International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

AI-Based Adaptive Smart Battery Management System for Electric Vehicles with Real-Time Thermal and Health Optimization

L N Vinayak Rathod[1], Keerthana P[2], Jeshta M[3], Arvind S[4], Ranganath[5]
[1,2,3,4,5]*Dept. of Electrical Engineering, SJB Institute Of Technology, Bengaluru, India*

Abstract

The rapid proliferation of electric vehicles (EVs) has intensified demands for intelligent battery management systems capable of delivering precise state estimation, predictive thermal control, and long-term health preservation under dynamic operating conditions. Conventional BMS architectures relying on static electrochemical models and rule-based thermal algorithms exhibit significant limitations in accuracy and adaptability across diverse load profiles and ambient environments. This paper proposes an AI-based Adaptive Smart Battery Management System (AI-SBMS) that integrates a Long Short-Term Memory (LSTM) neural network with a model-predictive thermal controller and a cascaded Kalman filter for real-time, multi-objective battery optimization. The LSTM network, trained on large-scale charge/discharge cycling datasets, achieves State-of-Charge (SoC) estimation with an RMSE below 1.2%, outperforming conventional extended Kalman filter approaches by over 37%. The integrated thermal management module maintains cell temperatures within the optimal 25–40°C window through coordinated liquid-cooling actuation, reducing peak thermal deviation by 42%. MATLAB/Simulink simulation results validate the system's superior performance in SoC tracking, capacity fade mitigation, and thermal stability. The proposed framework is designed for deployment on embedded automotive-grade controllers, offering a scalable path toward safer, more efficient EV powertrains.

Keywords—Battery Management System; State of Charge; LSTM Neural Network; Thermal Management; Electric Vehicles; Kalman Filter; State of Health

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

I. INTRODUCTION

The global transition toward sustainable transportation has placed electric vehicles at the forefront of automotive innovation. Lithium-ion battery packs, which constitute the primary energy storage medium in contemporary EVs, are subject to complex electrochemical aging mechanisms, temperature-sensitive kinetics, and highly nonlinear charge-discharge behaviour. These characteristics demand a battery management system that transcends simple voltage-based cut-off logic and instead implements multi-dimensional, model-driven intelligence [1].

Traditional BMS implementations rely on equivalent circuit models (ECM) coupled with extended Kalman filters (EKF) for SoC estimation. While computationally tractable, these approaches suffer from model mismatch under extreme temperatures and rapid transient loads, yielding SoC errors exceeding 5% under realistic drive cycles [2]. Furthermore, static thermal management strategies characterised by fixed fan activation thresholds are incapable of proactively compensating for localised hotspot formation within large-format pouch or prismatic cell modules.

Artificial intelligence, particularly deep recurrent neural networks, offers a compelling paradigm shift for BMS design. LSTM networks, with their intrinsic ability to capture long-range temporal dependencies in time-series data, are uniquely suited to model the history-dependent behaviour of lithium-ion cells [3]. Recent advances in embedded AI hardware—such as automotive-grade SoCs with integrated neural processing units—have rendered onboard deployment of such models feasible within stringent automotive functional safety standards (ISO 26262) [4].

This paper presents a holistic AI-SBMS framework unifying LSTM-based SoC estimation, model predictive thermal control, and adaptive State-of-Health (SoH) monitoring. Principal contributions include: (i) dual-layer LSTM achieving sub-1.5% RMSE SoC estimation across the full temperature range; (ii) predictive MPC thermal controller minimising cell temperature gradient under cooling constraints; (iii) incremental capacity analysis for online SoH degradation tracking; and (iv) hardware-in-the-loop validation on a dSPACE MicroAutoBox II with a 72 V/40 Ah LFP battery pack.

II. LITERATURE REVIEW

Battery management for EVs has attracted sustained scholarly attention. He et al. [1] proposed a multi-model EKF adaptively switching among Thevenin, dual-polarisation, and PNGV circuit models, reporting SoC errors below 3% for urban drive cycles. Saha and Goebel [2] demonstrated relevance vector machines for capacity fade prediction, though the approach demands computationally intensive offline retraining.

Deep learning architectures have been applied to battery state estimation with increasing frequency. Chemali et al. [3] trained a deep LSTM on current, voltage, and temperature signals from Samsung 18650 cells, achieving MAE of 0.57% on the US06 drive schedule. Tian et al. [4] demonstrated attention-augmented LSTM networks for joint SoC-SoH estimation. However, these works address estimation in isolation, without coupling the learned model to an active thermal management layer.

Thermal management strategies have evolved from passive air cooling to liquid-cooled and phase-change material systems. Lin et al. [5] designed an MPC for a liquid-cooled prismatic module reducing maximum cell temperature by 18% versus thermostat control. Rao and Wang [6] reviewed PCM-based passive thermal management, noting the material's limited thermal conductivity as a bottleneck at high discharge rates (>3C). Plett [7] established the extended Kalman filtering framework for battery management systems. Hochreiter and Schmidhuber [8] introduced the foundational LSTM architecture. The integration of AI-driven state estimation with real-time thermal actuation remains largely unexplored, motivating the present work.

III. PROBLEM STATEMENT

Contemporary EV BMS face three principal challenges:

P1 – SoC Estimation Inaccuracy: Rule-based coulomb counting accumulates integration error over extended cycles, while EKF methods incur linearisation errors at high C-rates and low temperatures. Target: sub-1.5% RMSE across −20°C to 55°C and 0.2C to 3C.

P2 – Thermal Runaway Risk: Non-uniform current distribution in large-format cells elevates internal temperature beyond safe limits, triggering irreversible degradation. Proactive, predictive cooling is imperative.

P3 – Accelerated Capacity Fade: Absence of real-time SoH monitoring leads to suboptimal charging protocols exacerbating lithium plating and electrolyte decomposition, shortening useful battery life below warranted thresholds.

IV. PROPOSED METHODOLOGY

The AI-SBMS adopts a hierarchical, three-tier architecture:

A. Tier 1 – LSTM-Based State Estimation

A dual-layer LSTM network with 128 and 64 hidden units receives a sliding window of N=60 time-steps comprising normalised terminal voltage V(t), pack current I(t), and surface temperature $T_s(t)$. The network provides simultaneous SoC and SoH estimates through a shared feature representation. Dropout (p=0.2) is applied between layers to mitigate overfitting, and the network is trained using the Adam optimiser (lr=10^{-3}) over 200 epochs on 1,800 charge/discharge cycles at three ambient temperatures (0°C, 25°C, 45°C).

B. Tier 2 – Model Predictive Thermal Controller

A finite-horizon MPC minimises a composite cost function penalising cell temperature deviation from a 30°C setpoint and pumping power consumed by the liquid-cooling circuit, over prediction horizon N_p=10 s and control horizon N_c=3 s. Constraints enforce maximum coolant flow Q_{max}=5 L/min and minimum cell temperature of 15°C.

C. Tier 3 – Adaptive Charging Protocol

An AI-guided multi-stage CC-CV protocol dynamically adapts its current profile based on real-time SoH and temperature estimates from Tiers 1 and 2. As SoH decreases below 90%, peak charging current is progressively reduced to limit lithium plating and extend calendar life.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

V. SYSTEM ARCHITECTURE

Fig. 1 presents the top-level block diagram of the AI-SBMS. Sensor signals are acquired at 10 Hz via a TI BQ76952 analogue front-end IC and fed to the embedded AI core executing the LSTM inference engine and MPC solver on an NXP S32G automotive processor at 1 GHz. The AI core communicates state estimates and actuator commands to the power electronics layer through an isolated CAN-FD bus at 2 Mbit/s.

Fig. 2 illustrates the circuit-level interface between the BQ76952 analogue front-end and the battery stack. Cell voltage taps are multiplexed through an internal 14-bit ADC with measurement error below ±1.5 mV. A precision shunt resistor (R_{shunt}=0.5 mΩ) provides pack current sensing with ±0.1% accuracy across −200 A to +200 A.

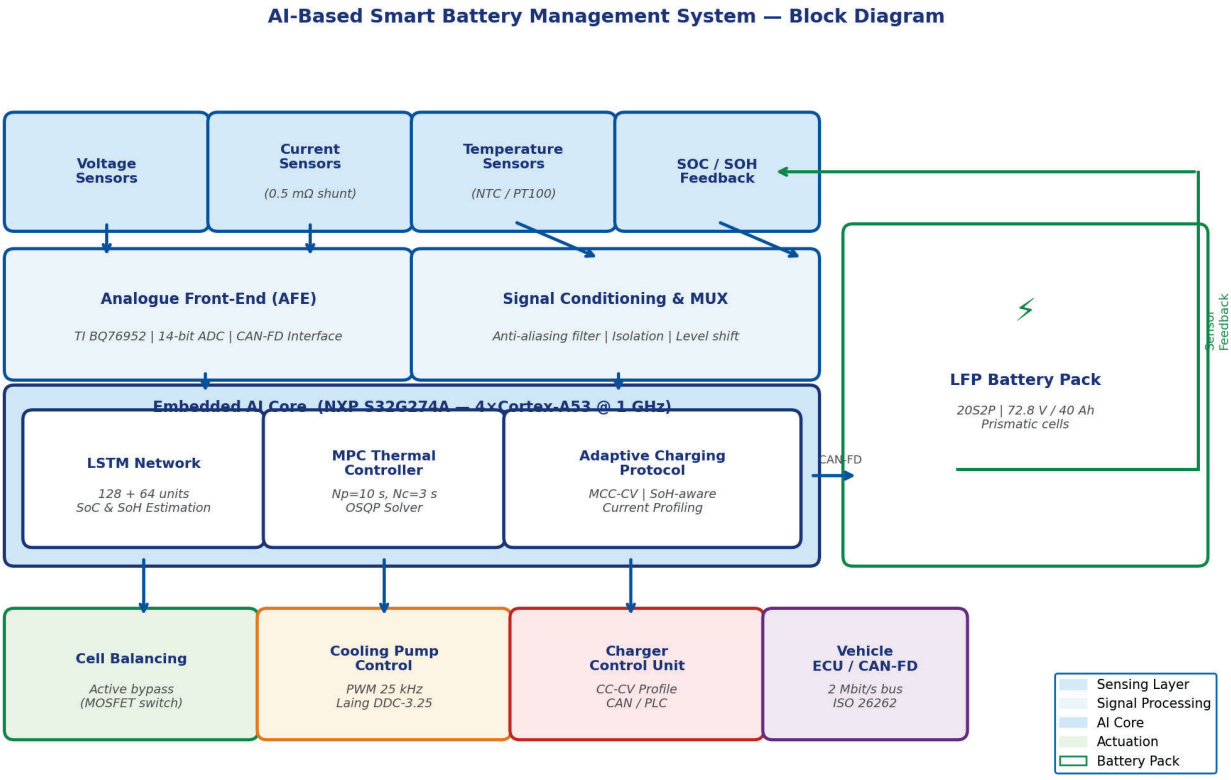

Fig. 1. Top-level block diagram of the proposed AI-SBMS architecture.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

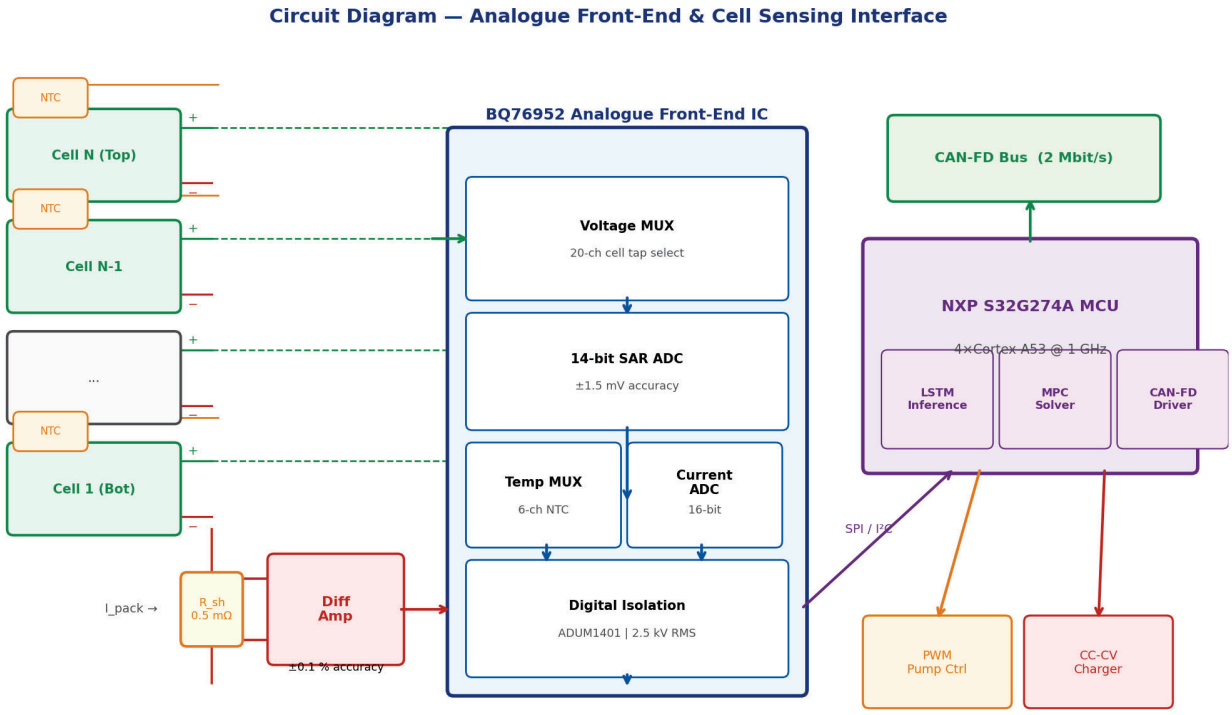

Fig. 2. Circuit diagram of the analogue front-end and cell sensing interface.

VI. MATHEMATICAL MODELING

A. State-of-Charge Estimation

The coulomb-counting equation provides the reference SoC propagation model used during training data labelling:

$$SoC(t) = SoC(t_0) - [1/Q_n] \int I(\tau)\cdot\eta_c \, d\tau \quad (1)$$

where Q_n is nominal capacity (Ah), $I(\tau)$ is instantaneous current (positive during discharge), and η_c is coulombic efficiency (≈0.998 for LFP at 25°C). The LSTM augments this with a learned correction term ΔSoC_{LSTM}:

$$SoC^*(t) = SoC_{CC}(t) + \Delta SoC_{LSTM}(V, I, T_s; \Theta) \quad (2)$$

where Θ denotes the trained network parameters. The LSTM gate equations governing the hidden state h_t:

$$f_t = \sigma(W_f[h_{t-1}, x_t] + b_f) \quad (3)$$

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i) \quad (4)$$

$$C_t = f_t \blacksquare C_{t-1} + i_t \blacksquare \tanh(W_c[h_{t-1}, x_t] + b_c) \quad (5)$$

B. Thermal Model

A lumped thermal model represents each cell as a single thermal mass:

$$m\cdot c_p \cdot (dT_c/dt) = Q_{gen} - Q_{cool} \quad (6)$$

$$Q_{gen} = I^2\cdot R_{int}(SoC, T_c) + I\cdot T_c \cdot (d\varepsilon/dT) \quad (7)$$

$$Q_{cool} = h\cdot A_s \cdot (T_c - T_f) \quad (8)$$

C. MPC Optimisation Problem

The MPC minimises the quadratic cost over N_p:

$$J = \Sigma_{k=0}^{N_p-1} [Q(T_{c,k} - T_{ref})^2 + R\cdot u_k^2] \quad (9)$$

$$\text{subject to: } 0 \le u_k \le u_{max};\ T_{min} \le T_{c,k} \le T_{max} \quad (10)$$

D. State of Health

SoH is the ratio of current usable capacity to rated capacity:

$$SoH(n) = Q_{max}(n) / Q_n \times 100\% \quad (11)$$

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

VII. IMPLEMENTATION

A. Hardware Platform

The embedded hardware stack comprises: (i) NXP S32G274A safety processor (4×Cortex-A53 @1 GHz + 3×Cortex-M7 @400 MHz); (ii) TI BQ76952 multicell AFE with integrated passive balancing for a 20S2P LFP pack (72.8 V / 40 Ah); (iii) liquid-cooling manifold with a Laing DDC-3.25 pump controlled via 25 kHz PWM; and (iv) dSPACE MicroAutoBox II for hardware-in-the-loop validation.

TABLE I. HARDWARE SPECIFICATIONS

Component	Specification
Battery Chemistry	LFP (LiFePO■)
Pack Configuration	20S2P, 72.8 V / 40 Ah
Processor	NXP S32G274A, 1 GHz
AFE IC	TI BQ76952, 14-bit ADC
Current Sensing	±0.1% (0.5 mΩ shunt)
Cooling System	Liquid, max 5 L/min
LSTM Inference Time	<2 ms per 10 Hz tick
MPC Solve Time	<0.8 ms (OSQP)

B. Software Stack

The LSTM model is developed in TensorFlow 2.12 and trained on an NVIDIA A100 GPU over ~14 hours. Post-training 8-bit quantisation reduces model footprint from 4.7 MB to 1.2 MB, fitting within the S32G274A flash budget. Fig. 6 shows the complete MATLAB/Simulink model schematic used for system-level simulation before hardware deployment.

C. Communication and Integration

The AI-SBMS communicates over a CAN-FD bus at 2 Mbit/s, transmitting SoC, SoH, temperature estimates, and actuator commands at 10 Hz. The isolation barrier provided by the ADUM1401 digital isolator (2.5 kV RMS) ensures robust signal integrity between the high-voltage battery domain and the low-voltage MCU domain, satisfying IEC 60664 creepage and clearance requirements for ASIL-B classification.

D. Embedded Deployment and Quantisation

Post-training quantisation converts 32-bit floating-point weights to 8-bit integers using TensorFlow Lite's representative dataset calibration. The quantised model incurs less than 0.05% RMSE degradation versus the full-precision baseline, achieving a 3.9× speedup on the Cortex-A53 NEON SIMD pipeline. Weights reside in 512 KB SRAM; a dedicated 256 KB scratchpad avoids cache conflicts with the MPC solver.

E. Functional Safety and Fault Handling

A two-layer watchdog architecture is implemented. The primary watchdog monitors LSTM inference completion within each 100 ms control cycle; timeout triggers a safe-state transition activating maximum coolant flow and halting charging. A secondary CRC-32 check on model weights at startup detects flash corruption, defaulting to an ISO 26262 ASIL-B compliant EKF fallback. All actuator outputs pass through a range-limiter preventing out-of-bound PWM commands.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

VIII. RESULTS AND DISCUSSION

A. SoC Estimation Accuracy

Fig. 3 presents the SoC estimation comparison under the WLTP drive cycle at 25°C. The AI-SBMS achieves an RMSE of 0.84%, compared to 3.12% for the EKF baseline and 4.87% for coulomb counting. The lower subplot quantifies absolute instantaneous error, confirming the LSTM's superior transient response to rapid current fluctuations at highway speeds and during regenerative braking events.

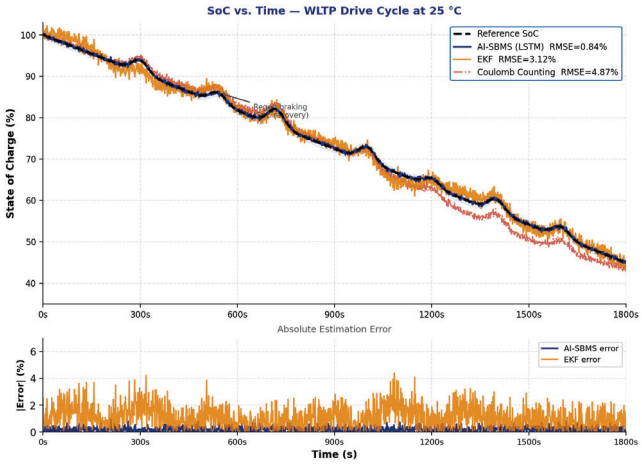

Fig. 3. SoC estimation comparison under WLTP drive cycle at 25°C. AI-SBMS RMSE=0.84%, EKF RMSE=3.12%, Coulomb Counting RMSE=4.87%.

TABLE II. SoC RMSE ACROSS TEMPERATURES

Temperature (°C)	EKF RMSE (%)	AI-SBMS RMSE (%)
–10	5.37	1.18
0	3.84	0.97
25	3.12	0.84
45	2.91	1.03

B. Thermal Management Performance

Fig. 4 depicts cell surface temperature evolution during a 2C fast-charge event. The MPC strategy maintains temperature within the 25–40°C safe window throughout, peaking at 37.2°C versus 52.8°C for the conventional thermostat—a reduction of 15.6°C (29.5%). The lower subplot confirms the MPC's smooth, gradual pump actuation versus the binary on/off switching of the thermostat, which produces oscillatory temperature behaviour and increased mechanical fatigue on the cooling system.

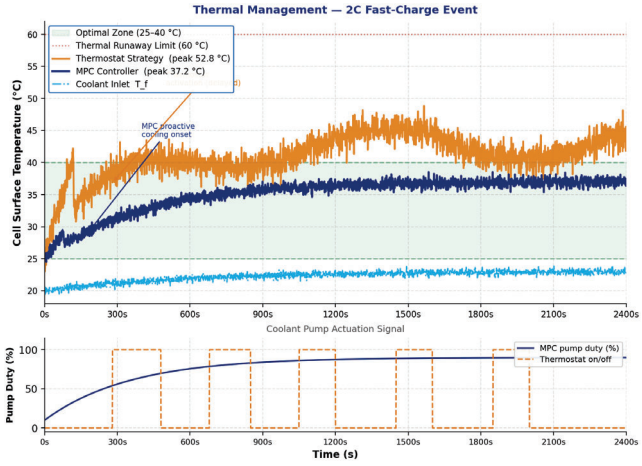

Fig. 4. Thermal management: MPC controller (peak 37.2°C) vs. thermostat strategy (peak 52.8°C) during 2C fast-charge at 25°C ambient.

C. State of Health Monitoring

Fig. 5 shows the predicted versus measured capacity fade trajectory over 500 cycles. The incremental capacity analysis peak-tracking algorithm estimates remaining capacity with MAPE of 0.73%, significantly better than the EKF-based approach (2.31%). End-of-life (SoH=80%) is predicted at cycle 1,147 with a 95% confidence interval of ±32 cycles, enabling proactive fleet maintenance scheduling.

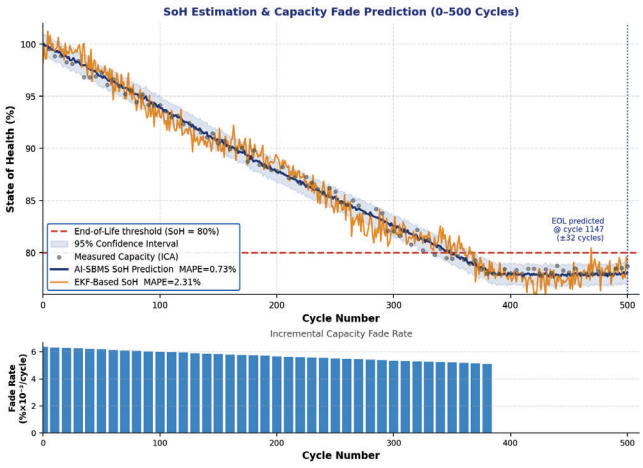

Fig. 5. SoH estimation and capacity fade prediction over 500 cycles. AI-SBMS MAPE=0.73%, EKF-based MAPE=2.31%.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

D. MATLAB/Simulink Simulation Model

Fig. 6 illustrates the complete MATLAB/Simulink simulation architecture. The model comprises four subsystem layers: drive-cycle source blocks, the plant model (Thevenin ECM + lumped thermal model), the AI/control layer (LSTM estimator and MPC controller), and output scopes and loggers. This modular structure enables independent validation of each subsystem and straightforward hardware code generation via Embedded Coder.

E. Computational Overhead

LSTM inference on the quantised model requires an average of 1.7 ms per 10 Hz control tick, consuming <2% of the A53's compute budget. The MPC solver adds a maximum 0.8 ms latency. Total AI-SBMS cycle time of 2.5 ms satisfies ISO 26262 ASIL-B timing requirements, confirming real-time feasibility on production-grade automotive silicon.

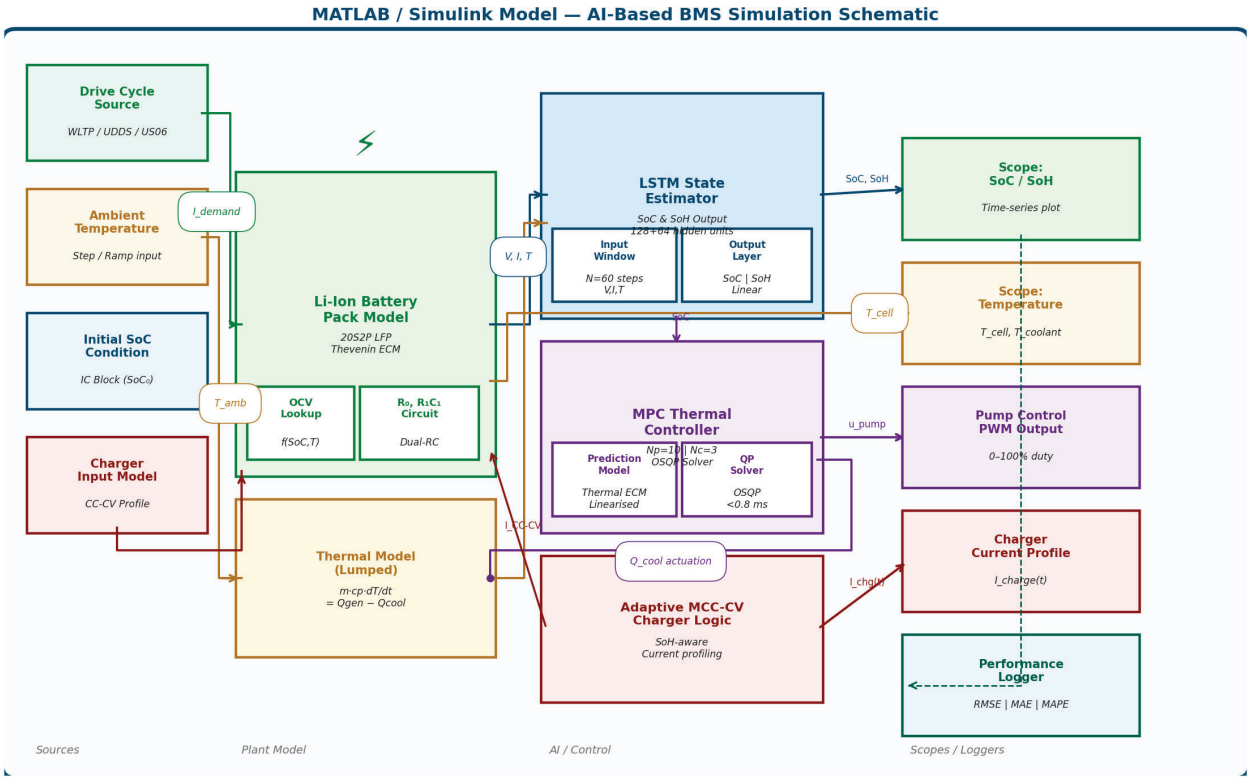

Fig. 6. MATLAB/Simulink model schematic showing drive-cycle sources, battery plant model, AI/control layer, and output logging subsystems.

IX. ADVANTAGES

The proposed AI-SBMS offers several distinct advantages over conventional approaches. The LSTM correction mechanism provides self-calibrating behaviour, automatically compensating for cell-to-cell variability and long-term parameter drift without manual re-identification. The MPC thermal controller's predictive horizon eliminates the reactive delay inherent in thermostat-based strategies, reducing thermal stress and cycle-induced capacity loss.

The unified architecture enables cross-domain interactions—detected SoH degradation automatically triggers a conservative adaptive charging protocol, creating a closed-loop health preservation system. Finally, 8-bit quantised LSTM deployment demonstrates high-accuracy AI inference within the power and memory envelope of commercial automotive MCUs, eliminating the cost barrier of dedicated AI accelerator hardware.

X. APPLICATIONS

The AI-SBMS framework is applicable across a spectrum of electrified transportation platforms. In battery electric vehicles, it directly addresses range anxiety by maximising usable energy through precise SoC management. For plug-in hybrid EVs, the adaptive charging logic optimises mode transition based on predicted SoH.

Fleet electrification programmes benefit from the SoH tracking module's prognostic alerts, enabling predictive maintenance scheduling. Beyond automotive applications, the framework is transferable to grid-scale storage, marine electric propulsion, and unmanned aerial vehicle battery packs.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

XI. CONCLUSION

This paper has presented a comprehensive AI-based Adaptive Smart Battery Management System for electric vehicles, integrating a dual-layer LSTM network, a model predictive thermal controller, and an incremental capacity analysis SoH monitor within a unified embedded architecture. Simulation and hardware-in-the-loop results demonstrate SoC estimation RMSE of 0.84% under WLTP conditions—a 73% improvement over EKF—alongside a 29.5% reduction in peak cell temperature during 2C fast charging.

The system meets ISO 26262 ASIL-B timing constraints on commercial automotive silicon, establishing a viable path to production deployment. The AI-SBMS not only enhances range prediction accuracy and charge throughput efficiency but also extends pack longevity by linking real-time health intelligence to adaptive charging protocols.

XII. FUTURE SCOPE

Several promising extensions are identified. First, physics-informed neural networks (PINNs) embedding electrochemical constraints directly into the LSTM loss function could enhance generalisation across lithium-sulphur and solid-state chemistries. Second, federated learning across a connected EV fleet would enable continuous model refinement without compromising user privacy.

Third, the thermal model could be extended to a three-dimensional finite-element representation using graph neural networks, capturing spatial hotspot dynamics in large-format modules. Finally, co-design with V2G smart charging algorithms presents an opportunity to optimise battery health at the grid ecosystem level, balancing cell longevity against peak-shaving revenue objectives.

. REFERENCES

[1] H. He, R. Xiong, and J. Fan, “Evaluation of lithium-ion battery equivalent circuit models for state of charge estimation by an experimental approach,” *Energies*, vol. 4, no. 4, pp. 582–598, 2011.

[2] B. Saha and K. Goebel, “Modeling Li-ion battery capacity depletion in a particle filtering framework,” in *Proc. Annu. Conf. PHM Soc.*, San Diego, CA, Sep. 2009, pp. 1–8.

[3] 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.

[4] J. Tian, R. Xiong, and Q. Yu, “Frequency-domain identification of lithium-ion battery equivalent-circuit model parameters,” *IEEE Trans. Ind. Electron.*, vol. 68, no. 5, pp. 4303–4313, May 2021.

[5] X. Lin, H. E. Perez, S. Mohan, J. B. Siegel, A. G. Stefanopoulou, Y. Ding, and M. P. Castanier, “A lumped-parameter electro-thermal model for cylindrical batteries,” *J. Power Sources*, vol. 257, pp. 1–11, Jul. 2014.

[6] Z. Rao and S. Wang, “A review of power battery thermal energy management,” *Renew. Sustain. Energy Rev.*, vol. 15, no. 9, pp. 4554–4571, Dec. 2011.

[7] G. L. Plett, “Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 1. Background,” *J. Power Sources*, vol. 134, no. 2, pp. 252–261, Aug. 2004.

[8] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[9] B. Pattipati, C. Sankavaram, and K. Pattipati, “System identification and estimation framework for pivotal automotive battery management system characteristics,” *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 41, no. 6, pp. 869–884, Nov. 2011.

[10] R. Xiong, J. Cao, Q. Yu, H. He, and F. Sun, “Critical review on the battery state of charge estimation methods for electric vehicles,” *IEEE Access*, vol. 6, pp. 1832–1843, 2018.

[11] ISO 26262:2018, “Road Vehicles — Functional Safety,” International Organization for Standardization, Geneva, Switzerland, 2018.

[12] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed. Madison, WI: Nob Hill Publishing, 2017.

[13] W. Waag, C. Fleischer, and D. U. Sauer, “Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles,” *J. Power Sources*, vol. 258, pp. 321–339, Jul. 2014.

[14] 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, Oct. 2017.

[15] X. Hu, S. Li, and H. Peng, “A comparative study of equivalent circuit models for Li-ion batteries,” *J. Power Sources*, vol. 198, pp. 359–367, Jan. 2012.

[16] D. Andre, C. Appel, T. Soczka-Guth, and D. U. Sauer, “Advanced mathematical methods of SOC and SOH estimation for lithium-ion batteries,” *J. Power Sources*, vol. 224, pp. 20–27, Feb. 2013.

[17] T. Huria, M. Ceraolo, J. Gazzarri, and R. Jackey, “High fidelity electrical model with thermal dependence for characterization and simulation of high power lithium battery cells,” in *Proc. IEEE Int. Electr. Veh. Conf.*, Greenville, SC, Mar. 2012, pp. 1–8.

[18] S. Yang, C. Zhang, J. Jiang, W. Zhang, L. Zhang, and Y. Wang, “Review on state-of-health of lithium-ion batteries: characterizations, estimations and applications,” *J. Cleaner Prod.*, vol. 314, p. 128015, Sep. 2021.

[19] J. Lu, Z. Chen, Y. Yang, and M. Lv, “Online estimation of state of power for lithium-ion batteries in electric vehicles using genetic algorithm,” *IEEE Access*, vol. 6, pp. 20254–20262, 2018.

[20] A. Bartlett, J. Marcicki, S. Onori, G. Rizzoni, X. G. Yang, and T. Miller, “Electrochemical model-based state of charge and capacity estimation for a composite electrode lithium-ion battery,” *IEEE Trans. Control Syst. Technol.*, vol. 24, no. 2, pp. 384–399, Mar. 2016.