International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(1s) 2026 pp. 1-10
https://doi.org/10.70917/ijcisim-2026-5993
Article

AI-Driven Battery Management System for Electric Vehicles Using Real-Time State-of-Charge and State-of-Health Estimation

Dr.Ch.Vinay Kumar[1], Dr.P.Lakshmi Supriya[2]

[1]Assistant Professor, Department of Electrical and Electronics Engineering Mahatma Gandhi Institute of Technology , chvinaykumar_eee@mgit.ac.in
[2]Assistant Professor , Department of Electrical and Electronics Engineering , Mahatma Gandhi Institute of Technology plaxmisupriya_eee@mgit.ac.in

Abstract: Accurate estimation of state of charge (SOC) and state of health (SOH) is fundamental to the safety, range prediction, power management, and lifetime optimization of electric-vehicle batteries. Neither state is directly measurable; both must be inferred from terminal measurements such as voltage, current, and temperature. Conventional coulomb counting is computationally simple but accumulates current-sensor and initialization errors, while model-based observers depend on parameter identification and may lose accuracy as the battery ages. This paper presents an AI-driven battery-management-system architecture that performs coupled, real-time-capable SOC and SOH estimation using two lightweight supervised regression models. The estimators consume causal features derived from voltage, current, temperature, elapsed time, cumulative ampere-hours, and voltage slope. A health-aware supervisory layer uses the resulting states to constrain charge, discharge, thermal, and balancing commands. A proof-of-concept evaluation was conducted on the NASA Ames lithium-ion battery-aging dataset. Cells B0005, B0006, and B0007 were used for training, while B0018 was held out entirely for testing to assess cell-to-cell generalization. On 34,866 held-out samples, the proposed estimator achieved an SOC mean absolute error (MAE) of 2.45 percentage points, SOC root-mean-square error (RMSE) of 3.39 percentage points, and coefficient of determination of 0.988. The corresponding SOH MAE and RMSE were 4.11 and 4.65 percentage points. A fixed-rated-capacity coulomb-counting baseline produced an SOC MAE of 11.19 percentage points on the same held-out cell, illustrating the importance of aging-aware estimation. Measured desktop inference time was approximately 7 microseconds per sample for each estimator, establishing computational feasibility but not replacing an embedded-controller timing certification. The results support lightweight AI as a practical correction and fusion layer for battery-state estimation, while also identifying cross-chemistry validation, uncertainty quantification, and hardware-in-the-loop testing as necessary next steps.

Keywords: Battery management system; electric vehicle; state of charge; state of health; lithium-ion battery; machine learning; real-time estimation; battery aging

1. Introduction

Lithium-ion battery packs are central to the performance and cost of modern electric vehicles. The battery management system (BMS) must protect cells against overcharge, over-discharge, overcurrent, and excessive temperature while estimating states that cannot be directly sensed. SOC describes the remaining usable charge relative to capacity, whereas capacity-based SOH expresses the available full-charge capacity of an aged cell relative to its fresh or nominal capacity. Reliable real-time estimation is difficult because voltage response depends on current history, temperature, hysteresis, relaxation, cell chemistry, and aging.[1]

Traditional SOC estimation commonly begins with coulomb counting. The method integrates measured current and is accurate when initial SOC and battery capacity are known, but its estimate drifts when those quantities or the

current measurement are inaccurate. Equivalent-circuit models combined with observers such as the extended Kalman filter can reduce drift and incorporate voltage feedback, but performance remains tied to the fidelity of the circuit parameters over temperature and aging. Recent reviews therefore identify data-driven and hybrid methods as important directions for automotive SOC, SOH, and state-of-energy estimation [2].

Deep and machine-learning estimators can learn nonlinear relationships among voltage, current, temperature, capacity loss, and internal resistance. Long short-term memory (LSTM), gated recurrent unit (GRU), convolutional neural network (CNN), autoencoder, and hybrid models are frequently studied because battery telemetry is sequential and strongly nonlinear [3]. However, many published models are validated on a single cell or a randomly shuffled subset of one dataset. Such protocols can leak cell-specific degradation patterns into the test set and overstate field generalization. Dataset scarcity, domain transfer, physical consistency, and interpretability remain important barriers to deployment [4].

This study addresses those issues with a transparent, reproducible proof of concept. Its principal contributions are:

- A dual-time-scale BMS architecture in which fast SOC estimation and slower SOH estimation share measured signals but retain separate update paths [5].
- A causal feature pipeline suitable for streaming implementation, using only voltage, current, temperature, elapsed time, accumulated charge, and voltage slope [6].
- A strict leave-one-cell-out evaluation on NASA cells, rather than a random sample split.
- A quantitative comparison with fixed-capacity coulomb counting on exactly the same held-out cell.
- A deployment analysis covering timing, fail-safe operation, model monitoring, and the limitations of transferring cell-level laboratory results to an EV pack.

2. Related Work

Conventional estimators

SOC is commonly defined as the ratio of remaining dischargeable capacity to fully charged capacity:

$$SOC(t) = \frac{Q_{rem}(t)}{C_{avail}(t)} \times 100\%. \#(1)$$

With the sign convention that discharge current is positive, discrete coulomb counting is written as:

$$SOC_k = SOC_{k-1} - \frac{\eta I_k \Delta t}{3600 C_{avail}} \times 100\%, \#(2)$$

where η is coulombic efficiency, I_k is current in amperes, Δt is the sampling interval in seconds, and C_{avail} is capacity in ampere-hours. Published analyses emphasize that the method relies on a known initial SOC and capacity. A fixed nominal capacity becomes increasingly inappropriate as the cell ages, which directly couples SOH error into SOC error [7].

Capacity-based SOH is defined as:

$$SOH_n = \frac{C_n}{C_0} \times 100\%, \#(3)$$

where C_n is the available capacity at cycle n and C_0 is the fresh or rated capacity. Internal-resistance-based definitions also exist, but capacity fade remains a widely used health indicator. Capacity and resistance do not necessarily age at identical rates, so an automotive BMS may maintain multiple health indicators rather than one scalar [8].

AI estimators

LSTM and GRU networks are attractive for battery-state estimation because gated recurrence can retain temporal information over long sequences. CNNs can extract local patterns, while attention mechanisms can weight the most informative time regions. Hybrid CNN-GRU-LSTM architectures have been proposed to combine spatial feature extraction with short- and long-term sequence learning. Reviews nevertheless note that model accuracy depends strongly on the dataset, chemistry, load profile, temperature range, and validation protocol [9].

Joint SOC-SOH estimation is preferable to treating the states as independent. Aging changes usable capacity, resistance, and voltage response; inaccurate SOH can therefore bias SOC, while errors in SOC can distort health features. One experimental NSSR-LSTM study reported SOC and SOH errors within 2.5% and 1.3%, respectively, and used the estimated SOH as an input to SOC estimation. Those values are not directly comparable with the present work because the dataset, input reconstruction, split protocol, cell configuration, and error definitions differ [10].

Real-time execution is also a design constraint. FPGA-based LSTM and bidirectional-LSTM implementations have demonstrated that hardware acceleration can support low-latency SOC/SOH inference, although bidirectional processing increases cost and is less naturally causal for streaming use. For safety-critical automotive deployment, accuracy alone is insufficient; deterministic execution, bounded memory, fault detection, physical plausibility, and graceful fallback must also be established [11].

Proposed BMS Architecture

Functional structure

The proposed architecture separates safety-critical measurement and actuation from AI inference. The signal path is:

1. Cell-monitoring integrated circuits measure cell voltage, pack current, and module temperature.
2. A deterministic preprocessing layer performs timestamp alignment, range checking, deglitching, and causal low-pass filtering.
3. A feature engine computes elapsed time, cumulative ampere-hours, voltage slope, and a validity mask.
4. The SOC estimator runs at the normal BMS measurement rate.
5. The SOH estimator runs at a slower rate or when sufficient excitation is present.
6. A state-fusion layer checks estimates against physical bounds and conventional observers.
7. The supervisory controller uses validated SOC and SOH values for power limits, charging limits, thermal control, balancing, range estimation, and diagnostic reporting.

The AI output must never bypass independent hardware protections. Overvoltage, undervoltage, short-circuit, overcurrent, and overtemperature responses remain deterministic and do not depend on a learned model [12].

Dual-time-scale estimation

SOC changes on a time scale of seconds, whereas SOH evolves over tens or hundreds of cycles. Running both estimators at the same rate wastes computation and can amplify noise [13]. The proposed design therefore uses a fast SOC channel and a slow health channel. At time k, the feature vector is:

$$x_k = [V_k, I_k, T_k, t_k, q_k, \dot{V}_k], \#(4)$$

with cumulative discharged charge:

$$q_k = q_{k-1} + \frac{max(I_k, 0)\Delta t}{3600}. \#(5)$$

The estimators are:

$$\widehat{SOC}_k = f_{SOC}(x_k), \widehat{SOH}_k = f_{SOH}(x_k). \#(6)$$

For the proof of concept, f_{SOC} and f_{SOH} are histogram-based gradient-boosted tree regressors. This choice tests whether a compact nonlinear model can provide useful state estimation without the memory and multiply-accumulate demand of a deep recurrent network. The production architecture is model-agnostic: the same interfaces can host quantized GRU, LSTM, temporal CNN, or physics-informed models after identical validation.

Health-aware correction

The available capacity used by an auxiliary coulomb counter is corrected with estimated SOH:

$$\hat{C}_{avail,k} = C_{rated} \frac{\widehat{SOH}_k}{100}. \#(7)$$

A fused SOC may then be calculated as:

$$SOC_k^{fused} = w_k \widehat{SOC}_k^{AI} + (1 - w_k) SOC_k^{CC}, \#(8)$$

where w_k decreases if inputs are missing, out of distribution, physically inconsistent, or associated with high model uncertainty. The experimental section reports the standalone AI estimate so that performance is not concealed by post-processing. Equation (8) is a deployment recommendation rather than part of the reported accuracy calculation [14].

Real-time algorithm

At each measurement interrupt, the controller executes the following sequence:

- Acquire synchronized voltage, current, and temperature.
- Reject impossible values and set diagnostic flags.
- Update the filtered signals, cumulative ampere-hours, and voltage slope.
- Normalize or map features using constants stored with the approved model version.
- Infer SOC; infer SOH only at its scheduled update instant.
- Clamp outputs to physical ranges and rate-limit implausible changes.
- Cross-check AI SOC against the auxiliary observer and open-circuit-voltage information when available.
- Publish validated estimates and confidence status to the supervisory controller.
- Fall back to the conventional observer if model inputs or outputs fail validation.

Dataset and Methods

NASA aging data

The NASA Ames Prognostics Center of Excellence dataset was selected because it contains run-to-failure lithium-ion time series collected under charge, discharge, and electrochemical-impedance-spectroscopy profiles at different temperatures. Its discharge records include measured voltage, current, temperature, time, and capacity, with an approximate acquisition rate of 10 Hz. The experiments used 18650-format cells and stopped at a stated end-of-life criterion of 30% fade from 2 Ah to 1.4 Ah [15].

Four cells from the FY08Q4 archive were analyzed. The dataset contains 168 discharge cycles each for B0005, B0006, and B0007, and 132 for B0018. After parsing valid discharge records, the first three cells contributed 150,855 training samples; B0018 contributed 34,866 test samples. The observed first and final discharge capacities were 1.856 and 1.325 Ah for B0005, 2.035 and 1.186 Ah for B0006, 1.891 and 1.432 Ah for B0007, and 1.855 and 1.341 Ah for B0018. These values were calculated directly from the downloaded NASA records.

Cell	Role	Discharge cycles	Samples	First capacity (Ah)	Final capacity (Ah)	Final rated-capacity SOH
B0005	Training	168	50,285	1.856	1.325	66.25%
B0006	Training	168	50,285	2.035	1.186	59.28%
B0007	Training	168	50,285	1.891	1.432	71.62%
B0018	Held-out test	132	34,866	1.855	1.341	67.05%

Ground-truth construction

For each discharge, charge throughput was obtained by trapezoidal integration of discharge current. Sample-level reference SOC was calculated using that discharge record's measured full capacity:

$$SOC_k^{ref} = max\left(0{,}1 - \frac{q_k}{C_n}\right) \times 100\%. \#(9)$$

Reference SOH was computed as $C_n/2.0 \times 100\%$, using the dataset's 2 Ah rated-capacity basis. This creates a capacity-based health label consistent with the dataset's stated end-of-life definition. The full-discharge capacity is used only to generate offline labels; it is not supplied as a model input during inference.

Preprocessing and features

Records were sorted in acquisition order. For every sample, six causal features were generated:

- Terminal voltage in volts.
- Measured current in amperes.
- Measured cell temperature in degrees Celsius.
- Elapsed time from the start of the discharge.
- Cumulative discharged ampere-hours.
- First-order voltage slope $\Delta V/\Delta t$.

No future voltage, future current, measured capacity, true SOC, or true SOH was included in the input vector. The split was performed by cell identity: every sample from B0018 remained unseen during training and

hyperparameter fitting. This protocol is more demanding than random sample splitting because adjacent samples from the same discharge are highly correlated.

Learning model

Separate histogram gradient-boosting regressors were trained for SOC and SOH. Each used a learning rate of 0.07, a maximum of 250 boosting iterations, 31 terminal leaves per tree, an L_2 regularization coefficient of 0.1, and a fixed random seed of 42. The outputs were clipped to the physically meaningful interval from 0 to 100%.

The loss minimized during training was mean squared error. Evaluation used:

$$MAE = \frac{1}{N}\sum_{i=1}^{N} |y_i - \hat{y}_i|, \#(10)$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N} (y_i - \hat{y}_i)^2}, \#(11)$$

$$R^2 = 1 - \frac{\sum_{i=1}^{N} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{N} (y_i - \bar{y})^2}. \#(12)$$

MAE and RMSE are reported in percentage points, not as percentages of the reference value. A fixed-capacity coulomb counter using 2 Ah was evaluated as the SOC baseline. Its initial SOC was set to 100% at the beginning of each discharge, which favors the baseline by removing initialization uncertainty.

3. Results

Estimation accuracy

The held-out-cell results are shown below. The AI SOC estimator reduced MAE from 11.19 percentage points for fixed-capacity coulomb counting to 2.45 percentage points. Its RMSE was 3.39 percentage points and R^2 was 0.988. The SOH estimator produced an MAE of 4.11 percentage points, RMSE of 4.65 percentage points, and R^2 of 0.647.

Estimator	Target	MAE (percentage points)	RMSE (percentage points)	R^2
Proposed AI model	SOC	2.45	3.39	0.988
Proposed AI model	SOH	4.11	4.65	0.647
Fixed 2 Ah coulomb counting	SOC	11.19	13.86	0.802

The approximately 78% reduction in SOC MAE relative to the baseline follows from the model's ability to use voltage, temperature, and nonlinear aging signatures in addition to accumulated current. The baseline divides charge throughput by the original 2 Ah rating even when the aged test cell's usable capacity is substantially lower. This comparison demonstrates the coupling between capacity loss and SOC error, not a universal superiority claim over calibrated, capacity-adaptive coulomb counting.

SOH estimation was less accurate than SOC estimation. This is expected because the instantaneous measurements contain many samples from a small number of physical cells, while health is a slowly varying cycle-level property. The held-out B0018 degradation trajectory differs from those of the three training cells, and the model

had no direct impedance or pack-history input. Recent SOH reviews similarly identify data scarcity, domain shift, and feature transferability as major challenges.

Computational performance

Training required approximately 1.40 seconds for the SOC model and 1.22 seconds for the SOH model in the analysis environment. Batch inference on 34,866 test samples required approximately 7.0 microseconds per sample for SOC and 7.6 microseconds per sample for SOH. These timings show that the selected regressors are lightweight on a general-purpose processor.

The timings must not be interpreted as proof of deadline compliance on an automotive microcontroller. Embedded performance depends on clock rate, memory hierarchy, numeric representation, compiler, runtime, model serialization, task scheduling, and safety diagnostics. A production implementation should measure worst-case execution time on the target electronic control unit and include preprocessing, communication, plausibility checks, and fallback logic.

4. Discussion

Why joint estimation matters

The experiment confirms a central BMS design principle: SOC and SOH should not be estimated in isolation. Coulomb counting requires capacity in its denominator, but available capacity declines with aging. If rated capacity remains fixed, the SOC trajectory becomes progressively optimistic for an aged cell. Joint methods address this interaction by adapting capacity or by feeding health information into the charge-state estimator. Prior joint-estimation work has likewise found that accurate SOH improves SOC estimation.

The proposed architecture is described as coupled rather than fully end-to-end. Separate estimators simplify verification, permit different update rates, and allow either channel to fall back independently. A future implementation can feed validated SOH into the SOC model or fuse the AI output with a capacity-adaptive coulomb counter. The coupling must be bounded so that a transient SOH error does not immediately destabilize SOC.

Safety and confidence management

A learned estimator can extrapolate incorrectly when exposed to a chemistry, temperature, aging mechanism, or sensor fault absent from training. Therefore, each estimate should carry an operating-domain status. Practical checks include input-range monitoring, residual comparison with an equivalent-circuit model, ensemble disagreement, rate-of-change limits, and consistency with open-circuit voltage after sufficient rest.

A safety-oriented BMS should implement three modes:

- **Normal mode:** AI estimate passes all plausibility and domain checks and is fused with the conventional observer.
- **Degraded mode:** confidence is low; power and charging limits are derated while the conventional observer receives higher weighting.
- **Fail-safe mode:** invalid measurements, memory errors, or estimator divergence cause the AI output to be rejected and independent protection thresholds to remain active.

Model artifacts should be versioned with the training-data identity, chemistry, feature definition, scaling constants, software checksum, validation evidence, and rollback image. This traceability is as important as average test-set accuracy.

Pack-level extension

The present experiment estimates single-cell states. An EV pack includes series and parallel groups, unequal temperatures, cell-to-cell variation, balancing currents, contact resistance, and sensor offsets. A deployable pack estimator should preserve cell-level safety while controlling computational cost. One approach is to run simple observers on every cell and apply the AI estimator to representative, hottest, coldest, highest-voltage, and lowest-voltage cells. Another is a shared model with cell-specific latent or calibration parameters.

Pack-level SOH should not be reduced blindly to the mean cell SOH. Usable pack energy and power are often limited by the weakest cell, thermal gradient, and voltage constraints. The supervisory layer should therefore retain minimum, maximum, dispersion, and confidence statistics across cells.

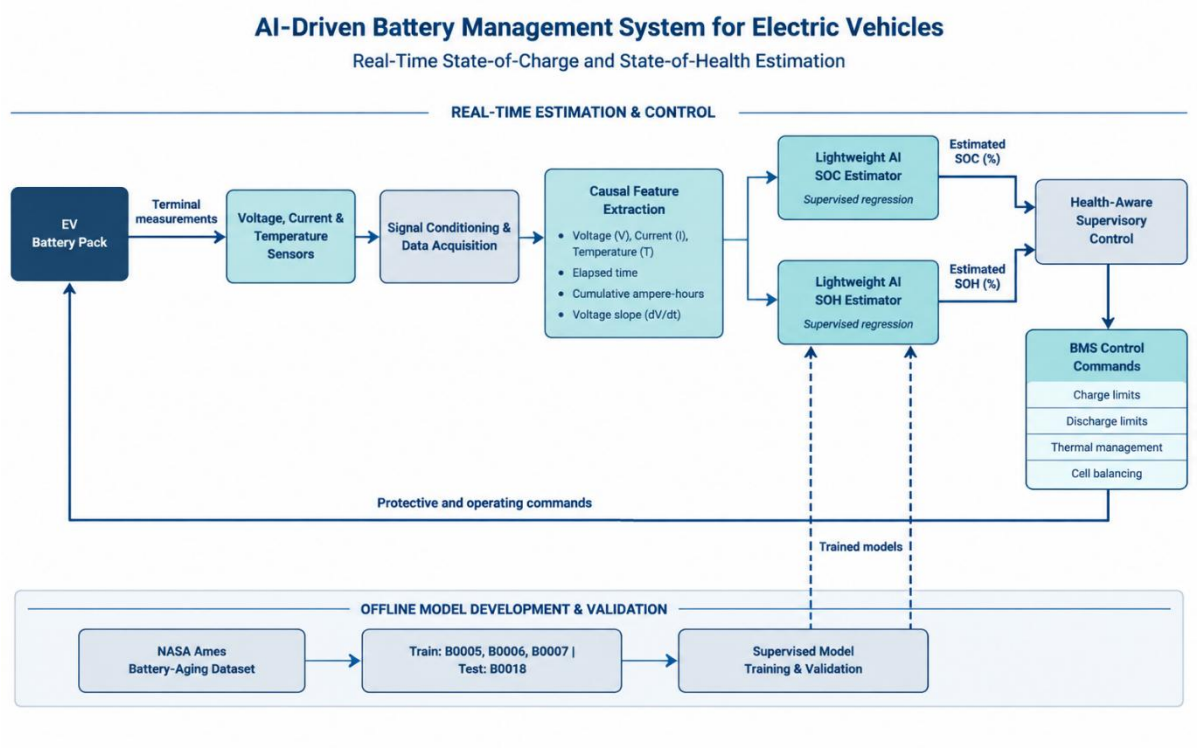


Figure : 1 Block diagram

Comparison with deep learning

The gradient-boosted model does not explicitly retain hidden temporal state as an LSTM or GRU does. Instead, it receives engineered causal history through elapsed time, accumulated charge, and voltage slope. This produces a small model with fast inference and strong SOC performance in the selected leave-one-cell-out test. It also sacrifices the ability to learn long-range patterns directly from raw sequences.

LSTM-based models have shown strong SOC and SOH performance, but computational cost, overfitting, and dataset dependence remain concerns. Attention and physics-informed approaches can improve temporal focus and physical consistency, but they add design and verification complexity. The appropriate production choice is therefore not determined by model accuracy alone. It should be selected through a common benchmark covering unseen cells, temperatures, drive cycles, sensor perturbations, memory use, worst-case latency, and out-of-distribution behavior.

Limitations

This study has five principal limitations.

- The NASA cells and laboratory cycling conditions do not reproduce the full chemistry, thermal gradients, vibration, fast charging, regenerative braking, and current transients of a modern EV pack.

- Only four cells were used, and one held-out cell cannot establish fleet-level generalization.
- Capacity-derived SOC labels assume that the capacity measured for a full discharge is the correct denominator throughout that discharge.
- SOH was defined only from capacity; resistance growth, power capability, lithium plating, and safety degradation were not modeled.
- The reported timing was measured on a general-purpose analysis processor, not an automotive-grade microcontroller, FPGA, or BMS system-on-chip.

Published literature cautions that high laboratory accuracy can degrade under different chemistries, operating conditions, and datasets. Consequently, the numerical results should be interpreted as a reproducible proof of concept and architecture validation, not as certification evidence for a road vehicle.[17][8]

Future Work

The next experimental phase should combine the NASA aging data with open datasets spanning LCO, LFP, NMC, and NCA chemistries. CALCE provides open-access battery data including full and partial cycling, storage, dynamic driving profiles, open-circuit voltage, and impedance measurements across several form factors and chemistries. Oxford also provides long-term degradation datasets with current, voltage, temperature, capacity, and impedance information. Cross-dataset testing should train on one institution's cells and test on another's to quantify domain shift.

Further technical work should include:

- A physics-informed loss that penalizes violations of charge conservation and implausible SOH increases.
- Uncertainty estimation using ensembles, quantile regression, or probabilistic sequence models.
- Comparison of gradient boosting, quantized GRU, temporal CNN, and hybrid equivalent-circuit-plus-AI models under identical splits.
- Temperature-stratified evaluation, including low-temperature and fast-charge conditions.
- Fault-injection tests for voltage bias, current offset, temperature-sensor failure, missing samples, and communication delay.
- Hardware-in-the-loop evaluation with worst-case execution-time and memory measurements.
- Pack-level validation incorporating imbalance, balancing currents, thermal gradients, and weakest-cell constraints.

Physics-informed and hybrid methods are especially promising because they combine learned nonlinear correction with conservation laws and interpretable battery models. Such constraints may reduce unphysical extrapolation and improve transfer when labeled aging data are limited.

5. Conclusion

A lightweight AI-driven BMS architecture for real-time-capable SOC and SOH estimation has been developed and evaluated using public NASA lithium-ion aging data. Strict leave-one-cell-out testing produced an SOC MAE of 2.45 percentage points and SOH MAE of 4.11 percentage points on 34,866 unseen samples. The AI SOC model

substantially outperformed a fixed-rated-capacity coulomb-counting baseline, demonstrating why health information must be incorporated into charge estimation as cells age.

The result does not eliminate the need for electrochemical knowledge or deterministic observers. The most credible deployment path is a hybrid architecture in which AI corrects nonlinear and aging-dependent behavior, physical models provide consistency checks, and independent hardware protections remain authoritative. Broader cross-chemistry data, uncertainty-aware inference, pack-level testing, and embedded hardware-in-the-loop validation are required before automotive use.

REFERENCES

1. Saikumar, K., & KIRAN, K. D. D. (2025). DESIGN AND IMPLEMENTATION OF ANFIS CONTROLLER BASED UPQC WITH RENEWABLE ENERGY BASED DISTRIBUTED GENERATION SYSTEM. *INDUSTRIAL ENGINEERING JOURNAL*, *54*(03), 22-37.
2. Saikumar, K., Siva, D., Srivalli, D., Basha, S. S., Kumar, B. S., & Mehbodniya, A. (2025). Deep learning driven heart disease prediction using ECG signal classification. *Panamerican Mathematical Journal*, *35*(2s), 42-56.
3. Saikumar, K., Patakottu, P., Baza, M., Rao, K. S., Rasheed, A., & Lalouani, W. (2024, December). RFID and IoT-enabled sign language to speech conversion using deep learning. In *2024 IEEE International Conference on RFID Technology and Applications (RFID-TA)* (pp. 161-164). IEEE.
4. Karpurapu, S., Saikumar, K., Kocharla, L., Rao, P. S., & Govathoti, S. (2024). Scalable innovative factors for shaping consumer intentions on electric two-wheelers adoptions. *Scalable Computing: Practice and Experience*, *25*(6), 4507-4517.
5. Poornima, R., Nagavarapu, S., Navva, S., Katkoori, A. K., Mohsen, K. S., & Saikumar, K. (2024, July). Multi-modal meta multi-task learning for social media rumor detection. In *2024 2nd World Conference on Communication & Computing (WCONF)* (pp. 1-6). IEEE.
6. Saikumar, K., Ravindra, P. S., Sravanthi, D., Mehbodniya, A., Webber, J. L., & Bostani, A. (2025). Heart Disease Prediction using Machine Learning and Deep Learning Approaches.
7. Neerumalla, S., Hanuman, A. S., Reddy, B. V., & Saikumar, K. (2025). ACDPSNet: Adaptive cross domain polarity aspect level learning scalable computing model for sentiment classification and quantification. *Scalable Computing: Practice and Experience*, *26*(4), 1671-1683.
8. Kumar, G. N. K., Unhelkar, B., Vani, K. S., Chakrabarti, P., & Saikumar, K. (2025). A cloud edge based heart disease detection using densenet convoluted radial basis neural network for diabetic patients. *Traitement du Signal*, *42*(3), 1481.
9. Saikumar, K., Bhavana, M., Prasanthi, R., Mallika, S. S., Kamidi, D., Malik, N., & Joshi, K. (2025). Enhanced air quality prediction using AI: A comparative study of GRU, CNN, and XGBoost models. *Advance Sustainable Science Engineering and Technology*, *7*(3), 02503012-02503012.
10. Thangaraj, K., Rao, K. S., Saikumar, K., & Garg, S. (2026). A software-defined wrapper discriminant federated learning-reinforcement attention adversarial regression approach for privacy and task management in cloud edge computing. *Knowledge-Based Systems*, 115430.
11. V. Daggupati, “Power Quality Disturbance Detection using CNN and Deep Learning,” International Research Journal of Engineering and Technology, vol. 13, no. 5, pp. 54–60, 2026. [Online]. Available: https://www.irjet.net/archives/V13/i5/IRJET-V13I0508.pdf.
12. V. Daggupati and S. N. S. Jammula, "Intelligent Fault Classification in Transmission Lines Using Deep Learning Models," International Research Journal of Engineering and Technology (IRJET), vol. 13, no. 4, pp. 3826–3837, Apr. 2026
13. V. Daggupati, Power Quality Disturbance Classification in IEEE 9-Bus System Using STFT and Deep Learning. Lubbock, TX, USA: Texas Tech University, 2025. [Online]. Available: https://hdl.handle.net/2346/103168.
14. Swarnalatha, T., Supraja, B., Akula, A., Alubady, R., Saikumar, K., & Prasadareddy, P. (2024, July). Simplified framework for diagnosis brain disease using functional connectivity. In *2024 2nd World Conference on Communication & Computing (WCONF)* (pp. 01-06). IEEE.
15. Rao, K. M., Kumar, P. S., Reddy, T. V., Nilima, D., Saikumar, K., & Khlaif, A. M. (2024, April). Ultra low power high speed DFT implementation for ASIC SoC. In *2024 IEEE 9th International Conference for Convergence in Technology (I2CT)* (pp. 1-6). IEEE.