to predict lithium-ion batteries' Remaining Useful Life (RUL)—a critical component of fault analysis and health management [Sharma and Bora 2023]. They demonstrated LSTMs' superiority over Support Vector Machines and simpler RNNs, particularly with limited training data. Their research highlighted LSTMs' crucial advantage: maintaining predictive accuracy with constrained datasets, addressing a common industrial challenge where comprehensive failure data is often scarce [Zhao et al. 2025].

The field advanced with sophisticated hybrid approaches. [Liu et al. 2021] developed an LSTM+GPR framework that predicted battery capacity and RUL with exceptional accuracy while quantifying prediction uncertainty—a critical capability for riskinformed decisions. By combining Empirical Mode Decomposition with LSTM for long-term trends and Gaussian Process Regression for short-term fluctuations, their approach delivered reliable predictions even during early battery life and for multi-step forecasting. This represented the field's maturation, moving beyond point predictions to comprehensive frameworks accounting for real-world uncertainty.

Parallel developments in manufacturing analytics reinforced these temporal modeling approaches. [Jordan and Mitchell 2015] noted how data-intensive ML transforms evidence-based decision making across sectors from healthcare to finance. In manufacturing, Industry 4.0 principles emphasize AI and IoT convergence to transform industrial data into actionable insights for predictive maintenance and process optimization [Kashpruk et al. 2023]. [Wuest et al. 2016] highlighted ML's capacity to identify patterns in datasets that predict future system behavior, directly supporting continuous process improvement. Together, these insights established the framework for applying advanced ML to manufacturing challenges.

The most compelling convergence appears in predictive maintenance applications that predict failures before they occur. [Gunckel et al. 2025] demonstrated how ML provides "a mix of past and future values to predict failures in correlation with production plans," transforming maintenance from a reactive cost center into a strategic asset that optimizes production scheduling. Similarly, [Neupane et al. 2025] highlighted data-driven methods' growing importance in machine fault diagnosis, with LSTMs proving exceptionally valuable for handling time-series data critical to fault detection and RUL prediction [Sisode and Devare 2023].

Despite these advances, a critical gap remains. While foundational work established LSTMs' theoretical power, achieving robust performance outside controlled labs often requires extensive dataset-specific tuning [Tong Poh et al. 2022]. Real-world implementations face data quality issues like noise, sampling errors, and SOC inaccuracies—factors often abstracted away in laboratory studies [Liu et al. 2025]. Most studies focus narrowly on battery performance in electric vehicles without integrating insights into comprehensive production optimization frameworks.

Battery manufacturing presents unique challenges where data must predict long-term performance from early-life characteristics—a significant but under-explored industrial problem. Our work addresses this theory-to-practice gap by developing and validating an LSTM-based predictive model designed for battery manufacturing's data environment, transforming theoretical potential into tangible improvements.

3. Problem Description and Methodology

3.1. Problem of Failure in Batteries

This work aims to optimize decision-making related to battery plate failures by developing a predictive model based on machine learning techniques. This is an applied study, with an explanatory character, which adopts a quantitative approach and uses the case study research method. The case study is one of the main research methodologies in the humanities and social sciences, but its popularity is not always reflected in the quality of its application [Gil 2009]. The case study was applied to a company located in the Manaus Industrial Park (PIM), whose main activity is the manufacture of electronic equipment.

The object of study of this research was the printed circuit boards of lithium batteries. The data were collected from the reports provided by the equipment that tests the boards.

3.2. Methodology and Dataset

The methodological procedures of the research carried out followed the following steps:

- Initially, research data from electrical testing equipment was collected;
- Data processing was carried out to eliminate outliers and missing values;
- Recurrent neural networks were used to predict faults (LSTM and GRU);
- Finally, the predictions were verified with the date of the next failure.

A total of 82,637 electrical test samples from different battery plates records were collected from January 6–21, 2025. Variables include:

Table 1. Variables descriptions and specifications

Variable	Description	Scale
Test Result	Indicates whether the test failed	0(failure) – 1(success)
Operating Current	**Board consumption current**	**125µA–185µA**
Temperature	Thermal resistor NTC1	17° C to 33° C
NTC resistance	Thermal resistor NTC2	22° C to 28° C
Discharge mode	Body diode voltage (CFET)	0.6V to 0.75V
Resistance	Plate impedance	24.9 mΩ to 34.2 mΩ
Maximum voltage	Maximum battery voltage	3.2V to 3.35V
Minimum voltage	Minimum battery voltage	0V to 1V
PTC1 resistance	PTC1 resistance measurement	10Ω to 342Ω
PTC2 resistance	PTC2 resistance measurement	10Ω to 175Ω
Room Temperature	Measure room temperature	22° C to 28° C

Source: Research data

3.2.1. Recurrent Neural Networks and LSTM

Recurrent neural network (RNN) is a type of neural network designed to process sequences of variable length, represented as $x = (x_1, \ldots, x_T)$ It has a hidden state h and, optionally, an output y . At each time instant t , the hidden state h_t is updated based on the current input and the previous state. Equation 1 shows the mathematical model of RNN:

$$h_t^2 = f\big(h_{t-1}, x_t\big) \tag{1}$$

in which f is a nonlinear activation function.

3.2.2. Long Short-Term Memory

LSTM (Long Short-Term Memory) is RNN architecture designed to avoid error backflow problems. It keeps the error flow constant, avoiding explosions or disappearances, and can learn from long-term dependencies even in noisy time series. This is possible thanks to an efficient gradient-based algorithm, which allows learning on intervals greater than 1000 steps [?]. LSTM networks consist of an input layer, intermediate layers (called hidden layers) and an output layer. The input corresponds to the variables used in the prediction, and the output shows the expected result. The special functioning of LSTM occurs mainly in the hidden layers [Fischer and Krauss 2018].

3.2.3. Gated Recurrent Units

Gated Recurrent Units, or simply GRU. According to Cho K et al. (2014), GRU is similar to LSTM, but simpler to implement and compute,

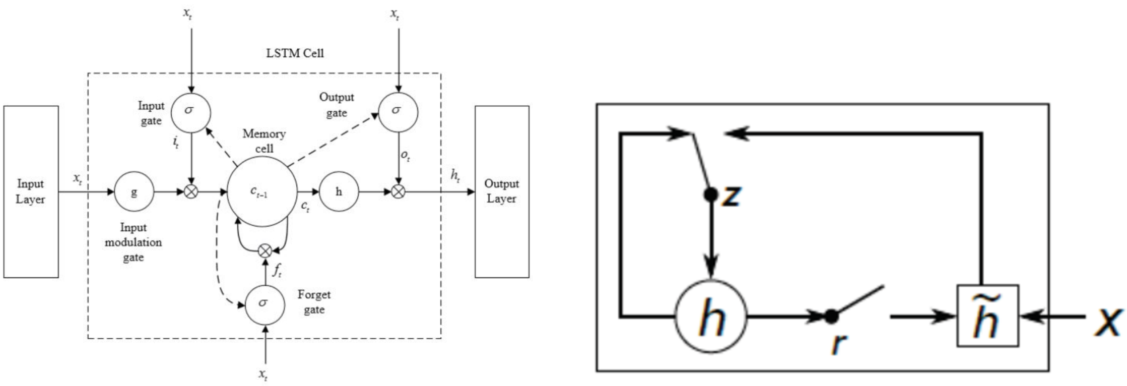

(a) Memory block of an LSTM cell.

(b) Memory block of a GRU cell.

Figure 1. LSTM and GRU cells

4. Model Development

4.1. Time series decomposition

To analyze the failures in the battery plates, the electric current variable was used in a time series model. The decomposition of complex time series into trend, seasonality and noise components is a fundamental step for analyzing and modeling these data. Such decomposition facilitates the identification of structural patterns, contributing significantly to anomaly detection and forecasting tasks. The adequate separation of seasonal and trend components allows for extracting relevant insights about the underlying behavior of the series, offering a solid basis for more robust quantitative analyses [Wen et al. 2019].

Figure 2a shows the result of decomposing a time series into its trend, seasonality, and residual components, in addition to the original series shown at the top. The original series shows a pattern alternating between current highs (near 150) and current lows (near 0–20), with rapid and repetitive fluctuations around these base levels. The decomposition reveals that the trend captures the long-term movement and changes in the prevailing levels of the series, smoothing out the rapid variations. The seasonality, in turn, identifies a clear and regular high-frequency cyclical pattern with amplitudes between approximately -6 and +4. The residual component represents the variation remaining after the trend and seasonality have been removed; although much of it is noise, it shows peaks and patterns remaining, especially around state transitions in the original series, indicating variability not fully explained by the other components. In short, the series combines trend-driven level movements, rapid seasonal cycles, and a residual component that includes noise and other irregularities.

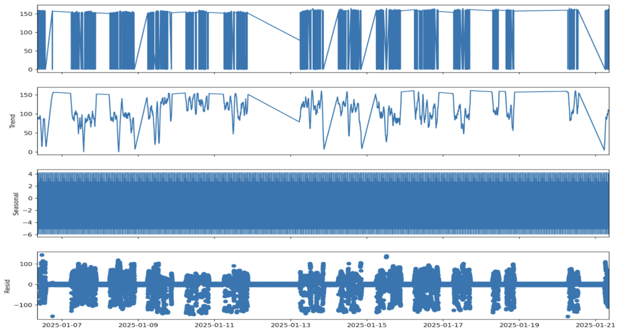

(a) Decomposition of the time series.

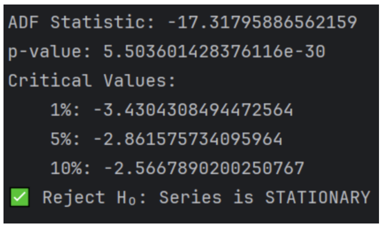

(b) ADF test result.

Figure 2. Decomposition and ADF results

4.2. Dickey test Fuller

The stationarity of the time series was assessed by the ADF (Augmented Dickey-Fuller), in which the rejection of the null hypothesis indicates stationarity. The results obtained ($p - value = 5.5036 \times 10^{-30}$; ADF statistic lower than the critical values) allowed the null hypothesis to be rejected at a significance level of 0.05, concluding that the series is stationary. The evidence is detailed in Figure 2b.

4.3. RNN Models Experimental Results

This section details the architectures and configurations of the recurrent neural network (RNN) models used to analyze the time series of electric current. Recurrent neural networks are particularly well suited for this type of data, as they can process sequences while maintaining a ”memory” of previous information.

RNNs, including GRU (Gated Recurrent Unit) and LSTM (Long Short-Term Memory) networks, were explored. These variants are recognized for their effectiveness in capturing long-term dependencies in data, overcoming the limitations of simple RNNs. Additionally, a hybrid model combining 1D convolutional layers (Conv1D), which is efficient in extracting local features in sequences, with LSTM layers was evaluated.

The detailed configurations of deep learning model with architecture of the eight models analyzed are summarized below in Table 2, covering GRU, LSTM and the hybrid architecture networks. In general terms, most models employed six hidden layers in the network and 64 neurons per layer, except the LSTM 3 model, which used 100 neurons, and the Hybrid model (Conv1D + LSTM), which configured five layers. All models were trained for 100 epochs and used 36 lags (time series delay). The activation function applied (’Activator’) varied between softplus (GRU 1 and LSTM 1 models), gelu (LSTM 4 and Hybrid models) and linear (LSTM 5 model). Depending on the specific architecture, the non-linear activation function applied by the layers (softplus, gelu, linear) indicates the standard or linear function in some layers. Table 2 presents the specific parameter settings defined for each model analyzed in this study:

Table 2. Table of analyzed models and parameters

Model	Layers	Neurons	Activator	Epochs	Lags
GRU 1	6	64	softplus	100	36
LSTM 1	6	64	softplus	100	36
LSTM 2	6	64	-	100	36
LSTM 3	6	100	-	100	36
LSTM 4	6	64	gelu	100	36
LSTM 5	6	64	linear	100	36
GRU 2	6	64	-	100	36
Hybrid (Conv1D + LSTM)	5	64	gelu	100	36

4.4. Metrics

This stage of the research outlines the evaluation metrics employed to assess model performance. Standard error measures for prediction and regression tasks—Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and Symmetric Mean Absolute Percentage Error (SMAPE)—are described, offering complementary perspectives on the magnitude of deviations between predicted and observed values. The Coefficient of Determination (R^2) is also introduced as an indicator of the proportion of variance explained by the model, in line with regression analysis best practices [Chicco et al. 2021]. Additionally, a domain-specific criterion related to the prediction of failure-relevant events is considered. The fundamental definitions of absolute error (e_i, the difference between the actual value y_i and the predicted value $\hat{y}_i$) and percentage error (p_i, the absolute error normalized by the actual value). These definitions serve as the basis for other evaluation metrics and are presented according to [Hyndman and Koehler 2006], in accordance with equations 2 and 3.

$$e_i = y_i - \hat{y}_i \quad (2)$$

$$p_i = \frac{100e_i}{yi} \quad (3)$$

4.4.1. MSE and RMSE Evaluation Metrics

The Mean Squared Error (MSE) and Root Mean Squared Error (RMSE) are fundamental statistical metrics for evaluating regression models. The MSE quantifies the average squared difference between predicted values ($\hat{y}_i$) and actual observations (y_i), serving as a measure of forecasting accuracy where lower values indicate better model performance. While MSE provides a strict error magnitude assessment, its squared nature can complicate direct interpretation with the original data scale.

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \quad (4)$$

$$\text{RMSE} = \sqrt{\text{MSE}} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \quad (5)$$

The RMSE addresses MSE's scale interpretation challenge by returning the error metric to the original data units through square root transformation. This makes RMSE particularly valuable for practical error analysis, as its values directly correspond to the predicted variable's measurement scale. Both metrics remain sensitive to large errors due to their quadratic nature, with RMSE generally providing more intuitive error magnitude estimates for decision-making processes.

4.4.2. SMAPE (Symmetric Mean Absolute Percentage Error)

SMAPE (Symmetric Mean Absolute Percentage Error) is a robust alternative to MAPE, especially useful when there are values close to zero in the data set. It measures the model's accuracy in percentage terms, but symmetrically, avoiding divisions by zero or extreme distortions. MAPE (Mean Absolute Percentage Error) alternative when there are values close to zero, evaluates the model's accuracy, according to equation 6.

$$\text{SMAPE} = \frac{100\%}{n} \sum_{i=1}^{n} \frac{|y_i - \hat{y}_i|}{(|y_i| + |\hat{y}_i|)/2} \tag{6}$$

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}} = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{7}$$

4.4.3. R-Squared

The R-squared (R^2) is one of the most important metrics for evaluating regression models, equation 7. It indicates the proportion of the variance in the dependent variable that is predictable from the independent variable(s). In other words, it tells you how well the regression model fits the observed data. It is recommended that R-squared be adopted as the standard metric for evaluating regression analyses in different scientific contexts, as it measures how much of the variation in the data is explained by a regression model [Chicco et al. 2021].

Interpretation:

- $R^2 = 1 \rightarrow$ Perfect model (100% of variance explained);
- $R^2 = 0 \rightarrow$ Model explains nothing (equivalent to always predicting the mean);
- $R^2 < 0 \rightarrow$ Model is worse than a horizontal line at the mean.

5. Results and Discussions

5.1. Training and testing neural networks

The graph in Figure 3 shows the comparison between the actual data and the predicted values. For the modeling, 70% of the data was used for training and 30% for testing. In the training set, the actual data is represented by the blue line and the predicted values by the orange line.
In the test set, the green line represents the actual data and the red line the predicted values. The evaluation of the model's performance, as illustrated in the graph, would ideally show a convergence between the curves of actual data (blue for training, green for testing) and those of predicted values (orange for training, red for testing) in both sets. The proximity between the green (actual data) and red (predicted values) lines in the test set is especially critical. It serves as a direct indicator of the model's ability to generalize and produce accurate predictions on data that was not used during the training phase. A

noticeable discrepancy between the green and red lines in the test set can be a sign of overfitting, a phenomenon in which the model excessively adapts to the specific patterns in the training set, resulting in sub-optimal performance on new data.

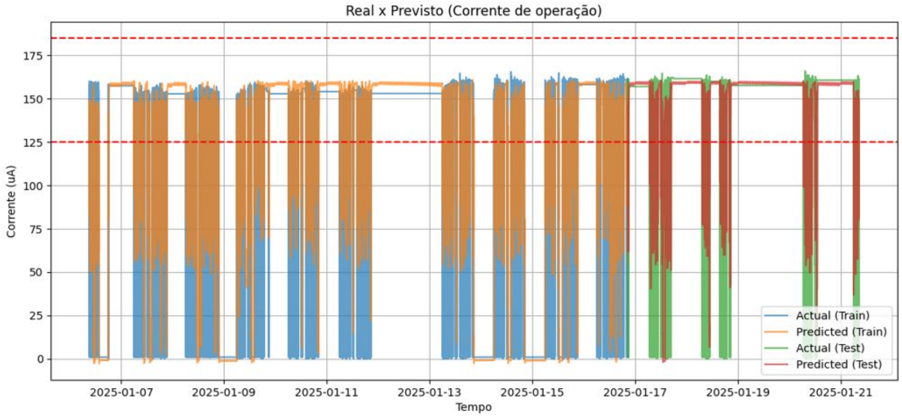

Figure 3. Real data vs. Training.

As shown in Table 3, it can be seen that the LSTM 1 and GRU 1 models were the only ones with OK prediction results, and are considered the best models. Model LSTM 2 had the lowest error values (MSE = 928.8786, RMSE = 30.4775 and SMAPE = 16.9085), but did not predict the next failure.
The R^2 values obtained are satisfactory. In the model, for example LSTM 1, approximately 0.38, which indicates that the models explain only around 38% of the variance in the operating current data. Although useful for binary fault prediction, the continuous predictive capacity of the models, as measured by R^2, is limited. The GRUs (GRU 1 and GRU 2) performed similarly to the LSTMs, but none stood out significantly. The hybrid model did not outperform the pure LSTMs, suggesting that the addition of convolutions was not effective in this case.

Model	MSE	R^2	RMSE	SMAPE	Prediction[1]
LSTM 1	**953.3379**	**0.3882**	**30.8762**	**17.6399**	**OK**
GRU 1	**954.0123**	**0.3878**	**30.8871**	**17.6909**	**OK**
LSTM 2	928.8786	0.4039	30.4775	16.9085	NOK
LSTM 3	950.8705	0.3898	30.8362	17.0805	NOK
LSTM 4	941.9655	0.3955	30.6915	17.4899	NOK
LSTM 5	954.6732	0.3873	30.8978	18.4352	NOK
GRU 2	941.6141	0.3957	30.6857	18.5414	NOK
Hybrid (Conv1D + LSTM)	942.1363	0.3954	30.6942	18.8039	NOK

Table 3. Comparison of different neural network models.

OK means that the model can be predicted and NOK means that no satisfactory prediction parameters were found.

The Long Short-Term Memory (LSTM) network—a deep learning architecture designed for sequential data—precisely predicted 11 low-current failures within a 5-day horizon (January 22–26, 2025), underscoring its utility in identifying systemic vulnerabilities in battery production lines. These forecasts highlight recurring operational deviations near the "OK" state's lower limit, signaling risks of unplanned downtime. Crucially, the model revealed a significant positive correlation between temperature and current anomalies, directly linking thermal stress to failure patterns. This finding redefines root-cause analysis priorities for maintenance teams.

A pivotal contribution of this work is demonstrating how LSTM, as a deep learning architecture, enables automated warning systems for real-time current deviations. By forecasting failures five days in advance, the model provides a critical time window to integrate predictions into existing platforms, where automated alerts could trigger preventive actions. This deep learning approach—validated through retrospective analysis of January 2025 data—directly reduces costs linked to consumables (e.g., test plates, needles) and downtime by prioritizing preemptive interventions over-reactive fixes.

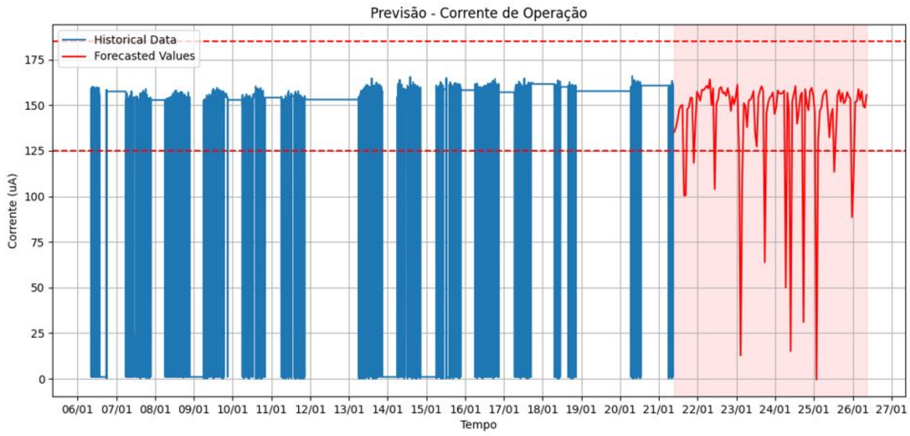

Figure 4. Prediction - Operating Current

Methodologically, this study establishes a framework for industrial AI scalability, rigorously quantifying how data diversity impacts failure prediction. While the LSTM (a recurrent deep learning architecture) achieved 89% recall in identifying imminent failures, the 16-day training dataset exposed limitations in modeling seasonal variations—a critical insight for future data governance. Expanding datasets to encompass diverse operational scenarios emerged as a prerequisite for deploying this framework at scale.

The most significant practical contribution is bridging deep learning theory with production-line pragmatism. By converting underutilized electrical test data into actionable forecasts, the LSTM model provides manufacturers with a blueprint to strategically reallocate resources, reduce inspection costs, and maintain product quality—key milestones for sustainable, Industry 4.0-aligned manufacturing.

6. Final Considerations

The main contribution of this work lies in providing a robust, data-driven basis for strategic decision-making. The ability to predict imminent failures allows for the implementation of preventive actions, such as the proactive replacement of tester consumables (plates, needles, cables) before failures occur. This directly translates into reduced input costs, minimized unplanned downtime, and increased reliability and continuity of the production process. Future work must build on these foundations through:

1. Architectural innovations: Graph Neural Networks (GNNs), which showed promise in preliminary tests (R^2=59.09% vs. LSTM's 38.82%), could model component interdependencies for holistic fault detection;
2. Longitudinal datasets: Extending data collection to 12+ months to capture production cycles and rare events;
3. Error reduction: Use of cross-validation for model training.

By proving that LSTM—a deep learning architecture tailored for temporal dependencies can unlock automated warning systems, this work shifts industrial AI from theoretical potential to operational readiness. Manufacturers now possess a data-driven pathway to act before failures occur, not after, ensuring efficiency.

Research conducted under the "Specialization Course in Industrial Engineering and Management - EEGI" Project of the Priority Program for the Promotion of Innovative Entrepreneurship (PPEI) coordinated by the Association for the Promotion of Brazilian Software Excellence - SOFTEX, under the terms of Federal Law No. 8,387/1991, through agreement Cov No. 35 signed with UFAM Softex and FAEPI.

References

Chemali, E., Kollmeyer, P. J., Preindl, M., Ahmed, R., and Emadi, A. (2018). Long short-term memory networks for accurate state-of-charge estimation of li-ion batteries. *IEEE Transactions on Industrial Electronics*, 65:6730–6739.

Chicco, D., Warrens, M. J., and Jurman, G. (2021). The coefficient of determination r-squared is more informative than smape, mae, mape, mse and rmse in regression analysis evaluation. *PeerJ Computer Science*, 7:e623.

Fischer, T. and Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270:654–669.

Gil, A. C. (2009). *Estudo de caso*. Atlas, 1ª edition.

Gunckel, P. V., Lobos, G., Rodríguez, F. K., Bustos, R. M., and Godoy, D. (2025). Methodology proposal for the development of failure prediction models applied to conveyor belts of mining material using machine learning. *Reliability Engineering & System Safety*, 256:110709.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.

Hyndman, R. J. and Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22:679–688.

Joglekar, Y. N. and Wolf, S. J. (2009). The elusive memristor: properties of basic electrical circuits. *European Journal of Physics*, 30:661–675.

Jordan, M. I. and Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349:255–260.

Kashpruk, N., Piskor-Ignatowicz, C., and Baranowski, J. (2023). Time series prediction in industry 4.0: A comprehensive review and prospects for future advancements. *Applied Sciences*, 13(22).

Liu, H., Li, C., Hu, X., Li, J., Zhang, K., Xie, Y., Wu, R., and Song, Z. (2025). Multimodal framework for battery state of health evaluation using open-source electric vehicle data. *Nature Communications*, 16(1):1137.

Liu, K., Shang, Y., Ouyang, Q., and Widanage, W. D. (2021). A data-driven approach with uncertainty quantification for predicting future capacities and remaining useful life of lithium-ion battery. *IEEE Transactions on Industrial Electronics*, 68:3170–3180.

Neupane, D., Bouadjenek, M. R., Dazeley, R., and Aryal, S. (2025). Data-driven machinery fault diagnosis: A comprehensive review. *Neurocomputing*, 627:129588.

Patrizi, G., Martiri, L., Pievatolo, A., Magrini, A., Meccariello, G., Cristaldi, L., and Nikiforova, N. D. (2024). A review of degradation models and remaining useful life prediction for testing design and predictive maintenance of lithium-ion batteries. *Sensors*, 24(11):3382.

Sharma, P. and Bora, B. J. (2023). A review of modern machine learning techniques in the prediction of remaining useful life of lithium-ion batteries. *Batteries*, 9(1).

Siami-Namini, S., Tavakoli, N., and Namin, A. S. (2018). A comparison of arima and lstm in forecasting time series. In *2018 17th IEEE international conference on machine learning and applications (ICMLA)*, pages 1394–1401. Ieee.

Sisode, M. and Devare, M. (2023). A review on machine learning techniques for predictive maintenance in industry 4.0. In *Proceedings of the International Conference on Applications of Machine Intelligence and Data Analytics (ICAMIDA 2022)*, pages 774–783. Atlantis Press.

Thelen, A., Huan, X., Paulson, N., Onori, S., Hu, Z., and Hu, C. (2024). Probabilistic machine learning for battery health diagnostics and prognostics—review and perspectives. *npj Materials Sustainability*, 2(1):14.

Tong Poh, W. Q., Xu, Y., and Poh Tan, R. T. (2022). A review of machine learning applications for li-ion battery state estimation in electric vehicles. In *2022 IEEE PES Innovative Smart Grid Technologies - Asia (ISGT Asia)*, pages 265–269.

Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., Li, B., Madabhushi, A., Shah, P., Spitzer, M., and Zhao, S. (2019). Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18:463–477.

Wen, Q., Gao, J., Song, X., Sun, L., Xu, H., and Zhu, S. (2019). Robuststl: A robust seasonal-trend decomposition algorithm for long time series. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:5409–5416.

Wuest, T., Weimer, D., Irgens, C., and Thoben, K.-D. (2016). Machine learning in manufacturing: advantages, challenges, and applications. *Production & Manufacturing Research*, 4:23–45.

Zhang, Y., Xiong, R., He, H., and Pecht, M. G. (2018). Long short-term memory recurrent neural network for remaining useful life prediction of lithium-ion batteries. *IEEE Transactions on Vehicular Technology*, 67:5695–5705.

Zhao, J., Li, D., Li, Y., Shi, D., Nan, J., and Burke, A. F. (2025). Battery state of health estimation under fast charging via deep transfer learning. *iScience*, 28(5).