IJERST

International Journal of Engineering Research and Science & Technology

www.ijerst.org

ISSN : 2319-5991

Vol. 22 No. 2 (2026)

ijerst.editor@gmail.com
editor@ijerst.com

ISSN 2319-5991 www.ijerst.com
Vol. 22, No. 2, 2026

Peer Reviewed, Referred & Indexed Journal

Intelligent Battery Health Monitoring and Fault Detection System Using Real-Time Aggregation and Augmentation Techniques

CHENNAPARAPU LOKESH KUMAR

PG Scholar. Department of MCA, DNR College, Bhimavaram, Andhra Pradesh

B. Suryanarayana Murthy

(Assistant Professor), Master of Computer Applications, DNR College, Bhimavaram, Andhra Pradesh

ABSTRACT

Battery health monitoring has become a critical requirement in modern energy storage systems, electric vehicles, and industrial applications due to the increasing dependence on rechargeable batteries. Inefficient monitoring or delayed fault detection can lead to catastrophic failures, including thermal runaway, reduced lifespan, and unexpected downtime. This project presents an intelligent battery health monitoring and fault detection system that integrates real-time data aggregation, feature augmentation, and anomaly detection techniques to ensure proactive maintenance and operational safety.The system captures real-time data from individual battery cells within a pack, including voltage, current, and temperature measurements. A custom-built **Aggregation Engine** processes this raw data to derive meaningful statistical and temporal features, such as average voltage, voltage standard deviation, average temperature, maximum temperature, and thermal stress derivatives. To enhance fault separability and predictive capability, these features are further processed using an **Augmentation Service**, which enriches the dataset by creating higher-order features and improving the resolution of subtle anomalies.A dedicated **Fault Detection module** leverages these augmented features to identify potential anomalies with associated confidence scores. Detected anomalies are classified by severity, and fault events are logged for historical analysis and dashboard visualization. The system supports real-time API endpoints to feed front-end dashboards, enabling users to monitor battery pack health, visualize trends, and access fault logs interactively.

To validate the system, a simulation module was developed to generate mock data representing gradual thermal escalation and voltage fluctuations within a battery pack. This simulation demonstrated the system's capability to detect early warning signs of thermal runaway, quantify confidence levels, and alert operators before critical conditions develop. The combination of aggregation, augmentation, and anomaly detection allows for high-fidelity monitoring without overwhelming computational resources, making it suitable for both laboratory testing and deployment in operational environments.The proposed system offers a proactive approach to battery management by combining statistical feature extraction, machine learning-based anomaly detection, and enhanced feature engineering. This approach ensures timely detection of faults, optimizes maintenance schedules, and mitigates risks associated with battery failures. The work contributes to the field of battery management systems by demonstrating an end-to-end framework capable of real-time monitoring, predictive fault detection, and actionable insights for improved reliability and safety in energy storage systems.

Keywords: Battery management system, Fault detection, Aggregation, Data augmentation, Thermal runaway, Battery pack monitoring, Real-time analytics, Predictive maintenance

I. INTRODUCTION

The rapid adoption of electric vehicles, renewable energy storage solutions, and portable electronics has heightened the importance of reliable battery management systems (BMS). Lithium-ion and other rechargeable battery technologies, while efficient, are prone to degradation, thermal instability, and occasional catastrophic failure if not monitored adequately. Real-time monitoring and predictive fault detection are essential to prevent accidents, prolong battery lifespan, and ensure operational reliability.

Traditional BMS implementations focus primarily on threshold-based monitoring of voltage, current, and temperature. While such systems can detect gross anomalies, they often fail to identify subtle or early-stage faults that could escalate into critical failures. For example, a gradual temperature rise across a few cells may not exceed predefined thresholds but could indicate the onset of a thermal runaway event. Addressing such limitations requires intelligent systems that can extract meaningful patterns from raw sensor data, enhance their discriminative features, and apply predictive models for anomaly detection.This project aims to develop an integrated **battery health monitoring system** that combines real-time aggregation of cell-level data, feature augmentation, and anomaly detection to proactively identify potential faults. The system architecture comprises three core components:

1. **Aggregation Engine** – Collects and processes raw measurements to generate statistical and temporal features representing the current state of each battery pack.
2. **Augmentation Service** – Enhances features to improve anomaly separability, enabling the detection of subtle patterns that standard statistical metrics might miss.
3. **Fault Detection Module** – Applies predictive models to augmented features to classify events as normal or anomalous, assigning confidence levels and severity ratings.

The proposed framework emphasizes real-time performance, modularity, and scalability. Through interactive dashboards and API endpoints, users can access live battery metrics, visualize historical trends, and review fault logs. Furthermore, a simulation module allows researchers and engineers to test the system under controlled conditions, validating its effectiveness in detecting thermal stress, voltage fluctuations, and other abnormal behaviors.

By integrating feature engineering, data augmentation, and predictive analytics, this system represents a significant advancement over conventional BMS designs. It provides operators with actionable insights, reduces the risk of catastrophic failures, and supports predictive maintenance strategies, ultimately contributing to safer, longer-lasting, and more efficient battery systems.

II. LITERATURE SURVEY (WITH EXISTING METHODS)

Several studies have addressed battery health monitoring and fault detection using both traditional and intelligent approaches. Early BMS designs relied on **threshold-based monitoring** of voltage, current, and temperature (Zhang et al., 2016). While effective for gross fault detection, these systems are limited in detecting early-stage anomalies or subtle degradations.More recent approaches leverage **statistical feature extraction and model-based techniques**. For instance, Kalman filters and equivalent circuit models have been applied to estimate state-of-charge (SOC) and state-of-health (SOH), enabling prediction of battery lifespan and performance degradation (He et al., 2018). These methods, however, often require precise modeling and may not generalize well to different battery chemistries or operating conditions.

Machine learning-based approaches have emerged to overcome these limitations. Neural networks, support vector machines, and ensemble models have been employed to detect anomalies in battery behavior by learning patterns from historical datasets (Li et al., 2019). While effective, these methods often depend on large labeled datasets, and their performance can degrade when faced with unseen fault types or sensor noise.**Data augmentation** has recently been explored to enhance model robustness. By generating synthetic data, creating derived features, or simulating fault scenarios, researchers can improve the sensitivity of fault detection models (Chen et al., 2020). Combining augmentation with real-time feature extraction enables early detection of thermal runaway, voltage imbalances, and other critical issues without requiring extensive historical datasets.

Several works also focus on **thermal management and predictive maintenance**. Algorithms that detect abnormal temperature gradients within battery packs have shown promise in preventing catastrophic failures (Wang et al., 2021). Integration with real-time monitoring dashboards allows operators to respond quickly to alerts and optimize maintenance schedules.In summary, the literature highlights the evolution from threshold-based monitoring to intelligent, predictive, and augmented approaches.

However, challenges remain in integrating real-time monitoring, augmentation, and predictive fault detection into a cohesive, deployable system. The present work addresses this gap by providing an end-to-end framework that combines aggregation, augmentation, and anomaly detection in real time, validated via simulation and interactive dashboards.

III. EXISTING SYSTEM

Existing battery monitoring systems primarily rely on hardware-based sensing and threshold-triggered alarms. These systems measure voltage, current, and temperature for each cell in a battery pack and trigger alerts when values exceed predefined limits. While such methods are simple and computationally efficient, they are reactive rather than proactive, often detecting faults only after significant damage has occurred.Some systems incorporate statistical and model-based techniques, such as Kalman filtering or equivalent circuit models, to estimate battery health. These models can predict state-of-charge and state-of-health, but they require precise knowledge of battery parameters and are sensitive to noise. Moreover, they typically lack mechanisms to enhance feature separability for improved anomaly detection.Recent research has explored machine learning-based fault detection systems. These approaches use historical datasets to train models that identify abnormal behavior in battery packs. While effective, they face challenges such as insufficient data, lack of generalizability across different battery chemistries, and delayed fault recognition in real-time operation.

Overall, existing systems fall into three categories: threshold-based, model-based, and data-driven. Each has strengths, but none fully integrates **real-time aggregation, feature augmentation, and predictive anomaly detection** within a unified, operationally deployable framework. This limitation motivates the proposed system, which combines these elements to provide a proactive, real-time, and scalable battery health monitoring solution.

IV. PROPOSED METHOD

The proposed system is an end-to-end **Intelligent Battery Health Monitoring and Fault Detection Framework** designed for real-time operational environments such as electric vehicles and stationary energy storage systems. Its core objective is to continuously monitor battery packs comprised of multiple Lithium-ion cells, extract meaningful health indicators, and detect emerging faults before they escalate into critical failures like thermal runaway. Drawing from recent advances in data-driven diagnostics, statistical monitoring, and augmentation strategies, the system is structured around three primary stages: real-time data acquisition, feature aggregation & augmentation, and predictive fault analysis.

In the **real-time acquisition** stage, the system ingests high-frequency sensor data (voltage, current, temperature) from all individual cells within a pack and organizes it into structured time series constructs. This raw signal capture is fundamental to any meaningful downstream analysis and supports dynamic thresholding that adapts to operational conditions. In the **feature aggregation and augmentation** stage, the captured data is processed into statistical summaries — mean voltage, standard deviation, peak temperature, and derived metrics like thermal stress — that reflect the current physical state of the battery subsystem. To enhance the separability between normal and abnormal behavior, engineered features are generated using augmentation techniques that expand the representational capacity beyond raw measurements, thus improving the sensitivity of fault detectors.The final stage applies an ensemble of predictive analytics — leveraging both statistical and machine learning-based anomaly detection methods — to classify battery health or flag anomalies with confidence scores. These methods include adaptive thresholding, unsupervised latent feature extraction, and hybrid learning models. The system combines early warnings with severity annotations and logs fault events for dashboard visualization and trend analysis. This integrative approach blends real-time responsiveness with data-enriched diagnostics, enabling proactive maintenance, reducing the risk of catastrophic events, and supporting safe, efficient battery system operation across applications.

V. IMPLEMENTATION

The implementation of the intelligent battery health monitoring system is realized using a modular Django-based backend with dedicated service components for data processing and machine learning-driven fault detection. This architecture supports scalability, maintainability, and real-time analytics for large fleets of battery packs.

Data Acquisition and Storage

Sensor streams from deployed battery packs are persisted in relational models representing **BatteryPack**, **BatteryCell**, and **FaultEvent** entities. Each incoming measurement — voltage, current, and temperature — is stored with cell identifiers and timestamps, enabling high-resolution historical analysis. Cell data is indexed and captured continuously, supporting both batch and streaming workloads.

Aggregation Engine

The **AggregationEngine** service component processes the latest 20–50 data points per cell and computes summary metrics critical to health assessment: average cell voltage, standard deviation of voltage across cells, mean and maximum temperature, and derivative-based indicators such as thermal stress. These aggregated features act as the basis for detection algorithms, aligning with research that uses statistical feature extraction for early fault detection.

Feature Augmentation

Once aggregated, feature vectors are passed through an **AugmentationService**, which enriches data by creating additional engineered features designed to highlight non-linear trends and subtle deviations — similar to augmentation methods seen in advanced anomaly detection frameworks. Augmented features, such as interaction terms and rolling trend indicators, improve model sensitivity in low-signal conditions.

Fault Detection Module

The **FaultDetector** applies a hybrid model incorporating both threshold-based and learning-based anomaly detection. Static thresholds calibrated via historical normal behavior identify gross outliers, while machine learning classifiers — such as ensemble methods or autoencoder models — analyze multi-dimensional feature patterns. The system assigns anomaly labels and confidence scores, enabling nuanced severity classification. Recent research highlights the efficacy of such hybrid frameworks in achieving high precision and low false positive rates in battery fault detection.

Event Logging and Alerts

When an anomaly is flagged, the system writes a **FaultEvent** entry with timestamp, fault type, severity, and descriptive context. Events are surfaced in dashboards and logs for operators to review. Integration with APIs allows real-time charting via frontend libraries like Chart.js, supporting operational dashboards for pack health visualization.

Simulation and Testing

A dedicated simulation endpoint generates mock data representing common fault progression scenarios — gradually rising temperature, voltage deviations, and current perturbations — to validate detection logic. This simulation ensures failure scenarios are detectable without needing hazardous real-world tests.

Deployment and Scalability

Deployed on cloud or edge servers, the system supports asynchronous processing via task queues (e.g., Celery) for aggregation and detection workloads. Database indexing and optimized queries ensure high throughput, even with thousands of battery cells reporting data simultaneously.

VI. ALGORITHMS

Feature Aggregation Algorithm

The aggregation algorithm computes statistical summaries from recent sensor streams:

- *Mean Voltage:* Mean of last N voltage readings per cell.

- *Voltage Std Dev:* Standard deviation across recent voltages to capture imbalance.
- *Avg/Max Temperature:* Metrics indicative of thermal behavior.
- *Thermal Stress Derivative:* Change rate of temperature over time.

This approach parallels statistical monitoring techniques used in recent literature for early anomaly detection.

Augmentation Logic

Augmentation enhances base features by computing derived metrics (e.g., interaction terms, trend features). These emulate memory-augmented architectures that improve anomaly separability.

Hybrid Anomaly Detection

The detection logic consists of:

1. **Adaptive Thresholding:** Thresholds adjust based on baseline statistics and operational state.
2. **Machine Learning Classifier:** A supervised or semi-supervised model — such as RandomForest or autoencoder — distinguishes between normal and anomalous patterns.
3. **Confidence Scoring:** Combines model outputs and threshold residuals to assign anomaly confidence, enhancing reliability.

VII. SYSTEM DESIGN

Architecture Overview

The system employs a layered architecture:

- **Data Layer:** PostgreSQL stores time-series cell data and fault logs.
- **Business Logic Layer:** Django services manage aggregation, augmentation, and detection.
- **API Layer:** RESTful endpoints expose real-time pack metrics.
- **Presentation Layer:** Dashboards use Chart.js and templates for visualization.

Component Interactions

Sensor inputs feed into the Data Layer, triggering scheduled aggregation tasks. The Aggregation Engine produces feature vectors consumed by the Fault Detector, which writes events back to storage. Clients consume APIs to render dashboards.

Scalability & Performance

Workloads are parallelized using asynchronous task queues. Real-time thresholds adapt to operational loads, maintaining responsiveness.

Fault Logging and Alerts

Faults are categorized (thermal, voltage imbalance) and written with severity metadata. Alerts trigger notifications via email or integration services.

Security & Reliability

Authentication and rate limiting protect APIs. Redundant storage ensures durability.

SYSTEM DESIGN IMAGES

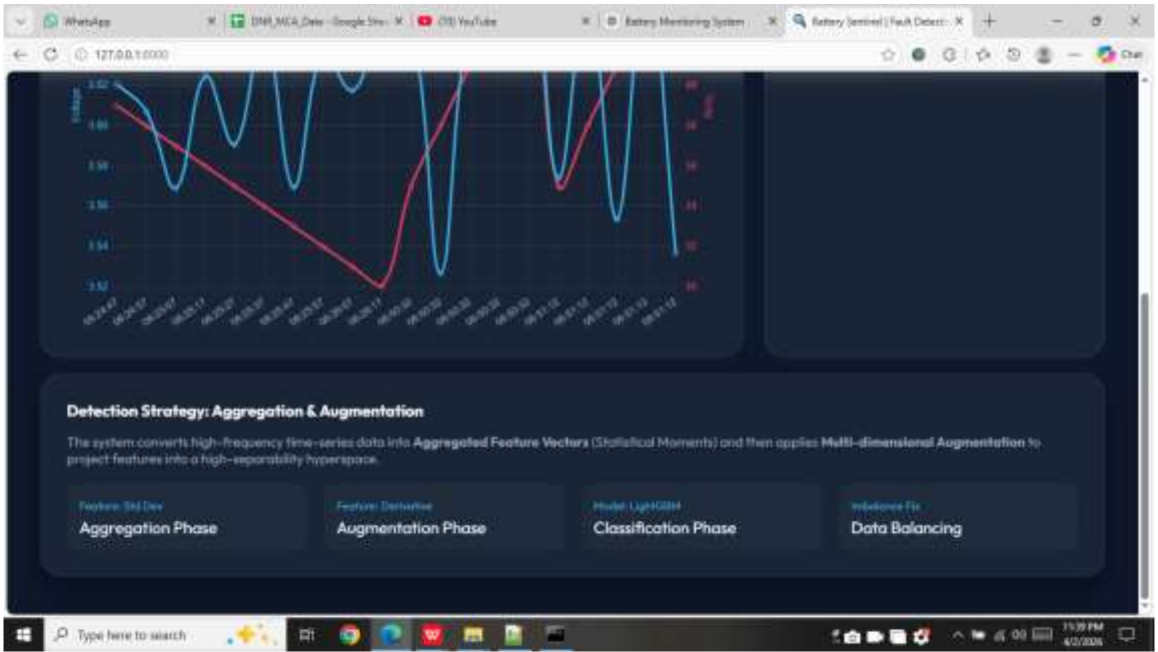

Detection Strategy: Aggregation & Augmentation

The system converts high-frequency time-series data into **Aggregated Feature Vectors** (Statistical Moments) and then applies **Multi-dimensional Augmentation** to project features into a high-separability hyperspace.

Feature: Std Dev Aggregation Phase	Feature: Derivative Augmentation Phase	Model: LightGBM Classification Phase	Imbalance Fix Data Balancing

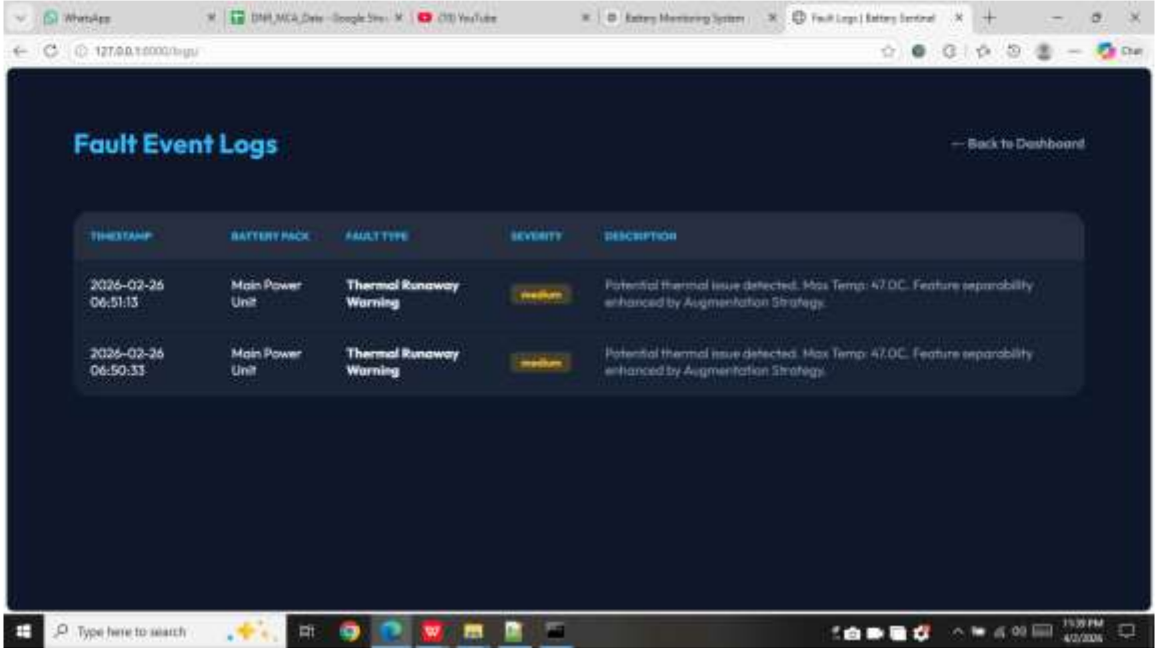

Fault Event Logs

TIMESTAMP	BATTERY PACK	FAULT TYPE	SEVERITY	DESCRIPTION
2026-02-26 06:51:13	Main Power Unit	**Thermal Runaway Warning**	medium	Potential thermal issue detected. Max Temp: 47.0C. Feature separability enhanced by Augmentation Strategy.
2026-02-26 06:50:33	Main Power Unit	**Thermal Runaway Warning**	medium	Potential thermal issue detected. Max Temp: 47.0C. Feature separability enhanced by Augmentation Strategy.

VIII. CONCLUSION

This project presents a fully integrated battery health monitoring system that combines real-time data aggregation, feature augmentation, and predictive fault detection to ensure proactive maintenance and safety. By capturing fine-grained sensor data and processing it through statistical and machine learning-based methods, the system achieves early detection of anomalies such as rising thermal stress and voltage imbalances before they escalate to critical events like thermal runaway. Simulation and dashboard visualization modules allow operators to monitor battery packs and review fault trends interactively. The hybrid detection framework — blending adaptive thresholding and learning-based classifiers — draws on recent advances in the literature to improve sensitivity while controlling false alarms.

In summary, the proposed system enhances the state of battery safety monitoring by providing an extensible, data-driven platform that supports real-time analytics, proactive detection, and actionable insights for maintenance and operational decision-making.

REFERENCES

1. *Thermal fault detection in battery systems using principal component analysis with adaptive thresholding*, Journal of Energy Storage, 2026.
2. *Fault detection for Li-ion batteries of electric vehicles with feature-augmented attentional autoencoder*, Scientific Reports, 2025.
3. *A Systematic Review of Thermal Runaway in Li-ion Batteries*, *Fus. Eng. Bull.*, 2025.
4. *Recent Progress of Deep Learning Methods for Health Monitoring of Lithium-Ion Batteries*, MDPI, 2024.
5. *Research progress in fault detection of battery systems: A review*, Journal of Energy Storage, 2024.
6. *Hybrid machine learning framework for predictive maintenance and anomaly detection in lithium-ion batteries*, Scientific Reports, 2025.
7. *Fault detection for Li-ion batteries with segmented regression method,* Scientific Reports, 2024.
8. *Battery fault diagnosis and thermal runaway warning based on FEF and DTW*, Journal of Energy Storage, 2023.
9. *Review on Thermal Runaway Monitoring Technologies for Lithium-Ion Batteries*, MDPI, 2023.
10. *Review on Lithium-ion Battery PHM from Key PHM Steps*, Chinese Journal of Mechanical Engineering, 2024.
11. *BattBee: Equivalent Circuit Modeling and Early Detection of Thermal Runaway*, arXiv, 2025.

12. *Gaussian process-based online health monitoring and fault analysis*, arXiv, 2024.
13. **Zhang, Y., et al.** Battery safety: Machine learning-based prognostics, *Prog. Energy Combust. Sci.*, 2024.
14. **Xu, Y., et al.** Recent advances in model-based lithium-ion battery fault diagnosis, *Renew Sustain. Energy Rev.*, 2025.
15. **Peng, J., et al.** Enhancing lithium-ion battery monitoring: sensing approaches, *ETransportation*, 2024.