International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(5s) 2026 pp. 1243-1252
https://doi.org/10.70917/ijcisim-2026-5995
Article

Digital Twin-Based Intelligent Battery Management System for Predictive Diagnostics and Lifetime Optimization of Electric Vehicle Batteries

Dr.Ch.Vinay Kumar[1], Dr.P.Lakshmi Supriya[2], Dr.P.Ram Kishore Kumar Reddy[3]

[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
[3]Professor, Department of Electrical and Electronics Engineering, Mahatma Gandhi Institute of Technology (Autonomous), Hyderabad, prkumarreddy_eee@mgit.ac.in

Abstract: Conventional electric-vehicle battery management systems primarily estimate present operating states and react to threshold violations. They have limited ability to forecast degradation, evaluate alternative charging and thermal strategies, or distinguish normal aging from incipient faults. This paper proposes a digital twin-based intelligent battery management system (DT-iBMS) that synchronizes an edge-resident battery observer with a cloud-assisted virtual battery. The twin combines electrical, thermal, and aging models with machine-learning residual correction, uncertainty estimation, predictive diagnostics, and constrained lifetime optimization. Its outputs include state of charge, state of health, state of power, remaining useful life, fault probability, and recommended operating limits.

A proof-of-concept capacity twin was evaluated using the NASA Ames lithium-ion battery-aging dataset. Cells B0005, B0006, and B0007 were used for training, while B0018 was kept completely unseen for testing. The adaptive nonlinear twin estimated the held-out cell's discharge capacity with a mean absolute error of 0.0426 Ah, root-mean-square error of 0.0524 Ah, mean absolute percentage error of 2.69%, and coefficient of determination of 0.885. A static cycle-based degradation baseline produced 0.1036 Ah MAE and 6.73% MAPE. Thus, the proposed twin reduced capacity MAE by 58.8% on the held-out cell. A residual-based diagnostic layer and a model-predictive supervisory controller are presented for early fault identification and lifetime-aware charging, cooling, balancing, and power limitation. The results demonstrate feasibility at cell level but do not constitute an automotive lifetime-extension claim; pack-level drive-cycle tests, uncertainty calibration, fault-injection studies, and hardware-in-the-loop validation remain necessary.

Keywords: Digital twin; intelligent battery management system; electric vehicle; predictive diagnostics; remaining useful life; state of health; lifetime optimization; lithium-ion battery

1. Introduction

Lithium-ion batteries determine much of an electric vehicle's range, charging time, power capability, safety, residual value, and lifecycle cost. Their internal electrochemical and thermal states are only partially observable from terminal voltage, current, and surface temperature. Aging further changes usable capacity, resistance, heat generation, and dynamic response. A conventional BMS can protect the pack with calibrated thresholds and observers, but fixed models cannot fully represent cell-to-cell variability or continuously evolving degradation.

A battery digital twin is a synchronized virtual representation of the physical battery that is updated using operational data and can predict its future behavior. Unlike an offline simulation, a functional twin maintains a data connection to the asset, assimilates measurements, updates model parameters, and supplies diagnostic or control information. Battery-twin research covers real-time state estimation, remaining-useful-life prediction, dynamic charging, thermal management, balancing, and fault diagnosis [1].

Predictive diagnostics moves the BMS from reacting to a limit violation toward estimating the probability and lead time of a developing problem. Lifetime optimization goes further by evaluating candidate actions—such as reduced charge current, increased coolant flow, cell balancing, or temporary power derating—before applying them. A twin can simulate these actions without experimenting directly on the physical battery, but a safety-certified edge controller must remain authoritative when connectivity or the cloud is unavailable [2].

This paper contributes:

- A five-layer edge-cloud digital-twin architecture for EV battery monitoring, diagnosis, prognosis, and control [3].
- A hybrid model combining electrical-thermal physics, online parameter adaptation, and data-driven residual learning [4].
- A residual and uncertainty-based predictive diagnostic method for sensor, electrical, thermal, and degradation faults [5].
- A constrained model-predictive framework for balancing charging speed, performance, temperature, and aging [6].
- A reproducible proof of concept using a strict leave-one-cell-out split of the NASA aging dataset.
- A deployment roadmap that separates demonstrated results from proposed pack-level functions [7].

2. Related Work

Battery digital twins

Digital twins couple physical assets, virtual models, and bidirectional data exchange. For batteries, sensors supply voltage, current, and temperature, while the virtual space contains geometric, electrical, thermal, and aging models. State estimates and predictions can then update charging, thermal, and balancing controls. Cloud platforms offer storage and computational capacity beyond that of a conventional embedded BMS, but communication delay and availability prevent the cloud from replacing local protection [8].

Battery twins range from low-order equivalent-circuit models to electrochemical-thermal models and neural surrogates. High-fidelity physics improves interpretability but increases computational cost and parameter-identification burden. Purely data-driven models can approximate nonlinear behavior efficiently but may extrapolate poorly outside their training domain. Hybrid twins combine physical state equations with learned residual correction, providing a practical compromise between computational cost, consistency, and adaptability [9].

Prognostics and maintenance

The NASA battery dataset contains run-to-failure charge, discharge, and electrochemical-impedance records at different conditions. Its discharge data include voltage, current, temperature, time, and measured capacity, and the tests use a 30% rated-capacity fade from 2 Ah to 1.4 Ah as an end-of-life criterion. These properties make it suitable for proof-of-concept prognostics, though not sufficient for EV pack certification [10].

A Bayesian adaptive digital twin previously reported life-prediction error near 5% and estimated maintenance-cost reductions of 62.0% and 52.5% for two tested cells. Those maintenance-cost values are study-specific and cannot be transferred directly to an EV fleet. Another AI-driven twin combined semi-empirical modeling with AdaBoost and LSTM methods on NASA aging data. Such studies support adaptive prognostics, but performance comparisons require identical datasets, splits, endpoints, and metrics [11].

Predictive control

Digital twins can test candidate charging currents and thermal-control actions before issuing a setpoint. Published smart-BMS frameworks describe dynamic charging, thermal management, balancing, and full-lifecycle state tracking. More recent multi-tier architectures extend the twin from visualization to prediction, prescriptive optimization, and autonomous control while quantifying uncertainty.

The central control problem is multi-objective: rapid charging and high power improve usability but can increase heat and degradation. Recent digital-twin reinforcement-learning work reported a 12.5% charging-time reduction relative to selected baselines while constraining temperature across standard, low-temperature, high-temperature, and aged scenarios. That result demonstrates potential, not a universal expectation; safe deployment requires chemistry-specific constraints and independent validation [12].

System Architecture

Five-layer framework

The proposed DT-iBMS contains five functional layers.

Layer	Main elements	Functions
Physical battery	Cells, modules, busbars, contactors, coolant loop, chargers	Stores and supplies energy; receives control actions
Sensing and edge	Cell-monitor ICs, current and temperature sensors, BMS MCU	Sampling, synchronization, filtering, hard protection, fallback estimation
Twin and data	Equivalent-circuit, thermal and aging models; asset history	State assimilation, parameter adaptation, model versioning
Intelligence	Residual learning, anomaly detection, RUL and uncertainty models	Diagnostics, prognosis, operating-domain detection
Decision and service	Constrained optimizer, maintenance scheduler, dashboard	Charging, cooling, balancing, derating, service recommendations

The architecture intentionally divides computation between edge and cloud. Fast SOC estimation, contactor control, overvoltage, undervoltage, overcurrent, and overtemperature protection remain onboard. Computationally intensive parameter identification, fleet learning, long-horizon RUL forecasting, and scenario simulation may run in the cloud. A cloud outage must not disable safe vehicle operation [13].

Twin synchronization

At sample k, the physical battery provides:

$$y_k = [V_k, I_k, T_k, V_{cell,1:k}, T_{module,1:k}]. \#(1)$$

The virtual model evolves according to:

$$x_{k+1} = f(x_k, u_k, \theta_k) + w_k, \#(2)$$

$$\hat{y}_k = h(x_k, u_k, \theta_k) + v_k, \#(3)$$

where x contains SOC, polarization voltage, internal temperature, available capacity, and resistance states; u contains current, coolant command, and ambient conditions; θ contains time-varying model parameters; and w and v represent process and measurement uncertainty.

A Kalman, unscented, particle, or moving-horizon estimator can assimilate measurements. Online parameter adaptation updates resistance, time constants, capacity, and thermal coefficients. Parameter updates are accepted only when excitation is sufficient and the estimate remains within chemistry-specific bounds. This prevents noise from being interpreted as degradation [14].

Hybrid residual model

The residual between the measured and physics-model outputs is:

$$r_k = y_k - \hat{y}_k^{phys}. \#(4)$$

A data-driven model learns the normal residual:

$$\hat{r}_k^{ML} = g(z_{k-L:k}), \#(5)$$

where z includes recent voltage, current, temperature, SOC, SOH, resistance, ambient conditions, and balancing state. The corrected prediction is:

$$\hat{y}_k^{DT} = \hat{y}_k^{phys} + \hat{r}_k^{ML}. \#(6)$$

The physics model supplies conservation structure and interpretable states, while residual learning compensates for unmodeled nonlinearities. If the machine-learning input is out of distribution, the twin rejects or down-weights the learned correction rather than extrapolating without constraint.

Predictive Diagnostics

Residual health index

A normalized residual score is calculated as:

$$J_k = r_k^T S_k^{-1} r_k, \#(7)$$

where S_k is the expected residual covariance. Persistent growth in J_k, rather than a single spike, indicates mismatch requiring diagnosis. The diagnostic engine uses residual signatures, parameter drift, spatial cell patterns, and uncertainty to distinguish possible fault classes [15].

Fault class	Candidate indicators	Twin response
Voltage sensor	Bias, stuck value, disagreement with neighboring cells	Isolate sensor; estimate virtual value; service alert
Current sensor	SOC drift, pack-energy imbalance, redundant-sensor mismatch	Recalibrate or switch sensor; conservative SOC
Cell imbalance	Increasing voltage dispersion at matched SOC	Targeted balancing; reduced charge ceiling
Resistance growth	Load-dependent voltage residual and excess heat	Update SOH/SOP; reduce peak power
Cooling fault	Thermal residual, abnormal gradient, slow cooldown	Increase cooling; derate; inspect pump/valve

Internal short risk	Self-discharge, localized heat, abnormal voltage decay	Urgent isolation and shutdown logic
Capacity fade	Declining identified capacity and shortened discharge duration	Update RUL and service forecast

Battery digital-twin literature emphasizes diagnosis across sensor, electrical, thermal, and degradation faults. Twins can explore thermal-runaway boundary conditions through simulation, but they do not remove the need for direct hardware protection and multi-source sensing. Because rare failures are poorly represented in ordinary fleet data, a diagnostic model should not claim validated thermal-runaway prediction without controlled abuse-test evidence.[9][12]

Remaining useful life

Capacity-based SOH is:

$$SOH_n = \frac{C_n}{C_{rated}} \times 100\%. \#(8)$$

Given an end-of-life threshold C_{EOL}, the cycle-domain RUL is the smallest future horizon h satisfying:

$$RUL_n = min\{h > 0 : \hat{C}_{n+h} \leq C_{EOL}\}. \#(9)$$

A useful RUL output is a distribution, not a single cycle count. The twin should report median RUL, prediction interval, assumed future duty cycle, and confidence. Long-term degradation depends on temperature, SOC window, charge rate, depth of discharge, calendar time, and prior path; Oxford's path-dependent dataset specifically shows why calendar and cyclic aging cannot always be treated as independent.[13][14]

Maintenance decision

The maintenance scheduler minimizes expected lifecycle cost:

$$C(a) = C_{service}(a) + P_{fail}(a)C_{failure} + C_{downtime}(a) + C_{degradation}(a), \#(10)$$

where a is a candidate action. Actions may include continued operation, power derating, balancing, coolant-system inspection, module replacement, or pack replacement. The system should communicate the cause, uncertainty, urgency, and expected consequence rather than issuing an unexplained alarm.

Lifetime Optimization

Multi-objective formulation

At each decision instant, the twin predicts outcomes for candidate control sequences over horizon N. The controller minimizes:

$$\min_{u_{0:N-1}} \sum_{j=0}^{N-1} \left(\alpha t_{charge,j} + \beta \Delta SOH_j + \gamma E_{loss,j} + \delta \Delta T_j + \epsilon \sigma_{cell,j}\right), \#(11)$$

subject to:

$$V_{min} \leq V_{cell} \leq V_{max}, T_{min} \leq T_{cell} \leq T_{max}, |I| \leq I_{max}(SOC, SOH, T), \#(12)$$

$$SOP \geq P_{request}, P_{fault} \leq P_{safe}. \#(13)$$

Here, ΔSOH is predicted degradation, E_{loss} is energy loss, ΔT is thermal stress, and σ_{cell} is imbalance. The weights are selected for operating mode: fast charge, daily longevity, cold weather, emergency mobility, or fleet cost minimization.

Candidate actions

The optimizer can evaluate:

- Charge-current shaping rather than fixed current limits.
- Temperature-dependent preconditioning before rapid charging.
- Cooling-flow scheduling based on predicted internal rather than only measured surface temperature.
- Active balancing scheduled when it produces the greatest lifetime benefit.
- Regenerative-braking limits based on cell-level charge acceptance.
- Peak-power derating targeted to the weakest or hottest cell.
- User-selectable charge ceilings informed by forecast duty cycle.

The twin must not issue an unverified setpoint directly. A safety filter checks every proposed action against hard constraints, rate limits, actuator status, and edge-estimator confidence. The action is rejected if the cloud result is stale, model uncertainty is excessive, or the predicted trajectory conflicts with independent onboard models.

Experimental Method

Dataset

The proof of concept uses NASA Ames cells B0005, B0006, B0007, and B0018. The source dataset was collected on a custom prognostics testbed using 18650 lithium-ion cells, programmable loads and supplies, temperature sensing, an environmental chamber, and electrochemical-impedance equipment. Discharge records supply measured voltage, current, temperature, time, and capacity.[^5]

The processed cycle summary is:

Cell	Role	Discharge cycles	First capacity (Ah)	Final capacity (Ah)
B0005	Training	168	1.8565	1.3251
B0006	Training	168	2.0353	1.1857
B0007	Training	168	1.8911	1.4325
B0018	Held-out test	132	1.8550	1.3411

The entire B0018 trajectory was excluded from training. This leave-one-cell-out protocol evaluates transfer to an unseen physical cell and avoids leakage caused by randomly splitting correlated cycles from the same battery.

Feature extraction

For each discharge cycle, the following causal or cycle-completion features were extracted:

- Cycle index.
- Mean terminal voltage.

- Minimum terminal voltage.
- Mean measured temperature.
- Maximum measured temperature.
- Discharge duration.

Measured cycle capacity is the target. A histogram gradient-boosting regressor serves as the adaptive twin estimator. It used 300 maximum iterations, learning rate 0.05, 15 leaf nodes, L_2 regularization 0.2, and random seed 42. A pooled linear regression using cycle index alone serves as a static degradation baseline.

The model is called a proof-of-concept capacity twin because it updates the virtual health state from operational measurements. It is not a complete multiphysics pack twin. In deployment, the same prediction would be combined with online state assimilation, uncertainty, thermal states, and bidirectional control.

Metrics

Performance was assessed using:

$$MAE = \frac{1}{N}\sum_{i=1}^{N} |C_i - \hat{C}_i|, \#(14)$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N} (C_i - \hat{C}_i)^2}, \#(15)$$

$$MAPE = \frac{100}{N}\sum_{i=1}^{N} \left|\frac{C_i - \hat{C}_i}{C_i}\right|. \#(16)$$

The coefficient of determination R^2 was also calculated. All metrics were computed only on B0018.

3. Results

Capacity prediction

The adaptive capacity twin achieved 0.0426 Ah MAE, 0.0524 Ah RMSE, 2.69% MAPE, and $R^2 = 0.885$ on 132 unseen discharge cycles. The static cycle-only model achieved 0.1036 Ah MAE, 0.1104 Ah RMSE, 6.73% MAPE, and $R^2 = 0.488$.

Method	MAE (Ah)	RMSE (Ah)	MAPE	R^2
Adaptive digital-twin estimator	0.0426	0.0524	2.69%	0.885
Static cycle-based baseline	0.1036	0.1104	6.73%	0.488

The adaptive model reduced MAE by 58.8% and MAPE by 60.1% relative to the baseline. The improvement arises because voltage, temperature, and discharge-duration features carry information about present battery condition that cycle count alone cannot represent. The result supports synchronized, measurement-informed health tracking rather than a fixed fleet-average aging curve.

Diagnostic residuals

The capacity residual $e_n = C_n - \hat{C}_n$ is useful for diagnostic monitoring. A sudden residual shift can indicate a changed load protocol, sensor issue, capacity-estimation error, or abnormal degradation; a persistent trend can signal model drift. However, a residual is an anomaly indicator rather than a fault diagnosis by itself. Classification requires corroboration from voltage, current, temperature, impedance, balancing, and operating-context signals.

A practical alarm uses an adaptive threshold such as:

$$|e_n - \mu_e| > k\sigma_e \#(17)$$

for multiple consecutive updates, where μ_e and σ_e describe validated normal residuals. The value of k and persistence duration must be selected from fault-injection data to balance detection delay and false alarms. They were not optimized in this proof of concept.

Lifetime interpretation

The adaptive twin more accurately tracks the state variable used for capacity-based end-of-life decisions. That should improve RUL forecasts compared with a static aging curve, but this experiment does not demonstrate actual lifetime extension because no alternative charging or cooling action was applied to a controlled test group. Lifetime extension can only be quantified through prospective A/B cycling or validated high-fidelity simulation followed by experimental confirmation.

4. Discussion

Digital twin versus ordinary AI

A regression model alone is not a complete digital twin. A genuine operational twin additionally requires persistent asset identity, synchronized measurements, model-state updates, uncertainty handling, lifecycle data, and a channel through which validated insights influence maintenance or control. The proposed architecture supplies those missing system elements and uses machine learning as one component rather than treating prediction as the entire twin.

The leave-one-cell-out test is more credible than a random-cycle split, but cell B0018 still belongs to the same NASA campaign. Cross-laboratory validation is required. CALCE provides open lithium-ion data covering full and partial cycling, storage, dynamic driving profiles, open-circuit voltage, impedance, several form factors, and LCO, LFP, and NMC chemistries. Oxford provides long-term pouch-cell degradation and path-dependent NCA aging data. These datasets can test chemistry and duty-cycle transfer.

Edge-cloud partition

Cloud computation enables fleet learning and detailed simulation, but network latency, bandwidth, and interrupted connectivity can degrade synchronization.

- Hard protection and fast state estimation remain on the vehicle.
- Cloud outputs carry timestamps, model version, expiry time, and confidence.
- Control proposals pass through an onboard safety filter.
- The edge twin continues in degraded mode during disconnection.
- Fleet updates are staged, signed, monitored, and rollback-capable.

This arrangement prevents an unavailable or compromised cloud service from becoming a single point of failure.

Cybersecurity

Connectivity creates attack surfaces involving false sensor data, malicious model updates, unauthorized control commands, ransomware, and privacy leakage. Research on digital-twin governance identifies fragmented standards, interoperability, and cybersecurity as major deployment barriers. Battery twins can also help detect false-data injection by comparing measurements with physically consistent estimates.

A production system requires authenticated telemetry, encryption in transit and at rest, hardware roots of trust, signed models, secure boot, least-privilege cloud access, anomaly detection, audit logging, and safe rollback. Commands must be bounded locally even if they carry valid credentials. Cybersecurity evidence should be maintained alongside functional-safety and model-validation evidence.

Explainability and uncertainty

Maintenance personnel need actionable explanations. A diagnostic output should identify the affected cell or module, observed residual, likely causes, confidence, expected time to limit, recommended inspection, and consequence of delay. Feature attribution can support interpretation but does not prove physical causation.

Uncertainty should increase when operating outside the training envelope, after sensor faults, or when future duty cycle is unknown. The optimizer should become more conservative as uncertainty grows. Reporting a precise RUL without a prediction interval can mislead fleet operators and service planners.

Limitations

The study has important limitations:

- It evaluates cell-level laboratory data, not an EV pack under road loads.
- Only four cells from one test campaign were used.
- Capacity labels depend on full-discharge characterization unavailable during normal driving.
- The model does not estimate internal temperature, lithium plating, resistance-based SOH, or thermal-runaway probability.
- Predictive diagnostics are architecturally specified but not validated with labeled fault-injection data.
- Lifetime optimization is formulated but not experimentally applied to a control group.
- Communication, embedded execution, and cybersecurity were not tested on automotive hardware.

High-fidelity twins also face cost, sensing, processing, standardization, and data-availability barriers. Reviews note that fully onboard high-fidelity execution may be impractical under strict memory, power, latency, and reliability constraints, and that a unified battery-twin reference architecture has not yet emerged.

Implementation Roadmap

A rigorous continuation should proceed in five stages:

1. **Cross-dataset model validation:** Train and test across NASA, CALCE, and Oxford cells; report chemistry-specific and cross-chemistry errors.

2. **Hybrid pack twin:** Add a second-order equivalent-circuit model, lumped thermal network, capacity/resistance aging, and machine-learning residual correction.
3. **Fault validation:** Inject sensor bias, open-wire, cooling degradation, imbalance, abnormal self-discharge, and contact-resistance faults; measure detection probability, false-alarm rate, and lead time.
4. **Lifetime-control experiment:** Cycle matched cells under baseline and twin-optimized charging/thermal policies until a predefined endpoint; compare energy throughput, time, temperature, and capacity retention.
5. **Hardware-in-the-loop:** Measure worst-case execution time, network behavior, failover, model rollback, and control safety using the intended BMS controller and cloud interface.

An academically strong follow-up experiment should preregister the endpoint, number of cells, operating profiles, optimization objective, and statistical analysis. This prevents post-hoc claims and makes any reported lifetime improvement defensible.

5. Conclusion

A digital twin-based intelligent BMS has been proposed for predictive diagnostics and lifetime-aware EV battery management. The framework combines an edge safety controller, synchronized physical and data-driven models, residual-based diagnostics, probabilistic RUL forecasting, and constrained optimization of charging, cooling, balancing, and power limits.

On a held-out NASA cell, the proof-of-concept capacity twin achieved 2.69% MAPE and reduced MAE by 58.8% compared with a static cycle-based degradation model. This demonstrates that measurement-informed adaptation can track health more effectively than cycle count alone. The result supports the proposed architecture, but pack-level diagnosis and lifetime extension require prospective experiments. The recommended deployment path is a hybrid, uncertainty-aware edge-cloud twin in which AI enhances prediction while deterministic onboard protection remains authoritative.

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.