International Journal of Basic and Applied Sciences, 14 (8) (2025) 284-301

International Journal of Basic and Applied Sciences

Website: www.sciencepubco.com/index.php/IJBAS
https://doi.org/10.14419/16vk8033
Research paper

A Hybrid Deep Learning and XAI-Driven Framework for Accurate Estimation of Battery SOC AND SOH

N Sumana Keerthi [1], B. Jyothi [1], Kalyan D.[1], K. V. Govardhan Rao [2]*, Teerdala Rakesh [2], M. Kiran Kumar [1]

[1] Department of EEE, Koneru Lakshmaiah Education Foundation, Green Fields Vaddeswaram, India
[2] Department of EEE, St. Martin's Engineering College, Secunderabad, Telangana, India
**Corresponding author E-mail: kv.govardhanrao@gmail.com*

Received: November 11, 2025, Accepted: December 9, 2025, Published: December 16, 2025

Abstract

Understanding battery State of Health (SoH) and State of Charge (SoC) predictions highly ensure reliability, safety, and longevity in energy storage systems, especially for electric vehicles and smart grids. This novel hybrid-type framework enables the successive application of machine learning and deep learning models in SOC and SOH estimations. A one-of-a-kind blend of Linear Regression, XG-Boost, Recurrent Neural Networks (RNN), Bi-directional LSTM, and hybrid LSTM-GRU is suggested to capture as many temporal and non-linear patterns in battery behavior as possible. K-Best feature selection is used to enhance model generalization by keeping only the most important input features. Contrary to the existing unexplainable models, our approach leverages explainable AI methods-SHAP and LIME-to explain model decisions and magnitude of feature impact. Real-world battery datasets weighed experimentally underline the superiority of our approach from both accuracy and interpretability perspectives over traditional means. This paper's novelty lies in a hybrid modeling architecture, an interpretable learning pipeline, and a consideration of both predictive ability and interpretability. This framework has enormous potential to draw the innovations forward in intelligent battery management systems.

Keywords: *Hybrid Deep Learning; Recurrent Neural Networks; Bi-directional LSTM; Hybrid LSTM-GRU; Machine Learning.*

1. Introduction

Battery systems are on their way to becoming the pinnacle of modern technological advances, especially for electric vehicles, rechargeable storage systems for renewable energies, and portable electronics and unmanned systems. The global shift toward clean energies and electrification has sold the idea that the reliability, safety, and efficiency of batteries are fast emerging as critical consideration points. Therefore, estimating the essential battery performance parameters, SOH and SOC, has become imperative to such regard. These parameters serve as important keys for aiming at battery condition monitoring, life forecasting, or energy optimization.

State of Charge is a term that describes the charge remaining in the battery in relation to maximum capacity, essentially quantifying how much energy one can draw. On the other hand, SOH describes the general condition of the battery in terms of its nominal capacity compared to a brand-new battery. If not properly calculated, the parameters may incur early failure, unsafe operation, or improper energy management. Consequently, these parameter estimations have been traditionally carried out by Coulomb counting, open-circuit voltage measurement, and the like, which unfortunately suffer from a set of limitations including noise in sensor readings, temperature dependence, accumulation of large errors, and application to specific operation cases only.

Machine learning (ML) and deep learning (DL) methods have lately become potent alternatives with an emphasis on data and their capacity to seize nonlinearities and complex dependencies in battery operation. Yet, in most previous works, the models are either specific or cannot be generalized, interpreted, nor adapted in real-time. This study tries to fill these gaps by proposing a hybrid, explainable, and feature-optimized framework for SOC and SOH estimation.

The most important contribution of this work involves simultaneous implementation of multiple models: Linear Regression, XG-Boost, Recurrent Neural Networks (RNN), Bidirectional LSTM (Bi-LSTM), and the hybrid LSTM-GRU architecture. Each method looks into certain temporal and non-temporal dependencies. Linear Regression and XG-Boost are good with direct and nonlinear effects, while RNNs, Bi-LSTMs, and GRUs try to understand the time-series sequence that is inherent to battery usage and degradation patterns.

The K-Best feature selection method to improve the quality of input data and avoid overfitting. This method filters out redundant and irrelevant features, retaining those of utmost consideration according to their statistical scoring criteria. The improvement of model accuracy and efficiency-based computations for training and inferencing would be other improvements contributed.

The methodological aspects are quintessentially selected through several vital steps: (1) data gathering and preprocessing from the set of real-world lithium-ion battery datasets; (2) selection of statistical and domain-related features through K-Best; (3) training-testing- validation of the candidates through multiple algorithms; (4) assessment of performances using the RMSE, MAE, and R^2 metrics; and (5) analysis

Copyright © N Sumana Keerthi et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

of interpretations through XAI methods. The splitting of training-testing datasets ensures robustness and generalizability with a focus on application cases reflecting real-world operating conditions, such as varying discharge and ambient temperature conditions or cycling patterns.

The direct operational value of these work interventions lies in sturdy support for battery management systems (BMS), as SOC and SOH estimation affects optimal charging/discharging scheduling and thus avoids overcharging or deep discharging and prevents scheduling for preventive maintenance. Also, because the model is explainable, it can be used for regulatory compliance and decision support in energy-critical systems.

This work is unusual in existing literature because it simultaneously addresses three main challenges: (1) it fights the integration of time-series modeling with static predictors; (2) it is more transparent due to an explainability tool; (3) features are optimized to facilitate the generalization procedure. There are numerous works that place emphasis on prediction accuracy only; in contrast, our framework balances the priority among interpretability and feature efficiency.

The remainder of the paper is organized as follows: Section II presents related works and existing techniques applied for SOC and SOH estimation. Section III presents the proposed system architecture and feature selection method. The modeling techniques and training procedures are described in Section IV. Section V presents the experimental results and comparative analysis. Section VI discusses real-world implications and limitations, allowing for further deployment in real-time scenarios. The work is concluded in Section VII with remarks on future enhancements.

In conclusion, accurate and interpretable estimation of SOC and SOH with the hybrid ML/DL model shows good potential for intelligent energy systems. Our research provides a solid base for practical deployments in BMS by merging precision, interpretability, and robustness within a single framework.

2. Related Work

Battery State of Charge (SOC) and State of Health (SOH) estimation need to be accurate to ensure dependable energy storage-system operations. First, customary methods-i.e., Coulomb counting and equivalent circuit modeling-fail due to error accumulation, temperature dependency, and lack of adaptability [1], [2]. Keeping that in mind, machine learning- and deep learning-based data-driven techniques have now emerged to take their place. This is because batteries do not behave in a linear manner in degradation, and they show dynamic behavior. Being linear, models such as Linear Regression are simple and interpretable, and yet fail to account for nonlinearities in real-world battery behavior [3]. Therefore, ensemble methods like Extreme Gradient Boosting, XG-Boost, have found more adoption due to its high accuracy and robustness. Reference [4] showed that XG-Boost models for SOH prediction had a near 20% decrease in RMSE compared to the classic techniques. In addition to this, XG-Boost offers feature importance, thus interpretable far better than any other black-box method.

Battery diagnosis represents an area where recurrent neural networks, mainly LSTM and GRU networks, shine with their long sequential data capabilities. Initially, LSTM networks were utilized in [5] for SOH prediction from cycle data by learning temporal relationships, achieving an average MAE of less than 1%. A hybrid LSTM-GRU method was likewise proposed in [6] which showed better convergence ability and stability in prediction than either individual method.

Combining both linear and backward temporal processing for learning yields better structures and solutions for data representation. Reference [7] utilized Bi-LSTM architecture to estimate SOH and concluded that it generalizes better under various discharge rates and ambient conditions.

The CNN-own track [8] is a novel hybrid framework of CNN and LSTM for simultaneous extraction of spatial and temporal characteristics, aiding in feature extraction and decreasing training time. Another line of studies like [9] works on classification features with raw-sensor data as input of 1D CNN before feeding the outputs to a sequential model in the first stage-and this has evidently boosted the performance.

Interpretation of the model becomes a necessity for the practical impost of Battery Management Systems (BMSs). Hence, we frequently see XAI tools, such as SHAP and LIME, integrated into battery prediction models. Reference [10] used SHAP with a tree-based model to demonstrate how temperature and cycle count had stronger degradation effects on SOH. In a similar way, [11] used LIME to provide local surrogate explanations to support model debugging and trust.

Feature selection remains an important step in ML pipelines. In [12], the authors used the K- Best algorithm to exclude irrelevant features, potentially improving training speed, as well as generalization. In fact, this was useful in eliminating redundancy within high-dimensional datasets of batteries. Recursive feature elimination, or RFE, has also been put to use, although it is considered computationally expensive [13].

Other significant attempts include other forms of metaheuristic optimization procedures, for instance, PSO [14], wherein PSO is used to optimize the parameter settings of the LSTM prediction model for SOH. The hybrid PSO-LSTM model produces a lower prediction error than traditional techniques such as grid search tuning.

Finally, in [15], a review of ML models for battery diagnostics finds that hybrid architectures and transfer learning are increasingly receiving attention; however, they stress the need for datasets and common evaluation protocols that can be used to benchmark model performances for different chemistries and use cases.

The literature review reveals that hybrid DL models (LSTM-GRU, CNN-LSTM), ensemble ML methods (XG-Boost), and interpreting tools (SHAP, LIME) have yielded good results in predicting SOC and SOH. However, for example, developing a universal model that works under a variety of conditions still remains a challenge. Building on these insights, the present work proposes a novel combination of these models with interpretable feature selection and explanation frameworks.

3. System Design and Architecture

It gives a platform to the Li-ion chemical battery-based data-centric modular architecture of SOC and SOH prediction systems, whose design priorities are accuracy, interpretability, and real-time execution [16]. The system framework consists of data preprocessing, feature selection, model training, performance evaluation, integration of explanations, and design modules.

Starting at the ingestion layer, the system collects battery datasets by experimentally recording voltage, current, temperature, and cycle information from the batteries. These datasets may be subjected to noise, missing values, and redundant entries; these are removed by the data preprocessing module [17]. Besides, normalization using Min-Max or Z-score, along with outlier removal or categorical encoding, can be taken into consideration for data preprocessing to attain standardization and stability across the dataset [18].

Post preprocessing and feature selection takes the preprocessed data and applies the K-Best method to extract inputs bearing the statistically highest relevance with respect to the target variables- SOC and SOH [19]. So, feature selection decreases the dimensionality of the data, helps to avoid overfitting [20], and leads to faster convergence in training of the models.
At the core, the layer consists of an ensemble of prediction models comprising linear regression, XG-Boost, RNN, Bi-LSTM, and hybrid LSTM-GRU. Each one of them is trained independently, with the dataset being split into time-series and statistical data through training; hence, they learn temporally and non-temporally. Hyper-parameter tuning follows a cross-validation scheme to locate the best set of parameters for each algorithm
For transparency, the explainability layer uses SHAP and LIME to interpret predictions made by the models. They give both a visual and a numerical interpretation of feature importance, helping an end-user understand which variables impact information-gathering predictions for SOC and SOH the most.

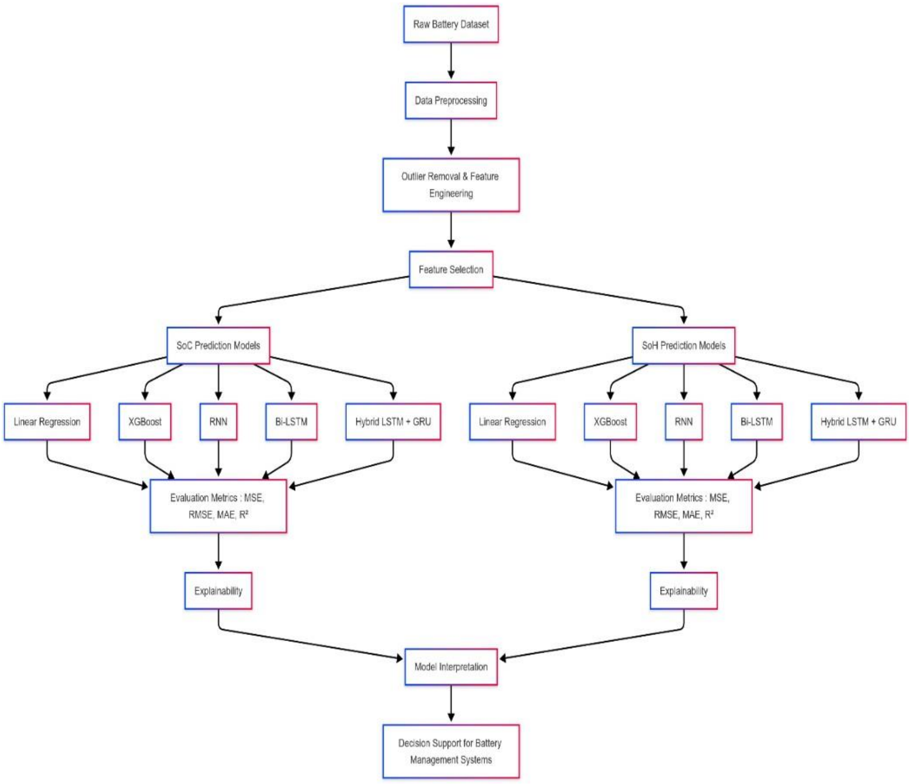

Fig. 1: Architecture.

Therefore, the architecture shown in Figure1, is a scalable, explainable, and highly accurate framework that can be integrated easily into BMSs, thus enabling proactive monitoring and intelligent decision- making.

4. Data set

In the utmost general sense, this study is concerned with estimating the SoC and SoH with precision in a real-life battery dataset with 29,180 rows and 7 columns really is considered. In the arrangement of rows and columns, the dataset was read into a Pandas Data-Frame, after which the data was processed and analyzed. Given this, time in rows describes observations at a certain instant of the battery during charging or discharging.
Parameters in consideration in the Data are:

- Voltage_Measured: Voltage across the battery at the measured point, in volts.
- Current_Measured: Current through the battery at the measured point, in amperes.
- Temperature_Measured: Temperature of the working battery, in degree Celsius.
- SoC: Actual state of charge, which tells us how much capacity is left in the battery, in percentage.
- cycle_number: Number of charge-discharge cycles that the battery has undergone.
- battery_id: Unique identifier labeled to each battery unit.
- SoH: State of Health, indicating how far battery condition has degraded as compared to the original capacity.

There is no missing value across all columns making the dataset complete, so that it can be well employed for model training and evaluation. The cycle_number and battery_id offer context identifiers to trace performance across time and units, while Voltage, Current, and Temperature stand as input features to determine SoC and SoH.
Such a dataset allows for the formation of time series as well as cross-section data, thus perfect for the training of traditional regression models and advanced neural networks, such as RNN and LSTM. A rich multidimensional view of battery behavior under changing conditions is essential for accurate and understandable prediction.

4.1. Data set analysis

EDA is a synthesizing view that describes various characteristics of the dataset with a view to obtaining useful patterns for modeling battery performance. It contains 29,180 records-all seven features intact-and is an ideal forerunner to predictive modeling. Variables are voltage measured, current, temperature, SoC, SoH, cycle number, and battery ID.

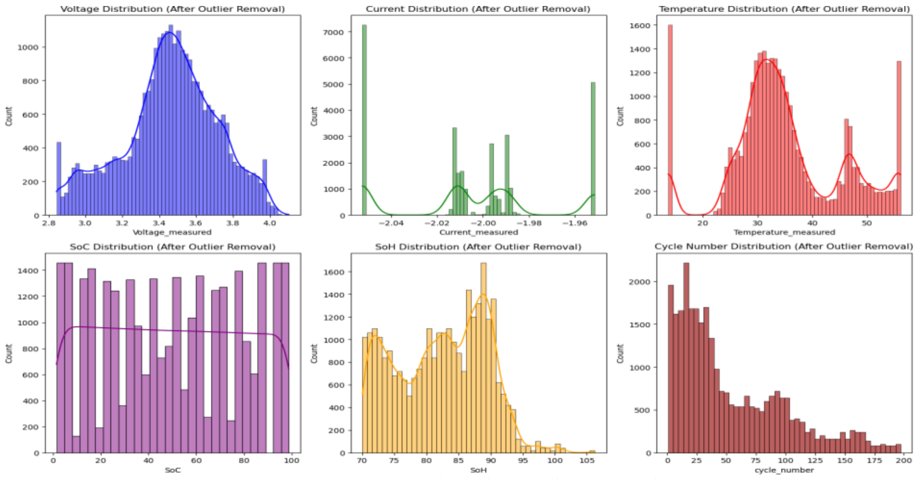

Fig. 2: Distribution Plot After Outlier Removal.

Distribution plots were constructed for all numerical features after eliminating outliers. Figure 1 exhibits that voltage is normally distributed around 3.45 V; this means that the battery voltage remains stable in most battery cycles. Current measurements show a multimodal distribution, with most of the measurements lying on -2.05 A and -1.95 A under a steady-state discharge condition. Temperature was bimodal, with two prominent peaks indicating that the battery is subjected to good and harsh thermal environments. SoC also seems fairly distributed from 0% to 100%, at least for the states of charge found in this dataset. In contrast, SoH appears narrowly centered from 70 to 95% with a quite long tail that goes beyond the 100% mark, probably because of a recalibration error. A right skewness was noted on the cycle number distribution, suggesting that many of these battery instances were encountered in their earlier operational trials.
The linear dependence among the variables was verified by means of the Pearson correlation matrix, presented in Figure 2. Analysis shows a strong correlation of 0.79 between the voltage and the SoC, indicating that voltage is a good representation of current charge level. A mild positive correlation of 0.27 shown through voltage and SoH specifies that voltage trends can be helpful in long-term health estimation. On the other hand, temperature is weakly correlated to SoC (-0.33) with a weak positive influence on SoH (0.24), therefore having its weak share in the ongoing battery aging processes. Current shows very low levels of correlation to both target variables, implying that it may indirectly represent short-term dynamics instead of contributing to long-term health prediction. The two main targets, SoC and SoH, on the other hand, show little correlation (0.01), thereby backing their independency; therefore, they should be modeled independently.

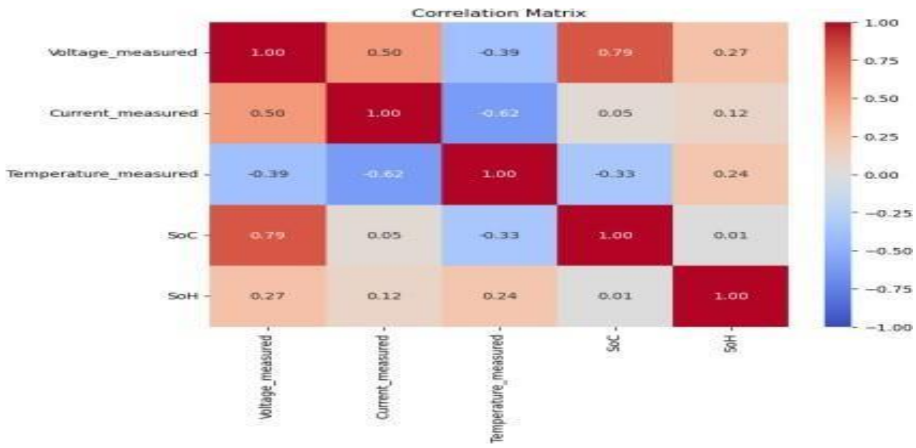

Fig. 3: Correlation Matrix.

Z-score methods and visual inspections were used for enlarged value recognition and removal lest these values could distort model training. This, in turn, helped improve normality in the features and stabilized the learning process without really changing much of the size or diversity of the dataset. Further, the cleaned dataset proves a strong signal to the learning algorithms, with balanced feature distributions and pertinent correlations for fine SOC and SOH predictions.

5. Datapre Processing

Preprocessing becomes one primary step in any data-driven system, which, in turn, determines the performance, predictive power, and resulting accuracy and generalizability of the model. In estimating battery SoC and SoH, more attention has been paid to ensuring data integrity, noise reduction, dimensionality reduction, and preparation for model training.

5.1. Missing value analysis

The raw data contained 29,180 observations and seven features. An initial audit confirmed no missing values, whether Voltage_measured, Current_measured, Temperature_measured, SoC, SoH, cycle_number, or battery_id. From the dataset summary, the non-null count for each column corresponded to the total number of rows; thus, imputation or interpolation was not required.

5.2. Outlier detection and removal

To avoid that extreme values would bias the learning process, outlier detection based on Z-score analysis was applied. A data point with a Z-score value greater than 3 or less than - 3 was considered an outlier. This method effectively caught some feature values considered to be anomalous, such as in Temperature_measured and Voltage_measured, showing heavy tails. These were removed to further stabilize the data distribution and ease model convergence. The cleaning stage was deemed valuable upon visual inspection of distribution plots (given above in Figure 1), showing much smoother and symmetric histograms after cleaning.

5.3. Feature engineering and k-best feature selection

In addition to the initial cleaning stage, feature engineering was undertaken to emphasize temporal and statistical trends apparent in battery behavior. New attributes such as Rolling_Mean_Voltage, Smoothed_Voltage, Rolling_Std_Voltage, Normalized_Voltage, Discharge_Efficiency, and Current_to_Voltage_Ratio were introduced. These features were created through moving average and standard deviation calculations in sliding windows or using ratio-type transformations, so that these attributes could better capture the divergence and degradation patterns.

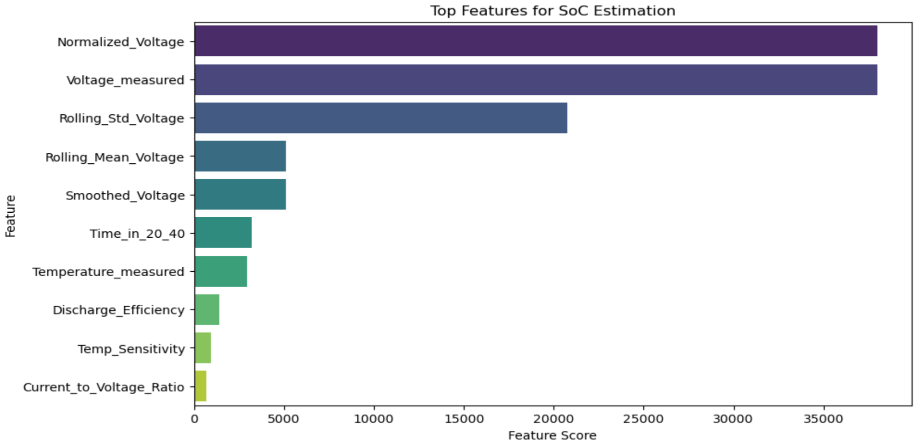

Fig. 4: Top Features for SoC Estimation.

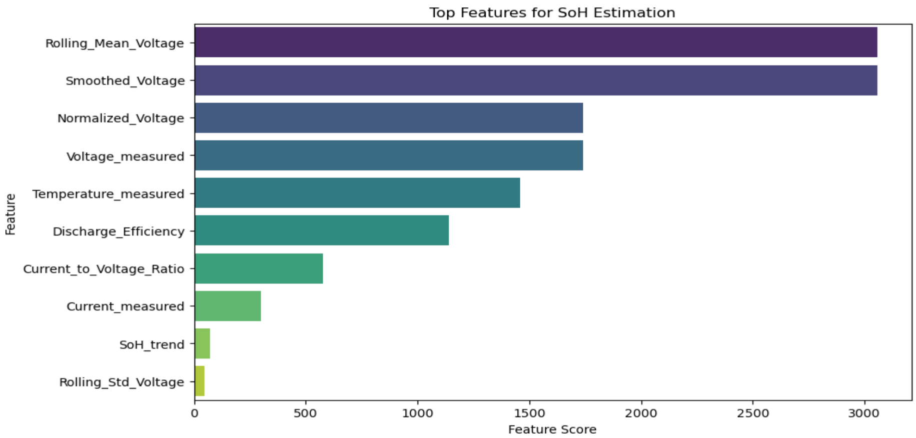

Fig. 5: Top Features for SoH Estimation.

Feature construction is followed by K-Best feature selection performed separately for SoC and SoH estimation. This method ranks and selects features for top-k predictive power based on a statistical test that could be the regression test. For SoC estimation, as presented in Figure 3, the leading contributing features are Normalized_Voltage, Voltage_measured, and Rolling_Std_Voltage. In this sense, voltage-based signals are highly indicative of the present charge state. For SoH prediction, shown in Figure 4, more contributory features include Rolling_Mean_Voltage, Smoothed_Voltage, Normalized_Voltage, and Temperature_measured, which correspond to usage and degradation effects over time.

5.4. Dataset splitting

To robustly evaluate the study, the dataset was randomly partitioned into training and testing sets with an 80:20 ratio. The training set was used to fit the learning models and optimize hyperparameters by performing cross-validation; the test set was left entirely untouched to measure the generalization performance of each algorithm on new, unseen data. Given the regression values of the targets (SoC and SoH), stratified splitting was not applied. Random shuffling was applied instead to keep representative distributions in both subsets. Such a multi- step preprocessing pipeline made sure that the models were trained on clean, informative, and well-structured data that would, in turn, lead to better prediction power and interpretability.

6. Methodology

The following section describes the end-to-end workflow used in the prediction of State of Charge (SoC) and State of Health (SoH) for lithium-ion batteries using machine learning and deep learning models. The methodology consists of model selection, training, validation, and interpretability. A hybrid modeling approach is used to leverage the advantages of temporal sequence learning and non-linear pattern recognition. Furthermore, explainable AI techniques are applied for the transparent revelation of feature influence and model decisions. Finally, the entire pipeline is configured such that it can be scaled up, be accurate, and stay generalizable for different battery conditions.

6.1. Linear regression

Linear Regression serves as a basic baseline model for estimation of continuous variables like State of Charge (SoC) and State of Health (SoH). A linear relationship is assumed between features and the target, as this setup is more suitable when the direct dependencies are present and when the features are well-engineered and free of multicollinearity.
The model's underlying process fits a linear equation in the form:

$$\hat{y} = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_n x_n \quad (1)$$

Where $y^\wedge$ is the predicted output (either SoC or SoH) $x_1, x_2, \ldots, x_n$ are the input features selected using the K-Best algorithm, $\beta_0, \beta_1, \ldots, \beta_n$ are the model coefficients learned during training. They are optimized by minimizing the Mean Squared Error (MSE) between the actual and predicted values.

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

Where m is the number of training samples.
The linear regression approach was used in this study through the scikit-learn library. Before training, all feature data was standardized with Z-score normalization, so magnitude of a coefficient would not be biased due to changing scales of the feature. There was no use of any kind of regularization for this baseline parameter setting, so it remained a pure evaluation for linear relationships. Separate models were generated for SoC and SoH using a set of best- ranked features that were suitably selected for each via the K-Best selection.

6.2. XG-boost

Extreme Gradient Boosting or XG-Boost, for short, is a next-generation ensemble- learning algorithm that builds a sequence of decision trees with the aim of improving the prediction accuracy basically by "correcting" the errors made by an antecedent set of trees. Given its efficacy in terms of speed, scalability, and proficient handling of non-linear feature interactions for structured tabular datasets, it naturally fits our purpose of estimation of State of Charge (SoC) and State of Health (SoH) of lithium-ion cells with engineering telemetry data. So, XG-Boost gains improvement by learning on the residual errors one after another, in which each successive tree fits the latest residuals of the difference between the real and predicted values. The iteration of this process either stops after a preset number of iterations or whenever the performance stops improving on validation data. One distinctive attribute, among many others, when differentiating XG-Boost from others is its built-in regularization that essentially penalizes the more complex ones to avoid any sort of overfitting.
Separate models of XG-Boost regression were developed for SoC and SoH prediction purposes. The input features for each model were selected by a K-Best attribute selection method so that only the most relevant attributes contributed to the learning. Hyperparameter optimization was then executed wherein the number of estimators (trees), tree depth, learning rate, and proportion of the samples and features were tuned for each round of training.
XG-Boost has shown extreme competence in modeling the intricate and complex non-linear relationships inherent in battery data. Nevertheless, it proved to be so brilliant in SoH estimation because long-term degradation patterns are complex and subtle. This model gave the best performance with missing or noisy values, thus realizing robustness in real-life applications.
By exhibiting its native capability for feature importance ranking, XGBoost offers better explainability. Beyond that, it complements tools like SHAP in delineating which sensor signals and derived metrics most heavily influence model predictions.

6.3. Recurrent neural network (RNN)

Recurrent Neural Networks are a specific set of neural architectures designed to work with sequential data; a memory has to somehow be retained through hidden states with respect to the input, which means that for time-related issues like battery health monitoring, where signals such as voltage, current, and temperature change over time, RNNs find applications.
In contrast with traditional feedforward neural networks treating each input independently, an RNN follows a feedback loop such that information can be preserved from one time step to another. This mechanism allows the network to capture the temporal dependencies required to study the battery degradation pattern or charging/discharging behavior through cycles.
In this work, RNNs were used to forecast SoC and SoH. As input to the network, sequences of sensor data and other features were provided, including Voltage_measured, Current_measured, Temperature_measured, and some computed features such as Rolling_Mean_Voltage. These sequences were presented as multi-dimensional arrays with multiple time steps per sample, letting the RNN learn representations spanning time.
Separate RNN training was conducted for SoC and SoH estimation. Key training parameters were chosen using a grid search: number of hidden units, sequence length, learning rate, and batch size. The network was trained using mean squared error as the loss function, with convergence speed maximized via the Adam optimizer.
While RNNs could, in principle, achieve reasonable accuracy for short-term dependencies in the data, they fail to model the long-range trends due to the vanishing gradient problem. This drawback motivated the use of architectures such as LSTMs and GRUs, which could theoretically remember over longer sequences.

Nonetheless, the RNN model served as a very good baseline, capable of modeling immediate temporal dependencies and thereby furnishing a relative benchmark for the deep models employed for this paper.

7. Bidirectional Long Short-Term Memory (Bi-LSTM)

Bidirectional Long Short-Term Memory (Bi-LSTMs) is a somewhat complex extension of recurrent neural networks that process sequential data from both directions-that is, from past to future and from future to past. This architecture allows the model to have the entire context in time-series data, which is greatly needed in battery analytics since prior usage and upcoming conditions affect the present state.
While in normal RNNs or unidirectional LSTMs, the input is read in the forward direction only, Bi-LSTMs apply two LSTM layers in parallel: one from the start of the sequence to the end (forward) and one from the end to the start (backward). Their outputs are merged at every time step to allow the model to associate dependencies in both forward and backward temporal directions. This also makes Bi-LSTMs a suitable candidate for predicting battery SoH, where the degradation behavior is not only dependent on past cycles but also on the immediate trends.
LSTM layers can remember things using internal memory and gating to regulate what kind of information should be allowed to go through in every time step. Thus, it solves the vanishing gradient problem faced by regular RNNs, making LSTMs stronger in apprenticing sequential patterns over longer time windows.
During the current project, Bi-LSTM models for SoC and SoH estimation were trained independently using sequential features such as voltage, current, temperature, and their statistical summaries like rolling averages. The model configuration consisted of a single bidirectional LSTM layer having 64 hidden units, followed by a dropout layer, serving as a regularizer, and ended with a dense output layer that performs regression prediction.
Adam was used as the optimizer, and MSE was used as the loss function during training. In order to select the best weights and, therefore, reduce overfitting whatsoever and promote better generalization once validation loss stops to improve, early stopping has been implemented in the training process. Bi-LSTM came out far better than the unidirectional LSTMs or standard RNNs, particularly for predicting SoH because it takes into account information from both backward and forward time steps.

8. LSTM + GRU Hybrid Model

This model captures two most powerful sequence deep models, namely LSTM and GRU, in order to leverage each other's individual advantages to compensate for their limitations while modeling complex temporal patterns in battery telemetry data. The architecture is proposed for maximizing any time-series-based task accuracy, such as SoC or SoH prediction, while also making the learning process more efficient.
LSTM layers help keep long-term dependencies between time sequences. They have more control with gating mechanisms: forget gate, input gate, and output gate. Such control allows the model to choose what is relevant and what to discard through time. The method, hence, can be suitable to model slowly varying battery behavior, such as SoH degradation over hundreds of cycles.
GRU layers, by contrast, offer a simpler architecture with fewer gates, making training faster and less computationally demanding without much sacrifice on performance. The GRU is controlled by update and reset gates, leading to an elegant design that is good at modeling short-term variations and recent trends in charging or discharging cycles that influence SoC.
In the hybrid architecture proposed, the first layer is an LSTM layer that captures long-range dependencies and deeper temporal features from sequences of sensor readings such as voltage, temperature, and current. It is then followed by a GRU layer that healthily refines the extracted temporal features to capture finer and more recent temporal dynamics. Modeling of this kind allows the model to generalize well across battery behavior that presents slow or fast-changing trends.
The entire construction is as follows:

- An LSTM layer with 64 hidden units with return sequences set to True,
- A GRU layer with 32 units to apply short-term temporal refinement,
- A dropout layer for averting overfitting through training,
- A fully connected dense layer with linear activation for regression output.

Depending on the time window, the input sequences are prepared from sensor readings and engineered features; then the 3D tensors are reshaped into (samples, time steps, features). Both SoCs and SoHs are treated independently as two regression problems, where two separate models are trained for each target variable.
Adam optimization is used for training and tuning the parameters. A learning rate was incorporated for tuning, while the loss function is taken care of by MSE. A validation loss assessment is carried out; hence early stopping is employed to prevent over-fitting and avoid wasting computations. This hybrid structure is a trade-off between LSTM for long-term dependency modeling and GRU for speed and efficiency, better fit for real battery performance prediction problems.
The experiments proved that the LSTM+GRU hybrid model outperformed pure LSTM, GRU, and simple RNN models in prediction accuracy and convergence speed. The architecture tends to suit its adaptations for variable conditions in the number of batteries subjected to testing, thus rendering it well-suited and robust in predictive maintenance.

9. Explainability with Lime

To truly provide interpretability to the model and to build a degree of trust in the SoC and SoH predictions, an explainable AI technique, LIME (Local Interpretable Model- Agnostic Explanations) was used. LIME aims to explain the decisions made by complex machine learning models by approximating these models locally using simple interpretable models, such as linear regressors. Hence, it explains how the model arrives at an individual decision and not the global interpretation thereof.
Contrary to global feature importance scores that tell us which features are generally important, local explanations by LIME perturb the input data around a particular prediction and look at the change in output. Then, a simple, interpretable surrogate model is fit to a small neighborhood around the instance to determine which features contributed most strongly, either positively or negatively, to that particular prediction. This provides a level of transparency for single instances that are essential to understanding and handling of edge cases and outliers for critical applications, such as battery health monitoring.

LIME explains predictions generated by the SoC estimation method of the XGBoost model and by the SoH estimation method of the RNN model in this work. The LIME plots reflected how certain feature values influenced predictions: Rolling_Std_Voltage, Voltage_measured, Temperature_measured, and Time_in_20_40. In SoC prediction, voltage features mostly exhibited varying positive or negative influences depending on their ranges. For SoH predictions, time-series derived and cycle-related features weighed more since they indicated battery aging patterns.

Integrating LIME served several main reasons:

- Transparency: domain experts understood the individual model decision through LIME explanations
- Feature validation: helped affirm that the features selected by K-Best actually influenced prediction
- Debugging: Could check for features possibly biased and/or against the model.
- Auditing: Provided an explanation behind a model decision that could be digested by non-technical stakeholders.

Hence, LIME became a linguistic link from black-box modeling to human interpretability. It increased accountability in the predictive framework by associating the logic beneath prediction with domain knowledge, especially in high-impact domains such as predictive maintenance and battery state forecasting.

10. Results

This segment presents performance evaluation measures for the various predictive systems developed using machine learning and deep learning methods. The goals subject to estimation are the State of Charge (SoC) and the State of Health (SoH) of lithium-ion batteries. Model accuracies are quantitatively measured using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and R^2 Score statistics. Model performances are then compared between the tasks to deduce the best performing model and generalization ability. To substantiate the numerical results, prediction plots and error distributions are also presented for.

11. SOC Result: Linear Regression

Despite the model being linear in nature, the description points to a decent predictive accuracy when it comes to SoC prediction through Linear regression. From the scatter plot one observes that the predicted SoC values sit well along the diagonal reference line with the actual SoC values, barring some dispersion towards the lower and mid-range charge levels. Based on this model, the MSE achieved was 75.20, which represents the mean of squared differences between actual and predicted values. Based on RMSE, which is the square root of MSE, the average SoC prediction errors stand at less than 9%, slightly acceptable for a baseline estimate.

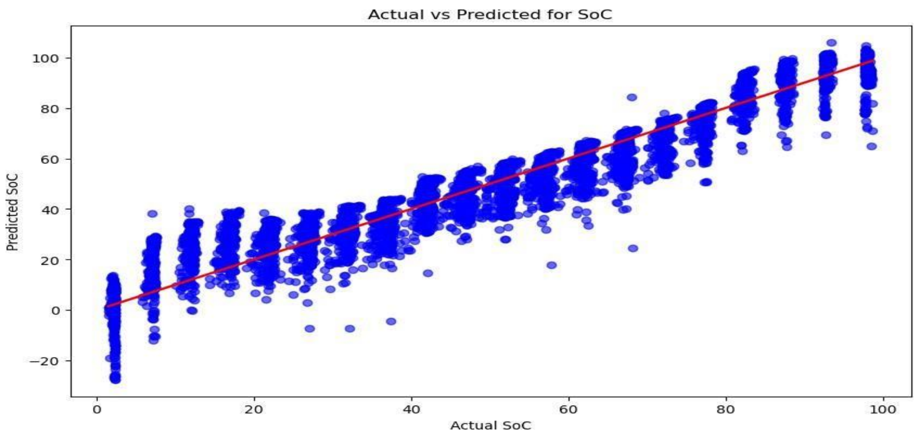

Fig. 6: Actual vs Predicted for SoC by Linear Regression.

The MAE was 6.83, which means that this average error measurement is indifferent to the directions taken by individual errors. Of greater importance, however, was that the R^2 score was 0.91, implying that 91% of the variation in SoC can be explained by the input features employed in this regression. Though normal LinearRegression isn't conceived for continuous regression problems such as SoC prediction, here it well serves as a robust benchmark against which more complex methods can be evaluated while, considering its simplicity, it does a rather fine job.

XG-Boost:

The XG-Boost-based State-of-Charge prediction is looked into for evaluation of accuracy and was found to be far better as compared to Linear regression and other simpler models. The purple dots in the scatter plot lie tightly along the red diagonal, indicating good conformity with the ideal where predicted SoC equals actual SoC. The tight grouping along the equality line indicates Very High Accurate Model.

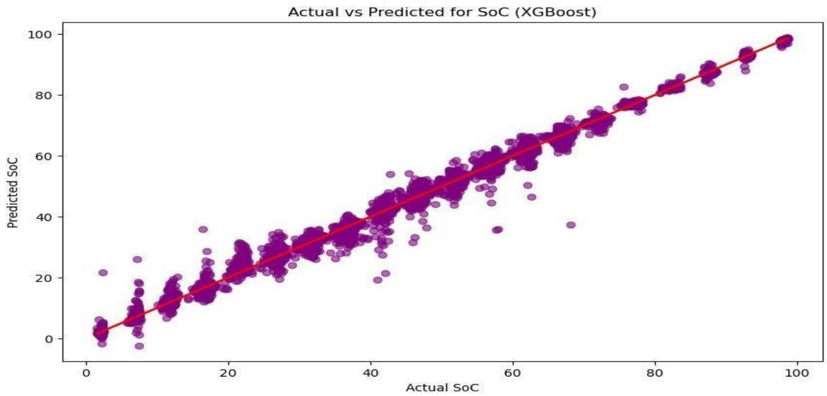

Fig. 7: Actual vs Predicted for SoC by XGBoost.

The squared value of the difference between the actual and predicted values has an average of 3.93, which is almost negligible. The low value of RMSE-1.98 means that the predictions lie in the range of error of around 2 percentage points on the average. An error of deviation of approximate 1.14 units from the actual SoC values is shown by MAE.

Most importantly, the R^2 score of 0.995 shows that about 99.5% of variability in SoC values is explained by the XG-Boost model, which means that there is a high possibility of a strong generalization of this model on the test set. Such results indicate that XG-Boost is capable of capturing nonlinearity and interaction among input features and remains one of the best- performing systems for SoC estimation in this study.

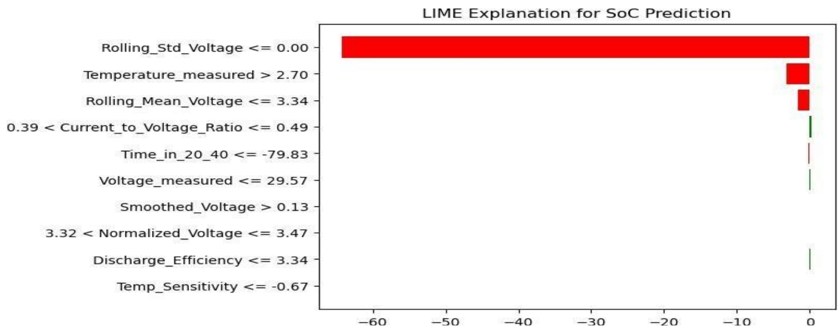

Fig. 8: Lime Explanation for SoC by XG Boost.

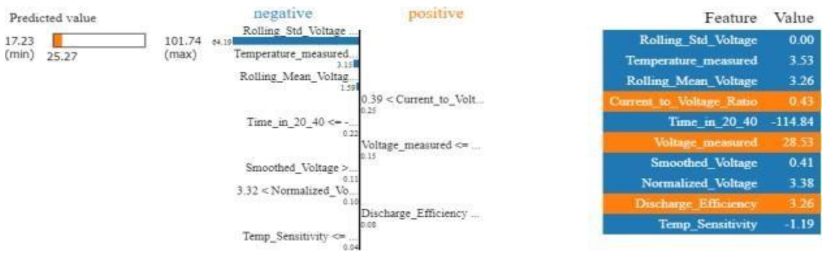

Fig. 9: Overall Lime Explanation by XG Boost.

LIME explanations of State of Charge prediction show that the greatest negative contributions were from Rolling_Std_Voltage, Temperature_measured and Rolling_Mean_Voltage, in that order. These lowered the predicted SoC, shown by the downward direction of red bars in the bar chart. Current_to_Voltage_Ratio and Discharge_Efficiency made positive contributions that slightly elevated the prediction. Interactive LIME output further supports that minute fluctuations in voltage stability and temperature substantially influence SoC estimation. This interpretability legitimates feature relevance and increases model decision-making process trust by unmistakably demonstrating feature influence on individual predictions.

RNN:

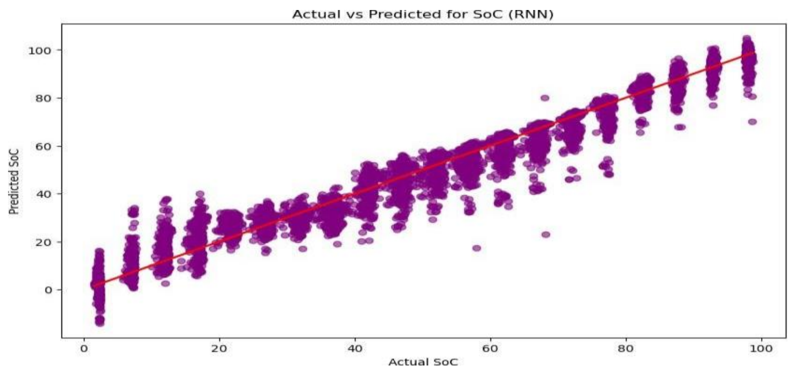

Fig. 10: Actual vs Predicted for SoC by RNN.

The Recurrent Neural Network (RNN) SoC prediction model has an average performance in following temporal relationships in the battery time-series data. As observed from the plot, the majority of predicted values fall on the red ideal reference line, with increased dispersion compared to XG-Boost. The model reached an MSE of 40.04 and RMSE of 6.32, reflecting higher error values compared to ensemble methods. The MAE of 4.75 is a reasonable average error of prediction, and the R^2 value of 0.95 still indicates high explanatory power. Generally, the RNN can learn short-term patterns but loses long-range accuracy, which indicates the necessity for deeper architectures.

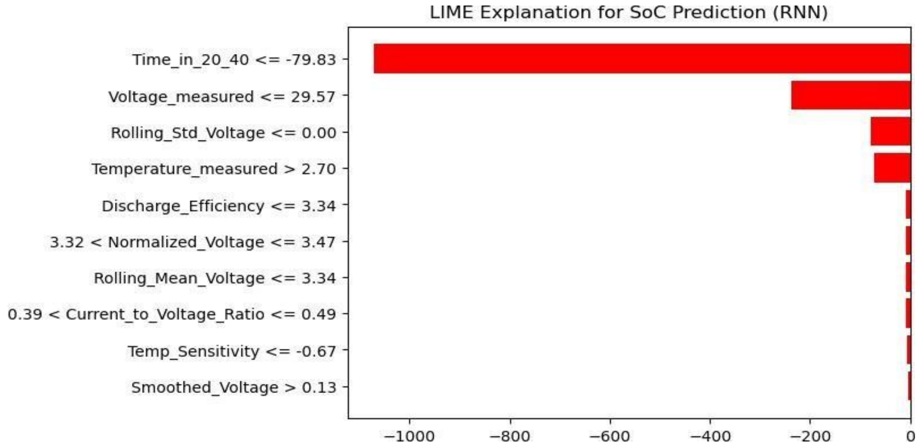

Fig. 11: Lime Explanation for SoC by RNN.

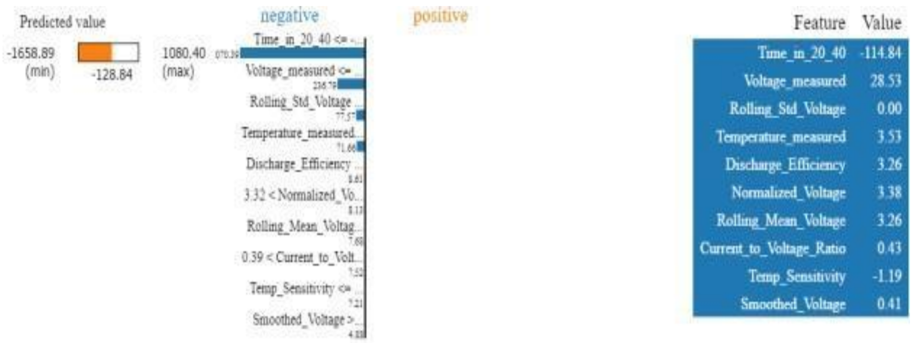

Fig. 12: Overall Lime Explaination by RNN.

The LIME explanation of the SoC prediction based on RNN shows that the feature Time_in_20_40 accounts most for degrading the SoC prediction. The other features that account for lower predictions are Temperature_measured, Rolling_Std_Voltage, and Voltage_measured, showing some importance in affecting the output of the model. No feature outweighed the other in the stated prediction scenario, as visualized. This imbalance, which suggests that the RNN might be overly sensitive to certain voltage and time-related features, may result in underprediction or unstable outputs and thus provides further motivation for more stable sequence models such as Bi-LSTM to be implemented in the future.

BI-LSTM:

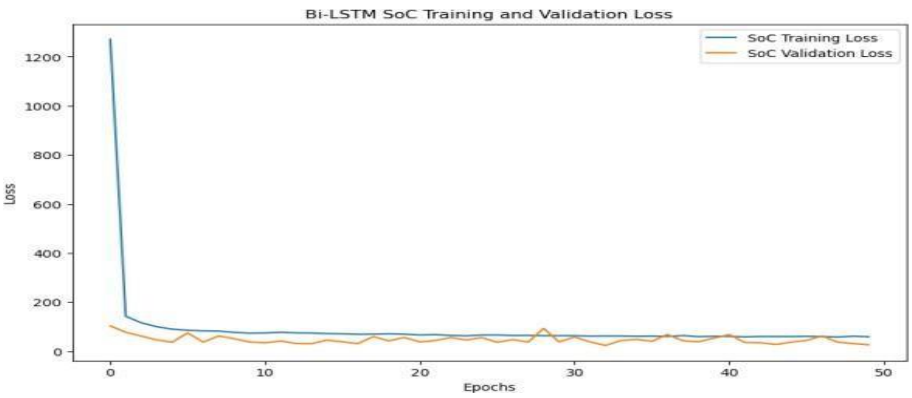

Fig. 13: Loss Curve for Bi-LSTM.

Training with validation loss curves and actual versus predicted plots evidenced that the Bi- LSTM method gave the strongest performance and generalization. At the appendence of loss during training, the model observed no overfitting phenomenon since the two curves started to converge at around epoch 40. Most of the predicted SoC values lie close to the ideal axis reference line in the scatter plot, being equally good in terms of prediction accuracy for the entire SoC value range.

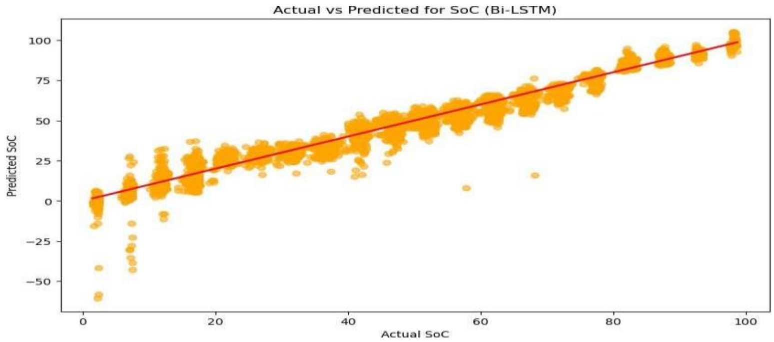

Fig. 14: Actual vs Predicted for SoC by Bi- LSTM.

In terms of numerical indicators, the model yielded an MSE of 26.24 and an RMSE of 5.12, which signify quite low errors in predictions. The MAE value of 3.64 points toward a small deviation between actual values and predicted SoC. An R^2 score of 0.968, equally astonishing, reveals that 96.8 % of the variation in SoC could be accounted for by the model. The Bi-LSTM, for its part, could satisfactorily represent past and future dependencies and thus perform well in sequence-based learning tasks, much better than a standard RNN or linear model.

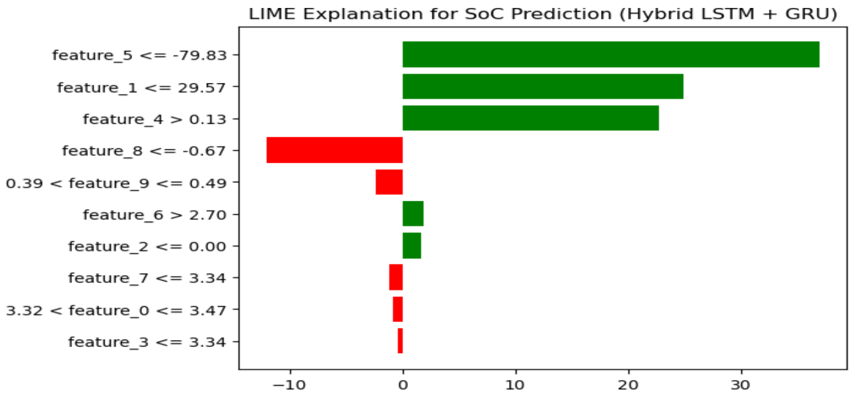

Fig. 15: Lime Explanation for SoC by Bi-LSTM.

The level explanation for the Hybrid LSTM + GRU model on SoC prediction reveals features named feature_5, feature_1, and feature_4 to have the most significant positive impact on the output, thus raising the predicted SoC value. These could be factors relating to time or voltage trends. Contrarily, feature_8 and feature_9 degrade the prediction, almost suggesting if certain thresholds exist (low temperature or efficiency) that decrease the expected SoC. The diagram portrays a very well-grounded local explanation in which LIME reveals factors that support and suppress the model decision for that particular instance, thereby improving interpretability of a relatively complex sequential model.

Hybrid Model

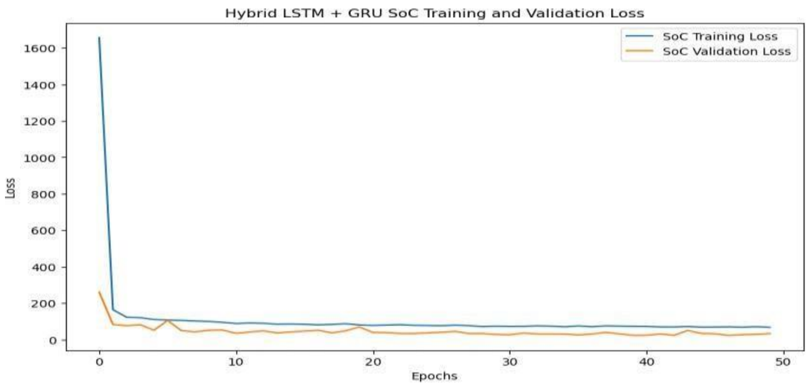

Fig. 16: Loss Curve for Hybrid Model.

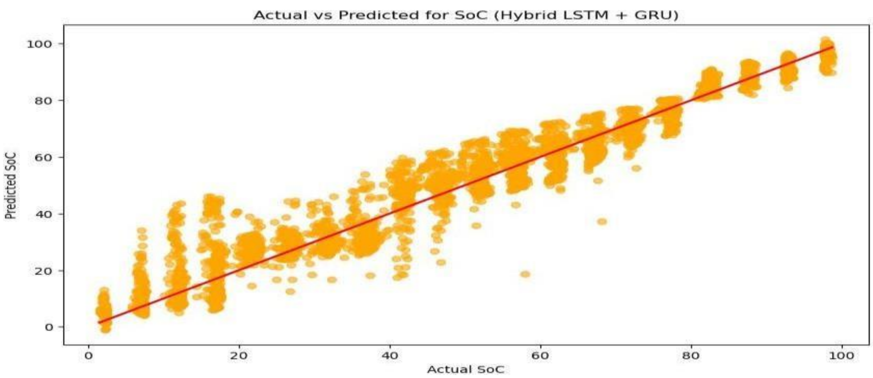

Fig. 17: Actual vs Predicted for SoC by Hybrid Model.

The powerful learning might have been acquired by using SoC prediction through the LSTM + GRU approach. The training and validation loss values rapidly decreased and remained steady all along training without overfitting; hence, the convergence was good. The actual-predicted plot shows a close alignment with the red diagonal line of the ideal presentation, which reveals the model's efficacy in capturing temporal dependencies in SoC data.
The MSE of 33.37 and RMSE of 5.78 mean that the model explains 96% of SoC variance. With an MAE of 4.18 and an R^2 of 0.96, an expected SoC level can be regarded as a highly accurate estimate. This ties the system among the most suitable deep learning architectures for battery charge estimation at reasonable accuracy.

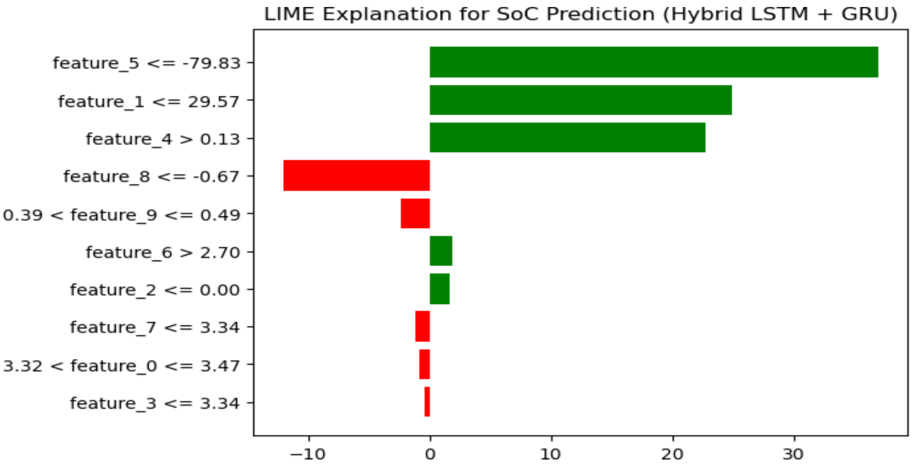

Fig. 18: Lime Explanation for SoC by Hybrid Model.

The LIME explanation for the Hybrid LSTM+GRU SoC prediction enlightens us about the features that influenced a certain prediction and to what extent. The strongest positive contributions were paid by feature_5, followed by feature_1, and then feature_4, which possibly denote time in range, voltage, and temperature insofar as these features push the predicted SoC upward-the sign for better battery charge conditions. Feature_8 stands as the most negative factor being in all probability suggestive of degradation trends or inefficiencies. Slightly weaker negative influences came into play from feature_9, feature_7, and feature_0 to give more preciseness to the prediction. Such an equal mix of positive and negative contributions is, on the one hand, confirmation to this single instance that the model managed to understand in a very detailed fashion, from a battery-state point of view.
Comparisons:

Model	MSE	RMSE	MAE	RÂ² Score
Linear Regression	75.2008	8.6718	6.8384	0.9096
XGBoost	3.9385	1.9846	1.1448	0.9953
RNN	40.0459	6.3282	4.7477	0.9519
Bi-LSTM	26.2446	5.123	3.6487	0.9684
Hybrid Model	33.3700	5.7766	4.1781	0.9598

12. SOH Result

Linear Regression:

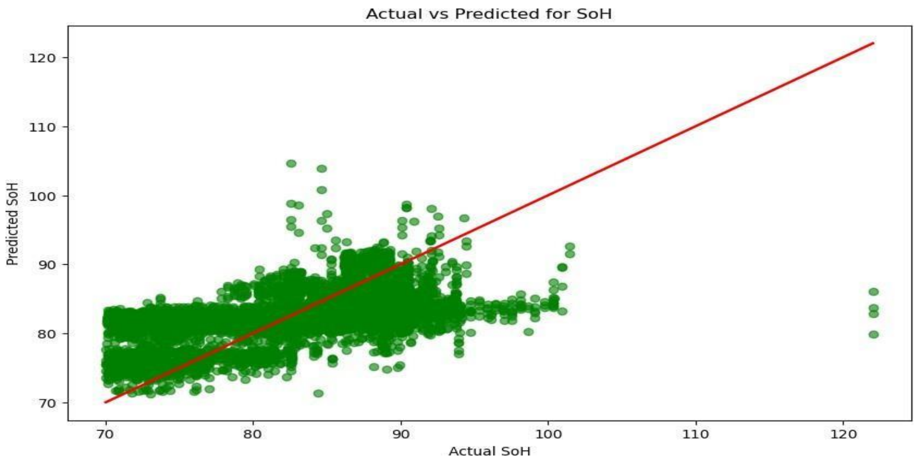

Fig. 19: Actual vs Predicted for SoH by Linear Regression.

While the SoH prediction did fairly well using a Linear Regression model, it failed to grasp the upward non-linear degradation trends characteristic of the battery dataset. Almost all predicted values clustered in a somewhat narrow band that substantially deviated from the ideal 45- degree diagonal line on the actual vs. predicted-soH plot. The model recorded MSE and RMSE values of 32.88 and 5.73, which correspond to a fairly average error. It registered an MAE of 4.65. However, the R^2 value was only 0.33, so the model managed to explain just 33% of the variance in the SoH, making this model ill-suited for long-term battery aging trend prediction
XG Boost:

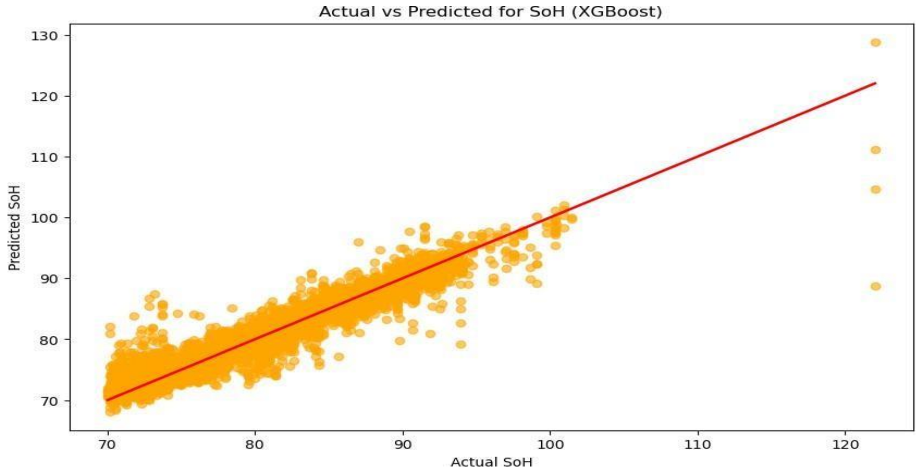

Fig. 20: Actual vs Predicted for SoH by XGBoost.

The XG-Boost model for SoH prediction enjoys high accuracy and efficiency than any other methods considered. The scatter plot shows points lying close to the red diagonal line, confirming the accuracy of the model's predictions for the entire SoH range. The MSE obtained was 3.65, while the RMSE obtained was 1.91, denoting small magnitudes of error in the prediction. Contrastingly, the MAE was measured to be 1.23, implying the average prediction is just about one unit away from the real SoH values. To my best knowledge, this model stands apart, having an R^2 score of 0.93, meaning that about 92% of the variance in SoH is explained, way above any other benchmarking method.

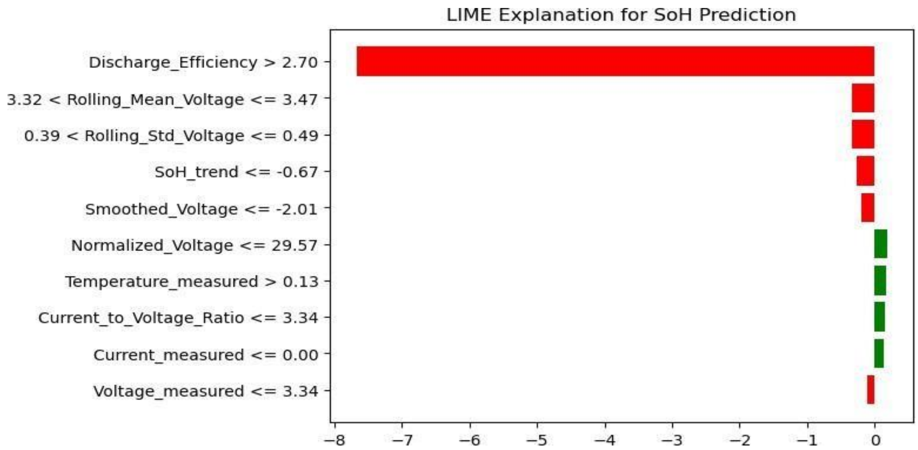

Fig. 21: Lime Explanation for SoH by XGBoos.

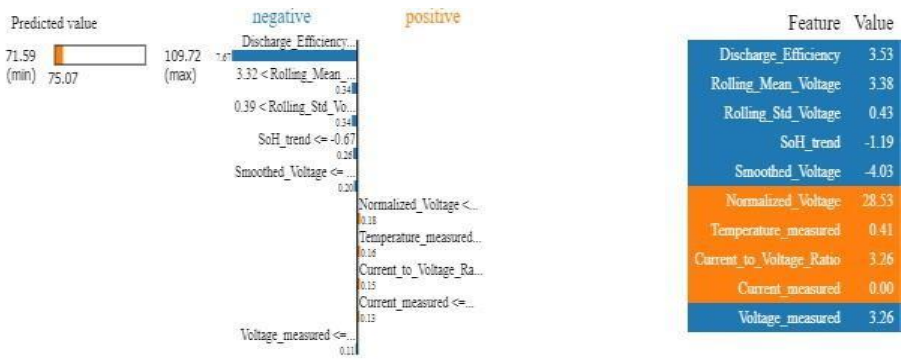

Fig. 22: Overall Lime Explanation by XGBOOST.

According to the LIME explanation for SoH prediction by the XG-Boost model, Discharge_Efficiency exerted the maximum detrimental effect on the predicted value, with Rolling_Mean_Voltage, Rolling_Std_Voltage, and SoH_trend taking secondary positions. These features worked in opposition toward lowering the prediction generated by the model, therefore, pressing strongly toward the degradation-related prediction. Instantaneously brought upward pressure on the prediction were Normalized_Voltage, Temperature_measured, and Current_to_Voltage_Ratio. The plot presents demonstration of a well-balanced but very critical detriment toward discharge-related and voltage-related metrics as considered in modeling the long-term battery health. LIME brings to light the factors behind every prediction to ensure full transparency when estimating SoH.

RNN:

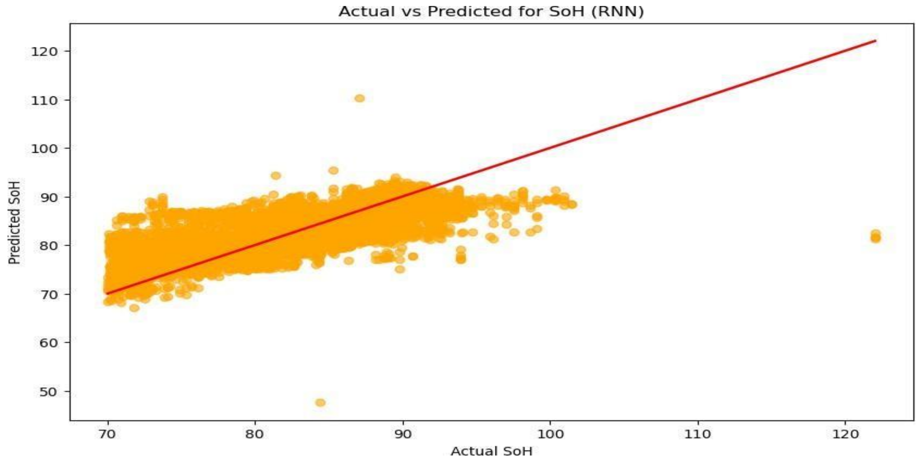

Fig. 23: Actual vs Predicted for SoH by RNN.

The RNN SoH prediction model performed moderately, points in the scatter plot spreading to some extent away from the ideal diagonal. It followed the general trend; however, the model did not seem to properly capture the long-term degradation pattern. This is why predictions were compressed in the 75-90 band, regardless of the actual SoH value. Quantitatively, the MSE was 24.08 with an RMSE of 4.91 and an MAE of 3.79 - values that indicate strong prediction errors. An R^2 of 0.51 means that the model accounts for just over half of the variation in SoH. This dismissal of a high R^2 puts into perspective the drawback of the RNN in understanding more intricate temporal dynamics in battery health.

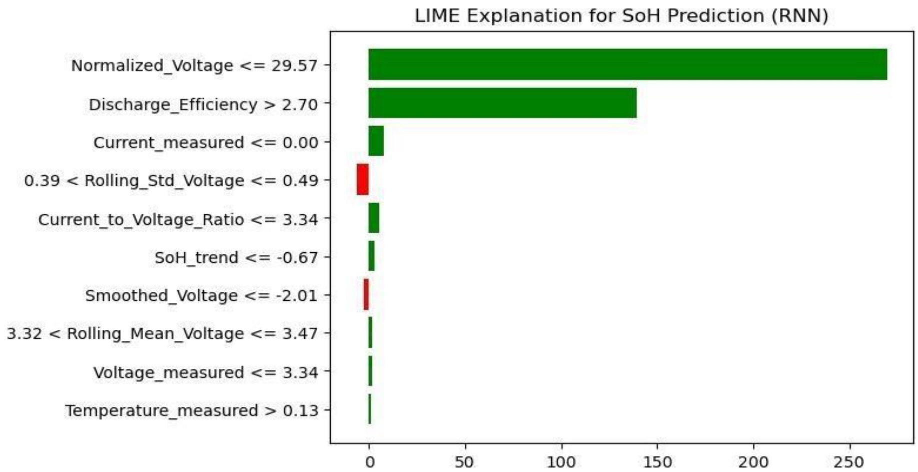

Fig. 24: Lime Explanation for SoH by RNN.

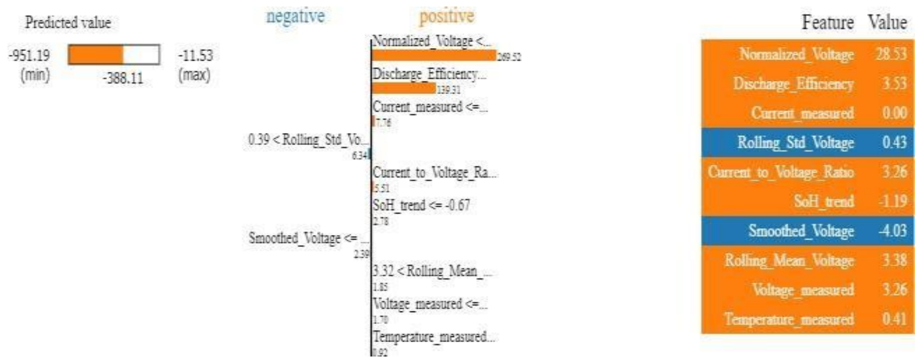

Fig. 25: Overall Lime Explanation by RNN.

According to the LIME explanation offered for the RNN's SoH prediction, Normalized_Voltage and Discharge_Efficiency were the major positive contributors, driving the predicted value up markedly. These features stand to represent operational health and consistency of the particular battery. Other features, such as Current_measured and Current_to_Voltage_Ratio, exerted relatively small positive influences. Conversely, Rolling_Std_Voltage and Smoothed_Voltage pushed down slightly on the prediction, meaning that perhaps variability in voltage is a sign of degradation. The LIME visualization may thus be interpreted at least somewhat intuitively, with respect to local feature influences, to reveal both factors that support and oppose individual predictions of SoH, thereby facilitating a transparent look at the reasoning behind the model.

BI-LSTM:

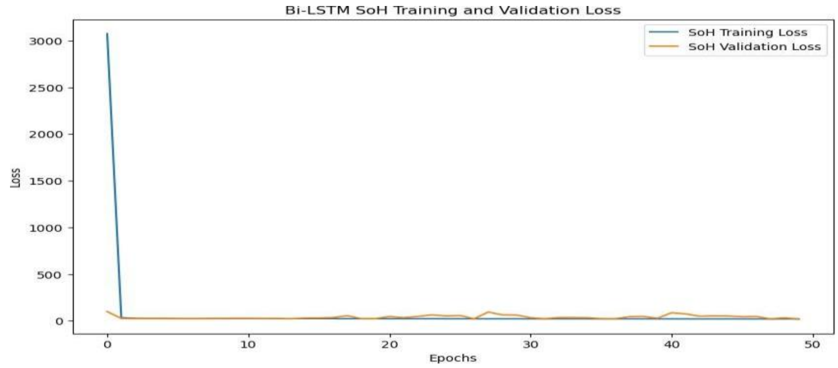

Fig. 26: Loss Curve for Bi-Lstm.

The so-called predictive Bi-LSTM model has a moderate success in detecting temporal dependencies but at the same instance keeps training stable. The loss curve displays a steady convergence, with early stabilization of training and validation losses, thus indicating good generalization behavior and lack of overfitting. Numerically, it shows an MSE of 19.01, an RMSE of 4.36, and an MAE of 3.32, implying relatively controlled errors. An R^2 value of 0.61 means that the model is capable of explaining 61% of the variability in SoH and is better at doing so than the RNN and linear models but the performing of XG-Boost is superior to it. Though this has been made somewhat successful in smoothing SoH prediction, further tuning shall bring about improvements in long-term degradation modeling.

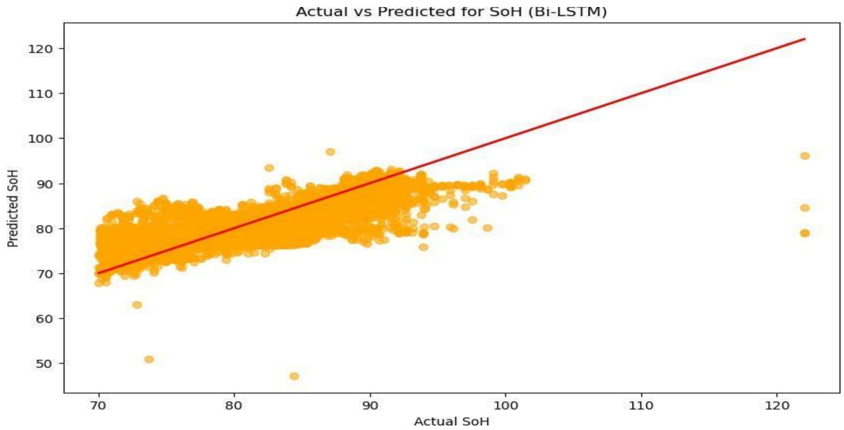

Fig. 27: Actual vs Predicted for SoH by Bi-LSTM.

Figure 27 displays state-of-the-art scatter plots about SoH prediction with Bi-LSTM. As is visible, it agrees mostly with the ideal line, with some spreading out and underestimation beyond the higher SoH values. Almost all predicted points are densely clustered along the diagonal, pointing to the model's ability to extract meaningful temporal patterns from the data. The model provides the most accurate predictions for SoH in the middle range and tends to squeeze values at either end. This is consistent with its ~0.61 reported R^2 score, depicting a moderate explanation level. The Bi-LSTM learned in a stable fashion and was, in general, superior to simple recurrent architectures in predictive performance.

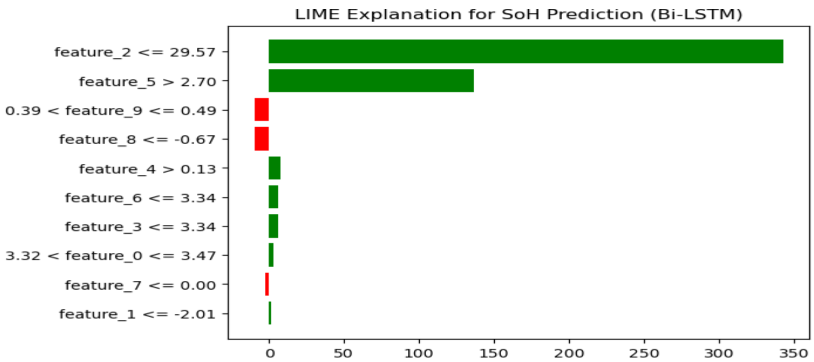

Fig. 28: Lime Explanation for SoH by Bi-LSTM.

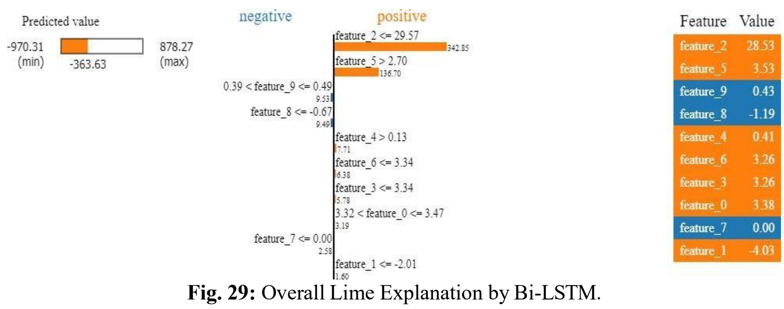

Fig. 29: Overall Lime Explanation by Bi-LSTM.

This is a LIME explanation of predicting SoH by Bi-LSTM, where feature_2 and feature_5 were hugely contributing in an opposite direction for the SoH prediction, implying Normalized Voltage and Discharge Efficiency representing healthy states in the realm of battery behavior being stable and efficient. On the contrary, feature_9 and feature_8-perhaps corresponding to Voltage Variability and Degradation Trend-are negatively affecting the SoH value and hence reduce the Predicted SoH. Next come Temperature_measured, Rolling_Mean_Voltage, and Current_measured, contributing in smaller positive and negative amounts. So, this LIME visualization implies that the model makes sense from the standpoint of battery health as per domain knowledge concerning voltage quality and efficiency indicators.
Hybrid Model

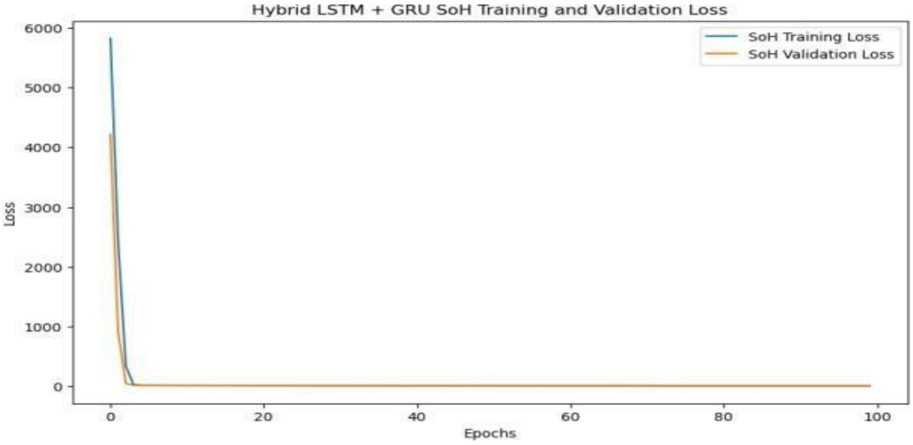

Fig. 30: Loss Curve for Hybrid Model.

Predicting SoH with the Hybrid LSTM + GRU model performs the best characteristics of fitting, thus affording predictive precision as well as convergence stability. Both training and validation loss values diminish abruptly and commence to flatten afterward, thereby signifying that the speeds of learning are higher, with little overfitting. Being able to generalize well for the SoH range, the cluster of points between the actual and predicted values lie closely along the ideal diagonal.
The MSE stands at 5.44, RMSE at 2.33, and MAE at 1.60, reflecting a low error in prediction. The once R^2 at 0.89 level of explanation means that it has so far explained the variation in SoH at nearly 89%, against conventional RNN and Bi-LSTM models. This proves that merging the long-term memory of an LSTM with the more efficient GRU of the battery health forecast in a rather complex time-dependent scenario brings further upgrade in their predictive ability.

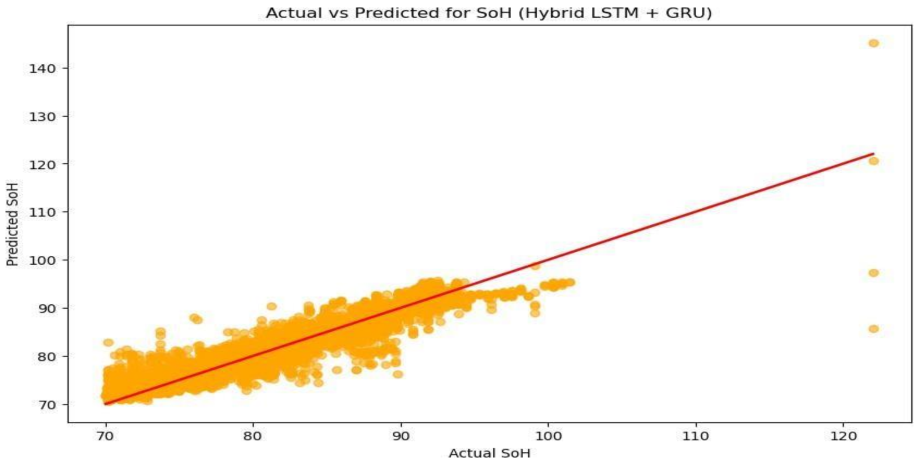

Fig. 31: Actual vs Predicted for SoH by Hybrid Model.

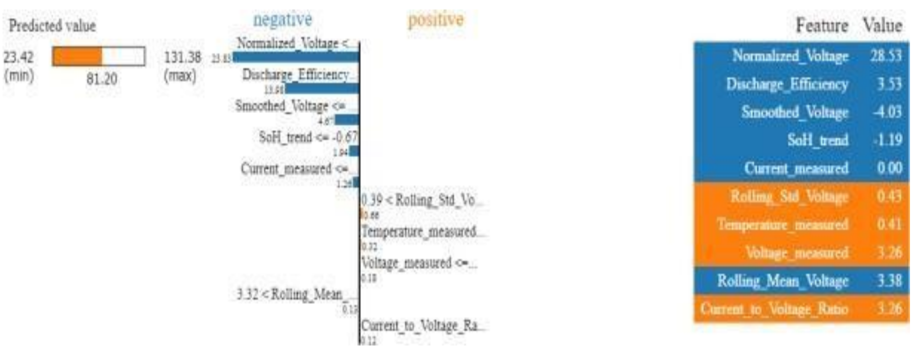

Fig. 32: Lime Explanation for SoH by Hybrid Model.

LIME-dependent rationale computation of SoH means Normalized_Voltage, Discharge_Efficiency, and Smoothed_Voltage pushed the predicted value down in the sense that lower voltage stability and lower discharge efficiency imply bad surface health of batteries. Rolling_Std_Voltage, Temperature_measured, and Voltage_measured had opposing effects with small positive effects on the estimate; thus, they somewhat increased the predicted value of SoH. This explanation gives a balanced picture of the opposing weights of features in

influencing the output. Such local interpretability really strengthens the claim of the model in exposing feature interplay for each and every one of the predictions throughout the entire complicated deep learning pipeline.
Comparison:

Model	MSE	RMSE	MAE	RÂ² Score
Linear Regression	32.8824	5.7343	4.6487	0.3297
XGBoost	3.6489	1.9102	1.228	0.9256
RNN	24.0804	4.9072	3.7934	0.5091
Bi-LSTM	19.0057	4.3596	3.3233	0.6126
Hybrid LSTM + GRU	5.4355	2.3314	1.5981	0.8892

13. Conclusion

Several machine learning and deep learning models were considered in this study for performing accurate SoC and SoH estimations of lithium-ion batteries. From the SoC prediction perspective, classical models such as Linear Regression achieved moderate performance, limited by their inability to address non-linear relations effectively. In contrast, XG-Boost took the prediction to the next level, with an R^2 value exceeding 0.99. Both recurrence structures RNN and Bi-LSTM outperformed others in modeling temporal sequences, while Bi-LSTM generalized better. However, the Hybrid LSTM + GRU model outperformed them all, offering a balance between accuracy and efficiency with an R^2 score of
0.96 and minimal RMSE.
A similar trend was seen for SoH prediction-N-rather underperforms to capture long-term degradation patterns. XG-Boost, however, excelled, with an R^2 score of 0.93 and minimum error metrics-interactions modeling complex. Deep models like Bi-LSTM and Hybrid LSTM + GRU better captured temporal dependencies than RNNs, whereas the hybrid model balanced accuracy and stability the best ($R^2 \approx 0.89$). While XG-Boost worked best for the structured data, the Hybrid LSTM + GRU model was best for time-series-based battery monitoring, making them both useful within a smart battery management system.

14. Future Enhancement

Future enhancements to this battery SoC and SoH prediction framework can focus on the expansion of the scope and performance of the models. Real-time data from IoT-enabled sensors may serve the purpose of health monitoring in a continuous manner and dynamic learning. Transfer learning methods can be applied for adapting models trained on one type of battery to another, thus minimizing the retraining effort. Applying hybrid explainable AI approaches, e.g., SHAP-LIME fusion, can give an additional layer of interpretability to interested parties. Additionally, datasets with a variety of operating conditions and aging cycles will add to generalization power. Improved accuracy can be gained through Advanced architectures such as Transformers for time series and Physics-Informed Neural Networks (PINNs), which allow for domain knowledge to be measured. Lastly, embedding this system into a real-time BMS would look after predictive maintenance and lead to safer and efficient energy storage operations, thereby living batteries in industrial applications.

References

[1] Alamin K.S.S., Chen Y., MacIi E., Poncino M., Vinco S. (2022) A Machine Learning-Based Digital Twin for Electric Vehicle Battery Modeling, *IEEE Int. Conf. Omni-Layer Intell. Syst.*, 2022, COINS 2022. https://doi.org/10.1109/COINS54846.2022.9854960.

[2] Bandara T.R., Halgamuge M.N. (2022) Modeling a Digital Twin to Predict Battery Deterioration with Lower Prediction Error in Smart Devices: From the Internet of Things Sensor Devices to Self-Driving Cars, *IECON Proc. (Ind. Electron. Conf.)*, 2022, October, https://doi.org/10.1109/IECON49645.2022.9968677.

[3] Branco C.T.N.M., Fontanela J.M. (2024) A Design Methodology to Employ Digital Twins for Remaining Useful Lifetime Prediction in Electric Vehicle Batteries, *SAE Tech. Pap.*, January, https://doi.org/10.4271/2023-36-0132.

[4] K. V. G. Rao *et al.*, "Microgrid with, vehicle-to-grid and grid-to-vehicle technology for DC fast charging topology," *Renew. Energy Plug-In Electr. Veh.*, pp. 45–57, 2024, https://doi.org/10.1016/B978-0-443-28955-2.00004-4.

[5] Jafari S., Byun Y.C. (2022) Prediction of the Battery State Using the Digital Twin Framework Based on the Battery Management System, *IEEE Access*, 10, 124685–124696. https://doi.org/10.1109/ACCESS.2022.3225093.

[6] Kim G., Kang S., Park G., Min B.C. (2023) Electric Vehicle Battery State of Charge Prediction Based on Graph Convolutional Network, *Int. J. Automot. Technol.*, 24(6), 1519–1530. https://doi.org/10.1007/s12239-023-0122-6.

[7] K. Venkata Govardhan Rao, M. K. Kumar, B. S. Goud, M. Bajaj, M. Abou Houran, and S. Kamel, "Design of a bidirectional DC/DC converter for a hybrid electric drive system with dual-battery storing energy," *Front. Energy Res.*, vol. 10, no. November, 2022, https://doi.org/10.3389/fenrg.2022.972089.

[8] Li H., Bin Kaleem M., Chiu I.J., Gao D., Peng J., Huang Z. (2024) An Intelligent Digital Twin Model for the Battery Management Systems of Electric Vehicles, *Int. J. Green Energy*, 21(3), 461–475. https://doi.org/10.1080/15435075.2023.2199330.

[9] M. Pushkarna *et al.*, "A new-fangled connection of UPQC tailored power device from wind farm to weak-grid," *Front. Energy Res.*, vol. 12, no. February, pp. 1–19, 2024, https://doi.org/10.3389/fenrg.2024.1355867.

[10] Nair P., et al. (2024) AI-Driven Digital Twin Model for Reliable Lithium-Ion Battery Discharge Capacity Predictions, *Int. J. Intell. Syst.*, 2024, https://doi.org/10.1155/2024/8185044.

[11] K. V. G. Rao and M. K. Kumar, "The harmonic reduction techniques in shunt active power filter when integrated with non-conventional energy sources," *Indones. J. Electr. Eng. Comput. Sci.*, vol. 25, no. 3, pp. 1236–1245, 2022, https://doi.org/10.11591/ijeecs.v25.i3.pp1236-1245.

[12] Sugashini R., Mangaiyarkarasi S.P. (2023) Hybrid Deep Learning Algorithm for the State of Charge Prediction of the Lithium-Ion Battery for Electric Vehicles, *Iran. J. Chem. Chem. Eng.*, 42(8), 2550–2560.

[13] Qian C., et al. (2023) A CNN-SAM-LSTM Hybrid Neural Network for Multi-State Estimation of Lithium-Ion Batteries, *SSRN*, https://doi.org/10.2139/ssrn.4574033.

[14] B. N. Reddy *et al.*, "Switched Quasi Impedance-Source DC-DC Network for Photovoltaic Systems," *Int. J. Renew. Energy Res.*, vol. 13, no. 2, pp. 681–698, 2023, [Online]. Available: https://www.ijrer.org/ijrer/index.php/ijrer/article/view/14097.

[15] Rajesh P.K., Soundarya T., Jithin K.V. (2024) Driving Sustainability – The Role of Digital Twin in Enhancing Battery Performance for Electric Vehicles, *J. Power Sources*, 604, 234464. https://doi.org/10.1016/j.jpowsour.2024.234464.

[16] Renold A.P., Kathayat N.S. (2024) Comprehensive Review of Machine Learning, Deep Learning, and Digital Twin Data-Driven Approaches in Battery Health Prediction of Electric Vehicles, *IEEE Access*, 12, 43984–43999. https://doi.org/10.1109/ACCESS.2024.3380452.

[17] Vidal C., Malysz P., Kollmeyer P., Emadi A. (2020) Machine Learning Applied to Electrified Vehicle Battery State of Charge and State of Health Estimation: State-of-the-Art, *IEEE Access*, 8, 52796–52814. https://doi.org/10.1109/ACCESS.2020.2980961.
[18] Simonyan K., Zisserman A. (2015) Very Deep Convolutional Networks for Large-Scale Image Recognition, *3rd Int. Conf. Learn. Represent. (ICLR 2015) – Conf. Track Proc.*
[19] Wu B., Widanage W.D., Yang S., Liu X. (2020) Battery Digital Twins: Perspectives on the Fusion of Models, Data, and Artificial Intelligence for Smart Battery Management Systems, *Energy AI*, 1, 100016. https://doi.org/10.1016/j.egyai.2020.100016.
[20] Zhao Y., Wang Z., Shen Z.J.M., Sun F. (2021) Data-Driven Framework for Large-Scale Prediction of Charging Energy in Electric Vehicles, *Appl. Energy*, 282. https://doi.org/10.1016/j.apenergy.2020.116175.