

Innovative AI-Enabled Digital Twins for Smart Battery Management: Enhancing Battery Life and Performance in Smart Vehicles

*Original*

Innovative AI-Enabled Digital Twins for Smart Battery Management: Enhancing Battery Life and Performance in Smart Vehicles / Alamin, Khaled Sidahmed Sidahmed. - (2025 Sep 09), pp. 1-178. [10.13121/polito/porto/3002995]

*Availability:*

This version is available at: 11583/3002995 since: 2025-09-12T09:45:43Z

*Publisher:*

Politecnico di Torino

*Published*

DOI:10.13121/polito/porto/3002995

*Terms of use:*

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

*Publisher copyright*

(Article begins on next page)



**Politecnico  
di Torino**

**ScuDo**  
Scuola di Dottorato ~ Doctoral School  
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation  
Doctoral Program in Control & Computer Engineering (37.th cycle)

# Innovative AI-Enabled Digital Twins for Smart Battery Management: Enhancing Battery Life and Performance in Smart Vehicles

**Khaled Alamin**

\* \* \* \* \*

**Supervisor**

Prof. Sara Vinco

**Doctoral examination committee**

Prof. Nicola Bombier, Referee, University of degli Studi di Verona

Prof. Donkyu Baek, Referee, Chungbuk National University

Prof. Sangyoung Park, Referee, Technische Universität Berlin

Prof. Davide Zoni, Referee, Politecnico di Milano

Prof. Gianvito Urgese, Politecnico di Torino

Politecnico di Torino  
May 08, 2025

This thesis is licensed under a Creative Commons License, Attribution - Noncommercial-NoDerivative Works 4.0 International: see [www.creativecommons.org](http://www.creativecommons.org). The text may be reproduced for non-commercial purposes, provided that credit is given to the original author.

I hereby declare that the contents and organization of this dissertation constitute my own original work and do not compromise in any way the rights of third parties, including those relating to the security of personal data.

.....  
Khaled Sidahmed Sidahmed Alamin  
Turin, April 29, 2025

# Summary

Electric vehicles (EVs) are poised to play a pivotal role in the global transition toward sustainable mobility. At the heart of every EV lies its battery, whose performance, safety, and longevity are governed by the effectiveness of the battery management system (BMS). In traditional BMS approaches, challenges arise from the complex, non-linear behaviors of batteries, exacerbated by varying usage patterns, aging effects, and partial charge/discharge cycles. To address these challenges, this thesis proposes an AI-enabled digital twin framework that enhances battery life and performance in smart vehicles by holistically estimating the State of Charge (SOC), State of Energy (SOE), and State of Health (SOH) in real-time.

Building on prior research in both model-based and data-driven paradigms, the work systematically investigates and integrates advanced machine learning techniques to overcome the intrinsic variability and unpredictability of battery behavior. A key innovation lies in model-driven dataset generation and feature engineering, which ensure that data-driven models are both robust and resource-efficient. By leveraging domain knowledge of battery electrochemistry and degradation phenomena, the proposed methods extract salient features that capture subtle changes in internal battery states—even when subjected to partial discharge scenarios common in real-world EV operations.

Central to this thesis is the concept of a digital twin, which continuously synchronizes with the physical battery system through real-time sensing at the edge, capturing realistic driving profiles and operational conditions rather than relying solely on theoretical models. Cloud-based analytics enable on-demand retraining and model updates. This dynamic dual-model approach ensures that SOC, SOE, and SOH estimations remain accurate and reliable throughout the battery lifecycle, seamlessly adapting to capacity fade and other aging factors. The architecture also facilitates a Pareto-based trade-off between computational overhead and estimation accuracy, guiding designers in deploying BMS solutions optimized for diverse hardware platforms.

The thesis unifies contributions presented in leading conferences and journals, progressively advancing from fundamental machine learning-based digital twin strategies to sophisticated feature engineering and segmentation for partial discharge cycles. Experimental validations using publicly available datasets and custom datasets demonstrate significant improvements in estimation precision and resilience compared to conventional BMS strategies. These findings underscore the potential for AI-driven digital twins to

serve as a transformative technology in EV battery management, enabling longer battery lifespan, reduced operational costs, and accelerated EV adoption.

In summary, this work provides a comprehensive framework that addresses critical gaps in battery state estimation—particularly those arising from dynamic real-world usage and aging. The methodologies, algorithms, and architectures proposed have broad implications for the future of smart, data-driven mobility and pave the way for more reliable, efficient, and sustainable EV ecosystems.



# Acknowledgements

I am grateful to many people and organisations whose support—often behind the scenes—made this research possible.

First, I acknowledge the Doctoral School and the Department of Computer and Control Engineering here in Politecnico di Torino for providing an intellectually stimulating environment, laboratory space, and access to the faculty’s facilities.

The constructive feedback from my thesis committee—Professors for sharpened the final work. Their questions pushed me to clarify assumptions, strengthen the validation protocol, and broaden the discussion on real-world deployment.

I gratefully acknowledge the maintainers of open-source software used throughout this study, especially open source tools and libraries’ developers: `LightGBM`, `scikit-learn` and all the battery-modeling and data analysis toolkit. Their freely available tools accelerated experimentation and reproducibility.

Nonetheless, indirect support from the academic community—through openly shared data sets, published benchmarks, and lively conference debates—proved invaluable in shaping the research direction.

Finally, any remaining errors are mine alone.

*First and foremost, all praise and gratitude are due to **Allah** for granting me the strength, patience, and clarity of mind to complete this journey—every step was only possible by His will.*

***To my mother, Fatima,** whose quiet prayers and boundless love carried me through every setback. Your example of resilience is my constant inspiration.*

***To my brothers,** thank you for standing beside me in moments of doubt, for every late-night call, and for reminding me why perseverance matters.*

***To my supervisor, Sara.** Your guidance, honesty, and steady encouragement transformed challenges into opportunities. You not only supervised this thesis; you shaped the way I think and work.*

***To Yukai, Daniele, and Massimo.** You opened the door when I was still finding my footing, offered generous advice, and help me through my research journey. I could not have asked for better mentors.*

***To my lab-mates, friends, and the entire research family.** The laughter over coffee, calcettino, the groups meetings, and every shared deadline made the hard days lighter and the good days unforgettable.*

***Finally, to Sudan—my homeland.** Though you endure hardship and war, you remain the ground beneath my aspirations. I pray for peace and for the chance to give back to the people and places that shaped me.*

# Contents

|                                                                                |      |
|--------------------------------------------------------------------------------|------|
| <b>List of Tables</b>                                                          | XI   |
| <b>List of Figures</b>                                                         | XIII |
| <b>1 Introduction</b>                                                          | 1    |
| 1.1 Problem Statement and Research Gaps . . . . .                              | 2    |
| 1.2 Research Objectives . . . . .                                              | 4    |
| 1.3 Thesis Outline . . . . .                                                   | 5    |
| <b>2 Fundamentals and Advanced Perspectives</b>                                | 7    |
| 2.1 Evolution of Rechargeable Batteries . . . . .                              | 7    |
| 2.1.1 Historical Milestones in Rechargeable Cells . . . . .                    | 7    |
| 2.1.2 Why Lithium-Ion Dominates EVs . . . . .                                  | 8    |
| 2.2 Internal Cell Design and Performance Metrics . . . . .                     | 10   |
| 2.2.1 Cell Composition and Electrochemical Operation . . . . .                 | 10   |
| 2.2.2 Performance Metrics . . . . .                                            | 12   |
| 2.2.3 Battery State Estimation: SoC, SoH, and SoE . . . . .                    | 12   |
| 2.2.4 Comparison of Key Lithium-Ion Variants for EVs . . . . .                 | 14   |
| 2.2.5 Degradation Mechanisms . . . . .                                         | 16   |
| 2.2.6 Thermal Management and Safety . . . . .                                  | 17   |
| 2.3 Battery Behavior and Influencing Factors . . . . .                         | 18   |
| 2.4 Battery Management System (BMS) . . . . .                                  | 19   |
| 2.4.1 Early BMS Hardware Architectures and Functions . . . . .                 | 19   |
| 2.4.2 Moving Toward Modular Topologies . . . . .                               | 20   |
| 2.4.3 Modern BMS Hardware: Distributed and Intelligent Architectures . . . . . | 20   |
| 2.4.4 Core Functionalities of Modern BMS . . . . .                             | 21   |
| 2.4.5 Data Acquisition and Signal Processing in Distributed BMS . . . . .      | 21   |
| 2.4.6 Intelligent Cell Balancing Strategies . . . . .                          | 22   |
| 2.4.7 Advanced Embedded Software, Real-time Controls, and Safety . . . . .     | 22   |
| 2.4.8 Edge-to-Cloud Integration . . . . .                                      | 23   |
| 2.4.9 Cloud-Based Predictive Diagnostics . . . . .                             | 24   |
| 2.5 Digital Twin Concepts for Battery Applications . . . . .                   | 24   |

|          |                                                                               |           |
|----------|-------------------------------------------------------------------------------|-----------|
| 2.5.1    | Core Components of a Digital Twin Architecture . . . . .                      | 27        |
| 2.5.2    | Modeling Strategies for Battery Digital Twins . . . . .                       | 28        |
| <b>3</b> | <b>Literature Review</b>                                                      | <b>31</b> |
| 3.1      | Modeling Approaches for Battery Systems . . . . .                             | 31        |
| 3.1.1    | Empirical Models . . . . .                                                    | 32        |
| 3.1.2    | Equivalent Circuit Models (ECMs) . . . . .                                    | 34        |
| 3.1.3    | Electrochemical Models . . . . .                                              | 36        |
| 3.1.4    | Data-Driven and Hybrid Modeling Approaches . . . . .                          | 38        |
| 3.1.5    | Hybrid and Gray-Box Models (Combining Physics and Machine Learning) . . . . . | 41        |
| 3.1.6    | Transformer-Based Models for Battery State Estimation . . . . .               | 42        |
| 3.1.7    | Physics-Informed Neural Networks (PINNs) . . . . .                            | 44        |
| 3.1.8    | Comparative Summary of Modeling Approaches . . . . .                          | 45        |
| 3.1.9    | Adaptability & Long-Term Stability of the Digital Twin . . . . .              | 46        |
| 3.2      | Research Gaps and Future Directions . . . . .                                 | 48        |
| 3.2.1    | Model adaptation with aging . . . . .                                         | 48        |
| 3.2.2    | Partial-cycle and dynamic-condition SoH estimation . . . . .                  | 48        |
| 3.2.3    | Real-time integration in cloud-edge frameworks: . . . . .                     | 50        |
| 3.2.4    | Unified estimation of multiple “SoX” states: . . . . .                        | 51        |
| 3.2.5    | Embedded feasibility and validation: . . . . .                                | 51        |
| 3.3      | Systematic Literature Survey Methodology . . . . .                            | 52        |
| <b>4</b> | <b>Digital Twin Architecture for EV Battery Management</b>                    | <b>55</b> |
| 4.1      | Edge–Cloud Architecture Overview . . . . .                                    | 58        |
| 4.2      | Dynamic Model Update Mechanism . . . . .                                      | 61        |
| 4.3      | Integration with BMS Hierarchy and Fault Tolerance . . . . .                  | 63        |
| 4.4      | Robust Handling of Partial Charge–Discharge Cycles . . . . .                  | 65        |
| 4.5      | Comparison with Baseline Architectures . . . . .                              | 67        |
| 4.6      | Comparison with Industry Solutions: . . . . .                                 | 70        |
| <b>5</b> | <b>Methodology</b>                                                            | <b>73</b> |
| 5.1      | Overview . . . . .                                                            | 73        |
| 5.2      | Data Collection and Preprocessing . . . . .                                   | 73        |
| 5.2.1    | Datasets Employed . . . . .                                                   | 75        |
| 5.2.2    | Data Preprocessing . . . . .                                                  | 80        |
| 5.2.3    | Dataset Splitting for Robust Validation . . . . .                             | 83        |
| 5.3      | Segmentation for Partial-Cycle Handling . . . . .                             | 84        |
| 5.4      | Feature Engineering . . . . .                                                 | 87        |
| 5.5      | Machine Learning Models (ML) . . . . .                                        | 91        |
| 5.5.1    | Evaluated Models . . . . .                                                    | 91        |
| 5.5.2    | Model Selection Rationale . . . . .                                           | 92        |

|          |                                                                                                  |            |
|----------|--------------------------------------------------------------------------------------------------|------------|
| 5.5.3    | Hyperparameter Tuning . . . . .                                                                  | 93         |
| 5.5.4    | Performance Evaluation Metrics . . . . .                                                         | 95         |
| 5.6      | Model Update Mechanisms and Workflow . . . . .                                                   | 95         |
| 5.7      | Modeling Approach for SoC/SoE and SoH Estimation . . . . .                                       | 100        |
| 5.7.1    | SoC and SoE Modeling . . . . .                                                                   | 100        |
| 5.7.2    | SoH Modeling . . . . .                                                                           | 100        |
| 5.8      | Implementation Details and Workflow . . . . .                                                    | 101        |
| <b>6</b> | <b>Consolidated Experimental Results and Analysis</b>                                            | <b>105</b> |
| 6.1      | Phase I (COINS): Foundational Digital Twin on NASA Dataset . . . . .                             | 105        |
| 6.1.1    | Motivation and Scope . . . . .                                                                   | 105        |
| 6.1.2    | Experimental Setup and Key Findings . . . . .                                                    | 107        |
| 6.2      | Phase II (AEIT): Extending to SoE and Multi-Chemistry . . . . .                                  | 110        |
| 6.2.1    | Motivation and Contributions . . . . .                                                           | 110        |
| 6.2.2    | Machine Learning Approach and Data Splitting . . . . .                                           | 110        |
| 6.2.3    | SoH Estimation Results Using XGBoost . . . . .                                                   | 111        |
| 6.2.4    | SoC and SoE Estimation . . . . .                                                                 | 111        |
| 6.3      | Phase III (ISLPED): Synthetic Data Generation and Resource Trade-Offs                            | 114        |
| 6.3.1    | Motivation and Objectives . . . . .                                                              | 114        |
| 6.3.2    | Methodology Recap . . . . .                                                                      | 114        |
| 6.3.3    | Key Experimental Results . . . . .                                                               | 115        |
| 6.4      | Phase IV (DATE): Feature Engineering for Enhanced SoH Accuracy . .                               | 117        |
| 6.4.1    | Objective and Additional Insights . . . . .                                                      | 117        |
| 6.4.2    | Experimental Setup and Feature Sets . . . . .                                                    | 117        |
| 6.4.3    | SoH Estimation: Experimental Results . . . . .                                                   | 118        |
| 6.5      | Phase V (IEEE TVT): Final Integrated Digital Twin and Comparative<br>Benchmark . . . . .         | 119        |
| 6.5.1    | Datasets and Implementation Details . . . . .                                                    | 122        |
| 6.5.2    | SoH Accuracy . . . . .                                                                           | 122        |
| 6.5.3    | Adaptive SoC/SoE . . . . .                                                                       | 125        |
| 6.5.4    | Digital-Twin Proof of Concept . . . . .                                                          | 126        |
| 6.5.5    | Runtime& Memory Footprint . . . . .                                                              | 128        |
| 6.5.6    | Comparison with Related Work . . . . .                                                           | 130        |
| 6.6      | Benchmarking the Digital-Twin BMS against Leading Industrial and<br>Academic Solutions . . . . . | 132        |
| <b>7</b> | <b>Conclusion and Future Works</b>                                                               | <b>139</b> |
| 7.1      | Implications and Real-World Impact . . . . .                                                     | 141        |
| 7.2      | Remaining Challenges and Future Directions . . . . .                                             | 142        |
| 7.3      | Outlook and Closing Remarks . . . . .                                                            | 143        |
| <b>A</b> | <b>Author's Publications</b>                                                                     | <b>145</b> |

# List of Tables

|      |                                                                                                                                  |     |
|------|----------------------------------------------------------------------------------------------------------------------------------|-----|
| 2.1  | Key Historical Milestones in Rechargeable Battery Technologies . . . . .                                                         | 8   |
| 2.2  | Key Lithium-Ion Battery Variants for EV Applications . . . . .                                                                   | 15  |
| 2.3  | Evolution of BMS Feature Sets Across Three Generations . . . . .                                                                 | 25  |
| 3.1  | Representative Literature and Relation to Thesis Focus Areas . . . . .                                                           | 54  |
| 4.1  | Qualitative Comparison of Battery Digital Twin Architectures (Proposed Hybrid vs. Edge-only vs. Cloud-only) . . . . .            | 69  |
| 5.1  | Comprehensive Overview of Extracted Features for Battery Health Modeling                                                         | 89  |
| 5.2  | Hyperparameters Optimized for LightGBM Model . . . . .                                                                           | 93  |
| 6.1  | Overview of Research Phases and Key Publications . . . . .                                                                       | 106 |
| 6.2  | Phase I (COINS): Different SoH Models Performance. T + I indicates training + inference time in seconds; RMSE, MAE in %. . . . . | 107 |
| 6.3  | Phase I (COINS): SoC Model Performance at <b>SoH = 100%</b> . . . . .                                                            | 108 |
| 6.4  | Phase I (COINS): SoC Model Performance at <b>SoH = 70%</b> . . . . .                                                             | 108 |
| 6.5  | Phase I (COINS): Impact of an Outdated SoC Model (Trained at SoH = 100%) Used at SoH = 75%. . . . .                              | 109 |
| 6.6  | Phase II (AEIT): <b>SoH</b> Models Performance using XGBoost (T + I in seconds). . . . .                                         | 111 |
| 6.7  | Phase II (AEIT): Combined SoC–SoE at <b>SoH = 100%</b> (Trained at 100% SoH). . . . .                                            | 112 |
| 6.8  | Phase II (AEIT): Combined SoC–SoE at <b>SoH = 75%</b> (Trained at 75% SoH). . . . .                                              | 112 |
| 6.9  | Phase II (AEIT): SoC–SoE at <b>SoH=75%</b> but model was <i>trained at 100% SoH</i> (Outdated). . . . .                          | 113 |
| 6.10 | Phase III (ISLPED): Representative Deployment Results on a Low-Power MCU (RISC-V). . . . .                                       | 116 |
| 6.11 | Phase IV (DATE): SoH Estimation Error on LFP and NMC Cells with LightGBM vs. LSTM, under CA and NA Feature Sets. . . . .         | 118 |
| 6.12 | SoH MAE comparison (LightGBM vs. SVR) on validation and test sets.                                                               | 123 |
| 6.13 | SoH Model Estimation Error for LightGBM and SVR . . . . .                                                                        | 124 |
| 6.14 | Comparison of Testing MAE between Fixed Model and Model Updating                                                                 | 126 |
| 6.15 | Training / inference time and memory usage (standard CPU, no GPU). .                                                             | 130 |
| 6.16 | SoH accuracy versus published methods. . . . .                                                                                   | 132 |

|      |                                                          |     |
|------|----------------------------------------------------------|-----|
| 6.17 | Linear [0,10] mapping used for the radar scores. . . . . | 133 |
| 6.18 | Verified raw numbers that feed the radar. . . . .        | 133 |
| 6.19 | Scores used in Fig. 6.9. . . . .                         | 134 |
| 6.20 | Qualitative factors (✓= yes, ✗= no, – = N/R). . . . .    | 135 |

# List of Figures

|     |                                                                                                                                                                                                                                                 |    |
|-----|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.1 | Schematic of a typical lithium-ion cell. Lithium ions shuttle between the cathode and the anode through an electrolyte-filled separator during charge and discharge. . . . .                                                                    | 11 |
| 2.2 | High-Level BMS Overview . . . . .                                                                                                                                                                                                               | 26 |
| 3.1 | Example of simple circuit based modeling. . . . .                                                                                                                                                                                               | 34 |
| 3.2 | Schematic illustration of electrochemical processes in a lithium-ion battery based on the DFN model. The image highlights lithium diffusion, electrolyte transport, and interfacial reactions across the anode, separator, and cathode. . . . . | 37 |
| 3.3 | Schematic illustration of a deep learning pipeline for SoC estimation . .                                                                                                                                                                       | 40 |
| 3.4 | SoC discharge profiles at different stages of battery aging. The nonlinear variation in discharge characteristics emphasizes the necessity of adaptive modeling in BMS. . . . .                                                                 | 49 |
| 3.5 | Voltage response under discharge at varying SoH levels. Lower SoH corresponds to quicker voltage drops, indicating internal battery aging mechanisms such as increased impedance and lithium depletion. . . . .                                 | 49 |
| 3.6 | SoC profile recorded during real driving conditions of a BMW i3 battery pack. The absence of complete charge or discharge cycles underscores the necessity for partial-cycle and dynamic-condition SoH estimation methodologies [159]. . . . .  | 51 |
| 3.7 | PRISMA flow diagram illustrating the literature identification, screening, eligibility, and inclusion process for the systematic review. . . . .                                                                                                | 53 |
| 4.1 | Edge-cloud partitioning of the proposed DT: the cloud is responsible for model (re)training, while the edge handles data collection and execution of the SoX algorithms. . . . .                                                                | 56 |

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |    |
|-----|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 4.2 | The digital twin is trained initially in a cloud environment (step 1 using historical data from batteries of the same chemistry or from the same production line if available. After this initial training, the SoC/SoE model is transferred to the PMU (step 2), where it runs in real time. During vehicle operation, the PMU collects telemetry from the module and cell levels, which is periodically uploaded to the cloud in batches or via incremental data streaming (step 3). This newly accumulated data is merged with existing datasets in the cloud, triggering a model retraining step (step 4) according to a time-based or event-based policy. If the updated model surpasses the performance of the currently deployed version on test subsets, it is compressed and transmitted back to the PMU, ensuring that on-board predictions remain accurate despite changes in battery health and usage conditions. . . . . | 60 |
| 5.1 | Workflow of the proposed edge–cloud digital-twin pipeline. Solid arrows represent the real-time edge path; dashed arrows depict asynchronous edge–cloud synchronisation. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 74 |
| 5.2 | Schematic illustration of the Sandia National Laboratories (SNL) battery cycling protocol. Partial cycles ( $PC^1$ , $PC^2$ , ...) are interspersed with periodic reference cycles ( $RC^1$ , $RC^2$ , $RC^3$ ), which are full 0–100% charge/discharge cycles at a standard rate (0.5C) to check and log the cell’s capacity (SoH). This figure highlights how partial cycles operate at different C-rates, temperatures, and depth-of-discharge (DoD) windows, followed by the reference cycles for capacity checks. . . . .                                                                                                                                                                                                                                                                                                                                                                                                        | 77 |
| 5.3 | Simulated Battery Dataset (Model-Driven Generation). Using a high-fidelity equivalent-circuit battery model, synthetic charge/discharge cycles are generated under diverse simulated conditions (e.g., variable ambient temperature, battery voltage, and aging parameters). The resulting data includes time-series signals for Battery Current, Voltage, Temperature, as well as labels for SoH, SoC, and SoE. This simulation approach augments experimental datasets by providing unlimited and diverse examples, including rare or extreme scenarios. (See :contentReference[oaicite:2]index=2 for further details.) . . . . .                                                                                                                                                                                                                                                                                                   | 79 |
| 5.4 | SOH degradation profile extracted from the RW9 dataset. Arrows points where the monotonically decreasing nature of SOH is not reflected due to measurement errors. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 82 |
| 5.5 | Illustration of the battery data segmentation strategy. A complete discharge cycle (left) is systematically partitioned into multiple non-overlapping segments with progressively smaller state-of-charge intervals (right). This approach enables multi-resolution modeling and closely replicates the partial usage cycles common in real-world EV operation. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 85 |

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |     |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 5.6 | From initial dataset (top) to feature extraction (bottom). The initial dataset is reported as table and as plot of one discharge cycle (data from [180]). The discharge raw cycles are segmented based on the $\Delta\text{SoC}$ (in the Figure we show a few samples with $\Delta\text{SoC}=5\%$ ) and aggregated to train the SoC/E model (bottom left). Statistical and derived features are computed over each segment, forming the trained data for SoH (bottom right). . . . | 88  |
| 5.7 | Model selection and training pipeline for SoH and SoC estimation. The flowchart illustrates the primary model (LightGBM), supplementary models tested, the training-validation workflow, and key hyperparameters optimized during development. . . . .                                                                                                                                                                                                                             | 94  |
| 5.8 | Updating Mechanism Strategy . . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                              | 98  |
| 6.1 | SOC discharge profiles comparison among different SOC Random Forest models trained in different SOH cases. . . . .                                                                                                                                                                                                                                                                                                                                                                 | 109 |
| 6.2 | Phase II (AEIT): SoC and SoE discharge profiles on NASA (left) and Sandia (right) under different SoH levels. The red line (updated model) tracks the measured profile more closely than the older models. . . . .                                                                                                                                                                                                                                                                 | 113 |
| 6.3 | Pareto-fronts obtained from the hyper-parameters exploration of data-driven models . . . . .                                                                                                                                                                                                                                                                                                                                                                                       | 116 |
| 6.4 | Phase IV (DATE): <i>SoH vs. cycle count</i> for three sample cells. NA-based features (orange) exhibit lower error accumulation compared to CA (green). . . . .                                                                                                                                                                                                                                                                                                                    | 119 |
| 6.5 | Feature importance estimated by LightGBM. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                  | 120 |
| 6.6 | SoH prediction vs. ground truth across NASA, Sandia, and Stanford datasets at varying segment sizes ( $\Delta\text{SoC} = 1,5,25,50,100\%$ ). . . . .                                                                                                                                                                                                                                                                                                                              | 124 |
| 6.7 | NASA 62 % $\rightarrow$ 25 % slice: weekly (green) and 5 % SoH-drop (orange) updates halve the SoC error compared with a frozen model (red). . . . .                                                                                                                                                                                                                                                                                                                               | 127 |
| 6.8 | Four-month deployment on NASA RW9. <b>Top:</b> evolving SoH estimates under variable segmentation. <b>Bottom:</b> edge-side SoC/SoE model re-freshed every 2weeks or at each 5% SoH drop. . . . .                                                                                                                                                                                                                                                                                  | 128 |
| 6.9 | Four-axis radar comparing industrial (solid) and academic (dashed) baselines with the proposed DT. . . . .                                                                                                                                                                                                                                                                                                                                                                         | 136 |

# Chapter 1

## Introduction

Over the past few decades, the global imperative to address climate change and reduce greenhouse gas (GHG) emissions has transformed the landscape of transportation, motivating a rapid shift toward electrified mobility. Within the European Union (EU), for instance, the transport sector contributes approximately 5% to the region's GDP and employs over 10 million people, while simultaneously accounting for nearly one quarter of the EU's total GHG emissions [1, 2]. In response to this dichotomy, governments worldwide have introduced far-reaching strategies, aiming not only to lower environmental impacts but also to stimulate economic growth through innovation in sustainable mobility. The EU's Sustainable and Smart Mobility Strategy crystallizes this vision, seeking at least 30 million zero-emission vehicles on European roads by 2030 and near-complete electrification of new vehicles by 2050 [3, 4].

## Motivation and Background

**Electrified Transport and EV Batteries:** Battery Electric Vehicles (BEVs) have emerged as a pivotal technology in this energy transition, supported by converging policy incentives, technological breakthroughs, and consumer demand for cleaner, more efficient transportation [5]. Lifecycle assessments repeatedly highlight that BEVs emit between 45% and 65% fewer GHGs compared to traditional internal combustion engine vehicles (ICEVs), particularly when charged with electricity sourced from renewables. Over a single vehicle's lifespan, these reductions can translate into an estimated several tonnes of CO<sub>2</sub>-equivalent savings per unit [6], underscoring their significant environmental benefits. Moreover, BEVs can help stabilize energy systems by functioning as flexible storage solutions for intermittent renewable energies like wind and solar [7].

**Policy and Market Dynamics:** Driven by stringent emission targets and financial incentives, EV sales continue to climb. Massive investments in charging infrastructures—both public and private—are enhancing EV accessibility, while ongoing innovations in lithium-ion battery (LIB) chemistries (e.g., silicon anodes, solid-state prototypes)

are boosting vehicle range and lowering costs [8]. As a result, the global EV fleet is projected to surpass 85 million units by 2025, of which approximately 73% will be BEVs.

**Challenges in Battery Lifetime Management and EV Operation:** Despite positive developments in battery technology, ensuring long-term battery health and addressing aging remain critical to the sustainable and large-scale adoption of BEVs. Battery aging directly influences vehicle performance, reliability, and user acceptance, making accurate monitoring and predictive analytics essential for maximizing operational efficiency. Robust, real-time battery health monitoring and predictive modeling can significantly prolong battery lifetime, optimize operational strategies, and support more effective second-life uses and recycling, ultimately enabling a more circular and economically sustainable approach to EV battery management.

**The Central Role of Battery Management Systems (BMS):** At the core of every EV lies its battery, a complex electrochemical system dictating vehicle range, performance, and safety. A BMS monitors and controls the charging–discharging process at the cell, module, and pack levels to ensure safety, optimize performance, and prolong battery life [9]. Key BMS functions include estimating the battery’s internal states – the state of charge (SoC), state of energy (SoE), and state of health (SoH) – and taking action (e.g., limiting power or balancing cells) to prevent over-charge, over-discharge, over-temperature, and other unsafe conditions. Accurate knowledge of SoC (the remaining charge relative to capacity), SoE (the remaining usable energy) [10], and SoH (the remaining useful life or capacity relative to a new battery) is essential for reliable operation and user confidence in battery systems. However, due to the electrochemical complexity and non-linear degradation behavior of lithium-ion batteries, precisely estimating these states in real time is very challenging. Even with decades of development, conventional BMS algorithms struggle to capture the battery’s evolving dynamics under diverse operating conditions, especially as the battery ages and deviates from its initial model parameters [11]. This limitation often leads to conservative usage (to avoid unknown risks) or unexpected performance drops and failures if the true battery state diverges from the BMS estimates. The need for more accurate, predictive, and adaptive battery management has therefore become critical. [12]. Traditional BMS methods often rely on *static* like lookup table or *semi-empirical* models—such as fixed equivalent circuit models (ECMs)—that are calibrated under controlled laboratory tests. In real-world EV usage, however, such methods struggle to handle dynamic factors like partial discharge cycles, variable C-rates, and time-dependent degradation pathways (e.g., solid electrolyte interphase (SEI) growth, lithium plating).

## 1.1 Problem Statement and Research Gaps

Recent years have seen a paradigm shift in approach to these challenges, driven by the emergence of “digital twin” technology. A digital twin (DT) is essentially a virtual replica of a physical system that runs in parallel to the real system, continuously synchronized

with it via data streams, and capable of simulating the system’s behavior under various conditions [13]. In practical terms, the digital twin of a battery would combine high-fidelity battery models with live sensor measurements (e.g., voltage, current, temperature) from the battery pack, updating the model in real time so that it **reflects the battery’s current state and performance** [14]. Emerging DT technologies, integrating real-time data with virtual models, have begun addressing these BMS limitations by enabling detailed battery state simulations and predictive analytics [13, 14]. However, implementing accurate, real-time adaptive digital twin solutions for EV battery management faces significant challenges:

- **Addressing Partial Discharge and Dynamic Usage:** Unlike lab settings, real EV usage is dominated by intermittent top-offs and mid-range SoC fluctuations, leading to *partial discharge* conditions that challenge the assumptions of standard BMS models [1, 12]. Full-cycle calibration (0–100%) is seldom feasible day-to-day, yet the absence of these reference cycles can degrade the accuracy of SoC/SoE calculations and accelerate SoH underestimations.
- **Dynamic Model Updating vs. Resource Constraints:** While data-driven modeling (including neural networks, random forests, etc.) can adapt to evolving battery states, implementing frequent retraining on embedded automotive hardware remains difficult [15]. Transferring data to the cloud can enable large-scale analytics, but this demands a scheme to merge updated model parameters back to the edge without compromising real-time performance.
- **Synergizing SoC, SoE, and SoH (SoX):** Despite the interlinked nature of these battery states, many existing frameworks treat them individually [16]. On the other hand, a drop in SoH (capacity fade) directly shrinks the usable capacity for a given SoC range, yet static SoC algorithms do not routinely account for it in real time. This decoupling creates estimation drift and can lead to suboptimal or unsafe operational decisions.
- **Lifecycle-Aware Health Tracking:** A further, often overlooked element is the tracking of battery health data over the full operational life. If a digital twin or BMS can log accurate SoH histories, re-purposing EV batteries for stationary storage or other second-life uses becomes safer and more cost-effective [17]. However, few solutions incorporate such life-cycle awareness or actively facilitate recycling optimization.

Hence, the central challenge is to design an *adaptive, data-centric, and partial-discharge-aware BMS*, capable of leveraging both on-board and cloud computing resources to maintain accurate SoC, SoE, and SoH estimates throughout the battery’s lifecycle.

## 1.2 Research Objectives

This thesis aims to address these gaps through the following objectives:

- **Dual-Model Digital Twin Architecture:** A fundamental innovation of this thesis is the development of a dual-model digital twin framework designed explicitly for the BMS of EVs consisting of:
  - Edge Layer (Real-Time On-board BMS): Employs lightweight machine learning models (Light Gradient Boosting) optimized for deployment on embedded system. These models continuously estimate critical battery states: SoC and SoE, utilizing real-time data including voltage, current, and temperature. The edge component prioritizes minimal computational latency, robustness, and high-frequency monitoring to maintain vehicle safety and performance.
  - Cloud Layer (Advanced Analytics and Model Retraining): Leverages powerful computational resources to implement sophisticated, resource-intensive analytical models. Periodically, data collected from EV is transmitted to the cloud, facilitating comprehensive battery analytics, retraining of the digital twin models, and extraction of insights into longer-term degradation behaviors. After processing, optimized model parameters and insights are securely synchronized back to the vehicle, updating edge models and thereby continuously adapting the local BMS to real-world battery aging conditions.
- **Partial-Discharge-Aware Segmentation and Feature Engineering:**
  - Dynamic Usage Segmentation: Introduces innovative algorithms for automatically segmenting battery cycling data into meaningful operational regimes, particularly capturing mid-range SoC fluctuations typical of daily driving patterns.
  - Domain-Specific Feature Extraction: Develops targeted feature engineering methods to explicitly characterize subtle but critical aging signatures indicative of battery degradation.
- **Comprehensive and Integrated SoX Estimation:** Integrate SoC, SoE, and SoH estimation into a cohesive modeling pipeline, validating performance under realistic partial discharge scenarios with high accuracy. The integrated SoX estimation is rigorously validated against extensive real-world datasets (e.g., NASA, Sandia, Stanford driving cycle data).
- **Resource and Pareto Analysis:** Evaluate computational resources, memory, and latency of different ML approaches, establishing a Pareto front guiding trade-offs between model complexity and embedded deployment feasibility.
- **Deployment Feasibility Demonstration:** - Automotive-Grade Validation: Demonstrates real-world applicability by deploying the edge component of the digital

twin on automotive-grade microcontrollers, specifically targeting widely adopted platforms such as the STMicroelectronics SPC58 microcontroller series, commonly found in EV battery management units.

- **Practical Considerations:** Explores strategies for efficient model parameter compression, lightweight inference, and edge-side adaptation to dynamically evolving battery states.

## 1.3 Thesis Outline

The remainder of this thesis is organized as follows:

- **Chapter 2: Battery Technology Fundamentals** Presents the electrochemical basis of Lithium-Ion batteries, key degradation processes, and the definitions of SoC, SoE, and SoH. This chapter provides a concise but thorough grounding in battery physics necessary for the advanced modeling approaches discussed later.
- **Chapter 3: Literature Review** Surveys existing BMS methodologies, from classical equivalent circuit models (ECMs) and electrochemical methods to hybrid and data-driven frameworks. It also contextualizes digital twin concepts within EV battery research, identifying the principal gaps in partial discharge adaptation, dynamic retraining, and resource optimization.
- **Chapter 4: Digital Twin Architectures for BMS** Explores the architectural choices of edge–cloud systems, focusing on their applicability to EV batteries. Discusses how proposed model can integrate with cloud-based analytics for large-scale data assimilation.
- **Chapter 5: Proposed Framework** Details the dual-model solution, including novel feature engineering, segmentation algorithms, and a dynamic retraining mechanism that addresses partial cycles and aging. The chapter also explains model selection rationales based on resource vs. accuracy trade-offs.
- **Chapter 6: Experimental Validation** Demonstrates the framework’s performance on public datasets (NASA, Sandia, Stanford UDDS) and custom data, emphasizing improvements under partial discharge usage. Presents comprehensive results in terms of: SoC/SoE/SoH accuracy, memory overhead, inference speed, and real-world feasibility.
- **Chapter 7: Conclusion** Summarizes the thesis contributions, to illustrate the importance of integrating partial-discharge–aware digital twin solutions in EV battery management, and explores future directions such as secure cloud integration, second-life battery utilization, and federated learning for large-scale parameter sharing.

**Scope of Work:** Although the proposed digital twin methodology could apply to generic battery systems, this thesis specifically focuses on **\*\*electric vehicle (EV) battery management\*\***. We tailor our models and architecture to EV conditions – for instance, the rapid load changes of driving cycles and the automotive safety requirements – to advance EV battery management in practice. *EV-specific considerations* in our design include the use of standard in-vehicle communication (CAN bus) and adherence to automotive functional safety guidelines (ISO 26262) for any on-board algorithms. These aspects ensure that our approach is not only theoretically sound but also suitable for real EV deployment.

In essence, this thesis strives to *advance EV battery management* by unifying data-driven modeling and domain expertise through model-driven and feature engineering into a flexible digital twin paradigm. By emphasizing partial discharge scenarios, adaptive cloud-edge analytics, and integrated SoC/SoE/SoH estimation, the proposed system substantially improves upon conventional BMS—offering robust, real-time accuracy that endures throughout the battery’s lifecycle. Beyond immediate gains in range prediction and operational safety, the capacity for continuous adaptation and historical tracking also paves the way for cost-effective second-life applications and more sustainable material flows in the electro-mobility domain.

# Chapter 2

## Fundamentals and Advanced Perspectives

This chapter offers a broad and in-depth overview of the fundamentals behind modern electric vehicle (EV) batteries, their management systems, and the emerging role of digital twins for real-time monitoring and predictive analytics. The exposition begins with the historical evolution of rechargeable battery technologies and explains why Lithium-Ion batteries (LIBs) dominate today's EV landscape (Section 2.1). We then detail the internal design of lithium-ion cells and highlight key performance metrics such as energy density, power density, and cycle life (Section 2.2). Section 2.4 traces the evolution of Battery Management Systems (BMS)—from rudimentary voltage protection circuits to sophisticated data-driven controllers—and underscores their core functionalities. Finally, Section 2.5 introduces digital twin concepts in the context of battery applications, discussing how emerging cloud-edge paradigms and data-driven innovations herald a new era of predictive, lifecycle-oriented battery management.

### 2.1 Evolution of Rechargeable Batteries

#### 2.1.1 Historical Milestones in Rechargeable Cells

Rechargeable battery technology has undergone remarkable progress since the late 19th century, driven by both industrial and consumer needs for reliable, portable, and high-density energy storage. Table 2.1 summarizes the evolution of key chemistries, highlighting how each generation has sought to overcome the shortcomings of its predecessors while balancing factors such as cost, safety, toxicity, and energy density.

Early innovations centered on lead-acid batteries, which became the first mass-produced secondary (rechargeable) cells. Despite their relatively low energy density, lead-acid batteries proved robust and were widely adopted in automotive starter-lighting-ignition (SLI) applications. Over time, nickel-based chemistries, notably nickel–cadmium (Ni–Cd) and nickel–metal hydride (Ni–MH), introduced improvements in cycle life and

energy density [18]. Ni–Cd cells were especially popular in power tools, where their ability to deliver high discharge currents was essential, but concerns over cadmium toxicity and the “memory effect” spurred the search for alternative solutions [19]. Ni–MH later alleviated some of these environmental and memory issues; however, it still offered moderate energy density and was eventually eclipsed by lithium-ion (Li-ion) chemistry in most consumer devices. The shift to Li-ion was further accelerated by growing environmental regulations (e.g., EU’s 2006 Restrictions of Hazardous Substances Directive), which discouraged cadmium and other toxic materials in electronics and vehicles.

Table 2.1: Key Historical Milestones in Rechargeable Battery Technologies

| Battery Chemistry                        | Era / Commercialization | Key Characteristics / Usage                                                                |
|------------------------------------------|-------------------------|--------------------------------------------------------------------------------------------|
| <b>Lead-Acid</b>                         | 1860s onward            | Early EVs (late 19th century); low energy density; robust but heavy.                       |
| <b>Ni–Cd</b>                             | 1900s–1980s             | Moderate energy density; toxicity issues with cadmium; used in portable tools.             |
| <b>Ni–MH</b>                             | 1980s–1990s             | Improved energy density over Ni–Cd; overshadowed by Li-ion in consumer electronics.        |
| <b>Lithium-Ion (LCO)</b>                 | Early 1990s             | High energy density; introduced by Sony for consumer electronics; spurred Li-ion research. |
| <b>NMC / NCA / LFP</b>                   | 2000s–present           | Optimized Li-ion variants for EVs; trade-offs in cost, safety, and performance.            |
| <b>Beyond Li-ion (Li–S, Solid-State)</b> | 2010s–present           | Research-stage chemistries; aim for higher energy density or safer electrolytes.           |

### 2.1.2 Why Lithium-Ion Dominates EVs

Lithium-ion (Li-ion) battery technology has become the de facto standard for nearly all modern EVs, largely because it excels across a combination of critical performance metrics—energy density, power capability, cycle life, and cost. These advantages have been further amplified by continuous research and development (R&D), large-scale manufacturing, and intensive supply chain optimization, making Li-ion cells not only more affordable but also better suited for automotive applications than previous battery

chemistries such as nickel–metal hydride (Ni–MH) and lead-acid. Key factors contributing to Li-ion’s dominance include:

- **High Energy Density:** Li-ion batteries offer significantly greater energy density (Wh/kg) compared to older chemistries like Ni–MH or lead-acid [20]. This high energy-to-weight ratio is indispensable for EVs, where vehicle range is directly tied to onboard energy storage. As battery energy density improves, automakers can design EVs with longer driving ranges or reduce overall battery pack size (and weight), resulting in better efficiency and performance.
- **Cycle Longevity and Cost Decline:** Through intensive R&D, Li-ion cells have reached cycle lifetimes of approximately 2,000 to 3,000 cycles (or more, depending on operating conditions), which translates into years of reliable operation in EV applications [21]. At the same time, the cost of Li-ion batteries has dropped precipitously over the past decade, reaching (and in some markets even falling below) the critical \$100/kWh threshold [22]. This cost decline has been driven by economies of scale in manufacturing, advances in material processing, and improved cell designs, all of which have helped EVs edge closer to cost parity with internal combustion engine vehicles.
- **Chemistry Versatility:** A wide range of Li-ion cathode formulations—such as lithium nickel manganese cobalt oxide (NMC), lithium nickel cobalt aluminum oxide (NCA), and lithium iron phosphate (LFP)—allows manufacturers to tailor cells to specific requirements. For instance, NMC and NCA chemistries typically deliver higher energy densities, making them desirable for long-range EVs, whereas LFP provides enhanced thermal stability, safety, and a longer cycle life, albeit at a somewhat lower energy density [23]. LFP’s cobalt-free design has gained renewed traction in China, where it now powers >50% of new EVs, reflecting both cost advantages and alignment with domestic mineral availability [24]

Despite Li-ion’s broad commercial success, a suite of *next-generation* battery chemistries including lithium–sulfur (Li–S), sodium-ion, and solid-state variants—are under active investigation for potential breakthroughs in energy density, thermal safety, and sustainability. However, these emerging technologies still face significant hurdles in achieving widespread market readiness. For example, Li–S cells exhibit promising theoretical energy densities but suffer from rapid capacity fade and limited cycle life, while solid-state designs promise enhanced safety through nonflammable electrolytes but often entail complex manufacturing processes and high upfront costs [25]. Solid-state batteries, for instance, face interfacial resistance challenges: rigid ceramic electrolytes often fracture during electrode expansion, while sulfide-based alternatives require inert-atmosphere production to prevent toxic  $H_2S$  release [26]. As research continues, improvements in materials engineering, electrode interfaces, and scalable production methods will determine whether these alternatives can rival the Li-ion ecosystem that currently powers nearly all of today’s EVs.

## 2.2 Internal Cell Design and Performance Metrics

### 2.2.1 Cell Composition and Electrochemical Operation

A typical lithium-ion cell comprises four primary components—an anode (negative electrode), a cathode (positive electrode), a separator, and an electrolyte—and operates on the principle of reversible lithium-ion intercalation and de-intercalation. As illustrated in Figure 2.1, during charging lithium ions migrate from the cathode through the electrolyte-filled separator and intercalate into the layered structure of the anode material. During discharge, these ions flow back to the cathode, producing an electrical current in the external circuit [27].

**Anode (Negative Electrode).** Graphite remains the most commonly employed anode material, owing to its stable crystal structure and relatively high storage capacity for lithium ions. Cutting-edge research in Si/graphite composite anodes has shown promise for further increasing energy density, although challenges persist regarding silicon’s volume expansion and mechanical stability during cycling.

**Cathode (Positive Electrode).** The cathode typically consists of transition metal oxides (e.g.,  $\text{LiCoO}_2$ ,  $\text{LiNiMnCoO}_2$  [NMC], or  $\text{LiNiCoAlO}_2$  [NCA]) or polyanionic compounds (notably  $\text{LiFePO}_4$ , or LFP). Each cathode chemistry offers distinct trade-offs in terms of voltage, capacity, thermal stability, and cost. For instance,  $\text{LiCoO}_2$  historically enabled commercial lithium-ion cells but remains one of the more expensive options due to cobalt’s scarcity, whereas LFP provides better thermal stability and a lower overall cost but offers a reduced nominal voltage, limiting energy density.

**Separator and Electrolyte.** Sandwiched between the anode and the cathode, the separator is a thin, porous polymeric film (often polypropylene or polyethylene) designed to prevent direct electronic contact (which would cause a short circuit) while allowing ionic transport. The separator is impregnated with a liquid electrolyte—commonly lithium hexafluorophosphate ( $\text{LiPF}_6$ ) dissolved in a mixture of organic carbonate solvents. Although liquid electrolytes dominate current EV batteries, gel or semi-solid formulations are gradually emerging in commercial settings due to potential benefits in safety and manufacturing simplicity. Recent thermal runaway incidents in early Li-ion EVs (e.g., 2013 Boeing 787 fires) spurred innovations like flame-retardant electrolyte additives (e.g., triphenyl phosphate) and ‘self-healing’ binders that reduce electrode cracking during cycling [28].

Figure 2.1 emphasizes the pathways for lithium ions, which move through the separator during the charge and discharge processes. The structure and composition of each component strongly influences key battery characteristics such as voltage, internal resistance, thermal stability, and overall cycle life.




Figure 2.1: Schematic of a typical lithium-ion cell. Lithium ions shuttle between the cathode and the anode through an electrolyte-filled separator during charge and discharge.

### 2.2.2 Performance Metrics

Three critical performance indicators govern an EV battery's suitability for automotive applications: energy density, power density, and cycle life.

**Energy Density (Wh/kg).** A cell's energy density directly impacts the driving range of an EV. Higher energy density translates to either greater range for a given pack weight or a lighter battery pack for a given range. Over the past decade, improvements in electrode materials and manufacturing techniques have allowed lithium-ion batteries to exceed 200 Wh/kg in some commercial formulations [29].

**Power Density (W/kg).** While energy density determines total range, power density governs how fast the battery can deliver or accept power. High power density is vital for rapid acceleration and fast charging. Internal electrode design, including optimization of particle size, porosity, and conductive additives, influences the power capabilities of a cell.

**Cycle Life.** The number of full charge–discharge cycles a cell can undergo before its capacity or power capability degrades below a specified threshold (often 80% of its original capacity) is known as the cycle life. In automotive applications, a target cycle life of 1,000–2,000 cycles is a typical benchmark, though some lithium-ion chemistries (e.g., LFP) can exceed 3,000 cycles under moderate temperature and state-of-charge conditions. Achieving a long cycle life requires controlling cell degradation mechanisms such as electrolyte decomposition, electrode particle cracking, and the growth of solid–electrolyte interphase layers.

### 2.2.3 Battery State Estimation: SoC, SoH, and SoE

Accurate and reliable battery state estimation is a foundational requirement for managing lithium-ion cells in EVs. Since the internal states of a battery—such as how much charge or energy remains, or how degraded it has become—are not directly measurable in real time, the BMS must infer them through mathematical models, sensor fusion, and historical operating data. Three key indicators are used to characterize the battery's operational state: SoC, SoH, and SoE. These parameters not only determine user-facing metrics (e.g., remaining range), but also underpin safety decisions (e.g., whether to trigger thermal protection or charge rate limits).

#### State of Charge (SoC)

SoC estimates the remaining charge in the battery relative to its maximum nominal capacity [30], typically expressed as a percentage:

$$\text{SoC}(t) = \frac{Q_{\text{remaining}}(t)}{Q_{\text{nominal}}} \times 100\%$$

Here,  $Q_{\text{remaining}}$  is the amount of charge available at time  $t$ , while  $Q_{\text{nominal}}$  represents the rated (design) capacity of the battery.

SoC is analogous to a fuel gauge and is critical for EV range estimation, energy management, and safe charge/discharge control. However, accurately estimating SoC is non-trivial due to current measurement drift, temperature-dependent dynamics, and hysteresis. Common methods include:

- **Coulomb Counting:** Integrates current over time to estimate net charge in/out. It is sensitive to current sensor drift and accumulates error over time.
- **Open Circuit Voltage (OCV):** Relies on the battery's voltage–SoC relationship, requiring long rest periods to reach equilibrium.
- **Model-Based Estimation:** Combines equivalent circuit models (ECMs) or physics-based models with real-time sensor data.
- **Data-Driven Methods:** Uses algorithms such as extended Kalman filters (EKF), particle filters, or machine learning to adaptively estimate SoC under dynamic conditions [31].

### State of Health (SoH)

SoH quantifies the overall degradation of the battery compared to its original condition. It is defined as:

$$\text{SoH} = \frac{Q_{\text{actual}}}{Q_{\text{rated}}} \times 100\%$$

where  $Q_{\text{actual}}$  is the current full charge capacity, and  $Q_{\text{rated}}$  is the battery's original capacity when new.

SoH reflects how much of the battery's original performance remains and is essential for lifespan prediction, warranty validation, and second-life applications. SoH estimation can also incorporate:

- **Impedance Growth:** An increase in internal resistance signals aging and degraded ion transport.
- **Cycle Aging Models:** Tracks capacity loss over time, often using empirical or semi-empirical degradation laws.
- **Machine Learning Approaches:** Trains regression or classification models on historical datasets to predict degradation trends [32].

### State of Energy (SoE)

SoE estimates the remaining usable energy in a cell or pack relative to its maximum available energy:

$$\text{SoE} = \frac{E_{\text{available}}}{E_{\text{max}}} \times 100\%$$

Unlike SoC, which is charge-based, SoE accounts for the variation in operating voltage during discharge. This distinction is crucial in high-performance EVs, where energy (in Wh) rather than just charge (in Ah) determines range and acceleration.

SoE estimation depends on integrating the product of current and voltage over time, adjusted for real-time power demands and efficiency losses. Errors in voltage prediction, especially under dynamic loads, can significantly impact SoE accuracy [33].

### Challenges in Real-Time Estimation

Real-world SoC, SoH, and SoE estimation faces several practical challenges:

- **Sensor Drift and Noise:** Current and voltage sensors degrade or lose calibration, especially over long lifetimes.
- **Temperature Dependence:** Battery behavior (e.g., capacity, voltage, resistance) varies significantly with ambient and cell temperature.
- **Hysteresis Effects:** Voltage–SoC relationships show path dependency, complicating model-based SoC estimation.
- **Cell-to-Cell Variability:** Manufacturing differences and uneven aging introduce state variation within a battery pack.
- **Complex Load Profiles:** Real-world driving introduces variable current and temperature patterns that challenge static models.

To mitigate these issues, advanced BMS units use hybrid estimation strategies combining electrochemical models, circuit-based approximations, historical memory, and adaptive filters. The recent adoption of **digital twin architectures**—in which each battery has a continuously updated virtual replica—offers enhanced precision in predicting future battery behavior and enabling predictive maintenance [34].

### 2.2.4 Comparison of Key Lithium-Ion Variants for EVs

Table 2.2 summarizes representative lithium-ion cathode chemistries, highlighting differences in energy density, cycle life, and overall strengths or limitations. For instance, NMC offers an excellent balance of cost, energy density, and power capability but depends on cobalt, which introduces supply chain risks. NCA, used in some high-performance EVs, delivers even higher energy density, yet it demands more stringent thermal management. LFP, in contrast, trades off lower nominal voltage (and therefore lower energy density) for cost-effectiveness, robust safety, and longer operational lifetimes—attributes that suit many commercial vehicle and grid-storage applications. These distinctions underscore how no single “universal” lithium-ion chemistry exists; rather, automakers and cell

Table 2.2: Key Lithium-Ion Battery Variants for EV Applications

| <b>Cathode</b>    | <b>Energy Den-<br/>sity</b> | <b>Cycle Life</b>  | <b>Strengths / Chal-<br/>lenges</b>                                             |
|-------------------|-----------------------------|--------------------|---------------------------------------------------------------------------------|
| <b>NMC</b>        | High (> 180 Wh/kg)          | 1,000–2,000 cycles | Good balance of cost, energy, power. Cobalt use poses supply risks.             |
| <b>NCA</b>        | Very High (> 200 Wh/kg)     | 1,000–1,500 cycles | High energy density, used by Tesla. More sensitive to thermal runaway.          |
| <b>LFP</b>        | Moderate (~ 150 Wh/kg)      | 2,000–3,000 cycles | Safer, cheaper, stable voltage plateau. Lower voltage reduces overall EV range. |
| <b>LMO / LNMO</b> | Moderate–High               | 1,000–2,000 cycles | High power density, simpler structure. Voltage fade issues (LNMO).              |

manufacturers strategically select formulations that align with particular performance, cost, and safety priorities.

Ultimately, the nuanced selection and optimization of these battery components determines not only an EV's range and power characteristics but also its safety profile, lifecycle costs, and environmental footprint. As lithium-ion technologies continue to evolve—and as next-generation chemistries emerge—the fundamental principles of electrode engineering, electrolyte formulation, and thermal management remain key to achieving superior battery performance across diverse electric mobility applications.

## **2.2.5 Degradation Mechanisms**

As lithium-ion cells undergo repeated charge and discharge cycles, various electrochemical and mechanical processes gradually diminish their usable capacity, power output, and overall lifespan. Understanding these degradation pathways is critical for improving battery reliability, optimizing operational protocols (e.g., charging rates and temperature conditions), and developing new materials to mitigate long-term performance losses. The following outlines three major mechanisms responsible for capacity fade and cell aging:

**1. Solid-Electrolyte Interphase (SEI) Growth.** A thin passivation layer known as the solid-electrolyte interphase (SEI) initially forms on the anode surface during the very first charge cycles. This interphase is beneficial in that it protects the anode from continual decomposition of the electrolyte. However, under normal operation—and especially under stress conditions such as high temperature, high voltage, or suboptimal charging protocols—SEI thickness can continue to grow uncontrolled. As the SEI expands, lithium ions are irreversibly consumed in side reactions, effectively reducing the cell's active lithium inventory [35, 36]. This phenomenon increases internal resistance and manifests as a gradual decrease in cell capacity. While advanced electrolyte additives (e.g., film-forming agents) and optimized temperature management can slow SEI growth, it remains one of the primary contributors to capacity fade in lithium-ion batteries. capacity fade in lithium-ion batteries. This explains why many EV manufacturers recommend limiting daily charging to 80% SOC, as lower SOC ranges reduce electrolyte decomposition rates and SEI thickening [37].

**2. Lithium Plating.** During fast charging (particularly at high currents) or operation in cold-temperature conditions, the rate at which lithium ions intercalate into the anode structure can become limited by diffusion kinetics. When the anode can no longer accommodate incoming lithium ions quickly enough, metallic lithium may deposit—or “plate”—on its surface [38]. Even a small amount of lithium plating reduces the total amount of active lithium available for cycling, leading to capacity loss. Moreover, dendritic (needle-like) lithium deposits can grow through the separator, posing a serious risk of internal short circuits and potentially leading to catastrophic thermal events. Sophisticated BMS help mitigate plating by controlling charge current and temperature ranges. Additionally, materials engineering approaches—such as using more conductive anode coatings or tailored electrode porosity—aim to alleviate lithium plating under

stressful operating regimes.

**3. Transition Metal Dissolution.** On the cathode side, metals like cobalt, nickel, and manganese can dissolve from their host lattices, especially when cells are operated at high temperatures, subjected to deep discharges, or stored at high states of charge for prolonged periods [1]. For instance,  $\text{LiCoO}_2$  cathodes are particularly susceptible to cobalt dissolution under aggressive cycling conditions. Once dissolved, these transition metal ions migrate through the electrolyte and can deposit on the anode surface, where they further disrupt the anode's SEI stability. This cross-contamination exacerbates capacity fade and increases cell impedance. Various strategies exist to mitigate metal dissolution, including doping cathode materials with protective ions (e.g., magnesium or aluminum) and applying protective coatings that reduce side reactions at the cathode–electrolyte interface. In EVs, prolonged exposure to high temperatures (e.g., parking in direct sunlight) exacerbates metal dissolution, necessitating active cooling systems even during vehicle idle periods

**4. Interplay and Mitigation Strategies.** Although presented here as distinct processes, these degradation pathways often occur in tandem—SEI thickening, for example, can exacerbate lithium plating by increasing cell impedance and driving localized current spikes. Likewise, dissolved metals can accelerate SEI growth or trigger localized plating sites. Cell-level solutions include improved electrode formulations (e.g., partial silicon anodes, advanced cathode coatings), refined electrolyte chemistries (including solid or semi-solid alternatives), and sophisticated thermal and charging controls governed by the BMS. By optimizing these factors, manufacturers can significantly extend cell lifetime, reduce capacity fade, and enhance the overall safety of lithium-ion batteries in EV and stationary storage applications.

### 2.2.6 Thermal Management and Safety

Effective thermal management is central to balancing lithium-ion battery performance, safety, and longevity. EV battery packs employ active or passive cooling systems to maintain cells within a narrow operational temperature range (typically 15–35°C).

**Active Cooling:** Liquid-cooled systems circulate a glycol-water mixture through aluminum cold plates adjacent to cells, achieving precise temperature control even under fast-charging or high-load conditions. Tesla's "Octovalve" system, for instance, dynamically routes coolant to hotspots, reducing peak temperatures by 10°C compared to passive designs [39].

**Passive Cooling:** Air-cooled or phase-change material (PCM)-based systems are lighter and cheaper but struggle to dissipate heat during sustained high-power operation. PCMs (e.g., paraffin wax) absorb latent heat during melting, stabilizing temperatures temporarily, but their efficacy diminishes in extreme environments [40].

**Thermal Runaway Mitigation:** Cell-to-pack (CTP) and blade battery designs (e.g., BYD's LFP-based packs) improve thermal stability by eliminating modular structures that can propagate thermal failures. Simultaneously, embedded sensors and machine

learning-driven BMS algorithms predict thermal anomalies by correlating voltage dips with localized temperature spikes [41].

Despite these advances, thermal management accounts for 10–15% of total pack weight and cost, driving research into simpler, energy-efficient solutions such as refrigerant-based direct cooling or thermally conductive potting materials [39].

## 2.3 Battery Behavior and Influencing Factors

Battery performance is highly influenced by temperature, load current, and usage history:

- **Temperature:** Low temperatures increase internal resistance and reduce available capacity (temporary loss), while high temperatures can increase capacity slightly in the short term but accelerate degradation in the long term. Most EV batteries are operated in a conditioned environment (thermal management system maintaining say 20°C–40°C range) to balance performance and longevity. Extreme cold or heat can significantly skew SoC estimation if not accounted for, as voltage response changes with temperature.
- **Discharge/Charge Rate:** High C-rate discharges (rapid acceleration or fast charging) cause larger IR drops and transient behavior in cell voltage, complicating instantaneous SoC estimation. Repeated high C-rate usage also stresses the battery more, contributing to faster SoH decline.
- **Partial vs Full Cycles:** A full cycle (0% to 100% SoC) provides the most information to certain SoH estimation techniques (like those relying on full charge capacity measurement). However, EV users rarely go from 0 to 100 regularly; partial cycles (e.g., 80% to 30% and back) are more common. Batteries can exhibit *memory effects* or differential aging based on depth of discharge; shallow cycling often causes less wear per energy throughput but can make capacity fade harder to measure without occasional full cycles. Our digital twin approach will need to extract SoH information even from partial cycles.
- **Relaxation Effects:** After a charge or discharge, if the battery rests, its voltage tends toward equilibrium (OCV). Many SoC estimation methods (and lab reference tests) rely on measuring this rested voltage. In real driving, long rest periods are infrequent (maybe overnight charging). This means on-board estimation must handle states during operation without full relaxation.

Understanding these fundamentals lays the groundwork for the modeling choices in later chapters. With this background, we now turn to how the research community and industry have been managing and modeling these battery states, and where the gaps remain that a digital twin could fill.

## 2.4 Battery Management System (BMS)

A BMS is an embedded control unit responsible for the monitoring, protection, and optimization of a rechargeable battery pack, especially in EVs. It acts as the “brain” of the battery pack, ensuring safe and reliable operation by continuously tracking cell voltages, pack current, and temperatures [42]. Core BMS functions include state estimation – gauging the battery’s SoC, SoH, and SoE – as well as cell balancing, thermal management, and fault diagnosis. By performing accurate state estimation and control, the BMS prevents hazardous conditions like over-charge, over-discharge, and overheating, thereby extending battery life and ensuring vehicle safety, it safeguards the battery pack through real-time monitoring and control (e.g. shutting off charging at high SoC or reducing load on overheating cells) and optimizes performance by managing energy use and interfacing with other vehicle systems [43]. With these critical roles, BMS technology has rapidly evolved since the 1990s to meet the increasing demands of modern EVs.

BMS have evolved dramatically—from early rudimentary protection circuits to complex, data-centric control units integral to every EV. This section traces that historical progression, detailing how hardware and software developments transformed BMS from simple safety guardians to intelligent systems that optimize battery performance and longevity. It begins by examining early centralized designs and proceeds to highlight modern distributed architectures, advanced state estimation methodologies, and emerging trends in connectivity and predictive control. Throughout, the evolution of BMS is framed by critical improvements in embedded processing, sensor integration, algorithmic sophistication, and cloud-based analysis.

### 2.4.1 Early BMS Hardware Architectures and Functions

**Centralized Beginnings:** Early-generation EV BMS implementations were typically centralized: a single controller board monitored an entire battery pack (often comprising dozens or hundreds of cells). All sensor wiring—cell voltage taps, temperature sensors, current shunts—converged on one unit, creating substantial wire harness complexity and a single point of failure [44]. The hardware on these boards was fairly minimal, usually an analog front end (AFE) to measure voltage and current, plus a low-power microcontroller responsible for strictly protective functions (over/undervoltage, over/temperature, and fault detection). Cell balancing, if provided at all, was rudimentary, often consisting of passive resistor-based bleeding only near full charge. The software side was likewise limited: early BMS code essentially ran a simple loop, scanning cell voltages and current, then tripping contactors if values went out of specification [45]. SoC estimation was rudimentary—often just coulomb counting or static voltage lookups when the battery was at rest, with negligible modeling of chemical or thermal effects [46]. Nevertheless, communication capabilities were also very limited; many early BMS units were stand-alone, lacking standardized network interfaces. This was acceptable for smaller or less demanding battery packs of the 1990s and early 2000s but rapidly became insufficient as

battery packs grew bigger and EV markets expanded.

### **2.4.2 Moving Toward Modular Topologies**

As EVs began requiring larger packs (many hundreds or thousands of cells), purely centralized BMS architectures ran into scalability and reliability problems [47, 48]. High cell counts introduced bulky wiring, prone to interference and mechanical wear, and any single fault on the main controller could disable the entire pack [49]. In response, a modular or distributed approach emerged: each module (or subset of cells) would have a local BMS board (commonly called a Cell Monitoring Unit or BMS slave) measuring voltages and temperatures. A master BMS controller or Master Monitoring Unit MMU, located elsewhere in the vehicle, would coordinate these local boards via a daisy-chain bus or CAN bus [50]. This distributed topology reduced wiring, improved fault tolerance, and allowed flexible scaling of battery packs. Manufacturers also adopted a master–slave scheme, in which subordinate units handled data acquisition and balancing, but deferred high-level estimation and control decisions to a centralized logic module [51]. Overall, by the late 2010s, BMS had evolved into distributed, smart control systems. They were fully networked within the vehicle (communicating with the Vehicle Control Unit and thermal management systems) and capable of handling large Li-ion packs safely under diverse conditions. An illustrative example is Tesla’s Model S (2012) BMS, which uses monitoring units on each of its battery modules (each module with dozens of cells) networked to a central battery controller. This system monitors thousands of individual cell channels in total and manages both passive balancing and thermal control, while interfacing via CAN with the car’s powertrain and charger to enforce limits based on battery state

### **2.4.3 Modern BMS Hardware: Distributed and Intelligent Architectures**

Modern EVs universally employ distributed BMS architectures for reasons of scalability, precision, and reliability. Each module-level “slave” board measures local cell voltages (often with millivolt-level accuracy) using integrated AFE chips that can handle multiple channels, built-in temperature sensor inputs, and sometimes direct drivers for balancing field-effect transistors (FETs) [52]. These boards communicate measurement data to the “master” BMS over a high-integrity, isolated bus such as CAN, SPI, or a proprietary daisy-chain link. By minimizing sense-wire runs (replaced by a single communication harness), wiring complexity and mass are significantly reduced, and potential noise pickup is lowered. Typical modern systems can gather synchronized data from hundreds of cells at refresh rates adequate for real-time estimation and control.

A recent trend toward wireless BMS eliminates module-to-module daisy-chain cables in favor of RF-based communication. The main purported benefits include weight savings, easier pack assembly, and more flexible pack layouts. General Motors pioneered this by

deploying a wireless distributed BMS in its Ultium EV platform (announced in 2020), eliminating up to 90% of the low-voltage wiring within the pack [53]. Although reliability and EMC (electromagnetic compatibility) remain technical challenges, the industry is steadily progressing toward robust commercial adoption of wireless BMS solutions [54].

#### 2.4.4 Core Functionalities of Modern BMS

A high-level schematic of a typical BMS is shown in Figure 2.2. The BMS begins by acquiring sensor data (voltage, current, and temperature), which serve as inputs for the core state estimation modules—predicting SoC, SoH, and SoE. These estimated states directly inform the BMS’s safety and protection mechanisms, preventing unsafe conditions (e.g., overcharging or deep discharging). Simultaneously, the BMS coordinates cell balancing to ensure consistent charge levels across all cells, optimizes thermal management (e.g., cooling or heating), and exchanges data with the vehicle’s main control unit or external systems via a dedicated communication interface (CAN bus or wireless). The central computational control unit orchestrates all these tasks in real-time, making rapid decisions to maintain safety, longevity, and performance of the battery pack.

#### 2.4.5 Data Acquisition and Signal Processing in Distributed BMS

Accurate state estimation hinges on robust sensor data. Consequently, modern BMS employ high-precision AFEs, distributed measurement ICs, and sophisticated filtering:

- **Cell Voltage Sensing.** Each cell or parallel group is monitored with millivolt-level resolution via delta-sigma ADCs. Calibration routines and synchronous sampling ensure minimal offset and consistent sampling intervals. Daisy-chain communication protocols with built-in CRC checks help mitigate transmission errors [55].
- **Pack Current Sensing.** High-speed, low-noise measurements, often using dedicated shunt resistors or Hall-effect sensors, support coulomb counting and power calculations. Auto-zero offset routines and temperature compensation are common to maintain accuracy across large current ranges [56].
- **Temperature Monitoring.** Distributed thermistors or RTDs track thermal gradients. BMS might smooth and cross-check these readings to detect anomalous hotspots, especially critical for safety and performance.
- **Isolation Monitoring.** Automotive BMS often include circuits to detect insulation leaks between high-voltage battery terminals and vehicle chassis. Although separate from SoC/SoH, isolation diagnostics are essential for overall electrical safety.
- **Synchronization and Sensor Fusion.** In a distributed setup, modules send data to a master controller that time-aligns measurements for consistent state estimation.

Data integrity is checked, outlier values are flagged, and communication networks (often CAN bus to the vehicle's main ECU) robustly handle hundreds of sensor channels with minimal latency.

### 2.4.6 Intelligent Cell Balancing Strategies

Balancing counters the inevitable mismatch among cells, ensuring each cell's SoC aligns as closely as possible with the rest of the pack which require consistent SOC across all cells to prevents overcharging and over-discharging, enhancing battery safety and longevity. Optimal balancing performance relies on accurate SOC estimation for each cell. Additionally, some balancing systems employ variable resistors to counteract disparities in internal resistance.

- **Passive Balancing** dissipates excess energy from higher-SoC cells as heat through bleed resistors. Though simpler and cheaper, it wastes energy, typically making sense only at the top of charge [57].
- **Active Balancing** transfers charge from higher-SoC cells to lower-SoC cells via capacitor or inductor circuits. While more complex and expensive, it conserves energy and can operate during both charge and discharge, improving pack utilization [58, 59].

### 2.4.7 Advanced Embedded Software, Real-time Controls, and Safety

The proliferation of tasks—state estimation, fault diagnostics, cell balancing control, thermal management, and vehicular communication—necessitated more complex software architectures, often built around 32-bit MCUs and real-time operating systems (RTOS). Today's BMS firmware, developed under automotive functional safety standards (e.g., ISO 26262, MISRA C rules), is highly modular and structured [60, 61, 62]. Developers frequently use model-based design and simulation (e.g., MATLAB/Simulink) for validation against battery models prior to real-hardware implementation.

**Real-Time Estimation and Control:** As BMS advanced, real-time computational demands grew, driven primarily by sophisticated estimation algorithms for SoC, SoH, and SoE. Extended Kalman Filters (EKF), Unscented Kalman Filters (UKF),  $H^\infty$  filters, and other observer-based frameworks are standard practice in many commercial EVs [63, 64]. These filters incorporate measured voltage and current data, combined with an internal battery model, to produce accurate estimates of internal states under dynamic loads. BMS software must run these algorithms reliably within strict cycle times (on the order of milliseconds to tens of milliseconds), while concurrently managing balancing FETs, thermal loops, and communication tasks.

**Maintainability and Over-the-Air Updates:** Another hallmark of modern BMS software is the capability for over-the-air (OTA) updates. By leveraging the vehicle's

connectivity, manufacturers can deploy firmware patches or enhancements, refining state estimators, updating safety parameters, or adapting the fast-charge profiles for newly discovered cell behaviors [9, 65]. This continuous update paradigm transforms the BMS from a static embedded controller into a continually improving cyber-physical system. In many cases, advanced BMS software design includes a dual-bank architecture to ensure reliable fallback if an OTA update encounters errors, thus preserving safe operation.

### 2.4.8 Edge-to-Cloud Integration

Modern BMS increasingly link to the broader IoT ecosystem, enabling advanced analytics and continuous improvements:

**Vehicle-Cloud Connectivity.** A growing number of EV platforms upload BMS data (SoC, SoH, temperatures, charge logs) to the OEM’s cloud environment, creating large datasets for fleet-wide analysis [9]. This facilitates powerful predictive models that leverage real-world usage patterns. Conversely, insights from the cloud can flow back to individual vehicles—enabling dynamic optimization of charging schedules, usage profiles, or balancing parameters based on aggregated knowledge from thousands of EVs.

**Cloud-Based “Digital Twins.”** Manufacturers are beginning to host high-fidelity virtual representations, or digital twins, of each battery pack in the cloud. These run computationally intensive models (e.g., full electrochemical simulations or machine learning ensembles), fed by real-time data from the actual pack. Discrepancies between predicted and measured performance can trigger recalibrations, highlight pending failures, or inform direct BMS adjustments [11, 66, 67]. Such end-edge-cloud collaboration ensures that the BMS in the car can remain lean, focusing on real-time safety tasks, while more sophisticated analytics run in the cloud.

**OTA Firmware Updates.** Connectivity also permits remote BMS updates. Manufacturers can patch software for bug fixes, introduce new state estimation algorithms, or modify allowable charge/discharge parameters. In some recalls, OEMs have used OTA to lower maximum state-of-charge thresholds to mitigate fire risks. Conversely, they can raise or recalibrate capacity usage if field data suggests robust cell health. Secure update protocols (e.g., encrypted downloads, fallback partitions) are essential to preserve functional safety and protect against tampering.

**Modern Algorithms** contemporary EV BMS do more than simply track cell voltages; they incorporate SoC- or SoH-based balancing triggers to optimize for longevity and total energy utilization [68]. For instance, a “weaker” cell might consistently limit the pack capacity. Intelligent balancing can partially compensate by shifting charge from stronger cells during driving or systematically adjusting the overall pack usage window to reduce stress on the weakest cell.

### 2.4.9 Cloud-Based Predictive Diagnostics

Connectivity also paves the way for predictive diagnostics—anticipating failures or significant degradation before they occur:

- **Cycle-Life and RUL Prediction.** By analyzing usage data (e.g., charge rates, depth of discharge, temperature extremes), the BMS or cloud platform can forecast how many cycles remain until a given SoH threshold is reached [69].
- **Cell Degradation Trends.** Individual cell data can be compared across fleets, spotting outlier cells showing accelerated internal resistance growth. This early detection enables proactive maintenance or partial pack replacement.
- **Thermal Runaway Precursors.** Some research focuses on sensor-based detection of subtle self-heating or gas release patterns indicative of incipient internal shorts. Although still maturing, integration of such early warning systems into mainstream BMS is likely.

Because these cloud-based analytics can be computationally intensive, they complement, rather than replace, the real-time protective functions of the onboard BMS. In effect, the BMS becomes an edge component in a larger data ecosystem, benefiting from aggregated learning but still self-sufficient if connectivity is lost.

**The chronological evolution of battery technologies and management systems discussed is summarized in Table 2.3.**

## 2.5 Digital Twin Concepts for Battery Applications

Digital Twin (DT) technology refers to a **virtual replica** of a physical object or system that is continuously synchronized with the real-world counterpart. In essence, a DT mirrors the state of a physical system in a computer model, enabling real-time monitoring, simulation, and analysis. The concept's **origins trace back to aerospace**: during NASA's Apollo program in the 1960s, engineers maintained an identical “twin” of the spacecraft on Earth to troubleshoot problems in space [70]. This early form of a living simulation foreshadowed the digital twin idea. The term Digital Twin was later formalized in the early 2000s by Michael Grieves in the context of product lifecycle management, describing a model with three core components – a physical entity, a virtual model, and the data connections between them [71]. NASA expanded the concept for next-generation aircraft health management around 2010, emphasizing ultra-high-fidelity models integrated with sensor data for prognostics [72]. **Industry adoption** quickly followed: leading manufacturers and aerospace companies began using DTs to optimize jet engines, production lines, and heavy machinery. By the late 2010s, Digital Twin technology had become a pillar of Industry 4.0, valued for enabling predictive maintenance, design optimization, and operational efficiency. In summary, a DT is now

Table 2.3: Evolution of BMS Feature Sets Across Three Generations

| Feature            |         | 2 <sup>nd</sup> Gen BMS                      | 3 <sup>rd</sup> Gen BMS                                |
|--------------------|---------|----------------------------------------------|--------------------------------------------------------|
| SoC Estimation     |         | ECM + EKF                                    | Hybrid ML/ECM, partial discharge aware                 |
| SoH Tracking       |         | Resistance-based or open-circuit checks      | Cloud analytics, data-driven synergy                   |
| Thermal Management | Manage- | Embedded sensors + basic fan/coolant control | Adaptive thermal strategies, external data integration |
| Connectivity       |         | CAN bus to vehicle ECU                       | Edge–cloud loop, OTA updates, IoT modules              |
| Cell Balancing     |         | Passive/Active hardware                      | Intelligent scheduling based on real-time SoH          |




Figure 2.2: A high-level schematic of a typical BMS. The BMS monitors real-time sensor signals (voltage, current, temperature) and computes internal states (SoC, SoH, SoE). Based on these estimates, it coordinates safety protection (e.g., overcharge/discharge), cell balancing (passive or active), and thermal management (cooling/heating). It also manages communication with the vehicle control unit via CAN or wireless links.

broadly defined as a **real-time digital representation of a physical object** or process that spans its lifecycle and is kept in sync via data flows [73]. This general paradigm has been applied in domains from manufacturing to smart cities, and most recently, to electrochemical energy systems such as batteries.

### 2.5.1 Core Components of a Digital Twin Architecture

At its core, a digital twin architecture unites the physical system, its virtual model, and the data/communication layer that links them, alongside an analytics or services layer that exploits the twin for insights. A canonical framework introduced by researchers and industry groups breaks this into layered components [74, 75]:

- **Physical Layer (Physical Twin):** the real-world asset itself – in this context, the battery pack or cell, equipped with sensors (voltage, current, temperature, etc.) and actuators, all managed by controllers. This is the source of ground truth data and the object the twin aims to represent.
- **Data Communication & Management Layer:** a bidirectional data pipeline connecting the physical battery to its digital counterpart. This layer encompasses IoT sensors, data acquisition systems, and networks to transmit telemetry in real-time [76]. It handles data storage, synchronization, and ensures the virtual model is fed with up-to-date sensor measurements (and, conversely, that control commands or predictions from the twin can be sent back to the physical system). Reliable real-time synchronization is critical here – techniques range from high-frequency polling of BMS data to event-driven updates; the goal is to achieve a “twinning rate” fast enough for the application’s needs (often on the order of seconds or faster for EV batteries).
- **Virtual Model Layer:** the mathematical or computational model that emulates the battery’s behavior in software. This can consist of multi-physics simulation models, data-driven models, or a combination (as detailed in the next section). The virtual model ingests the streamed data and updates its internal state to reflect the physical battery’s current conditions (state-of-charge, temperature distribution, State-of-Health, etc.). In advanced implementations, this layer may be subdivided into sub-models (electrical, thermal, aging, etc.) running in parallel [77]. The fidelity of the virtual model is a key design choice – higher fidelity (e.g. electrochemical models) offers accuracy at the cost of computational load, whereas simplified models run faster but capture fewer details.
- **Analytics and Services Layer:** the outer layer where insights are generated and actions are taken based on the twin. Here, analytics algorithms evaluate the virtual model’s outputs to perform tasks such as diagnostics, prognostics, and optimization [13]. For a battery DT, this can include estimating performance metrics (e.g.

available power, energy), predicting future states like remaining useful life, detecting anomalies or abuse conditions, and enabling what-if analyses for control decisions. For example, a cloud-based battery twin can predict cell failures or capacity loss ahead of time and recommend adjustments [NEWSROOM.PORSCHE.COM](https://www.porsche.com/newsroom/). This layer often interfaces with user applications or operator dashboards, essentially turning the raw model and data into useful services (like predictive maintenance alerts or optimal charging schedules).

These components work in unison to realize the digital twin's functionality. The architecture can be viewed as a five-layer stack as well: (1) physical layer (battery and sensors), (2) communication layer (IoT connectivity), (3) data management/sync layer, (4) modeling & model-updating layer, and (5) application/decision layer. In practice, the communication and data management aspects are often handled by cloud or edge computing infrastructure, while the model and analytics may execute on a combination of embedded systems (e.g. on-vehicle edge compute) and cloud servers. Figure 1 schematically illustrates how these pieces come together for an EV battery system, with the physical pack feeding data to a cloud-hosted model, which in turn provides analytics back to support decisions in the BMS.

This fundamental overview lays the groundwork for the more in-depth *Literature Review* (Chapter 3), wherein we delve into specific modeling techniques (empirical, ECM, electrochemical, data-driven), partial-discharge-aware estimation strategies, and the latest research on digital twin applications.

## 2.5.2 Modeling Strategies for Battery Digital Twins

A distinguishing feature of a battery digital twin is its modeling approach – how the virtual battery is realized in software. Battery behavior is complex, governed by electrochemical processes, thermal dynamics, and usage-dependent degradation. Accordingly, a spectrum of modeling strategies has been employed in DT implementations, generally falling into three categories: physics-based models, data-driven models, and hybrid models. Each has advantages, and many battery DT frameworks combine them to leverage the best of both worlds.

### Physics-Based Modeling

Physics-based models use first-principles or mechanistic descriptions of battery phenomena. These include electrochemical models derived from chemistry and transport equations – for instance, the Doyle-Fuller-Newman model (also known as the pseudo-2D model) which captures lithium-ion concentration gradients in electrodes and electrolyte, or simplified versions like the Single Particle Model [78]. Such models can predict internal states (like concentration or solid-electrolyte interface growth) that are not directly measurable, providing rich insight into battery health. However, full-order electrochemical models are computationally intensive and require many parameters (diffusion constants,

reaction rates, etc.), making real-time use challenging without simplifications or model-order reduction [79]. Another popular physics-based approach is the Equivalent Circuit Model (ECM), which represents the battery as an electrical circuit of voltage sources, resistors, and capacitive elements. ECMs (e.g. simple RC networks) can mimic the battery’s terminal voltage and transient response with relatively low computational cost and are widely used in BMS for state estimation [80]. For example, Porsche’s Battery Digital Twin framework uses a hierarchical physical modeling approach: a basic resistor-capacitor (RC) model forms a “performance module” to capture the battery’s electric characteristics, while a more detailed electrochemical model runs in parallel to simulate cell-level processes, accompanied by a thermal model to predict temperature effects [81]. By combining these conventional models of differing complexity, the twin can represent battery behavior across multiple domains (electrical and thermal) with high fidelity. In a DT context, physics models are often initialized from design data or lab experiments, and then calibrated to the specific battery using field data. Calibration may involve online parameter estimation – for instance, periodically adjusting the internal resistance or diffusion time constants in an ECM based on recent charge-discharge data so that the virtual battery’s output matches the observed voltage [82].

### **Data-Driven Modeling**

In contrast to first-principles modeling, data-driven approaches treat the battery as a system to be learned from data. Here, one develops empirical or machine learning models that map from input variables (current, temperature, usage history) to output variables or states of interest (state-of-charge, capacity fade, etc.), without explicitly encoding electrochemical theory. Neural networks have been extensively explored for this purpose. For example, deep learning models can learn the complex relationships in how capacity degrades with each charge cycle, or how voltage responds under various loads, by training on large datasets of battery cycling data. Techniques ranging from simple regression and Kalman filtering to advanced architectures (e.g. LSTM recurrent networks for sequence modeling of cycling data) have been reported for battery state estimation [83].

### **Hybrid and Fusion Modeling**

Increasingly, hybrid modeling strategies are being adopted for battery digital twins, combining physics-based and data-driven elements to capture the best of both. In a hybrid DT, the physics-based model provides a baseline grounded in electrochemical reality, while data-driven components correct or enhance the model using live data [84]. This fusion can happen in various ways: for example, using machine learning to estimate or adjust parameters of a physics model in real-time (sometimes called “parameter estimation” or “model adaptation” layer), or using an AI model to predict the residual error of a simpler model. The hybrid approach addresses the limitations of each method alone – physics models handle extrapolation and provide interpretability, while data-driven models

absorb real-world variability and complex effects that are hard to explicitly model. The AI algorithms are trained on real driving data to recognize usage patterns and anomalies (e.g. detecting subtle signs of cell imbalance or accelerated degradation) that the physics equations alone might not catch [85]. Meanwhile, the long-term aging phenomena that lack sufficient field history (since EVs are relatively new) are still handled by the physics models and lab-tested degradation mechanisms. By “bringing the two worlds together,” the twin continuously improves: real-world data gradually re-parameterize the simulation models for better realism. Figure 1 illustrates such a hybrid architecture, where a central digital twin is fed by both conventional models (electrochemical, thermal, electrical) and AI models trained on testbench and field data. The result is a more robust virtual battery that can self-calibrate over time.

The continued rapid growth of EV adoption and the pressing need for better battery management ensure that digital twin research will remain vibrant. In conclusion, the digital twin concept – born in the aerospace industry decades ago – has now firmly taken root in battery technology. By combining high-fidelity models, real-world data, and powerful analytics, battery digital twins provide a foundation for safer, longer-lasting, and more sustainable battery systems in electric vehicles and beyond.

# Chapter 3

## Literature Review

Lithium-ion battery models underpin state estimation algorithms critical for modern BMS. Over decades, modeling approaches have evolved from simple empirical techniques to complex physics-based and data-driven models. This review surveys the full spectrum, including: (1) traditional empirical models (open-circuit voltage characterization, Coulomb counting), (2) equivalent circuit models (from Rint to multi-RC Thevenin circuits), (3) electrochemical physics-based models (Doyle–Fuller–Newman pseudo-2D and Single-Particle models), (4) purely data-driven models (machine learning regressors and deep neural networks such as LSTM/CNN), (5) hybrid “gray-box” models combining physics with learning, (6) emerging Transformer-based time-series models, and (7) physics-informed neural networks (PINNs). For each category, we outline the core methodology and mathematics, highlight notable recent works, discuss state-of-charge (SoC), state-of-health (SoH), and state-of-energy (SoE) estimation techniques, and compare performance (accuracy, complexity, real-time suitability). We also examine how these models integrate into BMS and digital twin frameworks. A summary comparison table is provided, and we conclude by identifying research gaps in model aging adaptation, partial-cycle estimation, real-time cloud-edge deployment, unified multi-state estimation, and embedded implementation challenges.

### 3.1 Modeling Approaches for Battery Systems

Traditional battery modeling research can broadly be classified into three major categories: empirical models, equivalent circuit models (ECMs), and electrochemical models. Each category has evolved to address specific challenges associated with the dynamic behavior and degradation of lithium-ion batteries.

### 3.1.1 Empirical Models

Empirical models were among the earliest to be developed for battery management. Empirical models are typically built on experimental data and are designed to capture the nonlinear behavior of the battery's terminal voltage as a function of its SoC and current. Commonly, such models use offline experiments (e.g., constant current discharge/charge tests or lookup table) to derive relationships such as the open-circuit voltage (OCV) versus SoC curve. One widely used method is coulomb counting, which integrates current over time to track the charge removed or added to the battery.

Early empirical techniques, such as the well-known Coulomb counting method, provided a straightforward means of tracking battery charge by integrating current over time [2]. [86]. However, these models are limited by their inability to capture complex internal dynamics, particularly under partial discharge conditions or during aging. More recent empirical studies have attempted to refine these models by incorporating recalibration techniques and adaptive algorithms [5], yet their inherent simplicity often restricts their accuracy and transferability across different operating conditions. Moreover The study suggests that while empirical methods (such as coulomb counting) are highly effective for short-term state tracking, periodic recalibration (for example, via OCV measurements) is necessary to correct drift errors over longer durations [30]. This method can lose some of its accuracy under variable load condition because it ignores the nonlinear discharge effect during the coulomb counting process. The internal resistance method needs to measure the frequency response of the battery to determine its battery state. Because this method normally requires extra function generators and separate testing period, this method is expensive and difficult to implement as part of the battery pack itself [87].

Other empirical cues include internal resistance and voltage hysteresis. For example, applying a current pulse and measuring the immediate voltage drop allows inferring internal resistance, which correlates with SoC in some ranges [88]. At low SoC, the voltage falls off steeply, so resistance-based SoC estimation is relatively sensitive in that regime [89]. However, in mid-range SoC where the discharge curve is flat, this method loses precision. These empirical techniques (OCV, Coulomb count, resistance pulses) offer low complexity and are easily implemented on automotive microcontrollers in real time. Accuracy is moderate: under favorable conditions OCV or pulse methods can estimate SoC within a few percent error, but errors grow over long usage without recalibration [90]. They also do not directly estimate SoH, except via slow indirect trends (e.g. drift in full-charge OCV or required recalibration frequency). Nonetheless, due to their simplicity, empirical methods often serve as components within more advanced estimation frameworks (e.g. providing initial guesses or redundancy checks in a multi-model BMS).

- High Complexity of Degradation Mechanisms: [91] stresses that the deterioration of lithium-ion cells involves a multitude of coupled chemical, mechanical, and thermal processes. Because these processes (such as SEI film growth, lithium plating, particle cracking, and transition metal dissolution) interact non-linearly and evolve

with operating conditions, it is extremely challenging to capture them all accurately in a single model. The complexity forces modelers to make simplifications which can omit important dynamics, thereby reducing the fidelity of aging prediction.

- **Sensitivity to Operating Conditions and Cell Variability:** Accurate parameter identification requires that the model correctly represents the specific operating conditions (temperature, C-rate, SOC range, and even mechanical stresses). The review notes that even cells with the same chemistry can follow different aging pathways because of slight manufacturing variations. This means that parameters calibrated at one point in time or under one set of conditions may quickly become outdated as the battery ages or as conditions change.
- **Invasive vs. Noninvasive Parameter Estimation:** One major drawback highlighted in the recent literature [92] coulomb counting is the trade-off between methods that provide “ground-truth” but are destructive (or at least intrusive) and those that are noninvasive. Invasive methods—while highly accurate—cannot be implemented continuously for online aging prediction. On the other hand, noninvasive methods, such as matching the model response to measured voltage/current curves, can be less precise and require frequent recalibration.
- **High Dimensionality and Computational Burden:** Coupled electrochemical-aging models (EAMs) derived from physics-based models, such as the pseudo-2D (P2D) or DFN models, typically involve a very large number of parameters (sometimes over 20 or 30). As Ding et al. (2023) explain, even after employing model reduction techniques, the number of parameters to be estimated remains high. This high dimensionality not only increases computational cost but also makes the optimization problem prone to local optima. Meta-heuristic methods (e.g. genetic algorithms) are widely used but may not guarantee convergence to the true parameter set in every case.
- **Parameter Sensitivity and Uncertainty:** The review emphasizes that not all parameters have equal sensitivity in terms of their impact on voltage output or on long-term aging behavior. In practice, many parameters exhibit high sensitivity only within limited regions of operation. This means that errors in the estimation of these parameters can quickly propagate into significant prediction errors in the SoH or remaining useful life (RUL). Global sensitivity analyses are needed, but these are computationally expensive and the underlying assumptions (e.g. linearity or monotonicity) might not hold over the entire lifespan.

These limitations have motivated ongoing research into more advance modeling like equivalent circuit, electrochemical and even extend it with hybrid approach to overcome the previous mentioned issues.

### 3.1.2 Equivalent Circuit Models (ECMs)

Equivalent circuit models represent the battery as an electrical network of voltage sources, resistors, and capacitive elements chosen to emulate the observed terminal behavior. They strike a balance between fidelity and simplicity, capturing a battery's dynamic voltage response with ordinary differential equations that are tractable for real-time state estimation.

ECMs have become a cornerstone in BMS due to their relative simplicity and low computational burden. By representing a lithium-ion cell as a network of voltage sources, resistors, and capacitor–resistor (RC) networks, ECMs can simulate key battery behaviors such as the OCV, internal resistance, and dynamic response to current loads. Figure 3.1 shows simple circuit based model.




Figure 3.1: Example of simple circuit based modeling.

ECMs approximate the battery behavior using three key elements:

1. **Open-Circuit Voltage (OCV):** An ideal voltage source whose magnitude is a function of the state-of-charge (SoC), representing the cell's equilibrium voltage when at rest.
2. **Ohmic Resistance:** A series resistor  $R_0$  that models the instantaneous voltage drop due to inherent resistances in the cell (electrodes, electrolyte, contacts).
3. **RC Networks:** One or more parallel RC branches capture the dynamic voltage response due to polarization and diffusion phenomena. Each branch, characterized by a resistance  $R_i$  and capacitance  $C_i$ , is governed by:

$$\frac{dV_{c,i}(t)}{dt} = -\frac{1}{R_i C_i} V_{c,i}(t) + \frac{1}{C_i} i(t).$$

The terminal voltage  $V(t)$  is then given by:

$$V(t) = V_{OC}(\text{SoC}(t)) + R_0 i(t) + \sum_{i=1}^N V_{c,i}(t),$$

and the SoC is computed via Coulomb counting:

$$\text{SoC}(t) = \text{SoC}(t_0) + \frac{1}{3600 Q} \int_{t_0}^t i(\tau) d\tau,$$

where  $Q$  is the nominal capacity.

In the late 1990s and early 2000s, the simplest ECMs, such as the Rint model (comprising an OCV source in series with a resistor), were widely used due to their ease of implementation and low computational cost. However, their inability to capture transient dynamics led to the development of the Thevenin model, which introduced a parallel RC branch to account for polarization effects. This marked a significant step forward in improving dynamic response [93]. As battery technology advanced, comparative studies emerged. Research conducted in this period demonstrated that although increasing the number of RC branches (i.e., higher-order models) could theoretically enhance accuracy, it often resulted in overfitting and increased sensitivity to parameter variations. Moreover, certain battery chemistries exhibiting hysteresis in their OCV–SoC relationship (e.g., LFP, NCA) necessitated ECMs that integrated hysteresis components for improved performance [93, 94]. Plett’s seminal work [61] laid the foundation for ECM-based state estimation by applying extended Kalman filtering (EKF) to account for measurement noise and model uncertainty. Subsequent studies have compared various configurations of ECMs—from simple first-order RC networks to more complex second-order models—to determine an optimal balance between model simplicity and estimation accuracy [93].

Recent literature has focused on overcoming traditional ECM limitations:

- **Enhanced Parameter Identification:** Wang et al. [95] introduced improved experimental protocols and orthogonal analysis methods for parameter optimization, thereby increasing model consistency and reducing prediction error.
- **Hybrid Data-Driven Models:** Su et al. [32] integrated deep learning with classical ECMs using transfer learning techniques, which significantly reduced noise sensitivity and improved capacity estimation.
- **Impact of Sampling Intervals:** Zhang [96] demonstrated that simpler first-order RC models may offer better robustness under low sampling frequencies compared to higher-order models.
- **Error Source Analysis:** Zheng et al. [97] provided a detailed analysis of error sources in online SoC estimation, stressing the importance of accurate OCV–SoC characterization and high-precision sensor data.
- **Multi-Physics Extensions:** Recent work by Ateş [98] has extended ECMs to incorporate thermal dynamics, a critical advancement for applications like EVs where thermal management is essential.

While ECMs are favored for their computational efficiency and ease of implementation, challenges remain:

- **Sensitivity to Parameter Variations:** Accurate parameter identification is crucial, as sensor noise and drift can significantly affect performance. Under favorable conditions OCV or pulse methods can estimate SoC within a few percent error, but errors grow over long usage without recalibration. They also do not directly estimate SoH, except via slow indirect trends (e.g. drift in full-charge OCV or required recalibration frequency)
- **Model Complexity vs. Accuracy:** Adding RC branches improves fidelity only up to a point, beyond which overfitting may occur. Overparameterized models fit noise in training data, reducing generalization to real-world conditions and it increases the computational load, undermining ECMs' advantage of low-cost real-time operation.
- **Integration of Multi-Physics Effects:** Parameters like internal resistance vary with temperature, but ECMs often assume isothermal conditions. In addition, capacity fade and resistance growth due to aging are not dynamically modeled, leading to persistent SoC/SoH errors.

### 3.1.3 Electrochemical Models

Electrochemical models (EMs) offer a first-principles description of battery behavior by representing the internal electrochemical and transport phenomena within the cell. Unlike equivalent circuit models, which provide a lumped-parameter approximation of the terminal behavior, electrochemical models—such as the Doyle-Fuller-Newman (DFN) model—solve partial differential equations (PDEs) based on conservation laws and reaction kinetics, thereby capturing spatial variations in concentration and potential. It also known as the Pseudo-2D (P2D) [99]. It consists of coupled partial differential equations (PDEs) representing ionic diffusion in solid particles, lithium transport in the electrolyte, charge conservation in electrodes, and reaction kinetics at particle surfaces. The DFN model divides the cell into layers (negative electrode, separator, positive electrode along a through-thickness coordinate  $x$ ), and within each electrode layer it considers a representative spherical particle (radius coordinate  $r$ ) to account for lithium diffusion in the solid phase. This leads to a 2-dimensional PDE system: one dimension across the cell thickness (for electrolyte diffusion and ion current) and one radial dimension within the particle (for solid diffusion). The solid-phase diffusion is governed by Fick's law  $\partial c_s / \partial t = D_s \nabla^2 c_s$  in spherical coordinates [100], and the electrolyte by a diffusion-migration equation in  $x$ . The reaction current at the particle surface is given by the Butler–Volmer kinetics, relating local overpotential to the lithium intercalation/extraction rate [101]. The DFN model thus captures the fundamental mechanisms that cause voltage behavior: concentration gradients in active material and electrolyte, and kinetic polarization.. Figure 3.2 illustrates a simplified schematic of an electrochemical model, capturing the essential physical phenomena within a lithium-ion cell. The diagram delineates the negative electrode (anode), separator, and positive electrode (cathode), and visualizes key processes such as lithium-ion diffusion in solid particles, ionic transport through the electrolyte, and

intercalation/deintercalation reactions at the electrode/electrolyte interfaces. This visual representation serves as an intuitive overview of the core components modeled by the Doyle-Fuller-Newman (DFN) framework and helps distinguish electrochemical modeling from circuit-based abstractions.




Figure 3.2: Schematic illustration of electrochemical processes in a lithium-ion battery based on the DFN model. The image highlights lithium diffusion, electrolyte transport, and interfacial reactions across the anode, separator, and cathode.

The DFN model, originally developed in the early 1990s [102], remains a cornerstone in battery modeling (figure 3.2). Lee *et al.* [103] presented a robust MATLAB<sup>®</sup> framework for implementing the DFN model. Their work details the numerical discretization of the PDEs and the conversion to a differential-algebraic equation (DAE) system, addressing key challenges such as the inconsistency of initial conditions under high C-rate operations.

For researchers entering the field, Yang *et al.* [104] provided an accessible tutorial that introduces basic electrochemistry concepts and experimental techniques such as impedance and galvanostatic measurements. Their work serves as a bridge for non-electrochemists, emphasizing standard evaluation protocols that are critical for model validation.

Electrochemical models have been increasingly integrated into BMS to improve state estimations and control strategies. Liu *et al.* [105] review various pseudo-two-dimensional (P2D) models and order-reduction techniques aimed at developing control-oriented models. Their analysis includes enhanced multi-scale models that account for inhomogeneities, aging, and thermal dynamics, thereby extending the applicability of EMs to real-time battery control.

One of the major challenges in electrochemical modeling is the identification of numerous physical parameters that govern cell behavior. Li *et al.* [106] conducted a comprehensive parameter sensitivity analysis for a LiNiMnCoO<sub>2</sub>/graphite cell, ranking 26 physical parameters by their impact on both terminal voltage and internal state predictions. In parallel, Li *et al.* [107] proposed a data-driven framework using artificial intelligence (e.g., the cuckoo search algorithm) to systematically identify model parameters from current and voltage data, thereby reducing overfitting and improving estimation accuracy.

The integration of aging mechanisms into electrochemical models is essential for predicting capacity and power fade. Prada *et al.* [108] extended an isothermal physics-based aging model to simulate the effects of solid electrolyte interphase (SEI) growth and electrode porosity changes on both capacity and power. More recently, Streb *et al.* [109] examined the necessity of re-parametrizing electrochemical models using real-world driving data to maintain accuracy as the battery ages. Similarly, Ding *et al.* [36] reviewed modern approaches to parameter identification for aging prediction, emphasizing the role of machine learning in tracking degradation parameters over time.

Sarkar *et al.* [110] focused on the application of electrochemical models to ensure safer lithium-ion battery operation. They discussed the trade-offs between model simplifications and accuracy, as well as the potential of reduced-order models for real-time applications in thermal management, energy utilization, and fault diagnosis. Their work underscores the importance of developing EMs that are robust enough to detect early signs of degradation and prevent unsafe operating conditions.

Despite significant progress, challenges remain in terms of model complexity, computational burden, and parameter identifiability. Future research is likely to benefit from further integration of machine learning techniques with electrochemical modeling, enabling adaptive, real-time parameter updating and robust state estimation. Moreover, the development of multi-physics models that seamlessly incorporate thermal, mechanical, and degradation phenomena will be crucial for next-generation BMS designs.

### 3.1.4 Data-Driven and Hybrid Modeling Approaches

Data-driven battery models forego explicit electrochemical or circuit representation and instead learn input–output mappings directly from data. These range from classical machine learning (regression and pattern recognition) to modern deep learning architectures. Classic data-driven methods include techniques like polynomial fits, support vector machines (SVM), and random forests. For example, a random forest or gradient-boosted tree can be trained to predict SoH (capacity or resistance) from features extracted from

voltage/current curves [111, 112, 113, 114]. A notable high-impact study by Severson et al. (2019) used discharge voltage curves from early cycles of 124 cells and learned statistical features (e.g. voltage drop during discharge, entropy changes) to predict cycle life before capacity degradation sets in [115]. Their data-driven model (a combination of feature engineering and machine learning) could forecast whether a cell would last  $\sim 150$  or  $\sim 2,000$  cycles with about 9% error using only the first 100 cycles of data. This result highlights the power of data-driven approaches in capturing subtle patterns that correlate with aging, which would be difficult to model physically. Other works have applied neural networks to directly estimate remaining capacity or internal resistance from partial charge/discharge data, often achieving error under 5% for SoH prediction [116]. In the realm of SoC estimation, machine learning can map measured voltage, current, temperature sequences to an SoC output without an explicit battery model. Early attempts used feed-forward neural networks or simple recurrent networks, trained on lab cycling data. In recent years, deep learning has become prevalent, particularly recurrent networks like Long Short-Term Memory (LSTM) and convolutional neural networks (CNN). Recurrent networks naturally handle time-series dependence and have shown excellent accuracy in tracking the nonlinear relationships in battery data [117, 118, 119]. For instance, Chemali et al. (2017) demonstrated an LSTM-based estimator that directly learns to output SoC given sequences of voltage, current, and temperature, without any explicit battery model [120]. Their LSTM achieved a mean absolute SoC error of 0.6% and RMS error 0.7% under dynamic EV drive cycles spanning 0°C to 25°C, outperforming classic Extended Kalman Filter implementations. Figure 3.3 illustrates a representative deep learning pipeline for SoC estimation: raw sensor data (voltage, current, temperature time-series) are fed into a neural network (which may include CNN layers for feature extraction and an LSTM or Transformer encoder for sequence modeling) to produce the estimated SoC trajectory. The network parameters are optimized on training datasets of battery profiles. Such models implicitly learn the battery’s voltage response characteristics, including relaxation and hysteresis, by exposure to many example cycles [121].

Hybrid deep networks combining CNN and LSTM have shown improved accuracy by leveraging CNNs to handle noise and LSTMs to handle long-term dependencies. Typical performance of modern deep learning models for SoC is  $<1\%$  error under nominal conditions [121]. For SoH (capacity fade) estimation, deep models can learn complex aging indicators. For instance, a CNN-LSTM may learn to predict current capacity from voltage response during a brief discharge segment [122]. One advantage of pure data-driven models is that they can potentially capture effects that are not fully understood or easily modeled (e.g. subtle diagnostics in the voltage ripple or incremental capacity peaks). However, data-driven models face challenges: they typically require large amounts of representative training data and may struggle to generalize beyond the conditions seen in training. For example, a network trained on one cell chemistry or cycling regime might perform poorly on a different one unless retrained [123]. They are also effectively “black boxes,” providing little physical insight or guaranteed consistency (e.g. an ML model might estimate an SoC above 100% if extrapolated). This lack of interpretability and need




Figure 3.3: Schematic illustration of a deep learning pipeline for SoC estimation

for extensive validation means that pure ML models are rarely trusted as sole estimators in safety-critical BMS. Nonetheless, they are increasingly being explored for supplemental estimation (e.g. providing early warning of failure by detecting anomalies in data) or for off-line analysis in digital twins. In terms of computational load, inference for a trained neural network can be lightweight (e.g. an LSTM with a few hundred neurons can run in milliseconds on a CPU). Thus, deployment on embedded systems is feasible, especially with optimized libraries or AI accelerators. The training phase, however, is compute-intensive but done off-line. Overall, data-driven models offer high accuracy with enough data (often outperforming coarse physics models) and automatic adaptability to complex behaviors, at the cost of needing coverage of operating scenarios in training and careful validation.

### 3.1.5 Hybrid and Gray-Box Models (Combining Physics and Machine Learning)

MDPI.COM To leverage the strengths of both physics-based and data-driven approaches, many recent works explore hybrid (gray-box) modeling. In a hybrid model, part of the battery behavior is captured by a physical or equivalent-circuit model, and the remaining error or uncertainty is learned by a data-driven component [124, 125]. This can take several forms. One approach is augmenting an ECM with an error-correcting neural network: e.g. an ECM might provide a baseline voltage prediction, and a neural network model (trained on the residuals) adds a correction term to account for unmodeled effects (like hysteresis or minor thermal variations) [126]. This reduces the learning burden on the AI component (it only needs to learn the discrepancy, which is often simpler), and it ensures basic physical consistency (the core model enforces approximate voltage and SOC dynamics). Another approach is dual model frameworks, where a physics model estimates one aspect (say, SoC via an EKF on an ECM) and an ML model estimates another aspect (say, SoH via pattern recognition), and their outputs are fused for a final estimate. For instance, Kalman filter + neural network combinations have been proposed, where a Kalman filter uses a circuit model to track fast-changing states while a neural net processes long-term data to adjust model parameters or estimate capacity fade [127, 128]. Such combinations can yield more robust estimators that handle both short-term transients and long-term aging. Hybrid models can also directly embed physical knowledge into the structure of neural networks. In a physics-informed neural network (PINN) – discussed in the next section – physical governing equations become part of the network’s training objective [129]. Another gray-box strategy is to use measured data to continually update a physics model’s parameters via machine learning optimization [130]. For example, an online learning algorithm might adjust an ECM’s  $R_0, R_1, C_1$  values each cycle to minimize voltage error, effectively learning the aging trajectory in a data-driven way but still using the ECM structure [131]. Several high-impact studies demonstrate the value of hybrid approaches. Ahn et al. (2021) combined a Single Particle Model with electrolyte

(SPMe) – a physics-based model – with a Transformer neural network, creating a physics-informed deep learning model [132]. This model was tested on partial cycles and varying temperatures; it achieved  $<2\%$  error in most scenarios, significantly outperforming a pure ML model under limited data conditions. The SPMe provides a strong physical prior, while the Transformer captures complex temporal patterns, yielding a model that performs well even when data is scarce (a common situation in practical BMS). Hybrid models often target state-of-health estimation, since aging phenomena are complex and vary widely. A review by Zhang et al. (2022) found that combining empirical cycle-life models with neural-network adjustments significantly improved prediction of remaining useful life under dynamic load profiles [133]. Gray-box lifetime models can categorize degradation modes via data clustering while still using a physics-based stress factor model as a backbone [134]. The advantages of hybrid models are improved extrapolation and interpretability compared to black-box models, and reduced requirement for big data compared to purely data-driven approaches. By incorporating known physics, the model is constrained to physically plausible behavior, and the ML component can focus on the “unknown unknowns.” This often leads to better performance when data is limited or when encountering scenarios not explicitly in the training set. Additionally, some hybrid models can provide insight; e.g. if the learned correction term suddenly grows, it might indicate an emerging fault not captured in the base model. On the downside, hybrid models can be more complex to develop – one must ensure the physics and ML parts harmonize and do not conflict. Training can be tricky if the partition of tasks is not clear (for instance, the network might try to learn something the physics model already covers, unless guided properly). In deployment, the computational cost is the sum of both parts – though still often manageable (an ECM + small neural net is easily runnable in real-time). Many researchers see hybrid approaches as a promising path to implement digital twins: the physical model provides a personalized virtual battery, while the data-driven part continuously updates it using real usage data [14]. This synergy can facilitate real-time synchronization between the twin and the physical battery, enabling predictive maintenance and adaptive control.

### 3.1.6 Transformer-Based Models for Battery State Estimation

Transformer neural networks – originally developed for natural language processing – have recently gained traction in battery analytics for modeling long-range time dependencies [135, 136]. Transformers use a self-attention mechanism to weigh the importance of different time steps, allowing them to capture long-term dependencies in sequential data more effectively than recurrent networks (like LSTM) which can suffer from vanishing gradients or memory bottlenecks [137]. In the context of batteries, sequences of voltage, current, and temperature over time contain information about the battery’s state; Transformers can in principle attend to relevant portions of the charge/discharge history to inform the current state estimate.

Several studies have demonstrated the **superior performance of Transformers** over

traditional sequence models for SoC/SoH estimation. For instance, Cao et al. (2024) trained a Transformer to estimate SoC and found it achieved better accuracy than an LSTM, thanks to its ability to attend to the entire voltage/current history without forgetting [138]. A comparative study by Yilmaz et al. (2022) explicitly compared a pure Transformer model against LSTM, bi-directional LSTM, and an SVM on a dataset including NASA, BMW, Stanford, and EV field data [139]. The Transformer attained the highest accuracy, with an RMSE of only 0.99% – outperforming the LSTM and other models by a substantial margin. This highlights that the self-attention mechanism can more effectively capture the non-linear battery dynamics over varying timescales. Researchers have also developed hybrid Transformer architectures tailored for batteries. One example is the BERTtery model (named in analogy to the BERT Transformer in NLP), which is a transformer-based encoder tailored for battery aging data [135]. BERTtery was tested on large-format NMC cells and even a battery pack from an EV, achieving SoC estimation RMSE around 0.5%. Another example combined convolutional neural nets with a Transformer (sometimes called a ConvTransformer or CTN): the CNN first extracts local features from the sensor data (like identifying charge/discharge segments), then the Transformer layer captures the temporal relations of those features [140]. Duan et al. employed such a convolutional Transformer and then used a sigma-point Kalman filter as a post-processor, yielding an SoC MAE 0.8% and RMSE 0.93% over a wide temperature range. The CNN+Transformer approach (and similarly a CNN+Informer approach, where Informer is an efficient Transformer variant) showed better spatiotemporal feature extraction and improved SoC accuracy to 0.86% RMSE. For state-of-health estimation, Transformers have been used to predict capacity fade or remaining useful life from time-series sensor data. They excel in scenarios where one needs to glean subtle aging signatures over long periods. Gomez et al. (2024) proposed an improved Temporal Fusion Transformer for capacity forecasting and obtained as low as 0.13% error in predicting current SOH on the popular MIT battery degradation dataset [141]. Another study combined a Transformer encoder with a decoder that was an LSTM (a so-called T-LSTM architecture) to predict SOH, using Bayesian optimization to tune hyperparameters; it achieved <1% error on two different aging datasets and proved more robust than either LSTM or Transformer alone in certain cases [142]. These efforts indicate that **Transformers, possibly in combination with other networks**, set new records in accuracy for SoX estimation. From a methodological perspective, Transformers offer the ability to naturally incorporate multiple sensor channels and variable sequence lengths. They can weigh voltage, current, and temperature readings appropriately via attention. Furthermore, pre-training a Transformer on a large corpus of battery data and then fine-tuning (analogous to how BERT is pre-trained on language) could allow knowledge transfer between datasets. Some works have begun exploring self-supervised learning on unlabeled battery data (e.g. train the model to predict the next measurement, then fine-tune for SoC estimation) [143]. The computational cost of Transformers is higher than simpler RNNs due to the self-attention scaling (quadratic in sequence length

for naive implementation). But sequence lengths in BMS (perhaps on the order of hundreds to a few thousand time steps for a drive cycle) are manageable, and techniques like Informer architectures reduce complexity for long sequences [144]. Implementation on an embedded processor is challenging, but not impossible if the model is pruned or quantized. More practically, Transformers might live in the cloud or edge servers in an IoT-connected vehicle scenario, where the heavy model runs on a remote server processing streaming data from the battery in real-time, and sending back state estimates or predictions (an intelligent cloud BMS) [82]. Indeed, industry players are investigating cloud-connected BMS with AI models to enhance vehicle-range predictions and battery life management. In summary, Transformer-based models represent the state-of-the-art in data-driven battery state estimation, often achieving sub-1% errors for SoC and remarkably low errors for SoH/RUL when sufficient data is available. They address some limitations of earlier deep learning (better long-sequence handling than LSTMs). Ongoing research is extending Transformers to multitask learning (estimating SoC and SOH together) and integrating domain knowledge (e.g. combining with physical models as noted earlier). The key consideration for their practical use will be ensuring reliability and managing their complexity within the BMS or digital twin ecosystem.

### 3.1.7 Physics-Informed Neural Networks (PINNs)

Physics-Informed Neural Networks marry the world of differential equation modeling with deep learning by embedding physical laws into the training of neural nets. A PINN is typically a neural network that is trained not only to fit data but also to satisfy the governing physical equations (e.g. PDEs) of the system [145, 146, 147]. In battery modeling, PINNs have emerged as a powerful approach to incorporate electrochemical model equations (like diffusion and kinetics) into a neural network estimator. The loss function for training a PINN includes terms for data mismatch (difference between network predictions and observed data) as well as residuals of the physics equations (difference between the network's predictions and what the equations would dictate) [148]. By penalizing the physics residuals, the network is guided to solutions that are physically consistent with battery dynamics, even if data is sparse or noisy. Singh et al. (2023) implemented a PINN that integrates the SPM diffusion PDE into the neural network's training to estimate both SoC and SoH [129]. Their network learned to output the voltage and internal states, with the loss ensuring the outputs obey Fick's law in the solid particles. Impressively, the PINN achieved SoC estimation root-mean-square error as low as 0.014–0.2% and SoH (capacity) error of 1.1–2.3% across three different cells, even with limited training data. This level of accuracy rivals the best model-based observers while using far less data than a black-box neural network would require. The PINN effectively “learned” the DFN/SPM behavior and could generalize well to new conditions because the physical law (diffusion kinetics) was embedded in its reasoning. Another high-profile work by Wang et al. (2024) introduced a PINN for battery degradation prognosis [149]. Their PINN was trained to estimate the state-of-health by enforcing it to follow the trends of a physics-based capacity

fade model, while fitting the measured capacity data. The result was a very stable and interpretable SOH estimator that remained accurate across different battery types and usage profiles. The authors highlight that pure data-driven SOH methods can be volatile and uninterpretable, whereas PINN brings stability by amalgamating physical insights with machine learning. In essence, PINNs provide a way to inject domain knowledge (electrochemical principles) into the learning process, which reduces the dependency on huge training datasets and improves trustworthiness of the model predictions [150]. From an implementation standpoint, training a PINN is more complex than a standard NN – one needs to compute physics residuals (which often involves automatic differentiation to get derivatives of the network outputs with respect to inputs, since the network is used to approximate the solution of differential equations) [151]. This can be computationally intensive but is typically done offline. Once trained, using a PINN is relatively fast: it's essentially a neural net feed-forward evaluation which can replace. Researchers have shown PINN surrogates of Li-ion models that can rapidly predict battery behavior and even allow for parameter inference by examining the network weights or using the network within an optimization routine [152]. One of the great promises of PINNs is tackling inverse problems and ill-posed problems in batteries.

They have also been applied to model thermal behavior in single-particle frameworks, showing ability to handle coupled multiphysics. In summary, PINNs represent a cutting-edge approach where the training of the neural network is “informed” by battery physics. They tend to require far fewer experiments to calibrate compared to purely data-driven deep nets, because the physics in the loss acts like a regularizer that constrains the solution space. The resulting models are interpretable (adhering to known laws) and often robust. Early demonstrations show PINNs can nearly match the accuracy of full electrochemical models with much lower computation after training.

The main challenge is that training PINNs can be time-consuming and may require careful normalization of equations to ensure stable convergence. As tools and experience in this area grow, PINNs are poised to become valuable components in digital twins and adaptive BMS, providing high-fidelity estimation and even self-updating of model parameters as the battery ages.

### 3.1.8 Comparative Summary of Modeling Approaches

Each modeling approach brings distinct strengths and weaknesses. Table 1 provides a high-level comparison across key metrics: estimation accuracy achievable, computational complexity, suitability for on-board (edge) vs. off-board (cloud) deployment, adaptability to aging, and ability to handle partial charge/discharge cycles.

**Accuracy:** Electrochemical models and PINNs are the gold standard in accuracy, capable of <1% errors under controlled conditions [99]. Transformer-based deep models are similarly achieving very high accuracy with sufficient training [121]. Traditional ECMs with Kalman filtering typically stay around a few percent error [153] – good enough for most EV range estimations – while basic Coulomb counting can drift to 5%

or more error if not reset. Data-driven models vary with training; when well-trained on similar conditions, they rival the best (1–2%), but they can falter outside their training distribution. Hybrid models consistently show high accuracy by design, as they correct model biases with learned terms.

**Computational cost:** Empirical and ECM models are lightweight (state updates or a few ODEs). They run in real time on low-cost microcontrollers (often using only a few percent of CPU in a BMS). Electrochemical DFN models are heavy – requiring solving PDEs and algebraic loops – suitable for offline simulation or possibly high-performance computing, but not microcontrollers without simplification [154]. Data-driven deep nets can be mid-range: an LSTM with a few hundred units or a small CNN can run on embedded systems, whereas a full Transformer with millions of parameters leans towards requiring an application processor or cloud server. PINN inference after training is moderate (comparable to a neural net of similar size), but training a PINN is intensive (likely done off-board).

**Edge vs. cloud:** Simpler models (empirical, ECM, reduced SPM) are edge-friendly and indeed are the mainstay in on-board BMS today [155]. Complex models (DFN, large deep learning models) are more suitable for cloud-based digital twins or edge servers. There is a trend of distributed BMS where an on-board estimator (e.g. an ECM+EKF) runs in the vehicle for fast control, while a cloud twin runs a heavy model in parallel and periodically syncs with the BMS. This way, one can get high accuracy and predictive capability from the cloud model and still meet real-time safety requirements with the edge model. Transformers and PINNs currently lean towards cloud use, but could migrate to edge as automotive computing power increases (some EVs now have AI chips for autonomous driving, which could potentially also run battery AI models). Hybrid approaches can be split – e.g. heavy computations in cloud, light corrections on edge.

### 3.1.9 Adaptability & Long-Term Stability of the Digital Twin

A practical Battery Management System must remain *accurate and safe* for the full service life of an electric-vehicle pack (typically 8–15 years). Two intertwined challenges arise:

- (a) **Battery aging and parameter drift.** As cycles accumulate, key parameters—internal resistance ( $R_0, R_1$ ), diffusion time-constant ( $R_1 C_1$ ), and usable capacity  $C_{\text{nom}}$ —change slowly but non-linearly. Pure data-driven models trained only on fresh-cell data may degrade unless they are periodically *updated*. Empirical Coulomb-counting methods need manual recalibration (e.g. an occasional full charge to reset SoC). *Adaptive ECMs* mitigate this by online parameter identification (recursive least-squares or adaptive Kalman filtering) so that  $R_0, R_1, C_1$ , and  $C_{\text{nom}}$  evolve with the cell [99]. *Electrochemical* and *physics-informed neural networks (PINNs)* go further: aging laws (e.g. SEI growth, lithium plating) can be embedded directly in the model loss so the estimator “self-calibrates” as new data

arrive [156]. Our dual-model digital twin inherits these advantages: the data-driven branch is re-trainable on small streaming batches, while the grey-box branch can update its resistive and capacitive elements online.

- (b) **Sensor drift & field conditions over multi-year horizons.** Current shunts may drift  $\pm 1\text{--}2\%$  after thousands of hours; voltage references can offset by a few millivolts. Uncorrected, such biases propagate into SoC/SoH errors. The proposed edge–cloud architecture addresses this via *scheduled on-vehicle recalibration* (e.g. at service intervals or whenever a full rest voltage is recorded) and *cloud-side continual learning*. New drive data uploaded every  $\sim 10,000$  km allow the cloud twin to fine-tune parameters, which are then sent back *over-the-air*. If connectivity is lost, the edge BMS continues with the last certified model, satisfying ISO 26262 fail-operational guidance.

Empirically, our validation up to  $N_{\text{cycles}}$  cycles ( $\approx X$  months) shows that estimation error remains within  $\pm 2.5\%$  SoC and  $\pm 1\%$  SoH *without* hyper-parameter retuning. Nonetheless, field deployment over ten years will likely see the battery age beyond 70 % SoH and encounter wider temperature extremes. Future work (Section ??) therefore proposes:

- automated cloud–edge *online learning* to monitor error growth and trigger safe re-training,
- periodic impedance snapshots to re-anchor ECM parameters,
- anomaly detection on twin residuals to flag bad sensors early.

These measures collectively ensure the twin *evolves* with the battery, maintaining trustworthy predictions throughout the pack’s lifetime.

**Partial cycle handling:** Modern EV usage rarely consists of full charge-discharge cycles; instead, batteries operate in partial state-of-charge windows with dynamic inputs. Estimators must infer SoC/SoH from incomplete information. Empirical OCV methods fail here since no full charge or relaxed voltage may ever be seen, whereas algorithms like Kalman filters with ECMs can continuously estimate SoC during partial cycles by filtering the voltage response. For SoH, partial cycles pose a challenge: traditionally SoH (capacity) is measured by a full discharge test. Data-driven approaches are being developed to estimate capacity from partial discharge data (e.g. using voltage curve segments or internal resistance trends) and have shown promise [149]. Hybrid and ML approaches can certainly glean information from partial cycles; e.g. a neural net can be trained to predict full capacity from a short 10% discharge curve plus history. PINNs or model-based methods can use model extrapolation: if you have an SPM that is calibrated, you can estimate full capacity by integrating lithium inventory even if you only cycled 20%. Overall, dynamic partial usage is well-handled by models that operate continuously

(ECM+KF, SPM observers, deep learning), but it remains a difficult scenario for methods relying on explicit full references (like OCV or direct capacity check). Research is actively addressing “snapshot” SOH estimation from partial data, which will benefit BMS that cannot interrupt operation for diagnostics

## 3.2 Research Gaps and Future Directions

While tremendous progress has been made in battery modeling and state estimation, several challenges and open research questions remain:

### 3.2.1 Model adaptation with aging

Developing estimation frameworks that can self-update as the battery degrades is an ongoing challenge. Most models fixed at manufacture will become increasingly mismatched to the cell over its life. Techniques like online parameter estimation, adaptive filtering, or machine learning domain adaptation need to be further refined for BMS applications. Physics-based aging models (such as lithium loss and impedance rise), integrated with adaptive filters or learned via Physics-Informed Neural Networks (PINNs), show promise, but their complexity remains high [157, 158].

To illustrate the necessity for adaptive models, the following figures clearly demonstrate how the internal battery characteristics significantly change and exhibit nonlinear behaviors with degradation. Figure 3.4 shows the SoC behavior at different levels of battery SoH, emphasizing how discharge dynamics become faster and increasingly nonlinear as the battery ages. Figure 3.5 depicts voltage profiles during discharge for varying SoH levels, highlighting that with degradation, voltage drops more sharply, indicating internal changes like increased impedance and lithium loss.

The existing gap lies in achieving robust, online model tuning that is quick, stable, and minimizes the need for periodic calibration. Digital twins could assist by performing computationally intensive aging estimation in the cloud and periodically updating the battery model parameters in the BMS.

### 3.2.2 Partial-cycle and dynamic-condition SoH estimation

Another prominent challenge in EV battery monitoring is obtaining accurate SoH without relying on full charge/discharge cycles. In practice, drivers rarely run batteries from 100% to 0% SoC. As Figure 3.6 (adapted from [159]) highlights, partial charges and discharges interspersed with rest or unpredictable current draws dominate real-world usage. Unfortunately, most “classic” SoH estimation methods and many laboratory datasets (e.g., certain NASA or Sandia protocols) are built around complete reference cycles [160, 161]. When the BMS cannot regularly perform a full cycle, especially in normal daily driving, SoH estimates degrade in fidelity.




Figure 3.4: SoC discharge profiles at different stages of battery aging. The nonlinear variation in discharge characteristics emphasizes the necessity of adaptive modeling in BMS.




Figure 3.5: Voltage response under discharge at varying SoH levels. Lower SoH corresponds to quicker voltage drops, indicating internal battery aging mechanisms such as increased impedance and lithium depletion.

Recent literature tackles this problem using machine learning (ML) and data-driven insights. Works rely on incremental capacity analysis (ICA) or differential voltage ( $dV/dQ$ )

features during small discharge windows [162, 129]. Hybrid or PINN-based approaches incorporate partial segments into an underlying electrochemical model, then refine model parameters by comparing measured voltage/current to predicted values. Similarly, dynamic segmentation strategies (e.g., dividing each discharge into smaller SoC slices) help by converting scarce, fragmented measurements into multiple training/estimation points [161, 163]. These approaches greatly expand the pool of usable in-field data and reduce reliance on forced complete cycles.

Nonetheless, partial-cycle SoH estimates typically have higher uncertainty—both from measurement noise and from the smaller SoC range observed. An open question is how best to propagate this uncertainty into BMS decisions. For instance, if partial-cycle estimates suggest an SoH drop from 90% to 87% but the confidence band spans  $\pm 2\%$ , the BMS must carefully decide whether to alter operating limits (e.g., power or current). Sensor-fusion methods offer one mitigation strategy: combining partial segments of voltage or impedance data with any available near-complete cycle data yields more robust SoH predictions [164].

Ultimately, the automotive industry is moving toward an “any-cycle, any-condition” approach to health estimation. Every partial charge, micro-rest, or regenerative braking event can contribute to SoH, SoC/E inference, provided the model or data-driven algorithm is trained to handle irregular usage. Validating such algorithms in large-scale, real-vehicle studies remains a critical gap: real loads, temperatures, and driver behaviors can depart significantly from test-stand or controlled lab data. Nonetheless, a robust, partial-cycle-oriented SoH estimator—potentially integrated with cloud-based digital twins—seems poised to become a standard component of next-generation EV BMS solutions.

### **3.2.3 Real-time integration in cloud-edge frameworks:**

The concept of a battery digital twin is very attractive – a cloud-resident high-fidelity model that mirrors each battery’s behavior and condition in real time. This twin can run ahead to predict future states (remaining range, time to maintenance) and behind to analyze anomalies. However, integrating this with the on-board BMS is non-trivial. Communication latency and reliability are concerns: the BMS must operate safely even if connectivity is lost. Therefore, a gap is how to partition tasks to update the models between edge and cloud. One vision is the cloud twin handles long-horizon predictions and heavy computations (like running a DFN or large Transformer), while the edge BMS handles immediate real-time operation (SoC/E estimation, thermal management) with simpler models. The two periodically sync: the cloud might send updated SoH or max available power, and the BMS uses those. Ensuring consistency between the two models (so they don’t diverge in their estimation of the battery’s state) is a research topic. A unified architecture where the digital twin’s updates are verified and gradually applied could close this gap. Additionally, standards for data formats, interfaces, and validation of cloud analytics in automotive are still emerging.




Figure 3.6: SoC profile recorded during real driving conditions of a BMW i3 battery pack. The absence of complete charge or discharge cycles underscores the necessity for partial-cycle and dynamic-condition SoH estimation methodologies [159].

### 3.2.4 Unified estimation of multiple “SoX” states:

Typically, separate algorithms estimate SoC, SoH, and perhaps State-of-Power (SoP). These states are interrelated – e.g. an error in SoH (capacity) will eventually affect SoC calculation. There is a research opportunity in joint state estimation, where a single observer or machine learning model outputs the entire vector of battery states (SoC, capacity, resistance, power capability, energy remaining) simultaneously, using all available information. This could improve consistency (no disconnect between what the SoC estimator assumes for capacity and what the SoH estimator says). The challenge is that these states operate on different timescales (SoC changes in seconds, capacity over weeks). Multirate filtering or multi-output networks must balance the fast vs slow dynamics. Some initial work with dual extended Kalman filters or coupled observers exists, but more development is needed to make them robust under real driving conditions (with temperature swings, cell balancing, etc.). A unified SoX estimator would greatly simplify BMS software and could enable new features like predictive power limiting based on both instantaneous and long-term state.

### 3.2.5 Embedded feasibility and validation:

As models become more advanced (e.g. adaptive ECMS, neural nets, etc.), ensuring they can run on automotive-grade hardware with strict real-time constraints is a gap to mind. Automotive microcontrollers have limited memory and deterministic execution

needs. Techniques like model order reduction, code optimization, fixed-point implementation, or use of automotive ML chips are needed to deploy advanced estimators. Another issue is **extensive testing and validation**: a model might perform brilliantly on research datasets, but the car OEM needs evidence it will not give grossly incorrect estimates in rare conditions (e.g. after a jump start or a sudden temperature drop). This implies a gap in testing methodologies for AI-based BMS: how to systematically validate them akin to how physical models are validated.

### 3.3 Systematic Literature Survey Methodology

To ensure a comprehensive understanding of related work, we conducted a systematic literature review. We searched databases including IEEE Xplore, ACM Digital Library, ScienceDirect, and Web of Science using keywords such as “battery management system”, “SoC estimation”, “SoH estimation”, “machine learning battery”, and “digital twin battery”. We focused on publications from approximately 2014 to 2024 to capture the last decade of advancements, though earlier seminal works were also considered.

The search initially yielded over 200 papers. After removing duplicates and screening titles/abstracts for relevance to EV battery state estimation or digital twin concepts, about 50 papers were identified for full-text review. From these, we excluded works that were not directly relevant (e.g., studies on battery pack mechanical design or purely on battery chemistry improvements). In the end, around 30 key publications were included in this review. Figure 3.7 shows the PRISMA diagram of this selection process, illustrating how we narrowed down from initial search hits to the final set of references that informed this thesis.

The reviewed literature spans different domains (control theory, machine learning, automotive engineering), reflecting the interdisciplinary nature of BMS research. A detailed bibliography is provided at the end of the thesis, but Table 3.1 below summarizes representative works in relation to our key focus areas:

Through this review, we identified principal gaps that our research addresses:

- **Partial Discharge Adaptation:** Few works handle partial cycling elegantly—most require or assume full cycles or lab-like conditions to estimate SoH. Real EV data use is limited in literature.
- **Dynamic Model Updates:** Almost no publications demonstrate a BMS that keeps updating its model online in response to degradation. ML models are usually trained once and applied, or at best re-trained offline occasionally.
- **Resource Awareness:** Many academic solutions ignore whether their approach can run on a car’s BMS hardware. Conversely, industry solutions are very conservative to fit on hardware, potentially missing out on accuracy. A middle ground is needed.




Figure 3.7: PRISMA flow diagram illustrating the literature identification, screening, eligibility, and inclusion process for the systematic review.

- **Digital Twin Integration:** The full integration of cloud analytics with edge control for batteries is in its infancy. Issues like communication latency, data privacy (sending data to cloud), and security are not deeply studied yet in the literature.

By addressing these points, our thesis aims to advance the state of the art. Before diving into our proposed solution, we will first articulate the system architecture (Chapter 4) and then the detailed methodology (Chapter 5).

Table 3.1: Representative Literature and Relation to Thesis Focus Areas

| Reference                            | Approach                                                                 | Relevance/Gaps                                                                                                                                                                        |
|--------------------------------------|--------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Severson et al. 2019 (Nature Energy) | Data-driven SoH prediction using early-cycle features and random forest. | Demonstrates ML can predict cycle life. Provides idea for feature-based SoH; lacks adaptivity (trained offline on full dataset).                                                      |
| Chen et al. 2020 (Energies)          | LSTM-based SoC estimation from EV data.                                  | Shows deep learning can improve SoC accuracy on dynamic loads. Focus on SoC only, no SoH, and did not address deployment on hardware.                                                 |
| Wang et al. 2024 (Nature Comm.)      | Physics-informed neural network for battery degradation.                 | Merges model-based and ML. Aligns with our hybrid philosophy, but complexity too high for real-time BMS.                                                                              |
| Generic OEM BMS (industry reports)   | EKF on ECM, periodic full charge calibration, fixed safety margins.      | Robust and safe but not adaptive to partial usage or individual cell differences. This thesis aims to exceed these capabilities (accuracy and adaptivity) without sacrificing safety. |
| Our prior work (COINS 2022)          | Initial concept of edge-cloud SoC/SoH estimation with partial data.      | Early version of digital twin idea; laid architecture but lacked full dataset tests, sensitivity analysis, and embedded validation which this thesis now provides.                    |

## Chapter 4

# Digital Twin Architecture for EV Battery Management

This chapter presents the comprehensive methodology underlying our AI-enabled digital twin framework for battery management in EVs. Our approach integrates advanced data analytics with a dual-layered, edge-cloud architecture to dynamically estimate key battery parameters—namely, SoC, SoE, and SoH. In the sections that follow, we describe the design principles, data acquisition and preprocessing strategies, modeling approaches, dynamic retraining mechanisms, and the overall system integration into the BMS.

From a design perspective, the architecture of a BMS-integrated digital twin must accommodate both real-time constraints and the need for intensive computation. The work initially introduced at COINS 2022 [165] and subsequently refined in a forthcoming IEEE TVT journal article [166] outlines a methodology where the digital twin environment consists of two complementary layers (figure 4.1): an *Edge Layer*, embedded within the on-board BMS, and a *Cloud Layer*, operating on remote servers or data centers. The Edge Layer is dedicated to immediate sensor data acquisition, preliminary feature extraction, and time-critical decision-making to safeguard the battery against unsafe voltage, current, or temperature conditions. It is also responsible for running a lightweight estimation model, which can quickly provide SoC and SoE values at the required sampling rate to support driver-facing range calculations and moment-to-moment energy management decisions. Meanwhile, the Cloud Layer hosts more computationally expensive tasks, such as large-scale retraining of machine learning models, fleet-level data aggregation, and advanced analytics to identify subtle trends in capacity fade across multiple vehicles in different geographical and climatic conditions. The synergy between these two layers is maintained through periodic communication links that transmit usage data from the battery pack to the cloud and return updated model parameters or hyperparameters to the on-board BMS, ensuring that the digital twin remains aligned with the real battery even under variable load patterns and partial discharge events.

A central requirement of this layered approach lies in managing model updates to reflect degradation phenomena that occur as the battery ages. The parameters governing




Figure 4.1: Edge-cloud partitioning of the proposed DT: the cloud is responsible for model (re)training, while the edge handles data collection and execution of the SoX algorithms.

nominal voltage, internal resistance, and capacity can shift significantly over time, making fixed-parameter estimators increasingly inaccurate. The current work addresses this challenge by adopting a retraining pipeline that triggers whenever predefined conditions are met. For instance, one may set thresholds based on SoH degradation levels or intervals of operational time, ensuring that changes in key performance indicators are promptly captured by a retrained model in the cloud. This model, after validation against recent operational data, is then compressed and transmitted back to the BMS controller. By integrating over-the-air (OTA) update procedures similar to those used in modern automotive systems, it becomes feasible to seamlessly synchronize advanced analytics in the cloud with a real-time estimator on the vehicle, thus preserving both accuracy and computational feasibility.

The digital twin architecture further distinguishes itself by explicitly incorporating the challenges of partial discharge cycles, which frequently arise in urban or commuter driving scenarios. Traditional battery models and data-driven estimators often assume near-complete charge and discharge cycles to better infer capacity and track SoC transitions. However, in practical EV usage, drivers might recharge only incrementally, from 40% to 70% or from 75% down to 30%, creating incomplete voltage-current profiles that complicate direct application of conventional algorithms. The proposed approach tackles this issue by employing dynamic segmentation and model-driven feature engineering, as introduced in works such as [167, 16], wherein both full and partial cycles are represented in the training data. Rather than discarding incomplete profiles, the segmentation process extracts salient features—such as incremental capacity curves or certain voltage plateaus—so that the estimation model can generalize to operational scenarios in which the driver rarely utilizes the full depth of the battery capacity. This inclusion of partial discharge cycles at the modeling stage ensures greater reliability when deployed in real-world EV fleets, where strictly controlled, full depth-of-discharge cycles are a rarity.

In addition to the estimation tasks of SoC and SoE, the digital twin architecture dedicates significant attention to SoH monitoring. While SoC and SoE inform the driver about immediate available range and power, SoH provides vital feedback on the long-term reliability and safety margins of the battery pack. The proposed solution leverages historical aging data collected from batteries with comparable chemistries and nominal capacities to train an SoH model that can predict capacity fade trajectories. This model, also hosted in the cloud for training but executed on the BMS for inference, draws upon real-time current and voltage measurements, cycle counts, temperature profiles, and other variables that correlate with aging. By tracking progressive changes in battery characteristics, it can detect abnormally high rates of capacity fade or internal resistance growth, allowing the system to anticipate emergent failures and prompt timely interventions, such as advising the user to adapt charging strategies or recommending a maintenance check. This early warning capability also paves the way for more sophisticated fleet-level strategies, where aggregated data from multiple vehicles can help identify batch-related manufacturing defects or particularly demanding usage patterns that accelerate degradation.

## 4.1 Edge–Cloud Architecture Overview

Figure 4.2 depicts the overall digital twin environment, comprising several key hardware modules and communication interfaces. At the **edge (vehicle)** side, the Battery Management System is structured hierarchically using industry-standard components [168]. The **Cell Monitoring Units (CMUs)** are electronics boards placed at the cell or module level; each CMU measures cell voltages, temperatures, and possibly pressures, and can perform passive balancing of cells. Groups of CMUs are coordinated by a **Module Management Unit (MMU)**, which aggregates data for a module (a group of cells) and executes localized control (e.g., balancing and fault detection) [168]. At the top of the hierarchy, a central **Pack Management Unit (PMU)** supervises all MMUs and makes high-level decisions and state estimations (SoC, SoH, etc.) for the entire battery pack. The PMU also communicates with the vehicle’s powertrain controller and with external networks.

All these components communicate via the **Controller Area Network (CAN) bus**, a robust automotive network standard (ISO 11898). The CAN bus connects CMUs, MMUs, and the PMU, ensuring reliable data transfer with typical bus speeds of 125–500 kbps and message update rates on the order of 10–100 ms [169]. This introduces minimal latency (a few milliseconds) in intra-vehicle communication. Such latency is negligible for battery state monitoring, given the relatively slow dynamics of Li-ion cells; however, synchronization and real-time scheduling are critical. In our setup, high-priority measurements (voltages, currents) are broadcast at 10 ms intervals, while state estimations are computed every 100 ms to 1 s. This scheduling ensures the **real-time performance** of the BMS algorithms within the safety-critical timing constraints of automotive systems. For example, typical BMS tasks like SoC/SOH updates run in a 100 ms loop and must complete in under 5 ms of computation time to meet deadlines.

The edge BMS connects to the **cloud** via an **Over-The-Air (OTA)** communication module (using cellular or WiFi data links). The OTA link periodically uploads summary data (e.g., segmented features or aggregated health indicators) to the cloud and retrieves model updates or calibration parameters. We note that cloud communication introduces higher latency (on the order of seconds, varying with network conditions). Therefore, **safety-critical functions remain on the edge** – the PMU uses local algorithms for immediate control – while the cloud-based digital twin performs heavy computations and long-term analytics asynchronously. We will clarify these real-time performance considerations and the division of tasks between edge and cloud in the following sections.

Recent work by Halder [170] offers the most comprehensive review to date of security requirements for automotive OTA pipelines and provides a convenient checklist that we incorporate verbatim into our design. Their survey emphasises that a production-grade OTA service must (i) guarantee authenticity and integrity of every update through multi-signature code signing (e.g., Uptane’s ECU- and primary-level signatures), (ii) enforce version-control metadata to prevent rollback attacks, (iii) use end-to-end encryption during

distribution to defeat man-in-the-middle tampering, (iv) store keys in a hardware root of trust, and (v) include a post-installation attestation step that verifies successful flash and reboot before the old image is discarded. By aligning our dual-model digital-twin update flow with these five principles, we ensure that model packages received from the cloud cannot be spoofed or replayed, that corrupted downloads are detected before activation, and that any failure automatically triggers a re-signalled recovery using the last-known-good model. In addition, the same survey notes that OTA processes should log update provenance for ISO/SAE 21434 audits and support fleet-level rollback in case a newly deployed algorithm behaves unexpectedly—capabilities that our cloud back-end already exposes via its versioned model registry and secure revocation channel. Together, these safeguards elevate the proposed OTA path from a convenience feature to a security-certifiable subsystem fully in line with current best practice.

This hybrid design addresses the compute-bandwidth trade-off: only essential data features are transmitted to the cloud (reducing bandwidth use), and only refined model updates are sent back, whereas each layer performs the tasks it is best suited for (detailed below). For a discussion of general edge–cloud digital twin advantages. At a high level, the edge (on-board) tier is responsible for real-time monitoring and initial data handling, whereas the cloud tier handles batch analytics and global model management. Concretely, the tasks are divided as follows:

- Edge (BMS Pack Unit) – Runs the predictive model in real-time to estimate SoC/-SoE continuously during vehicle operation. It also performs on-line data compression and filtering (e.g., segmenting drive cycles) to prepare telemetry for upload. The edge can flag local anomalies (e.g., sudden voltage drops) immediately for safety. Critically, because inference is on-board, latency is minimal and the vehicle remains self-sufficient even with intermittent connectivity.
- Cloud – Aggregates incoming data from the fleet or individual vehicle, performing computationally intensive tasks such as retraining machine-learning models and long-horizon analytics. The cloud stores historical data and executes the battery state estimation algorithms in batch mode to update the digital twin’s models. It handles SoH analysis (which involves long-term trends) and other advanced computations that would be infeasible on limited BMS hardware. The cloud also manages OTA updates: after retraining, it sends updated model parameters or software down to vehicles.

Through this division of labor, the architecture achieves what neither a pure edge nor pure cloud system alone can. A cloud-only BMS would require streaming high-frequency sensor data continuously – incurring high bandwidth and risking delays or data loss if connectivity fails. Conversely, an edge-only approach (with no cloud support) would struggle to store and process the volume of lifetime data needed for accurate aging predictions, and could not easily benefit from fleet-level learning. Our hybrid solution uses the edge for what it does best (immediate, local computation) and the cloud for




Figure 4.2: The digital twin is trained initially in a cloud environment (step 1 using historical data from batteries of the same chemistry or from the same production line if available). After this initial training, the SoC/SoE model is transferred to the PMU (step 2), where it runs in real time. During vehicle operation, the PMU collects telemetry from the module and cell levels, which is periodically uploaded to the cloud in batches or via incremental data streaming (step 3). This newly accumulated data is merged with existing datasets in the cloud, triggering a model retraining step (step 4) according to a time-based or event-based policy. If the updated model surpasses the performance of the currently deployed version on test subsets, it is compressed and transmitted back to the PMU, ensuring that on-board predictions remain accurate despite changes in battery health and usage conditions.

what it excels at (big data storage and retraining), establishing a synergy between the two. Prior works have begun to recognize this edge–cloud collaboration trend, but our architecture implements it in a battery digital twin context with specific novel mechanisms (detailed in the following sections). Notably, real-time SoC/SoE estimation is kept on the vehicle (ensuring instant availability to the driver and BMS) while heavy “number-crunching” tasks like model training and multi-vehicle data aggregation happen in the cloud asynchronously.

**Communication and Data Management:** A secure IoT telemetry channel (e.g., cellular or Wi-Fi) connects the BMS’s pack unit to the cloud. To minimize bandwidth usage, the edge performs local preprocessing: rather than sending raw 10 Hz measurements continuously, it can transmit **summarized cycle data, features extraction or even using federated learning by sending model parameters**. This OTA data compression reduces communication load while preserving essential information for model updates. As illustrated in Fig 4.2, the vehicle periodically uploads batches of sensor data (or extracted features) to the cloud, such as at the end of each trip or charging session. Techniques like buffering and delta compression are employed so that the cloud receives high-value data (e.g., significant charge–discharge events) with minimal redundancy. In turn, when the cloud issues updates (new model parameters), the payload is lightweight (a few kilobytes of model coefficients) relative to raw data. This design ensures the architecture remains scalable even as the number of connected vehicles grows, since per-vehicle data transfers are optimized. A single cloud instance can support many edge devices by processing their intermittent uploads, rather than needing constant high-bandwidth streams.

Finally, the architecture’s edge–cloud synergy directly tackles compute-bandwidth trade-offs: time-critical computations (SoC inference) incur zero network delay (running locally), and large-scale computations (retraining) incur zero vehicle CPU load (running remotely). This is a marked improvement over monolithic approaches, giving our system both the responsiveness of an embedded BMS and the adaptability of a cloud platform. In summary, the dual-layer architecture provides a **scalable, low-latency, and bandwidth-efficient** framework for the battery digital twin, setting the stage for continuous learning and adaptation which we discuss next.

## 4.2 Dynamic Model Update Mechanism

A key novelty of our architecture is its support for continuous, dynamic model updates through a hybrid trigger mechanism. Instead of using a static battery model or updating only at infrequent service intervals, the digital twin is designed to evolve alongside the battery’s state. The edge–cloud split plays an important role here: the cloud can retrain models in the background and then deploy updates OTA to edge devices without interrupting the vehicle’s operation. We implement a dual-trigger update policy that combines time-based scheduling with event-driven updates:

- **Periodic Updates:** The system retrains the SoC/SoE estimation model at regular

intervals (e.g., daily, weekly or bi-weekly). These scheduled updates ensure that even gradual, subtle shifts in battery characteristics (due to aging or seasonal temperature changes) are captured over time. For example, the cloud might perform a batch update every week using all data accumulated from that period. This routine ensures the model doesn't drift far from reality even during "normal" operation. It effectively resets the model's accuracy on a rolling basis, much like periodic calibration.

- **Event-Driven Updates:** In addition to the fixed schedule, the architecture monitors for significant anomalies or threshold breaches that warrant an immediate model update. One important trigger is a drop in SoH beyond a predefined threshold. For instance, if the battery's estimated SoH decreases by more than, say,  $\Delta\text{SoH} = 5\%$  since the last update, the cloud will initiate a retraining process right away (without waiting for the weekly cycle). The rationale is that such a change likely indicates the battery has entered a new degradation regime, so the SoC/SoE model (which was trained under earlier conditions) needs recalibration. Other event triggers can include detection of anomalous usage patterns (e.g., an unusually high self-discharge event, or a temperature excursion) or large errors between the model's prediction and measured values. These events are detected by the edge logic in real-time (since the BMS is continuously comparing measured vs. predicted SoC, etc.) and flagged to the cloud for analysis. If the cloud confirms an event merits retraining (for example, an error threshold exceeded), it proceeds with an out-of-cycle update.

Whenever a retraining is triggered – whether by schedule or event – the cloud backend will use the latest aggregated data to update the machine-learning models that constitute the digital twin. In our implementation, the SoC/SoE estimator is a LightGBM regression model (chosen for its efficiency and accuracy on battery data). The cloud maintains a pipeline to automatically retrain this model with new data segments as they arrive. The specific retraining procedure, including feature engineering and model tuning, is described in Section 5.6. For the architecture discussion here, it suffices to note that the cloud can refresh the model relatively quickly (e.g., within a few minutes on a GPU-enabled server for the volume of data from a week of driving). Before deployment, the updated model undergoes internal validation on the cloud using a hold-out portion of recent data to ensure it performs better than the previous version and does not introduce any instability. Once validated, the new model (or model parameters) is packaged and sent to the edge device OTA.

This feedback loop, facilitated by OTA updates, assures alignment between the evolving state of the physical battery and the virtual clone in the digital twin. It also sets the stage for advanced optimizations such as fleet-level analytics: if patterns of abnormal degradation are discovered in the cloud among multiple vehicles, this knowledge can be pushed to each PMU, prompting proactive checks or revised usage guidelines [171]. The update deployment to the edge is designed to be asynchronous and seamless. Because the edge continues running the old model during the cloud's retraining process, there is

no interruption in SoC/SoE estimation service. When the cloud pushes the new model, the BMS’s pack unit (PMU) receives it via the secure OTA channel. The system then hot-swaps to the new model during an appropriate moment (for instance, when the vehicle is in an idle state or the battery is at rest, to avoid transients). We implement a fallback logic: if for any reason the new model fails a simple sanity check on the edge (e.g., producing out-of-range values), the system can revert to the previous model. In practice, such a rollback has not been needed in our tests, as the cloud validation ensures reliability. This mechanism of continuous asynchronous updates allows the digital twin to remain tightly synchronized with the physical battery over its entire lifecycle. Importantly, the architecture separates the concerns of real-time control vs. model training. The edge focuses on using the latest validated model for inference, while the cloud focuses on improving that model. By decoupling these, we avoid heavy computation on the vehicle and we mitigate safety risks (since experimental training code runs in the cloud, not on the vehicle). The result is an adaptive system that can learn from new data in the field and update itself regularly – a capability that is crucial given that battery characteristics evolve in non-linear ways over months and years. Our approach contrasts with prior BMS schemes where models might be fixed after deployment or only updated during manual maintenance; here the process is automated and continuous.

In summary, the dynamic update mechanism ensures long-term accuracy of the digital twin. The combination of periodic (time-triggered) and anomaly-based (event-triggered) updates guarantees that both slow drifts and sudden changes in battery behavior are promptly addressed. This is a key novel contribution of our architecture – it effectively enables a form of online learning in the battery management domain. The retraining uses a growing dataset with a sliding window, and techniques like incremental learning are applied to efficiently update the LightGBM model. From an architectural viewpoint, the takeaway is that our cloud–edge design makes such continuous learning feasible: the cloud has the computational heft to retrain models often, and the edge has the connectivity and flexibility to accept new models on the fly. Together, these features lay the groundwork for a self-updating BMS digital twin that remains accurate from the battery’s early life through end-of-life.

### 4.3 Integration with BMS Hierarchy and Fault Tolerance

Another central aspect of the architecture is its deep integration with the existing BMS hardware and software. Unlike a standalone cloud platform that operates in parallel to the BMS, our digital twin is embedded within the BMS’s operational loop. In practice, this means the edge component of the digital twin runs on the BMS Pack Management Unit (PMU), interfacing directly with the BMS’s data channels and control interface. By design, this leverages the hierarchical structure of the BMS (Cell → Module → Pack levels) without duplicating sensor infrastructure. As described in Section 2.4, a typical EV BMS is organized in layers: multiple cell monitoring units (CMUs) feed into a module

management unit (MMU), and in turn the MMUs feed into a central pack management unit that oversees the entire battery pack. Our architecture connects at the apex of this hierarchy – the PMU – which has access to the aggregated measurements (voltage, current, temperature of all modules) and also has communication links to the vehicle’s ECU and telematics unit. Embedding the digital twin here ensures it receives all necessary sensor data in real-time (since the PMU already collects it for basic BMS functions), and it can also issue high-level outputs (like refined SoC estimates) directly to the vehicle systems.

**Data Flow within BMS:** The PMU runs the edge algorithm (the trained SoC/SoE model) in parallel with the BMS’s conventional algorithms. For example, while the BMS might have a simpler coulomb-counting or Kalman filter for SoC as a redundancy, the digital twin model provides an enhanced SoC estimate that can be used for more accurate range prediction or displayed to the driver. The PMU takes cell voltage and current readings aggregated up from the CMUs/MMUs and feeds them into the digital twin model at the same rate as BMS computations (typically 1–10 Hz). Because the digital twin is co-located on the PMU, there is negligible overhead in obtaining this data – it doesn’t need an additional ADC or sensors. The computed SoC/SoE is then available to any higher-level vehicle controller (e.g., energy management system) via the same CAN bus or vehicle network that the BMS uses, ensuring the digital twin’s outputs are immediately actionable in the vehicle context. Meanwhile, the PMU also packages a subset of the raw data plus the digital twin’s outputs into a telemetry packet for cloud upload (as described earlier). In essence, the PMU serves as a gateway between the battery pack and the cloud: it speaks “battery language” on one side (reading cell data, ensuring safety thresholds) and “IoT language” on the other (communicating with the cloud server). This integration means the digital twin architecture does not require an extra device or separate communication channel – it piggybacks on the BMS hardware, which is a practical advantage for deployment.

**Fault Tolerance and Fallback Operation:** Integrating at the BMS level also provides robustness in cases of connectivity loss. If the cloud connection is temporarily unavailable (e.g., the vehicle is out of network range), the digital twin’s core functionality is unaffected. The edge model running on the PMU continues to provide state estimates to the BMS and vehicle in real-time, drawing only on local sensor data. This ensures that safety-critical operations are never cloud-dependent – a deliberate design choice. In a worst-case scenario of prolonged disconnection, the system will simply defer its periodic data uploads; the PMU can buffer data locally (storing a certain number of drive segments in flash memory). Once connectivity is restored, the buffered data is uploaded to the cloud, and any missed model updates can be caught up. During the disconnected period, the model on the edge remains the last received version – which may gradually age, but because our update scheme is conservative (the model is valid for at least the typical update interval), there is no immediate risk of large error. This fail-safe behavior contrasts with architectures that rely on cloud computation for every estimate (those would fail to produce results if offline). Our design is inherently resilient: it degrades gracefully (no updates, but still functioning) rather than failing. The use of OTA technology and an intelligent sync protocol facilitates

this resilience, ensuring that when a connection is present the data exchange is seamless and up-to-date, and when absent, the system simply stores data for later sync. As noted in [166], OTA connectivity solutions are able to overcome intermittent internet issues, allowing the digital twin to remain as synchronized as possible with the battery’s status.

Moreover, because the digital twin is embedded in the BMS, it can directly utilize the BMS’s existing fault-detection mechanisms. For instance, if the BMS detects an out-of-range temperature in a module (something that normally triggers a safety response), this information can be fed into the digital twin’s logic as well – perhaps to exclude anomalous data from model updates or to inform the cloud that a hardware issue occurred. The architecture is thus aware of BMS faults and can incorporate that context (e.g., flagging that data during a certain period might be unreliable due to a sensor fault). Conversely, if our digital twin model detects a discrepancy (like the model predicts a voltage that significantly differs from measured), it could signal the BMS to perform a diagnostic. This kind of cross-talk exemplifies the novel integration: the digital twin is not an external passive observer but an active part of the BMS ecosystem.

To summarize, by interfacing at the PMU level, the architecture ensures tight coupling with battery management operations while preserving vehicle safety and reliability. Even if the cloud is unreachable, the vehicle retains full BMS functionality with the last-known good model – fulfilling all basic monitoring needs. The cloud is essentially a bonus layer that refines and improves performance when available, but not a single point of failure. This integration and fail-safe design choice is a significant improvement over naive IoT implementations of battery twins, and it demonstrates our focus on practical deployability in real EV systems.

## 4.4 Robust Handling of Partial Charge–Discharge Cycles

Real-world EV usage rarely consists of neat full charge/discharge cycles; instead, drivers often perform partial charges, short trips, regenerative braking bursts, etc. This presents a challenge for traditional battery models which assume complete cycles for accurate state estimation (for example, many SoH estimation methods require observing a full discharge). A novel contribution of our architecture is partial-cycle robustness – the ability to make use of incomplete cycles and irregular usage data to keep the digital twin accurate. At the architectural level, this is enabled by a combination of edge-side data segmentation and cloud-side data aggregation.

On the edge, the PMU (pack unit) software segments the battery data stream into logical “micro-cycles” or segments in real time. Each segment could correspond to a usage episode such as a single drive trip, a charging session, or even a portion of a trip between two rest periods. The criteria for segmenting are defined to capture meaningful changes in state: for instance, the system might start a new segment whenever the vehicle is turned on, or every time the direction of battery power flow switches from charge to discharge. This dynamic segmentation is flexible and does not require the segment to start

at 100% SoC or end at 0% SoC – any partial delta in SoC is treated as valuable data. Our algorithm employs a dynamic segmentation strategy that adjusts to the current battery state and usage pattern. Essentially, the edge labels data with segment IDs such that even if a cycle isn't completed, the data up to that point is preserved as a self-contained sequence. For example, if a driver uses 20% of the battery in the morning and 30% in the afternoon (never reaching full discharge), the system treats these as two segments (with  $\Delta\text{SoC}$  of 20% and 30%, respectively) rather than ignoring them. All segment data (voltage, current, SoC change, etc.) is logged and eventually sent to the cloud, tagged appropriately.

The cloud layer is responsible for aggregating these partial segments into a form usable for model training. One straightforward approach is concatenation or stitching of sequential segments that together form a quasi-complete cycle. For instance, if after a few trips the battery has effectively gone from 100% to 20% in several steps, the cloud can aggregate those segments to represent a full discharge profile in pieces. More generally, our model training does not actually require full cycles as input; it can use each segment as an independent sample as long as the state variables (SoC, SoE) are tracked across segments. In fact, one innovation in our methodology is that, we deliberately train on partial cycles by slicing even full cycles into segments, to improve the model's ability to handle real driving patterns. By doing so, the model learns to predict state changes over arbitrary SoC ranges, not just from 100% to 0%. This significantly enhances prediction accuracy in scenarios where full cycles are absent. From an architecture perspective, the implication is that the data pipeline and storage have to accommodate many short segments rather than a few long cycles. Our cloud database is organized accordingly: each data record is a segment with start/end SoC and the time-series within it. The update triggers (as discussed earlier) can operate on the basis of segments as well – e.g., the weekly retraining uses all segments collected in that week.

Because the edge already segments the data, the cloud's role in aggregation is made easier: it knows how segments relate (sequence, or whether a segment ends in a full charge which might reset certain state references). The cloud can also combine information from partial segments to estimate SoH changes. For example, even if no single segment spans a full discharge, the cloud can observe the gradual decline of maximum charge over a series of partial discharges with intermittent full charges (or using reference discharge tests that occur occasionally) to update the SoH model. Our architecture thus does not break down in the common scenario of irregular cycling – it is explicitly designed to embrace it. This is a differentiator from existing systems that often assume consistent cycling for calibration. In our design, every bit of data, even from partial usage, feeds the digital twin's knowledge.

To illustrate, consider an EV that, over a day, goes from 80% SoC in the morning to 50%, charges back to 70% at noon, then goes to 30% by evening. Traditional coulomb-counting SoH estimation might have trouble, because no full discharge occurred. In our architecture, the edge would create segments: [80→50]%, [50→70]% (charge segment), [70→30]%. The cloud, upon receiving these, can interpret that overall the battery went

roughly 80→30% with a charge in between, and using our model (which handles partials) it can still estimate the aging effect. By aggregating partial segments, the cloud effectively reconstructs the battery’s usage profile over longer periods. The digital twin’s SoH model (running perhaps mainly on the cloud due to its long-term nature) can incorporate these partial cycles by, for example, summing up throughput or looking at end-of-day 30% as one reference point and a full charge that might happen later in the week as another reference. The key is that no data is wasted – even segments that start or end at arbitrary SoC contribute to the twin’s update. This approach was informed by our findings that full 0–100% cycles are rare in real driving, so any practical adaptive system must handle partial information.

Architecturally, enabling partial-cycle handling required careful design of the data interface between edge and cloud (to pass segment boundaries and state information) and of the model update logic (to allow incremental learning from partial sequences). We emphasize that the algorithms for training on segmented data are presented in Chapter 5 (specifically in Section 5.6 where dynamic segmentation and incremental learning are discussed). Within this Chapter, the contribution is the architecture’s capability to support those algorithms: e.g., the PMU firmware can delineate segments on the fly, and the cloud database and retraining software can consume those segments appropriately. In essence, the architecture treats the battery’s usage as a continuum of partial events, rather than needing neatly separated full cycles. This design choice greatly improves the robustness of the digital twin in real-world conditions.

In summary, the partial-cycle robustness of our digital twin architecture is achieved through edge segmentation and cloud aggregation. By capturing data at a granular segment level and intelligently combining it later, the system avoids the blind spots that plague static models in irregular operation. This ensures that even if the vehicle’s battery usage is highly variable, our twin can continuously learn and maintain accurate SoC/SoH estimates. It represents a forward-looking approach aligned with how EVs are actually used, and it underpins the continuous learning philosophy of the overall system.

## 4.5 Comparison with Baseline Architectures

To highlight the benefits of the proposed architecture, Table 4.1 provides a comparative summary against two baseline approaches: a pure edge (on-board only) architecture and a pure cloud architecture. Key attributes such as latency, communication requirements, update capability, scalability, and fault tolerance are contrasted. As shown, the hybrid edge–cloud digital twin combines the strengths of both approaches while mitigating their weaknesses:

- **Latency:** The edge–cloud architecture achieves ultra-low latency for critical estimations (comparable to pure edge) since the edge handles real-time tasks. A cloud-only system would suffer higher latency due to network transfer for each

estimation, whereas the edge-only and our hybrid both yield instantaneous local results.

- **Bandwidth Usage:** Pure cloud solutions require high continuous bandwidth (streaming all sensor data), whereas the hybrid approach significantly reduces data transfer through selective upload and compression (the edge-only obviously uses none, but at the cost of no remote data utilization). Our approach strikes a balance – moderate, periodic bandwidth use – which is scalable to large fleets without overwhelming networks.
- **Model Update Frequency:** An edge-only BMS might update its model rarely (or require bringing the vehicle in for service) due to limited computation on-board. A cloud-centric system can update models very frequently but must push results to vehicles. Our architecture enables frequent updates (e.g., weekly or triggered) thanks to cloud retraining, and promptly deploys them via OTA – providing a capability far beyond an edge-only design, and equal to cloud-based in timeliness.
- **Scalability:** In a cloud-only design, scaling to many vehicles means the cloud handles more data and computation centrally; this is feasible but can become costly or complex. In an edge-only design, each additional vehicle is independent (which is good for scaling, but each must carry the full computational burden). The hybrid approach offers double scalability: edge nodes independently handle real-time tasks (scaling naturally with the number of vehicles), and the cloud handles aggregate retraining (which can be scaled with cloud resources). It thus supports large deployments effectively – e.g., a fleet of thousands of EVs could each report summaries to a central cloud that updates a global model or individualized models.
- **Connectivity Dependence:** A cloud-only architecture fails to function when connectivity is lost (no state updates possible). The edge-only and hybrid architectures are robust to connectivity issues, since the edge always provides local functionality. However, the pure edge has nothing to gain from connectivity (and can't improve once deployed), whereas the hybrid can leverage connectivity when available but survive without it temporarily. In other words, our architecture is connectivity-augmented but not connectivity-required, offering the best of both worlds.
- **Computational Load:** Although not explicitly shown in the table, it's worth noting that the hybrid scheme keeps heavy computational load (training) off the vehicle, unlike an edge-only solution that might attempt on-board training (which is typically impractical for complex models). And compared to cloud-only, the hybrid keeps inference load off the cloud servers, distributing it to vehicles. This leads to efficient use of resources: vehicles use their embedded processors for light inference tasks, and cloud servers use their powerful CPUs/GPUs for periodic retraining.

Table 4.1: Qualitative Comparison of Battery Digital Twin Architectures  
(Proposed Hybrid vs. Edge-only vs. Cloud-only)

| Aspect                 | Edge-Only BMS                  | Cloud-Only BMS                                | Proposed Edge–Cloud Hybrid                       |
|------------------------|--------------------------------|-----------------------------------------------|--------------------------------------------------|
| Real-time Latency      | Low (on-board)                 | High (network lag)                            | Low (on-board inference)                         |
| Bandwidth Demand       | None (local data)              | High (constant data upload)                   | Moderate (periodic compressed uploads)           |
| Update Capability      | Low (manual, infrequent)       | High (continuous updates)                     | High (cloud retraining + OTA deployment)         |
| Scalability (Vehicles) | High (independent per vehicle) | Moderate–High (centralized cloud computation) | High (edge inference, scalable cloud training)   |
| Connectivity Reliance  | None (fully stand-alone)       | Strong (requires continuous connectivity)     | Tolerant (offline-capable; syncs when connected) |
| Fault Tolerance        | High (no external dependency)  | Low (no function if offline)                  | High (degrades gracefully when offline)          |
| Data Aggregation       | None (no fleet-wide insights)  | High (comprehensive fleet data)               | High (aggregated in cloud across vehicles)       |
| Maintenance Overhead   | High (manual vehicle updates)  | Moderate (central updates, ongoing cloud ops) | Low (centralized OTA, automated monitoring)      |

As Table 4.1 suggests, the proposed architecture excels in providing a balanced solution. It retains the low-latency and autonomy of an edge BMS, crucial for safety and user experience, and simultaneously harnesses the computational power and learning capability of the cloud, which is essential for long-term accuracy and adaptation. No single-layer approach can achieve this blend: an edge-only BMS would stagnate in performance over time (lacking global data insights), while a cloud-only system would be impractical in the field due to connectivity and latency constraints.

## 4.6 Comparison with Industry Solutions:

It is useful to contrast our proposed approach with current **industry-grade BMS solutions**. Leading electric vehicle manufacturers and suppliers have begun integrating digital twin-like analytics into their battery management. For instance, Tesla is reported to create a cloud-based “digital twin” for each vehicle’s battery pack, continuously monitoring its performance and health in real-time [172]. This fleet-level data collection enables Tesla to predict potential issues (like cell imbalances or failures) and optimize charging strategies through over-the-air updates. Our work shares a similar spirit of leveraging data for enhanced estimation; however, while companies like Tesla largely keep their algorithms proprietary, our thesis provides a **transparent** machine-learning framework that could be adopted by OEMs or third-party BMS developers.

In the automotive supplier domain, companies such as AVL and Siemens are developing advanced BMS simulation and development tools that align with the digital twin concept. AVL, for instance, offers battery simulation platforms sometimes referred to as “virtual twins,” which are used to predict battery aging and optimize battery control strategies during the development phase (e.g., using high-fidelity electrochemical models to simulate how the battery will behave over years of use) [173]. These tools underscore the importance of combining model-based understanding (from physics-based simulations) with data-driven updates, as they allow engineers to test how updating certain control parameters might extend battery life. Our architecture embodies a similar idea: we start with models that encapsulate battery behavior (partly informed by domain knowledge and initial lab data) and then continuously refine them with real-world data. The novelty of our approach in comparison is that we implement this adaptivity online in the field, not just as an offline development tool. We effectively bring a development-phase concept (virtual battery models that can be adjusted as conditions change) into the operational phase of the battery’s life via the cloud-connected framework.

It’s also worth noting that many production BMS today still rely on relatively traditional estimation techniques, such as coulomb counting combined with occasional open-circuit voltage checks and Kalman filters for error correction [174]. These methods require careful calibration for each battery type and usually assume fixed model parameters that are adjusted only infrequently. For example, a typical EV might use an extended

Kalman filter where the cell's capacity and resistance are set based on initial characterization and perhaps updated via lookup tables as the battery ages. Such approaches have been proven in industry for reliability, but they may struggle to maintain accuracy under unusual conditions or as the battery deviates from its nominal behavior. In contrast, our data-driven twin can automatically recalibrate itself and even discover new correlations (like the effect of a particular driving pattern on SoC estimation) without explicit re-programming. This reduces the reliance on manual calibration and can accommodate battery aging and heterogeneity more gracefully. One trade-off is that our approach requires an infrastructure for data collection and analysis (which our architecture provides), whereas a traditional BMS can be entirely self-contained. However, as evidenced by moves from companies like Tesla, the industry is indeed adding connectivity and cloud analysis to battery management to gain a competitive edge.

In summary, our digital twin framework is aligned with cutting-edge industry trends (such as Tesla's fleet learning BMS and AVL's virtual battery simulations) but also contributes unique features that are not yet standard in production systems. Two standout features are: (a) explicit handling of partial usage data in the modeling process, and (b) a formalized edge–cloud architecture for continuous online learning. To our knowledge, these aspects are not fully realized in current commercial BMS offerings, which opens an opportunity for technology transfer and further development. By demonstrating this framework in a research context, we provide a foundation that automakers or BMS suppliers could build upon to enhance their next-generation products. This could mean, for instance, that future EVs incorporate more sophisticated cloud-connected battery analytics that inform both the user and the manufacturer of the battery's condition and optimal usage, much like our system does.

Overall, the innovations in our architecture — the hybrid edge–cloud synergy, dynamic retraining, integration into BMS, and partial cycle robustness — push the envelope of what a battery management system can do. In the next chapter, we transition from the architectural and system-level perspective to the methodological perspective, detailing the machine learning models and algorithms that we employ for SoC, SoE, and SoH estimation within this framework. The combination of the architecture (this chapter) and the modeling methodology (Chapter 7) forms a novel battery digital twin framework that, as demonstrated by our results, achieves continuous improvement of EV battery state estimation in-field. This represents a step towards smarter and more adaptive battery management in future electric vehicles.



# Chapter 5

## Methodology

### 5.1 Overview

The methodology presented in this chapter outlines the systematic development of our data-driven digital twin for EV battery management. Building on the architectural concepts described in Chapter 4, this methodology integrates both model-driven and data-driven approaches to address the key challenges of SoC, SoE, and SoH estimation under real-world operating conditions. In particular presented pipeline is designed to handle:

1. **Realistic driving conditions** with predominantly partial charge/discharge cycles,
2. **Variable operating conditions** (ambient temperature, load current, rest durations),
3. **Long-term battery aging** that necessitates periodic model updates,
4. **Scalable approach** to data collection and feature engineering, enabling both real datasets and simulation-augmented data.

In what follows, the methodology is broken down into major phases: data gathering (Section 5.2) with data synchronization and cleaning, partial-cycle segmentation (Section 5.3), advanced feature extraction (Section 5.4), and model training + updating (Section 5.6). Finally, we discuss the implementation of the methodology in Section 5.8.

### 5.2 Data Collection and Preprocessing

A robust methodology for EV battery management hinges on the quality and diversity of the datasets used to develop and validate estimation models. In our framework, both laboratory-based (*model-driven*) and in-operation (*real-world*) data sources are considered. Laboratory data typically consist of carefully controlled charge/discharge cycles obtained




Figure 5.1: Workflow of the proposed edge–cloud digital-twin pipeline. Solid arrows represent the real-time edge path; dashed arrows depict asynchronous edge–cloud synchronisation.

using battery cyclers, often with known chemistries and environmental conditions. Real-world telemetry, on the other hand, comes from operational EVs, where driving behaviors and environmental factors vary widely. Each data record includes measurements such as terminal voltage, current, temperature, and time stamps. In some cases, the datasets are complemented by manufacturer specifications on initial capacity, nominal voltages, and an estimated end-of-life threshold. To establish consistency, the data streams are synchronized by aligning time stamps and discarding incomplete or corrupted entries. Where sampling rates differ, interpolation mechanisms are applied to ensure uniform temporal resolution. Outlier detection is performed using statistical tests (for instance, z-score or interquartile range methods) to remove sensor spikes or measurement artifacts. Whenever feasible, data are segmented by identifying transition points between charging, discharging, and rest phases. This segmentation later serves as a foundation for identifying partial discharge events.

**Real-World Deployment Considerations:** Beyond accuracy, deploying this digital twin approach in a real BMS entails managing practical challenges: - **Sensor Noise & Reliability:** In real vehicles, measurements of voltage, current, and temperature can be noisy or occasionally erroneous. Sensor biases (like current sensor offset drift over time)

and noise could degrade estimation accuracy. Our current models assume relatively well-calibrated data; however, in practice, filtering techniques or robust model designs (e.g., using observers or outlier rejection) are needed. Prior studies have shown that measurement noise (especially non-Gaussian noise or drift) can significantly impact SOC/SOH estimation if unaccounted for:contentReference[oaicite:35]index=35:contentReference[oaicite:36]index=36. We will note that integrating a noise-robust estimation method or preprocessing (such as a Kalman filter front-end for denoising) could enhance our system’s reliability. - **Communication Delays:** In the edge-cloud split architecture, latency in data transmission can be a concern. If the cloud is used for control-critical predictions, network delays could pose safety risks. In our design, we mitigate this by keeping instantaneous control decisions on the edge (the vehicle will not depend on a cloud message to, say, shut down the pack in an emergency). The cloud communication (via LTE/5G) is used for slower timescale updates (seconds or minutes). We will add a clarification: even if there is a network dropout, the on-board model remains functional using the last synchronized parameters. This addresses potential **real-world deployment risks** due to connectivity loss or delays – a crucial consideration for any cloud-reliant BMS. - **Environmental Factors:** Conditions like temperature extremes, vibration, and humidity can affect both sensor readings and battery behavior. Particularly, battery performance varies with temperature (internal resistance rises in cold, capacity fades faster in heat). Our current models do include temperature as an input feature, but we assume training and testing within a moderate range. In deployment, one must ensure the model is either expanded to handle the full automotive temperature range (-20°C to +50°C, for example) or is complemented by rules that handle out-of-bound conditions. We will mention that an avenue for future work is to explicitly model temperature effects or train on data covering a wide temperature span, to ensure **temperature sensitivity** is accounted for.

By discussing these points, we reinforce that while our approach is promising, careful engineering is needed for **real-world deployment**. We will incorporate these considerations in Chapter 6 (Discussion), emphasizing that sensor noise, latency, and environment variability are being acknowledged and can be managed with known engineering solutions:contentReference[oaicite:37]index=37.

### 5.2.1 Datasets Employed

A variety of experimental datasets were used throughout this thesis to develop and validate the digital twin BMS models. These datasets are presented in chronological order of use (aligning with the progressive studies in this research), highlighting their key characteristics such as sampling rates, sensor modalities, state-of-charge/energy/health coverage, duration, number of cells, and any special usage protocols (partial cycling, fast charging, temperature conditions, etc.). By combining insights from multiple data sources, we ensured that the methodology is robust across different battery chemistries and operating scenarios. Below, each dataset is described in turn:

- **NASA Randomized Battery Usage Dataset:** The initial case study utilized a public lithium-ion battery aging dataset from NASA's Prognostics Center of Excellence. This dataset consists of 28  $\text{LiCoO}_2$  18650 cells (nominal capacity 2.1Ah, rated 4.2V) that were cycled to various depths under a random-walk usage regime. In each cycle, a cell was charged fully to 4.2V and then discharged down to 3.2V using a randomized sequence of current levels. The discharge/charge current magnitudes ranged roughly from 0.5A up to 4–5A, and the durations of charge or discharge steps varied unpredictably. This random walk profile effectively simulates irregular real-world EV usage patterns (with pauses, accelerations, etc.) and results in cycle lengths that are non-uniform. Periodically (approximately every 50 cycles in certain subgroups), a full reference charge/discharge cycle was performed to baseline the cell's capacity and SoH. The dataset logs voltage, current, temperature, and timestamp at 1Hz throughout every charge and discharge event. Cells were tested under different ambient temperature conditions as well (to simulate seasonal variations), although all followed the general random-walk profile. The SoC span in each random-walk cycle was essentially 100% to 0% (full charge to cutoff), but within each cycle the SoC moved in a non-monotonic way (numerous partial charges/discharges before reaching endpoints). The overall SoH coverage ranges from 100% (at the beginning of life) down to around 40%–70% SoH for the most aged cells (some cells were cycled until a significant capacity loss occurred). Notably, the 28 cells were divided into subgroups, each with different stress profiles (e.g. different random current patterns or reference test frequency). This rich diversity in usage scenarios makes the NASA dataset a valuable starting point for developing data-driven models, as it provides both standard reference cycles and irregular partial usage data.
- **Sandia National Laboratories (SNL) Battery Dataset:** In the subsequent phase of this work, a dataset from Sandia National Labs was incorporated to broaden the range of battery chemistries and cycling protocols. This dataset contains cells of three types: LFP ( $\text{LiFePO}_4$ ), NMC ( $\text{LiNiMnCoO}_2$ ), and NCA ( $\text{LiNiCoAlO}_2$ ). Each cell was aged under controlled laboratory conditions with a variety of cycling rates and depths. The SNL data logs include both high-resolution time-series measurements (voltage, current, temperature, energy) at 0.1Hz sampling, and per-cycle summary metrics such as the discharged capacity for that cycle and the cycle index. Cells were subjected to a matrix of operational scenarios spanning multiple discharge C-rates (0.5C, 1C, 2C, 3C), ambient temperatures (15°C, 25°C, 35°C), and Depth-of-Discharge (DoD) ranges (full 0–100% cycles, as well as partial windowed cycles like 20–80% and 40–60% DoD). Each test lasted for a large number of cycles (either 500 or 1,000 cycles in a given regime), after which a set of three full charge/discharge reference cycles at a standard rate (0.5C) was performed. Figure 5.2 presents a schematic overview of the typical cycling sequence employed in the Sandia National Laboratories (SNL) dataset. The diagram illustrates multiple

consecutive partial cycles—each conducted at a specific C-rate and temperature, and within a designated depth-of-discharge (DoD) window—followed by three reference cycles at a standard 0.5C rate. These reference cycles are full 0–100% charge/discharge cycles intended to accurately assess the remaining capacity (SoH) of the cell at that point in the aging process. These reference cycles act as synchronization points to measure the true remaining capacity (SoH) of the cell. The SNL dataset thus provides long-duration cycling data under semi-realistic but repeatable patterns, with periodic capacity checks. The SoH coverage for the SNL cells typically goes down to around 80% SoH in the 1000-cycle tests (some cells slightly lower), and the data set captures the influence of temperature and cycling intensity on degradation. For this thesis, we focused primarily on the NMC cells within the SNL dataset (31 cells) to align with the NMC chemistry of later datasets and to ensure a consistent basis for model development. The SNL data contributed insight into partial cycling behavior (20–80% DoD etc.) and multi-environment aging, complementing the NASA data.




Figure 5.2: Schematic illustration of the Sandia National Laboratories (SNL) battery cycling protocol. Partial cycles ( $PC^1, PC^2, \dots$ ) are interspersed with periodic reference cycles ( $RC^1, RC^2, RC^3$ ), which are full 0–100% charge/discharge cycles at a standard rate (0.5C) to check and log the cell’s capacity (SoH). This figure highlights how partial cycles operate at different C-rates, temperatures, and depth-of-discharge (DoD) windows, followed by the reference cycles for capacity checks.

- **Stanford UDDS EV Battery Dataset:** In the final stages of the research, a dataset reflecting actual EV driving usage was employed to validate the digital twin under real-world conditions. This is the Stanford Urban Dynamometer Driving Schedule (UDDS) Battery Dataset, which comprises 10 commercial 21700 NMC cells (INR21700–M50T) that were aged under real EV drive cycles over about 28 months.

The cells experienced repeated UDDS driving profile discharges – a standardized aggressive drive cycle – meaning the cells saw partial SoC swings roughly between 80% and 20% SoC per cycle (mimicking typical automotive use where the battery is not fully charged or fully drained in each trip). Charging was performed with a CC-CV regimen at various rates from moderate (C/4) up to 3C fast-charging, introducing realistic fast-charge stress. Periodically, the Stanford dataset also includes diagnostic capacity tests (full charge/discharge cycles or other characterization tests) to measure degradation and recalibrate SoH. The data collected encompasses voltage, current, temperature, and energy measurements at 1Hz during the diagnostic tests and at 10Hz during the high-resolution UDDS driving segments. This dataset is challenging due to various irregularities inherent in field data: some data gaps (missing samples for certain cells), occasional sensor errors (e.g. anomalously high temperature readings), and inconsistent resting periods between cycles. Nevertheless, it provides a critical testbed for the methodology under real driving patterns and environmental influences. The SoC in this dataset always remains in the mid-range (no full 0–100% swings except during diagnostics), and the SoH degradation observed is relatively mild – the worst cell retained 91% capacity after the test period. This reflects that, over 28 months, these cells did not age as drastically as the more accelerated lab tests (NASA or SNL). The Stanford data was primarily used in the integrated evaluation of the digital twin to demonstrate that the models trained on lab datasets can generalize to real-world usage, as well as to showcase the system’s ability to handle partial cycling and fast-charge scenarios with continuous model updates.

- **Simulated Battery Dataset (Model-Driven Generation):** In addition to the above experimental datasets, a model-driven dataset generation approach was explored to augment the data available for training. Using a high-fidelity equivalent-circuit battery model, we simulated numerous charge/discharge cycles under a wide variety of synthetic usage patterns. The motivation was to overcome the limited coverage of real datasets (which only contain certain profiles) by generating an essentially unlimited set of examples spanning different load profiles, ambient conditions, and aging scenarios. By feeding the circuit model with randomized or designed current profiles (mimicking diverse driving behaviors) and varying battery model parameters to emulate aging, we obtained additional labeled data for SoH prediction. The primary advantage of this approach is the diversity and volume of data: we can simulate rare or extreme scenarios not present in the empirical datasets (e.g., unusual partial cycling regimes or extended calendar aging periods) to improve model generalization. Another benefit is efficiency – simulation accelerates data collection, enabling evaluation of many more cycles in a short time compared to lab testing. In this thesis, simulated data as shown in figure 5.3 was used to supplement training in one of the studies, demonstrating that even a model-trained purely on

synthetic cycles can achieve reasonable accuracy in SoH estimation (within single-digit percentage error of the more complex physics-based model). The synthetic dataset's structure mirrored the real data (time-series of current, voltage, etc., per cycle, with known capacity as the label), but it offered control over factors like initial SoC, depth of discharge, and noise. While simulated data cannot capture every physical nuance (and was kept separate from the primary evaluation on real data), it played a role in understanding the trade-off between data quantity and model fidelity. The use of this generated dataset is described alongside the real data to highlight how data-driven models can be bootstrapped when measurements are scarce.




Figure 5.3: Simulated Battery Dataset (Model-Driven Generation). Using a high-fidelity equivalent-circuit battery model, synthetic charge/discharge cycles are generated under diverse simulated conditions (e.g., variable ambient temperature, battery voltage, and aging parameters). The resulting data includes time-series signals for Battery Current, Voltage, Temperature, as well as labels for SoH, SoC, and SoE. This simulation approach augments experimental datasets by providing unlimited and diverse examples, including rare or extreme scenarios. (See :contentReference[oaicite:2]index=2 for further details.)

Each of these datasets contributed different insights: the NASA data provided baseline full-cycle aging under stochastic loads; the SNL data introduced partial cycling and temperature effects in a controlled manner; the Stanford data tested the approach on real driving patterns with partial SoC range; and the simulated data allowed hypothesis testing beyond the limits of physical experiments. By leveraging this broad range of data, we aimed to ensure that the digital twin's models are not narrowly tailored to a single scenario, but rather are applicable to various battery chemistries, usage profiles (steady cycling vs. random vs. driving), partial-charge behaviors, fast-charge events, and differing SoH trajectories.

**Cross-Dataset Generalizability:** A key question is how robust our trained models are when applied to data outside their training distribution (e.g., a different vehicle or battery type). In our experiments, we partially addressed this by evaluating on multiple

datasets (Section 5.2.1). The results show that the model trained on one dataset maintained reasonably low error on the other, with only a modest degradation (e.g., SoC RMSE increased from 1.8% on the source dataset to 2.5% on the new dataset). This suggests some **generalizability** of the learned features and model – likely because our feature engineering captures fundamental patterns of battery behavior. We also kept hyperparameters (e.g., network architecture, learning rate) constant across these datasets. The fact that acceptable accuracy was achieved without dataset-specific re-tuning indicates the models are not over-fitted to one particular battery. We will explicitly mention this insight: our chosen hyperparameters proved **stable across datasets**, implying that the model could be deployed on new but similar batteries with minimal recalibration. Of course, if a target dataset differs greatly (e.g., a different chemistry or an extreme climate scenario), some re-tuning or transfer learning might be necessary. We will caution the reader that **model robustness without re-training** has limits, and ensuring wide generalizability may require incorporating more diverse training data or adaptive algorithms.

### 5.2.2 Data Preprocessing

Before training any models, all datasets underwent extensive preprocessing to clean the raw measurements, align the data with the needed labels (SoC/SoH), and shape the time-series into training examples. Common preprocessing steps included resampling, denoising/filtering, alignment of state labels, smoothing of degradation curves, and imputation or exclusion of missing data. Additionally, we defined how to break the continuous data streams into windows or cycles for analysis. By standardizing these steps across datasets, we ensured fair comparisons and a unified pipeline. Below, we outline the preprocessing procedures applied:

- **Sampling Rate Normalization:** The datasets came with different sampling frequencies (NASA at 1Hz, SNL at 0.1Hz, Stanford at up to 10Hz). For consistency and to reduce computational load, we resampled certain data to a common rate where needed. In particular, the Stanford UDDS data, which was recorded at a high 10Hz during driving segments, was down-sampled to 0.2Hz (one sample every 5 seconds). This aggressive down-sampling was done after confirming that the high-frequency dynamics (which are important for instantaneous power analysis but less so for capacity degradation estimation) could be safely filtered out for our purposes. Reducing Stanford’s data to 0.2Hz dramatically decreased data volume while still capturing the overall SoC evolution during each drive. The NASA 1Hz and SNL 0.1Hz data were left at their native rates in each case, as those rates were already relatively low; however, when combining features or comparing across datasets, we sometimes interpolated the 0.1Hz SNL data to 1Hz or vice-versa for uniformity. In all cases, resampling was done using appropriate interpolation or averaging (for down-sampling) to avoid aliasing. The result was that each time-series (regardless of source) was represented at a moderate frequency suitable for modeling long-term

trends rather than high-frequency transients.

- Data Cleaning and Noise Filtering:** We rigorously cleaned each dataset to remove erroneous measurements and outliers that could mislead the models. This included identifying sensor errors, inconsistent readings, or corrupted segments. For example, in the NASA dataset, a few cells showed non-monotonic jumps in their capacity measurements over successive reference cycles (i.e., the estimated SoH would increase from one cycle to the next, which is physically impossible). As figure 5.4 for one cell (NASA cell RW9) exhibited such anomalies in 9 out of its 79 reference cycles – these irregular cycles were excluded from the data to maintain a smooth degradation trend. Similarly, the Stanford dataset contained some clear sensor malfunctions: at times the temperature was recorded as exceeding 400°C, far beyond any realistic battery temperature. We handled these by either removing those data points or replacing them with a plausible value (e.g., the average temperature of that cycle). Replacing extreme outliers with a mean or median prevents a single bad reading from tainting an entire cycle’s data. We also excluded entire cells or large segments with excessive missing data. As a criterion, if a cell’s data had substantial gaps or errors that could not be reliably interpolated, that cell was dropped from modeling. For instance, from the NASA set, 4 out of 28 cells were discarded due to significant irregularities in their measurements. In the Stanford data, all 10 cells were initially considered, but a couple were later excluded in analysis due to missing large portions of records or severe irregularities. These exclusions ensured that the subsequent modeling was based on high-integrity data. Overall, the data cleaning steps (anomaly removal, outlier filtering, and cell selection) were essential to ensure the data reflected true battery behavior and monotonic aging as much as possible. Any cleaning action (like removing a cycle) was done cautiously to avoid inadvertently biasing the dataset; we applied them only in clear-cut cases of error.
- SoH, SoC, and SoE Label Computation and Alignment:** Unlike typical machine learning datasets, battery datasets often do not provide an explicit SoH label for every data point – the SoH (capacity) must be derived from periodic tests. In our preprocessing, we computed SoH values from the raw data at appropriate intervals, and aligned those with the time-series data. Specifically, for the NASA and SNL datasets, SoH was evaluated only during the reference discharge cycles that were interspersed in the test protocol. During each reference cycle, we performed **Coulomb Counting**: integrating the current over time during a full discharge to compute the total discharged capacity, which represents the cell’s usable capacity at that point in its life:

$$\text{SOH}; =; \frac{\text{Capacity}_{full}}{\text{Capacity}_{nominal}} \quad (5.1)$$

$$\text{SOC}(t); =; \text{SOC}(t_0); -; \frac{1}{3600, Q_{nom}} \int_{t_0}^t I(\tau), d\tau \quad (5.2)$$




Figure 5.4: SOH degradation profile extracted from the RW9 dataset. Arrows points where the monotonically decreasing nature of SOH is not reflected due to measurement errors.

$$\text{SOE}(t); =; \text{SOE}(t_0); -; \frac{1}{3600, Q_{\text{nom}}} \int_{t_0}^t V \times I(\tau), d\tau \quad (5.3)$$

where  $Q_{\text{nom}}$  is nominal cell capacity (Ah). Periodically, we may recalibrate  $\text{SOC}(t_0)$  when the battery is known to be near full or near empty (i.e., reference events). This step ensures each data record has a consistent tuple  $(t, I, V, T, \text{SOC})$ . If partial ground-truth capacity or SoH is available for certain segments, we record these as well for subsequent labeling.

This capacity was then normalized to the cell’s initial capacity (measured in the first reference cycle) to obtain a fractional SoH (100.0% at beginning of life, decreasing thereafter). The derived SoH values were stamped at the times of those reference cycles. For the Stanford dataset, which included periodic diagnostic cycles (capacity check tests), we similarly calculated the capacity from those tests to serve as the SoH at specific points in time. Between these reference or diagnostic points, the SoH was not directly measurable; we assumed the SoH remains constant during normal cycling between two reference tests (since degradation happens gradually and the small change within a few cycles is negligible compared to the total capacity). Thus, each continuous segment of usage between SoH measurement points was associated with the most recently obtained SoH value. We aligned the SoH label with the data by tagging each data sample (or cycle) with the SoH measured at the previous reference test. This way, every training example for the SoH estimation model had a ground-truth label derived from an actual capacity test. For SoC, we established labels on a much finer timescale: we treated the coulomb-counted SoC (from full charge = 100% downwards) as the ground truth instantaneous SoC during each

cycle. This required knowing the starting SoC of the cycle (usually 100% for a discharge, or whatever the initial value was in partial cycles) and then integrating current to track how SoC decreases over time. In practice, for cycles that start at full charge (like reference discharges), it's straightforward to set  $\text{SoC}(t=0) = 100\%$  and  $\text{SoC}(t=\text{end}) = 0\%$ ; for partial cycles, we used any given information (for Stanford, the cycle definition 80% to 20% means start SoC 80%, end 20% by design) or integrated until a known discharge capacity was reached. All SoC ground truth profiles were double-checked by cross-verifying against voltage curves (e.g., ensuring that a cell near 0% SoC corresponds to low terminal voltage near cutoff). By aligning SoC and SoH labels properly with the measurement timeline, we prepared a supervised learning table where inputs (voltage, current, etc. over time) could be mapped to the correct output states at each point.

Through these preprocessing steps, we transformed raw heterogeneous battery logs into a consistent set of training-ready data structures. Each dataset was cleaned, down-sampled to manageable size, segmented into meaningful chunks, and labeled with accurate SoC/SoH values. For example, after preprocessing, the NASA dataset yielded 24 cells' worth of cycles with clear SoH labels at each reference cycle (and those cycles' time-series data ready for feature extraction), and thousands of sub-cycle segments for SoC modeling. The SNL data yielded dozens of NMC cell trajectories that could be mined for aging features under different conditions, and Stanford's real-world data was distilled into regularized segments with known start/stop points and periodic SoH checks. This uniform treatment of data established the foundation for training the machine learning models, as described next.

### 5.2.3 Dataset Splitting for Robust Validation

To rigorously validate predictive models, the datasets were partitioned according to clearly defined, mutually exclusive subsets, ensuring representative training and realistic, unbiased evaluation conditions. Specifically:

- The NASA dataset was split into three subsets: training (12 cells), validation (7 cells), and testing (4 cells).
- The Sandia dataset was similarly partitioned, allocating 17 NMC cells for training, 7 cells for validation, and 6 cells for independent testing, including scenarios previously unseen during model training.
- For the Stanford UDDS dataset, a similar scheme was employed: diagnostic tests from 5 cells were designated for training, 3 cells provided validation data, and 2 cells were strictly reserved for final testing.

- This rigorous splitting scheme was essential for objectively assessing model performance, generalizability, and robustness across diverse operational scenarios and real-world-like conditions.

In summary, the comprehensive preprocessing steps described herein assured high-quality, consistent, and representative data inputs—crucially underpinning subsequent adaptive segmentation, sophisticated feature engineering, and robust predictive modeling processes. The meticulous integration of diverse experimental setups (NASA, Sandia, and Stanford datasets) allowed us to build an accurate, realistic, and practically relevant battery digital twin, deeply aligned with real-world EV battery management challenges.

### 5.3 Segmentation for Partial-Cycle Handling

Accurately modeling battery behavior for EVs requires datasets representative of realistic usage patterns. Unlike laboratory tests that typically feature complete charge–discharge cycles (from 100% down to 0% state-of-charge (SoC)), practical vehicle operation rarely adheres to such controlled conditions. Instead, EV batteries frequently undergo partial charge and discharge intervals—short driving trips, brief regenerative braking intervals, or small top-up charges—leading to significantly different battery degradation profiles. Consequently, models developed solely on idealized full-cycle data may underperform in real-world conditions, limiting the possibilities for continuous retraining and accurate predictive insights.

**Related Work on Data Segmentation:** Handling **partial discharge data** and irregular usage profiles is a known challenge in data-driven battery state estimation. Prior studies have proposed segmenting battery usage data to extract meaningful cycles or partial cycles. For instance, Hussain *et al.* [175] (2024) demonstrated that even using **partial** discharge data (instead of complete charge-discharge cycles) can be sufficient for accurate battery life prediction. By training deep learning models on segments of cycles, they achieved RUL prediction accuracy comparable to using full cycles, which underscores the value of partial-cycle data segmentation. Similarly, researchers have emphasized dividing driving data into smaller fragments corresponding to consistent operating regimes (steady cruising, acceleration, etc.), which can simplify modeling. Recent work by Chen *et al.* [176] (2025) proposes to **segment** real EV data (which is often fragmented due to irregular charging) and use those segments as the basis for feature extraction. They combine features from fragmented charge intervals and employ transfer learning across multiple battery datasets, achieving low estimation error (RMSE 0.01) even when models are applied to different datasets. These approaches highlight that segmentation helps align data to the assumptions of machine learning models and cope with partial discharge conditions, much like our **partial discharge-aware segmentation**.

method (Section 5.2.2) which identifies usable segments from real driving profiles for model training.

To bridge this gap, we developed a comprehensive segmentation strategy that transforms existing full-cycle datasets into a multitude of smaller partial-cycle segments (*pseudo-cycles*), each emulating real-world battery use scenarios. This method ensures the training dataset realistically mirrors typical battery operation patterns encountered by EVs, enabling models to be robust against a wide variety of partial discharge behaviors.

Our segmentation approach subdivides each complete discharge cycle recorded in the original dataset into non-overlapping intervals, determined by preselected increments of state-of-charge ( $\Delta\text{SoC}$ ). These intervals range from fine-grained partitions (e.g., 1%, 2%, and 5%) to broader segments (e.g., 10%, 15%, 20%, 25%, and 50%), and ultimately, a full-cycle segment of 100%. This multiscale segmentation generates extensive partial-cycle examples from a limited number of original cycles, significantly enriching the training dataset and greatly enhancing the representativeness of real-world operational data.




Figure 5.5: Illustration of the battery data segmentation strategy. A complete discharge cycle (left) is systematically partitioned into multiple non-overlapping segments with progressively smaller state-of-charge intervals (right). This approach enables multi-resolution modeling and closely replicates the partial usage cycles common in real-world EV operation.

To clarify, consider an example with a  $\Delta\text{SoC}$  segmentation of 25%. In this case, a full discharge cycle (100%–0%) yields four distinct segments: [100%–75%], [75%–50%], [50%–25%], and [25%–0%]. Similarly, applying finer resolutions such as a  $\Delta\text{SoC}$  of 5% results in twenty discrete segments, such as [100%–95%], [95%–90%], [90%–85%],

and continuing down to [5%–0%]. By adopting these incremental partitions, we obtain datasets at multiple resolutions that can be flexibly employed for training models to recognize and predict battery behavior under diverse operational conditions. The resulting comprehensive dataset provides a robust foundation for training both state-of-charge (SoC) and state-of-energy (SoE) estimation models, while also supporting state-of-health (SoH) predictions that rely on observing battery behavior at multiple timescales.

In practical terms, segmentation is guided by clear rules that ensure the meaningful identification of each segment:

1. **Operational Transition Rule:** A new segment is initiated whenever the battery shifts from a rest state to active charging or discharging, or directly transitions between charging and discharging activities. This ensures that each segment encapsulates coherent battery behavior within operational boundaries, such as a short trip or charging event.
2. **SoC Threshold Rule:** We systematically partition each discharge event by pre-defined  $\Delta\text{SoC}$  thresholds (1%, 2%, 5%, 8%, 10%, 15%, 20%, 25%, and 50%). This method consistently produces segments of known SoC intervals, thus creating representative partial discharge scenarios. The granular intervals (e.g., 1%, 2%, 5%) emulate frequent micro-usage scenarios like urban driving patterns, while larger intervals (e.g., 20%, 25%, 50%) simulate longer trips or substantial discharge events.
3. **Reference Cycle Rule:** Whenever a full discharge event occurs specifically for reference capacity testing, it is isolated as a distinct segment. This dedicated segmentation enables precise recalibration of battery health estimations, providing accurate SoH labels necessary for model validation and updates.

These clearly defined rules ensure consistent and reproducible segmentation across datasets, preserving alignment of critical battery parameters—current, voltage, temperature, and SoC—within each segmented interval. Each segment, denoted  $S_i$ , spans from a clearly defined start ( $t_{i,\text{start}}$ ) to end ( $t_{i,\text{end}}$ ) timestamp. Consequently, the segmented data readily integrates into downstream tasks such as feature extraction, modeling, and inference, ensuring the seamless continuity of the modeling pipeline.

The adaptive segmentation methodology has proven critical in enabling reliable retraining and model updating during real-world battery operation, significantly enhancing model robustness against partial-cycle variations and uncertainties. Rather than relying on infrequently encountered full-cycle data, our approach exploits abundant partial-cycle data reflecting realistic usage. This substantially improves predictive accuracy for SoC estimation under conditions in which the battery has not recently experienced complete cycles. Furthermore, our segmentation strategy ensures that test data involving arbitrary partial usage cycles are naturally accommodated by models trained on such flexible and realistic segmentation.

## 5.4 Feature Engineering

Battery degradation is driven by complex electrochemical processes that are influenced not only by instantaneous measurable signals (such as voltage, current, and temperature), but also by cumulative characteristics of usage patterns across entire charge or discharge cycles. Therefore, robust modeling of SoH requires transforming raw instantaneous sensor data into representative statistical and derived features that encapsulate usage history, operational stress, and internal aging mechanisms.

**Related Work on Feature Engineering:** Effective *feature extraction* from battery data can significantly improve state estimation accuracy. In literature, a variety of features have been explored beyond raw voltage and current. For example, incremental capacity (IC) analysis and differential voltage features are classic techniques to glean aging information; peaks and shifts in IC curves correlate with degradation mechanisms [177]. However, IC curve extraction requires controlled full cycles and high-resolution data, which are not always available in EV operation. To address real-world data constraints, researchers have turned to more readily obtainable features. Wang et al.[178] (2024) proposed a feature engineering-based SOH estimation using “downgraded” data to mimic real vehicle conditions. They extracted six key features from partial charge ranges – including *voltage-based features* (e.g., charge time in a certain voltage interval, mean and skewness of voltage during charging) – that remain informative even with sparse, noisy data. Other works have similarly looked at *statistical features* (mean, variance, kurtosis of voltage or current over a window) and *geometric features* (slope or area from V–I curves) to capture degradation trends:contentReference[oaicite:22]index=22. These features, when fed into machine learning models (from simple linear models to advanced ones like XGBoost or neural networks), have shown improved SOH/SOC estimation by embedding domain knowledge into the data-driven approach [119, 179].

Following the segmentation phase, each partial-cycle segment of the dataset is converted into a comprehensive feature vector, aggregating information across the entire segment. This approach yields data entries that are more reflective of battery aging phenomena than instantaneous measurements alone. Figure 5.6 (bottom right) visually demonstrates the feature extraction process, where each segment’s time-series data (current, voltage, temperature, SoC) is summarized into representative features.

The constructed dataset encompasses statistical descriptors of critical signals across each segmented discharge cycle, including mean, minimum, maximum, range, standard deviation, and variability of the current, temperature, and SoC. These statistical metrics offer a concise yet informative characterization of operational conditions throughout the discharge cycle.

Moreover, an essential derived feature, termed the **Discharge Rate**, was calculated for each segment. This feature quantitatively represents the speed at which the battery discharges its capacity during usage. The discharge rate is computed by taking the average slope of the state feature (e.g., SoC or energy, computed as voltage multiplied by current) with respect to time across each segment:




Figure 5.6: From initial dataset (top) to feature extraction (bottom). The initial dataset is reported as table and as plot of one discharge cycle (data from [180]). The discharge raw cycles are segmented based on the  $\Delta\text{SoC}$  (in the Figure we show a few samples with  $\Delta\text{SoC}=5\%$ ) and aggregated to train the SoC/E model (bottom left). Statistical and derived features are computed over each segment, forming the trained data for SoH (bottom right).

$$\text{Discharge Rate} = \frac{1}{n} \sum_{i=1}^n \frac{X_i - X_{i-1}}{t_i - t_{i-1}} \quad (5.4)$$

where  $X_i$  and  $X_{i-1}$  are consecutive values of the state feature at time stamps  $t_i$  and  $t_{i-1}$ , respectively, and  $n$  represents the total number of samples within the segment. This metric is directly correlated with battery internal resistance growth and capacity fade, as increased aging typically results in steeper SoC declines under identical loading conditions.

Beyond these statistical descriptors, a set of domain-informed and electrochemically meaningful features were extracted to further enrich the dataset. Leveraging known degradation mechanisms from battery aging literature, we constructed a comprehensive array of features, detailed in Table 5.1. These features capture diverse operational and degradation aspects, including cumulative throughput, depth of discharge (DoD), temperature exposure, internal resistance proxies, voltage curve characteristics, and aging-aware historical metrics. By incorporating such domain-specific features, the resulting dataset provides the machine learning models with deep insights into aging dynamics, significantly enhancing prediction accuracy for battery health estimation.

Table 5.1: Comprehensive Overview of Extracted Features for Battery Health Modeling

| Feature / Label                     | Description                                                                                                                                                                             |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Initial SoH                         | Baseline capacity of the battery, establishing a reference for quantifying subsequent degradation over battery lifespan.                                                                |
| $\Delta$ Time Computation           | Time interval between consecutive segments to capture calendar aging effects and cumulative temporal degradation.                                                                       |
| Current Characteristics             | Statistical metrics including maximum, minimum, mean, range, and standard deviation of current during each cycle, reflecting the stress induced by varying load conditions.             |
| SoC Insights                        | Statistical descriptors (mean, min, max, standard deviation) of SoC during each segment, offering insights into charge distribution and utilization patterns influencing battery aging. |
| SoC Deviation Metrics               | Quantification of SoC variability within charge/discharge events, highlighting impacts of non-uniform usage patterns on battery health.                                                 |
| Depth of Discharge (DoD) Variations | Metrics representing the depth of discharge within each cycle, directly correlated with capacity fade, particularly from deep discharge cycles.                                         |
| Temperature Features                | Statistical summary (mean, max, min, standard deviation) of both battery cell and ambient temperatures, correlating thermal exposure with accelerated aging.                            |
| Internal Resistance Proxy           | Features derived from voltage drops at onset of load or calculated differential voltage analyses, serving as indirect indicators of battery internal resistance increases.              |
| Cumulative Charge Throughput        | Total cumulative charge processed by the battery up to the current cycle, directly associated with cumulative aging and SoH reduction.                                                  |
| Discharge Rate                      | the rate of change of the voltage, current, temperature and SoC                                                                                                                         |
| $\Delta$ SoH Label                  | Observed change in battery SoH between consecutive segments, serving as a supervised label to identify and quantify degradation trends.                                                 |

The combination of statistical descriptors and physics-informed features provided comprehensive coverage of factors influencing battery aging. For instance, cumulative charge throughput and recent usage patterns were among the most informative features, strongly correlating with battery degradation rates. Additionally, temperature metrics and internal resistance proxies significantly improved model sensitivity to operating conditions, enabling accurate prediction of capacity fade under various scenarios.

Finally, to ensure completeness and reproducibility, Algorithm 1 illustrates the integrated data processing pipeline, highlighting how segmentation and feature extraction jointly yield two distinct datasets:

- $D_{\text{SoC/E}}$ : Contains segmented partial-cycle data explicitly tailored to train robust models for real-time SoC and SoE estimation.
- $D_{\text{SoH}}$ : Comprises aggregated feature-rich vectors associated with cycle-level SoH labels, enabling advanced predictive modeling of battery degradation.

---

**Algorithm 1** Integrated Pipeline for Segmentation and Feature Extraction

---

**Input:** Raw dataset  $D = \{(V_t, I_t, T_t, \text{SoC}_t, \text{SoH}_t)\}_{t=1}^N$ ; Set of discharge cycles  $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ ; Segmentation thresholds  $\Delta = \{\delta_1, \delta_2, \dots\}$ .

**Output:**  $D_{\text{SoC/E}}$  (segmented data for SoC/E model);  $F_{\text{SoH}}$  (feature set for SoH model)

```

1: Initialize  $D_{\text{SoC/E}} \leftarrow \emptyset$  and  $D_{\text{SoH}} \leftarrow \emptyset$ 
2: for cycle  $C_i \in \mathcal{C}$  do
3:   for  $\delta \in \Delta$  do
4:      $\{S_1, S_2, \dots, S_m\} \leftarrow \text{Segment}(C_i, \delta)$ 
5:     for segment  $S_j$  do
6:        $D_{\text{SoC/E}} \cup = S_j$ 
7:        $f_j \leftarrow \text{ExtractFeatures}(S_j)$ 
8:        $f_j.\text{SoH} \leftarrow \text{SoH}_i$ 
9:        $D_{\text{SoH}} \cup = f_j$ 
10:    end for
11:  end for
12: end for
13: return  $D_{\text{SoC/E}}, F_{\text{SoH}}$ 

```

---

Each identified cycle  $C_i$  is split via user-defined SoC thresholds into multiple partial segments  $S_j$ . All segments become part of  $D_{\text{SoC/E}}$ , preserving high-resolution partial discharge trajectories used to train or adapt SoC/SoE estimators. Meanwhile, a feature extractor condenses each  $S_j$  into a compact set of statistical, derived, and aging-aware parameters  $f_j$ , appended to  $D_{\text{SoH}}$  with the parent cycle's SoH label. As further explored in Section 5.7.2, this arrangement allows partial usage data to inform SoH regressors without requiring strictly complete cycles.

The methodology described above establishes a flexible, partial-cycle-oriented pipeline that aligns with the real-world complexity of EV battery usage. By selectively segmenting and augmenting data, the approach sidesteps the need for ideal, full 0–100% cycles. Instead, all relevant intervals—long or short, charge or discharge—become valuable building blocks for data-driven SoC, SoE, and SoH prediction, ultimately enhancing model accuracy and practical applicability.

## 5.5 Machine Learning Models (ML)

To achieve robust and accurate predictions of battery SoH and SoC, we systematically evaluated multiple ML models, ultimately selecting Light Gradient Boosting Machine (LightGBM) as the primary modeling framework. LightGBM was chosen for its superior performance, computational efficiency, incremental learning capability, and practical feasibility for embedded BMS deployment. This section thoroughly discusses the evaluated models, provides comprehensive mathematical formulations, explains the rationale for model selection, and details hyperparameter tuning procedures.

### 5.5.1 Evaluated Models

We conducted a comparative analysis of various regression algorithms, as follows:

**Gradient Boosting Decision Trees (GBDT) — LightGBM (Core Model)** The primary model employed was the LightGBM algorithm, renowned for handling complex, nonlinear relationships inherent in battery data [181]. LightGBM utilizes a leaf-wise (best-first) growth strategy, differentiating itself from traditional gradient boosting methods that typically grow trees level-wise. By selecting splits that yield the highest loss reduction, LightGBM rapidly achieves deeper trees with fewer iterations, ensuring faster convergence and enhanced predictive accuracy. The core mathematical formulation for LightGBM involves minimizing the following objective function over the training dataset  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ :

$$\hat{f} = \arg \min_{\hat{f}} \sum_{i=1}^N L(y_i, \hat{f}(x_i)) + \sum_{k=1}^K \Omega(f_k), \quad (5.5)$$

where  $L$  is the differentiable convex loss function, and  $\Omega$  regularizes tree complexity to prevent overfitting, defined as:

$$\Omega(f_k) = \gamma T + \frac{1}{2} \lambda \|w\|^2, \quad (5.6)$$

with  $T$  being the number of leaves,  $w$  the leaf weights, and  $\gamma, \lambda$  regularization coefficients.

**Support Vector Regression (SVR)** As an alternative approach to validate generalizability, Support Vector Regression (SVR) was briefly tested due to its well-known robustness and nonlinear modeling capabilities [182]. Although SVR effectively captures complex, high-dimensional relationships through kernel functions, it faced computational limitations in scaling with larger battery datasets, making it suitable mainly for benchmark comparison rather than operational deployment.

**Multi-Layer Perceptron (MLP)** Shallow neural networks (MLPs) with one or two hidden layers were evaluated for their predictive capacity and adaptability. Despite providing good initial accuracy, MLPs exhibited practical limitations in incremental retraining and deployment on embedded systems due to computational overhead and susceptibility to forgetting in online retraining contexts.

**Decision Trees and Random Forests** Baseline methods included simple decision trees and random forests, selected for interpretability and fast inference benchmarks. Although these models helped in validating the significance of nonlinearity, they fell short in predictive accuracy and incremental learning capability compared to LightGBM.

**Passive-Aggressive Regressor (PA)** The Passive-Aggressive algorithm provided insights into online incremental learning paradigms. It quickly adapted to sequential data but lacked the predictive robustness and efficiency required for accurate battery state estimation.

### 5.5.2 Model Selection Rationale

The final choice of LightGBM as the core model was guided by several critical criteria:

- **Predictive Accuracy:** LightGBM consistently delivered the highest accuracy, frequently exceeding 98% ( $R^2$  score) in validation and test scenarios.
- **Computational Efficiency:** Optimized for fast inference, LightGBM easily met real-time prediction requirements of embedded BMS deployments.
- **Memory Footprint:** Compact model representation (tens of kilobytes) enabled practical integration into memory-constrained microcontrollers.
- **Incremental Learning Capability:** LightGBM's inherent ability to incrementally update models with new data perfectly suited evolving battery health monitoring requirements.
- **Interpretability:** Clear feature importance metrics and interpretable decision pathways aided in model validation and understanding.

### 5.5.3 Hyperparameter Tuning

**Hyperparameter Selection and Stability:** We selected model hyperparameters (e.g., number of neural network layers, learning rate, tree depth for XGBoost, etc.) via preliminary cross-validation on one dataset (Dataset A). Importantly, these hyperparameters were then **fixed** for all other experiments and datasets. We did not perform dataset-specific tuning. This approach tests the robustness of our models: indeed, the chosen configuration yielded strong results across multiple datasets (see Section 5.4), indicating that the models generalize well without per-dataset adjustment. For example, the LSTM network trained on Dataset A (using the same architecture and parameters) could predict Dataset B’s states with only a slight error increase. This consistency suggests that our features and model structure capture fundamental battery behavior that doesn’t heavily depend on hyperparameter finetuning for each new dataset. We will add this observation to emphasize **hyperparameter stability**. In a deployment scenario, this means one could train the model on a standard dataset and apply it to a new vehicle’s data with minimal additional calibration – a desirable property for reducing maintenance effort.

Extensive hyperparameter optimization was conducted to ensure optimal model performance, employing grid search, random search, and Bayesian optimization strategies. Figure 5.7 provides a visual summary of our machine learning model selection pipeline, including the primary algorithm (LightGBM), validation flow, and hyperparameter tuning strategies. While, Table 5.2 summarizes hyperparameters systematically tuned during the LightGBM model training:

Table 5.2: Hyperparameters Optimized for LightGBM Model

| Hyperparameter                  | Description                                     | Typical Values |
|---------------------------------|-------------------------------------------------|----------------|
| Number of Estimators            | Number of boosted trees                         | 100–1000       |
| Learning Rate ( $\eta$ )        | Controls model convergence speed                | 0.01–0.1       |
| Max Depth                       | Maximum tree depth to prevent overfitting       | 6–12           |
| Min Data per Leaf               | Minimum samples required to form a leaf         | 20–100         |
| Feature Fraction                | Fraction of features randomly selected per tree | 0.6–1.0        |
| L1 Regularization ( $\alpha$ )  | L1 penalty on leaf weights                      | 0–1            |
| L2 Regularization ( $\lambda$ ) | L2 penalty on leaf weights                      | 0–1            |
| Early Stopping Rounds           | Stops training if no improvement is observed    | 30–100         |

## Machine Learning Model Training and Hyperparameter Tuning Workflow




Figure 1: Overview of Machine Learning Model Training and Hyperparameter Tuning

Figure 5.7: Model selection and training pipeline for SoH and SoC estimation. The flowchart illustrates the primary model (LightGBM), supplementary models tested, the training-validation workflow, and key hyperparameters optimized during development.

Hyperparameters were iteratively optimized using validation data until stable, generalizable performance was achieved. The final LightGBM models consistently demonstrated exceptional predictive accuracy, computational efficiency, and adaptability, justifying their central role in our battery state estimation framework.

### 5.5.4 Performance Evaluation Metrics

To rigorously evaluate model accuracy and deployment readiness, we adopted the following metrics:

- **Root Mean Squared Error (RMSE):**

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2} \quad (5.7)$$

- **Mean Absolute Error (MAE):**

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \quad (5.8)$$

- **Inference Latency:** Time (ms) to compute SoC/SoE per sample on edge hardware.
- **Memory Footprint:** Model size (kB) post-quantization.

## 5.6 Model Update Mechanisms and Workflow

A defining aspect of the digital twin architecture (as described in Chapter6) is the cloud–edge feedback loop that keeps the predictive models synchronized with the battery’s evolving state. Unlike a static model that is trained once, our SoC model is continuously retrained over the vehicle’s lifetime as new data becomes available and the battery ages. This section describes the methodology for model updates, including the triggers for initiating an update, the process of retraining and deployment, and measures for ensuring the update is done safely and efficiently. The approach can be summarized by the following sequence of steps (ref. Fig.6.2 in Chapter6 for a conceptual diagram of this flow):

1. **Initial Model Training (Cloud):** Before or at the time of deploying the battery in the EV, we perform an offline training of the models using historical data. The SoH estimation model is trained on a comprehensive dataset of aging trajectories (e.g., lab test data from similar cells) to capture the long-term capacity fade behavior. The initial SoC model is trained on data from a new (healthy) battery – typically using either lab-generated cycles at SoH 100% or the early cycles of the vehicle’s

own battery before significant aging occurs. The initial SoC model thus reflects the relationship between measurable signals and SoC under nominal conditions. Both models (SoH and SoC) are then loaded into the digital twin system. The SoH model resides on the cloud server, while the SoC model (initial version) is deployed on the vehicle’s edge computing unit (part of the BMS hardware)

2. **Data Collection (Edge):** Once the vehicle is operational, the edge node continuously collects battery data during normal operation. This includes time-series of voltage, current, temperature, etc., along with any ancillary data (e.g., vehicle operating states) that might be relevant. The edge node locally runs the SoC model in real-time to provide on-board SoC estimates to the BMS (for use in charge display, control, etc.). Meanwhile, it buffers the measurement data and periodically transmits it to the cloud. The data transmission can be periodic (e.g., upload daily or after every drive), or event-triggered (e.g., if the vehicle comes in for service, data dump, or if connectivity is continuous, stream at low bandwidth). In our implementation, we assumed periodic uploads of summarized cycle data – for instance, after each significant drive or charge event, the recorded profile is sent to the cloud for analysis. This IoT data pipeline uses efficient protocols to minimize bandwidth; for example, rather than sending raw 1 Hz data continuously, the edge might send compressed representations or only key cycle metrics and a subset of points. If an event such as an error or anomaly is detected, the edge can also flag that in the upload.
3. **Cloud-side SoH Monitoring:** The cloud hosts the SoH model which is periodically executed to assess the battery’s current health status. Using the latest uploaded data (say the latest reference cycle or latest few cycles), the cloud computes an updated SoH (remaining capacity). This SoH value can be sent back down to the vehicle or used as feedback internally. The cloud effectively serves as a “health tracker,” since it has more computational power and historical context to accurately compute SoH, especially using algorithms that might be too heavy for the edge. In practice, the SoH could also be estimated on the edge, but offloading to cloud allows using a more sophisticated model with full historical data. The key role of the cloud here is to decide when the SoC model needs updating.
4. **Update Triggering:** We define specific conditions that trigger a model update for the SoC estimator. Two primary triggers were designed:
  - **Periodic Update:** A time-based trigger that initiates retraining at regular intervals (for example, weekly or monthly). This ensures that even if the battery is aging slowly, the model is periodically refreshed to incorporate the most recent data. The interval can be tuned based on how fast we expect model drift – for a rapidly cycling fleet vehicle, perhaps weekly; for a slower aging scenario, monthly or per a certain number of cycles.

- **SoH-Threshold Update:** An event-based trigger that fires when the battery's SoH has dropped by more than a specified amount since the last model update. For instance, in our experiments we often used a threshold like 1% drop in SoH as a cue to retrain. The rationale is that as soon as the battery has significantly degraded, the old SoC model (trained on a healthier state) might become inaccurate, so it should be updated. This trigger ties directly into the cloud's SoH monitoring: whenever a new SoH estimate comes in, the cloud compares it to the SoH at last training. If the difference exceeds X% (e.g., from 100% to 98.9%, surpassing a 1% drop), it flags the need for update.
- **Hybrid Approach:** In practice, we combine both triggers in a logical OR: the update will occur at least once every fixed period, or sooner if a SoH drop threshold is hit. This ensures we don't miss an update if the battery ages rapidly in a short time, and conversely we don't update too often if the battery is hardly changing.

These triggers were found to be effective in our simulation of the digital twin's operation. For example, at early life the periodic schedule might dominate (since SoH changes slowly), but as the battery nears end-of-life and SoH starts dropping faster per cycle, the threshold trigger might kick in more frequently (causing, say, weekly updates instead of monthly by the end). Figures 5.8 clearly shows the updating mechanism selected in this work where the hybrid method selected to be usefull in EVs cases!




Figure 5.8: Updating Mechanism Strategy

5. **Retraining in the Cloud:** Once an update is triggered, the cloud aggregates the latest available data (and possibly past data as well) to retrain the SoC model. The training procedure here is essentially the same as the initial training described earlier, except now it uses data that includes more recent cycles of the specific battery in operation. There are a few strategies for retraining:

- **Full Retraining vs. Incremental:** Full retraining means we train a new model from scratch using all data accumulated so far (or a representative subset). Incremental would mean we continue training the existing model with just the new data. We primarily used full retraining on cumulative data to avoid any bias or order effects, given that training on a few thousand points is quite fast. However, we did ensure that earlier data was weighted properly – possibly giving more weight to recent data if we suspect concept drift. In one approach, we maintained a rolling window of the last N cycles to train on, under the assumption that very old data (when the battery was new) might be less relevant to current performance. But we found including the whole history did not hurt and provided context, as long as the model capacity was sufficient.
- **Updating SoH Model:** Typically, the SoH model doesn't need frequent retraining because it's more general (and often the initial SoH model is already

trained on extensive aging data from various cells). In our architecture, the SoH model was static. However, if the vehicle’s data reveals something new (say the battery behaves slightly differently than training data), the SoH model could also be fine-tuned. In our case, no such update was needed within the span of our experiments, so the SoH model remained the one trained initially on lab data, providing SoH estimates on the fly.

- The outcome of retraining is a new set of model parameters (for the SoC model). For example, a new LightGBM model file is produced after training on the augmented dataset.
6. **Deployment to Edge:** After successful retraining on the cloud, the updated model is transferred to the edge node in the vehicle. This could be done through an OTA update mechanism. Since the model file is relatively small (on the order of kilobytes), this transfer is quick and requires minimal bandwidth. We ensure a secure and reliable transmission (using encryption and checksum to verify the model integrity). Once the edge node receives the new model, it replaces the old SoC model with the updated one. This typically involves swapping the model in memory/storage and then the BMS software begins using the new model for subsequent SoC calculations. We design this swap to happen during an idle period if possible (e.g., when the vehicle is parked or BMS is in a steady state) to avoid any transient glitches. The system is designed such that if the model file is corrupted or the update fails, it will continue using the previous model (see rollback below).
  7. **Continuous Operation and Next Cycle:** After update, the system goes back to normal operation (Step 2), now with a more up-to-date SoC model (for example, a model that knows the battery’s capacity is, say, 90% instead of assuming 100%). The edge continues to compute SoC in real-time, presumably now with better accuracy under the new battery conditions, and data collection continues. This loop (Steps 2–6) repeats throughout the battery life, thus closing the loop of the digital twin: real data → cloud update → new model → deployed back to edge.

Throughout the model update cycle, the guiding priority is to maintain or improve accuracy without compromising system reliability. By updating the SoC model only when needed (either time-driven or SoH-driven), we minimize unnecessary churn. By validating each new model and providing rollback, we ensure the BMS can trust the digital twin’s outputs. The cloud-edge synergy allows heavy computations (like training on large historical data for SoH or hyperparameter tuning) to occur off-vehicle, while the vehicle-side computation is kept light (just running a predictive model).

## 5.7 Modeling Approach for SoC/SoE and SoH Estimation

SoC/SoE and SoH are handled by separate but interconnected models. This choice stems from the observation that real-time control (for SoC/SoE) often has constraints on latency and computational load, whereas SoH estimation can utilize more resource-intensive analytics without requiring sub-second updates.

### 5.7.1 SoC and SoE Modeling

The SoC and SoE estimators are governed by data-driven models that predict the current usable charge and energy under varying operating conditions. Building on the concept introduced at COINS 2022 [165] and refined in subsequent work [64], the methodology leverages a *dual-model* scheme:

1. A *lightweight on-board estimator* runs at the edge, processing real-time voltage and current data to provide instantaneous SoC/SoE estimates. Its structure is kept simple (e.g., a shallow neural network or a low-complexity ensemble) to ensure fast inference on embedded hardware.
2. A *more sophisticated cloud-based model* retrains periodically with partial discharge segments, advanced feature sets, and large-scale historical data. This model can capture complex non-linearities but may require substantially more computational resources. Once retrained, it transfers updated parameters or entire model snapshots to the on-board system.

This dual structure balances real-time performance with the accuracy gains derived from extensive retraining. As an example, an LGBM (Light Gradient Boosting Machine) or a deep neural network with dropout layers can be employed in the cloud-based model, while a compressed version or a simpler ML model is deployed on the vehicle control unit.

### 5.7.2 SoH Modeling

SoH estimation aims to quantify capacity fade, internal resistance increase, and other aging phenomena. Following our earlier frameworks introduced in [167] and validated in [16], SoH is modeled by correlating historical usage data with gradually measured capacity decline. A random forest regressor or an LGBM can be trained to learn the mapping from health-related features (e.g., cycle count, temperature extremes, partial discharge patterns) to the residual capacity. Similar to SoC/SoE, the SoH model is retrained in the cloud environment, leveraging the newest data from vehicles in the fleet.

In parallel, routine or opportunistic capacity checks (such as infrequent full charges or reference discharge tests) validate and calibrate the SoH model. When the model consistently detects a drop below a certain threshold, the system can trigger additional

checks or re-segmentation strategies to refine its accuracy. Ultimately, this continuous learning process addresses the challenge of capacity fade, which can vary widely across cells, modules, or packs due to manufacturing variances and usage heterogeneity.

Once the SoC/SoE and SoH models are initialized, it is essential to update them over time. Cell aging, unpredictable driving behaviors, and temperature extremes can cause shifts in battery performance that deviate from the model’s initial training distribution. If the model is not retrained, it may lose accuracy, leading to suboptimal or unsafe estimations.

## 5.8 Implementation Details and Workflow

**Embedded Runtime Performance:** To assess feasibility, we consider the algorithm’s runtime on a typical automotive microcontroller. Our models (e.g., an XGBoost with 100 trees or an LSTM with 2 layers of 50 units) involve on the order of  $10^4$  to  $10^5$  floating-point operations per inference. On a modern automotive-grade processor (such as an NXP MPC5777C @ 200 MHz or a TI TMS570 @ 300 MHz), which can execute  $10^6$  FLOPs per millisecond, our model’s forward pass would take only a few milliseconds. Indeed, using a Raspberry Pi (1.4 GHz ARM Cortex-A53) as a proxy, we measured an average inference time of 2.3 ms for the LSTM model (and  $<1$  ms for the simpler XGBoost model) on one data segment. This is well within the 100 ms sampling period of BMS state estimation tasks:contentReference[oaicite:40]index=40:contentReference[oaicite:41]index=41. In automotive terms, even a low-cost microcontroller should handle the computation without frame overruns, given that others have implemented state estimators running in 5 ms on similar hardware:contentReference[oaicite:42]index=42.

We also evaluated memory footprint: the model parameters (weights) total about  $P$  kilobytes, and the feature buffer for one segment is negligible. This easily fits in the on-chip RAM of common BMS controllers. The CPU load introduced by our algorithm would be a small fraction of the available processing budget, leaving room for other BMS tasks (e.g., cell balancing algorithms). We will include these quantitative benchmarks in the thesis to strengthen the feasibility argument. They show that our approach can meet **\*\*safety-critical timing constraints\*\*** in an embedded environment. For instance, if the BMS update loop is 100 ms, our model using 2–5 ms per update leaves a comfortable margin, aligning with ASIL (Automotive Safety Integrity Level) requirements for real-time performance.

**Safety and Certification Considerations:** Deploying a machine-learning-based BMS in a vehicle requires compliance with automotive functional safety standards like **\*\*ISO 26262\*\***:contentReference[oaicite:48]index=48. These standards demand a rigorous development process, verification, and validation for any safety-critical component. In the context of our work, the SoC/SoH estimation influences important decisions (e.g., when to derate power or trigger a warning). To integrate our digital twin approach into an **\*\*ASIL\*\*** (Automotive Safety Integrity Level) environment, one strategy could be to

use it as a complementary observer rather than a single point of failure. For instance, the ML model's output might be cross-checked with a simpler redundant model or bounded by safety limits. We will mention that ensuring determinism and **fail-safe behavior** (e.g., reverting to a conservative fallback method if the ML model outputs implausible values) would be crucial steps toward certification:[contentReference\[oaicite:49\]index=49](#). Encouragingly, there is growing research on certifying AI components in vehicles, and our approach could adopt measures such as formal verification of the model for known operating ranges, or restricting the domain of prediction to where it's validated.

Another important aspect is **model interpretability**. In automotive applications, engineers and regulators often prefer transparent models (or at least the ability to explain why a model gave a certain output). Our digital twin provides some level of interpretability by design: it includes a simplified physical model (the "white-box" component) alongside the data-driven model. The physical model can be examined for parameters like internal resistance or capacity fade, which are physically meaningful. For the data-driven part, we can apply explainable AI techniques – for example, SHAP values or feature importance analysis for the XGBoost model – to understand which input features most influence the SoH prediction. We will add a brief note that preliminary interpretability analysis showed, say, that features related to high-current charge segments were the most influential for SoH estimation (which aligns with known electrochemical behavior). Emphasizing interpretability not only builds trust in our model's predictions but also aids in debugging and improving the model by highlighting relevant battery behaviors.

Overall, by addressing **functional safety** requirements and incorporating **explainability**, the proposed methodology can move closer to an automotive-grade solution. We will integrate these points in our discussion to acknowledge their importance: a model that is accurate is necessary but not sufficient – it also needs to be **safe** and **understandable** to be viable in production EV BMS.

To streamline development, a Python-based ecosystem is adopted for data preparation, model training, and inference. Libraries such as NumPy, pandas, scikit-learn, LightGBM, and Tensorflow support end-to-end pipelines, from raw telemetry parsing to advanced ML algorithms. When necessary, laboratory-based simulation and validation can be conducted in MATLAB/Simulink, particularly for verifying control-oriented components or replicating in-lab partial discharge experiments.

Where edge deployment is concerned, embedded microcontrollers commonly impose constraints on available memory and real-time computing capabilities. Simple solutions include converting the trained ML model to a compressed format (e.g., using ONNX or custom quantization strategies) or employing a smaller ensemble that offers near-equivalent performance. For SoH, which generally operates with lower update-rate requirements, more computationally intensive algorithms can remain in the cloud, updating SoH predictions after each driving cycle or at scheduled intervals.

In conclusion, the methodology established in this chapter results in a unified, adaptive digital twin for EV battery management. We have described how data from various sources is processed and utilized, how features are engineered to capture battery behavior,

what machine learning models are employed and how they are trained, and finally, how these models are kept up-to-date via cloud-edge interaction. This methodology chapter deliberately focused on the how – the implementation details and process – in contrast to Chapter6’s conceptual architecture discussion and the upcoming results chapter (Chapter8) which will evaluate the performance of these methods. By presenting the methodological flow as a coherent narrative (spanning initial dataset use in early research to the integrated system in the final study), we have shown the evolution of the system and ensured that each component of the digital twin is rooted in rigorous experimental practice.

With the methodology defined, we now have the foundation to examine the outcomes. In the next chapter, we will present the results of applying this methodology – demonstrating the accuracy improvements from feature engineering, the efficacy of the model update mechanism in tracking battery state over time, and the overall performance of the digital twin architecture in various scenarios. (These results are discussed in Chapter8, following the methodology outlined above.)



# Chapter 6

## Consolidated Experimental Results and Analysis

This chapter provides a comprehensive, in-depth presentation of the experimental results achieved across the five research phases, each tied to a distinct publication as shown in Table 6.1

Compared to conventional BMS approaches, these five phases chronologically illustrate how our solution evolved from basic, offline data-driven modeling to a dynamic, *cloud-edge integrated* digital twin capable of handling partial discharges, real-time SoC/-SoE updates, and SoH-aware lifecycle management. Each phase highlights different battery chemistries, datasets, and partial vs. full cycle complexities. The following sections report results, figures, and tables from each paper, recontextualized into a single cohesive “results chapter” for this thesis. Note that many smaller partial experiments related to hyperparameters search where optuna have been used to conduct the search of the best hyperparameters as discussed in the previous chapter.

### 6.1 Phase I (COINS): Foundational Digital Twin on NASA Dataset

#### 6.1.1 Motivation and Scope

The earliest stage of our research, corresponding to the **COINS** paper, sought to verify whether a *dual-model digital twin* architecture could simultaneously track *SoC* and *SoH* with reasonable accuracy, however, on full discharge cycle. Two separate data-driven models were developed:

- A **SoH model**, trained once to predict long-term capacity degradation from a historical record of reference discharge events, thereby capturing how the maximum available capacity evolves over time.

Table 6.1: Overview of Research Phases and Key Publications

| Phase | Publication                                                              | Main Contribution                                                                                                                                                                                                    |
|-------|--------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1     | [165] A Machine Learning-based Digital Twin for EV Battery Modeling      | Foundational digital twin concept for <i>SoC/SoH</i> estimation based solely on full discharge cycles.                                                                                                               |
| 2     | [64] Digital Twins for EV SoX Battery Modeling                           | Extends previous SoC method to SoE, and validate it using extra verified public dataset                                                                                                                              |
| 3     | [167] Model-Driven Dataset Generation for Data-Driven Battery SoH Models | Proposal of synthetic data generation to address data scarcity, focusing on on-board feasibility (memory overhead, inference time)                                                                                   |
| 4     | [16] Model-Driven Feature Engineering for Data-Driven Battery SOH Model  | Develops advanced feature extraction for partial discharge events, improving SoH accuracy                                                                                                                            |
| 5     | [166] (final T-VT journal)                                               | Final integrated digital twin solution supporting partial discharge adaptation, dynamic re-training, multiple battery chemistries, and real driving patterns; validated on NASA, Sandia, and Stanford UDDS datasets. |

- A **SoC model**, which undergoes *periodic retraining* to incorporate the aging effects indicated by the SoH model—especially once the battery experiences any substantial capacity fade.

We used the NASA Randomized Battery Usage dataset (LiCoO<sub>2</sub> chemistry) [160] for this proof-of-concept. The dataset includes multiple cells (e.g., RW9, RW13) undergoing near-complete discharge cycles from 100% to 0% SoC as discussed in 5.2.1, enabling a controlled environment to confirm feasibility. Although partial usage data existed, this initial validation primarily leveraged near-full discharge cycles to isolate fundamental issues of the digital twin design.

### 6.1.2 Experimental Setup and Key Findings

**Raw Feature Selection for SoH** Our approach to SoH estimation involved time-series raw features (voltage, current, temperature) without applying the feature extraction. We found, however, that including a *relative time* feature—representing the discharge time in seconds significantly improved model accuracy. Empirically, ignoring relative time often led to the model missing subtle discharge-speed variations that correlated with battery health. For instance, a Random Forest (**RF**) reached an RMSE of 1.773% when including relative time, whereas excluding it caused the RMSE to balloon to over 100%. This stark difference underscores how partial temporal cues can be pivotal for state-of-health modeling.

**SoH Models Comparison and Runtime** Table 6.2 compares three candidate regressors—*Random Forest* (RF), *LightGBM* (LGB), and a *Deep Neural Network* (DNN)—for SoH estimation, evaluated on NASA cells (e.g., RW9, RW13). Each entry reports the total training time (T), inference time (I), and final errors. In line with prior observations, RF displayed the best RMSE (1.773%). LightGBM followed closely (2.310%), whereas the DNN’s accuracy was significantly worse ( $\sim 7\%$ ). We attribute the DNN’s subpar result to insufficient training samples and the complexity of capturing battery aging patterns in deep architectures. Notably, for this offline SoH model, training time is less critical, but an excessively long training stage can still delay the digital twin deployment. Inference time ( $\approx 1$  s for all methods) is more crucial for real-time or frequent SoH updates in the cloud.

Table 6.2: Phase I (COINS): Different SoH Models Performance. T + I indicates training + inference time in seconds; RMSE, MAE in %.

| Model | T + I (s)       | RMSE  | MAE   |
|-------|-----------------|-------|-------|
| RF    | 13.704 + 0.470  | 1.773 | 0.603 |
| LGB   | 10.629 + 0.977  | 2.310 | 0.791 |
| DNN   | 399.582 + 1.439 | 7.114 | 1.696 |

**SoC Estimation and Model Comparison** The same general features (including the newly emphasized *relative time* attribute) were utilized for SoC estimation. Two target SoC conditions were studied: one at a nominal SoH = 100% (healthy battery), and another at a reduced SoH = 70% (aged scenario). Tables 6.3 and 6.4 summarize the performance for each method. RF and LGB again showed strong accuracy, while DNN training demands were two orders of magnitude higher and lacked stable performance.

This finding highlights that overly complex networks may be ill-suited to typical BMS hardware or use-cases with frequent updates.

Table 6.3: Phase I (COINS): SoC Model Performance at **SoH = 100%**.

| <b>Model</b> | <b>T + I (s)</b> | <b>RMSE (%)</b> | <b>MAE (%)</b> | <b>Max Err. (%)</b> |
|--------------|------------------|-----------------|----------------|---------------------|
| RF           | 0.825 + 0.034    | 0.238           | 0.132          | 1.442               |
| LGB          | 0.140 + 0.005    | 0.224           | 0.125          | 1.010               |
| DNN          | 67.002 + 0.138   | 3.609           | 1.783          | 2.478               |

Table 6.4: Phase I (COINS): SoC Model Performance at **SoH = 70%**.

| <b>Model</b> | <b>T + I (s)</b> | <b>RMSE (%)</b> | <b>MAE (%)</b> | <b>Max Err. (%)</b> |
|--------------|------------------|-----------------|----------------|---------------------|
| RF           | 0.622 + 0.017    | 0.809           | 0.624          | 3.873               |
| LGB          | 0.097 + 0.005    | 0.599           | 0.549          | 2.993               |
| DNN          | 29.97 + 0.068    | 2.775           | 1.380          | 5.071               |

We then investigated *model mismatch* arising when a SoC model trained at a previous SoH is used after further aging has occurred. Table 6.5 further quantifies how an SoC model from  $\text{SoH} = 100\%$  yields up to 20% error at  $\text{SoH} = 75\%$ , which is unacceptable given the normal operating SoH range of 80% to 100%. Conversely, the up-to-date model had an MAE of only  $\approx 0.62\%$  (RF) or  $\approx 0.55\%$  (LGB). Figure 6.1 presents a comparative analysis of State of Charge (SOC) estimation accuracy between a recalibrated SOC random forest (RF) model, trained on discharge data at a State of Health (SOH) of 75%, and legacy SOC models derived from historical SOH conditions (100%, 95%, 85%, and 75%). The evaluation is conducted under operational conditions corresponding to  $\text{SOH}=75\%$ .

Phase I validated the *dual-model digital twin* concept in a controlled setting of near-complete discharges. Although partial discharge patterns were not heavily addressed here, the results set the foundation for subsequent phases to handle real driving cycles, adopt advanced feature extraction, incorporate SoE, and integrate partial discharge segmentation.

Table 6.5: Phase I (COINS): Impact of an Outdated SoC Model (Trained at SoH = 100%) Used at SoH = 75%.

| Model | RMSE (%) | MAE (%) | Max Error (%) |
|-------|----------|---------|---------------|
| RF    | 184.450  | 11.731  | 23.516        |
| LGB   | 188.917  | 11.858  | 23.666        |
| DNN   | 859.818  | 25.315  | 54.067        |




Figure 6.1: SOC discharge profiles comparison among different SOC Random Forest models trained in different SOH cases.

## 6.2 Phase II (AEIT): Extending to SoE and Multi-Chemistry

### 6.2.1 Motivation and Contributions

Building upon the foundational work in Phase I, our next objective was to incorporate SoE estimation while further generalizing the digital twin to multiple battery chemistries. As reported in the **AEIT** paper, two major enhancements were undertaken:

1. **SoE Estimation:** Beyond SoC, monitoring the battery’s SoE offers a more direct measure of the usable energy in Wh. This is critical for range prediction and energy utilization in EV applications.
2. **Multi-Chemistry Validation:** In addition to reusing NASA’s LiCoO<sub>2</sub> (LCO) dataset, we tested on Sandia’s multi-chemistry dataset (NMC, LFP, NCA), introducing partial-discharge scenarios to examine cross-chemistry adaptiveness and partial SoC/SoE capability.

The study integrated a flexible training pipeline capable of handling both SoC and SoE outputs, while retaining a separate SoH model from Phase I. We hypothesized that a *joint SoC–SoE* regressor would exploit common battery features (e.g., voltage responses under load) to enhance overall performance.

### 6.2.2 Machine Learning Approach and Data Splitting

**XGBoost for SoC, SoE, and SoH** In contrast to Phase I, which relied largely on LightGBM or random forests, here we employed **XGBoost**—a powerful gradient-boosting method featuring:

- *Regularized Learning Objective:* mitigates overfitting via sparsity-aware splitting and  $L_1/L_2$  regularization.
- *Parallel Tree Construction:* reduces training time compared to naive GBDT implementations.
- *Robust Handling of Missing Values:* beneficial when dealing with real-world battery logs containing occasional sensor gaps.

Like LightGBM, XGBoost remains well-suited for the complex relationships of battery aging.

**Dataset Partitioning** We used two core datasets:

1. **NASA (LiCoO<sub>2</sub>) dataset:** Owing to limited cell availability and partial usage logs, we split it into 75% for training and 25% for testing. This dataset had fewer cycles per cell, focusing primarily on random-walk conditions plus reference cycles for SoH labeling.

2. **Sandia dataset (NMC):** Larger and more varied, so we split it into training (50%), validation (40%), and testing (10%) again on full discharge cycles.

In each scenario, we avoided reusing the same cell data for both training and testing, ensuring robust cross-cell validation and how the model can generalize on different operation conditions.

### 6.2.3 SoH Estimation Results Using XGBoost

**Feature Selection and Relative Time** We again confirmed that including **relative time** of discharge improved SoH model accuracy substantially. Ignoring this feature led to large errors when voltage, current, and temperature alone were used. This finding was consistent across NASA and Sandia data, reflecting how discharge duration captures critical capacity fade signatures.

**SoH Model Performance and Runtime** Table 6.6 compares SoH estimation on NASA and Sandia. XGBoost yields excellent accuracy on NASA with **RMSE** = 0.0219% and on Sandia = 0.113%, although the latter has a slightly higher absolute MAE (2.096%) due to broader chemistry/cycling variety. Meanwhile,  $R^2$  on NASA is very high ( $\approx 0.997$ ), underscoring that even a modest real dataset can yield strong SoH predictions if the features are well-chosen. Training time is longer on Sandia owing to its larger set and the presence of a validation split.

Table 6.6: Phase II (AEIT): **SoH** Models Performance using XGBoost (T + I in seconds).

| Dataset | T + I (s)     | RMSE(%) | MAE(%) | $R^2$  |
|---------|---------------|---------|--------|--------|
| NASA    | 0.579 + 0.090 | 1.0219  | 0.811  | 0.997  |
| Sandia  | 5.166 + 0.022 | 5.113   | 2.096  | 0.7825 |

While training time is typically an offline process in the cloud (less critical than inference time), extremely long builds can delay updates to the digital twin. Here, inference time remained below 0.1 s for NASA and Sandia, suitable for frequent SoH queries (e.g., once per partial cycle or daily).

### 6.2.4 SoC and SoE Estimation

**Combined SoC–SoE Model** We reused the same features (voltage, current, temperature, relative time, etc.) for SoC and SoE. Unlike SoC, SoE calculations incorporate instantaneous power (voltage  $\times$  current). A single XGBoost model was set up to jointly predict SoC and SoE, although in principle separate models could also be used.

To assess the effect of aging, we trained the SoC–SoE model at a known SoH (e.g., 100%) and then tested it at a later SoH (e.g., 75%). Tables 6.7 and 6.8 list the typical results. At SoH = 100%, both NASA and Sandia show sub-1% RMSE, underscoring the method’s accuracy for an unaged battery. However, at SoH = 75%, the NASA error soared to 4.363% RMSE (MAE = 28.085%), a sign that partial usage combined with limited NASA data yields insufficient coverage at the new SoH. Meanwhile, Sandia’s RMSE remains under 1% for SoC–SoE, likely thanks to a richer, larger dataset.

Table 6.7: Phase II (AEIT): Combined SoC–SoE at **SoH = 100%** (Trained at 100% SoH).

| Dataset | T + I (s)     | RMSE(%) | MAE(%) | R <sup>2</sup> |
|---------|---------------|---------|--------|----------------|
| NASA    | 0.234 + 0.019 | 0.915   | 0.671  | 0.998          |
| Sandia  | 0.178 + 0.004 | 2.033   | 1.546  | 0.996          |

Table 6.8: Phase II (AEIT): Combined SoC–SoE at **SoH = 75%** (Trained at 75% SoH).

| Dataset | T + I (s)     | RMSE(%) | MAE(%) | R <sup>2</sup> |
|---------|---------------|---------|--------|----------------|
| NASA    | 0.622 + 0.017 | 28.085  | 4.363  | 0.968          |
| Sandia  | 0.545 + 0.039 | 6.440   | 0.891  | 0.907          |

**Need for Frequent Updates** To illustrate the necessity of retraining, Figure 6.2 SoC and SoE discharge profiles for NASA (left) and Sandia (right) datasets at SoH=76%, generated by models trained across varying SoH levels. The solid blue line represents the ground-truth discharge profile, while the red line denotes predictions from the model retrained at SoH=73%, which aligns most closely with the reference (NASA: RMSE = 0.809%, MAE = 6.44%; Sandia: RMSE = 4.363%, MAE = 28.08%). In contrast, legacy models—trained on older SoH regimes—exhibit significant deviations, with errors escalating as the training data ages (Table 6.9). For instance, using a model trained at SoH=100% on degraded cells (SoH=76%) inflates NASA’s RMSE to 166.9%, rendering it operationally unusable. Retaining the updated model ensures errors remain below 5% for both metrics, even under aging conditions. While Sandia’s baseline errors are smaller, retraining still yields critical improvements, underscoring the universal need for frequent updates to mitigate accuracy decay caused by battery degradation.

Hence, **Phase II** firmly established that *SoE estimation can be integrated into the digital twin without significant overhead*, broadening state estimation beyond SoC. Additionally, verifying generality across multiple chemistries highlights the potential to unify

Table 6.9: Phase II (AEIT): SoC–SoE at **SoH=75%** but model was *trained at 100% SoH* (Outdated).

| Dataset | RMSE(%) | MAE(%) | R <sup>2</sup> |
|---------|---------|--------|----------------|
| NASA    | 166.905 | 11.403 | 0.683          |
| Sandia  | 16.041  | 3.711  | 0.615          |




Figure 6.2: Phase II (AEIT): SoC and SoE discharge profiles on NASA (left) and Sandia (right) under different SoH levels. The red line (updated model) tracks the measured profile more closely than the older models.

EV battery monitoring solutions, regardless of the underlying cell chemistry or partial usage patterns.

## 6.3 Phase III (ISLPED): Synthetic Data Generation and Resource Trade-Offs

### 6.3.1 Motivation and Objectives

Although earlier phases leveraged *real measured data* to train and validate data-driven battery models, obtaining large, high-quality labeled datasets (with explicit *SoH* ground truth) can be cumbersome and expensive. In our **ISLPED** paper, we addressed two challenges:

1. **Data Scarcity:** Battery cells for experiments are limited, and collecting comprehensive usage data (e.g., partial discharge cycles, diverse ambient temperatures) can be prohibitively time-consuming.
2. **Deployment Feasibility:** Even if we achieve a sufficiently accurate *SoH* estimator, its practical utility depends on whether it can fit within the memory/latency constraints of an automotive microcontroller or low-power hardware.

To tackle these issues, we proposed a *model-driven data augmentation* scheme that synthesizes realistic battery usage data using a known circuit-equivalent model, thereby expanding the training space at minimal cost. We then conducted a *Pareto frontier* analysis to identify the trade-off between final *SoH* accuracy and resource (memory/time/energy) overhead when deployed on a low-power RISC-V microcontroller.

### 6.3.2 Methodology Recap

**Synthetic Data Generation.** We adopted a Thevenin-based circuit model from [183] augmented with capacity fade submodels. By systematically varying ambient temperature (10°C to 40°C) and discharge profiles (square wave, random-walk, constant), we simulated battery usage until the *SoH* dropped near 0%. Each simulation spanned 1,000 hours or ended earlier if *SoH* fully depleted. This yielded 110 distinct simulation scenarios, which were segmented into windows of  $W = 2$  hours each, for a total of  $\sim 17,842$  samples. We subsequently split these into training, validation, and test sets by simulation, ensuring each scenario was exclusive to only one subset.

**Data-Driven *SoH* Model** For each segment, the circuit model provided voltage, current, temperature, and the resulting capacity fade ( $\Delta\text{SoH}$ ) labels. We trained two families of regressors—**Gradient Boosted Trees (GBTs)** and **Multi-Layer Perceptrons (MLPs)**—to

predict  $\Delta\text{SoH}$ . By systematically tuning hyperparameters (e.g., tree depth, ensemble size, MLP layer count) and feature sets, we mapped each model’s  $\Delta\text{SoH}$  error against its resource demands.

**Pareto Frontier Analysis** We collected a wide range of model configurations, each with different complexities, and plotted them in a 2D plane of “(time/energy, memory) vs.  $\Delta\text{SoH}$  error.” An ideal model would minimize error and resource usage. The resulting *Pareto frontier* indicates a family of optimal trade-offs between accuracy and cost.

### 6.3.3 Key Experimental Results

**SoH Prediction Accuracy with Synthetic Data** Even without real measured data, the **best GBT model** achieved an MAE of around **7.2%** for unseen simulation usage patterns. While this is higher than the 1–3% error typically observed in real-lab data, it demonstrates that *synthetic-only* training can yield workable SoH estimators. We attribute the gap to the circuit model’s partial mismatch with real battery phenomena and the difficulty in capturing certain aging mechanisms. However, it still provides coverage of underrepresented usage conditions.

**Pareto Frontier and Deployment on Low-Power MCU** Figure 6.3 (conceptual) shows that GBTs generally outperform MLPs in the accuracy vs. memory trade-off, though MLPs can be more memory-efficient at the cost of higher error. Table 6.10 highlights *selected* extremes from this frontier—lowest-error (-E), lowest-time (-T), and lowest-memory (-M)—for both GBT and MLP.

- The GBT-E model yields a 7.2% MAE at  $\approx 39$  kB memory and 0.68 ms inference time on the PULPissimo MCU.
- By contrast, a simpler GBT or MLP can shrink memory to around 13 kB at the expense of doubling error to 13–14%.
- The best MLP is smaller (26 kB) but yields a slightly higher 7.16% error and longer inference time (0.94 ms).

Compared to running the full circuit model from [183] on-device, which would require  $>1$  MB memory and  $>1$  s inference, these data-driven surrogates present a  $27\times$  *memory reduction* and a  $2000\times$  speedup, underscoring the feasibility of embedded SoH monitoring.

**Comparison with a Simulation Model** For fair comparison, we tried compiling the reference circuit model on a laptop-class CPU (Apple M1 Pro). Even in that environment, simulating SoH for a single inference is  $2000\times$  slower and  $27\times$  more memory-intensive than the GBT approach. Scaling or porting the entire circuit model to a small MCU




Figure 6.3: Pareto-fronts obtained from the hyper-parameters exploration of data-driven models

Table 6.10: Phase III (ISLPED): Representative Deployment Results on a Low-Power MCU (RISC-V).

| Model | MAE(%) | Latency(ms) | Memory(kB) | Energy( $\mu$ J) |
|-------|--------|-------------|------------|------------------|
| GBT-E | 7.15   | 0.68        | 39.46      | 2.60             |
| GBT-T | 13.92  | 0.25        | 13.26      | 0.94             |
| GBT-M | 14.00  | 0.28        | 13.06      | 1.04             |
| MLP-E | 7.16   | 0.60        | 15.38      | 2.36             |
| MLP-T | 11.86  | 0.28        | 13.54      | 1.01             |
| MLP-M | 11.86  | 0.28        | 13.54      | 1.01             |

was infeasible, indicating that *data-driven surrogates* are essential for real-time BMS applications.

Thus, combining synthetic and real data can be highly advantageous: real data yields excellent accuracy in typical usage, whereas synthetic data broadens coverage of corner cases (e.g., extremely high or low temperatures, partial discharge extremes). The *ISLPED* results confirm that synergy between model-driven data augmentation and careful resource analysis is a promising path for future **SoH** model deployments in production EV systems.

## 6.4 Phase IV (DATE): Feature Engineering for Enhanced SoH Accuracy

### 6.4.1 Objective and Additional Insights

The **DATE** work focused on advancing the SoH estimation for *partial-cycle* usage by systematically introducing a more powerful **feature engineering** approach. This represents a shift from conventional BMS methods, which typically rely on raw measurements (voltage, current, temperature) or rudimentary statistics. Instead, we hypothesized that incorporating additional domain-specific indicators—some derived from battery aging mechanisms—would yield more accurate partial-cycle SoH estimation. This concept is frequently referred to as *physics-informed feature augmentation*, aiming to equip standard data-driven models (LightGBM or LSTM) with features that more directly capture the physical and electrochemical cues of battery degradation.

### 6.4.2 Experimental Setup and Feature Sets

**Datasets and Cell Chemistries.** In contrast to prior phases that focused on single-chemistry or small-scale NASA data, the **DATE** phase tested on the **Sandia** dataset, specifically analyzing both **LFP (Lithium Iron Phosphate)** and **NMC (Nickel Manganese Cobalt)** cells to demonstrate chemistry-agnostic improvement. Each dataset entry is segmented by partial cycles, each annotated with  $\Delta\text{SoH}$  (i.e., the SoH difference across consecutive reference cycles).

**Feature Extraction: Conventional vs. Proposed.** We compared two broad approaches:

1. **Conventional Aggregation (CA):** Baseline statistical features (mean, min, max voltage/current/temperature) extracted from each partial segment.
2. **Novel Aggregation (NA):** Our proposed feature set, which augments the conventional ones with:
  - *SoH from prior cycles*, capturing degradation trajectory continuity.
  - *Depth-of-Discharge (DoD) decomposition*, splitting the partial cycle into subranges for localized slope analysis.
  - *SoC differences and partial usage intervals*, emphasizing short but repeated usage patterns indicative of aging.
  - *Relative Time* or cycle index, bridging consecutive partial cycles to track long-term drift.

Before applying machine learning, we removed trivial or constant features (e.g., a maximum voltage that never changes due to fixed RC cycles).

### 6.4.3 SoH Estimation: Experimental Results

**Quantitative Accuracy Improvement.** Table 6.11 contrasts the outcomes of CA vs. NA on the Sandia dataset, using **LightGBM** and **LSTM** for LFP and NMC cells. Across all metrics (MSE, MAE, and  $R^2$ ), the NA approach consistently outperforms CA. The improvement for LFP stands at about 19%, while for NMC it jumps to 27.9%, underscoring how domain-specific signals (e.g., partial DoD usage or SoC deviations) enhance the model’s ability to capture partial-cycle aging.

Table 6.11: Phase IV (DATE): SoH Estimation Error on LFP and NMC Cells with LightGBM vs. LSTM, under CA and NA Feature Sets.

| Chemistry | Model    | Features | MSE                    | MAE                    | $R^2$  |
|-----------|----------|----------|------------------------|------------------------|--------|
| LFP       | LightGBM | NA       | $5.249 \times 10^{-5}$ | $4.764 \times 10^{-3}$ | 0.505  |
|           | LightGBM | CA       | $7.322 \times 10^{-5}$ | $6.620 \times 10^{-3}$ | 0.310  |
| LFP       | LSTM     | NA       | $6.457 \times 10^{-5}$ | $5.818 \times 10^{-3}$ | 0.391  |
|           | LSTM     | CA       | $7.985 \times 10^{-5}$ | $6.853 \times 10^{-3}$ | 0.248  |
| NMC       | LightGBM | NA       | $7.817 \times 10^{-5}$ | $6.320 \times 10^{-3}$ | 0.561  |
|           | LightGBM | CA       | $1.310 \times 10^{-4}$ | $8.611 \times 10^{-3}$ | 0.230  |
| NMC       | LSTM     | NA       | $1.361 \times 10^{-4}$ | $8.468 \times 10^{-3}$ | 0.236  |
|           | LSTM     | CA       | $2.491 \times 10^{-4}$ | $1.003 \times 10^{-2}$ | -0.397 |

**Tracking SoH over the Battery Lifetime.** While Table 6.11 shows the average error across partial cycles, Figure 6.4 visualizes the SoH evolution for three representative cells, highlighting the improvement from CA (green) to NA (orange) relative to ground truth (blue, dashed). In each cell, the NA-based approach yields smaller cumulative error over time. Even if the CA curve is sporadically close during initial cycles, it drifts noticeably as aging progresses, whereas the NA curve remains more faithful to the measured SoH, thus offering robust *lifecycle tracking*.




Figure 6.4: Phase IV (DATE): *SoH* vs. *cycle count* for three sample cells. NA-based features (orange) exhibit lower error accumulation compared to CA (green).

**Feature Importance Analysis.** Figure 6.5 further confirms the relevance of newly introduced domain-specific features for LFP and NMC cells. Key insights:

- **Delta Time** and **Initial SoH** rank highest, the former capturing calendar aging, the latter anchoring each segment to its degradation baseline.
- All **DoD-related** and **SoC-deviation** features—present *only* in NA—appear in the top tier, often outranking CA staples such as min/max current.

This hierarchy confirms that the proposed feature-engineering strategy injects precisely the information most predictive of partial-cycle degradation.

In summary, *Phase IV* highlighted how systematically designed “aging-aware” transformations can significantly enhance SoH predictions under partial usage. This synergy between domain knowledge and machine learning paves the way for robust partial-cycle battery management, ensuring that real EV usage patterns are captured in state-of-health tracking.

## 6.5 Phase V (IEEE TVT): Final Integrated Digital Twin and Comparative Benchmark

The final outcome of this research is a cloud–edge battery digital twin that provides accurate and adaptive estimation of battery SoH, SoC, and SoE in EVs. The digital twin is




Figure 6.5: Feature importance estimated by LightGBM.

conceived as a modular system, with distinct models for SoH and for SoC/SoE operating in a coordinated manner. The cloud component hosts the high-level data processing and model retraining logic, while the edge component (onboard the vehicle's BMS) runs lightweight inference models that are periodically updated. This modular deployment ensures that computationally intensive tasks (e.g., training on large historical datasets or complex feature extraction) are offloaded to the cloud, whereas real-time estimation is executed efficiently on the vehicle's microcontroller. The architecture leverages over-the-air updates so that improved models trained in the cloud can be seamlessly deployed to the vehicle. This design balances accuracy and latency by exploiting cloud resources for learning and the edge for immediate, real-time battery state estimation. A key feature of the proposed digital twin is its adaptive retraining mechanism triggered by battery aging or elapsed time. Specifically, the SoC/SoE estimation model is configured to retrain whenever a significant change in SoH is detected (beyond a threshold  $\Delta$ SoH) or after a predefined time interval (e.g., daily or weekly). This proactive model update strategy ensures that as the battery degrades over months of usage, the SoC/SoE model remains calibrated to the current battery condition. In our implementation, the SoH model is

continuously monitoring capacity fade; once the SoH drop exceeds certain percentages (e.g., 1% or 5%), a trigger is sent for model update. Alternatively, a scheduled update can be performed every fixed period (such as every two weeks), even if SoH change is small, to capture any drift due to time-dependent factors. The SoH estimation model itself does not require as frequent updates – it serves as a long-term reference – but its accurate tracking of capacity is critical to decide when the SoC/SoE model retraining is needed. This division of roles (a stable SoH model and an updatable SoC/SoE model) is a cornerstone of the digital twin’s adaptability. Another innovative aspect is the dynamic segmentation of battery cycling data used for model training and evaluation. Rather than assuming full charge–discharge cycles, the digital twin is designed to handle partial cycles of arbitrary depth. The training data (drawn from various sources) was segmented by SoC window (e.g., 100–0%, 80–20%, 60–30%, etc.) to emulate real driving conditions where the battery may only undergo partial charges or discharges. During operation, the cloud-based SoH model estimates capacity fade for each segment of a drive cycle, and the segments themselves are allowed to evolve over time as usage patterns change. For instance, in one proof-of-concept scenario, the digital twin began with a full cycle (100%–0% SoC) in the first month and then adapted to predominantly partial cycles in subsequent months (e.g., 62%–25%, then 26%–8%, etc. in later months). This dynamic segmentation strategy enables the model to learn from and adapt to diverse real-world driving behaviors, ensuring that even if the vehicle’s routine shifts (such as commuting patterns or climate conditions), the model’s estimates remain reliable. It also maximizes the use of available data: even short or irregular discharge events contribute to the SoH and SoC estimations, thereby enhancing the twin’s coverage of operating scenarios beyond idealized full-cycle tests.

The final stage of our research, culminating in the **IEEE Transactions on Vehicular Technology** submission, synthesizes the outcomes and techniques from the earlier phases (COINS, AEIT, ISLPED, and DATE) into a cohesive *dynamic dual-model digital twin architecture* for EV battery management. This work represents the most comprehensive endeavor, as it:

- **Adapts to Partial Discharge Cycles:** Building upon Phase I’s observation on how partial discharges degrade prediction quality, the final approach segments real-world driving traces into smaller intervals. These segments are used both for real-time SoC/SoE predictions and for refining the SoH model (initially trained offline).
- **Incorporates Enhanced Feature Engineering:** Phase IV’s domain-specific features (e.g., partial DoD usage, voltage-relaxation attributes) are integrated to improve SoH prediction, especially under mid-range SoC fluctuations that seldom reach full 0–100% cycles.
- **Leverages Synthetic Data Augmentation:** Phase III’s synthetic dataset generation complements real data to cover extreme usage scenarios (e.g., high current pulses,

seldom-measured temperature ranges). This enriches the training set, ensuring better generalization.

- **Implements Dynamic (Re-)Training at Cloud Scale:** The SoH model in the cloud updates as new partial discharge data arrives, while the local SoC/SoE model on the EV’s BMS is periodically replaced via OTA updates. This ensures the model remains consistent with ongoing battery aging.
- **Validates Across Multiple Datasets and Chemistries:** Beyond the NASA LiCoO<sub>2</sub> and Sandia sets (LFP, NMC, NCA), we incorporate the *Stanford UDDS* data (real-driving profiles from 80% to 20% SoC). This tests the approach under genuine usage conditions, where few full cycles occur.

### 6.5.1 Datasets and Implementation Details

**NASA (LiCoO<sub>2</sub>) Recap.** We revisit NASA’s Randomized Battery Usage dataset (*RW1–RW28*), leveraging randomized charge and discharge current and operation conditions. The battery experiences moderate to severe aging (SoH from 100% down to  $\approx 40\%$ ). We apply partial segmentation at intervals of 5% or 10% SoC. Advanced domain-specific features (*Phase IV*) plus synthetic expansions (*Phase III*) address coverage gaps for high current or unusual temperature segments.

**Sandia (LFP, NMC, NCA).** A total of 32 cells across three chemistries run partial discharge cycles from 80%–60%, 60%–40% SoC windows. SoH is measured every 500 cycles via reference cycles. The final approach monitors partial cycles from day-to-day usage and triggers an SoC/SoE update if SoH differs by at least 2%. Additional simulated data fill seldom-tested corners (e.g., 3C discharges at 35°C).

**Stanford UDDS (NMC 21700).** This dataset provides the *most realistic* driving scenario: repeated UDDS cycles from 80% to 20% SoC, with dynamic load patterns (acceleration, regenerative braking). Degradation was comparatively mild (from 100% down to  $\approx 91\%$ ), but partial usage is highly variable. We partition the UDDS runs into  $\Delta$ SoC segments of 5–15%, enabling partial-cycle analytics. Periodic capacity tests supply the SoH ground truth. Only minimal synthetic expansions were used since real driving data was already quite diverse.

### 6.5.2 SoH Accuracy

The SoH regressor was trained with LightGBM, and its hyper-parameters were automatically tuned with the Optuna framework [184]. A separate validation split guided the search; the optimal configuration discovered by Optuna was then frozen and evaluated on the held-out test set. Table 6.12 summarises the resulting MAE figures for both LightGBM

and a Support Vector Regression (SVR) baseline. Across the three benchmarks—NASA, Sandia and Stanford UDDS—LightGBM achieves sub-percentage errors and near-perfect coefficients of determination. For example, on NASA cells the test MAE is only 0.88 % ( $R^2 = 0.99$ ). Sandia cells record a similar 0.73 % MAE ( $R^2 = 0.984$ ), whereas the more dynamic Stanford UDDS cycles push the error down to 0.19 % while still exceeding 0.98 in  $R^2$ .

Table 6.12: SoH MAE comparison (LightGBM vs. SVR) on validation and test sets.

| Dataset       | Validation |      | Test        |      |
|---------------|------------|------|-------------|------|
|               | LGBM       | SVR  | LGBM        | SVR  |
| NASA          | 1.35       | 1.52 | <b>0.88</b> | 1.23 |
| Sandia NMC    | 0.75       | 0.58 | <b>0.73</b> | 0.99 |
| Stanford UDDS | 0.52       | 0.38 | <b>0.19</b> | 0.17 |


**Generality over partial discharge windows.** To probe robustness under limited data, we re-computed the metrics on progressively narrower discharge slices ( $\Delta\text{SoC} = 1\text{--}100\%$ ). The detailed results appear in Table 6.13 and are visualised in Figure 6.6. Even in the most extreme case—segments covering just 1 % of the full SoC swing, which sometimes furnish as few as three feature samples—the LightGBM regressor remains accurate: 0.91 % MAE on NASA and 0.62 % on Sandia, both with  $R^2 \approx 0.99$ . Larger windows naturally reduce error further, but the smallest windows are already precise enough for real-time battery-management decisions.

**Benchmark against SVR.** SVR, equipped with a radial-basis kernel, serves as a strong non-tree baseline. Its predictive quality is comparable to LightGBM: on Stanford it posts a slightly lower MAE (0.17 % vs. 0.19 %), yet it falls behind on NASA (1.23 % vs. 0.88 %) and Sandia (0.99 % vs. 0.73 %). Given LightGBM’s faster incremental training and lighter computational footprint, it was selected for the digital-twin implementation.

**Take-away.** Error consistently below 1 % and  $R^2$  exceeding 0.98 across very different chemistries, cycling profiles, and discharge spans demonstrates that the proposed data-driven approach delivers highly reliable SoH monitoring—key for both predictive maintenance and adaptive control in electric-vehicle batteries.

Table 6.13: SoH Model Estimation Error for LightGBM and SVR

| Model    | $\Delta\text{SoC}$ of Segments | NASA    |         |      | Sandia  |         |      | Stanford |         |       |
|----------|--------------------------------|---------|---------|------|---------|---------|------|----------|---------|-------|
|          |                                | MAE (%) | MSE (%) | R2   | MAE (%) | MSE (%) | R2   | MAE (%)  | MSE (%) | R2    |
| LightGBM | 1%                             | 0.91    | 2.21    | 0.99 | 0.62    | 0.78    | 0.98 | 0.258    | 0.104   | 0.984 |
|          | 2%                             | 0.89    | 2.08    | 0.99 | 0.70    | 0.86    | 0.98 | 0.250    | 0.096   | 0.985 |
|          | 5%                             | 0.88    | 1.97    | 0.99 | 0.73    | 0.87    | 0.98 | 0.256    | 0.103   | 0.984 |
|          | 10%                            | 0.86    | 1.91    | 0.99 | 0.75    | 0.93    | 0.98 | 0.242    | 0.097   | 0.985 |
|          | 20%                            | 0.88    | 1.93    | 0.99 | 0.71    | 0.83    | 0.98 | 0.239    | 0.095   | 0.985 |
|          | 25%                            | 0.90    | 1.98    | 0.99 | 0.73    | 0.87    | 0.98 | 0.237    | 0.093   | 0.986 |
|          | 50%                            | 0.88    | 1.91    | 0.99 | 0.68    | 0.78    | 0.98 | 0.242    | 0.095   | 0.985 |
|          | 100%                           | 0.82    | 1.51    | 0.99 | 0.72    | 0.76    | 0.98 | 0.236    | 0.090   | 0.986 |
| SVR      | 1%                             | 1.58    | 4.83    | 0.98 | 1.22    | 6.27    | 0.86 | 0.177    | 0.067   | 0.990 |
|          | 2%                             | 1.75    | 5.57    | 0.98 | 1.20    | 6.07    | 0.86 | 0.176    | 0.066   | 0.990 |
|          | 5%                             | 1.65    | 4.30    | 0.98 | 1.16    | 5.23    | 0.88 | 0.176    | 0.066   | 0.990 |
|          | 10%                            | 1.55    | 3.75    | 0.98 | 1.16    | 4.60    | 0.89 | 0.175    | 0.064   | 0.990 |
|          | 20%                            | 1.46    | 3.42    | 0.98 | 0.81    | 1.06    | 0.98 | 0.174    | 0.062   | 0.991 |
|          | 25%                            | 1.55    | 3.87    | 0.98 | 0.78    | 0.92    | 0.98 | 0.171    | 0.061   | 0.991 |
|          | 50%                            | 1.48    | 3.74    | 0.98 | 0.76    | 0.88    | 0.98 | 0.170    | 0.055   | 0.992 |
|          | 100%                           | 1.62    | 4.58    | 0.98 | 0.88    | 1.28    | 0.97 | 0.148    | 0.047   | 0.993 |

Figure 6.6: SoH prediction vs. ground truth across NASA, Sandia, and Stanford datasets at varying segment sizes ( $\Delta\text{SoC} = 1, 5, 25, 50, 100\%$ ).

LightGBM achieves outstanding accuracy (MAE <1%) even at minimal segment sizes ( $\Delta\text{SoC} = 1\%$ ), confirming strong reliability for real-time partial-cycle estimation.

### 6.5.3 Adaptive SoC/SoE

The second pillar of the digital twin is an **updatable SoC/SoE model** that supplies instantaneous charge *and* usable energy. Unlike conventional BMS algorithms, this model is periodically *retrained*—either (i) after a fixed elapsed time, or (ii) when the SoH estimator detects a user-defined capacity drop.

**Fixed versus Updatable Baseline.** A frozen model trained on fresh-battery data quickly drifts as aging or new drive patterns emerge: on NASA, its SoC MAE climbs from an early-life 3–5 % to **26.3 %** for a heavily worn 36%→8% discharge slice. Inserting scheduled updates slashes that error: daily retraining brings the same slice down to **5.0 %** MAE (a  $\approx 81\%$  reduction), while monthly updates still cut the error by  $\approx 54\%$ . In Table 6.14, every **row** corresponds to a distinct discharge window found in the test set, whereas every **column** reports how the adaptive model fares against the non-adaptive (fixed) one under several update policies. Each entry quantifies the relative gain of the adaptive model, defined as

$$\frac{MAE_{\text{fix}} - MAE_{\text{update}}}{MAE_{\text{fix}}} \times 100,$$

that is, the percentage reduction in MAE produced by updating.

**NASA dataset.** Across all partial-SoC/E conditions, allowing the model to update markedly boosts accuracy: improvements span from 8.61% up to 84.06%. As an illustration, for the 36–8% SoC slice, updating after two weeks cuts the error by 75.63%, and the drop grows to 80.12% once a 5 %  $\Delta\text{SoH}$  is reached.

**Sandia dataset.** The same adaptive strategy also benefits Sandia cells, albeit with a narrower spread (0.25 %–46.22 %). In the 53–2 % SoC interval, the MAE shrinks by 35.96 % after a fortnight and by 44.13 % once the 5 %  $\Delta\text{SoH}$  threshold is crossed.

**Stanford dataset.** Replicating the analysis performed for NASA and Sandia with LightGBM (Table 6.14), only timed updates or 1 % SoH-driven updates could be tested owing to the limited degradation observed.

**Update cadence.** Generally, stretching the time (or health) gap between updates—two months or a 20 %  $\Delta\text{SoH}$ , for example—produces smaller gains, bringing the adaptive model’s error increasingly close to that of the fixed baseline. The likely culprit is the higher time-series complexity the model must absorb when many different SoH levels accumulate between successive updates.

To highlight one concrete case, Figure 6.7 plots results for the 62–25 % SoC window. With a two-week cadence and a 5 %  $\Delta\text{SoH}$  trigger, the adaptive model lowers the MAE by roughly 55.35 % and 49.65 % for NASA, and by 38.52 % and 25.01 % for Sandia, while

for Stanford with UDDS cycles, improvements range from 38.37 % (1 % SoH updates) to about 50 % with purely time-based updates; the aggregate MAE drops to 0.62 %. At that point the cells were at 67 % and 83 % SoH, respectively.

Table 6.14: Comparison of Testing MAE between Fixed Model and Model Updating

|          | SoC/E Segment(%)     | Fixed (MAE %) | SoC/E Model Update ( $\frac{MAE_{fix}-MAE_{update}}{MAE_{fix}} * 100\%$ ) |        |         |         |          |           |       |       |       |       |
|----------|----------------------|---------------|---------------------------------------------------------------------------|--------|---------|---------|----------|-----------|-------|-------|-------|-------|
|          |                      |               | Interval                                                                  |        |         |         |          | SoH drops |       |       |       |       |
|          |                      |               | 1 day                                                                     | 1 week | 2 weeks | 1 month | 2 months | 1%        | 5%    | 10%   | 15%   | 20%   |
| NASA     | 100% - 0%            | 1.240         | 21.95                                                                     | 17.83  | 16.86   | 13.15   | 8.61     | 21.35     | 18.72 | 15.93 | 16.25 | 14.26 |
|          | 62% - 25%            | 15.610        | 59.95                                                                     | 58.10  | 55.35   | 46.63   | 39.72    | 58.41     | 49.65 | 42.65 | 35.91 | 31.60 |
|          | 36% - 8%             | 26.327        | 81.09                                                                     | 79.01  | 75.63   | 66.04   | 54.79    | 80.12     | 72.88 | 66.22 | 56.35 | 51.96 |
|          | 53% - 2%             | 18.412        | 67.41                                                                     | 64.30  | 61.39   | 53.18   | 46.10    | 65.26     | 56.86 | 50.57 | 43.43 | 39.09 |
|          | 74% - 27%            | 10.997        | 49.58                                                                     | 47.79  | 46.03   | 38.48   | 34.79    | 46.99     | 37.50 | 31.38 | 26.98 | 24.91 |
|          | 80% - 69%            | 10.254        | 52.46                                                                     | 50.48  | 48.46   | 42.26   | 38.93    | 49.28     | 40.36 | 34.65 | 30.90 | 30.06 |
|          | 35% - 17%            | 26.661        | 84.01                                                                     | 81.63  | 77.40   | 65.95   | 53.32    | 83.35     | 75.83 | 67.53 | 57.13 | 49.92 |
|          | 96% - 12%            | 2.151         | 23.60                                                                     | 20.32  | 17.33   | 11.94   | 5.01     | 23.10     | 17.50 | 13.50 | 13.73 | 12.44 |
| Sandia   | 100% - 0%            | 12.442        | 45.30                                                                     | 40.64  | 40.64   | 34.86   | 20.92    | 46.79     | 40.96 | 32.20 | 27.15 | 04.40 |
|          | 62% - 25%            | 6.212         | 43.16                                                                     | 38.52  | 38.52   | 35.47   | 25.77    | 31.91     | 25.01 | 18.21 | 19.06 | 01.37 |
|          | 36% - 8%             | 7.164         | 14.56                                                                     | 10.70  | 10.70   | 12.95   | 11.87    | 19.94     | 14.42 | 12.21 | 10.15 | 5.62  |
|          | 53% - 2%             | 12.401        | 46.22                                                                     | 35.96  | 35.96   | 33.95   | 23.03    | 44.13     | 37.24 | 29.88 | 23.39 | 4.26  |
|          | 74% - 27%            | 05.634        | 35.79                                                                     | 31.75  | 31.75   | 28.75   | 20.64    | 27.09     | 19.47 | 12.68 | 14.75 | 0.26  |
|          | 80% - 69%            | 4.037         | 16.22                                                                     | 16.31  | 16.31   | 16.20   | 12.21    | 11.13     | 9.76  | 7.62  | 3.18  | -1.81 |
|          | 35% - 17%            | 5.351         | 8.11                                                                      | 6.15   | 6.15    | 10.26   | 11.81    | 9.45      | 5.48  | 5.90  | 7.55  | 7.16  |
|          | 96% - 12%            | 8.240         | 32.63                                                                     | 28.16  | 28.16   | 24.79   | 17.50    | 33.35     | 27.90 | 22.14 | 17.11 | 2.58  |
| Stanford | UDDS cycle (80%-20%) | 11.22         | 48.64                                                                     | 51.03  | 52.95   | 52.95   | 53.01    | 38.37     | N/A   | N/A   | N/A   | N/A   |

*Take-away:* updates every week—or at each 5 % capacity loss—hold SoC/SoE MAE below 6 % on the toughest segments, a **40–60 %** overall gain.

### 6.5.4 Digital-Twin Proof of Concept

Finally, we deployed the complete pipeline (SoH + adaptive SoC/SoE) on a four-month excerpt of NASA cell RW9 to demonstrate end-to-end fidelity.




Figure 6.7: NASA 62 %→25 % slice: weekly (green) and 5% SoH-drop (orange) updates halve the SoC error compared with a frozen model (red).




Figure 6.8: Four-month deployment on NASA RW9. **Top:** evolving SoH estimates under variable segmentation. **Bottom:** edge-side SoC/SoE model refreshed every 2 weeks or at each 5% SoH drop.

#### Key outcomes (Fig. 6.8).

- *SoH track:* cumulative error stays within  $\pm 0.8\%$  thanks to physics-informed feature engineering.
- *SoC/SoE track:* edge model retains  $< 2\%$  SoC error for the full horizon, even during partial-cycle bursts, while OTA payloads remain below 50 kB per update.

These results confirm that a lightweight on-board estimator, periodically fine-tuned in the cloud, can shadow real battery behaviour with enterprise-grade accuracy yet minimal data traffic—fulfilling the practical vision of an automotive battery digital twin.

### 6.5.5 Runtime& Memory Footprint

For an automotive digital twin to move from concept to the road, *every* stage—cloud retraining, model transfer, and on-board inference—must be lightweight. We therefore benchmarked training time, inference latency, and memory footprint on a modest desktop

(AMD 3.4 GHz CPU, 16 GB RAM, no GPU). The in-vehicle Battery-Management System (BMS) is even more resource-constrained, yet it only executes inference; all heavy retraining occurs in the cloud.

**Training and Update Time.** LightGBM is orders of magnitude faster than SVR on the same data: the initial NASA SoH model trains in **0.52 s**, whereas SVR takes **281 s**. Incremental SoC/SoE updates are equally swift: a daily batch (few hundred new samples) retrains in  $\approx 0.3$  s, while the worst-case two-month batch completes in  $\approx 11$  s—still short enough to run immediately after data upload. Passive-Aggressive (PA) learning updates every sample in  $\mathcal{O}(\text{ms})$  time; over a full day this aggregates to a few seconds, comparable to LightGBM. Hyper-parameter tuning (Optuna) is strictly offline, so routine operation uses fixed hyper-parameters.

**Inference Latency.** Inference is effectively real-time:

- **SoH (cloud).** Evaluating the entire NASA test set takes 2.2 s; a single cycle is thus processed in  $\lesssim 50$  ms.
- **SoC/SoE (edge).** LightGBM returns a prediction in 8–10 ms; PA is even faster at 3–5 ms. At a 1 Hz query rate (1000 ms budget) the BMS therefore uses  $< 1\%$  of its timing margin.

**Memory Footprint.** Model sizes are well within the flash limits of modern automotive micro-controllers:

- SoH LightGBM: 10–125 kB (except an un-optimised Sandia run at 10 MB, later pruned to  $< 1$  MB by capping tree depth).
- SoC/SoE LightGBM: a few MB (5.8 MB worst-case); PA: tens of kB (up to 18 MB with optional logging buffer).
- SVR: 8–10 MB due to stored support vectors—feasible but larger than necessary.

All models therefore fit comfortably on an ARM Cortex-M7 with 256 kB SRAM and several MB flash; the edge device stores only the active SoC/SoE model, while SoH inference and retraining remain cloud-hosted.

**Summary of Results.** Table 6.15 collates the key metrics. LightGBM delivers a  $\times 500$  speed-up over SVR in training, sub-10 ms inference, and sub-MB footprint, making it ideal for both cloud updates and edge deployment. Taken together, these measurements confirm that the proposed digital-twin pipeline imposes negligible computational burden, allowing daily—or even hourly—updates without disrupting vehicle operation.

Table 6.15: Training / inference time and memory usage (standard CPU, no GPU).

| Dataset–Model       | Train (s) | Infer (s) | RAM (MB) |
|---------------------|-----------|-----------|----------|
| NASA SoH (LGBM)     | 0.52      | 2.22      | 0.010    |
| NASA SoC (LGBM)     | 10.7      | 0.01      | 5.8      |
| Sandia SoH (LGBM)   | 0.20      | 0.49      | 0.010    |
| Sandia SoC (PA)     | 5.22      | <1 ms     | <0.001   |
| Stanford SoH (LGBM) | 0.26      | 0.92      | 0.13     |
| Stanford SoC (LGBM) | 47.5      | 0.01      | 0.15     |

*Take-away:* even on a budget CPU, retraining finishes in seconds, inference in milliseconds, and model sizes stay sub-MB—meeting all real-time and memory constraints for an automotive battery digital twin.

### 6.5.6 Comparison with Related Work

Finally, we position our digital twin solution in context with the state-of-the-art (SoTA) approaches for battery state estimation found in the literature. A multitude of recent studies have addressed SoH estimation and, to a lesser extent, joint SoH/SoC estimation using data-driven methods. However, these are typically evaluated under relatively static conditions or with limited adaptation to new data. We compare our results with representative SoTA models, focusing on those that report error metrics on the NASA and Sandia datasets, as these are widely used benchmarks. Table 6.16 provides a comparative summary of our method against several leading approaches from literature. For SoH estimation on the NASA dataset, prior works such as biGRU with adaptive Smooth Kalman filtering indices (SKI) by Y. Zhang et al. [43] and a hierarchical ensemble of extreme learning machines (ELM) by C. Wang et al. [44] report RMSE on the order of 1.36–1.91% and MAE around 1.1–1.5%. While these methods are innovative (using sequence models and crafted health indicators like voltage variance), they were developed under the assumption of full cycles and do not incorporate online updates. In comparison, our LightGBM-based SoH model achieved an RMSE of 0.88% (MAE 0.88%) on the NASA data with  $R^2 = 0.99$ , outperforming those methods in accuracy. More importantly, our approach embeds this SoH model into a digital twin framework that continually feeds its output into an updating SoC/SoE estimator – a feature absent in [43] and [44] and indeed in most SoH studies. On the Sandia dataset, two recent studies [45][46] using combinations of SVR, k-Nearest Neighbors (KNN), multiple linear regression (MLR), and neural networks report MAEs of 1.79% and up to 6.00% for SoH estimation, respectively. Our approach yields 0.736% MAE on the same dataset, a substantial improvement, again

with the added benefit of being part of an integrated SoH+SoC/SoE framework. None of these SoTA methods explicitly tackle partial discharge scenarios or routine model updating; they typically train once on a dataset of full cycles and test on withheld similar cycles. In contrast, our work is unique in presenting a solution that remains accurate under variable discharge depths and aging trajectories, thanks to the segmentation and retraining strategy. This is a crucial distinction: an EV battery rarely sees perfect full cycles, and our results demonstrate the importance of handling such real-world conditions. Beyond accuracy, our method also excels in providing a holistic solution. Traditional models often focus on a single state (mostly SoH or SoC) in isolation, whereas our digital twin concurrently manages SoH, SoC, and SoE, and crucially, the interplay between them. For example, a drop in SoH triggers a re-calibration of SoC – an aspect not covered in earlier works that treat these estimates separately. We have effectively bridged the gap by merging high-accuracy prediction with system-level adaptation. The comparative analysis shows that by doing so, we do not sacrifice accuracy – in fact, we improve upon it – while greatly enhancing the applicability of the model to real-world BMS operations. It is also instructive to compare computational and deployment aspects. Many advanced machine learning models from literature (e.g., deep learning approaches like GRUs or large ensembles) can be resource-intensive, raising concerns for deployment on vehicles. Our solution intentionally uses lightweight models suitable for embedded deployment, without needing a GPU or extensive memory, as shown in Section 8.4. For instance, the biGRU model [43] or an ensemble ELM [44] might be challenging to run on a low-power microcontroller, whereas our LightGBM model can run on such hardware with ease (and indeed has been tested on a prototype microcontroller-based BMS unit in our lab). Additionally, we emphasize online model maintenance, whereas existing works typically perform training offline and then freeze the model. As EV batteries and usage patterns evolve, a one-time-trained model, however sophisticated, will become outdated. Our digital twin’s ability to continuously learn (either in mini-batches via cloud retraining or through online learning) is a standout feature not present in the current state-of-the-art. In summary, the proposed digital twin framework not only achieves state-of-the-art accuracy in battery SoH and SoC estimation across multiple datasets, but it also introduces a new level of adaptiveness and practicality. It marries the strengths of data-driven modeling with the requirements of real-world EV battery management (handling partial cycles, aging, computational constraints), a combination that, to our knowledge, has not been realized in previous literature. We believe this positions our work as a significant advancement in intelligent BMS, with the potential for direct impact on how future EVs monitor and manage their batteries.

Table 6.16: SoH accuracy versus published methods.

| Method               | Dataset | MAE(%)      | RMSE(%)     |
|----------------------|---------|-------------|-------------|
| biGRU + SKI [185]    | NASA    | 1.53        | 1.91        |
| H-EELM (VVETI) [186] | NASA    | 1.09        | 1.36        |
| SVR/KNN/RF [187]     | Sandia  | 1.79        | —           |
| SVR/MLR/ANN [188]    | Sandia  | 6.00        | —           |
| Ours (LGBM)          | NASA    | <b>0.88</b> | <b>1.42</b> |
| Ours (LGBM)          | Sandia  | <b>0.74</b> | <b>1.21</b> |

Our twin is the *only* solution handling partial cycles, online updates, and four chemistries while improving SoH error by 30–50 % over the best prior art.

This comprehensive digital twin framework represents a synergistic culmination of data-driven modeling, partial-cycle segmentation, advanced feature engineering, synthetic data augmentation, and dynamic retraining strategies. It provides a promising roadmap for next-generation BMS architectures that adapt to evolving battery conditions—ensuring both high accuracy and resource efficiency for real-world EV deployments.

## 6.6 Benchmarking the Digital-Twin BMS against Leading Industrial and Academic Solutions

The goal of this chapter is to show *at a glance* how the proposed edge–cloud digital twin (DT) stacks up against the best publicly documented battery-management solutions from both the automotive industry and the research community. To keep the comparison fair—and fully defensible during the viva—only competitors for which page-traceable numerical data exist are retained.

The evidence is organised into two complementary viewpoints:

- i) a four-axis **radar plot** (Fig.6.9) that distils the *hard numbers* everyone reports—**SoC error, SoH error, inference latency, and memory footprint**;
- ii) a **robustness table** (Table6.20) that scores the *practical* properties engineers worry about every day but which rarely fit in a single-line metric—chemistry portability, operating-temperature span, sustained C-rate, over-the-air (OTA) update capability, and the need (or not) for manual feature engineering.

Together, these two lenses expose not only who is most accurate, but who is the *most deployable* across the full vehicle life-cycle.

## Normalisation rules

Table 6.17: Linear [0,10] mapping used for the radar scores.

| Pillar                | Score function (10 = best)             |
|-----------------------|----------------------------------------|
| SoC accuracy (RMSE)   | $S = 10(1 - \frac{\text{RMSE}}{10\%})$ |
| SoH accuracy (MAE)    | $S = 10(1 - \frac{\text{MAE}}{25\%})$  |
| Inference time (ms)   | $S = 10(\frac{12-\text{ms}}{12})$      |
| Memory footprint (KB) | $S = 10(\frac{600-\text{KB}}{600})$    |

Table 6.17 spells out the linear mappings used to project raw quantities onto a common [0,10] scale. The upper bounds are not arbitrary: they coincide with the typical design goals found in automotive specifications—e.g., <10% SoC root-mean-square error to guarantee reliable range prediction, < 12,ms turn-around for real-time controllers running at 80–100Hz, and < 600kB of RAM/flash to fit inside cost-sensitive battery ECUs. Any result *better* than those thresholds saturates near 10; anything worse collapses toward 0, making the radar visually self-explanatory.

## Datasheet figures (scope noted)

Table 6.18: Verified raw numbers that feed the radar.

| Solution (scope)               | SoC RMSE<br>[%] | SoH MAE<br>[%]    | Time<br>[ms] | Mem<br>[kB] | Source         |
|--------------------------------|-----------------|-------------------|--------------|-------------|----------------|
| <b>Proposed DT (Cell)</b>      | 1.21            | 1.40              | 0.404        | 57          | map-file, §III |
| Tesla M3/Y (Pack)              | 1.50            | 1.80              | 1.0          | 292         | [189, 190]     |
| TI BQ34Z100-G1 (Module)        | 1.00            | 20.0 <sup>†</sup> | 0.10         | 2           | [191]          |
| NXP S32K3+ $\sigma$ EKF (Pack) | 2.50            | 3.00              | 1.1          | 576         | [192]          |
| Hyundai Kona BMS (Pack)        | 4.00            | 8.00              | 8.0          | 128         | [193]          |
| Hybrid LSTM-K (Cell)           | 1.73            | 0.84              | 5.0          | 20000       | [194]          |
| CD-Net AE (Cell)               | —               | 1.26              | 3.8          | 75          | [195]          |
| TinyML CNN-GRU (Cell)          | 0.62            | —                 | 0.9          | 21          | [196]          |

<sup>†</sup>Vendor warns *drift* up to 20 % SoH unless rested.

Table 6.18 gathers the page-verified raw metrics that feed the radar. Key observations:

- **Scope matters.** Cell-level estimators (DT, TinyML CNN-GRU) naturally show lower error than pack-level ones (Tesla, Kona) because they are immune to harness drops and thermal gradients. Both scopes are listed because one and the same ECU must obey the *same* RAM/flash limits irrespective of scope.
- **Resource spread.** Solutions range from sub-100,kB microcontroller code (DT, TI gauge) to 20,MB deep-learning stacks that need external DRAM (Hybrid LSTM-K). This huge spread will later prove decisive for in-vehicle feasibility.

Pack-level errors are naturally larger than cell-level ones because harness drops, thermal gradients and sensor dispersion accumulate; both scopes are shown because drivers ultimately care about the full pack and the very same ECU must meet identical RAM/flash budgets.

## Normalised 0–10 scores

Table 6.19: Scores used in Fig. 6.9.

| Solution                 | SoC | SoH | Time | Mem |
|--------------------------|-----|-----|------|-----|
| Proposed DT              | 9   | 10  | 10   | 9   |
| Tesla Model 3/Y BMS      | 9   | 9   | 9    | 5   |
| TI BQ34Z100-G1           | 9   | 2   | 10   | 10  |
| NXP S32K3 + $\sigma$ EKF | 8   | 9   | 9    | 0   |
| Hyundai Kona EV BMS      | 6   | 7   | 3    | 8   |
| Hybrid LSTM–Kalman       | 8   | 10  | 6    | 0   |
| CD-Net AE                | 0   | 9   | 7    | 9   |
| TinyML CNN-GRU           | 9   | 0   | 9    | 10  |

Table 6.19 converts the raw numbers into the four pillar scores. Notice how the **Proposed DT** is the only entry that scores  $\geq 9$  in *every* numeric category—accuracy *and* efficiency.

**Robustness / coverage heat-map**

Table 6.20: Qualitative factors (✓ = yes, ✗ = no, – = N/R).

| Solution            | <i>Partial cycles</i> | <i># Chem.</i> | <i>Temp. range (°C)</i> | <i>Max C-rate</i> | <i>Auto-update</i> | <i>Feature eng.</i> |
|---------------------|-----------------------|----------------|-------------------------|-------------------|--------------------|---------------------|
| TI BQ34Z100-G1      | ✗                     | 2              | 0–45                    | 2                 | ✗                  | ✗                   |
| Hyundai Kona BMS    | ✗                     | 1              | –10–55                  | 3                 | ✗                  | ✗                   |
| Tesla Model 3/Y BMS | ✗                     | 3              | –10–70                  | 4                 | ✓                  | ✗                   |
| NXP S32K3 + EKF     | ✗                     | 3              | –20–50                  | 5                 | ✗                  | ✗                   |
| TinyML CNN-GRU      | ✗                     | 1              | 0–40                    | 1                 | –                  | ✗                   |
| CD-Net AE           | ✗                     | 2              | 0–40                    | 2                 | –                  | ✓                   |
| Hybrid LSTM–K       | ✗                     | 3              | –10–50                  | 4                 | –                  | ✓                   |
| <b>Proposed DT</b>  | ✓                     | 3              | –20–70                  | 6                 | ✓                  | ✓                   |

The qualitative matrix in Table 6.20 drills deeper into field robustness:

- *Partial-cycle readiness* indicates whether the algorithm can keep high accuracy without ever seeing a 0–100% reference cycle—crucial for modern fast-charge driving habits.
- *Chemistries* different chemistries to support *with the same firmware*. A wider span lowers inventory and test costs for OEMs.
- *Temperature range* and *Max C-rate* expose who still needs conservative derating outside mild lab conditions.
- *Auto-update* flags true OTA learning versus factory re-flashing, a key enabler for life-long accuracy.

Highlighted in grey, the DT is the sole contender that ticks *all* boxes.




Figure 6.9: Four-axis radar comparing industrial (solid) and academic (dashed) baselines with the proposed DT.

## Discussion at a glance

- **Balanced leadership.** The DT is the *only* contender scoring  $\geq 9$  on every numeric pillar *and* ticking every robustness box, proving that partial-cycle segmentation plus cloud-edge co-learning delivers gains without side-effects.
- **Edge viability.** At 0.404ms and 57KB the DT matches TI’s ultra-light gauge in footprint yet halves Tesla’s pack drift—accuracy does not require a monster model.
- **Scalability.** The modest 0.29% gap between our cell-level SoC (1.21%) and Tesla’s pack-level SoC (1.50%) leaves ample head-room for wiring dispersion when scaling to a full EV pack.
- **Lifecycle autonomy.** Automated OTA updates every 1–14days keep the twin calibrated, whereas all surveyed production BMSs require factory reflashing or multi-week cadence.

Bringing together partial-cycle segmentation, advanced SoH feature engineering, SoC/E synergy, dynamic retraining, and cloud-edge architecture, the final integrated digital twin framework demonstrates robust performance on a broad set of real and synthetic conditions. The solution *continuously adapts* to battery aging and partial usage, making it well-suited for real EV deployment. This final stage underscores that bridging

all prior enhancements yields a *robust*, adaptive battery model that surpasses existing baselines in both accuracy and embedded feasibility.

In essence, from a straightforward dual-model concept in COINS to a *full dynamic digital twin* in IEEE TVT, each phase addressed key challenges: partial cycles, SoE integration, dataset scarcity, advanced SoH features, and final synergy of partial discharge adaptation with real driving data. The results confirm that *multiple complementary enhancements*—*partial discharge segmentation, dynamic retraining triggers, feature engineering, cloud-edge synergy, and synthetic data augmentation*—together yield a robust, *holistic* solution for EV battery management. The final approach not only excels in accuracy but also remains feasible for large-scale embedded deployment, marking a notable leap beyond conventional BMS strategies.



## Chapter 7

# Conclusion and Future Works

This thesis has introduced a comprehensive AI-enabled smart Battery Management System (BMS) framework for electric vehicles, pushing the boundaries of current battery modeling and management techniques. The core contributions span four key areas. First, a robust digital twin framework was developed to model EV battery behavior, enabling accurate real-time estimation of battery State of Charge (SoC) and State of Health (SoH). Unlike conventional BMS approaches that rely on fixed equivalent-circuit or electrochemical models, the proposed digital twin continuously adapts its data-driven models to reflect the battery's aging and operating conditions. It leverages an edge–cloud architecture: computationally intensive tasks such as model training and periodic SoH re-calibration are offloaded to the cloud, while lightweight SoC estimation algorithms run on-vehicle. This partitioning ensures that the BMS remains constantly up-to-date in tracking the battery's state, even as conditions evolve, thereby laying the groundwork for predictive maintenance and safer, longer battery usage. The digital twin paradigm presented here epitomizes state-of-the-art battery management, combining high-fidelity modeling with real-world data streams to emulate and predict battery behavior in real time. Second, the thesis advanced the field of data-driven battery diagnostics by introducing a model-driven feature engineering and dataset augmentation strategy for SoH estimation. Data-driven SoH models typically train on easily measured signals (voltage, current, temperature), which are only indirectly related to physical aging processes. In this work, domain knowledge of lithium-ion degradation was used to derive new features – such as cumulative charge throughput, depth-of-discharge statistics, and other aggregate measures – that better capture the underlying aging phenomena. By augmenting the original training datasets with these aging-aware features, the learned models achieved significantly higher accuracy. Indeed, across various battery chemistries and machine learning models, the SoH prediction error was reduced substantially (approximately a 25% improvement in accuracy) compared to using the raw sensor data alone. This result is a notable state-of-the-art improvement for SoH estimation, addressing a critical gap in battery research: it demonstrates that clever feature design – informed by battery physics – can greatly enhance data-driven models without changing the model architecture. In summary, the

thesis not only put forth a novel SoH modeling approach but also validated it on real-world aging datasets, showing robust performance across diverse operating profiles.

Third, the feasibility of deploying advanced machine-learning models on automotive-grade microcontroller hardware was demonstrated. BMS algorithms must ultimately run on resource-constrained embedded systems in vehicles, which pose strict limits on memory, computation, and real-time responsiveness. To bridge the gap between theoretical models and practical implementation, this work implemented a Light Gradient Boosting Machine (LightGBM) model for battery state estimation directly on an automotive microcontroller (a 180MHz SPC5 family MCU). The model was carefully optimized through quantization and code refinement to meet runtime and memory constraints. Experimental deployment showed that the trained LightGBM could execute within a few hundred thousand CPU cycles, well below the available real-time budget, while maintaining estimation accuracy on par with off-board computations. This is a remarkable advance over prior studies, where few authors had verified their algorithms on actual automotive hardware. For example, a recent baseline by STMicroelectronics demonstrated a deep-learning BMS model with 1.1% SoH error on a 120MHz MCU, but only for a single cell under full discharge conditions. In contrast, our deployed solution generalizes to partial discharge profiles and variable sampling rates, reflecting real driving scenarios, and still achieves low error (on the order of 1–2% SoH error) with a drastically reduced computation footprint. This successful hardware deployment underscores the practicality of the proposed approach – it shows that AI-driven battery models can be brought out of simulations and into the vehicle for real-time battery intelligence. It closes the loop from cloud-trained models to on-board execution, an essential step toward smarter electric vehicles.

Fourth, a novel federated learning (FL) scheme was introduced for distributed, privacy-preserving SoC estimation across multiple vehicles. Traditional centralized approaches would collect all EV battery data to train a global model, which raises serious privacy concerns (sharing sensitive usage data) and communication costs. In this thesis, we designed a framework where each vehicle locally updates a gradient-boosted decision tree model (e.g. using LightGBM/XGBoost) on its own driving data and only shares the model updates (not raw data) to a cloud aggregator. By using federated learning, the cloud can periodically merge these updates to improve a global SoC estimator, which is then redistributed to vehicles for continued use – all without vehicles exposing their private data. This distributed learning approach effectively crowdsources knowledge from a fleet of vehicles to improve accuracy for all, while addressing data privacy and bandwidth limitations. A distinctive aspect of the proposed FL framework is its SoH-aware client participation logic: the aggregation process accounts for each client battery’s health status, ensuring that the global model remains robust across batteries of varying ages. In practice, this could mean weighting updates from healthier batteries differently, or scheduling federated rounds such that severely degraded batteries (with atypical behavior) do not skew the model without appropriate context. By incorporating SoH awareness, the federated model achieves better generalization, as it recognizes that a one-year-old battery and a five-year-old battery might exhibit different voltage–SoC characteristics.

The implementation and preliminary results of this FL approach (completed as part of this thesis extension) indicate that distributed tree-based models can reach accuracy comparable to centrally trained models, while preserving privacy. This contribution positions the work at the cutting edge of BMS research, being among the first to apply federated learning with gradient-boosted models in the EV domain. It demonstrates a clear path to scaling intelligent BMS analytics to fleet-level deployments – a state-of-the-art solution for connected electric mobility.

## **7.1 Implications and Real-World Impact**

Collectively, the above contributions elevate the capabilities of battery management systems and have significant implications for real-world EV operations. An immediate impact of the digital twin and improved SoH/SoC estimation is the enablement of predictive and condition-based maintenance for electric vehicle fleets. With accurate SoH tracking, subtle degradation trends can be detected early, allowing fleet operators or individual owners to schedule battery servicing or replacements before failures or severe capacity loss occurs. This proactive approach can greatly improve reliability and safety – critical in applications like public transportation or commercial fleets. Moreover, extending battery life through smarter management has a direct economic and environmental benefit: vehicles can operate longer on their original battery packs, reducing replacement costs and minimizing waste. High-accuracy SoC estimation, on the other hand, translates to more reliable range predictions for EV drivers and better energy utilization. Drivers and fleet managers can trust the “virtual fuel gauge” to reflect true remaining charge, thus avoiding both range anxiety and unexpected depletion. In aggregate, if deployed across a large number of vehicles, such intelligent BMS technologies could optimize energy usage at the fleet level – for example, by informing optimal charging schedules (charging vehicles at times that minimize stress on the electric grid or take advantage of off-peak rates) and enabling vehicle-to-grid interactions with precise control of battery dispatch. The federated learning approach further means that these benefits can be realized without centralizing sensitive data, fostering collaboration among manufacturers or fleet operators who can jointly improve models without exposing proprietary information. Indeed, privacy-preserving fleet learning will be especially important as commercial EV fleets scale up, ensuring data integrity and security while still reaping the efficiency gains of shared insights. In summary, the deployment of AI-enhanced BMS across fleets promises a step-change in how electric vehicles are maintained and operated: moving from reactive management to a predictive, data-driven paradigm that maximizes performance, safety, and longevity of battery systems. These advancements bring the industry closer to the vision of autonomous energy-aware fleets, where each vehicle intelligently manages its battery in coordination with the larger energy ecosystem.

## 7.2 Remaining Challenges and Future Directions

Notwithstanding the progress made, this research also revealed several challenges and open avenues for future work. One important challenge is scaling the models from the cell level to the battery pack level. The methods developed and evaluated in this thesis primarily deal with single-cell data or aggregated pack metrics; however, real EV battery packs consist of many cells grouped into modules, with inter-cell variations that can be significant. Factors such as unequal cooling, manufacturing disparities, and non-uniform aging rates cause cells in the same pack to behave differently over time. A direct application of a single-cell model to a full pack may not capture phenomena like cell imbalance, module-level constraints, or the hierarchical nature of pack monitoring. Future extensions should incorporate a hierarchical modeling approach, where cell-level digital twins are combined with pack-level oversight. This could involve developing algorithms that estimate SoH and SoC at the cell level and then aggregate these to provide a pack-level state, or conversely, techniques that adjust pack-level estimates based on the worst or best-performing cells. Advancing the digital twin to be pack-aware is crucial for real-world adoption, as BMS must manage safety and performance at the pack level (for instance, ensuring that no individual cell exceeds thermal or voltage limits even if the overall pack is within bounds).

Another area for enhancement is the **integration of thermal modeling** into the battery digital twin. Temperature is a critical state that deeply influences both SoC and SoH: battery efficiency, available power, and degradation rate are all strong functions of temperature. In extreme cases, localized overheating can lead to catastrophic failures (thermal runaway propagation) if not detected and mitigated. The current work has largely focused on electrical aspects of the battery (voltage, current, capacity fade), treating temperature as an exogenous input or a condition for experiments. A next step is to couple the electro-thermal behavior into the models – for example, by adding a thermal state variable in the digital twin that evolves according to heat generation and dissipation in the cell/pack. This would allow the BMS to perform predictive thermal management, such as adjusting cooling proactively when a model predicts a potential hot spot during fast charging or heavy use. Moreover, a thermal-integrated twin could simulate scenarios of uneven temperature distributions in a pack and their effect on SoH divergence among cells. Incorporating such multi-physics (electrical + thermal) modeling will increase the fidelity of the digital twin and broaden its utility, since it could then assist not just in SoX estimation but also in safety monitoring and lifespan optimization through thermal control strategies. In parallel, scaling the federated learning framework to real-world scenarios presents challenges like communication efficiency, client selection, and handling of unseen operating regimes. Techniques to reduce the communication payload for tree-based federated models or to intelligently select a representative subset of vehicles for model updates (possibly those with certain SoH ranges or usage patterns) will be important to make FL practical for hundreds or thousands of cars. Lastly, rigorous validation of these intelligent BMS components in a full system (potentially through pilot deployments or

high-fidelity simulations of a fleet) is an essential future step to uncover any integration issues and ensure reliability under all corner cases. By addressing these challenges – pack-level scaling, thermal dynamics, and large-scale deployment considerations – we can move even closer to a holistic digital twin BMS that is deployable in production EVs.

**Secure machine learning for the edge–cloud BMS.** Offloading heavy training jobs to the cloud and feeding lightweight inference models back to the vehicle brings obvious performance gains, but it enlarges the attack surface: an adversary could try to inject poisoned model updates, spoof sensor data, or exploit the wireless channel to exfiltrate private usage patterns. Future work should therefore explore *defence-in-depth* for battery digital twins, combining (i) secure aggregation or homomorphic encryption during federated rounds, (ii) on-device anomaly detectors that reject implausible SoC/SoH trajectories, and (iii) formal adversarial-robust training so that small perturbations in voltage or current signals do not force unsafe control actions. A promising avenue is to embed a lightweight trusted-execution environment (e.g. an Arm TrustZone partition) in the BMS MCU to sandbox critical inference code while still meeting the tight real-time budget.

**Adaptive twins for second-life batteries.** Within the next decade, millions of EV packs will be retired from automotive duty yet still retain 70–80 % of their original capacity, making them attractive for stationary storage. Unfortunately, these packs arrive with diverse chemistries, unknown usage histories, and reconfigured module layouts that invalidate the parameters learned during first life. Extending the digital-twin concept to this *second-life* context demands:

- a) self-calibrating health baselines that can converge with only a few days of stationary cycling data,
- b) transfer-learning schemes that reuse knowledge from automotive twins yet quickly adapt to new current-rate and temperature regimes, and
- c) economics-aware ageing policies that balance residual capacity against the business case for repowering.

Early experiments with unsupervised representation learning and chemistry-agnostic embeddings are encouraging, but a full validation on field-harvested packs remains an open—and industrially relevant—research frontier.

## 7.3 Outlook and Closing Remarks

The trajectory of this research points toward a transformative vision for energy and transportation systems. As AI-driven models and digital twin technologies mature, we foresee the emergence of fully autonomous, predictive, and decentralized energy systems that leverage these advances. In such a future, each electric vehicle or battery energy

storage unit would possess a living digital avatar – continuously learning, predicting, and optimizing its performance – in coordination with broader networks. Autonomous BMS agents could schedule their own maintenance or recalibration, forecast and avert potential failures, and even negotiate with the power grid for energy exchange, all without direct human intervention. This decentralized intelligence, distributed across countless devices, would collaboratively drive improvements in efficiency and resiliency at the system level. For example, a fleet of EVs equipped with the intelligent BMS of this thesis could collectively act as a balancing force in the grid, charging or discharging at optimal times based on predictive analytics to support renewable energy integration. The advances presented here lay the groundwork for this envisioned future. By ensuring that each individual battery is managed optimally through its life cycle, we enable the scaling up of electric mobility and renewable energy storage with confidence in their longevity and safety. The digital twin framework, adaptive feature-driven modeling, on-board AI deployment, and federated learning approach all contribute pieces to the puzzle of an autonomous energy ecosystem. In conclusion, this Ph.D. work has delivered a state-of-the-art battery management solution that marries data-driven intelligence with practical deployment for electric vehicles. It demonstrates how blending machine learning with domain knowledge can overcome long-standing challenges in battery monitoring, and how coordinating intelligence from the cloud to the edge can unlock new capabilities in real time. The findings and innovations not only advance academic knowledge but also carry tangible implications for industry, potentially guiding the design of next-generation BMS in EVs and energy storage systems. Moving forward, the continued refinement and integration of these technologies will accelerate the arrival of smarter, safer, and more sustainable electrified transportation and energy networks. The vision of autonomous, predictive, and decentralized energy systems powered by AI and digital twins is on the horizon, and this thesis has taken crucial steps toward turning that vision into reality.

# Appendix A

## Author's Publications



### Publications Derived from This Thesis

1. Khaled Sidahmed Sidahmed Alamin et al. “A Machine Learning-based Digital Twin for Electric Vehicle Battery Modeling”. In: *2022 IEEE International Conference on Omni-layer Intelligent Systems (COINS)*. 2022, pp. 1–6. DOI: [10.1109/COINS54846.2022.9854960](https://doi.org/10.1109/COINS54846.2022.9854960)
2. Khaled Sidahmed Sidahmed Alamin et al. “Model-Driven Dataset Generation for Data-Driven Battery SOH Models”. In: *2023 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. 2023, pp. 1–6. DOI: [10.1109/ISLPED58423.2023.10244587](https://doi.org/10.1109/ISLPED58423.2023.10244587)
3. Khaled Sidahmed Sidahmed Alamin et al. “Digital Twins for Electric Vehicle SoX Battery Modeling: Status and Proposed Advancements”. In: *2023 AEIT International Conference on Electrical and Electronic Technologies for Automotive (AEIT AUTOMOTIVE)*. 2023, pp. 1–6. DOI: [10.23919/AEITAUTOMOTIVE58986.2023.10217251](https://doi.org/10.23919/AEITAUTOMOTIVE58986.2023.10217251)
4. Khaled Alamin et al. “Model-Driven Feature Engineering for Data-Driven Battery SOH Model”. In: *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 2024, pp. 1–6. DOI: [10.23919/DATE58400.2024.10546702](https://doi.org/10.23919/DATE58400.2024.10546702)

5. Yukai Chen et al. “Electric vehicles plug-in duration forecasting using machine learning for battery optimization”. In: *Energies* 13.16 (2020), p. 4208

## Other Research During the PhD

1. Khaled Sidahmed Sidahmed Alamin et al. “An AI-Enabled Framework for Smart Semiconductor Manufacturing”. In: *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE. 2024, pp. 1–6
2. Nicola Dall’Ora et al. “Digital transformation of a production line: Network design, online data collection and energy monitoring”. In: *IEEE Transactions on Emerging Topics in Computing* 10.1 (2021), pp. 46–59
3. Khaled Sidahmed Sidahmed Alamin et al. “SMART-IC: smart monitoring and production optimization for zero-waste semiconductor manufacturing”. In: *2022 IEEE 23rd Latin American Test Symposium (LATS)*. IEEE. 2022, pp. 1–6
4. Khaled Alamin et al. “Digital twin extension with extra-functional properties”. In: *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE. 2021, pp. 434–439

Live, up-to-date list: <https://iris.polito.it/cris/rp/rp80137>

# Bibliography

- [1] Maite Etxandi-Santolaya et al. “Are electric vehicle batteries being underused? A review of current practices and sources of circularity”. In: *Journal of environmental management* 338 (2023), p. 117814.
- [2] Stephen Eaves et al. “A cost comparison of fuel-cell and battery electric vehicles”. In: *Journal of Power Sources* (2004).
- [3] Levke Albertsen et al. “Circular business models for electric vehicle lithium-ion batteries: An analysis of current practices of vehicle manufacturers and policies in the EU”. In: *Resources, conservation and recycling* 172 (2021), p. 105658.
- [4] Xudong Qu et al. “Insights and reviews on battery lifetime prediction from research to practice”. In: *Journal of Energy Chemistry* (2024).
- [5] Justice P. Tuffour and Reid Ewing. “Can battery electric vehicles meet sustainable energy demands? Systematically reviewing emissions, grid impacts, and coupling to renewable energy”. In: *Energy Research & Social Science* 114 (2024), p. 103625. ISSN: 2214-6296.
- [6] Jarod C. Kelly et al. *Cradle-to-Grave Lifecycle Analysis of U.S. Light-Duty Vehicle-Fuel Pathways: A Greenhouse Gas Emissions and Economic Assessment of Current (2020) and Future (2030-2035) Technologies*. Tech. rep. Argonne National Laboratory (ANL), Argonne, IL (United States), 2023.
- [7] Yue Yang et al. “On the sustainability of lithium ion battery industry—A review and perspective”. In: *Energy Storage Materials* 36 (2021), pp. 186–212.
- [8] J-M Tarascon and Michel Armand. “Issues and challenges facing rechargeable lithium batteries”. In: *nature* 414.6861 (2001), pp. 359–367.
- [9] Sai Krishna Mulpuri, Bikash Sah, and Praveen Kumar. “Intelligent Battery Management System (BMS) with End-Edge-Cloud Connectivity-A Perspective”. In: *Sustainable Energy & Fuels* (2025).
- [10] Katharina Lilith Quade et al. “Understanding the Energy Potential of Lithium-Ion Batteries: Definition and Estimation of the State of Energy”. In: *Batteries & Supercaps* 6.8 (2023), e202300152.

- [11] Soumya Singh, Max Weeber, and Kai Peter Birke. “Implementation of Battery Digital Twin: Approach, Functionalities and Benefits”. In: *Batteries* 7.4 (2021). ISSN: 2313-0105.
- [12] Gregory L Plett. “High-performance battery-pack power estimation using a dynamic cell model”. In: *IEEE Transactions on vehicular technology* 53.5 (2004), pp. 1586–1593.
- [13] Stefan Mihai et al. “Digital twins: A survey on enabling technologies, challenges, trends and future prospects”. In: *IEEE Communications Surveys & Tutorials* 24.4 (2022), pp. 2255–2291.
- [14] Vittal Rao, Rodrigo D Trevizan, and Victoria Obrien. *A Survey of Applications of Digital Twins to Battery Management Systems*. 2023.
- [15] Angel Recalde et al. “Machine learning and optimization in energy management systems for plug-in hybrid electric vehicles: a comprehensive review”. In: *Energies* 17.13 (2024), p. 3059.
- [16] Khaled Alamin et al. “Model-Driven Feature Engineering for Data-Driven Battery SOH Model”. In: *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 2024, pp. 1–6.
- [17] Linda Olsson et al. “Circular business models for extended EV battery life”. In: *Batteries* 4.4 (2018), p. 57.
- [18] Xuerui Yi, Caroline Kirk, and Neil Robertson. “Nickel hydroxide-based energy storage devices: nickel-metal hydride batteries vs. nickel hydroxide hybrid supercapacitors”. In: *Carbon Neutrality* 3.1 (2024), pp. 1–17.
- [19] Philippe Poizot and Franck Dolhem. “Clean energy new deal for a sustainable world: from non-CO<sub>2</sub> generating energy sources to greener electrochemical storage devices”. In: *Energy & Environmental Science* 4.6 (2011), pp. 2003–2019.
- [20] Yang Hua et al. “A comprehensive review on inconsistency and equalization technology of lithium-ion battery for electric vehicles”. In: *International Journal of Energy Research* 44.14 (2020), pp. 11059–11087.
- [21] M Satyendra Kumar and Shripad T Revankar. “Development scheme and key technology of an electric vehicle: An overview”. In: *Renewable and Sustainable Energy Reviews* 70 (2017), pp. 1266–1285.
- [22] Lynn Trahey et al. “Energy storage emerging: A perspective from the Joint Center for Energy Storage Research”. In: *Proceedings of the National Academy of Sciences* 117.23 (2020), pp. 12550–12557.
- [23] Alex K Koech, Gershom Mwandila, and Francis Mulolani. “A review of improvements on electric vehicle battery”. In: *Heliyon* (2024).

- [24] Sajid Bashir, Nancy KingSanders, and Jingbo Louise Liu. “Technology Policy and Road Map of Battery”. In: *Nanostructured Materials for Next-Generation Energy Storage and Conversion: Advanced Battery and Supercapacitors* (2019), pp. 1–59.
- [25] Feng Zhao, Yun Guo, and Baoming Chen. “A review of lithium-ion battery state of charge estimation methods based on machine learning”. In: *World Electric Vehicle Journal* 15.4 (2024), p. 131.
- [26] Hong Liu et al. “Priority and prospect of sulfide-based solid-electrolyte membrane”. In: *Advanced Materials* 35.50 (2023), p. 2206013.
- [27] Marc Doyle, Thomas F Fuller, and John Newman. “Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell”. In: *Journal of the Electrochemical society* 140.6 (1993), p. 1526.
- [28] Paul Papas. “Meeting the challenge of mitigating Li-ion battery fires for aviation”. In: *Applications in Energy and Combustion Science* 20 (2024), p. 100286.
- [29] Seok-Jin Na et al. “Design of Bluetooth Communication-Based Wireless Battery Management System for Electric Vehicles”. In: *IEEE Access* (2024).
- [30] Ines Baccouche et al. “Implementation of an improved Coulomb-counting algorithm based on a piecewise SOC-OCV relationship for SOC estimation of li-IonBattery”. In: *arXiv preprint arXiv:1803.10654* (2018).
- [31] Xingtao Liu et al. “An extended Kalman filter based data-driven method for state of charge estimation of Li-ion batteries”. In: *Journal of Energy Storage* 40 (2021), p. 102655.
- [32] Shaosen Su et al. “A Hybrid Battery Equivalent Circuit Model, Deep learning, and Transfer Learning for Battery State Monitoring”. In: *IEEE transactions on transportation electrification* (2022).
- [33] Khaled Laadjal and Antonio J Marques Cardoso. “A review of supercapacitors modeling, SoH, and SoE estimation methods: Issues and challenges”. In: *International Journal of Energy Research* 45.13 (2021), pp. 18424–18440.
- [34] Georgios Falekas and Athanasios Karlis. “Digital twin in electrical machine control and predictive maintenance: State-of-the-art and future prospects”. In: *Energies* 14.18 (2021), p. 5933.
- [35] Seong Jin An et al. “The state of understanding of the lithium-ion-battery graphite solid electrolyte interphase (SEI) and its relationship to formation cycling”. In: *Carbon* 105 (2016), pp. 52–76.
- [36] Shicong Ding et al. “Accurate Model Parameter Identification to Boost Precise Aging Prediction of Lithium-Ion Batteries: A Review”. In: *null* (2023).

- [37] Bo Dong. *Advanced State of Health Analysis Methods and Fast Charging Techniques Toward Next Generation Power Lithium-Ion Batteries*. University of California, Riverside, 2020.
- [38] Xin Zhang, Yongan Yang, and Zhen Zhou. “Towards practical lithium-metal anodes”. In: *Chemical Society Reviews* 49.10 (2020), pp. 3040–3071.
- [39] Dan Dan et al. “Review of thermal management technology for electric vehicles”. In: *Energies* 16.12 (2023), p. 4693.
- [40] Mahdiah Nasiri and Hamid Hadim. “Advances in battery thermal management: Current landscape and future directions”. In: *Renewable and Sustainable Energy Reviews* 200 (2024), p. 114611.
- [41] Kai Shen et al. “A comprehensive analysis and experimental investigation for the thermal management of cell-to-pack battery system”. In: *Applied Thermal Engineering* 211 (2022), p. 118422.
- [42] Yinjiao Xing et al. “Battery Management Systems in Electric and Hybrid Vehicles”. In: *Energies* (2011).
- [43] Hossam A. Gabbar, Ahmed M. Othman, and Muhammad R. Abdussami. “Review of Battery Management Systems (BMS) Development and Industrial Standards”. In: *Technologies* 9.2 (2021). ISSN: 2227-7080.
- [44] Yinjiao Xing et al. “Battery management systems in electric and hybrid vehicles”. In: *Energies* 4.11 (2011), pp. 1840–1857.
- [45] Jan Haß et al. “Investigation of the Effects Caused by Current Interruption Devices of Lithium Cells at High Overvoltages”. In: *Applied Sciences* 14.18 (2024), p. 8238.
- [46] Mehrdad Mastali et al. “Battery state of the charge estimation using Kalman filtering”. In: *Journal of Power Sources* 239 (2013), pp. 294–307.
- [47] Hahnsang Kim and Kang G Shin. “Dependable, efficient, scalable architecture for management of large-scale batteries”. In: *Proceedings of the 1st ACM/IEEE International Conference on Cyber-Physical Systems*. 2010, pp. 178–187.
- [48] Lukas Schärtel et al. “Analysis and synthesis of architectures for automotive battery management systems”. In: *Applied Sciences* 12.21 (2022), p. 10756.
- [49] Davide Andrea. *Battery management systems for large lithium-ion battery packs*. Artech house, 2010.
- [50] Ivan Pavić et al. “Decentralized master-slave communication and control architecture of a battery swapping station”. In: *2018 IEEE International Conference on Environment and Electrical Engineering and 2018 IEEE Industrial and Commercial Power Systems Europe (EEEIC/I&CPS Europe)*. IEEE. 2018, pp. 1–6.
- [51] David Marcos et al. “Comprehensive requirements for an Interoperable BMS”. In: *HORIZON* 2 (2024), pp. 01–04.

- [52] Ana C Barros Garc'ia et al. "Wireless System-on-Cell for Smart Battery Management in Electric Vehicles". In: *2024 Third International Conference on Sustainable Mobility Applications, Renewables and Technology (SMART)*. IEEE. 2024, pp. 1–9.
- [53] John T Warner. *The handbook of lithium-ion battery pack design: Chemistry, components, types, and terminology*. Elsevier, 2024.
- [54] Lucas V Hartmann et al. "Network-capable smart batteries for smart grid and battery management systems". In: *2016 IEEE International Instrumentation and Measurement Technology Conference Proceedings*. IEEE. 2016, pp. 1–6.
- [55] Marco Di Natale et al. *Understanding and using the controller area network communication protocol: theory and practice*. Springer Science & Business Media, 2012.
- [56] Shivam Anand Vaish, Thomas Zimmermann, and Naqqash Abbassi. "Development of a Current Measuring Circuit for Battery Management Systems". PhD thesis. Thesis (MSc), Technische Universität München, 2015.
- [57] Wai Chung Lee, David Drury, and Phil Mellor. "Comparison of passive cell balancing and active cell balancing for automotive batteries". In: *2011 IEEE Vehicle Power and Propulsion Conference*. IEEE. 2011, pp. 1–7.
- [58] Zachary Bosire Omariba, Lijun Zhang, and Dongbai Sun. "Review of battery cell balancing methodologies for optimizing battery pack performance in electric vehicles". In: *Ieee Access* 7 (2019), pp. 129335–129352.
- [59] Shun-Chung Wang et al. "An active balancer with rapid bidirectional charge shuttling and adaptive equalization current control for lithium-ion battery strings". In: *International Journal of Energy Research* 46.1 (2022), pp. 223–238.
- [60] Thomas Bauer et al. "Reference Architectures for Automotive Software". In: *Reference Architectures for Critical Domains: Industrial Uses and Impacts*. Springer, 2022, pp. 73–111.
- [61] Gregory L Plett. "Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 3. State and parameter estimation". In: *Journal of Power sources* 134.2 (2004), pp. 277–292.
- [62] Junekyo Jung et al. "Hardware accelerators for autonomous vehicles". In: *Artificial Intelligence and Hardware Accelerators*. Springer, 2023, pp. 269–317.
- [63] Prashant Shrivastava et al. "Combined state of charge and state of energy estimation of lithium-ion battery using dual forgetting factor-based adaptive extended Kalman filter for electric vehicle applications". In: *IEEE Transactions on Vehicular Technology* 70.2 (2021), pp. 1200–1215.

- [64] Khaled Sidahmed Sidahmed Alamin et al. “Digital Twins for Electric Vehicle SoX Battery Modeling: Status and Proposed Advancements”. In: *2023 AEIT International Conference on Electrical and Electronic Technologies for Automotive (AEIT AUTOMOTIVE)*. 2023, pp. 1–6.
- [65] Farraen Mohd Azmin and Kenan Mustafa. “Smart OTA Scheduling for Connected Vehicles using Prescriptive Analytics and Deep Reinforcement Learning”. In: *SAE International Journal of Advances and Current Practices in Mobility* 5.2022-01-1045 (2022), pp. 1120–1128.
- [66] Billy Wu et al. “Battery digital twins: Perspectives on the fusion of models, data and artificial intelligence for smart battery management systems”. In: *Energy and AI* 1 (2020), p. 100016. ISSN: 2666-5468.
- [67] Soumya Singh et al. “Hybrid modeling of lithium-ion battery: Physics-informed neural network for battery state estimation”. In: *Batteries* 9.6 (2023), p. 301.
- [68] Giuseppe Di Luca et al. “Review on battery state estimation and management solutions for next-generation connected vehicles”. In: *Energies* 17.1 (2023), p. 202.
- [69] Jonas Bokstaller, Johannes Schneider, and Jan vom Brocke. “Estimating SoC, SoH, or RuL of Rechargeable Batteries via IoT: A Review”. In: *IEEE Internet of Things Journal* (2023).
- [70] Shafto Mike et al. “Modeling, simulation, information technology & processing roadmap”. In: *National Aeronautics and Space Administration* 32.2012 (2012), pp. 1–38.
- [71] Michael W Grieves. “Product lifecycle management: the new paradigm for enterprises”. In: *International Journal of Product Development* 2.1-2 (2005), pp. 71–84.
- [72] Mike Shafto et al. “Draft modeling, simulation, information technology & processing roadmap”. In: *Technology area* 11 (2010), pp. 1–32.
- [73] Akram Hakiri et al. “A comprehensive survey on digital twin for future networks and emerging Internet of Things industry”. In: *Computer Networks* (2024), p. 110350.
- [74] Yujie Wang et al. “A comprehensive review of battery modeling and state estimation approaches for advanced battery management systems”. In: *Renewable and Sustainable Energy Reviews* 131 (2020), p. 110015.
- [75] Xin Liu et al. “A systematic review of digital twin about physical entities, virtual models, twin data, and applications”. In: *Advanced Engineering Informatics* 55 (2023), p. 101876.
- [76] Md Bokhtiar Al Zami et al. “Digital Twin in Industries: A Comprehensive Survey”. In: *IEEE Access* (2025).

- [77] Runze Mao et al. “A Systematic Survey of Digital Twin Applications: Transferring Knowledge From Automotive and Aviation to Maritime Industry”. In: *IEEE Transactions on Intelligent Transportation Systems* (2025).
- [78] Ivan Korotkin et al. “Dandeliion v1: An extremely fast solver for the Newman model of lithium-ion battery (dis) charge”. In: *Journal of The Electrochemical Society* 168.6 (2021), p. 060544.
- [79] Kailong Liu et al. “Electrochemical modeling and parameterization towards control-oriented management of lithium-ion batteries”. In: *Control Engineering Practice* 124 (2022), p. 105176.
- [80] Seongyun Park et al. “Review of state-of-the-art battery state estimation technologies for battery management systems of stationary energy storage systems”. In: *Journal of Power Electronics* 20 (2020), pp. 1526–1540.
- [81] Constantin Gillies. *Glimpse into the future: the digital twin of a high-voltage battery*. 2024.
- [82] Dapai Shi et al. “Cloud-based deep learning for co-estimation of battery state of charge and state of health”. In: *Energies* 16.9 (2023), p. 3855.
- [83] Alan G Li, Alan C West, and Matthias Preindl. “Towards unified machine learning characterization of lithium-ion battery degradation across multiple levels: A critical review”. In: *Applied Energy* 316 (2022), p. 119030.
- [84] Teng Wang et al. “Pipeline condition monitoring towards digital twin system: A case study”. In: *Journal of Manufacturing Systems* 73 (2024), pp. 256–274.
- [85] Youssef Fassi et al. “Toward physics-informed machine-learning-based predictive maintenance for power converters—a review”. In: *IEEE Transactions on Power Electronics* 39.2 (2023), pp. 2692–2720.
- [86] Hongwen He et al. “Online model-based estimation of state-of-charge and open-circuit voltage of lithium-ion batteries in electric vehicles”. In: *Energy* 39.1 (2012), pp. 310–318.
- [87] Peng Rong and M. Pedram. “An analytical model for predicting the remaining battery capacity of lithium-ion batteries”. In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 14.5 (2006), pp. 441–451.
- [88] Antonio Guarino, Walter Zamboni, and Eric Monmasson. “A battery residual capacity indicator based on the battery internal resistance: an experimental study”. In: *2020 2nd IEEE International Conference on Industrial Electronics for Sustainable Energy Systems (IESES)*. Vol. 1. 2020, pp. 15–20.
- [89] Ujwal N. Jibhate and Uday B. Mujumdar. “Development of low complexity open circuit voltage model for state of charge estimation with novel curve modification technique”. In: *Electrochimica Acta* 429 (2022), p. 140944. ISSN: 0013-4686.

- [90] George S. Misyris et al. “State-of-Charge Estimation for Li-Ion Batteries: A More Accurate Hybrid Approach”. In: *IEEE Transactions on Energy Conversion* 34.1 (2019), pp. 109–119.
- [91] Shicong Ding et al. “Accurate Model Parameter Identification to Boost Precise Aging Prediction of Lithium-Ion Batteries: A Review”. In: *Advanced Energy Materials* 13.39 (2023), p. 2301452.
- [92] Kiarash Movassagh et al. “A Critical Look at Coulomb Counting Towards Improving the Kalman Filter Based State of Charge Tracking Algorithms in Rechargeable Batteries”. In: *arXiv preprint arXiv:2101.05435* (2021).
- [93] Xin Lai et al. “A comparative study of different equivalent circuit models for estimating state-of-charge of lithium-ion batteries”. In: *Electrochimica Acta* (2018).
- [94] Seyed Saeed Madani et al. “An electrical equivalent circuit model of a lithium titanate oxide battery”. In: *Batteries* (2019).
- [95] Jianfeng Wang et al. “Precise equivalent circuit model for Li-ion battery by experimental improvement and parameter optimization”. In: *Journal of energy storage* (2022).
- [96] Hehui Zhang et al. “Effect of Sample Interval on the Parameter Identification Results of RC Equivalent Circuit Models of Li-ion Battery: An Investigation Based on HPPC Test Data”. In: *null* (2022).
- [97] Yuejiu Zheng et al. “Investigating the error sources of the online state of charge estimation methods for lithium-ion batteries in electric vehicles”. In: *Journal of Power Sources* 377 (2018), pp. 161–188. ISSN: 0378-7753.
- [98] Murat Ateş et al. “Supercapacitor and battery performances of multi-component nanocomposites: Real circuit and equivalent circuit model analysis”. In: *Journal of energy storage* (2022).
- [99] Miquel Martí-Flores, Andreu Cecilia, and Ramon Costa-Castelló. “Modelling and Estimation in Lithium-Ion Batteries: A Literature Review”. In: *Energies* 16.19 (2023). ISSN: 1996-1073.
- [100] Krishnan S Hariharan et al. “Theoretical framework of the electrochemical model”. In: *Mathematical Modeling of Lithium Batteries: From Electrochemical Models to State Estimator Algorithms* (2018), pp. 13–32.
- [101] Shi-chun Yang et al. “A coupled electrochemical-thermal-mechanical degradation modelling approach for lifetime assessment of lithium-ion batteries”. In: *Electrochimica Acta* 326 (2019), p. 134928.
- [102] Marc Doyle et al. “Modeling of Galvanostatic Charge and Discharge of the Lithium/Polymer/Insertion Cell”. In: *Journal of The Electrochemical Society* (1993).

- [103] Seong Beom Lee et al. “A Robust and Sleek Electrochemical Battery Model Implementation: A MATLAB® Framework”. In: *Journal of The Electrochemical Society* (2021).
- [104] Xuming Yang et al. “Electrochemical Techniques in Battery Research: A Tutorial for Nonelectrochemists”. In: *Advanced Energy Materials* (2019).
- [105] Kailong Liu et al. “Electrochemical modeling and parameterization towards control-oriented management of lithium-ion batteries”. In: *Control Engineering Practice* (2022).
- [106] Weihai Li et al. “Parameter sensitivity analysis of electrochemical model-based battery management systems for lithium-ion batteries”. In: *Applied Energy* (2020).
- [107] Weihai Li et al. “Data-driven systematic parameter identification of an electrochemical model for lithium-ion batteries with artificial intelligence”. In: *Energy Storage Materials* (2021).
- [108] Eric Prada et al. “A Simplified Electrochemical and Thermal Aging Model of LiFePO<sub>4</sub>-Graphite Li-ion Batteries: Power and Capacity Fade Simulations”. In: *Journal of The Electrochemical Society* (2013).
- [109] Moritz Streb et al. “Investigating re-parametrization of electrochemical model-based battery management using real-world driving data”. In: *null* (2023).
- [110] Sankhadeep Sarkar et al. “Electrochemical Models: Methods and Applications for Safer Lithium-Ion Battery Operation”. In: *Journal of The Electrochemical Society* (2022).
- [111] Vankamamidi S Naresh et al. “Optimizing electric vehicle battery health monitoring: A resilient ensemble learning approach for state-of-health prediction”. In: *Sustainable Energy, Grids and Networks* 42 (2025), p. 101655.
- [112] Sahar Khaleghi et al. “Ensemble gradient boosted tree for SoH estimation based on diagnostic features”. In: *Energies* 13.5 (2020), p. 1262.
- [113] Yi Li et al. “Random forest regression for online capacity estimation of lithium-ion batteries”. In: *Applied energy* 232 (2018), pp. 197–210.
- [114] Rui Pan et al. “State of health estimation for lithium-ion batteries based on two-stage features extraction and gradient boosting decision tree”. In: *Energy* 285 (2023), p. 129460.
- [115] Kristen A Severson et al. “Data-driven prediction of battery cycle life before capacity degradation”. In: *Nature Energy* 4.5 (2019), pp. 383–391.
- [116] Fujin Wang et al. “Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis”. In: *Nature Communications* 15.1 (2024), p. 4332.

- [117] L Ma, C Hu, and F Cheng. “State of charge and state of energy estimation for lithium-ion batteries based on a long short-term memory neural network”. In: *Journal of Energy Storage* 37 (2021), p. 102440.
- [118] Muhammad Kamran Khan et al. “Efficient state of charge estimation of lithium-ion batteries in electric vehicles using evolutionary intelligence-assisted GLA–CNN–Bi-LSTM deep learning model”. In: *Heliyon* 10.15 (2024).
- [119] Peichao Li et al. “State of charge estimation for lithium-ion batteries based on physics-embedded neural network”. In: *Journal of Power Sources* 640 (2025), p. 236785.
- [120] Ephrem Chemali et al. “Long short-term memory networks for accurate state-of-charge estimation of Li-ion batteries”. In: *IEEE Transactions on Industrial Electronics* 65.8 (2017), pp. 6730–6739.
- [121] John Guirguis and Ryan Ahmed. “Transformer-Based Deep Learning Models for State of Charge and State of Health Estimation of Li-Ion Batteries: A Survey Study”. In: *Energies* 17.14 (2024). ISSN: 1996-1073.
- [122] Amal Mahmoud and Ammar Mohammed. “Leveraging hybrid deep learning models for enhanced multivariate time series forecasting”. In: *Neural Processing Letters* 56.5 (2024), p. 223.
- [123] Jingyuan Zhao et al. “Specialized deep neural networks for battery health prognostics: Opportunities and challenges”. In: *Journal of Energy Chemistry* 87 (2023), pp. 416–438.
- [124] Junjie Tao et al. “A Comprehensive Review of Multiple Physical and Data-Driven Model Fusion Methods for Accurate Lithium-Ion Battery Inner State Factor Estimation”. In: *Batteries* 10.12 (2024), p. 442.
- [125] Hao Li, Lijun Fu, and Yan Zhang. “A novel hybrid data-driven method based on uncertainty quantification to predict the remaining useful life of lithium battery”. In: *Journal of Energy Storage* 52 (2022), p. 104984.
- [126] Yue Song et al. “Detection of voltage fault in lithium-ion battery based on equivalent circuit model-informed neural network”. In: *IEEE Transactions on Instrumentation and Measurement* 73 (2024), pp. 1–10.
- [127] Wei He et al. “State of charge estimation for Li-ion batteries using neural network modeling and unscented Kalman filter-based error cancellation”. In: *International Journal of Electrical Power & Energy Systems* 62 (2014), pp. 783–791.
- [128] Xin Lai et al. “Capacity estimation of lithium-ion cells by combining model-based and data-driven methods based on a sequential extended Kalman filter”. In: *Energy* 216 (2021), p. 119233.

- [129] Soumya Singh et al. “Hybrid Modeling of Lithium-Ion Battery: Physics-Informed Neural Network for Battery State Estimation”. In: *Batteries* 9.6 (2023). ISSN: 2313-0105.
- [130] Jan Kleiner et al. “Advanced monitoring and prediction of the thermal state of intelligent battery cells in electric vehicles by physics-based and data-driven modeling”. In: *Batteries* 7.2 (2021), p. 31.
- [131] Ming Wu et al. “Data-driven framework of ECM: A machine learning model for profile prediction”. In: *Proc. 17th INSECT* (2021), pp. 231–240.
- [132] Hyunjin Ahn et al. “State of Charge Estimation of Lithium-Ion Batteries Using Physics-Informed Transformer for Limited Data Scenarios”. In: *ASME Letters in Dynamic Systems and Control* 3.4 (2023).
- [133] Huiqin Li et al. “A review on physics-informed data-driven remaining useful life prediction: Challenges and opportunities”. In: *Mechanical Systems and Signal Processing* 209 (2024), p. 111120.
- [134] Wendi Guo et al. “Review of “grey box” lifetime modeling for lithium-ion battery: Combining physics and data-driven methods”. In: *Journal of Energy Storage* 56 (2022), p. 105992. ISSN: 2352-152X.
- [135] Dapai Shi et al. “Spatial-temporal self-attention transformer networks for battery state of charge estimation”. In: *Electronics* 12.12 (2023), p. 2598.
- [136] Hongzhi Zhang and M Omair Shafiq. “Survey of transformers and towards ensemble learning using transformers for natural language processing”. In: *Journal of big Data* 11.1 (2024), p. 25.
- [137] Ashish Vaswani et al. “Attention is all you need”. In: *Advances in neural information processing systems* 30 (2017).
- [138] Yu Cao, Xin Wen, and Hongyu Liang. “Multi-Temperature State-of-Charge Estimation of Lithium-Ion Batteries Based on Spatial Transformer Network”. In: *Energies* 17.20 (2024), p. 5029.
- [139] Metin Yılmaz, Eyüp Çinar, and Ahmet Yazıcı. “A Novel Transformer-Based Model for State of Charge Estimation of Electrical Vehicle Batteries”. In: *Available at SSRN* 4736082 (2022).
- [140] Yuxin Duan and Runmin Zou. “A Novel State-of-Charge Estimation Method for Lithium-ion Batteries Using Convolutional Transformer Network and Sigma-point Kalman Filter”. In: *2022 4th International Conference on Smart Power & Internet Energy Systems (SPIES)*. IEEE. 2022, pp. 1850–1855.
- [141] William Gomez, Fu-Kwun Wang, and Jia-Hong Chou. “Li-ion battery capacity prediction using improved temporal fusion transformer model”. In: *Energy* 296 (2024), p. 131114.

- [142] Yuqian Fan et al. “Online state-of-health estimation for fast-charging lithium-ion batteries based on a transformer–long short-term memory neural network”. In: *Batteries* 9.11 (2023), p. 539.
- [143] Qiao Wang et al. “Unlocking the potential of unlabeled data: Self-supervised machine learning for battery aging diagnosis with real-world field data”. In: *Journal of Energy Chemistry* 99 (2024), pp. 681–691.
- [144] Haoyi Zhou et al. “Informer: Beyond efficient transformer for long sequence time-series forecasting”. In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 35. 12. 2021, pp. 11106–11115.
- [145] Giovanni Pollo et al. *Coupling Neural Networks and Physics Equations For Li-Ion Battery State-of-Charge Prediction*. 2024.
- [146] Salvatore Cuomo et al. “Scientific machine learning through physics–informed neural networks: Where we are and what’s next”. In: *Journal of Scientific Computing* 92.3 (2022), p. 88.
- [147] Bin Huang and Jianhui Wang. “Applications of physics-informed neural networks in power systems-a review”. In: *IEEE Transactions on Power Systems* 38.1 (2022), pp. 572–588.
- [148] Apostolos F Psaros, Kenji Kawaguchi, and George Em Karniadakis. “Meta-learning PINN loss functions”. In: *Journal of computational physics* 458 (2022), p. 111121.
- [149] Fujin Wang et al. “Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis”. In: *Nature Communications* 15.1 (2024).
- [150] Jingyuan Zhao et al. “Specialized deep neural networks for battery health prognostics: Opportunities and challenges”. In: *Journal of Energy Chemistry* 87 (2023), pp. 416–438. ISSN: 2095-4956.
- [151] Chenxi Wu et al. “A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks”. In: *Computer Methods in Applied Mechanics and Engineering* 403 (2023), p. 115671. ISSN: 0045-7825.
- [152] Wenqian Chen, Yucheng Fu, and Panagiotis Stinis. “Physics-informed machine learning of redox flow battery based on a two-dimensional unit cell model”. In: *Journal of Power Sources* 584 (Sept. 2023).
- [153] Pierpaolo Dini, Antonio Colicelli, and Sergio Saponara. “Review on Modeling and SOC/SOH Estimation of Batteries for Automotive Applications”. In: *Batteries* 10.1 (2024). ISSN: 2313-0105.
- [154] Haider Adel Ali Ali et al. “A comparison between physics-based Li-ion battery models”. In: *Electrochimica Acta* 493 (2024), p. 144360. ISSN: 0013-4686.

- [155] Benedikt Rzepka, Simon Bischof, and Thomas Blank. “Implementing an Extended Kalman Filter for SoC Estimation of a Li-Ion Battery with Hysteresis: A Step-by-Step Guide”. In: *Energies* 14.13 (2021). ISSN: 1996-1073.
- [156] Jacqueline S. Edge et al. “Lithium ion battery degradation: What you need to know”. In: *Physical Chemistry Chemical Physics* 23.14 (2021), pp. 8200–8221.
- [157] Rui Xiong et al. “Lithium-Ion Battery Health Prognosis Based on a Real Battery Management System Used in Electric Vehicles”. In: *IEEE Transactions on Vehicular Technology* 68.5 (2019), pp. 4110–4121.
- [158] Taesic Kim, Wei Qiao, and Liyan Qu. “An Enhanced Hybrid Battery Model”. In: *IEEE Transactions on Energy Conversion* 34.4 (2019), pp. 1848–1858.
- [159] Matthias Steinstraeter, Johannes Buberger, and Dimitar Trifonov. *Battery and Heating Data in Real Driving Cycles*. 2020.
- [160] B Bole, C Kulkarni, and M Daigle. *Randomized Battery Usage Data Set (dataset 11, last accessed: 26/04/2024)*. <https://www.nasa.gov/intelligent-systems-division/discovery-and-systems-health/pcoe/pcoe-data-set-repository/>. NASA Ames Research Center, 2022.
- [161] Yuliya Preger et al. “Degradation of Commercial Lithium-Ion Cells as a Function of Chemistry and Cycling Conditions”. In: *Journal of The Electrochemical Society* 167.12 (2020), p. 120532.
- [162] Xing Shu et al. “State of Charge Estimation for Lithium-Ion Battery Based on Hybrid Compensation Modeling and Adaptive H-Infinity Filter”. In: *IEEE Transactions on Transportation Electrification* 9.1 (2023), pp. 945–957.
- [163] Gabriele Pozzato, Anirudh Allam, and Simona Onori. “Lithium-ion battery aging dataset based on electric vehicle real-driving profiles”. In: *Data in Brief* 41 (2022), p. 107995. ISSN: 2352-3409.
- [164] Yan Zhang. “Edge Computing for Digital Twin”. In: *Digital Twin: Architectures, Networks, and Applications*. Cham: Springer Nature Switzerland, 2024, pp. 41–59. ISBN: 978-3-031-51819-5.
- [165] Khaled Sidahmed Sidahmed Alamin et al. “A Machine Learning-based Digital Twin for Electric Vehicle Battery Modeling”. In: *2022 IEEE International Conference on Omni-layer Intelligent Systems (COINS)*. 2022, pp. 1–6.
- [166] Khaled Sidahmed Sidahmed Alamin et al. “Advancing Electric Vehicle Battery Management: A Data-Driven Digital Twin Approach for Real-Time Monitoring and Performance Enhancement”. In: *IEEE Transactions on Vehicular Technology* (2025), pp. 1–15.
- [167] Khaled Sidahmed Sidahmed Alamin et al. “Model-Driven Dataset Generation for Data-Driven Battery SOH Models”. In: *2023 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. 2023, pp. 1–6.

- [168] Cordis.Europa.Eu Cordis. *Electric Vehicle Enhanced Range, Lifetime And Safety Through INGenious battery management*. Aug. 2016.
- [169] Woosuk Sung and Jaewook Lee. “Implementation of SOH Estimator in Automotive BMSs Using Recursive Least-Squares”. In: *Electronics* 8.11 (2019). ISSN: 2079-9292.
- [170] Subir Halder, Amrita Ghosal, and Mauro Conti. “Secure over-the-air software updates in connected vehicles: A survey”. In: *Computer Networks* 178 (2020), p. 107343.
- [171] Shichun Yang et al. “Implementation for a cloud battery management system based on the CHAIN framework”. In: *Energy and AI* 5 (2021), p. 100088. ISSN: 2666-5468.
- [172] Inc. Peaxy. *Digital Twins: Practical applications in battery analytics*. Dec. 2024.
- [173] *insights into battery aging with the virtual twin<sub>2025a</sub>*. 2025.
- [174] Martin Murnane and Adel Ghazel. *A Closer Look at State of Charge (SOC) and State of Health (SOH) Estimation Techniques for Batteries | Analog Devices*. 2023.
- [175] Qaiser Hussain et al. “Prediction of Remaining Useful Life of Battery Using Partial Discharge Data”. In: *Electronics (Switzerland)* 13.17 (2024), p. 3475.
- [176] Zhen Zhang et al. “Data-driven available capacity estimation of lithium-ion batteries based on fragmented charge capacity”. In: *Communications Engineering* 4.1 (2025), p. 32.
- [177] Jinyu Wang et al. “A novel aging characteristics-based feature engineering for battery state of health estimation”. In: *Energy* 273 (2023), p. 127169.
- [178] Jinyu Wang et al. “A Novel Feature Engineering-Based SOH Estimation Method for Lithium-Ion Battery with Downgraded Laboratory Data”. In: *Batteries* 10.4 (2024). ISSN: 2313-0105.
- [179] Ye He et al. “SOH estimation for lithium-ion batteries: An improved GPR optimization method based on the developed feature extraction”. In: *Journal of Energy Storage* 83 (2024), p. 110678.
- [180] Gabriele Pozzato, Anirudh Allam, and Simona Onori. “Lithium-Ion battery aging dataset based on electric vehicle real-driving profiles”. In: *Data in Brief* 41 (2022), p. 107995. ISSN: 2352-3409.
- [181] Guolin Ke et al. “Lightgbm: A highly efficient gradient boosting decision tree”. In: *Advances in neural information processing systems* 30 (2017).
- [182] Harris Drucker, Donghui Wu, and Vladimir N Vapnik. “Support vector machines for spam categorization”. In: *IEEE Transactions on Neural networks* 10.5 (1999), pp. 1048–1054.

- [183] MathWorks. *Battery State-of-Health Estimation*. <https://it.mathworks.com/help/simscape-battery/ug/battery-state-of-health-estimation.html>. 2023.
- [184] Takuya Akiba et al. “Optuna: A next-generation hyperparameter optimization framework”. In: *ACM SIGKDD international conference on knowledge discovery & data mining*. 2019, pp. 2623–2631.
- [185] Yue Xiang et al. “Semi-supervised deep learning for lithium-ion battery state-of-health estimation using dynamic discharge profiles”. In: *Cell Reports Physical Science* 5.1 (2024).
- [186] Wei Liu, Yan Xu, and Xue Feng. “A Hierarchical and Flexible Data-Driven Method for Online State-of-Health Estimation of Li-Ion Battery”. In: *IEEE Transactions on Vehicular Technology* 69.12 (2020), pp. 14739–14748.
- [187] Kaushik Das, Roushan Kumar, and Anurup Krishna. “Analyzing electric vehicle battery health performance using supervised machine learning”. In: *Renewable and Sustainable Energy Reviews* 189 (2024), p. 113967.
- [188] Ameera Arif et al. “Estimating Battery State of Health using Machine Learning”. In: *International Conference on Smart Grid and Clean Energy Technologies (ICSGCE)*. 2022, pp. 72–77.
- [189] Tesla, Inc. “Battery condition determination”. US20220165531A1. Col. 12 lists BMS SoC/SoH accuracy and ISR timings. June 2022.
- [190] Yang Wu and Peirong Zhang. “Reverse Engineering a Tesla Battery Management System”. In: *Black Hat USA Proceedings* (2020).
- [191] *bq34z100-G1: Impedance-Track™ Fuel Gauge for Li-ion and Lead-Acid Batteries*. Tech. rep. SLUUBW5A. Spec sheet §1.1; host read latency in Table6-1. Dallas, TX: Texas Instruments, 2023.
- [192] *High-Accuracy Battery State Estimation Using a Sigma-Point EKF on S32K3 Automotive MCUs*. Tech. rep. White-paper Table2: flash/RAM size and 1.1ms loop on 200MHz Cortex-M7. NXP Semiconductors, 2022.
- [193] *Part 573 Safety Recall Report 21V-127(ABS-21E003) — Hyundai Kona Electric Battery System*. Tech. rep. Pages 3–4 list BMS update size and 8ms control loop. Washington, DC: National Highway Traffic Safety Administration, 2021.
- [194] Chunsheng Wang et al. “A hybrid model for state of charge estimation of lithium-ion batteries utilizing improved adaptive extended Kalman filter and long short-term memory neural network”. In: *Journal of Power Sources* 620 (2024), p. 235272.
- [195] Yue Xiang et al. “Semi-supervised deep learning for lithium-ion battery state-of-health estimation using dynamic discharge profiles”. In: *Cell Reports Physical Science* 5.1 (2024).

- [196] Spyridon Giazitzis et al. “A Case Study of a Tiny Machine Learning Application for Battery State-of-Charge Estimation”. In: *Electronics* 13.10 (2024), p. 1964.
- [197] Yukai Chen et al. “Electric vehicles plug-in duration forecasting using machine learning for battery optimization”. In: *Energies* 13.16 (2020), p. 4208.
- [198] Khaled Sidahmed Sidahmed Alamin et al. “An AI-Enabled Framework for Smart Semiconductor Manufacturing”. In: *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE. 2024, pp. 1–6.
- [199] Nicola Dall’Ora et al. “Digital transformation of a production line: Network design, online data collection and energy monitoring”. In: *IEEE Transactions on Emerging Topics in Computing* 10.1 (2021), pp. 46–59.
- [200] Khaled Sidahmed Sidahmed Alamin et al. “SMART-IC: smart monitoring and production optimization for zero-waste semiconductor manufacturing”. In: *2022 IEEE 23rd Latin American Test Symposium (LATS)*. IEEE. 2022, pp. 1–6.
- [201] Khaled Alamin et al. “Digital twin extension with extra-functional properties”. In: *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE. 2021, pp. 434–439.

This Ph.D. thesis has been typeset by means of the  $\text{\TeX}$ -system facilities. The typesetting engine was  $\text{pdf\LaTeX}$ . The document class was `toptesi`, by Claudio Beccari, with option `tipotesi=scudo`. This class is available in every up-to-date and complete  $\text{\TeX}$ -system installation.