

Article

Integrated Predictive-Maintenance Framework for EV Batteries Using Short-Horizon *SoH* Forecasting, Degradation Warning, and Acceleration Risk Detection

Ch. Hadassa Parimala ¹, P. Srinivasa Varma ¹ , Ch. Paul Bakht Singh ¹  and Alagar Karthick ^{2,*}

¹ Department of Electrical and Electronics Engineering, Koneru Lakshmaiah Education Foundation, Green Fields, Vaddeswaram 522302, Andhra Pradesh, India

² Department of Electrical and Electronics Engineering, Saveetha School of Engineering, Saveetha Institute of Medical and Technical Sciences, SIMATS, Chennai 602105, Tamil Nadu, India

* Correspondence: karthicka.sse@saveetha.com or karthick.power@gmail.com

Abstract

Precision battery-health monitoring and rapid degradation detection are essential for improving the security, durability, and efficacy of electric vehicles (EVs). By incorporating short-term State-of-Health (*SoH*) forecasting, mid-term deterioration alarms, and degradation acceleration risk modeling into a temporally consistent machine learning architecture, this research suggests a hierarchical predictive-maintenance framework. The rolling-origin cross-validation approach is implemented to maintain the chronological order of the data and prevent any potential information leaks. The predictive core employs an ensemble learning approach that integrates Random Forest, Extremely Randomized Trees, and Histogram-Based Gradient Boosting. Validation-driven model blending and training only feature selection are implemented to improve generalizability. The one-hour *SoH* forecasting model for short-horizon monitoring exhibits exceptional accuracy in an assessment of health prediction, with an R^2 of 0.9254, an RMSE of 0.0033, and a MAPE of 0.32%. Early detection of anomalies and the provision of a seven-day degradation warning may be achieved by a proactive maintenance scheduling model with an area under the curve (AUC) of 0.7838 and a recall of 0.8205. In addition, the degradation acceleration risk module could identify rapid health decline with a robustness of 0.8796 and a precision–recall AUC of 0.7101 when operating under significant stress. Reliability in critical domains is demonstrated through validation using scenarios that simulate severe temperature and stress conditions. Achieving intelligent predictive maintenance of electric vehicle battery packs is now feasible due to the proposed multi-layer ensemble structure.



Academic Editor: Grzegorz Sierpiński

Received: 14 April 2026

Revised: 21 May 2026

Accepted: 25 May 2026

Published: 28 May 2026

Copyright: © 2026 by the authors.

Published by MDPI on behalf of the

World Electric Vehicle Association.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

Keywords: electric vehicle battery management; SDG 11; state-of-health (*SoH*) prediction; ensemble learning; predictive maintenance; degradation risk modeling; gradient boosting

1. Introduction

The rapid growth of electric mobility has intensified the need for reliable battery-health monitoring and predictive-maintenance strategies. Figure 1a illustrates the continuous increase in global electric vehicle adoption, where annual EV sales increased from nearly 14 million units in 2023 to approximately 17 million units in 2024. The figure also highlights the increasing contribution of EVs to total global vehicle sales, reflecting the accelerating transition toward electrified transportation systems. This

sustained expansion directly increases the operational dependence on high-performance lithium-ion battery systems. Correspondingly, Figure 1b presents the projected growth in global electric vehicle battery demand. The demand reached nearly 1 TWh in 2024 and is expected to surpass 3 TWh by 2030 under current policy and market development trends. The figure clearly demonstrates that battery usage is expanding at a much faster rate than conventional energy-storage growth, thereby increasing concerns related to battery degradation, safety, lifecycle reliability, and maintenance costs. Consequently, accurate State-of-Health (*SoH*) estimation and degradation prediction have become critical requirements for ensuring safe, economical, and reliable operation of modern electric vehicle fleets. The trends illustrated in Figure 1a,b further emphasize the necessity for advanced battery-health estimation and predictive diagnostic techniques. Among the different battery indicators, State of Health (*SoH*) is regarded as one of the most informative metrics for assessing aging, as it reflects capacity reduction, increased resistance, and a gradual decline in performance relative to the original condition. Past assessments have indicated that estimating the State of Health and Remaining Useful Life encounters challenges stemming from nonlinear aging processes, variability among cells, reliance on operating temperature, impacts of current rates, and the differences between laboratory cycling and real-world electric vehicle duty cycles [1–3]. Additional studies emphasized the importance of efficient battery-health management, integrating data-driven learning with robust temporal features, given that field data are often noisy, irregular, and significantly affected by user behavior and environmental stress [4–6]. The findings presented here are highly pertinent to electric vehicle applications, as similar battery chemistries can experience notably different degradation processes influenced by diverse driving, charging, and thermal conditions.




Figure 1. Cont.




Figure 1. (a) Global electric car sales and EV sales share in 2023 and 2024 based on IEA official statistics [7]. (b) Global EV battery demand in 2024 and projected demand in 2030 under the IEA STEPS scenario [7].

Recent developments in machine learning have significantly enhanced the assessment of battery health and the prediction of their performance. Recent review articles emphasize a clear transition from conventional model-based estimation to hybrid and data-driven methodologies. This encompasses methods like ensemble learning, deep neural networks, transfer learning, and predictions that account for uncertainty [5–9]. At the same time, considerable studies have shown that deep learning methods can evaluate battery health across various cell types without requiring extensive labeling of targets. Furthermore, information gathered from actual electric vehicle usage can improve health-estimation models, resulting in accuracy that is truly advantageous. The literature distinctly shows that accurate short-term tracking, degradation notifications, and long-term predictions pose different levels of difficulty. Short-term State-of-Health trajectories often show significant persistence, and recognizing degradation during stressful events requires models that can detect subtle but operationally critical changes in degradation dynamics.

While the estimation of battery health has improved considerably, notable practical gaps remain. A considerable portion of the literature addresses short-term State-of-Health regression, long-term lifetime forecasting, and anomaly detection as distinct tasks. Maintenance choices for electric vehicles rarely operate in isolation. Fleet operators and battery management systems require a comprehensive strategy for intelligence: accurate monitoring of current health, notifications for medium-term deterioration, and swift identification of accelerated aging patterns in stressful conditions [10,11]. The 7-day degradation warning is crucial for operations, as it offers a practical window for carrying out preventive maintenance, redistributing loads, adjusting thermal management, or altering charging strategies. Unlike long-term lifetime forecasting, which often involves unpredictability and the risk of extrapolation errors, a one-week warning period offers a practical balance between predictability and a timeframe that is beneficial for taking action. Furthermore, identifying the acceleration of degradation is essential for understanding nonlinear aging patterns that arise from elevated temperatures, higher current levels, or frequent deep cycling conditions commonly encountered in field operations, which are often not sufficiently represented in

laboratory data. Recent studies often emphasize the value of point-wise estimation accuracy, while they frequently neglect the importance of early warning capabilities that prioritize recall. When considering safety and reliability, failing to recognize a rapid degradation event can lead to more severe repercussions than simply issuing a precautionary alert. Therefore, it is crucial to concentrate on modeling that takes acceleration into account and optimizing thresholds that are sensitive to recall for successful implementation. The primary objective of this research is to develop a unified predictive-maintenance framework capable of simultaneously addressing three complementary battery-health intelligence tasks: short-horizon *SoH* forecasting, mid-term degradation warning, and accelerated aging risk detection. The specific novelty of this work is the development of a temporally consistent predictive-maintenance framework in which short-horizon *SoH* forecasting, seven-day degradation warning, and acceleration-sensitive risk detection are jointly organized as connected maintenance-intelligence tasks rather than as independent prediction problems. In contrast to conventional ensemble-based or single-task battery-health studies, the proposed framework combines validation-weighted ensemble forecasting, training-only feature selection, rolling-origin validation, and recall-oriented warning calibration to support deployment-oriented battery monitoring. Thus, the contribution is centered on converting battery-health estimation into an integrated early warning and risk-screening workflow for electric vehicle battery maintenance.

2. Materials and Methods

The current research has been organized into four methodological streams: a review-based consolidation of State-of-Health/Remaining Useful Life estimation techniques, feature-engineered machine learning models, deep temporal architectures for partial-charge or partial-load data, and practical frameworks for electric vehicles to improve transferability beyond laboratory cycling. Since 2021, there has been an increasing emphasis on temporal convolution, recurrent learning, transfer learning, interpretable hybrid models, and estimation informed by domain knowledge. Furthermore, most of the research has mostly focused on point-wise *SoH* estimates, whereas subjects like degradation alerts, acceleration-aware risk screening, and action-oriented maintenance lead times have garnered relatively little systematic examination. This discrepancy provided potential for frameworks that integrate short-horizon tracking, mid-horizon warning, and stress-sensitive risk identification as a unified predictive-maintenance issue rather than treating them as distinct estimating jobs. The comprehensive literature evaluation of the ongoing research initiatives is delineated in Table 1 accordingly.

Table 1. Comprehensive literature survey of the existing research works.

Ref.	Main Objective	Core Methodology	Datasets	Key Significance	Gap Carried Forward
[12]	<i>SOH</i> /RUL review	Review of ML degradation modeling	Public benchmark + experimental datasets	Established the theoretical and technological foundation for ML-based battery prognostics	Event-oriented maintenance warning was not systematized
[13]	EV/BESS health-estimation review	Comparative review	EV + BESS datasets	Clarified links between EV health estimation and storage relevance	Actionable warning horizons were not emphasized
[14]	ANN-based BMS review	Review of ANN architectures	ANN-based datasets	Consolidated neural approaches in battery management	Real-world maintenance deployment issues remained open

Table 1. Cont.

Ref.	Main Objective	Core Methodology	Datasets	Key Significance	Gap Carried Forward
[15]	Partial-load <i>SoH</i> estimation	Temporal convolutional network	Partial-load dataset	Demonstrated that raw partial-load data can support <i>SoH</i> estimation	Mid-horizon warning and acceleration risk were not addressed
[16]	High-accuracy <i>SoH</i> regression	Improved LSTM with health-indicator extraction	Cycling dataset	Shown the value of deep recurrent learning for nonlinear degradation	Transferability to field operation remained limited
[17]	Hybrid estimation	Incremental-capacity analysis + SVR	Incremental-capacity dataset	Combined physical feature extraction with machine learning	Scalability to large EV fleets was not established
[18]	Multi-feature deep temporal learning	Multi-feature extraction + TCN	Multi-feature dataset	Strengthened representation learning through richer health indicators	Operational warning logic remained absent
[19]	Sequence modeling for capacity/health	Graph-enhanced LSTM	Sequence dataset	Advanced automatic feature extraction for temporal degradation learning	Interpretability and decision support remained limited
[20]	Physics-aware data-driven estimation	Domain-knowledge-guided ML	Physics-informed dataset	Improved robustness by embedding physically meaningful indicators	Maintenance lead-time modeling was not the central target
[21]	Data scarcity and transferability	Transfer learning from synthetic to measured partial charging	Synthetic + real dataset	Addressed domain mismatch and data scarcity	A full predictive-maintenance layer was not constructed
[22]	Generalization across datasets	Transfer-stacking ensemble	Multi-dataset (cross-domain)	Improved cross-domain generalization in <i>SoH</i> estimation	Warning-risk integration was not pursued
[23]	Real-world EV application	Feature extraction + XGBoost-type regression on EV field data	Real-world EV dataset	Strengthened confidence in field-deployable <i>SoH</i> estimation	Event-based degradation screening remained secondary
[24]	New sensing modality for <i>SoH</i>	Mechanical–electrical features + ML	Multi-modal dataset	Extended <i>SoH</i> estimation beyond purely electrical signatures	Practical warning horizons were not evaluated
[25]	Feature optimization for <i>SoH</i>	PCA/MDS/filtering with LR, GPR, SVR, and LSTM comparisons	Feature-optimized dataset	Clarified the effect of feature optimization on model accuracy	Predictive-maintenance decision logic was not incorporated
[26]	Interpretable real-world EV estimation	Interpretable hybrid NGBost/CatBoost-type framework	Large-scale, real-world EV time-series dataset	Highlighted interpretability and uncertainty-aware estimation in field data	Joint modeling of short-horizon tracking, warning, and acceleration risk remained open

The comparative analysis presented in Table 1 clearly demonstrates the progressive transition of battery-health research from conventional State-of-Health estimation toward more operationally adaptive prognostic frameworks. Earlier studies primarily concentrated on improving prediction accuracy through deep learning, temporal convolution, transfer learning, and hybrid machine learning architectures. However, most reported methodologies addressed forecasting, degradation analysis, and anomaly detection as isolated tasks. It was further observed that limited attention has been given to integrated maintenance-oriented frameworks capable of simultaneously supporting short-term health monitoring,

actionable warning generation, and acceleration-sensitive degradation screening. The summarized comparison therefore establishes the necessity for a unified predictive-maintenance architecture that can transform battery-health estimation into practical maintenance decision support for electric vehicle applications. The literature has gradually transitioned from solely point-estimation frameworks to more transferable and operationally informed battery-health analytics. Initial consolidation studies primarily focused on outlining the methodological framework for State-of-Health (*SoH*) and Remaining Useful Life (*RUL*) estimation. These studies emphasized machine learning, neural network-based battery management, and the increasing demand for health indicators pertinent to electric vehicles [12–14]. Subsequently, greater focus was directed towards temporal deep learning models, such as temporal convolutional networks, graph-enhanced recurrent architectures, enhanced LSTM frameworks, and hybrid physics-assisted estimators [15–21]. Recent contributions have shown a growing focus on transferability, interpretability, and the practical deployment of electric vehicles. This is evident in advancements such as transfer-stacking ensembles, State-of-Health estimation based on field data, studies on feature optimization, and interpretable hybrid models [22–26]. Nonetheless, the scope of research examined has largely focused on the accuracy of State-of-Health (*SoH*) estimation, while integrated approaches for short-term health monitoring, degradation alerts, and risk detection that accounts for acceleration have been explored only to a limited extent. Although machine learning-based predictive maintenance for EV batteries has received considerable attention, most existing studies remain concentrated on individual tasks such as *SoH* estimation, *RUL* prediction, anomaly detection, or feature-based health modeling. These studies have improved prediction accuracy, but they do not fully address the operational requirement of converting battery-health data into maintenance decisions with defined warning horizons and acceleration-sensitive risk information. In practical EV battery management, maintenance decisions require three connected capabilities: reliable short-term health tracking, early identification of degradation within an actionable lead time, and rapid detection of accelerated aging under stress conditions. Treating these tasks independently may limit decision support because accurate *SoH* prediction alone does not necessarily indicate when intervention is required or whether degradation is entering a risky accelerated phase. Therefore, the problem addressed in this study is the lack of an integrated, temporally consistent predictive-maintenance framework that jointly links short-horizon *SoH* forecasting, 7-day degradation warning, and acceleration risk detection for EV battery systems. Accordingly, the contribution of this work is not the isolated use of machine learning for battery-health prediction but the integration of forecasting, warning, and acceleration-aware risk screening into a single time-ordered maintenance decision framework. Based on the identified research gap, this work proposes a unified predictive-maintenance framework designed to improve operational battery intelligence through integrated forecasting and warning mechanisms. The principal research contributions are outlined in the following section.

- A cohesive machine learning framework has been created that concurrently executes short-term *SoH* forecasting, provides 7-day degradation alerts, and identifies risk with consideration for acceleration. This cohesive framework connects ongoing health monitoring with anticipatory maintenance guidance in electric vehicle battery systems.
- A rolling-origin validation strategy has been adopted to maintain temporal causality and avoid data leakage, thereby ensuring realistic deployment. A blending mechanism that adapts has been integrated to improve generalization under different operational conditions.
- A degradation warning and acceleration modeling scheme centered on recall has been developed to facilitate the early detection of unusual aging under thermal and load

stress. This design enhances battery prognostics by shifting from mere estimation to actionable predictive maintenance that can be implemented in real operations.

This paper is organized in the following manner: Section 2 offers an in-depth review of the literature, emphasizing the latest advancements in State-of-Health estimation, degradation modeling, and data-driven battery prognostics. In Section 3, the problem is mathematically formulated, and the regression and classification goals that support short-horizon forecasting, degradation warning, and acceleration risk modeling are established. Section 4 outlines the multi-layer machine learning framework that has been implemented. It covers aspects such as temporal feature engineering, ensemble regression, probabilistic classification, recall-oriented threshold optimization, and stress-conditioned validation. Section 5 presents the simulated results along with a discussion based on various scenarios, while Section 6 offers concluding remarks and outlines potential directions for future research.

3. Problem Formulation and Mathematical Modeling

The operation of electric vehicle batteries produces ongoing multivariate data that reflect electrical, thermal, and load-related behaviors. The measurements are depicted as a feature space indexed by time, as shown in Equation (1):

$$\mathcal{X} = \left\{ \mathbf{x}_t \in \mathbb{R}^d \mid t = 1, 2, \dots, T \right\} \quad (1)$$

where $\mathbf{x}_t$ denotes the feature vector at time instant t , d represents the dimensionality of the measurement space, and T denotes the total number of temporal samples. Every feature vector can encompass voltage, current, temperature, power demand, state of charge, and statistical descriptors derived from sliding windows. The measurements together define the current operating state of the battery. The State-of-Health (SoH) of the battery is described in Equation (2).

$$SoH_t = \frac{C_t}{C_{\text{rated}}} \quad (2)$$

where C_t represents the available discharge capacity at time t , C_{rated} denotes the nominal rated capacity at the beginning of life, and $SoH_t \in (0, 1]$ quantifies the relative health condition. This standardized approach guarantees consistency among batteries and their operational characteristics. The goal of predictive maintenance is organized into three interconnected learning tasks: forecasting health in the short term, issuing degradation warnings in the mid-term, and detecting risks based on acceleration.

3.1. Short-Horizon SoH Forecasting

The objective of short-term State-of-Health forecasting is to predict the battery's condition at a future time frame h . The final implementation considers a prediction horizon of one hour, as it is crucial for real-time monitoring operations. The forecasting function is outlined in Equation (3).

$$\widehat{SoH}_{t+h} = f(\mathbf{X}_t) \quad (3)$$

where $\widehat{SoH}_{t+h}$ denotes predicted health at time $t + h$, $f(\cdot)$ represents a nonlinear regression model, and $\mathbf{X}_t$ denotes a temporal feature matrix constructed from historical observations. The temporal input window is defined mathematically in Equation (4):

$$\mathbf{X}_t = [\mathbf{x}_{t-w+1}, \mathbf{x}_{t-w+2}, \dots, \mathbf{x}_t] \quad (4)$$

where w represents the sliding-window size, where each row corresponds to a time-aligned feature vector. This windowing strategy captures local temporal dynamics such as trends,

fluctuations, and stress patterns preceding time t . The regression problem is formulated through minimization of mean squared error, and it is mathematically expressed in Equation (5), respectively:

$$\mathcal{L}_{reg} = \frac{1}{N} \sum_{i=1}^N (SoH_{i+h} - \widehat{SoH}_{i+h})^2 \quad (5)$$

where N denotes the number of training samples, SoH_{i+h} is the ground truth health value, and $\widehat{SoH}_{i+h}$ is the predicted value. This formulation ensures that deviations between predicted and actual SoH trajectories are minimized in a statistically consistent manner.

3.2. Seven-Day Degradation Warning Formulation

Although short-term forecasting addresses current health trends, effective operational maintenance necessitates early alerts within significant intervention periods. The magnitude of mid-horizon degradation over a specified interval of Δ (7 days) is represented in Equation (6):

$$D_t = SoH_t - SoH_{t+\Delta} \quad (6)$$

where D_t represents cumulative health loss over 7 days, and Δ corresponds to the number of time steps equivalent to 7 days. A degradation warning event is triggered when the observed drop exceeds a predefined threshold τ_{7d} , which is mathematically expressed in Equation (7), respectively:

$$y_t^{(7d)} = \begin{cases} 1, & D_t \geq \tau_{7d} \\ 0, & D_t < \tau_{7d} \end{cases} \quad (7)$$

where $y_t^{(7d)} \in \{0, 1\}$ denotes the binary degradation label, and τ_{7d} defines the minimum health loss considered operationally significant. The classification function is defined in Equation (8), respectively:

$$\hat{y}_t^{(7d)} = g(X_t) \quad (8)$$

where $g(\cdot)$ outputs the probability of a degradation event. Optimization is performed using cross-entropy loss, which is mathematically defined in Equation (9). This probabilistic formulation enables recall-oriented threshold tuning, ensuring early identification of significant degradation patterns.

$$\mathcal{L}_{cls} = -\sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (9)$$

3.3. Acceleration-Based Risk Detection Framework

Gradual degradation may transition into accelerated aging under thermal or load stress. To quantify this transition, the first-order degradation rate is defined in Equation (10), which represents instantaneous health decay.

$$R_t = SoH_{t-1} - SoH_t \quad (10)$$

Acceleration is then expressed as the second-order difference, and it is mathematically represented in Equation (11), respectively:

$$A_t = R_t - R_{t-1} \quad (11)$$

where A_t quantifies change in the degradation slope, and positive A_t indicates increasing degradation severity. A risk label is constructed when acceleration exceeds the threshold τ_{risk} , and it is mathematically expressed in Equation (12), while the corresponding classifier is represented in Equation (13), respectively. A corresponding classifier is trained to identify

operational stress-induced rapid aging conditions. This formulation ensures sensitivity to abrupt slope transitions rather than absolute *SoH* values.

$$y_t^{(risk)} = \begin{cases} 1, & A_t \geq \tau_{risk} \\ 0, & A_t < \tau_{risk} \end{cases} \quad (12)$$

$$\hat{y}_t^{(risk)} = r(X_t) \quad (13)$$

4. Methodology: Implemented Machine Learning Framework

The research methodology was structured to ensure clarity, completeness, and reproducibility by explicitly organizing the predictive-maintenance framework into a coherent sequence of stages. Initially, multivariate battery data were chronologically arranged and preprocessed to remove inconsistencies and preserve temporal integrity. Subsequently, temporal feature engineering was performed using sliding-window techniques, through which statistically and physically meaningful descriptors such as rolling averages, gradients, and stress indicators were derived exclusively from historical observations to prevent information leakage. The problem was then formulated as a multi-objective learning task, where short-horizon State-of-Health forecasting, mid-term degradation warning, and acceleration-based risk detection were jointly addressed. A rolling-origin validation strategy was adopted to maintain causality and to simulate real-world deployment conditions. For forecasting, an ensemble regression model combining multiple tree-based learners was employed, with blending weights determined from validation performance to enhance generalization. Degradation warning and acceleration risk detection were modeled as probabilistic classification tasks, where thresholds were calibrated with emphasis on recall to ensure early identification of critical events. Furthermore, stress-conditioned evaluation was conducted to assess robustness under high-temperature and high-load scenarios. Through this systematically organized methodology, the framework was designed to transform battery-health data into actionable maintenance intelligence, thereby addressing the reviewer's concern regarding methodological clarity and completeness.

4.1. Data Preprocessing and Temporal Feature Engineering

The initial step involves converting the raw multivariate time-series data into organized temporal formats to effectively capture short-term dynamics and patterns of operational stress. The time-indexed feature space is depicted in Equation (14), while the application of a sliding window of length w to create temporal matrices is illustrated in Equation (15):

$$\mathcal{X} = \{x_t\}_{t=1}^T \quad (14)$$

$$X_t = [x_{t-w+1}, x_{t-w+2}, \dots, x_t] \quad (15)$$

For each window, various derived descriptors are calculated, such as rolling mean and variance, first-order gradients, temperature–current interaction terms, load magnitude statistics, and local degradation slope indicators. These embedding transformations convert immediate measurements into contextually relevant inputs that can effectively capture nonlinear changes in health over time. All operations for feature extraction are conducted without referring to future observations, thus maintaining the integrity of the chronological order. To avoid any risk of information leakage, feature selection is performed solely within the training subset of each validation fold. $\mathcal{D}_{train}^k$ denotes the training partition of fold k . Tree-based impurity measures are utilized to calculate feature importance scores. The identical transformation is subsequently implemented on the related validation partition.

This guarantees that the validation data stay statistically unobserved throughout the feature ranking process, as mathematically represented in Equation (16).

$$X_{train}^k \rightarrow X_{train}^{k*} \quad (16)$$

Battery-health data are inherently temporal; therefore, conventional random cross-validation is unsuitable. A rolling-origin scheme is implemented in Equation (17).

$$\mathcal{D}_1 \subset \mathcal{D}_2 \subset \dots \subset \mathcal{D}_K \quad (17)$$

Every fold extends the training period into the future while assessing performance on later, unobserved data. This framework simulates actual deployment scenarios in which future data are not accessible during the training phase. Chronological ordering is rigorously maintained, preventing any look-ahead bias.

4.2. Ensemble Regression for Short-Horizon SoH Forecasting

Short-horizon SoH forecasting is achieved using a convex ensemble of three nonlinear regressors: Histogram-Based Gradient Boosting Regressor (HGBR), Random Forest Regressor (RF), and Extremely Randomized Trees Regressor (ET) [22,26].

$$\widehat{SoH}_{t+h}^{(k)} = f_k(X_t; \theta_k) \quad (18)$$

The final blended prediction is defined in Equation (19).

$$\widehat{SoH}_{t+h} = \sum_{k=1}^3 w_k \widehat{SoH}_{t+h}^{(k)}, \text{ subject to } \sum_{k=1}^3 w_k = 1, w_k \geq 0. \quad (19)$$

The blending weights w_k are optimized based on validation R^2 performance:

$$w_k = \frac{R_k^2}{\sum_{j=1}^3 R_j^2} \quad (20)$$

In Equation (18), the short-horizon State-of-Health forecast generated by the k^{th} regression learner is defined as a nonlinear mapping from the temporally embedded feature matrix X_t to the predicted future health value $\widehat{SoH}_{t+h}^{(k)}$. Here, t denotes the present time index, h represents the prediction horizon, and θ_k encapsulates the internal parameters of the corresponding regression model, including tree structures, splitting rules, and boosting coefficients, depending on the learner type. This formulation establishes that multiple heterogeneous regressors independently approximate battery-health evolution using identical historical input representations while maintaining model diversity. Equation (19) then consolidates these independent predictions into a convex ensemble output, where the final estimate $\widehat{SoH}_{t+h}$ is obtained as a weighted linear combination of the individual forecasts. The blending weights w_k are required to be non-negative and normalized, ensuring that their total sums to one. This condition of convexity guarantees numerical stability and stops any individual learner from excessively influencing the final prediction. Equation (20) provides additional details indicating that the blending weights are obtained from validation-based coefficients of determination R_k^2 . These coefficients measure the extent to which each learner accounts for the variance in the observed State of Health during out-of-sample validation. By standardizing these validation scores, a greater emphasis is consistently placed on models that exhibit enhanced generalization ability. This combination effectively lowers model variance, reduces the risk of overfitting, and improves generalization across various operating conditions.

4.3. Seven-Day Degradation Warning Classifier

The degradation-warning mechanism is subsequently formalized through Equation (6), where cumulative health loss over a predefined mid-horizon interval Δ is computed as the difference between present and future State-of-Health values. The parameter Δ corresponds to the seven-day interval expressed in discrete time steps. Equation (7) transforms this continuous degradation magnitude into a binary event indicator by comparing it against a predefined operational threshold τ_{7d} . This threshold represents a clinically or operationally meaningful decline in capacity beyond which maintenance intervention is considered necessary. The probabilistic classifier defined in Equation (8) produces a degradation-event likelihood $\hat{p}_t^{(7d)}$ based on the same temporally embedded feature matrix. The learning objective associated with this classifier is expressed in Equation (9) through the cross-entropy loss function, which penalizes divergence between true binary labels and predicted probabilities. This probabilistic formulation enables flexible threshold adjustment and supports recall-sensitive decision strategies. Mid-horizon degradation is modeled as a probabilistic classification task.

4.4. Degradation Acceleration Risk Modeling

Equations (10) and (11) present the analysis of degradation based on acceleration. The first-order degradation rate is calculated by taking the difference between successive State-of-Health Values, which indicates the immediate rate of health decline. The second-order difference, referred to as acceleration, measures the variations in this degradation slope. Positive acceleration values suggest an increase in the severity of aging, frequently linked to stress-related factors like higher temperatures or increased current demands. Equation (12) transforms the calculated acceleration magnitude into a binary risk label by juxtaposing it with a specified acceleration threshold τ_{risk} . The classifier outlined in Equation (13) assesses the likelihood of rapid degradation onset, facilitating the early identification of nonlinear aging transitions that might not be apparent from absolute health values by themselves. Risk detection based on acceleration emphasizes the identification of nonlinear slope amplification.

4.5. Recall-Oriented Threshold Optimization

The classification outputs are determined by the decision rule outlined in Equation (21), which involves comparing predicted probabilities against an adjustable threshold parameter γ . A binary event is designated when the anticipated probability surpasses this threshold. Equation (22) establishes the threshold-selection strategy by enhancing the F1-score while adhering to a minimum recall requirement. The F1-score serves as the harmonic means of precision and recall, effectively balancing detection accuracy with sensitivity, while the recall constraint guarantees that early degradation events are not missed. Given that early detection takes precedence over conservative precision, probability thresholds are fine-tuned to enhance the F1-score while adhering to recall limitations. The rule for decision making is outlined in Equation (21).

$$\hat{y}_t = \begin{cases} 1, & \hat{p}_t \geq \gamma \\ 0, & \hat{p}_t < \gamma \end{cases} \quad (21)$$

The threshold γ is selected via validation search to satisfy with respect to Equation (22):

$$\max_{\gamma} F1(\gamma); \text{ subject to } Recall(\gamma) \geq R_{min} \quad (22)$$

This guarantees a responsive approach to both degradation and acceleration occurrences. The predictive-maintenance framework is assessed under clearly defined stress

conditions to guarantee reliability at the deployment level. The following equations broaden the modeling framework to include stress-conditioned validation. The operational state vector includes temperature and current magnitude, while stress regimes are established by comparing these variables against specific thresholds. Metrics for conditional evaluation are subsequently calculated independently for both stress and nominal subsets. A robust stability index is defined as the ratio of stress-regime performance to nominal-regime performance, effectively measuring the resilience of the predictive framework in challenging operating conditions.

4.5.1. Stress Regime Definition

The operational state at time t be defined as:

$$\mathbf{s}_t = [T_t, I_t] \quad (23)$$

where T_t denotes battery temperature, and I_t denotes current magnitude. Two stress thresholds are defined T_{crit}, I_{crit} : a stress indicator variable is constructed as:

$$\delta_t = \begin{cases} 1, & T_t \geq T_{crit} \text{ or } |I_t| \geq I_{crit} \\ 0, & \text{otherwise} \end{cases} \quad (24)$$

Thus, $\delta_t = 1$ represents the stress regime, and $\delta_t = 0$ represents the nominal regime.

4.5.2. Conditional Model Evaluation

Regression robustness is evaluated conditionally:

$$R_{stress}^2 = R^2(\{t : \delta_t = 1\}) \quad (25)$$

$$R_{normal}^2 = R^2(\{t : \delta_t = 0\}) \quad (26)$$

Similarly, classification metrics are computed:

$$AUC_{stress} = AUC(\{t : \delta_t = 1\}) \quad (27)$$

4.5.3. Robustness Stability Index

To quantify performance degradation under stress, a stability index is defined:

$$\Phi = \frac{M_{stress}}{M_{normal}} \quad (28)$$

where M represents any evaluation metric (R^2 , AUC, F1). Interpretation: $\Phi \approx 1 \rightarrow$ stress resilience, while $\Phi < 1 \rightarrow$ performance degradation. The modeling sequence then transitions to the uncertainty-aware neural architecture. The multi-layer perceptron (MLP) mapping is defined through successive affine transformations and nonlinear activation functions applied to the multi-fidelity input vector. Each hidden layer is characterized by a weight matrix and bias vector, while nonlinear activation functions such as rectified linear units introduce representational flexibility. Unlike deterministic regression models, the heteroscedastic formulation simultaneously predicts the mean state of charge $\mu(t)$ and the input-dependent variance $\sigma^2(t)$, enabling distributional output rather than a single-point estimate. The negative log-likelihood (NLL) loss function is employed to jointly penalize inaccurate mean predictions and underestimated variance. This formulation allows the model to express confidence in its predictions, thereby supporting safety-critical decision making. The fast-charging optimization problem is subsequently formulated as a constrained control objective. The charging current $u(t)$ serves as the decision variable,

and the cost function penalizes deviation between predicted state of charge and a desired reference trajectory while incorporating a regularization term to prevent excessively aggressive charging. Battery dynamics are represented through a first-order differential equation derived from Coulombic balance, where the rate of change of the state of charge depends on applied current, Coulombic efficiency, and nominal capacity. Operational constraints are imposed to ensure that the state of charge remains within physical bounds, current magnitude does not exceed hardware limits, and temperature remains below a maximum safety threshold. These constraints guarantee safe and feasible charging trajectories. The final group of equations defines evaluation metrics used for performance validation. The coefficient of determination R^2 measures explained variance relative to total variance. Mean squared error (MSE) quantifies the average squared deviation between predicted and actual values. Prediction Interval Coverage Probability (PICP) represents the proportion of true values contained within predicted confidence intervals, while Mean Prediction Interval Width (MPIW) quantifies the average span of those intervals. Together, these metrics provide complementary insight into accuracy, reliability, and uncertainty calibration. Equations (21)–(28) transform the predictive-maintenance framework from a simple regression architecture into a comprehensive, uncertainty-aware, stress-validated, and control-integrated battery intelligence system. The overall framework of the proposed research problem is depicted in Figure 2, respectively.




Figure 2. Overall framework of the proposed research work.

The suggested framework is following chronological preprocessing, short-term State-of-Health forecasting, a 7-day degradation alert, and acceleration risk identification, which

are collectively developed utilizing crafted temporal features. A rolling-origin validation approach is utilized to maintain causality, and feature selection along with model calibration is conducted solely during training to prevent information leakage. Short-term forecasting is achieved using validation-weighted ensemble regression, while warning and risk decisions are based on recall-constrained probabilistic classifiers. The final performance is assessed in both nominal and stress-conditioned operating conditions. The corresponding implemented approach pseudocode is tabulated in Algorithm 1, respectively.

Algorithm 1: Proposed Multi-Layer Predictive-Maintenance Framework for EV Batteries

Input:

Battery dataset $\mathcal{D} = \{x_t, SoH_t\}$, forecast horizon h ,
degradation horizon Δ , training window W , step size S ,
number of folds K , feature count p ,
thresholds τ_{7d} and τ_{risk}

Output:

Short-horizon SoH forecast, degradation warning,
acceleration risk prediction, performance metrics

1: Preprocess $\mathcal{D}$ and sort samples chronologically

2: Compute derived targets:

$$D_t = SoH_t - SoH_{t + \Delta}$$

$$R_t = SoH_{t-1} - SoH_t$$

$$A_t = R_t - R_{t-1}$$

3: Define labels:

$$y_t(7d) = 1 \text{ if } D_t \geq \tau_{7d}, \text{ else } 0$$

$$y_t(risk) = 1 \text{ if } A_t \geq \tau_{risk}, \text{ else } 0$$

4: Construct temporal feature matrix using sliding-window statistics,
lag terms, gradients, and stress indicators

5: for $k = 1$ to K do

6: Split data using rolling-origin validation

7: Fit scaler on training fold only

8: Select top- p features using training fold only

9: Train HGBR, RF, and Extra Trees regressors for SoH forecasting

10: Estimate validation R^2 for each regressor

11: Compute normalized ensemble weights from validation R^2

12: Generate final SoH forecast by weighted model blending

13: Train gradient-boosted classifier for 7-day degradation warning

14: Calibrate warning threshold on validation data under recall constraint

15: Predict degradation warning labels on test fold

16: Train gradient-boosted classifier for acceleration risk detection

17: Calibrate risk threshold on validation data under recall constraint

18: Predict acceleration risk labels on test fold

19: Evaluate regression and classification metrics

20: Compute stress-conditioned performance under high-temperature
and high-load regimes

21: end for

22: Aggregate fold-wise results and report mean performance

Return final forecasts, warning labels, risk labels, and robustness statistics

5. Simulated Results and Discussion

The dataset employed in this study comprises multivariate time-series measurements collected from an electric vehicle battery system under realistic operating conditions. The data were originally recorded at 15 min intervals and, where required for computational efficiency, were downsampled to 30 min intervals without altering temporal continuity.

The complete dataset contains approximately 175,000 observations prior to preprocessing and between 85,000 and 90,000 samples after downsampling and data cleaning procedures. Each record includes a timestamp and key battery state variables, namely State of Health (SoH), state of charge (SoC), battery temperature, terminal voltage, current magnitude, and cumulative charge cycle count. The SoH values span approximately 0.40 to 1.00, thereby representing the progressive aging trajectory of the battery from near beginning-of-life conditions to advanced degradation stages. The chronological structure of the dataset facilitates simultaneous modeling of short-term electrochemical dynamics and long-term degradation behavior. Operational variability reflected in the dataset includes fluctuations in load demand, repeated charging–discharging events, and temperature variations. These factors are known to influence capacity fade and aging acceleration. To support predictive-maintenance analysis, derived features are constructed exclusively from historical observations. These include lagged state variables, rolling statistical measures computed over hourly and daily windows, local degradation gradients, and stress indicators based on temperature and current thresholds. Such feature construction ensures that no future information is incorporated during model training. The modeling targets are formulated to address multiple maintenance objectives. These include short-horizon SoH prediction, mid- to long-horizon degradation estimation, and a binary acceleration-based risk indicator derived from temporal changes in degradation rate. The dataset therefore supports both regression-based health forecasting and classification-based early warning detection.

Scenario 1. *Short-Horizon Health Tracking:* The time-series comparison shown in Figure 3 demonstrates that the ensemble model successfully captures the fine-grained SoH dynamics over the final validation fold. The predicted trajectory closely overlaps the true SoH curve, including minor oscillations and transient variations. Although the persistence baseline also performs strongly due to high short-term autocorrelation in SoH, the ensemble maintains comparable precision while preserving structural consistency under temporal variation.




Figure 3. Temporal comparison of actual, ensemble prediction, and persistence baseline.

The quantitative evaluation confirms this visual observation, with an R^2 of 0.9254 and an RMSE of 0.0033. The very small absolute error indicates that short-term degradation behavior remains smooth and highly predictable. In practical deployment, this layer serves as a continuous health tracking mechanism, ensuring reliable short-horizon state awareness. As seen in Figure 4 the data points are closely grouped along the ideal line, indicating a robust linear correlation and negligible systematic bias. The lack of curvature or dispersion widening indicates that prediction stability is preserved throughout the entire SoH operating range.




Figure 4. Predicted vs. actual *SoH* for one-hour horizon.

The residual behavior illustrated in Figure 5 demonstrates that prediction errors are centered around zero, exhibiting no noticeable drift. No signs of heteroscedasticity or temporal clustering are apparent. The limited residual variation indicates that the regression model is statistically stable and does not suffer from structural overfitting. This layer guarantees that the battery management system has precise near-term insights into health changes, even amidst varying load and temperature conditions. The comparatively stable short-horizon forecasting behavior may be attributed to the strong temporal persistence exhibited by battery degradation trajectories over limited operational intervals. Since short-term electrochemical aging evolves gradually under normal driving and charging behavior, recent historical observations retain considerable predictive relevance for near-future *SoH* estimation. Furthermore, the ensemble learning configuration benefits from the complementary nonlinear learning characteristics of tree-based learners, allowing transient fluctuations caused by dynamic loading and thermal variation to be captured without excessive sensitivity to local noise. The rolling-origin validation strategy additionally contributes toward realistic temporal generalization by preventing information leakage between historical and future operating states.




Figure 5. Residual error distribution for one-hour *SoH* forecast: temporal stability assessment.

Scenario 2. The second aim is to identify degradation at a mid-horizon level. This layer needs to detect important degradation trends before noticeable health loss occurs, unlike the short-term task. Figure 6 illustrates that the classifier exhibits robust recall performance throughout various operating thresholds, as shown by the precision–recall curve. The system demonstrates a recall rate surpassing 0.82 at the chosen threshold, effectively recognizing most impending degradation events. This behavior is advantageous in maintenance applications, where overlooked events hold greater significance than moderate false alarms.




Figure 6. Precision–recall curve for seven-day degradation warning under class imbalance conditions.

Figure 7 illustrates the ROC curve, which further validates the discrimination capability, showing an AUC of 0.7838. The curve consistently stays well above the diagonal reference, demonstrating a clear distinction between degrading and non-degrading states. Figure 8 presents the confusion matrix, offering further operational insights. A total of 563 true-positive warnings were accurately recognized, whereas 151 degradation events were overlooked. While there are still false positives (1108), this compromise aligns with a strategy focused on recall. In predictive maintenance, this type of setup is suitable, as early inspections can help reduce the risk of further complications. This layer effectively converts raw *SoH* trajectories into practical maintenance insights, enabling scheduling modifications and inspection planning days ahead of time. From an operational perspective, the comparatively high recall achieved by the degradation-warning module is particularly significant because missed degradation events may result in accelerated capacity fade, reduced driving reliability, or unexpected maintenance interruptions in electric vehicle applications. In predictive-maintenance environments, conservative early-warning behavior is often preferable to delayed detection because preventive intervention can be scheduled before severe degradation propagation occurs. The observed classification behavior also suggests that cumulative degradation patterns become more distinguishable when temporal feature embedding captures sustained thermal and loading trends rather than isolated instantaneous measurements alone.




Figure 7. Receiver operating characteristic (ROC) curve for seven-day degradation warning.




Figure 8. Confusion matrix for seven-day degradation warning: event-level classification performance.

Scenario 3. *Degradation acceleration risk.* The third objective focuses on the rapid acceleration of degradation, which poses a greater safety concern compared to a gradual decline. The precision–recall curve illustrated in Figure 9 demonstrates a more robust performance profile in comparison to the seven-day warning module. The PR-AUC of 0.7101 indicates a better balance between precision and recall in the presence of class imbalance. A high recall rate of 0.8845 guarantees that instances of accelerated degradation are seldom missed. Figure 9 illustrates the ROC curve, which shows an AUC of 0.8796, indicating a strong capacity for discrimination.




Figure 9. Precision–recall curve for degradation acceleration risk prediction.

Figure 9 illustrates the ROC curve, which shows an AUC of 0.8796, indicating a strong capacity for discrimination. The curve rises sharply near its origin, signifying high true-positive rates while maintaining relatively low false-positive rates. This functionality is crucial for oversight in situations where safety is paramount. Figure 10 offers further insight into operational effectiveness by illustrating the confusion matrix. The low count of false negatives highlights the effectiveness of the acceleration-based detection method. This module functions as a proactive alert system, aimed at identifying atypical aging trends resulting from thermal stress, load fluctuations, or operational inconsistencies.




Figure 10. Confusion matrix for degradation acceleration risk classification.

The strong acceleration risk detection behavior may be physically associated with the sensitivity of lithium-ion degradation mechanisms to elevated thermal exposure and

aggressive load cycling. Under high-temperature and high-current operating conditions, electrochemical side reactions, internal resistance growth, and localized degradation processes may intensify nonlinearly, thereby increasing the degradation acceleration rate rather than only the absolute *SoH* decline. The incorporation of temporal degradation gradients and stress-sensitive operational indicators therefore improves the capability of the proposed framework to recognize abrupt degradation transitions before severe long-term health deterioration becomes fully observable. The analysis of feature importance illustrated in Figure 11 offers clarity regarding the predictive framework. The most influential features include lagged State-of-Health (*SoH*) indicators, rolling voltage averages, and short-term temporal windows. The slopes of *SoH* indicate that temporal memory and recent degradation gradients play a crucial role in short-horizon forecasting. The inclusion of temperature and current rolling statistics among the highest-ranked features strengthens their physical validity. The aging of batteries is recognized to be affected by thermal stress and the intensity of the load. The observed dominance of temporal degradation history features further indicates that battery aging behavior is strongly path dependent, meaning that recent operating history contributes substantially toward future degradation evolution. In addition, the importance of temperature-related variables supports established electrochemical understanding that elevated thermal exposure accelerates parasitic side reactions and capacity-loss mechanisms within lithium-ion battery systems. Similarly, current-related loading features likely capture stress-induced operational variability associated with repeated charging–discharging intensity and transient power demand behavior. Consequently, the hierarchy of learned features corresponds with the principles of electrochemical degradation instead of being based solely on statistical artefacts. Cyclic time encodings play a significant role, suggesting that operational periodicity affects short-term *SoH* behavior. The results compiled in Table 2 show that the proposed framework consistently performs well in both regression and classification layers while also preserving temporal integrity through rolling-origin validation. For the short-term forecasting task, an R^2 value of 0.9254 was obtained, along with an RMSE of 0.0033 and a MAPE of 0.32%. The data indicate that the one-hour *SoH* dynamics were monitored with great accuracy. The similar performance of the persistence baseline shows that short-term *SoH* evolution is still highly autocorrelated; nonetheless, the suggested ensemble maintained predictive stability while ensuring structural generalization.




Figure 11. Top feature importance scores for one-hour *SoH* forecasting model.

Table 2. Comprehensive performance evaluation metrics.

Task	Metric	Value
1 h <i>SoH</i> Forecast (Proposed Model)	R^2	0.9261
	RMSE	0.0033
	MAE	0.028
	MAPE	0.32%
7d Degradation Warning	AUC	0.7838
	PR-AUC	0.0033
	Recall	0.8205
Acceleration Risk Detection	AUC	0.8796
	PR-AUC	0.7101
	Recall	0.8845

The seven-day degradation warning task yielded an AUC of 0.7838, demonstrating a significant distinction between degradation and non-degradation states. A recall of 0.8205 indicates that most upcoming degradation events were effectively recognized. While the precision stayed at a moderate level, the configuration focused on recall guarantees that essential maintenance signals are not missed, which is beneficial in preventive maintenance environments. The highest performance was noted in the degradation acceleration risk module, achieving an AUC of 0.8796 and a PR-AUC of 0.7101. A recall of 0.8845 suggests that accelerated aging events were identified with considerable sensitivity. The minimal number of false negatives further validate the framework's appropriateness for safety-critical monitoring applications. The findings suggest that although short-term State-of-Health forecasting provides precise health visibility, the classification layers add significant operational benefits by converting degradation patterns into practical early-warning and risk-detection signals. The combined multi-layer design thus improves both forecasting precision and maintenance insight within a cohesive machine learning system. The obtained forecasting and classification results were further analyzed in comparison with previously reported battery-health prognostic studies. Earlier investigations based on temporal convolutional networks, recurrent neural architectures, transfer learning models, and hybrid machine learning frameworks mainly concentrated on improving State-of-Health estimation accuracy under controlled cycling conditions [15–26]. Several studies reported strong regression performance for *SoH* estimation; however, most methodologies were designed primarily for single-objective prediction tasks and did not simultaneously incorporate degradation-warning capability or acceleration-sensitive maintenance intelligence. In contrast, the proposed framework integrates short-horizon forecasting, seven-day degradation warning, and acceleration risk detection within a unified temporally consistent architecture. This integration improves operational interpretability by converting battery-health estimation into actionable maintenance-oriented decision support. The strong forecasting accuracy obtained in the present work is consistent with observations reported in deep temporal learning and ensemble learning studies, where recent degradation history and temporal memory were identified as dominant contributors to short-term State-of-Health prediction [15,18,22]. The high predictive stability achieved by the proposed ensemble framework can be attributed to the complementary behavior of Histogram-Based Gradient Boosting, Random Forest, and Extremely Randomized Trees models, which collectively reduced variance while preserving nonlinear learning capability. Unlike several existing approaches that rely heavily on deep neural structures requiring extensive training complexity and large-scale parameter optimization, the adopted ensemble learning configuration maintained high regression accuracy while improving robustness and interpretability under rolling-origin validation. For degradation warning and acceleration risk detection, direct comparison with the existing literature remains

limited because most prior studies primarily focused on point-wise State-of-Health estimation or Remaining Useful Life prediction without explicitly modeling operational warning horizons or acceleration-sensitive degradation transitions. The improved recall behavior achieved in the proposed framework indicates that the adopted recall-oriented threshold optimization successfully prioritized early identification of critical degradation events. The acceleration risk module further demonstrated strong discrimination capability under stress-conditioned scenarios, which may be attributed to the incorporation of temporal degradation gradients, temperature–current interaction indicators, and acceleration-aware feature representations. Therefore, the proposed framework extends beyond conventional battery-health estimation by integrating predictive forecasting with operational maintenance intelligence suitable for practical electric vehicle deployment environments. The performance results summarized in Table 2 indicate that the proposed framework achieved consistently strong behavior across both regression and classification tasks. In the short-horizon forecasting layer, the obtained coefficient of determination and low RMSE confirm that the ensemble structure was capable of accurately tracking near-term State-of-Health variations with minimal prediction deviation. The very small MAPE further demonstrates that the forecasting error remained operationally negligible throughout the validation period. For the seven-day degradation warning task, the achieved recall value indicates that various upcoming degradation events were successfully identified before severe health deterioration occurred. Although moderate false-positive behavior was observed, the adopted recall-oriented configuration is considered appropriate for predictive-maintenance environments where early detection is prioritized over conservative alarm suppression. The acceleration risk detection layer demonstrated the highest classification robustness among all operational modules. The elevated AUC and PR-AUC values confirm strong discrimination capability under class-imbalanced conditions, while the high recall value indicates reliable identification of rapidly degrading operating states. Collectively, the results demonstrate that the proposed multi-layer framework not only improves forecasting precision but also enhances operational maintenance intelligence through early-warning and stress-sensitive risk detection capabilities.

6. Conclusions

A comprehensive predictive-maintenance framework for electric vehicle battery systems has been created and thoroughly assessed using a temporally consistent rolling-origin validation approach. The system combines short-term State of Health (*SoH*) predictions, a seven-day alert for degradation, and the identification of risks associated with accelerated degradation into a cohesive machine learning workflow. For forecasting the State of Health (*SoH*) one hour ahead, a coefficient of determination (R^2) of 0.9254 was obtained, along with an RMSE of 0.0033 and a MAPE of 0.32%. The findings indicate that the short-term health dynamics were accurately recorded with a low margin of prediction error. The consistency of residual behavior and the strong correlation between predicted and actual trajectories suggest that the regression module maintained structural integrity without any noticeable bias. During the seven-day degradation warning task, an AUC of 0.7838 and a recall of 0.8205 were achieved. The values indicate that potential degradation events were detected with high sensitivity, which is crucial for scheduling preventive maintenance. The most significant performance was noted in the degradation acceleration risk module, achieving an AUC of 0.8796 and a PR-AUC of 0.7101. A recall of 0.8845 indicates that accelerated aging episodes were identified with great accuracy, while only a few critical events were overlooked. This behavior underscores the appropriateness of the framework for monitoring scenarios that are critical to safety. The suggested architecture does not act merely as a standalone forecasting model; instead, it serves as a cohesive maintenance-intelligence sys-

tem that can effectively support monitoring, planning, and risk mitigation all at once. The obtained findings suggest that the proposed predictive-maintenance framework is capable of supporting integrated battery-health monitoring through short-horizon forecasting, degradation-warning generation, and acceleration-sensitive risk analysis within a unified temporal learning architecture. The achieved forecasting and classification performance indicate promising capability for identifying degradation behavior under dynamically varying operating conditions. Although the framework demonstrated stable performance under the evaluated scenarios, additional large-scale validation using publicly benchmarked datasets and broader operating environments would further strengthen generalizability assessment and deployment confidence. Therefore, the proposed methodology should be interpreted as a practically oriented predictive-maintenance framework with potential applicability for intelligent electric vehicle battery monitoring and early-warning support systems. Future investigations may additionally incorporate systematic benchmarking against advanced deep learning and standalone machine learning prognostic models using standardized public battery datasets to further evaluate comparative performance and generalization behavior under diverse operating conditions.

Author Contributions: Conceptualization, methodology, software, validation, formal analysis investigation, resources, and data curation, C.H.P. and P.S.V.; writing—original draft preparation, writing—review and editing, visualization, and supervision A.K.; project administration, Investigation, validation, and Review and Editing C.P.B.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Hossain Lipu, M.S.; Hannan, M.A.; Hussain, A.; Hoque, M.M.; Ker, P.J.; Saad, M.H.M.; Ayob, A. A review of state of health and remaining useful life estimation methods for lithium-ion battery in electric vehicles: Challenges and recommendations. *J. Clean. Prod.* **2018**, *205*, 115–133. [CrossRef]
2. Li, Y.; Liu, K.; Foley, A.M.; Zülke, A.; Berecibar, M.; Nanini-Maury, E.; Van Mierlo, J.; Hoster, H.E. Data-driven health estimation and lifetime prediction of lithium-ion batteries: A review. *Renew. Sustain. Energy Rev.* **2019**, *113*, 109254. [CrossRef]
3. Chandran, V.; Patil, C.K.; Karthick, A.; Ganeshaperumal, D.; Rahim, R.; Ghosh, A. State of Charge Estimation of Lithium-Ion Battery for Electric Vehicles Using Machine Learning Algorithms. *World Electr. Veh. J.* **2021**, *12*, 38. [CrossRef]
4. Ge, M.-F.; Liu, Y.; Jiang, X.; Liu, J. A review on state of health estimations and remaining useful life prognostics of lithium-ion batteries. *Measurement* **2021**, *174*, 109057. [CrossRef]
5. Sui, X.; He, S.; Vilsen, S.B.; Meng, J.; Teodorescu, R.; Stroe, D.-I. A review of non-probabilistic machine learning-based state of health estimation techniques for lithium-ion battery. *Appl. Energy* **2021**, *300*, 117346. [CrossRef]
6. Wang, Z.; Feng, G.; Zhen, D.; Gu, F.; Ball, A. A review on online state of charge and state of health estimation for lithium-ion batteries in electric vehicles. *Energy Rep.* **2021**, *7*, 5141–5161. [CrossRef]
7. IEA. *Global EV Outlook 2024*; IEA: Paris, France, 2024. Available online: <https://www.iea.org/reports/global-ev-outlook-2024> (accessed on 21 May 2026).
8. Maheswari, N.; Sahaya, J.; Rama Mohan, P.; Joselin Jeya Sheela, J. Maximizing electric vehicle efficiency: Integrating direct torque control for two-wheel drive system using performance enhancing approach. *Energy* **2025**, *324*, 135237. [CrossRef]
9. Ramkumar, G.; Kannan, S.; Mohanavel, V.; Karthikeyan, S.; Titus, A. The future of green mobility: A review exploring renewable energy systems integration in electric vehicles. *Results Eng.* **2025**, *27*, 105647. [CrossRef]
10. Ramkumar, M.S.; Muthukrishnan, S.; Giri, J.; Sathish, T.; Fatehmulla, A. Optimizing battery and supercapacitor management in electric vehicles: A hybrid approach for enhanced performance and reduced harmonics. *Case Stud. Therm. Eng.* **2025**, *68*, 105815. [CrossRef]
11. Lu, Y.; Guo, D.; Xiong, G.; Wei, Y.; Zhang, J.; Wang, Y.; Ouyang, M. Towards real-world state of health estimation: Part 2, system level method using electric vehicle field data. *eTransportation* **2024**, *22*, 100361. [CrossRef]

12. Rauf, H.; Khalid, M.; Arshad, N. Machine learning in state of health and remaining useful life estimation: Theoretical and technological development in battery degradation modelling. *Renew. Sustain. Energy Rev.* **2022**, *156*, 111903. [[CrossRef](#)]
13. Urquizo, J.; Singh, P. A review of health estimation methods for lithium-ion batteries in electric vehicles and their relevance for battery energy storage systems. *J. Energy Storage* **2023**, *73*, 109194. [[CrossRef](#)]
14. Kurucan, M.; Özbaltan, M.; Yetgin, Z.; Alkaya, A. Applications of artificial neural network based battery management systems: A literature review. *Renew. Sustain. Energy Rev.* **2024**, *192*, 114262. [[CrossRef](#)]
15. Bockrath, S.; Lorentz, V.; Pruckner, M. State of health estimation of lithium-ion batteries with a temporal convolutional neural network using partial load profiles. *Appl. Energy* **2023**, *329*, 120307. [[CrossRef](#)]
16. Ma, Y.; Shan, C.; Gao, J.; Chen, H. A novel method for state of health estimation of lithium-ion batteries based on improved LSTM and health indicators extraction. *Energy* **2022**, *251*, 123973. [[CrossRef](#)]
17. Zhang, Y.; Liu, Y.; Wang, J.; Zhang, T. State-of-health estimation for lithium-ion batteries by combining model-based incremental capacity analysis with support vector regression. *Energy* **2022**, *239*, 121986. [[CrossRef](#)]
18. Liu, S.; Chen, Z.; Yuan, L.; Xu, Z.; Jin, L.; Zhang, C. State of health estimation of lithium-ion batteries based on multi-feature extraction and temporal convolutional network. *J. Energy Storage* **2024**, *75*, 109658. [[CrossRef](#)]
19. Xu, J.; Qu, J.; Xu, H. Capacity estimation of lithium-ion batteries with automatic feature extraction and graph-enhanced LSTM. *J. Energy Storage* **2024**, *85*, 111131. [[CrossRef](#)]
20. Lanubile, A.; Bosoni, P.; Pozzato, G.; Allam, A.; Acquarone, M.; Onori, S. Domain knowledge-guided machine learning framework for state of health estimation in lithium-ion batteries. *Commun. Eng.* **2024**, *3*, 168. [[CrossRef](#)]
21. Hofmann, T.; Hamar, J.; Mager, B.; Erhard, S.; Schmidt, J.P. Transfer learning from synthetic data for open-circuit voltage curve reconstruction and state of health estimation of lithium-ion batteries from partial charging segments. *Energy AI* **2024**, *17*, 100382. [[CrossRef](#)]
22. Wu, J.; Cui, X.; Meng, J.; Peng, J.; Lin, M. Data-driven transfer-stacking-based state of health estimation for lithium-ion batteries. *IEEE Trans. Ind. Electron.* **2024**, *71*, 604–614. [[CrossRef](#)]
23. Wu, J.; Fang, L.; Dong, G.; Lin, M. State of health estimation for lithium-ion batteries in real-world electric vehicles. *Sci. China Technol. Sci.* **2023**, *66*, 47–56. [[CrossRef](#)]
24. Gong, L.; Zhang, Z.; Li, X.; Sun, K.; Yang, H.; Li, B.; Ye, H.; Wang, X.; Tan, P. State of health estimation of lithium-ion batteries based on machine learning with mechanical-electrical features. *Batter. Supercaps* **2024**, *7*, e202400201. [[CrossRef](#)]
25. Mu, G.; Wei, Q.; Xu, Y.; Li, J.; Zhang, H.; Yang, F.; Zhang, J.; Li, Q. State of health estimation of lithium-ion batteries based on feature optimization and data-driven models. *Energy* **2025**, *316*, 134578. [[CrossRef](#)]
26. Wen, S.; Lin, N.; Huang, S.; Li, X.; Wang, Z.; Zhang, Z. Lithium battery state of health estimation using real-world vehicle data and an interpretable hybrid framework. *J. Energy Storage* **2024**, *96*, 112623. [[CrossRef](#)]

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