



Systematic Identification of Key Prognostic Features for Early-Stage Lithium-Ion Battery Degradation Prediction

Meilia Safitri^{1,2}, Teguh Bharata Adji^{1*}, Adha Imam Cahyadi¹

¹ Department of Electrical and Information Engineering, Universitas Gadjah Mada, Yogyakarta 55281, Indonesia

² Department of Medical Electronics Technology, Universitas Muhammadiyah Yogyakarta, Yogyakarta 55183, Indonesia

Corresponding Author Email: adji@ugm.ac.id

Copyright: ©2025 The authors. This article is published by IETA and is licensed under the CC BY 4.0 license (<http://creativecommons.org/licenses/by/4.0/>).

<https://doi.org/10.18280/jesa.581114>

ABSTRACT

Received: 25 September 2025

Revised: 30 October 2025

Accepted: 23 November 2025

Available online: 30 November 2025

Keywords:

LIBs, battery prognostics, feature engineering, LSTM, early-cycle degradation prediction

Accurate prognostics of lithium-ion batteries (LIBs) are essential for ensuring safe, reliable, and cost-effective operation in electric vehicles and grid-scale storage systems. While sequence models such as Long Short-Term Memory (LSTM) networks have shown strong predictive capability, their performance depends critically on the quality and structure of input features. Although prior studies have demonstrated the utility of statistical descriptors extracted from early-cycle data, no systematic investigation has established a reproducible and standardized feature subset suitable for benchmarking prognostic models across multiple degradation markers. This study addresses this gap by extracting 70 statistical descriptors from the first 100 cycles of voltage and current data, transforming them using multicycle statistics, and ranking their prognostic value using Pearson correlation with respect to three targets: knee onset, knee point, and end-of-life (EoL). Based on this ranking, two reduced subsets were evaluated using a stacked LSTM model under cross-validation. The results show that a Top-40 subset consistently outperforms both the full 70-feature set and a smaller Top-20 set, particularly in EoL forecasting. Because this subset is systematically derived and validated across three prognostic markers, it provides a reproducible benchmark for researchers and a practical, computationally efficient feature basis for deployment in battery management systems.

1. INTRODUCTION

Lithium-ion batteries (LIBs) are the dominant energy storage technology for electric vehicles (EVs), consumer electronics, and grid-scale applications, owing to their high energy density, long cycle life, and favorable performance characteristics [1-3]. As the demand for electrification increases, the reliability, safety, and cost-effectiveness of LIBs have become critical considerations [4, 5]. One of the central challenges in this context is the gradual degradation of battery performance over time [6-8]. Understanding and predicting this degradation is essential for ensuring safety, optimizing operational strategies, and extending the usable life of energy storage assets [9, 10]. Accurate prognostics enable battery management systems (BMS) to anticipate critical events, such as the “knee point” of rapid capacity fade, and to plan maintenance or replacement before failures occur [11, 12].

Over the past two decades, several prognostic approaches have been proposed. Model-based methods rely on electrochemical or equivalent circuit models, offering interpretability and physical grounding; however, they require extensive parameterization, are computationally intensive, and often fail to generalize across real-world conditions. In contrast, data-driven methods leverage statistical patterns or machine learning algorithms applied to cycling data [12, 13]. These approaches have demonstrated strong predictive

performance, particularly when early cycle information is used for long-term forecasting. Recent advances in recurrent neural networks, particularly Long Short-Term Memory (LSTM) architectures, have demonstrated the ability to capture long-term temporal dependencies in degradation trajectories [14, 15]. However, as several studies have emphasized, predictive accuracy is not determined by algorithmic sophistication alone [16-18]. However, it depends critically on the quality, relevance, and diversity of the input features.

Feature engineering has emerged as a key enabler of reliable prognostics in recent years. Approaches such as statistical descriptors [17], incremental capacity analysis [19, 20], principal component analysis (PCA) [21, 22], and autoencoder-based representations [23] have been extensively investigated. Each involves trade-offs among interpretability, dimensionality reduction and generalizability. For example, incremental capacity features are closely tied to electrochemical mechanisms but are sensitive to measurement noise; PCA offers compact representations but sacrifices interpretability; and autoencoders automatically extract latent variables but require large datasets. Studies by Severson et al. [16] and Attia et al. [17] demonstrated that simple statistical descriptors of early cycle voltage curves can strongly predict cycle life; however, no systematic ranking or reproducible benchmark subset has been established. Similarly, although descriptors such as variance, skewness, and kurtosis have been

highlighted as useful indicators, there is no consensus on which feature combination optimally balances the accuracy, efficiency, and interpretability [24-26].

Another limitation is that most prior investigations focus narrowly on predicting cycle life or end-of-life (EoL) [16, 27], while neglecting intermediate prognostic markers such as knee onset and knee points. These markers are of practical importance for system operators: the knee onset signals the acceleration of capacity fading, and the knee point marks the transition to rapid degradation [28, 29]. Therefore, a comprehensive prognostic framework should address all three markers: knee onset, knee point, and EoL. However, systematic evaluations of multiple targets are lacking. Table 1 summarizes representative prior studies and contrasts them with the contributions of this study.

This study addressed the absence of a standardized, reproducible feature subset that balances accuracy, efficiency, and interpretability across all prognostic markers. Although prior studies have demonstrated the feasibility of early cycle feature-based prognostics, they have not provided a systematic framework that ranks, selects, and validates features simultaneously for knee onset, knee point, and EoL predictions.

To bridge this gap, we offer three key contributions. First, it establishes a systematically ranked and validated Top-40 feature subset, which, to the best of our knowledge, has not been previously reported, and is optimized for prognostic accuracy and efficiency. Second, it evaluates this subset across three prognostic targets, providing a comprehensive view of degradation dynamics. Third, it demonstrates that features beyond the Top-20 provide complementary information that is critical for robust long-term forecasts, offering reproducible benchmarks for both researchers and practitioners.

The remainder of this paper is organized as follows: Section 2 details the methodology, including feature engineering, correlation-based ranking, and the model design. Section 3 presents the results and discussion of the study. Section 4 concludes with the key findings and outlines the directions for future work.

2. METHODOLOGY

In this study, we developed a systematic framework to identify and validate prognostic features using early cycle lithium-ion battery data. The methodology transforms raw voltage–current time-series data into compact, physically meaningful descriptors, systematically ranks and selects the most informative features, and evaluates their predictive values for forecasting long-term degradation.

An overview of the proposed framework is illustrated in Figure 1. As shown, raw cycling data are processed along two complementary tracks: prognostic target labeling, where capacity degradation curves are smoothed and annotated with prognostic markers, and feature engineering, where statistical descriptors are extracted and condensed using the Multicycle Feature (MCF) transformation before undergoing correlation-based ranking. These two tracks converge in the model training and evaluation stage, where a stacked LSTM network is used to map the engineered features to the prognostic targets, and the predictive performance is quantified.

The following subsections provide a detailed account of each stage in this framework, beginning with the labeling of ground truth degradation markers.

Table 1. Comparative overview of prior studies and the contribution of this work

Study	Features Used	Targets	Limitation
Severson et al. [16]	Voltage statistical descriptors	EoL	No feature ranking; only lifetime prediction
Attia et al. [17]	Statistical learning descriptors	EoL	No standardized subset; no intermediate markers
Ibraheem et al. [30]	Voltage-curve signatures / statistical descriptors	Knee onset, Knee point, EoL	No standardized feature subset; high computational cost of signature extraction; evaluated on limited datasets
Zhou et al. [31]	Curvature-based features	Knee onset	Single marker only
This work	70 engineered features → Top-40 subset	Knee onset, knee point, EoL	Provides reproducible Top-40 subset validated across all three markers

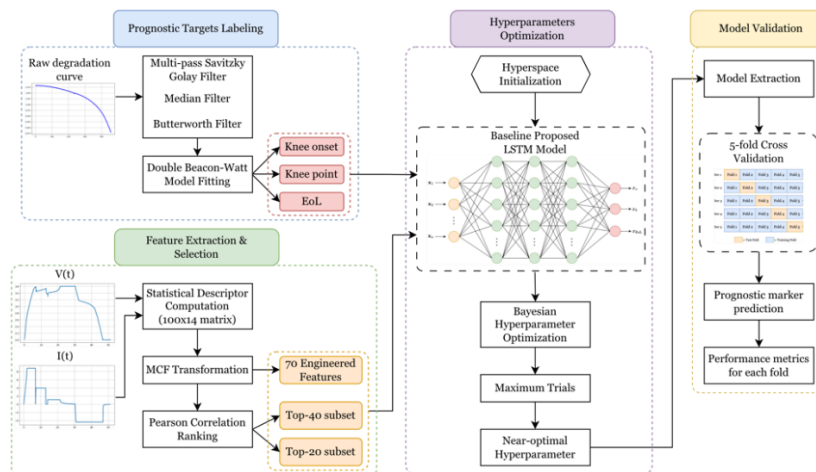


Figure 1. Overview of the proposed prognostic framework

2.1 Dataset and experimental setup

This study is based on the publicly available battery aging dataset introduced by Severson et al. [16], which has established itself as a reference benchmark for data-driven prognostic studies in lithium-ion battery health management. The dataset contains cycling data for 121 prismatic A123 Systems APR18650M1A cells. Each cell uses a lithium iron phosphate (LFP) cathode paired with a graphite anode, offering a nominal capacity of 1.1 Ah and a nominal voltage of 3.3 V. The cells were cycled in a temperature-controlled chamber maintained at 30°C to ensure uniform environmental conditions and to minimize confounding thermal effects. A 48-channel Arbin LBT potentiostat/galvanostat was used to precisely control and measure the electrochemical parameters, providing high-fidelity time-stamped current, voltage, and temperature data for each cycle.

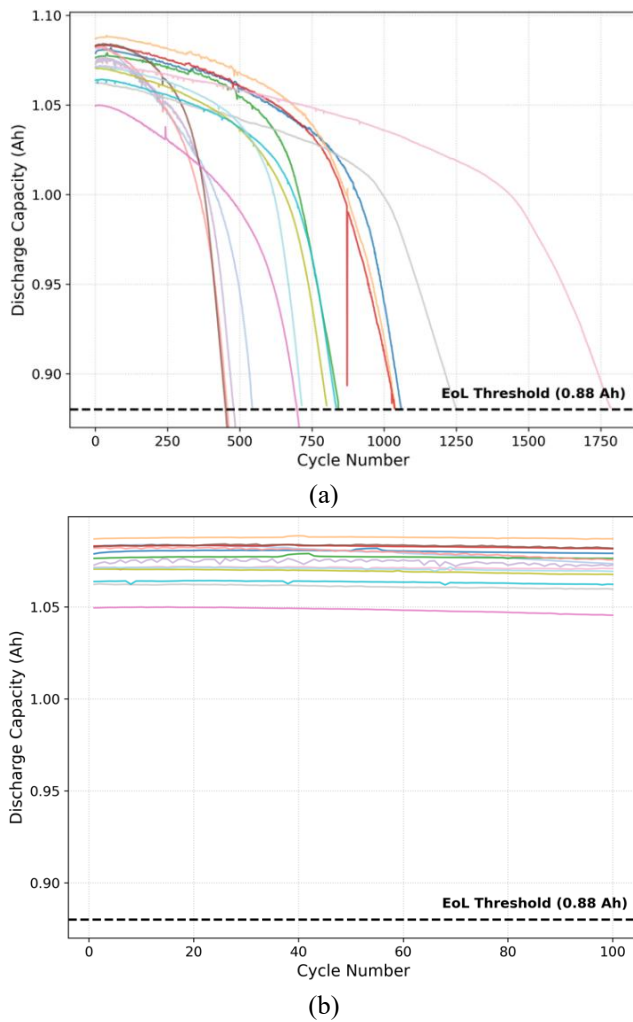


Figure 2. Discharge capacity fading curves from the dataset of Severson et al. [16]. (a) across multiple batteries and (b) from the first 100 cycles

The experimental protocol was explicitly designed to accelerate the aging process and emulate aggressive fast charging regimes, such as those anticipated in high-performance electric vehicles. Charging was performed using either a one- or two-step constant-current (CC) fast charge up to an 80% state of charge (SOC), followed by a constant-current–constant-voltage (CC–CV) charging phase to top off the cell. Discharge operations were performed at a constant

rate of 4 C until the voltage reached 2.0 V. This combination of aggressive charging and high-rate discharging produces a wide range of degradation trajectories and cycle life outcomes, ranging from early failures at ~150 cycles to exceptionally durable cells exceeding 2300 cycles before reaching an end-of-life (EoL) threshold of 80% nominal capacity.

The aggressive protocol not only accelerated capacity fading but also introduced diverse degradation mechanisms, including lithium plating, electrolyte decomposition, and structural changes in the active materials. These complex processes result in significant heterogeneity in the lifetime, making the dataset ideal for evaluating whether early cycle features can predict long-term performance. The availability of complete high-resolution time-series data for the first 100 cycles is particularly important because it enables the extraction of subtle diagnostic signals long before conventional indicators, such as direct capacity loss, become apparent.

Figure 2 illustrates two perspectives on the capacity-fade behavior of a subset of representative cells.

2.2 Prognostic targets labeling

The accurate identification of degradation markers is essential for supervised prognostics. In this study, three targets were defined: knee onset, marking the initial deviation from the nominal capacity trend; the knee point, representing the cycle of maximum acceleration in degradation; and end-of-life (EoL), defined as the first cycle at which the capacity falls below 80% of the nominal value.

Direct identification from raw capacity curves is challenging because of measurement noise, sensor drift, and stochastic degradation behavior. To obtain reliable targets, the capacity trajectories were processed using a multistage filtering pipeline consisting of a Savitzky–Golay filter (to preserve curvature), median filter (to remove spikes), and Butterworth filter (to suppress high-frequency noise). The smoothed curves were then fitted using the Double Bacon–Watts model, enabling consistent detection of the knee onset and knee point through its smooth, piecewise-transition structure. The EoL indices were extracted directly from the filtered curve using the 80% threshold criterion.

This labeling procedure does not claim to establish an absolute ground truth but rather provides operational prognostic targets that are reproducible, consistent, and sufficiently precise for evaluating the predictive capabilities of engineered features and learning models.

2.3 Comprehensive feature extraction and selection

The objective of this stage was to transform the raw high-resolution cycling data from the first 100 cycles into a compact and physically meaningful representation that maximized prognostic accuracy. Unlike prior studies that applied multicycle summarization to a limited set of variables, this study extends the framework to a broad set of statistical descriptors derived from both voltage and current signals and integrates them into a systematically ranked benchmark-feature subset.

2.3.1 Statistical descriptor computation

To characterize the electrochemical behavior within each cycle, a set of statistical descriptors was extracted from the voltage (V) and current (I) time-series data. For each of the first 100 cycles ($i \in \{1, \dots, 100\}$), standard statistical moments

were computed, including the minimum, maximum, mean (μ), and variance (σ^2), along with higher order descriptors such as skewness and kurtosis. In addition, the extremal values of the first-order derivatives were included to capture rapid changes in the voltage and current.

For the voltage, an eight-dimensional feature vector was generated for each cycle:

$$\mathbf{f}_V(i) = [V_{\min} \ V_{\max} \ \mu_V \ \sigma_V^2 \ \text{skew}_V \ \text{kurt}_V \ \frac{dV}{dt}_{\min} \ \frac{dV}{dt}_{\max}]^T \quad (1)$$

For the current, a six-dimensional vector is defined as

$$\mathbf{f}_I(i) = [I_{\min} \ I_{\max} \ \mu_I \ \sigma_I^2 \ \text{skew}_I \ \text{kurt}_I]^T \quad (2)$$

Concatenating these vectors yields a 14-dimensional feature vector for each cycle:

$$\mathbf{F}(i) = [\mathbf{f}_V(i)^T \ \mathbf{f}_I(i)^T]^T \in \mathbb{R}^{14} \quad (3)$$

Stacking $\mathbf{F}(i)$ across all 100 cycles produced a 100×14 feature matrix that compactly represented the early life electrochemical behavior of each cell. This served as the input for subsequent MCF transformations.

2.3.2 MCF transformation

Although the 100×14 feature matrix captured rich early life dynamics, its dimensionality was high and prone to redundancy. To condense the temporal information, each feature sequence was transformed into five summary statistics describing its level, drift, and curvature over the first 100 cycles.

For feature sequences $\mathbf{F}(i)$, $i \in \{1, \dots, j\}$, $j = 100$, the MCF parameters are defined as follows:

$$f_0 = \text{median}\{\mathbf{F}(i)\}_{i=1}^{10} \quad (4)$$

$$f_{\frac{j}{2}} = \text{median}\{\mathbf{F}(i)\}_{i=\frac{j}{2}-10}^{\frac{j}{2}+10} \quad (5)$$

$$f_j = \text{median}\{\mathbf{F}(i)\}_{i=j-10}^j \quad (6)$$

$$f_{j0} = f_j - f_0 \quad (7)$$

$$f_{\text{diff}} = f_j - 2f_{\frac{j}{2}} - f_0, \quad (8)$$

where, f_0 , $f_{j/2}$, and f_j denote early-, mid-, and late-cycle levels, respectively; f_{j0} captures the drift; and f_{diff} quantifies the curvature across the sequence. When applied to the 14 descriptors, this yielded 70 engineered features per cell. While MCF has been used in earlier studies, its application here to a broad, systematically designed statistical feature set represents an important extension, enabling both interpretability and computational efficiency of the model.

2.3.3 Correlation-based feature selection

To evaluate the relevance of the engineered features, we used the Pearson Correlation Coefficient (PCC) to rank their association with the three prognostic targets (knee onset, knee point, and EoL):

$$\rho_{x,y} = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \cdot \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}} \quad (9)$$

where, x denotes the feature vector, y is one of the prognostic targets, and N is the number of samples. Features were ranked by their absolute correlation values, $|\rho_{x,y}|$, and two subsets were defined: a Top-40 set, prioritizing predictive power while maintaining diversity, and a Top-20 set, emphasizing efficiency for resource-constrained applications.

Although Pearson's correlation was used as the primary ranking method owing to its interpretability and reproducibility, additional nonlinear analysis using Mutual Information (MI) was conducted to confirm ranking stability.

2.3.4 Feature normalization

Prior to model training, the selected features were normalized to reduce their sensitivity to outliers. Robust scaling was applied as follows:

$$x' = \frac{x - \text{median}_{\text{train}}(x)}{\text{IQR}_{\text{train}}(x)}, \text{IQR} = Q_3 - Q_1 \quad (10)$$

where, Q_3 and Q_1 are the first and third quartiles, respectively. This method is more resilient than min-max scaling for skewed distributions, which are commonly observed in early cycle descriptors. To prevent division by zero when $\text{IQR} \approx 0$, a small $\varepsilon = 10^{-6}$ was added. The targets were scaled for numerical stability, with inverse transformations applied before computing the evaluation metrics (i.e., MAPE and RMSE).

To ensure clarity in the subsequent analysis, each engineered feature is denoted using the convention descriptor_fstat, where the first part specifies the statistical descriptor (e.g., Vmax, Vmin, σ^2_V), and the subscript indicates the MCF statistic (f_0 , $f_{j/2}$, f_j , f_{j0} , and f_{diff}). For example, Vmax_fj refers to the late-cycle median value of the maximum voltage, and $\sigma^2_V f_{\text{diff}}$ denotes the curvature of the voltage-variance trajectory. The complete naming convention is summarized in Table 2, which provides a consistent mapping between raw electrochemical signals, their statistical characterization, and temporal patterns captured using the MCF.

2.4 Prognostic model architecture

To rigorously evaluate the predictive value of the engineered feature sets, we implemented a stacked LSTM network as the core prognostic model. LSTMs are particularly suited for battery health forecasting because their gated memory cells capture long-term temporal dependencies, allowing them to model gradual degradation patterns that unfold over hundreds of cycles. Unlike most prior studies that predict only end-of-life (EoL), our architecture was explicitly designed to forecast three prognostic targets (knee onset, knee point, and EoL) within a single framework.

2.4.1 Network design rationale

The proposed network was designed to balance the expressive power with computational efficiency. Stacking multiple LSTM layers enables the model to capture temporal dynamics at different scales; lower layers represent cycle-to-cycle fluctuations in the engineered features, whereas deeper layers encode slower degradation trends that unfold across

hundreds of cycles [32]. This hierarchical representation reflects the multiscale nature of lithium-ion battery aging. Importantly, by training the network on systematically ranked feature subsets (Top-40 and Top-20), this study directly

examined how compact yet diverse feature sets influence the architectural depth and complexity required for an accurate prognosis.

Table 2. Feature naming convention for statistical descriptors and MCF parameters

Descriptor	Meaning (Per Cycle)	MCF Suffix	Meaning (Across Cycles)	Example
Vmax	Maximum voltage	f ₀	Early-cycle level (first 10)	Vmax_f0 = early median of max voltage
Vmin	Minimum voltage	f _{j/2}	Mid-cycle level (around midpoint)	Vmin_fj/2 = midpoint median of min voltage
σ ² _V	Voltage variance	f _j	Late-cycle level (last 10)	σ ² _V_fj = late median of variance
skew_I	Current skewness	f _{j0}	Drift (late – early)	skew_I_fj0 = drift in skewness
kurt_V	Voltage kurtosis	f _{diff}	Curvature indicator	kurt_V_fdiff = curvature of kurtosis

2.4.2 Multi-head output structure

The central novelty of the proposed model lies in its multihead output design. The output from the final LSTM layer was first passed through a shared dense representation that captured the patterns common to all targets. This shared representation was then branched into three independent single-unit dense layers, each dedicated to one prognostic marker. This design enables parameter sharing across tasks while allowing task-specific specialization, ensuring that the knee onset, knee point, and EoL are predicted consistently within a single network. Figure 3. illustrates the overall architecture, from the 1×70 engineered feature input through the stacked LSTM encoder, dropout layers, shared dense layer, and multi-head outputs. This multi-head design enables parameter sharing across tasks while providing specialized predictions for each prognostic marker.

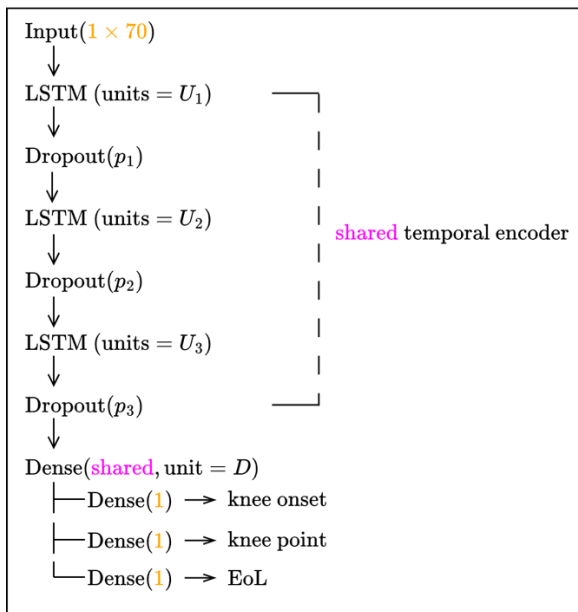


Figure 3. Stacked LSTM prognostic architecture with dropout, shared dense layer, and multi-head outputs

2.4.3 Training and optimization strategy

The model training strategy was designed to ensure robustness against noise, efficient hyperparameter exploration, and stable convergence of the model. We adopted the Huber loss function, which combines the sensitivity of the mean squared error with the robustness of the mean absolute error, thereby providing resilience to sporadic anomalies commonly observed in battery degradation data. Hyperparameters, including the number of stacked LSTM layers (2–4), units per layer (128–512), dropout rates (0.20–

0.50), dense layer width (32–128), and learning rate (10^{-3} – 10^{-2}), were tuned using Bayesian Optimization [33]. This probabilistic search efficiently navigates the high-dimensional design space and avoids manual bias.

Training was performed with the Adam optimizer [34] and a batch size of 32, augmented with two stabilization mechanisms: EarlyStopping, to terminate training when the validation loss stagnated for 10 epochs, and ReduceLROnPlateau, to adaptively lower the learning rate when progress slowed. Together, these choices ensured that the model converged efficiently, avoided overfitting, and generalized consistently across the folds.

2.5 Evaluation protocol

To ensure that the proposed framework was evaluated fairly and rigorously, we employed a training and validation strategy designed to test both predictive performance and generalizability. A five-fold cross-validation procedure was adopted: the dataset was partitioned into five equally sized subsets, and in each round, four folds were used for training, whereas the remaining fold served as the test set. This process was repeated until each fold was used as the test set. Importantly, all preprocessing steps, including robust scaling, were fitted only on the training folds and then applied to the held-out fold, thereby preventing information leakage and ensuring a realistic assessment.

The performance of all three prognostic targets, namely, knee onset, knee point, and end-of-life (EoL), was evaluated using two complementary error metrics. The Mean Absolute Percentage Error (MAPE),

$$\text{MAPE} = \frac{100\%}{n} \sum_{i=1}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \quad (11)$$

Provides an interpretable measure of the relative error across samples, allowing for comparison across targets with different magnitudes. The Root Mean Squared Error (RMSE),

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (12)$$

Quantifies the absolute deviations between the predicted and true values, with a higher penalty for larger errors. Together, these two metrics balance interpretability and sensitivity to severe mispredictions.

The final performance scores were computed as the average of the MAPE and RMSE across the five folds. A method was

considered effective if it achieved consistently low errors across the folds and for all three prognostic markers. In this study, these evaluation criteria not only test the accuracy of the predictions but also serve to demonstrate the robustness of the feature engineering framework under different levels of feature compactness (full set, Top-40, and Top-20 subsets).

The complete evaluation protocol is illustrated in Figure 4. Each row corresponds to one cross-validation iteration: blue blocks represent training folds, the orange block denotes the test fold, and the green box indicates the evaluation with the MAPE and RMSE. The final performance values were computed as the average of the MAPE and RMSE across all five iterations. A method is considered effective if it achieves consistently low errors across folds and all prognostic targets, thereby demonstrating robustness to data partitioning and reliability of the engineered feature subsets.

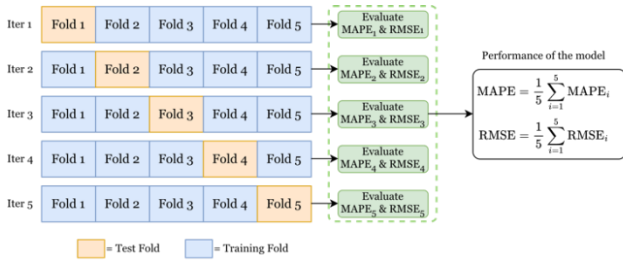


Figure 4. Evaluation protocol for five-fold cross-validation

3. RESULTS AND DISCUSSION

3.1 Feature ranking and selection

The prognostic significance of the engineered features was first examined using Pearson's correlation coefficient (PCC) analysis, which quantified the linear association between each of the 70 multicycle descriptors and three prognostic targets: knee onset, knee point, and end-of-life (EoL). The correlation structure is illustrated in Figure 5 as a heatmap, while Table 3 groups the Top-40 ranked features into statistical categories and highlights their representative forms and prognostic roles. Together, these results reveal distinct patterns in which different descriptors capture the long-term degradation behavior.

Table 3. Categorization of the Top-40 prognostic features ranked by absolute PCC

Descriptor Type	Selected Features (Rank Order)
Voltage Variance	$\sigma^2_V_fj/2$, $\sigma^2_V_fj$, $\sigma^2_V_fdiff$, $\sigma^2_V_f0$, $\sigma^2_V_fj0$
Voltage Maximum	V_{max_fj} , V_{max_f0} , $V_{max_fj/2}$, V_{max_fdiff}
Voltage Kurtosis	$kurt_V_fj/2$, $kurt_V_fj$, $kurt_V_fdiff$, $kurt_V_f0$
Current Kurtosis	$kurt_I_fdiff$, $kurt_I_f0$, $kurt_I_fj/2$, $kurt_I_fj$
Voltage Skewness	$skew_V_fj/2$, $skew_V_fdiff$, $skew_V_fj$, $skew_V_f0$
Voltage Derivative	dV/dt_max_f0 , dV/dt_max_fdiff , $dV/dt_max_fj/2$, dV/dt_max_fj , dV/dt_min_fj , dV/dt_min_fj0
Current Variance	$\sigma^2_I_fj/2$, $\sigma^2_I_fdiff$, $\sigma^2_I_fj$, $\sigma^2_I_f0$
Current Mean	μ_I_fj , $\mu_I_fj/2$, μ_I_fdiff

Current Minimum	$I_{min_fj/2}$, I_{min_fj} , I_{min_f0}
Voltage Mean	μ_V_fdiff , $\mu_V_fj/2$

Features derived from voltage variance consistently occupied the highest ranks, particularly those reflecting mid- and late-cycle behaviors, such as $\sigma^2_fj/2$ and σ^2_fj . Their dominance is physically meaningful because the variance in the voltage profile reflects increasing electrochemical heterogeneity and internal resistance, which are well-recognized precursors to accelerated degradation [16]. The maximum voltage and higher-order moments, including kurtosis in both the voltage and current signals, closely followed the rankings. These descriptors are particularly valuable because they are sensitive to rare or anomalous events within the cycling profile, allowing them to serve as early prognostic indicators [17, 25].

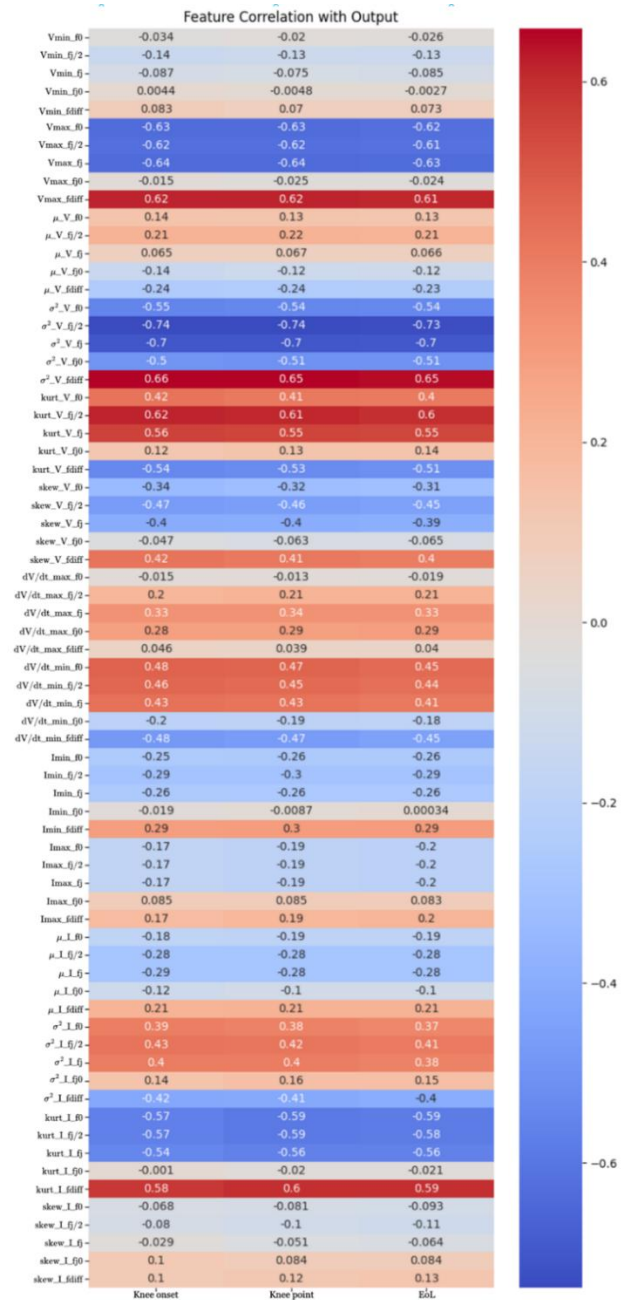


Figure 5. Heatmap of Pearson correlation between engineered features and prognostic targets (knee onset, knee point, and EoL)

The second tier of features demonstrated moderate correlations, including voltage skewness, current variance, and extremal derivatives of voltage. Although not as dominant as variance or kurtosis, these descriptors contribute to complementary perspectives. For instance, skewness quantifies the asymmetry in the charge–discharge dynamics, which may indicate imbalances in the charge transfer processes, whereas the extremal derivatives highlight the sharpness of the voltage transitions within a cycle.

Prior studies have shown that such moderately correlated features enrich the feature space and improve robustness when they are combined with stronger predictors.

At the lower end of the rankings, descriptors such as mean voltage, mean current, and minimum current consistently exhibited weak correlations with the prognostic markers. These measures largely capture steady-state operating points rather than dynamic changes in behavior, which explains their limited prognostic values. Their consistently low scores justified their exclusion from the optimized subsets.

To address the limitation that the PCC captures only linear associations, we complemented the analysis with a Mutual Information (MI)-based ranking of all 70 engineered features. The results showed a consistent 52.5% overlap between the MI-selected features and the PCC-based Top-40 subset across all three prognostic targets, as illustrated in the Venn diagram (Figure 6). This overlap indicates that many of the most influential descriptors exhibit both linear and nonlinear relevance, whereas the remaining MI-unique features highlight additional nonlinear dependencies not captured by the PCC alone. The substantial intersection between the two methods reinforces the robustness of the proposed Top-40 subset as a balanced and reproducible collection of features that preserves both linear and nonlinear prognostic information for subsequent analyses.

Taken together, the PCC analysis underscores the importance of maintaining feature diversity. While variance and kurtosis emerged as dominant signals, the presence of moderately correlated features provided complementary information that prevented overfitting to the narrow degradation aspects. This conclusion echoes prior findings on the value of combining descriptors from different statistical classes for robust prognostics [17, 35]. Based on this insight, two optimized feature sets were constructed: a Top-40 subset, designed to balance parsimony with predictive strength, and a Top-20 subset, designed to maximize computational efficiency. These subsets served as reproducible benchmarks for evaluating the model performance in subsequent sections.

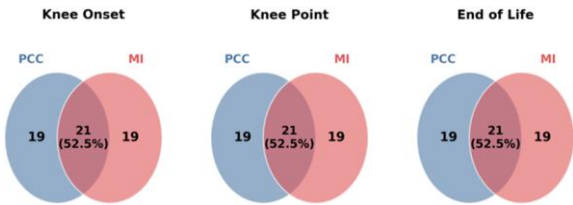


Figure 6. Overlap between MI-ranked features and the PCC-based Top-40 subset

3.2 Performance evaluation of feature subsets

3.2.1 Comparative accuracy across feature subsets

To examine the impact of feature subset selection on prognostic accuracy, the stacked LSTM model was trained and evaluated under three input configurations: the Full Set of 70

engineered features, the Top-40 subset derived from the PCC ranking, and the more compact Top-20 subset. The performance of all three prognostic targets—knee onset, knee point, and end-of-life (EoL)—was assessed using a five-fold cross-validation protocol. The model accuracy was quantified using the MAPE and RMSE. The aggregated results are reported in Table 4, illustrated in Figure 7 through comparative bar charts, and further examined in Figure 8, which presents the actual and predicted trajectories of the test set.

Table 4. Comparative performance of LSTM models trained on Full Set, Top-40, and Top-20 subsets

Target	MAPE (%)			RMSE (Cycles)		
	Full Set	Top-40	Top-20	Full Set	Top-40	Top-20
Knee Onset	12.02	11.73	11.07	86.77	80.45	80.134
Knee Point	11.51	11.28	10.68	116.57	104.77	111.012
EoL	12.17	10.59	10.69	133.33	116.32	124.428

A consistent trend emerged: the Top-40 subset outperformed the Full Set across nearly all prognostic targets, particularly in terms of EoL forecasting accuracy. Using all 70 features, the model achieved an average MAPE of 12.17% and an RMSE of 133.33 cycles for EoL prediction. In contrast, restricting the input to the Top-40 subset reduced these errors to 10.59% and 116.32 cycles, respectively. Comparable improvements were observed for the knee point, where the RMSE decreased from 116.57 to 104.77. These findings indicate that redundant or weakly informative features in the Full Set diluted the predictive signal and impaired generalization, a limitation well documented in both general machine learning [36] and battery prognostics [16]. By eliminating such features, the Top-40 subset provided a cleaner and more effective representation of the early cycle dynamics.

The comparison between the Top-40 and Top-20 subsets highlights the trade-off between parsimony and species diversity. For early prognostic markers, the Top-20 configuration achieved marginal improvements in MAPE (11.07% vs. 11.73% for knee onset; 10.68% vs. 11.28% for the knee point). However, these gains came at the cost of degraded EoL forecasting, with the RMSE increasing from 116.32 to 124.43 cycles. This decline indicates that the Top-20 subset suffers from information loss due to over-pruning, as features ranked between 21 and 40, although individually less correlated, collectively capture complementary information that is indispensable for robust long-term predictions. This observation is consistent with prior studies that emphasized the role of feature diversity in enhancing prognostic robustness [35].

Another important observation is the heightened sensitivity of EoL prediction to feature pruning. While the knee onset and knee point predictions remained relatively stable even under reduced feature sets, the EoL forecasts deteriorated markedly when the input dimensionality was restricted excessively. This reflects the fact that long-term degradation pathways evolve gradually and manifest through multiple complementary descriptors, thus requiring a broader feature base for accurate predictions.

3.2.2 Efficiency and convergence behavior

Beyond predictive accuracy, evaluating the computational footprint of the proposed feature subsets is essential for assessing their suitability for deployment in resource-

constrained Battery Management Systems (BMS). For the Top-40 feature configuration, which emerged as the best-performing subset, the stacked LSTM network contained approximately 1.36 million trainable parameters (≈ 5.7 MB of memory) and required approximately 57.8 thousand FLOPs

for each inference pass. This model size is well within the computational limits of modern embedded microcontrollers and edge inference accelerators typically used in BMS applications.

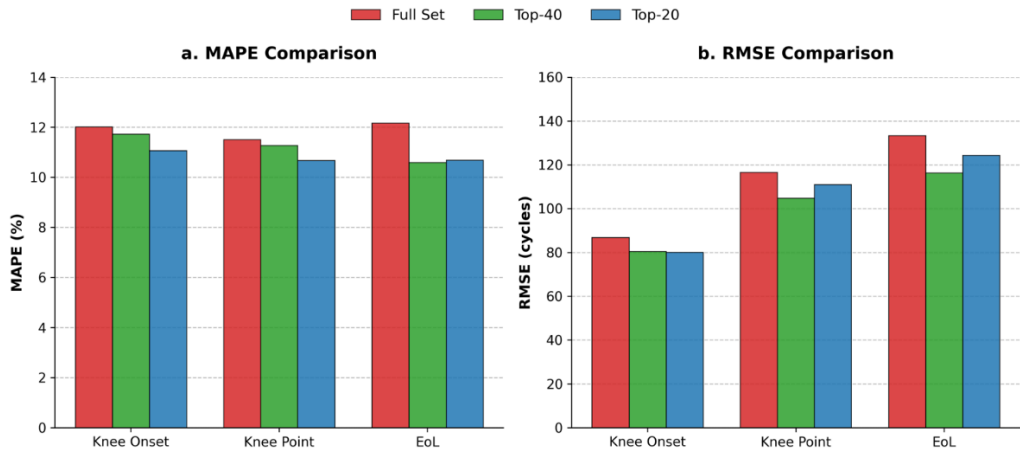


Figure 7. MAPE comparison of LSTM performance across feature subsets. The Top-40 subset demonstrated a superior balance between accuracy and efficiency

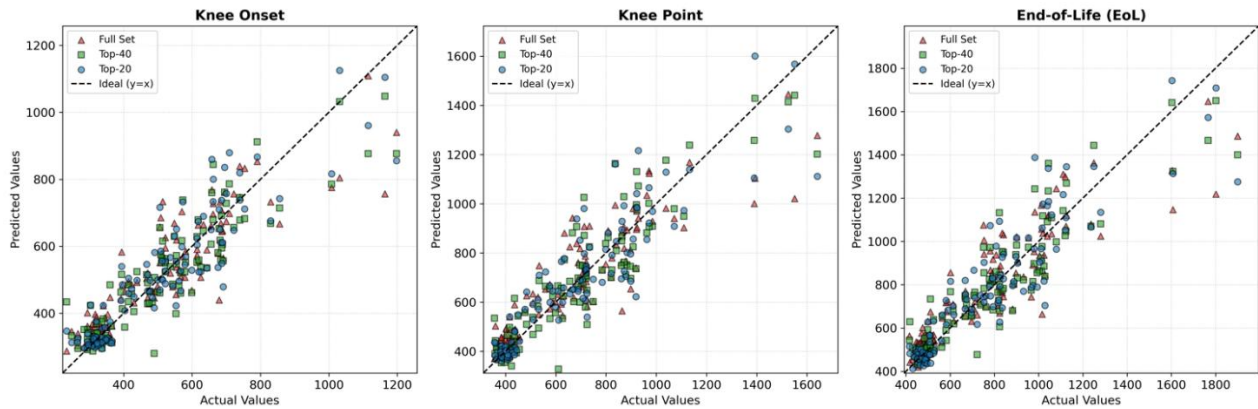


Figure 8. Actual versus predicted values for knee onset, knee point, and end-of-life using the Full Set, Top-40, and Top-20 feature subsets

The training and inference efficiencies were quantified using runtime measurements on a Google Colab T4 GPU with a batch size of 32 (Table 5). Although the Full Set converged the fastest (0.09 min) because it required only 14 epochs, this speed came at the expense of lower predictive accuracy. The Top-40 subset required a moderate 0.16 min to converge (28 epochs) but consistently achieved the best accuracy and the most stable convergence behavior among all configurations. In contrast, the smaller Top-20 subset exhibited the longest total runtime (0.26 min), caused by slower per-epoch validation steps and prolonged training (34 epochs), while also yielding weaker performance in the EoL prediction.

These findings illustrate that reducing dimensionality does not necessarily yield lower computational costs; under-pruned subsets may force the network to learn more abstract representations and converge more slowly. Conversely, the Top-40 subset provides a well-balanced trade-off: it reduces redundancy relative to the Full Set, avoids the over-pruning effects observed in the Top-20 subset, and delivers superior predictive performance with only a modest increase in training time. The representative training–validation loss curves shown in Figure 8 further confirm the stability and reproducibility of

this configuration across the folds.

Table 5. Training runtime comparison across features subsets (Google Colab T4, batch size = 32)

Configuration	Input Dimension	Avg. Epoch Time	Epochs to Convergence	Total Time (min)
Full Set	70	0.378	14	0.09
Top-40	40	0.338	28	0.16
Top-20	20	0.457	34	0.26

Figure 9 presents the representative training–validation loss curves for the Top-40 configuration (Fold 1), which illustrate smooth convergence and close alignment between the two trajectories. This agreement indicates that the model does not overfit and that the regularization mechanisms—dropout, EarlyStopping, and learning-rate scheduling—effectively stabilize the training. Similar loss patterns were observed across the other folds, reinforcing the reproducibility and robustness of the proposed framework.

3.3 Discussion

The results presented in Sections 3.1 and 3.2 demonstrate that systematic feature engineering is critical for advancing early stage lithium-ion battery prognostics. The proposed Top-40 subset consistently outperformed both the Full Set and Top-20 subsets, particularly in end-of-life (EoL) prediction, demonstrating that feature quality and diversity, rather than sheer quantity, determine the effectiveness of prognostic models. By striking a balance between parsimony and diversity, the Top-40 configuration provides both accuracy and efficiency, addressing a key challenge in developing reliable and scalable prognostic frameworks.

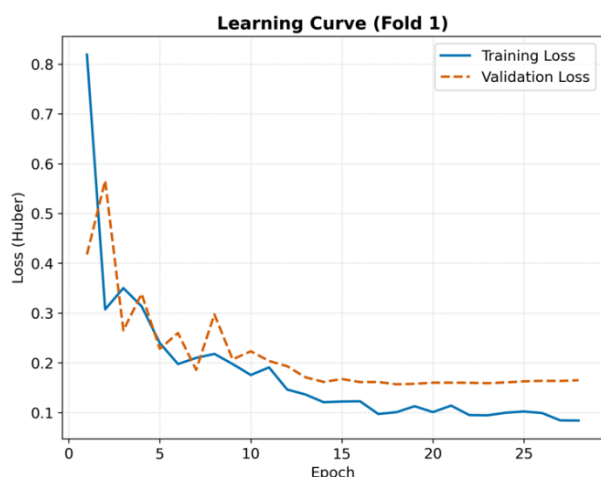


Figure 9. Training and validation loss curves for the Top-40 feature subset (fold 1)

These findings extend and refine prior contributions to this field. Severson et al. [16] established that voltage-derived statistical descriptors, such as variance, are strong predictors of cycle life, while Attia et al. [17] and other machine learning-enabled approaches [3] demonstrated the feasibility of using early-cycle data for long-term prediction. However, these efforts have not produced a standardized or systematically ranked feature subset that supports reproducibility across studies. Mechanistic work by Dubarry et al. [37] linked resistance growth to capacity fading, offering a physical explanation for the prominence of the variance-related descriptors. Grid-scale studies [27] further emphasize the utility of early prognostics in large-scale deployments; however, their insights have not been translated into reproducible cell-level benchmarks. In contrast, the present study contributes the first validated Top-40 feature subset that consistently predicts the knee onset, knee point, and EoL, providing a reproducible benchmark for research and practice.

Recent studies have explored end-to-end deep learning models that operate directly on raw cycling data [38, 39], achieving high accuracy but often requiring large datasets and substantial computational resources. Their limited interpretability also constrains their deployment in real-world Battery Management Systems (BMS). The present results show that a physics-informed feature engineering pipeline offers a complementary path: the proposed Top-40 subset reduces dimensionality, suppresses noise, and enables stable sequence modeling with LSTMs, effectively bridging raw-data neural approaches and interpretable physics-guided analysis.

This study makes three key contributions. First, it establishes a reproducible and systematically ranked Top-40 subset that balances accuracy, efficiency, and interpretability, which is a benchmark not previously available in the literature. Second, it validated this subset across three prognostic targets (knee onset, knee point, and EoL), extending prior work that focused mainly on lifetime prediction. Third, it shows that complementary descriptors beyond the Top-20 are essential for robust long-term forecasting, particularly for the EoL, underscoring the value of feature diversity rather than minimal-feature sets.

The broader implications of this study are twofold. For researchers, the Top-40 subset offers a standardized input space that facilitates fair benchmarking and improves reproducibility across studies. For practitioners, it reduces computational overhead without sacrificing accuracy, thereby supporting deployment in embedded or resource-constrained BMS environments. More generally, the findings reinforce that feature engineering and deep learning are not competing paradigms but complementary ones; physics-grounded descriptors can meaningfully strengthen temporal neural models.

Several limitations point to opportunities for future research. Although the dataset by Severson et al. [16] is widely regarded as a representative, high-resolution foundation for early cycle prognostics, it reflects a single chemistry (LFP) under controlled laboratory conditions. Validation of other chemistries (e.g., NMC and LCO) and real-world cycling regimes is essential. Additionally, five-fold cross-validation evaluates in-sample generalization but does not address cross-batch or cross-facility variabilities and degradation behavior may differ across datasets due to variations in manufacturing, cycling facilities, and operating protocols. Future studies should assess external datasets and their performance. The current feature ranking relies on a linear metric (PCC); nonlinear methods such as mutual information, recursive feature elimination, or LASSO could provide complementary insights. Finally, testing the Top-40 subset with alternative sequence architectures (GRUs, TCNs, transformers) and in online BMS environments with incomplete or noisy measurements will further clarify its practical robustness.

In summary, the proposed Top-40 feature subset advances early stage battery prognostics by improving the accuracy, reducing the computational cost, and enhancing the robustness. Its reproducibility and balanced design establish it as a practical benchmark for both the research community and real-world BMS deployment.

4. CONCLUSION

This study presents a systematic framework for identifying, ranking, and validating prognostic features from early cycle lithium-ion battery data. By applying multicycle transformations and correlation-based ranking, a reproducible Top-40 feature subset was established and evaluated using a stacked LSTM network.

The results demonstrate three central findings. First, the Top-40 subset consistently outperformed both the Full Set of 70 features and the more compact Top-20 subset, particularly in end-of-life prediction, confirming the importance of feature diversity for long-term forecasting. Second, the Top-40 configuration reduced redundancy and noise while maintaining computational tractability, yielding improved accuracy without incurring high training costs. Third, the

feature subset enhanced the interpretability and robustness of the sequence models, highlighting the complementary value of physics-informed feature engineering in conjunction with deep-learning architectures.

Because the Top-40 subset is systematically derived and validated across knee onset, knee point, and EoL, it serves not only as an empirical finding but also as a standardized and reproducible feature benchmark for future research. The balance between its accuracy, efficiency, and interpretability provides a strong foundation for developing scalable and reliable battery health monitoring systems.

ACKNOWLEDGMENT

This work was supported by the Indonesian Education Scholarship through the Center for Higher Education Funding and Assessment and the Indonesian Endowment Fund for Education (LPDP) under Decree Number 01882/J5.2.3/BPI.06/9/2022.

REFERENCES

- [1] Goodenough, J.B., Park, K.S. (2013). The Li-ion rechargeable battery: A perspective. *Journal of the American Chemical Society*, 135(4): 1167-1176. <https://doi.org/10.1021/ja3091438>
- [2] Tarascon, J.M., Armand, M. (2001). Issues and challenges facing rechargeable lithium batteries. *Nature*, 414(6861): 359-367. <https://doi.org/10.1038/35104644>
- [3] Megantoro, P., Prabowo, I.E., Shomad, M.A. (2020). Development of maximum power point tracking solar charge controller for 120 volt battery system at pandansimo hybrid power plant. In *Journal of Physics: Conference Series*, West Sumatera, Indonesia, p. 012047. <https://doi.org/10.1088/1742-6596/1471/1/012047>
- [4] Nykvist, B., Nilsson, M. (2015). Rapidly falling costs of battery packs for electric vehicles. *Nature Climate Change*, 5(4): 329-332. <https://doi.org/10.1038/nclimate2564>
- [5] Hasan, M.M., Haque, R., Jahirul, M.I., Rasul, M.G., Fattah, I.M.R., Hassan, N.M.S., Mofijur, M. (2025). Advancing energy storage: The future trajectory of lithium-ion battery technologies. *Journal of Energy Storage*, 120: 116511. <https://doi.org/10.1016/j.est.2025.116511>
- [6] Li, Y., Guo, W., Stroe, D.I., Zhao, H., et al. (2024). Evolution of aging mechanisms and performance degradation of lithium-ion battery from moderate to severe capacity loss scenarios. *Chemical Engineering Journal*, 498: 155588. <https://doi.org/10.1016/j.cej.2024.155588>
- [7] Rahman, T., Alharbi, T. (2024). Exploring lithium-Ion battery degradation: A concise review of critical factors, impacts, data-driven degradation estimation techniques, and sustainable directions for energy storage systems. *Batteries*, 10(7): 220. <https://doi.org/10.3390/batteries10070220>
- [8] Amifia, L.K., Kamali, M.A. (2023). Evaluating the battery management system's performance under levels of state of health (SOH) parameters. *Journal of Robotics and Control (JRC)*, 4(6): 808-817. <https://doi.org/10.18196/jrc.v4i6.20401>
- [9] Pugalenth, K., Park, H., Hussain, S., Raghavan, N. (2024). Remaining useful life estimation for lithium-ion batteries using physics-informed neural networks. In *2024 IEEE International Conference on Prognostics and Health Management (ICPHM)*, Spokane, WA, USA, pp. 67-73. <https://doi.org/10.1109/ICPHM61352.2024.10627352>
- [10] Wu, L., Fu, X., Guan, Y. (2016). Review of the remaining useful life prognostics of vehicle lithium-ion batteries using data-driven methodologies. *Applied Sciences*, 6(6): 166. <https://doi.org/10.3390/app6060166>
- [11] Khaleghi, S., Hosen, M.S., Van Mierlo, J., Berecibar, M. (2024). Towards machine-learning driven prognostics and health management of Li-ion batteries. A comprehensive review. *Renewable and Sustainable Energy Reviews*, 192: 114224. <https://doi.org/10.1016/j.rser.2023.114224>
- [12] Chen, K., Luo, Y., Long, Z., Li, Y., et al. (2025). Big data-driven prognostics and health management of lithium-ion batteries: A review. *Renewable and Sustainable Energy Reviews*, 214: 115522. <https://doi.org/10.1016/j.rser.2025.115522>
- [13] Liu, Q., Zheng, M., Li, P. (2023). A review of data-driven SOH and RUL estimation for lithium-ion batteries. In *2023 42nd Chinese Control Conference (CCC)*, Tianjin, China, pp. 8769-8774. <https://doi.org/10.23919/CCC58697.2023.10240526>
- [14] Khumprom, P., Yodo, N. (2019). A data-driven predictive prognostic model for lithium-ion batteries based on a deep learning algorithm. *Energies*, 12(4): 660. <https://doi.org/10.3390/en12040660>
- [15] Meng, H., Geng, M., Han, T. (2023). Long short-term memory network with Bayesian optimization for health prognostics of lithium-ion batteries based on partial incremental capacity analysis. *Reliability Engineering & System Safety*, 236: 109288. <https://doi.org/10.1016/j.res.2023.109288>
- [16] Severson, K.A., Attia, P.M., Jin, N., Perkins, N., et al. (2019). Data-driven prediction of battery cycle life before capacity degradation. *Nature Energy*, 4(5): 383-391. <https://doi.org/10.1038/s41560-019-0356-8>
- [17] Attia, P.M., Severson, K.A., Witmer, J.D. (2021). Statistical learning for accurate and interpretable battery lifetime prediction. *Journal of The Electrochemical Society*, 168(9): 090547. <https://doi.org/10.1149/1945-7111/ac2704>
- [18] Paulson, N.H., Kubal, J., Ward, L., Saxena, S., Lu, W., Babinec, S.J. (2022). Feature engineering for machine learning enabled early prediction of battery lifetime. *Journal of Power Sources*, 527: 231127. <https://doi.org/10.1016/j.jpowsour.2022.231127>
- [19] Liu, Y., Sun, G., Liu, X. (2023). Remaining useful life prediction of lithium-ion batteries based on peak interval features and deep learning. *Journal of Energy Storage*, 73: 109308. <https://doi.org/10.1016/j.est.2023.109308>
- [20] Long, B., Li, X., Gao, X., Liu, Z. (2019). Prognostics comparison of lithium-ion battery based on the shallow and deep neural networks model. *Energies*, 12(17): 3271. <https://doi.org/10.3390/en12173271>
- [21] Zeinali, S., Moqtaderi, H. (2024). Data-driven SOH early prediction of lithium-ion batteries. In *2024 12th RSI International Conference on Robotics and Mechatronics (ICRoM)*, Tehran, Iran, Islamic Republic of, pp. 173-

180.
https://doi.org/10.1109/ICRoM64545.2024.10903627
- [22] Pang, X.Q., Wang, Z.Q., Zeng, J.C., Jia, J.F., Shi, Y.H., Wen, J. (2019). Prediction for the remaining useful life of lithium-ion battery based on PCA-NARX. Transactions of Beijing Institute of Technology, 39(4): 406-412. https://doi.org/10.15918/j.tbit1001-0645.2019.04.012
- [23] Wang, Y., Zhang, T., Shi, Z., Cao, M., Zhu, W., Liu, Y. (2022). State of health estimation for Lithium-ion battery based on fusion features. In IET Conference Proceedings CP819, Sichuan, China, pp. 1462-1467. https://doi.org/10.1049/icp.2022.3073
- [24] Yang, F., Lu, Z., Tan, X., Tsui, K.L., Wang, D. (2024). Battery prognostics using statistical features from partial voltage information. Mechanical Systems and Signal Processing, 210: 111140. https://doi.org/10.1016/j.ymssp.2024.111140
- [25] Ibraheem, R., Strange, C., Dos Reis, G. (2023). Capacity and Internal Resistance of lithium-ion batteries: Full degradation curve prediction from Voltage response at constant Current at discharge. Journal of Power Sources, 556: 232477. https://doi.org/10.1016/j.jpowsour.2022.232477
- [26] Safitri, M., Adji, T.B., Cahyadi, A.I. (2025). Enhanced early prediction of Li-ion battery degradation using multicyle features and an ensemble deep learning model. Results in Engineering, 25: 104235. https://doi.org/10.1016/j.rineng.2025.104235
- [27] Apriowo, C.H.B., Hadi, S.P., Wijaya, F.D., Setyonegoro, M.I.B. (2024). Early prediction of battery degradation in grid-scale battery energy storage system using extreme gradient boosting algorithm. Results in Engineering, 21: 101709. https://doi.org/10.1016/j.rineng.2023.101709
- [28] Zhang, H., Altaf, F., Wik, T. (2024). Battery capacity knee-onset identification and early prediction using degradation curvature. Journal of Power Sources, 608: 234619. https://doi.org/10.1016/j.jpowsour.2024.234619
- [29] Kim, H., Kim, I., Kim, M., An, S., et al. (2025). Detection of the knee point in lithium-ion battery degradation using a state-of-charge-dependent parameter. Proceedings of the National Academy of Sciences, 122(23): e2424838122. https://doi.org/10.1073/pnas.2424838122
- [30] Ibraheem, R., Wu, Y., Lyons, T., Dos Reis, G. (2023). Early prediction of Lithium-ion cell degradation trajectories using signatures of voltage curves up to 4-minute sub-sampling rates. Applied Energy, 352: 121974. https://doi.org/10.1016/j.apenergy.2023.121974
- [31] Zhou, K.Q., Qin, Y., Yuen, C. (2023). Lithium-ion battery state of health estimation by matrix profile empowered online knee onset identification. IEEE Transactions on Transportation Electrification, 10(1): 1935-1946. https://doi.org/10.1109/TTE.2023.3265981
- [32] Hochreiter, S., Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8): 1735-1780. https://doi.org/10.1162/neco.1997.9.8.1735
- [33] Snoek, J., Larochelle, H., Adams, R.P. (2012). Practical bayesian optimization of machine learning algorithms. Advances in Neural Information Processing Systems, 25.
- [34] Adam, K.D.B.J. (2014). A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 1412(6): 980.
- [35] Roman, D., Saxena, S., Robu, V., Pecht, M., Flynn, D. (2021). Machine learning pipeline for battery state-of-health estimation. Nature Machine Intelligence, 3(5): 447-456. https://doi.org/10.1038/s42256-021-00312-3
- [36] Li, J., Cheng, K., Wang, S., Morstatter, F., Trevino, R.P., Tang, J., Liu, H. (2017). Feature selection: A data perspective. ACM Computing Surveys (CSUR), 50(6): 1-45. https://doi.org/10.1145/3136625
- [37] Dubarry, M., Truchot, C., Liaw, B.Y. (2012). Synthesize battery degradation modes via a diagnostic and prognostic model. Journal of Power Sources, 219: 204-216. https://doi.org/10.1016/j.jpowsour.2012.07.016
- [38] Strange, C., Dos Reis, G. (2021). Prediction of future capacity and internal resistance of Li-ion cells from one cycle of input data. Energy and AI, 5: 100097. https://doi.org/10.1016/j.egyai.2021.100097
- [39] Hong, J., Lee, D., Jeong, E.R., Yi, Y. (2020). Towards the swift prediction of the remaining useful life of lithium-ion batteries with end-to-end deep learning. Applied Energy, 278: 115646. https://doi.org/10.1016/j.apenergy.2020.115646

APPENDIX

A. Feature Naming Convention

Feature	Description
Vmin_f0	early median of min voltage
Vmin_fj/2	midpoint median of min voltage
Vmin_fj	late median of min voltage
Vmin_fj0	drift in min voltage
Vmin_fdiff	curvature of min voltage
Vmax_f0	early median of max voltage
Vmax_fj/2	midpoint median of max voltage
Vmax_fj	late median of max voltage
Vmax_fj0	drift in max voltage
Vmax_fdiff	curvature of max voltage
μ_V _f0	early median of mean voltage
μ_V _fj/1	midpoint median of mean voltage
μ_V _fj	late median of mean voltage
μ_V _fj0	drift in mean voltage
μ_V _fdiff	curvature of mean voltage
σ^2_V _f0	early median of voltage variance
σ^2_V _fj/2	midpoint median of voltage variance
σ^2_V _fj	late median of voltage variance
σ^2_V _fj0	drift in voltage variance
σ^2_V _fdiff	curvature of voltage variance
Kurt_V_f0	early median of voltage kurtosis
Kurt_V_fj/2	midpoint median of voltage kurtosis
Kurt_V_fj	late median of voltage kurtosis
Kurt_V_fj0	drift in voltage kurtosis
Kurt_V_fdiff	curvature of voltage kurtosis
Skew_V_f0	early median of voltage skewness
Skew_V_fj/2	midpoint median of voltage skewness
Skew_V_fj	late median of voltage skewness
Skew_V_fj0	drift in voltage skewness
Skew_V_fdiff	curvature of voltage skewness
dV/dt_max_f0	early median of max voltage derivative
dV/dt_max_fj/2	midpoint median of max voltage derivative
dV/dt_max_fj	late median of max voltage derivative
dV/dt_max_fj0	drift in max voltage derivative
dV/dt_max_fdiff	curvature of max voltage derivative

Feature	Description
dV/dt_min_f0	early median of min voltage derivative
dV/dt_min_fj/2	midpoint median of min voltage derivative
dV/dt_min_fj	late median of min voltage derivative
dV/dt_min_fj0	drift in min voltage derivative
dV/dt_min_fdiff	curvature of min voltage derivative
Imin_f0	early median of min current
Imin_fj/2	midpoint median of min current
Imin_fj	late median of min current
Imin_fj0	drift in min current
Imin_fdiff	curvature of min current
Imax_f0	early median of max current
Imax_fj/2	midpoint median of max current
Imax_fj	late median of max current
Imax_fj0	drift in max current
Imax_fdiff	curvature of max current
μ_I _f0	early median of mean current
μ_I _fj/1	midpoint median of mean current

Feature	Description
μ_I _fj	late median of mean current
μ_I _fj0	drift in mean current
μ_I _fdiff	curvature of mean current
σ^2_I _f0	early median of current variance
σ^2_I _fj/2	midpoint median of current variance
σ^2_I _fj	late median of current variance
σ^2_I _fj0	drift in current variance
σ^2_I _fdiff	curvature of current variance
Kurt_I_f0	early median of current kurtosis
Kurt_I_fj/2	midpoint median of current kurtosis
Kurt_I_fj	late median of current kurtosis
Kurt_I_fj0	drift in current kurtosis
Kurt_I_fdiff	curvature of current kurtosis
Skew_I_f0	early median of current skewness
Skew_I_fj/2	midpoint median of current skewness
Skew_I_fj	late median of current skewness
Skew_I_fj0	drift in current skewness
Skew_I_fdiff	curvature of current skewness