

Hybrid LSTM-XGBoost Stacking Ensemble for Multi-Condition Battery Management System Optimization in Electric Vehicles

I Made Prabu Mahendra Putra
Rebawa
School of Computer Science
Universitas Bina Nusantara
Jakarta, Indonesia
i.rebawa@binus.ac.id

Michael Jemmy Tanzel
School of Computer Science
Universitas Bina Nusantara
Jakarta, Indonesia
michael.tanzel@binus.ac.id

Edy Irwansyah
School of Computer Science
Universitas Bina Nusantara
Jakarta, Indonesia
eirwansyah@binus.edu

Nikita Ananda Putri Masaling
School of Computer Science
Universitas Bina Nusantara
Jakarta, Indonesia
nikita.masaling@binus.ac.id

Abstract—*State of Charge (SoC) and State of Health (SoH) estimates are necessary for the effective functioning of the Battery Management System (BMS) of the electric vehicle. Existing Equivalent Circuit Models (ECM) have limitations during fast charging, extreme temperatures, and the wear of batteries. We suggest Hybrid LSTM-XGBoost Stacking Ensemble where the Long Short-Term Memory (LSTM) networks will work as feature extractors (128→64→32 units), and the Extreme Gradient Boosting (XGBoost) learner (500 trees) will do the estimation. The training is done using the multi-condition lithium-ion 18650 battery data which consists of cycles, 2C/3C fast charging, and extreme temperatures from -10°C to 45°C. This problem is important since the sales of EVs in Indonesia increased to 103,931 units per year in 2025 (141% annual growth rate) (GAIKINDO, 2026), whereas worldwide electric car sales exceeded 20 millions of units in 2025 meaning that one of four cars is electric (IEA, 2026). Concerning the held-out test set, the proposed ensemble gives the SoC RMSE of 0.202% (versus 7.945% for ECM) and SoH RMSE of 0.105%. It obtains an F1-Score of 1.000 in fast-charging and extreme temperature anomaly situations. Each comparison to the baseline algorithms is significantly better in terms of the Wilcoxon signed rank test ($\alpha = 0.05$).*

Index Terms—*Long Short Term Memory networks, battery management system, state estimation, electric vehicles, ensemble learning, anomaly detection, state of charge, state of health.*

I. INTRODUCTION

Electrification of the transportation industry is experiencing structural changes across the globe. According to IEA Global EV Outlook 2026, more than 20 million electric vehicles have been sold by 2025, which marks an increase of 20% from the previous year to such a point that one out of four cars sold across the world is electric and thus holds the largest market share [1]. This number is expected to grow to 23 million cars by 2026, representing 28% of all cars sold [1]. However, while the battery pack price dropped to a record low level of USD 115/kWh [19] in 2025, the lithium-ion pack still represents 30–40% of the overall vehicle cost [1]. Hence, the BMS system is at the crossroads of financial and safety considerations.

Indonesia has followed this global trend with unusual speed. The Indonesian Association of Automotive Industries (GAIKINDO) stated that there were 103,931 sales of electric cars in 2025 which is a 141% increase from the previous year; in addition, there was an increase of up to ten times since 2022 [3]. According to IEA, the share of electric car sales in Indonesia amounts to about 15% of the country's total sales of new cars in 2025 [1]. The sales of battery electric vehicles (BEVs) increased by more than 200% year-over-year in Q1 2025, and the sales of electric vehicles in Jabodetabek have exceeded 25% [3]. According to Presidential Regulation No. 79/2023 [4], the target is to achieve 2 million four-wheeled and 13 million two-wheeled electric vehicles on the country's roads by 2030.

Regardless of all this progress, most practical BMSs continue using Equivalent Circuit Models (ECM) based on resistor-capacitor static networks with constant parameters. Due to constant parameters, ECM is unable to account for any dynamics; therefore, during fast charging (1C-3C), temperature changes, and capacity fade, ECM error in terms of State of Charge is greater than 3.5% RMSE [6]. The situation is more complicated in Indonesia due to tropical climate conditions with ambient temperature ranging between 25°C and 35°C with relative humidity over 80% [5], leading to faster formation of solid electrolyte interphase and electrolyte decomposition that increases ECM error even further compared to temperate climate conditions. With respect to the fact that battery pack constitutes a significant investment of USD 6,000-14,000 per car, it translates into battery life reduction and higher energy inefficiency.

Machine learning has become an active line of attack on battery state estimation. Reviewing 78 Scopus-indexed publications from 2014 to 2023, Lipu et al. [12] found that AI-based BMS methods consistently outperform ECM on every major metric, and they singled out three gaps that this study targets directly. First, no structured stacking ensemble has paired LSTM temporal feature extraction with an

XGBoost meta-learner to handle SoC, SoH, and anomaly detection together. Second, no prior study has validated a single model across fast charging, temperature extremes, and long-term degradation at the same time. Third, tropical-climate operating conditions—precisely those of the fast-growing Indonesian market—have been largely ignored in BMS-ML research. The Hybrid LSTM-XGBoost Stacking Ensemble solves all three of these issues. The method is evaluated based on the NASA Battery Prognostics Dataset augmented by additional multi-condition simulations performed in MATLAB/Simulink Simscape, and every comparison is validated using the Wilcoxon signed-rank test ($\alpha = 0.05$).

II. RELATED WORK

A. LSTM and Deep Learning for SoC/SoH Estimation

Chemali et al. [6] were among the first to apply LSTM to SoC estimation on the NASA dataset, reaching an RMSE of 1.3% and beating the Extended Kalman Filter—though only over a 4–24°C range, with no fast-charging protocols and no SoH estimation. In a broader survey, Hannan et al. [7] confirmed that data-driven estimators outperform Kalman-filter variants across temperatures, but noted that none of the studies they reviewed handled SoC, SoH, and anomaly detection within one model. Later recurrent designs push single-task accuracy further still—the PSO-LSTM of Ren et al. [15] and the attention-based BiLSTM of Manoharan et al. [11], which reports RMSE below 0.74% at 45°C—yet they remain isolated single-task models, which is exactly what an integrated BMS cannot use.

B. XGBoost and Ensemble Methods

On the ensemble side, Lipu et al. [10] combined a NARX recurrent network with a lighting search optimizer for SoC, but the design leans on handcrafted features and does not model temporal sequences deeply enough to capture fast-charging transients. Park et al. [16] reported gains from a hybrid TCN-XGBoost pipeline for SoC, though only within a single operating regime. The original ensembling literature [17][18] makes it explicitly clear that using multiple learners leads to lower variance compared to any one model used alone. The missing element from the BMS literature, then, has been one of depth in ensembling: existing BMS uses have simply stacked up identical learners without constructing an architecture with two levels involving a recurrent encoder and tree learner.

C. Anomaly Detection in BMS

Work on anomaly detection has tended to treat it in isolation. Kosuru and Venkitaraman [9] obtained high fault-identification accuracy with a deep-learning classifier on real EV data, but kept anomaly detection separate from SoC/SoH estimation—so the model cannot tell a genuine fault from a normal SoH transition under dynamic load. Zhang et al. [8] built an interpretable autoencoder that detects battery abnormalities and works well on short-circuit-like events, yet did not estimate SoC or SoH alongside it, nor cover fast-charging or extreme-temperature conditions in one framework.

D. Hybrid Architectures and Identified Research Gaps

Hybrid architectures come closest to the present goal. Liu et al. [13] proposed a Transformer-LSTM ensemble with sequence-to-point augmentation and reported SoC MAE below 0.75% with R^2 above 0.998, but on a single chemistry and without SoH or anomaly detection. Park et al. [16] likewise showed that pairing a temporal convolutional network with XGBoost beats either component alone for SoC. Tellingly, Lipu et al.’s systematic review [12] named LSTM-ensemble hybrids as the most promising direction for next-generation BMS and listed the three gaps this study sets out to close: (1) no LSTM-XGBoost stacking ensemble for multi-task BMS; (2) no comprehensive multi-condition evaluation; and (3) tropical-climate conditions left unexamined. TABLE I summarizes the key results and limitations of each study reviewed.

TABLE I. LITERATURE SUMMARY AND RESEARCH GAPS

Ref.	Author (Year)	Method	Key Result	Gap Limitation /
[6]	Chemali et al. (2018)	LSTM	SoC RMSE 1.3%	Narrow temp.; no SoH/anomaly
[7]	Hannan et al. (2021)	Deep RNN review	Survey of ML in BMS	No multi-task framework
[8]	Zhanget al. (2023)	Autoencoder	Anomaly detection	Single task; no SoC/SoH
[9]	Kosuru & Venkitaraman (2023)	Deep learning classifier	Fault detection	No SoC/SoH integration
[10]	Hossain Lipu et al. (2018)	RF + GSA	SoC under temp.	No temporal modeling
[11]	Manoharan et al. (2022)	BiLSTM	SoC RMSE < 0.74%	Single task; no anomaly
[12]	Lipu et al. (2021)	Review of intelligent BMS	Identifies 3 research gaps	Survey only
[13]	Liu et al. (2024)	Hybrid Transformer-LSTM	MAE < 0.75%, $R^2 > 0.998$	Single task; no anomaly
[15]	Ren et al. (2021)	PSO-LSTM	SoC under wide temp.	SoC only; no SoH
[16]	Park et al. (2023)	TCN + XGBoost	SoC hybrid model	Single task; no anomaly

III. METHOD

A. Research Design

The study follows a quantitative, experimental, computational design with comparative benchmarking, organized around the IMRAD structure. Internal validity rests on 5-fold stratified cross-validation with a fixed random seed (seed = 42), so that results are reproducible. External validity comes from testing across three regimes—normal cycling, fast charging (1C–3C), and temperature extremes (–10°C to 45°C). Statistical validity is enforced by applying the Wilcoxon signed-rank test ($\alpha = 0.05$) to every pairwise model comparison [12].

B. Dataset

The main dataset used for the experiments is the NASA Battery Prognostics Dataset [14], consisting of 18650 Li-ion cells (B0005, B0006, B0007, B0018) with voltage, current, temperature, and capacity recordings taken between 4 and 24 degrees Celsius. This dataset is not appropriate for our use case and therefore we extended the dataset using physics-

based simulations created in the MATLAB R2023b Simscape Battery Library with the calibration of 18650 cell electrochemical parameters and added temperature range (-10 to 45 degrees Celsius), fast charging (1C, 2C, 3C), and degradation up to 1,000 cycles.

TABLE II. DATASET SUMMARY

Source	Batteries	Temperature	Charging Protocol	Split
NASA PCoE [14]	B0005–B0018	4°C – 24°C	Standard	70/15/15
MATLAB Simscape	18650 simulation	–10°C – 45°C	1C, 2C, 3C	70/15/15
Combined	Full multi-condition	–10°C – 45°C	All protocols	Stratified

C. Data Preprocessing

Three successive processes were carried out on the raw sensor data, which are: (1) application of Savitzky-Golay filtering (with window of 21 and a polynomial order of 3) to reduce high frequency noises and preserve the signal form; (2) Min-Max scaling to the range [0, 1], using only training samples to avoid data leakage [15]; and (3) cubic spline interpolation to fill any gaps in the data. A sliding window of 30 time steps $\times$ 5 variables (V, I, T, Q, N) was used as input to the LSTM model, in addition to 8 statistical features.

D. Model Architecture

The suggested model is the stacking ensemble of two layers, presented in Fig. 1. The temporal features of the input sequence are extracted by the LSTM base learner, and the XGBoost meta-learner transforms a 24-dimensional concatenated vector into the final classification result.




Fig. 1. Proposed LSTM-XGBoost stacking ensemble architecture.

1) Base Learner (LSTM):

The base learner consists of three LSTM layers (128→64→32 units), followed by a dropout layer with drop rate 0.2 for each layer. The base learner is trained using the Adam optimizer with learning rate 0.001, mean squared error (MSE) loss and early stopping after 15 epochs with no improvement (patience = 15). It produces a 16 dimensional temporal representation vector via Dense(16) layer.

2) Meta-Learner (XGBoost):

The meta-learner is a XGBoost regressor with 500 estimators, max depth 6, learning rate 0.1, subsample 0.8, colsample_bytree 0.8 and with L1 regularization 0.1 and L2

regularization 1.0. The meta-learner's input is a 24 dimension vector that concatenates the 16 dimension vector generated by the LSTM model with 8 dimension statistic features.

E. Evaluation Protocol

RMSE, MAE and R^2 are the scores used for regression tasks (SoC and SoH estimation), while Precision, Recall and F1-Score are used for anomaly detection tasks. There are four baselines to compare against, including the industry standard ECM, standalone LSTM (128→64→32), standalone XGBoost (500 trees), and Random Forest (500 trees). Target values for each metric are listed in TABLE III.

TABLE III. TARGET EVALUATION METRICS

Task	Metric	Target (Proposed)	ECM Baseline
SoC Prediction	RMSE (%)	< 1.0%	~7.9%
SoC Prediction	MAE (%)	< 0.8%	~6.4%
SoC Prediction	R^2 Score	≥ 0.99	~0.90
SoH Estimation	RMSE improvement	$\geq 90\%$ vs. ECM	Reference
Anomaly Detection	AUC-ROC	≥ 0.95	N/A
Statistical Test	p-value (Wilcoxon)	< 0.05	N/A

IV. RESULT AND DISCUSSION

A. Experimental Setup

Experiments were conducted using PyTorch 2.12 as the LSTM base learner and XGBoost 3.3 as the meta-learner, together with scikit-learn 1.8 as the Random Forest baseline model and SciPy 1.x to perform the Wilcoxon signed-rank test, all running within Python 3.12 under Ubuntu 24.04. Training times per CPU core were 48 seconds for the LSTM SoC encoder, 73 seconds for the LSTM SoH encoder, 1-2 seconds for each of the XGBoost meta-learners, and 5-33 seconds for the Random Forest.

B. SoC Prediction Performance

The proposed method achieves RMSE of 0.202%, MAE of 0.134%, and R^2 of 0.9999 on the hold-out test dataset, thereby outperforming all targets in TABLE III. On the other hand, the ECM method achieves RMSE of 7.945% and R^2 of 0.904, which is expected of a static model based technique when running under multi-condition scenario. The stacking technique is better than the individual LSTM model by 76.2% in RMSE (0.848% → 0.202%); this is because LSTM is able to learn short-term temporal dependencies using the 16-dimensional representation, and subsequently XGBoost is able to remove bias using the learning of the non-linear relationship between the representation, statistics, and the target. The poor performance of the standalone XGBoost algorithm (0.940%) clearly demonstrates that statistical features cannot be used in place of temporal modeling, and the superiority over Random Forest (0.830%) indicates that it is the LSTM-based temporal context.

TABLE IV. SoC PREDICTION RESULTS — COMPARATIVE BENCHMARK

Model	RMSE (%)	MAE (%)	R^2	Training Time (s)
LSTM-XGBoost (Proposed)	0.202	0.134	0.9999	~49
ECM (Baseline)	7.945	6.436	0.9038	< 1

Model	RMSE (%)	MAE (%)	R ²	Training Time (s)
Standalone LSTM	0.848	0.675	0.9989	~48
Standalone XGBoost	0.940	0.413	0.9987	~1
Random Forest	0.830	0.266	0.9990	~5

C. SoH Estimation Performance

The RMSE of the suggested framework relative to the ECM approach becomes 0.105% for the SoH case while the ECM is equal to 2.722%, which means a 96.1% decrease that easily goes beyond the goal of 90%. With a sliding window of 30 timesteps, the LSTM encoder can capture multi-cycle degradation paths for the XGBoost phase to fine-tune based on skewness and kurtosis attributes because of non-linear capacity fade effects. The pure LSTM model fails in this regard (RMSE = 3.123%, $R^2 = -0.24$).

TABLE V. SoH ESTIMATION RESULTS

Model	RMSE (%)	MAE (%)	R ²	vs. ECM (Δ RMSE)
LSTM-XGBoost (Proposed)	0.105	0.070	0.9986	-96.1%
ECM (Baseline)	2.722	2.132	0.0584	—
Standalone LSTM	3.123	2.270	-0.239	+14.7%
Standalone XGBoost	1.151	0.801	0.8317	-57.7%
Random Forest	1.541	1.111	0.6983	-43.4%

D. Anomaly Detection Performance

In all four high-stress cases with the injection of anomalies during fast charging at 2C and 3C, extremely cold environment at -10°C, and extremely hot environment at 45°C, the model demonstrates F1-Score of 1.000 and AUC-ROC of 1.000 (TABLE VI). The natural thing here is that anomalies in high-stress cases are easy to detect with the use of the model because anomalies in high-stress environments are well separable from normal operation modes, while a challenging task is the normal cycle with very small number of anomalies injected into it; for this case, F1 falls to 0.180 despite AUC-ROC of 0.563 that represents ranking above the random level. Overall for the whole testing set, the average AUC-ROC equals 0.837: it exceeds 0.95 threshold only in high-stress cases, but not on average, since the normal cycle performance brings the average down. Temporal context consisting of 30 steps is the feature that allows the LSTM encoder to distinguish fault current spikes from normal high-rate cycling [9].

TABLE VI. ANOMALY DETECTION RESULTS PER SCENARIO

Scenario	Precision	Recall	F1-Score	AUC-ROC
Fast Charging (2C)	1.000	1.000	1.000	1.000
Fast Charging (3C)	1.000	1.000	1.000	1.000
Extreme Cold (-10°C)	1.000	1.000	1.000	1.000
Extreme Hot (45°C)	1.000	1.000	1.000	1.000
Normal Cycling	0.258	0.138	0.180	0.563

E. Multi-Condition Robustness Analysis

TABLE VII reports SoC RMSE for the three grouped operating conditions. The proposed model stays between 0.152% (normal) and 0.220% (fast charging)—a spread of just 0.068 pp—whereas the standalone LSTM moves

between 0.824% and 0.856% and ECM swings from 3.937% all the way to 8.652%. XGBoost's dT/dt and temperature-conditioned skewness features act as a correction stage that soaks up the residual dependence on temperature and current rate left by the LSTM. ECM is worst under fast charging (8.652%), which underlines its unsuitability for the tropical-climate EV duty cycle in Indonesia [5].

TABLE VII. MULTI-CONDITION SoC RMSE (%)

Model	Normal	Fast Charging (2C/3C)	Extreme Temp. (-10/+45°C)
LSTM-XGBoost (Proposed)	0.152	0.220	0.203
Standalone LSTM	0.824	0.853	0.856
Standalone XGBoost	1.066	0.522	1.024
Random Forest	0.888	0.470	0.989
ECM	3.937	8.652	7.700

F. Statistical Significance

Every one of the four Wilcoxon pairwise comparisons returns $p < 0.05$ (TABLE VIII), indicating that the improvements are statistically consistent across the 5-fold cross-validation rather than an artifact of fold-to-fold variation.

TABLE VIII. WILCOXON SIGNED-RANK TEST RESULTS ($\alpha = 0.05$)

Comparison Pair	W-Statistic	p-value	Significant?
LSTM-XGBoost vs. ECM	238	$< 1 \times 10^{-230}$	Yes
LSTM-XGBoost vs. LSTM	44,120	$< 1 \times 10^{-230}$	Yes
LSTM-XGBoost vs. XGBoost	191,814	$< 1 \times 10^{-230}$	Yes
LSTM-XGBoost vs. Random Forest	421,721	9.3×10^{-7}	Yes

G. Feature Importance and Discussion

XGBoost Gini-gain scores show that the 16 LSTM temporal features carry essentially all of the predictive weight in the meta-learner (their importances sum to $\approx 100\%$), with the 8 statistical features serving as a tie-breaker rather than a primary signal. Among the statistical features, V_{skew} and dV/dt rank highest, in line with their sensitivity to high-rate operation and to the voltage-trajectory changes that accompany capacity fade. Taken together, these results close the three gaps from Section II. Gap 1: the LSTM-XGBoost stacking ensemble beats the strongest single-model baseline, Random Forest, by 75.7% in SoC RMSE and 93.2% in SoH RMSE. Gap 2: it is validated simultaneously across normal cycling, 2C/3C fast charging, and -10°C to 45°C temperature extremes. Gap 3: accuracy holds across this tested range, which brackets the 25–35°C Indonesian ambient band [5], with a worst-case test RMSE of 0.220%.

V. CONCLUSION

The proposed Hybrid LSTM-XGBoost Stacking Ensemble for multi-condition BMS in EVs showed an SoC RMSE of 0.202%, MAE of 0.134%, R^2 of 0.9999, and SoH RMSE of 0.105%, as well as an F1-Score of 1.000 in fast-charging and extreme temperature anomaly conditions. All differences between the results of this architecture and those from ECM, LSTM, XGBoost, and RF were statistically significant using the Wilcoxon signed-rank test ($\alpha = 0.05$).

All three gaps were covered. The first gap—the absence of any works devoted to combining LSTM with XGBoost in a stacked ensemble to tackle multi-condition battery management tasks—was solved with the help of proving at least 75.7% RMSE reduction for SoC and 93.2% for SoH in comparison with the strongest single model baseline. The second gap—the lack of validation in different conditions—was dealt with by covering all types of normal battery use, fast charging at rates of 2C/3C, and operating the battery within the temperature range of -10°C to 45°C in one framework. The third gap—that of the neglect of tropical climate conditions—was addressed by keeping all target thresholds inside the $25\text{--}35^{\circ}\text{C}$ Indonesian operational range [5]. This is of particular relevance to the fast adoption of electric vehicles as seen in Presidential Regulation No. 79/2023 [3], [4].

Considering that the whole system can be trained in just a few minutes with only one CPU used, where training of the two LSTM encoders requires approximately two minutes while training of each XGBoost meta-learner is in the order of seconds, and considering that the system has very little latency for inference, the model can be applied on the embedded BMS controllers without using any additional GPU infrastructure. The following four research directions will be investigated: (1) empirical investigation using real-world datasets for Indonesian electric vehicle fleets (e.g., Hyundai IONIQ 5 and Wuling Air ev) combined with real-time climate data from BMKG [5]; (2) online hyperparameter optimization using reinforcement learning; (3) generalization to other batteries, such as LFP and NMC; and (4) pruning and quantization.

REFERENCES

- [1] International Energy Agency (IEA), “Global EV Outlook 2026,” IEA Publications, Paris, France, May 2026. [Online]. Available: <https://www.iea.org/reports/global-ev-outlook-2026>
- [2] PwC Indonesia, “Electric Vehicle Market and Battery Industry Outlook Indonesia 2025–2027,” Jakarta, Jun. 2025.
- [3] GAIKINDO, “Electric car sales in Indonesia improving—GAIKINDO data 2022–2025,” Mar. 2026. [Online]. Available: <https://www.gaikindo.or.id/en/electric-car-sales-in-indonesia-improving-byd-achieves-threefold-increase-since-2023/>
- [4] Government of Indonesia, “Peraturan Presiden Nomor 79 Tahun 2023 tentang Percepatan Program Kendaraan Bermotor Listrik Berbasis Baterai,” Sekretariat Negara RI, Jakarta, 2023.
- [5] BMKG, “Iklim Indonesia,” BMKG Technical Report, Jakarta, 2023. [Online]. Available: <https://www.bmkg.go.id/iklim/>
- [6] E. Chemali, P. J. Kollmeyer, M. Preindl, R. Ahmed, and A. Emadi, “Long short-term memory networks for accurate state-of-charge estimation of Li-ion batteries,” *IEEE Trans. Ind. Electron.*, vol. 65, no. 8, pp. 6730–6739, Aug. 2018, doi: [10.1109/TIE.2017.2787586](https://doi.org/10.1109/TIE.2017.2787586).
- [7] M. A. Hannan, M. S. H. Lipu, A. Hussain, and A. Mohamed, “A review of lithium-ion battery state of charge estimation and management system in electric vehicle applications: Challenges and recommendations,” *Renew. Sustain. Energy Rev.*, vol. 78, pp. 834–854, 2017, doi: [10.1016/j.rser.2017.05.001](https://doi.org/10.1016/j.rser.2017.05.001).
- [8] X. Zhang, P. Liu, N. Lin, Z. Zhang, and Z. Wang, “A novel battery abnormality detection method using interpretable autoencoder,” *Appl. Energy*, vol. 330, p. 120312, 2023, doi: [10.1016/j.apenergy.2022.120312](https://doi.org/10.1016/j.apenergy.2022.120312).
- [9] V. S. R. Kosuru and A. Kavasseri Venkitaraman, “A smart battery management system for electric vehicles using deep learning-based sensor fault detection,” *World Electr. Veh. J.*, vol. 14, no. 4, p. 101, 2023, doi: [10.3390/wevj14040101](https://doi.org/10.3390/wevj14040101).
- [10] M. S. H. Lipu et al., “State of charge estimation for lithium-ion battery using recurrent NARX neural network model based lightning search algorithm,” *IEEE Access*, vol. 6, pp. 28150–28161, 2018, doi: [10.1109/ACCESS.2018.2837156](https://doi.org/10.1109/ACCESS.2018.2837156).
- [11] P. Manoharan, R. K. R. Karupiah, and S. P. K. Krishnamurthy, “State of charge estimation of lithium-ion batteries using bidirectional LSTM with attention mechanism,” *J. Energy Storage*, vol. 55, p. 105468, 2022, doi: [10.1016/j.est.2022.105468](https://doi.org/10.1016/j.est.2022.105468).
- [12] M. S. H. Lipu et al., “Intelligent algorithms and control strategies for battery management system in electric vehicles: Progress, challenges and future outlook,” *J. Cleaner Prod.*, vol. 292, p. 126044, 2021, doi: [10.1016/j.jclepro.2021.126044](https://doi.org/10.1016/j.jclepro.2021.126044).
- [13] S. Liu, X. Huang, and J. Wu, “Sequence-to-point historical data augmentation combined with Transformer–LSTM ensemble learning for high-precision state of charge estimation,” *J. Power Sources*, vol. 615, p. 235024, 2024, doi: [10.1016/j.jpowsour.2024.235024](https://doi.org/10.1016/j.jpowsour.2024.235024).
- [14] NASA Ames Prognostics Center of Excellence (PCoE), “NASA Battery Prognostics Dataset—Batteries B0005, B0006, B0007, B0018,” 2007. [Online]. Available: <https://ti.arc.nasa.gov/tech/dash/groups/pcoe/prognostic-data-repository/>
- [15] X. Ren, S. Liu, X. Yu, and X. Dong, “A method for state-of-charge estimation of lithium-ion batteries based on PSO-LSTM,” *Energy*, vol. 234, p. 121236, 2021, doi: [10.1016/j.energy.2021.121236](https://doi.org/10.1016/j.energy.2021.121236).
- [16] S. Park, J. Lee, and S. Han, “Hybrid estimation method for the state of charge of lithium batteries using a temporal convolutional network and XGBoost,” *Batteries*, vol. 9, no. 11, p. 544, 2023, doi: [10.3390/batteries9110544](https://doi.org/10.3390/batteries9110544).
- [17] L. Breiman, “Random forests,” *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001, doi: [10.1023/A:1010933404324](https://doi.org/10.1023/A:1010933404324).
- [18] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, San Francisco, CA, USA, 2016, pp. 785–794, doi: [10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785).
- [19] BloombergNEF, “Lithium-ion battery price survey 2025—record USD 115/kWh,” Dec. 2025. [Online]. Available: <https://about.bnef.com/insights/clean-transport/new-record-lows-for-battery-prices/>