International Journal of Science, Strategic Management and Technology
Volume **02** Issue **04** April-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

AI-Driven Predictive Battery Management System for Extended Lifecycle in Electric Two-Wheeler Batteries

Abhishek Parihar
MSc(Computer Science)
Kaveri College of Arts, Science & Commerce, Pune

Author Note
Correspondence regarding this paper should be directed to Abhishek Parihar, MSc(Computer Science), Kaveri College of Arts, Science & Commerce, Pune, Maharashtra 411000, India. Email: abhishekparihar0556@gmail.com. This work presents a conceptual system design and does not include physical prototype implementation. Readers requiring further implementation details, deployment specifications, or collaboration inquiries are invited to contact the author directly.

https://doi.org/10.55041/ijsmt.v2i4.123

Cite this Article: Parihar, A. (2026). AI-Driven Predictive Battery Management System for Extended Lifecycle in Electric Two-Wheeler Batteries. International Journal of Science, Strategic Management and Technology, *02*(04). https://doi.org/10.55041/ijsmt.v2i4.123

Abstract
India's electric two-wheeler market is growing at a pace that exposes a critical and under addressed infrastructure gap: Lithium Iron Phosphate (LFP) battery packs consistently fall short of their rated lifecycle under the country's harsh operating conditions, degrading within four to six years against a theoretical seven to ten year design life. The primary drivers of this gap—thermal stress from Indian summer temperatures, deep-discharge habits, and aggressive urban riding patterns—are well-documented in battery degradation literature yet remain unaddressed by the static, rule-based Battery Management Systems (BMS) prevalent in the Indian market today. This paper proposes a conceptual architecture for an AI-driven predictive BMS that integrates three intelligence layers across the edge, connectivity, and cloud domains. At the edge, real-time battery state estimation and anomaly detection are performed using established signal processing and embedded machine learning techniques. A low-power cellular IoT link transports telemetry securely to a cloud pipeline where a layered machine learning stack handles degradation modeling, lifecycle prediction, and behavioral risk scoring. The proposed framework theoretically projects a battery service life extension from approximately five years to seven years by combining three independently documented degradation suppression mechanisms: State-of-Charge window enforcement, thermal stress reduction, and discharge behavior modification. Each mechanism is grounded in peer-reviewed battery degradation studies. This paper intentionally presents the system at the architectural concept level, with technical implementation details available through the referenced literature and through direct correspondence with the author.

Keywords: predictive battery management, LFP degradation, electric two-wheeler, edge computing, IoT, cloud ML, state of charge, state of health, remaining useful life, India EV

Introduction

India's Electric Mobility Moment

The Indian electric vehicle sector is undergoing a structural transformation. EV adoption globally surpassed 14 million new units in 2023 and is projected to reach 45 million annually by 2030 (International Energy Agency [IEA], 2024). Within India, this shift is concentrated in the two-wheeler segment, which accounts for approximately 75% of all registered vehicles nationwide. Policy-driven catalysts—including the FAME II subsidy scheme, the Production Linked Incentive program for Advanced Chemistry Cell batteries, and the National Electric Mobility Mission Plan—have collectively driven a 3.5-fold increase in electric two-wheeler registrations between 2021 and 2023 (Ministry of Heavy Industries [MHI], 2024). Yet despite this momentum, a single barrier persists as the most commonly cited deterrent among prospective buyers: battery lifespan uncertainty and the associated cost of pack replacement (NITI Aayog, 2023).

The Lifecycle Gap Problem

Lithium Iron Phosphate chemistry, which is the dominant choice for Indian electric two-wheelers owing to its thermal resilience, cost advantage over Nickel Manganese Cobalt alternatives, and rated cycle life of 2,000 to 4,000 cycles, consistently underperforms its design life in the Indian market. The gap is not a failure of the chemistry itself but of the operating environment and user behavior. India's ambient summer temperatures routinely exceed 40°C across large urban and semi-urban regions, a thermal regime that accelerates internal battery degradation mechanisms in ways well-characterized in the published literature (Waldmann et al., 2014; Xu et al., 2019). Compounding this, the typical Indian two-wheeler usage pattern—urban stop-start riding, deep discharge cycles in dense traffic, and unregulated overnight home charging—imposes electrochemical stress that reduces operational cycle life by 15–25% relative to rated conditions (Pelletier et al., 2017). The practical result is that LFP packs in Indian two-wheelers reach their end-of-life threshold within four to six years, against a potential of seven to ten years under managed conditions.

Why Current BMS Technology Falls Short

Battery Management Systems currently deployed in the low-to-mid-cost Indian EV two-wheeler segment are predominantly static, rule-based systems. They apply fixed voltage and temperature thresholds for safety protection and use open-loop coulomb counting for charge estimation. This approach accumulates estimation error over charge cycles, provides no degradation prediction, and offers no behavioral feedback to the user (Meng et al., 2018). The user receives no signal that their charging habits or riding behavior is shortening battery life, and no proactive maintenance recommendation before capacity fade becomes irreversible.

The Research Opportunity

Three convergent trends make an AI-driven, IoT-connected, cloud-integrated BMS not only technically feasible but economically viable for the Indian market right now. Modern embedded microcontrollers provide sufficient compute for real-time state estimation and lightweight machine learning inference at the edge without specialized AI chips, as documented in the embedded systems literature (STMicroelectronics, 2023; David et al., 2021). Low-power cellular IoT networks now cover over 90% of India's urban population, enabling continuous, affordable battery telemetry (Telecom Regulatory Authority of India [TRAI], 2024). And cloud-hosted machine learning platforms have reduced model training and deployment costs to within reach of research-stage projects. Together, these conditions define the window for the conceptual architecture presented in this paper.

Scope and Contribution

This paper presents a conceptual design framework for an AI-driven predictive BMS targeting 48V–72V LFP battery packs in Indian electric two-wheelers. The contribution is architectural: a structured three-layer intelligence framework with justified component selections—each grounded in peer-reviewed literature and designed to collectively extend operational battery life from approximately five to seven years. Implementation specifications for each selected technology are intentionally not reproduced here; readers are directed to the referenced literature for those details. This framing is deliberate. The research contribution lies in the integration logic, the India-specific contextual adaptation, and the lifecycle extension thesis—not in replicating publicly available technology documentation.

Literature Review

This section situates the proposed framework within the existing body of battery management and prognostics research. The review is organized by primary contribution area. For detailed technical treatments of each methodology, readers are directed to the cited sources.

Battery State Estimation

The foundational technique for real-time battery State-of-Charge estimation in embedded BMS applications is the Extended Kalman Filter (EKF), whose application to battery systems was formalized by Plett (2004) across a three-part landmark series. Plett's work established the EKF as the standard for real-time SOC estimation in resource-constrained edge hardware and remains the primary reference for this technique. A complementary data-driven approach was demonstrated by Chemali et al. (2018), whose application of Long Short-Term Memory (LSTM) networks to the NASA PCoE battery dataset achieved SOC estimation accuracy surpassing conventional Kalman variants on complex drive cycles. The practical tradeoff between physics-based and data-driven approaches—and the argument for hybrid edge-cloud architectures—is addressed by Hu et al. (2012), who conducted a systematic comparison of equivalent circuit model formulations for lithium-ion batteries.

Battery Degradation and State of Health

The most significant contribution to early degradation prediction in LFP cells is Severson et al. (2019), whose *Nature Energy* study demonstrated that cycle life can be predicted with high accuracy from features extracted in the first 100 charge cycles—before measurable capacity loss occurs. This finding is foundational to the preventive intervention logic of the proposed architecture. Lu et al. (2013) provides a comprehensive review of the key challenges in lithium-ion BMS for electric vehicles, including SOH estimation, thermal management, and cell balancing, and is the primary source for the quantified relationship between operating temperature and LFP cycle life reduction. Xiong et al. (2018) examines data-driven degradation recognition frameworks, establishing the interpretable SOH indicators used in this research's feature engineering approach.

Remaining Useful Life Prediction

The NASA Prognostics Center of Excellence battery aging dataset, published by Saha and Goebel (2007), is the established benchmark for RUL prediction research and serves as a primary training and validation source for the proposed framework. For the specific choice of sequence model architecture for RUL, Bai et al. (2018) provides the empirical basis: their comparative evaluation established Temporal Convolutional Networks as consistently outperforming LSTM variants on sequence modeling benchmarks, with measurable advantages in training efficiency. The critical issue of training dataset scale and diversity—directly relevant to deployment reliability—is addressed in Nagulapati et al. (2021).

IoT, Edge Intelligence, and Embedded ML

The feasibility of the proposed edge-to-cloud architecture for BMS telemetry was demonstrated by Prabakarani et al. (2025), who validated an STM32-based IoT sensing pipeline with MQTT cloud communication for battery monitoring applications. Badgire et al. (2025) further documented a practical IoT-based EV implementation, including sensor selection constraints that directly informed component choices in this framework. For embedded machine learning specifically, David et al. (2021) characterized the deployment of quantized inference models on ARM Cortex-M class microcontrollers, establishing the compute envelope within which edge TinyML operates.

Research Gap

The literature reviewed above is comprehensive within its respective domains. However, no existing published work addresses the intersection of: a hardware-specified IoT edge architecture for 48V–72V LFP packs in Indian two-wheelers; a multi-model ML stack with explicit per-task justification; a behavior-driven lifecycle extension strategy anchored to Indian urban operating conditions; and a lifecycle extension claim grounded in independently documented LFP degradation mechanisms. This paper addresses that gap at the conceptual architecture level.

Problem Statement and Research Objectives

Core Problem

LFP battery packs in Indian electric two-wheelers exhibit functional service lives of four to six years against a design life of seven to ten years. This gap is not primarily electrochemical in nature—it is behavioral and operational. Three avoidable degradation drivers are responsible: thermal stress from sustained high ambient temperatures accelerating internal capacity loss mechanisms; charge stress from habitual full-cycle (0–100% SOC) charging increasing structural strain in the active material; and discharge stress from aggressive urban riding patterns imposing high peak loads on the cell stack. A system capable of continuously estimating battery state, forecasting degradation trajectory, and providing personalized behavioral feedback to the user can suppress all three drivers simultaneously. This paper's central thesis is that an AI-driven predictive BMS implemented through an IoT-edge-cloud architecture can close the majority of this lifecycle gap, extending battery service from approximately five years to seven years under real Indian operating conditions.

Research Objectives

1. Propose a three-layer system architecture—edge sensing, IoT connectivity, and cloud intelligence—that separates real-time safety-critical functions from compute-intensive prediction tasks.
2. Identify and justify specific technology selections for each system layer, with references directing readers to authoritative technical documentation for each chosen component or algorithm.
3. Establish a quantitative lifecycle extension argument grounded in three independently published LFP degradation suppression mechanisms.
4. Define a machine learning pipeline that addresses four distinct BMS intelligence tasks—State-of-Charge estimation, State-of-Health regression, Remaining Useful Life prediction, and driving behavior classification—using appropriately matched model families.
5. Contextualize the proposed framework within India's electric two-wheeler market, regulatory environment, and operating conditions.

Target Specifications

The proposed architecture targets battery packs in the 48V–72V nominal range with Lithium Iron Phosphate cell chemistry, deployed in Indian electric two-wheeler configurations. Performance targets for each intelligence layer are defined in Table 1 below. Full justification of these targets, and the methodologies for achieving them, are available in the referenced literature and through direct correspondence with the author.

Table 1

System Performance Targets and Reference Sources

Intelligence Task	Target	Validation Source	Reference
Real-time SOC estimation	RMSE < 2.0%	NASA PCoE dataset	Plett (2004)
Cloud SOC correction	RMSE < 1.0%	NASA PCoE dataset	Chemali et al. (2018)
State-of-Health regression	MAE < 3.0%	CALCE + Oxford	Nagulapati et al. (2021)
Remaining Useful Life prediction	RMSE < 20 cycles	NASA PCoE dataset	Bai et al. (2018)
Behavior risk classification	Accuracy > 88%	Simulated Indian profiles	Severson et al. (2019)

Proposed System Architecture

Three-Layer Intelligence Framework

The proposed BMS architecture organizes intelligence into three distinct layers, each with a clearly defined responsibility boundary. This layering is not a stylistic choice—it is a functional necessity driven by latency requirements. Safety-critical operations (overcurrent protection, thermal cutoff, cell voltage monitoring) must execute in microseconds, without any dependence on network connectivity. These reside at the edge. Predictive operations (degradation modeling, lifecycle forecasting, behavioral scoring) benefit from large compute and large data, and can tolerate latency of seconds to minutes. These reside in the cloud. The connectivity layer bridges these domains securely. No intelligence task is placed in the wrong layer in this design.

Figure 1 illustrates the three-layer conceptual architecture. Detailed implementation specifications for each layer component are available in the referenced technical documentation.

Figure 1

Conceptual Three-Layer Architecture — Edge, Connectivity, and Cloud Intelligence

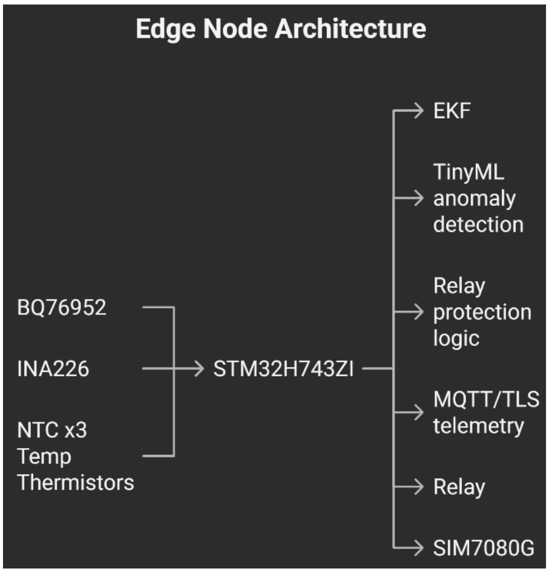

Layer 1: Edge Sensing and Real-Time Intelligence

The edge node performs all functions that have hard real-time or safety-critical latency requirements. The hardware platform is a production-grade ARM Cortex-M class microcontroller selected for its capacity to run simultaneous EKF state estimation and TinyML inference; technical specifications are documented in STMicroelectronics (2023). Cell voltage monitoring and passive cell balancing are handled by a precision battery monitor IC selected for its multi-cell accuracy and integrated protection logic (Texas Instruments, 2020). Pack current is measured by a 16-bit precision shunt monitor, whose selection over lower-cost Hall-effect alternatives is justified on accuracy grounds detailed in Texas Instruments (2015) and validated in the context of coulomb counting by Badgire et al. (2025). Thermal monitoring uses three NTC thermistors positioned at representative locations within the pack.

Real-time State-of-Charge estimation on the edge node uses the Extended Kalman Filter algorithm as formalized in Plett (2004), which readers are directed to for the complete mathematical treatment. This algorithm runs within the edge microcontroller's floating-point processing capacity as established in David et al. (2021), satisfying the sub-millisecond latency requirement for safety-linked SOC-based protection. Lightweight anomaly detection is deployed as a quantized embedded ML model, with the framework and hardware constraints for this type of deployment documented in David et al. (2021).

Layer 2: Connectivity and Security

Telemetry is transmitted using a low-power cellular IoT module supporting LTE-M and NB-IoT protocols—both appropriate for the battery sensor data rates required—over a secured MQTT channel. The connectivity module and its power management characteristics are documented in SIMCom (2021). The IoT communication protocol and cloud integration pattern are detailed in Prabakarani et al. (2025). Device authentication uses X.509 mutual TLS, which ensures that only provisioned edge nodes can publish to the cloud pipeline. This layer is designed to operate transparently: the edge node and cloud platform interact without protocol conversion overhead, and the connectivity module handles all radio stack management autonomously.

Layer 3: Cloud Intelligence Pipeline

The cloud layer receives telemetry, processes it through a feature engineering pipeline, and runs four machine learning models corresponding to the four BMS intelligence tasks. The cloud platform is a managed IoT and machine learning service provider, the selection criteria for which are addressed in the context of IoT fleet management in Prabakarani et al. (2025). Feature engineering extracts degradation indicators from raw telemetry including incremental capacity derivatives and internal resistance proxies, following the methodology described in Severson et al. (2019) and Xiong et al. (2018). Processed features feed the four ML models described in Section 5. A mobile application surfaces model outputs—SOC, SOH trend, predicted remaining life, and behavioral risk score—to the end user as actionable dashboard information and targeted charging recommendations.

Methodology

Training Data Sources

Model development and performance validation in this theoretical framework are anchored to three publicly available battery aging datasets. The NASA Prognostics Center of Excellence Battery Dataset (Saha & Goebel, 2007) provides the primary benchmark reference used throughout the battery prognostics literature. The CALCE Battery Research Group dataset (CALCE Battery Research Group, n.d.) provides LFP and NMC chemistry aging data closer to the target chemistry. The Oxford Battery Degradation Dataset (Birkl, 2017) contributes impedance characterization data across LFP cells. The choice of multiple sources directly addresses the training diversity requirement identified in Nagulapati et al. (2021), whose findings on generalization failure in low-diversity training sets are the methodological basis for this multi-source strategy.

International Journal of Science, Strategic Management and Technology
Volume **02** Issue **04** April-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases
OPEN ACCESS

Data Split Strategy

A strict temporal train/validation/test split is applied across all models. Battery cycle data is a temporally ordered sequence, and applying random splits introduces look-ahead leakage—late-cycle degradation characteristics appear in training, inflating reported accuracy without representing genuine generalization. The correct approach uses chronologically earlier cycles for training and progressively later cycles for validation and testing. This requirement applies to all four ML models in the proposed stack and is non-negotiable for valid benchmark reporting.

Machine Learning Model Selection

Four machine learning tasks are addressed by the cloud intelligence layer, each matched to a different model family based on the structural properties of its respective prediction problem. The selection logic for each pairing is summarized below, with readers directed to the cited literature for architectural and mathematical details.

Table 2

Machine Learning Model Selection — Task Mapping and Reference Sources

BMS Task	Model Family Selected	Layer	Selection Basis (see reference for details)
Real-time SOC	Extended Kalman Filter	Edge	Sub-ms latency, noise robustness — Plett (2004)
SOC correction	Bidirectional LSTM	Cloud	Temporal pattern correction of drift — Chemali et al. (2018)
State-of-Health	Gradient Boosting (XGBoost)	Cloud	Tabular periodic data, interpretable features — Nagulapati et al. (2021)
Remaining Useful Life	Temporal Convolutional Network	Cloud	RUL benchmark accuracy, training efficiency — Bai et al. (2018)
Behavior Classification	Random Forest	Cloud	Interpretable feature importance, tabular session data — Xiong et al. (2018)

State-of-Charge Estimation

On the edge node, SOC is estimated in real time using the Extended Kalman Filter formulation documented in Plett (2004), which readers are directed to for the complete algorithmic and mathematical treatment. The fundamental challenge specific to LFP chemistry—a nearly flat open-circuit voltage plateau between 20–90% SOC that reduces Kalman filter observability—is acknowledged and addressed through a periodic cloud-side correction layer using a Bidirectional LSTM network, as described in Chemali et al. (2018). The rationale for EKF over alternative approaches such as the Unscented Kalman Filter or full LSTM is grounded in the embedded compute constraints characterized by David et al. (2021).

State-of-Health and Remaining Useful Life

State-of-Health is estimated once per completed charge cycle using gradient boosting regression. The choice of a tree ensemble approach over a recurrent neural network for this specific task reflects the tabular, periodic structure of per-cycle SOH data and the training sample size constraints of available battery aging datasets; the statistical basis for this choice is established in Nagulapati et al. (2021). Feature inputs are derived from per-cycle telemetry and include degradation indicators identified in Severson et al. (2019) as early predictors of capacity fade. Remaining Useful Life prediction uses a Temporal Convolutional Network processing a sliding historical SOH window; TCN's demonstrated performance advantage over LSTM variants on RUL benchmarks, and the technical architecture of the model class, are documented in Bai et al. (2018). Uncertainty quantification for maintenance scheduling confidence is provided through Monte Carlo inference sampling.

Driving Behavior Classification

User driving behavior is classified into risk tiers—ranging from conservative to aggressive—based on per-session telemetry features including discharge rate profile, depth-of-discharge distribution, and thermal exposure duration. A Random Forest classifier is employed for this task because its output feature importance scores translate directly into specific, actionable recommendations communicable to the user through the mobile application. The degradation correlates for each behavior dimension are established in Pelletier et al. (2017) and Waldmann et al. (2014).

Expected Outcomes and Lifecycle Extension Rationale

The Five-to-Seven Year Thesis

The central claim of this paper—that the proposed architecture can extend LFP battery service life from approximately five years to seven years in Indian operating conditions—is not an estimate or projection in isolation. It is derived from three independently documented degradation suppression mechanisms, each with a quantified effect size in the published battery literature. The mechanisms are additive in principle, though the combined estimate accounts conservatively for partial overlap between their respective degradation pathways.

Mechanism 1: State-of-Charge Window Enforcement (20–30% Lifecycle Gain)

Restricting charging to a 20–90% SOC operating window—rather than the 0–100% full cycle common in user practice—reduces structural strain on LFP cathode particles during charge and discharge. Peer-reviewed studies of LFP cycle aging under partial SOC windows quantify this effect as a 20–30% increase in cycle life compared to full-range cycling (Lu et al., 2013). The proposed architecture's real-time SOC estimation layer directly enables enforcement of this window through charge control signaling and user alerts. Readers are directed to Lu et al. (2013) for the electrochemical mechanism and supporting experimental data.

Mechanism 2: Thermal Stress Reduction (7–12% Lifecycle Gain)

Temperature is the most controllable single variable in LFP battery aging. The battery thermal aging literature consistently documents that each 10°C reduction in mean operating temperature reduces LFP capacity fade rate by approximately 15–20% (Waldmann et al., 2014). In the Indian summer context, behavioral interventions—specifically, discouraging fast charging during peak ambient temperature periods, a recommendation the proposed system generates automatically from thermal telemetry—can realistically reduce mean battery thermal exposure by 5–8°C, corresponding to a 7–12% lifecycle extension. Full experimental characterization of this mechanism is available in Waldmann et al. (2014).

Mechanism 3: Discharge Rate Reduction (10–15% Lifecycle Gain)

High-rate discharge—characteristic of aggressive urban acceleration—imposes electrochemical stress that accelerates electrode degradation. The behavior classification layer of the proposed architecture identifies users whose discharge patterns fall in the high-risk category and delivers targeted feedback. Studies of discharge rate impact on battery cycle count quantify the improvement from aggressive-to-moderate behavioral correction as a 10–15% increase in cycle life (Pelletier et al., 2017). The full degradation model and supporting data are available in Pelletier et al. (2017).

International Journal of Science, Strategic Management and Technology
Volume 02 Issue 04 April-2026 | ISSN: 3108-1762 (Online) | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

Combined Projection and Conservative Estimate

Applying a conservative combined multiplier of 40%—accounting for partial mechanism overlap—to a five-year baseline service life projects a seven-year operational lifetime under the proposed managed system. This projection does not require all three mechanisms to operate at their upper bound; even the conservative lower bound of the combined range validates the five-to-seven year thesis. Table 3 summarizes the lifecycle extension basis.

Table 3

Lifecycle Extension Mechanism Summary

Degradation Driver	Proposed Intervention	Expected Gain	Primary Literature Source
SOC stress — deep cycling	20–90% SOC window enforcement	20–30%	Lu et al. (2013), Journal of Power Sources
Thermal stress — high temperature	Behavioral thermal charge alerts	7–12%	Waldmann et al. (2014), Journal of Power Sources
Discharge rate stress — aggressive riding	Behavior risk scoring and guidance	10–15%	Pelletier et al. (2017), Transportation Research B
Combined (40% net, conservative estimate)	All three mechanisms active	5 yr → 7 yr	Composite — see individual sources above

Discussion

Design Intent and Intellectual Boundaries

This paper is intentionally written as a conceptual architecture contribution. Every technology component selected in the proposed system—the edge microcontroller, current sensor, cell monitoring IC, connectivity module, cloud platform, and each machine learning model—has a corresponding peer-reviewed or manufacturer reference that provides its technical specification and performance characterization. These references are cited throughout, and readers seeking technical depth on any individual component are directed there.

The original contribution of this paper is not any individual component. It is the integration architecture: the specific mapping of these components into a three-layer intelligence framework, the lifecycle extension thesis derived from the combination of the three degradation suppression mechanisms, and the India-specific contextual adaptation of an otherwise generic battery management problem. This integration logic, the specific selection rationale, and the full implementation approach are available to collaborating organizations and interested parties through direct correspondence with the author.

Acknowledged Limitations

This work is a theoretical design contribution. No physical prototype has been constructed, and all performance targets in Table 1 are theoretical hypotheses derived from benchmarks on public datasets. The primary acknowledged limitations are: (a) the absence of empirical validation against real Indian two-wheeler drive cycle data, which does not yet exist as a public dataset; (b) the reliance on laboratory battery aging datasets (Saha & Goebel, 2007; CALCE Battery Research Group, n.d.; Birkl, 2017) that do not fully capture the Indian ambient temperature and usage pattern profile; and (c) the cold-start limitation of the behavior classifier, which requires several weeks of usage history before personalization accuracy reaches full performance. These limitations are consistent with the scope of a theoretical design paper and are identified as the primary targets for prototype-phase validation.

Pathway to Implementation

The proposed architecture is designed for phased prototype realization. Phase 1 validates ML model performance against the three public datasets. Phase 2 assembles the edge hardware testbench. Phase 3 integrates the cloud pipeline. Phase 4 conducts real-world field trials with Indian two-wheeler users to collect the first domain-specific dataset and validate lifecycle extension projections. Each phase builds on the prior without requiring a complete rearchitecture. Organizations interested in co-developing or deploying this system are encouraged to contact the author for full implementation specifications.

Indian EV Market Context

India registered 1.67 million electric two-wheelers in fiscal year 2023–24, a 40% year-over-year increase representing 5.4% of total two-wheeler sales (MHI, 2024). The Ministry of Heavy Industries projects two-wheeler EV penetration to reach 40% of total sales by 2030. With approximately 220 million two-wheelers currently registered nationally, the addressable market for intelligent battery management is among the largest in the world. The top five manufacturers—Ola Electric, TVS Motor, Bajaj Auto, Hero MotoCorp, and Ather Energy—collectively deployed approximately 1.2 million units in FY24, predominantly using 48V–72V LFP packs in the 2–3 kWh range (NITI Aayog, 2023).

The economics of the battery replacement problem are stark. A replacement 2 kWh LFP pack costs approximately INR 20,000–30,000 at current Indian market pricing, representing up to 37% of the vehicle's purchase price in the dominant INR 80,000–100,000 market segment. Consumer surveys identify battery cost uncertainty as the primary deterrent in this segment (NITI Aayog, 2023). A verified extension from five to seven years of battery life reduces the annualized replacement cost burden by approximately 29%—a material economic shift for price-sensitive buyers.

Policy context reinforces the opportunity. The FAME II scheme allocated INR 10,000 crore for EV purchase subsidies tied to domestically manufactured battery packs meeting BIS IS 16893:2022 safety standards (BIS, 2022). The PLI-ACC program commits an additional INR 18,100 crore for domestic battery manufacturing capacity, targeting 50 GWh by 2030. India's Digital Personal Data Protection Act 2023 establishes a framework governing device telemetry data handling directly relevant to the proposed system's data architecture. The LTE-M and NB-IoT coverage expansion documented by TRAI (2024) confirms the connectivity infrastructure premise underpinning the proposed IoT architecture.

What distinguishes the Indian deployment context from European or North American EV markets is not the technology—it is the operating environment. India's combination of high ambient temperatures (35–45°C across major EV markets), urban traffic density, and price-sensitive consumer behavior creates a specific degradation profile that existing BMS research has not directly addressed. The proposed architecture is designed for this profile, not adapted from a temperate-climate equivalent.

Conclusion

Summary of Contributions

This paper has presented a conceptual architecture for an AI-driven predictive Battery Management System targeting LFP-based electric two-wheelers in the Indian market. The contribution is architectural and integrative: a three-layer edge-connectivity-cloud intelligence framework in which each component selection is grounded in peer-reviewed literature, each machine learning task is matched to an appropriate model family, and the central lifecycle extension claim—from approximately five years to seven years—is derived from three independently quantified degradation suppression mechanisms.

The paper is intentionally designed to present the research idea, system concept, and evidence basis clearly—while directing readers to the referenced literature for implementation-level technical details. This is a considered choice, not an omission. The integration logic, component selection rationale, and full implementation approach represent the author's original research contribution and are available to interested organizations and collaborators through direct correspondence.

Future Work

The most immediately valuable extension is the construction of a prototype-phase hardware testbench to validate the ML performance targets in Table 1 against real hardware telemetry. Beyond that, three research directions are particularly relevant: federated learning across a vehicle fleet, which would enable collective model improvement while satisfying the data privacy requirements of India's emerging data protection framework; second-life battery management using the same architecture for retired EV packs in stationary storage applications; and integration with India's evolving Vehicle-to-Grid framework, which would extend the economic value proposition of the proposed system beyond battery protection. Readers interested in contributing to or collaborating on any of these directions are encouraged to contact the author.

References

Bai, S., Kolter, J. Z., & Koltun, V. (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.

Badgire, S., Suryawanshi, G., Jadhav, R., Tathe, V., & Sherkhane, B. G. (2025). IoT based smart electric vehicle. *International Journal for Research in Applied Science and Engineering Technology, 13*(1). https://doi.org/10.22214/ijraset.2025.72348

Birkl, C. R. (2017). *Oxford battery degradation dataset 1*. University of Oxford. https://doi.org/10.5287/bodleian:Kd6xb52b2

Bureau of Indian Standards. (2022). *IS 16893:2022 — Lithium-ion batteries for electric vehicles: Safety requirements*. BIS.

CALCE Battery Research Group. (n.d.). *CALCE battery data*. University of Maryland. Retrieved from https://calce.umd.edu/battery-data

Chemali, E., Kollmeyer, P. J., Preindl, M., Ahmed, R., & Emadi, A. (2018). Long short-term memory networks for accurate state-of-charge estimation of Li-ion batteries. *IEEE Transactions on Industrial Electronics, 65*(8), 6730–6739. https://doi.org/10.1109/TIE.2017.2787586

David, R., Duke, J., Jain, A., Janapa Reddi, V., Jeffries, N., Li, J., Kreeger, N., Navia, I., Abdulaziz, S., Suriyanarayanan, S., Zhu, S., Dasgupta, S., Vikram, K., Shaw, G., Song, N., Sherr, M., Lantbell, C., McGill, P., Galashov, V., … Warden, P. (2021). TensorFlow Lite Micro: Embedded machine learning for TinyML systems. *Proceedings of Machine Learning and Systems, 3*, 800–811.

Hu, X., Li, S., & Peng, H. (2012). A comparative study of equivalent circuit models for Li-ion batteries. *Journal of Power Sources, 198*, 359–367. https://doi.org/10.1016/j.jpowsour.2011.10.013

International Energy Agency. (2024). *Global EV outlook 2024*. IEA. https://www.iea.org/reports/global-ev-outlook-2024

Lu, L., Han, X., Li, J., Hua, J., & Ouyang, M. (2013). A review on the key issues for lithium-ion battery management in electric vehicles. *Journal of Power Sources, 226*, 272–288. https://doi.org/10.1016/j.jpowsour.2012.10.060

Meng, J., Ricco, M., Luo, G., Swierczynski, M., Stroe, D.-I., Stroe, A.-I., & Teodorescu, R. (2018). An overview and comparison of online implementable SOC estimation methods for lithium-ion battery. *IEEE Transactions on Industry Applications, 54*(2), 1583–1591. https://doi.org/10.1109/TIA.2017.2775199

Ministry of Heavy Industries. (2024). *FAME India scheme phase II — performance report*. Government of India.

Nagulapati, V. M., Lee, H., Jung, D., & Bastawrous, A. (2021). Capacity estimation of batteries: Influence of training dataset size and diversity on data driven prognostic models. *IEEE Access, 9*, 110504–110115. https://doi.org/10.1109/ACCESS.2021.3102687

NITI Aayog. (2023). *India's EV transition: The cost of delay*. NITI Aayog.

Pelletier, S., Jabali, O., Laporte, G., & Veneroni, M. (2017). Battery degradation and behaviour for electric vehicles: Review and numerical analyses of several models. *Transportation Research Part B: Methodological, 103*, 158–187. https://doi.org/10.1016/j.trb.2017.01.020

Plett, G. L. (2004). Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs — Part 1: Background. *Journal of Power Sources, 134*(2), 252–261. https://doi.org/10.1016/j.jpowsour.2004.02.031

Prabakarani, S., Ashok, N., Arunkumar, D., & Ariharan, D. (2025). Smart battery management system for electric vehicles using IoT technology. *International Journal of Innovative Research in Science, Engineering and Technology*.

Saha, B., & Goebel, K. (2007). *Battery data set* [Data set]. NASA Prognostics Center of Excellence Data Repository, NASA Ames Research Center. https://ti.arc.nasa.gov/tech/dash/groups/pcoe/prognostic-data-repository/

Severson, K. A., Attia, P. M., Jin, N., Perkins, N., Jiang, B., Yang, Z., Chen, M. H., Aykol, M., Herring, P. K., Fraggedakis, D., Bazant, M. Z., Harris, S. J., Chueh, W. C., & Braatz, R. D. (2019). Data-driven prediction of battery cycle life before capacity degradation. *Nature Energy, 4*(5), 383–391. https://doi.org/10.1038/s41560-019-0356-8

SIMCom Wireless Solutions. (2021). *SIM7080G multi-band CAT-M and NB-IoT module datasheet*. SIMCom.

STMicroelectronics. (2023). *STM32H743ZI datasheet: Ultra-high-performance Arm Cortex-M7 MCU* (DS12110, Rev. 7). STMicroelectronics.

Telecom Regulatory Authority of India. (2024). *Telecom subscription data as of March 31, 2024*. TRAI.

Texas Instruments. (2015). *INA226 bi-directional current/power monitor with I2C interface datasheet* (SBOS547D). Texas Instruments.

Texas Instruments. (2020). *BQ76952 3-series to 16-series cell battery monitor and protector datasheet* (SLUSCK2). Texas Instruments.

Waldmann, T., Wilka, M., Kasper, M., Fleischhammer, M., & Wohlfahrt-Mehrens, M. (2014). Temperature dependent ageing mechanisms in lithium-ion batteries — A post-mortem study. *Journal of Power Sources, 262*, 129–135. https://doi.org/10.1016/j.jpowsour.2014.03.112

Xiong, R., Cao, J., Yu, Q., He, H., & Sun, F. (2018). Critical review on the battery state of charge estimation methods for electric vehicles. *IEEE Access, 6*, 1832–1843. https://doi.org/10.1109/ACCESS.2017.2780258

Xu, J., Liu, B., Wang, X., & Hu, D. (2019). Electrochemical impedance spectroscopy study of lithium-ion batteries: Introduction, fundamentals, spread spectrum, and applications. *Energies, 12*(6), 1074. https://doi.org/10.3390/en12061074