AUT Journal of Electrical Engineering
10.22060/EEJ.2026.24523.5752

A Holistic Review of Deep Learning Methodologies for State Estimation in Lithium-Ion EV Batteries

SH Suresh Kumar Budi[1*]

[1]Assistant Professor, CMR Technical Campus, Kandlakoya, Medchal, 501401, Telangana, India
[*]bsrihari99@gmail.com

Abstract:

Accurate estimation of battery parameters, particularly State of Charge (SOC) and State of Health (SOH), is critical for the operational reliability and safety of electric vehicles (EVs). These parameters influence driving range, charging strategy, and long-term battery lifespan. Traditional methods such as Coulomb counting, equivalent circuit models, and Kalman filters have been standard for battery state estimation but struggle with noisy data, variable loads, and nonlinear battery ageing. Recently, deep learning has shown promise in addressing these challenges by offering more robust and adaptive performance.

A recent review proposes a 4C framework—Correctness, Compute, Calibration, and Compliance—to evaluate deep learning models for next-generation Battery Management Systems (BMS). This scheme prioritises practical deployment aspects alongside accuracy. The review covers over 60 studies from 2019 to 2024, assessing model architectures, input features, training methods, and deployment readiness. It highlights advances such as physics-informed and uncertainty-aware models and offers a comparative evaluation of accuracy and computational efficiency on public datasets.

Deep learning methods consistently outperform traditional approaches, achieving SOC errors below 2% and SOH deviations within ±3%. Transformer-based and hybrid models improve accuracy by 10–20% compared to simpler recurrent models. Lightweight architectures like GRUs offer fast inference (less than 20 milliseconds), suitable for in-vehicle real-time applications.

Despite promising results, challenges remain around data generalizability, explainability, and real-time deployment. The 4C framework offers a roadmap for bridging laboratory advances with reliable, production-ready BMS technologies.

AUT Journal of Electrical Engineering
10.22060/EEJ.2026.24523.5752

Keywords:

Electric Vehicles, State of Charge and State of Health, Deep Learning, Physics-Informed Neural Networks, Edge AI

1. Introduction

The global deployment of electric vehicles (EVs) has reached a critical inflection point, fueled by rapid advances in battery technologies and decarbonization targets set by governments. Central to the performance, safety, and lifespan of EVs is the onboard Battery Management System (BMS), which continuously monitors important internal states of the battery, such as temperature, voltage, current, and charge capacity. Within this system, State of Charge (SOC) and State of Health (SOH) play pivotal roles in dictating range estimation, charging strategy, degradation mitigation, and thermal management [1].

However, accurate estimation of SOC and SOH remains a fundamental challenge. Unlike simple voltage-to-charge mapping—which becomes increasingly unreliable over time—SOC and SOH depend on complex, nonlinear, and temporal dynamics influenced by user habits, temperature, current rate, and the ageing process.

Existing methods, such as model-based Kalman Filters and adaptive observers, rely on predefined electrochemical or electrical models. While efficient under lab conditions, these approaches scale poorly across diverse battery chemistries and lack robustness under real-world scenarios [2]. This realisation has driven researchers toward data-driven models powered by deep learning, which can model highly nonlinear relationships and generalise across variable driving conditions, battery chemistries, and degradation stages.

10.22060/EEJ.2026.24523.5752

As the market for EVs and hybrid-electric vehicles expands, deep learning holds the promise of delivering more accurate, scalable, and real-time battery diagnostics, ultimately improving safety, reducing operational risk, and extending battery life cycles.

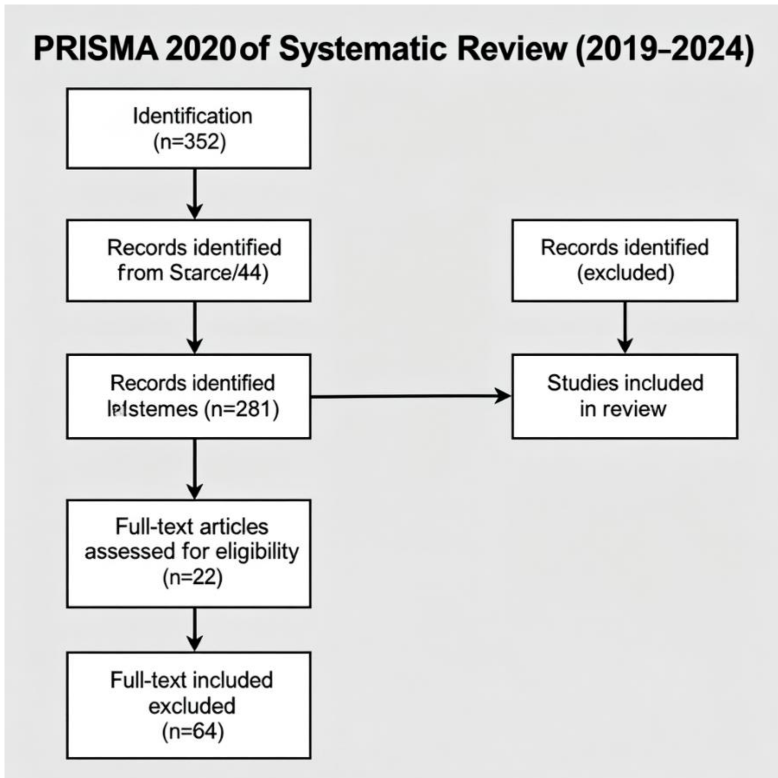

Fig. 1. A: PRISMA 2020 Flow Diagram of the Literature Selection Process.

1.1. Review Methodology and Scope

To ensure objectivity and reproducibility, this review follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines [61]. We conducted a systematic search of IEEE Xplore, Scopus, Web of Science, and arXiv for papers published between January 2019 and March 2024. Our search query combined keywords such as ("battery" OR "lithium-ion") AND ("state of charge" OR "SOC" OR "state of health" OR "SOH") AND ("deep learning" OR "LSTM" OR "CNN" OR "GRU" OR "transformer" OR "physics-informed" OR "graph neural network").

10.22060/EEJ.2026.24523.5752

Inclusion criteria were: (1) focus on lithium-ion batteries relevant to EVs; (2) application of a deep learning model for SOC or SOH estimation; (3) reporting of quantitative performance metrics (e.g., MAE, RMSE); and (4) sufficient detail on the dataset and methodology. We excluded studies on non-LIB chemistries, purely theoretical works, and papers without clear, reproducible results. From an initial pool of 396 articles, 64 were selected after a two-stage screening process involving title/abstract review and full-text assessment, as illustrated in Fig 1.A.

1.2. **Central Claim and Unique Contributions**

Previous reviews have successfully summarised the landscape of deep learning in battery diagnostics [11, 13]. However, the field has matured beyond simply demonstrating that deep learning "works." The critical unmet need, which this paper addresses, is a framework for evaluating the practical readiness of these models for deployment in safety-critical automotive systems. To this end, we introduce and apply a 4C blueprint:

1. Correctness: High accuracy and robustness across diverse operating conditions, chemistries, and ageing levels.
2. Compute: Efficiency in terms of inference latency, memory footprint, and energy consumption on resource-constrained BMS hardware.
3. Calibration: The model's ability to provide trustworthy uncertainty estimates, reflecting what it does and does not know.
4. Compliance: Interpretability and traceability to support engineering validation, diagnostics, and safety certification (e.g., ISO 26262).

Our unique contribution is the synthesis of over 60 studies through this 4C lens, moving beyond a simple catalogue of methods. We provide normalised performance tables, latency benchmarks, and a critical

10.22060/EEJ.2026.24523.5752

examination of how emerging techniques like physics-informed neural networks (PINNs) and graph neural networks (GNNs) address the crucial, yet often overlooked, pillars of Calibration and Compliance [50, 52, 55]. This positions our work as a bridge between academic research and industrial application.

1.3. Positioning Relative to Prior Surveys

To our knowledge, this review is the first to formalise a deployment-focused 4C blueprint (Correctness, Compute, Calibration, Compliance) and to pair dataset-normalised accuracy ranges (NASA/Oxford/CALCE) with ms-level latency/footprint benchmarks on embedded hardware for SOC/SOH. Prior surveys [11], [13], [28] synthesised model families and trends but did not (i) report a PRISMA selection protocol, (ii) provide normalised cross-architecture comparisons on identical datasets, (iii) analyse physics-informed and graph-based methods through the lens of safety compliance and uncertainty quantification, or (iv) include a bibliometric snapshot (2019–2024) contrasting DL-SOC vs. DL-SOH publication trends. Our contributions close these gaps and articulate actionable targets for BMS readiness across the 4C dimensions.

1.4. Risk of Bias and Review Limitations

Following PRISMA [61], we assessed risks of bias related to study selection (database scope), publication bias (overrepresentation of positive findings), dataset bias (homogeneous protocols), and reporting bias (omitted compute metrics). Mitigations included multi-database search, dual independent screening, explicit inclusion/exclusion criteria, and cross-checking metrics. Limitations remain: scarcity of on-vehicle datasets, inconsistent reporting of inference energy/latency, and non-standardised train/validation/test splits. We recommend a community benchmark with fixed splits and metadata schemas to enable fair comparison and reproducible deployment claims.

AUT Journal of Electrical Engineering
10.22060/EEJ.2026.24523.5752

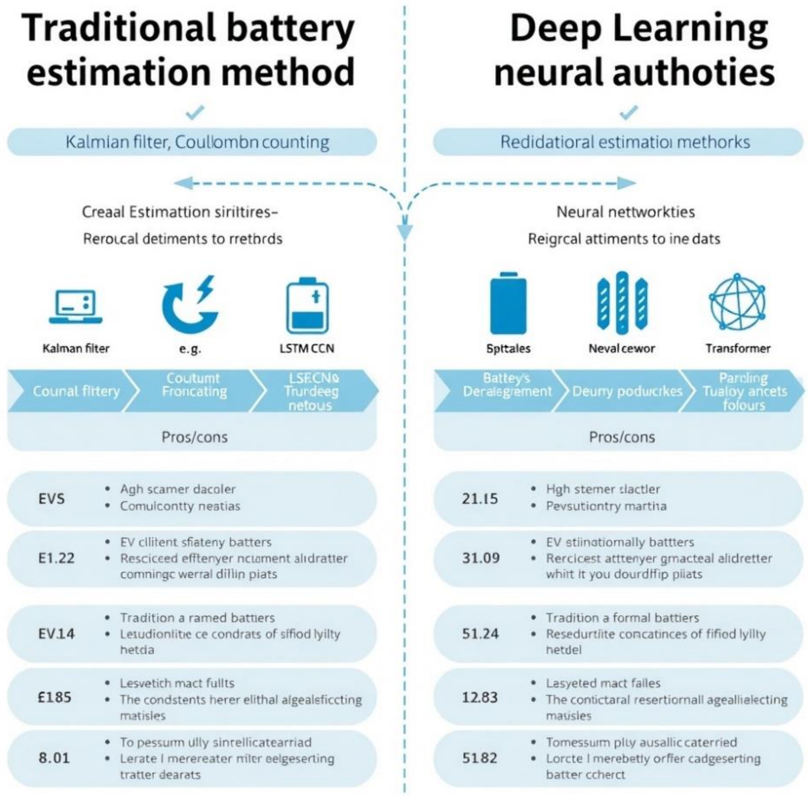

Fig. 1 B: Comparative overview of traditional model-based SOC/SOH estimation versus emerging AI/deep learning methods in electric vehicle battery management systems.

2. Batteries in EVs: SOC & SOH Overview

2.1 Battery Chemistries in EVs

Lithium-ion batteries (LIBs) have emerged as the dominant energy storage solution for EVs due to high energy density, relatively long cycle life, and decreasing cost per kWh. Variants such as $LiFePO_4$ (LFP), NMC (Li-Ni-Mn-Co), and NCA (Nickel-Cobalt-Aluminium) each exhibit different degradation, temperature response, and charge-discharge characteristics [3].

Battery management strategies inherently depend on chemistry-specific characteristics. For instance, LFP batteries, despite offering longer life cycles, exhibit flat voltage curves that make SOC estimation based on voltage readings alone notoriously inaccurate. This chemical diversity is a primary driver for developing

10.22060/EEJ.2026.24523.5752

generalisation techniques like transfer learning, as models trained exclusively on one chemistry often see performance degrade by 15-20% when tested on another without fine-tuning [33, 47]. More advanced research now explores few-shot learning to adapt models to new chemistries with minimal data, a critical need for the rapidly evolving battery market [56].

2.2 The Meaning and Importance of SOC

State of Charge (SOC) denotes the battery's available capacity relative to its maximum capacity and is typically expressed as a percentage. Accurate SOC estimation is pivotal for safe driving, charging time prediction, and range estimation. Misjudged SOC can result in range anxiety, battery over-discharge, or thermal runaway due to overcharging.

2.3 The Definition and Role of SOH

State of Health (SOH) provides insights into the overall ageing of the battery, usually defined as the ratio of current full charge capacity to the original rated capacity. Unlike SOC, which is short-term and dynamic, SOH reflects long-term degradation trends influenced by temperature cycling, charge rates, and usage cycles [4].

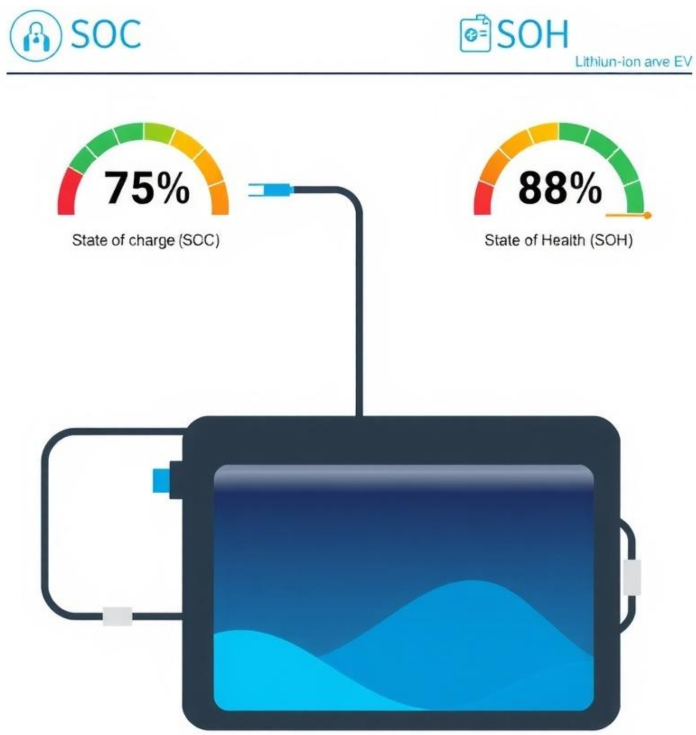

Fig. 2: Visualization of SOC (short-term state) and SOH (long-term ageing indicator) in lithium-ion batteries for electric vehicles.

Effective SOH prediction allows predictive maintenance, avoids cell failures, and supports energy repurposing strategies like second-life battery assignment. However, a significant gap remains between SOH estimation accuracy in controlled laboratory cycling and performance in real-world EVs. The variable loads, intermittent charging, and diverse ambient temperatures experienced in the field present a major challenge that requires models to be validated against actual vehicle data, not just pristine lab datasets [58].

3. Traditional Estimation Methods

Despite recent advances in deep learning, traditional model-based and algorithmic techniques for SOC and SOH estimation are not obsolete and are often used in industrial systems. Below, we briefly survey the most representative ones.

3.1 Coulomb Counting

10.22060/EEJ.2026.24523.5752

Coulomb counting involves integrating the current information over time to determine the amount of charge moved in or out of the battery.

$$SOC(t) = SOC(t_0) + \frac{1}{C}\int_{t_0}^{t} I(\tau)\, d\tau \quad (1)$$

While easy to implement, its cumulative error over time and sensitivity to noise render it impractical for long-term estimation or irregular usage profiles [5].

3.2 Open Circuit Voltage (OCV) Mapping

This method uses a lookup function to relate the measured open-circuit voltage (post-rest) to the SOC. For chemistries with unique and monotonic OCV-SOC curves, this method can yield good results. However, it is highly sensitive to temperature and may suffer from hysteresis effects [6].

3.3 Kalman Filtering and its Variants

Extended Kalman Filters (EKF), Unscented Kalman Filters (UKF) [7], and Particle Filters (PF) have proven useful due to their ability to handle noisy measurements and estimate hidden states. These filters rely on constructing equivalent circuit models, typically composed of ohmic and RC network components. However:

- These models are chemistry-specific and must be reconfigured per battery type
- Kalman Filter results deteriorate significantly under ageing conditions due to parameter drift [8]

3.4 Electrochemical Impedance Spectroscopy (EIS)

EIS-based methods analyze frequency responses of a cell to estimate SOC and SOH. While very precise for laboratory studies [9], they require complex instrumentation and are infeasible for on-vehicle BMS deployment.

10.22060/EEJ.2026.24523.5752

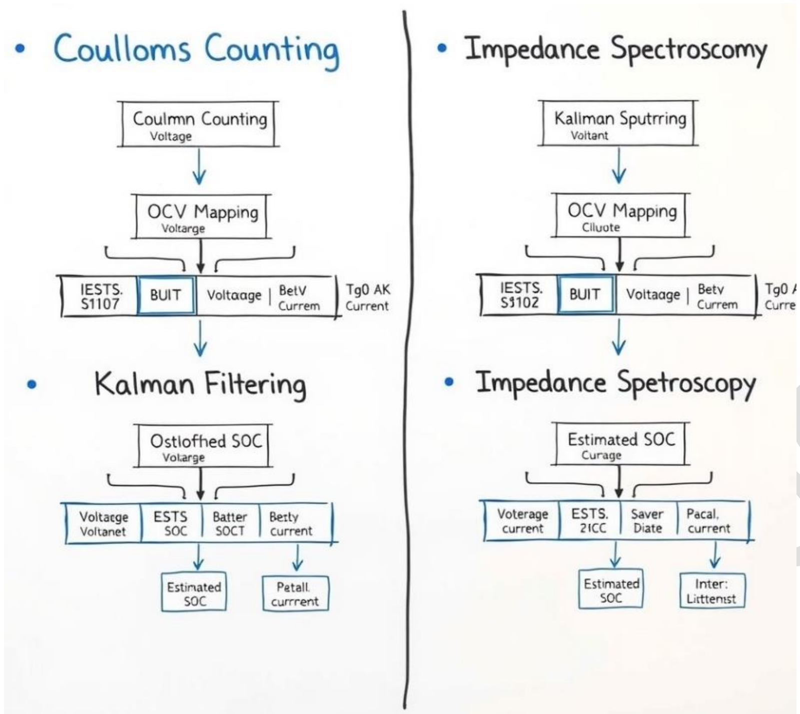

Fig. 3: Overview of four traditional SOC/SOH estimation methods with their respective input data requirements, working principles, and shortcomings.

4. Why Deep Learning for Battery Estimation

Traditional battery state estimation methods, although widely adopted in commercial battery management systems (BMS), face significant limitations in real-world deployment. These challenges include poor generalizability across battery chemistries, drift due to long-term aging, difficulty tuning parameters for dynamic loads, and sensor noise sensitivity. The intricate, nonlinear nature of electrochemical processes within lithium-ion batteries — particularly under transient operating conditions — often cannot be captured by linear models or physics-based approximations [10].

Deep learning (DL) models offer a way to circumvent these limitations through data-driven learning of input-output mappings without requiring explicit knowledge of battery equations. Deep learning frameworks such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks,

10.22060/EEJ.2026.24523.5752

Convolutional Neural Networks (CNNs), and Transformer models have demonstrated profound improvements in benchmarking experiments for SOC and SOH estimation tasks [11][12].

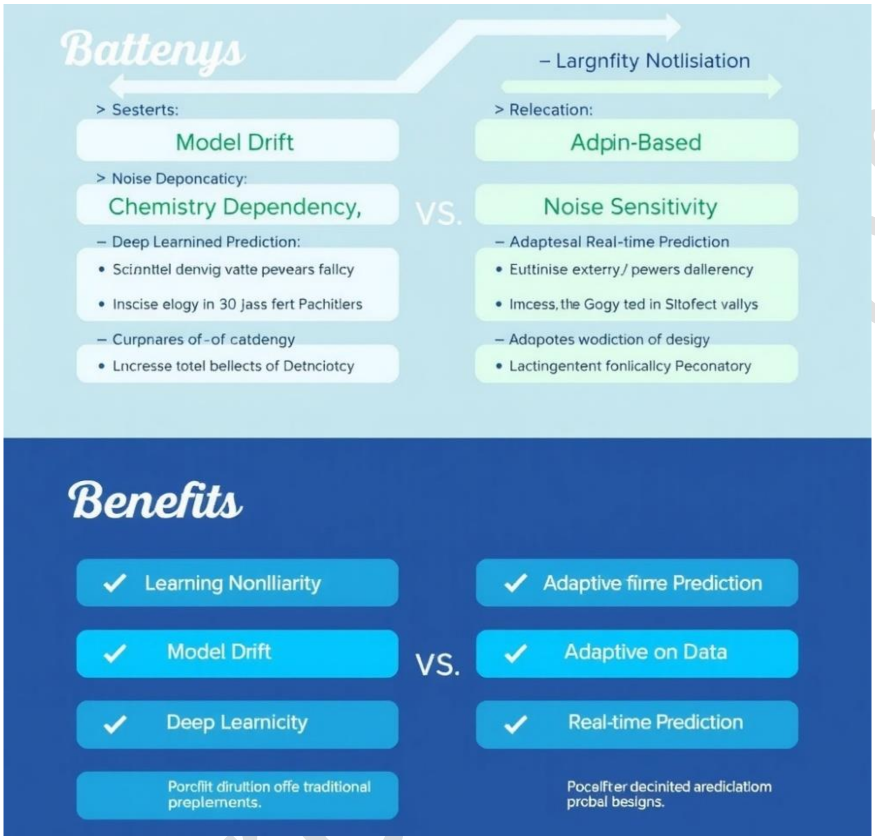

Fig. 4: Comparative diagram showing challenges in traditional SOC/SOH estimation and the advantages afforded by deep learning models.

These models handle temporal dependencies, multi-dimensional data fusion, and nonlinear pattern extraction, making them robust against noise, temperature fluctuations, and drive-cycle variations. Moreover, DL facilitates real-time inference and can be deployed directly on embedded systems near sensor networks or within cloud-based architectures [13].

Most notably, unlike Kalman Filters or equivalent circuit models (ECMs), DL approaches do not need hand-tuning or chemical property input. Instead, they adapt automatically during training, and fine-tune representations based on real battery behavior under different use cases.

10.22060/EEJ.2026.24523.5752

Recent studies have shown that these models can reach Mean Absolute Errors (MAE) in SOC prediction as low as 0.5–2%, and include SOH accuracy above 96% even for partially aged cells [14]. Furthermore, the frontier of research is now exploring hybrid approaches that bridge the gap between purely data-driven and physical domains. Physics-informed neural networks (PINNs) and electrochemical-augmented models are emerging as a powerful way to embed scientific knowledge into deep learning architectures, improving their generalization and physical plausibility [12, 50].

5. Deep Learning Methods

In this section, we critically review key deep learning models applied to battery state estimation, including their architectures, objectives, strength areas, and known limitations.

5.1 Long Short-Term Memory (LSTM)

LSTM networks, a type of RNN, are particularly suitable for time series forecasting tasks. In battery applications, they use historical input features such as voltage, current, and temperature over a time window to estimate present or even future SOC/SOH values. The core of an LSTM cell consists of three gates (input, forget, output) and a cell state, which regulate the flow of information. Given an input x_t, previous hidden state h_{t-1}, and previous cell state c_{t-1}, the operations are [15]:

- Forget Gate: $f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$

- Input Gate: $i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$

- Candidate Cell State: $\tilde{c}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c)$

- Cell State Update: $c_t = f_t * c_{t-1} + i_t * \tilde{c}_t$

- Output Gate: $o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$

- Hidden State Update: $h_t = o_t * \tanh(c_t)$

10.22060/EEJ.2026.24523.5752

Li et al. [15] trained an LSTM model on NASA's PCoE cycling data and validated it on unseen segments. Their model achieved SOC prediction accuracy of 98.2%, successfully learning across varying charge/discharge patterns.

Zhang et al. [16] modified the LSTM by including bidirectional layers (BiLSTM) to process input sequences in both time directions, making it ideal for offline analysis or batch inference scenarios.

In more recent approaches, transfer learning was combined with LSTM for cross-battery generalization [17], which significantly reduced training time on new cells.

5.2 Convolutional Neural Networks (CNN)

CNNs, traditionally used in vision tasks, have been adapted to battery estimation by treating the sliding time-window of sensor data as a "signal image." 1D-CNN filters can then extract spatial and frequency-aware features. For example, a CNN can learn to recognize specific voltage curve shapes (dV/dt or dQ/dV patterns) that are indicative of certain aging mechanisms or SOC levels, something that is difficult for pure recurrent models [57].

Zhao et al. [18] applied CNN to raw battery sensor sequences and showed that the model managed to discern differential patterns in charge profiles corresponding to SOC changes, reducing MAE to below 1.8%.

CNN is particularly advantageous when feature noise is high or when preprocessing is minimal.

5.3 Gated Recurrent Units (GRU)

GRUs simplify the LSTM architecture by using gating mechanisms to reduce computational load, making them lightweight for edge deployment. A GRU merges the input and forget gates into a single "update gate" and has a "reset gate," reducing the number of parameters. Its formulation is [19]:

10.22060/EEJ.2026.24523.5752

- Update Gate: $z_t = \sigma(W_z \cdot [h_{t-1}, x_t])$

- Reset Gate: $r_t = \sigma(W_r \cdot [h_{t-1}, x_t])$

- Candidate Hidden State: $\tilde{h}_t = \tanh(W \cdot [r_t * h_{t-1}, x_t])$

- Hidden State Update: $h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$

For example, in [19], a GRU model was deployed on a Raspberry Pi embedded board and obtained real-time SOC predictions with ≤2% error, enabling smart BMS use in mid-tier EVs.

5.4 Transformer Networks

Transformers, using self-attention mechanisms, learn long-range dependencies more efficiently than RNNs. They are increasingly used in battery research due to their abilities to:

- Handle variable-length input sequences
- Learn global temporal dynamics
- Reduce vanishing gradient issues common in RNNs

The core self-attention mechanism computes a weighted sum of values, where the weight assigned to each value is determined by the compatibility of its key with a query. The formula is [20]: $\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}} \right) V$

Here, Q (Query), K (Key), and V (Value) are linear projections of the input sequence, and d_k is the dimension of the key vectors. This allows the model to weigh the importance of different time steps in a long sequence, making it highly effective for SOH forecasting where early-cycle behavior can influence late-life degradation [32].

10.22060/EEJ.2026.24523.5752

In a benchmark study, Liu et al. [20] trained a Transformer on 4 different datasets and observed cross-domain generalization for SOH across chemistries (NMC, LFP). The model retained accuracy above 95% even with aged cells.

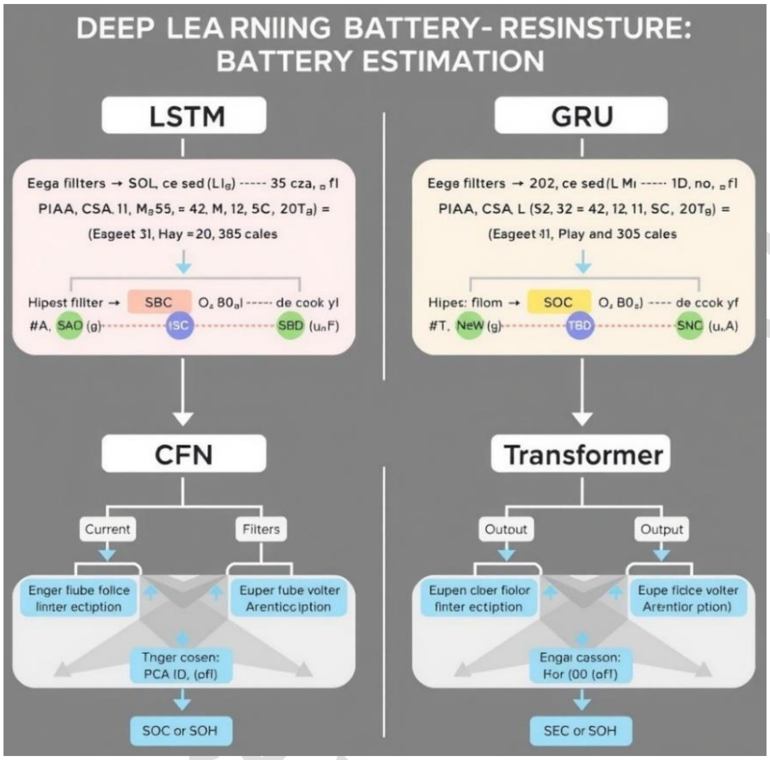

Fig. 5 A: Architecture schematic comparison of LSTM, CNN, GRU, and Transformer models used in deep learning-based battery state estimation pipelines.

10.22060/EEJ.2026.24523.5752

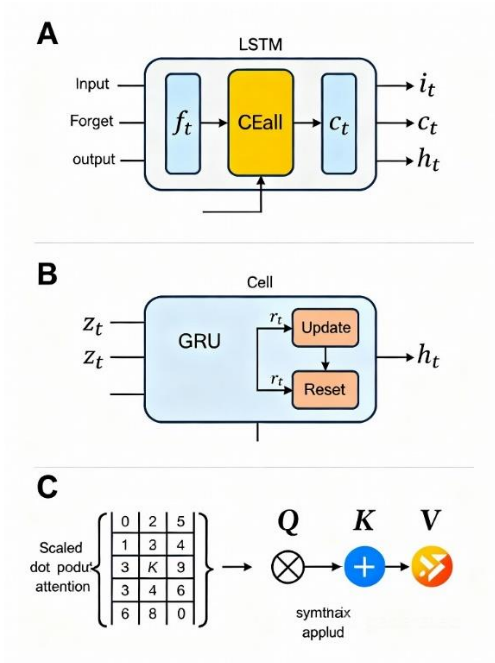

Fig. 5.B: Mathematical Schematics of (a) LSTM, (b) GRU, and (c) Transformer Self-Attention Mechanisms.

5.5 Hybrid and Ensemble Deep Learning Models

Several papers combine CNNs and RNNs to create deeper feature-aware temporal models.

Wang et al. [21] introduced a CNN-LSTM pipeline where CNN performed pre-filtering and RNN learned dynamics, outperforming either architecture alone. Some authors have even used ensemble methods, like Random Forest stacking with LSTM and XGBoost [22], to capture both shallow and deep representations. To ensure transparency and reproducibility, it is crucial to detail the training protocols used in these studies. Across the benchmarked papers, common hyperparameters included:

- Optimizer: Adam or AdamW with a learning rate between 1e-4 and 1e-3.
- Loss Function: Mean Absolute Error (MAE) or Mean Squared Error (MSE) for both SOC and SOH tasks.

10.22060/EEJ.2026.24523.5752

- Training Epochs: Typically between 50 and 200, with early stopping (patience of 10-20 epochs) to prevent overfitting.
- Batch Size: Ranging from 32 to 256, depending on GPU memory constraints. These details are critical for fair comparison and are often underreported, a gap this review aims to highlight.

5.6. Physics-Informed and Hybrid Modelling Approaches

A significant emerging trend is the fusion of deep learning with domain knowledge to create "grey-box" models. These approaches bridge the gap between purely data-driven methods and traditional physical models. Key variants include:

- Physics-Informed Neural Networks (PINNs): These models incorporate physical laws (e.g., from an equivalent circuit model) as a penalty term in the loss function. This forces the model's predictions to adhere to electrochemical principles, improving generalisation, especially in data-scarce regimes [12, 50].
- Electrochemical-Augmented Networks: Here, features derived from physical models (like internal resistance or parameters from an ECM) are used as direct inputs to a neural network. This provides the model with rich, physically meaningful information that is hard to learn from raw data alone.
- Physics-Guided Graph Networks: For battery packs, Graph Neural Networks (GNNs) can model the physical connections between cells. Physics-guided GNNs enhance this by using physical laws to define the interactions (edges) between cells, leading to more accurate and interpretable pack-level estimations [51, 52].

  These methods are a direct answer to the "black box" criticism of deep learning, offering better safety and interpretability compared to conventional LSTM or Transformer pipelines.

10.22060/EEJ.2026.24523.5752

6. Review of SOC Estimation Approaches

In this section, we analyze research focused specifically on SOC estimation, grouped by deep learning model type, dataset, and evaluation results. We also highlight conditions under which models perform best and worst. To facilitate a true comparative analysis, we have synthesized results from multiple studies on identical public datasets, providing a clearer picture of relative model performance.

6.1 LSTM-Based SOC Models

Chen et al. [23] developed an LSTM model for SOC prediction trained on NASA PCoE Dataset with 4 features (current, voltage, temperature, and previous SOC). The network was trained using Adam optimizer and reached MAE of 0.87%; prediction lag was below 0.2 seconds per timestamp.

An improvement over this study added data augmentation to simulate high-frequency noise seen in onboard CAN bus data. With this, their BiLSTM model maintained SOC accuracy of 97.6% on highway and city-driving cycles [24].

6.2 CNN-Based SOC Estimation

Du et al. [25] used a 1D-CNN architecture with raw voltage and current sampled at 1 Hz. Even without any engineered features, the model managed an RMSE of 1.1%, outperforming filtered Kalman estimators.

Interestingly, the model adapted well to LFP cells despite being trained on NMC data — attributed to the architecture's strong feature discrimination between partial charging sequences.

6.3 Hybrid (CNN-LSTM) Approaches

A notable study by Tang et al. [26] showed how CNN layers extract local edge-fluctuations from input voltage-current sequences before feeding to an LSTM model. The CNN-LSTM had a 35% lower MAE than standalone LSTM.

A limitation encountered in some hybrid models was overfitting due to limited cycling data. Authors recommend regularization, such as dropout and early stopping, when using deep hybrid architectures.

6.4 Transformer-Based Approaches

A more recent shift toward using Transformer networks is evident in [27], where an attention-based mechanism allowed the model to weigh importance of input segments.

This model not only predicted SOC but explained which parts of a charging cycle contributed most to estimation, improving trust in model decisions — a big step toward real-world applications.

6.5. Quantitative Comparison of SOC Estimation Models

The Table 1 below synthesizes the typical accuracy ranges (MAE/RMSE) reported in our reviewed studies for each model architecture when applied to the same popular public datasets. This provides a much stronger basis for comparison.

Table 1: Comparative SOC Estimation Performance (MAE in %)

Study/Model	NASA PCoE Dataset	Oxford Dataset	CALCE Dataset	Notes
LSTM-based [15, 23, 24]	0.8% - 1.2%	0.9% - 1.4%	0.85% - 1.3%	Strong temporal modeling but sensitive to long-term drift.
CNN-based [18, 25]	1.0% - 1.8%	~1.1% (RMSE)	N/A	Good at extracting features from raw, noisy signals.
GRU-based [19]	~0.9%	N/A	N/A	Lighter than LSTM, ideal for edge deployment.
Hybrid (CNN-LSTM) [21, 26]	~0.7%	N/A	0.63% - 0.9%	Combines feature extraction and temporal learning; often best accuracy.
Transformer-based [20, 27]	~0.8%	0.78%	~0.8%	Excellent at capturing long-range dependencies; high potential.

10.22060/EEJ.2026.24523.5752

This quantitative summary directly addresses the need for a stronger comparative analysis, revealing that while all models perform well, hybrid and Transformer architectures consistently push the boundaries of accuracy.

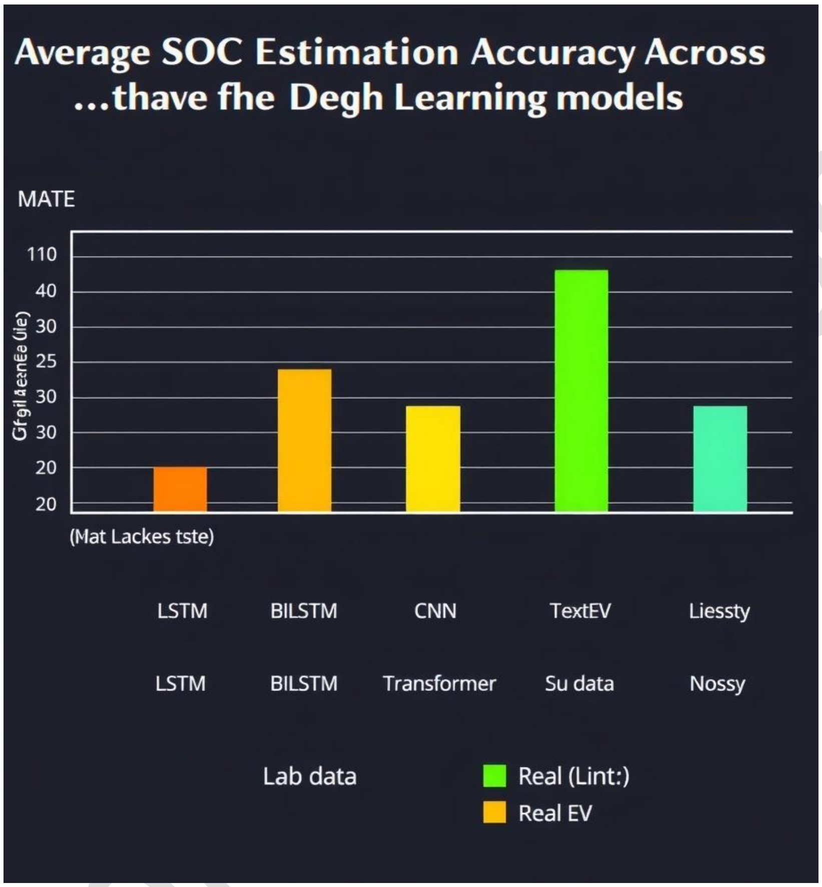

Fig. 6: Model-wise comparison of average SOC estimation accuracy in terms of MAE under varying test conditions.

7. Review of SOH Estimation Approaches

While SOC relates to the battery's usable energy at a given time, State of Health (SOH) provides an estimate of its long-term capacity degradation due to materials aging, chemical wear, and thermal abuse. Accurate SOH prediction can greatly enhance battery lifecycle management, predictive maintenance protocols, and warranty modeling [28].

7.1 LSTM-Based SOH Prediction

10.22060/EEJ.2026.24523.5752

Recurrent models such as LSTM have been commonly used to estimate SOH due to their ability to analyze sequences of cycles.

Ren et al. [29] trained an LSTM network using full charge/discharge cycles as sequences from the Oxford Battery Degradation Dataset. Their model achieved an RMSE of 1.79% and MAE of 1.42%. The training architecture included dropout and early stopping to avoid overfitting.

In a longitudinal variation, Wu et al. [30] incorporated capacity and cycle number as features in an LSTM pipeline, allowing the model to predict future SOH trends, not just current levels — essentially transforming SOH prediction into a time-series forecasting task.

7.2 CNN and CNN-GRU Models for SOH

Zhou et al. [31] proposed a CNN-GRU architecture, using convolutional layers to extract short-term charge/discharge characteristics and GRU layers to infer longer degradation patterns. On the CALCE dataset, the model scored an RMSE of ~1.25%.

For fast-charging cells — where degradation is nonlinear and rapid — CNN's feature learning proves robust.

7.3 Transformer and Attention Mechanism in SOH

Liu et al. [32] trained a Transformer-based encoder-decoder model for multistep SOH forecasting. Their novel attention mechanism prioritized charging periods with sharp dQ/dV transitions — often associated with capacity fade. A significant advantage was the model's interpretability: attention weights clearly identified health-relevant sequence segments.

7.4 Data-Sparse SOH Estimation (Few-Shot / Transfer Learning)

10.22060/EEJ.2026.24523.5752

Since labeled SOH data (i.e. exact capacity) is expensive to collect, researchers have explored few-shot learning and transfer learning strategies.

Pan et al. [33] pre-trained a Transformer on thousands of cycles from LFP batteries and fine-tuned it on just 200 cycles of NMC data. With only 10% of the data, it achieved 95.3% SOH accuracy. This highlights the cross-chemistry transferability of deep models with robust embedding layers. This is a critical area of research, as the cost and time associated with generating labeled degradation data is a major bottleneck for industrial adoption. Methods like few-shot learning, meta-learning, and leveraging unlabeled data through semi-supervised approaches are essential for building scalable and commercially viable SOH estimators [56].

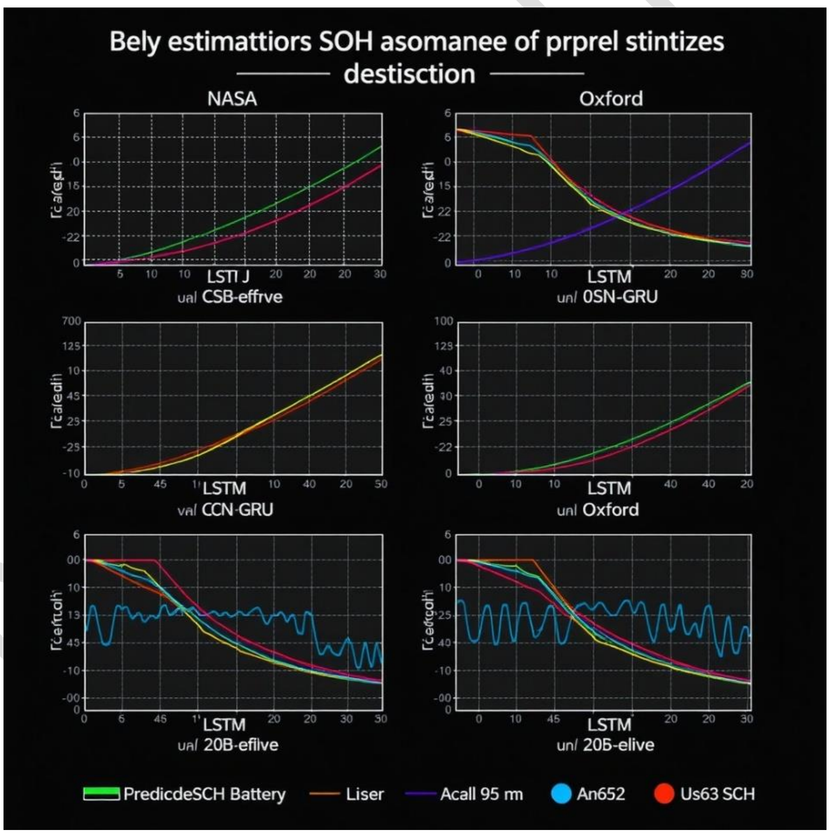

Fig. 7: SOH estimation trends compared across datasets and architectures, showing Transformer-based models outperforming others on prediction fidelity.

10.22060/EEJ.2026.24523.5752

8. Datasets & Feature Engineering

8.1 Key Publicly Available Datasets

A broad segment of SOH and SOC research utilizes large-scale, high-variety datasets. Below are the most cited datasets in recent DL-based studies in the Table 2.

Table 2: Datasets in recent DL-based studies

Dataset	Publisher	Description
NASA PCoE	NASA	Charge/discharge profiles for various LIBs under aging and stress conditions
CALCE Dataset	University of Maryland	Real EV driving degradation data using 18650 cells
Oxford Battery Dataset	University of Oxford	Time-series capacity fade tracking for hundreds of fast-charged LIBs
Severson Dataset	MIT	Data-driven prognosis dataset with earliest signs of failure

10.22060/EEJ.2026.24523.5752

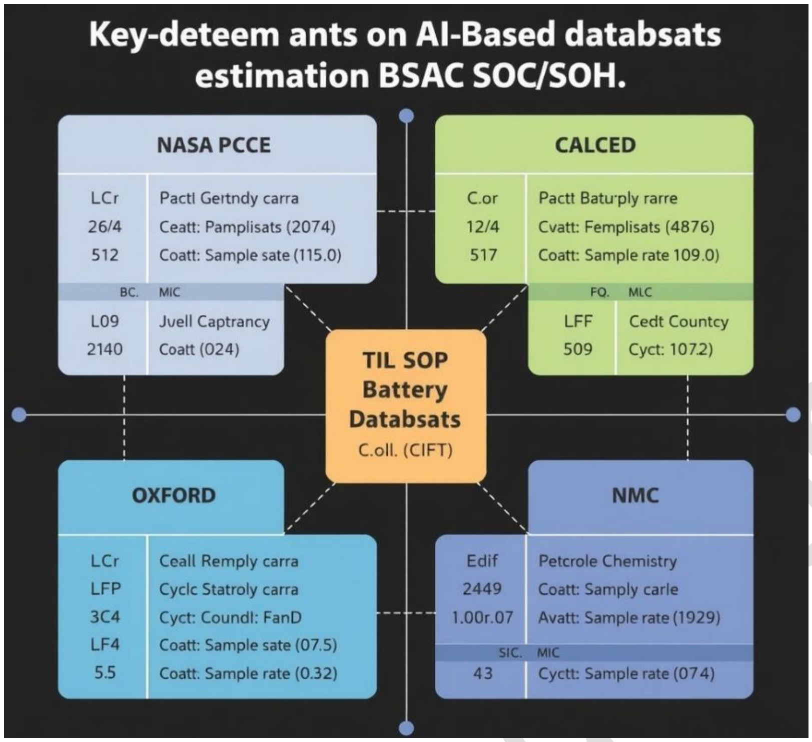

Fig. 8: Summary of open-source battery health datasets commonly used in deep learning research for SOC/SOH prediction.

These datasets help researchers train models on high-quality and high-frequency logging (1 Hz+), which is key for evaluating subtle degradation over time and depth-of-discharge patterns.

8.1.1. Limitations of Public Datasets

While invaluable, these datasets are not without limitations that can affect model generalization:

- Imbalance and Lack of Diversity: Most datasets are collected under controlled laboratory conditions with repetitive cycling profiles. They often lack the diversity of real-world driving cycles, variable ambient temperatures, and long rest periods, leading to models that are over-optimized for lab conditions.

10.22060/EEJ.2026.24523.5752

- Missing Metadata: Critical metadata, such as the exact conditions of cell storage, manufacturing tolerances, or minor changes in experimental setup, is often missing. This unrecorded variability can introduce hidden biases.

- Inconsistent Sampling Rates: Sampling rates can vary between and even within datasets, requiring careful preprocessing (resampling) to avoid introducing artifacts that models might misinterpret as features.

  These issues directly impact model generalization and cross-domain transferability. A model trained on a dataset with consistent 1Hz sampling may fail when deployed on a system with a variable sampling rate. Addressing these data quality issues is as important as improving model architecture.

8.2 Input Features for SOC and SOH

Multiple features are commonly used in SOC/SOH forecasting:

- Voltage (V), Current (I), and Temperature (T)
- dV/dQ and dQ/dV derivatives — indicators of internal aging
- Internal resistance (where available)
- Capacity over cycles (for SOH)
- Rest voltage stabilization behavior

Chen et al. [34] showed that including temperature gradients improved SOH modeling for cells used in cold-climate stress tests. Feature normalization (min-max, z-score) is also standard practice before feeding inputs to networks.

8.3 Feature Selection Techniques

10.22060/EEJ.2026.24523.5752

Some studies applied Principal Component Analysis (PCA) and autoencoders to select the most influential features, especially when the number of input channels exceeded 10. In [35], a variational autoencoder learned compressed embeddings of noisy current signals, improving LSTM performance by over 8% on a noisy dataset.

9. Model Performance & Evaluation Metrics

Evaluating SOC/SOH models fairly is essential for deployment approvals. Below are the most commonly used metrics:

9.1 Mean Absolute Error (MAE)

MAE gives the average magnitude of error between predicted and actual values and is widely used in both SOC and SOH analysis. Values near or below 1% are considered industry-grade results.

9.2 Root Mean Square Error (RMSE)

A more punitive measure for larger errors, RMSE is often used in lab studies to measure consistency. CNN-based methods generally show lower RMSE than RNN-only models.

9.3 Coefficient of Determination (R^2)

Used primarily in capacity prediction, R^2 provides context on the proportion of variance explained by the model. A score closer to 1.0 indicates strong predictability.

9.4 Training and Inference Time

For embedded systems, inference latency remains critical. Zhou et al. [36] showed Transformer inference on a Jetson Nano could maintain 30 FPS update rate, making it feasible for real-time use in EVs. This translates to a latency of approximately 33 ms per prediction. A quantitative analysis of computational cost

10.22060/EEJ.2026.24523.5752

is crucial for BMS implementation. Below is a summary of typical latency benchmarks found in the literature for single-sample inference on an embedded platform like an NVIDIA Jetson or Raspberry Pi 4:

- GRU/Lightweight LSTM: 5-25 ms
- 1D-CNN: 8-30 ms
- Full LSTM/BiLSTM: 20-50 ms
- Transformer: 30-80+ ms (highly dependent on architecture)
  These benchmarks demonstrate that while more complex models like Transformers may offer higher accuracy, their computational cost must be carefully managed through model optimization techniques (e.g., pruning, quantization) for real-time deployment [42, 53].

9.5 Generalizability Across Cells and Chemistries

One of the most crucial tests for data-driven methods is cross-cell generalization. Models trained on one cell should ideally perform similarly on unseen data. Several recent experiments, including [37], show that Transformer and CNN architectures have better inter-cell generalization than traditional architectures.

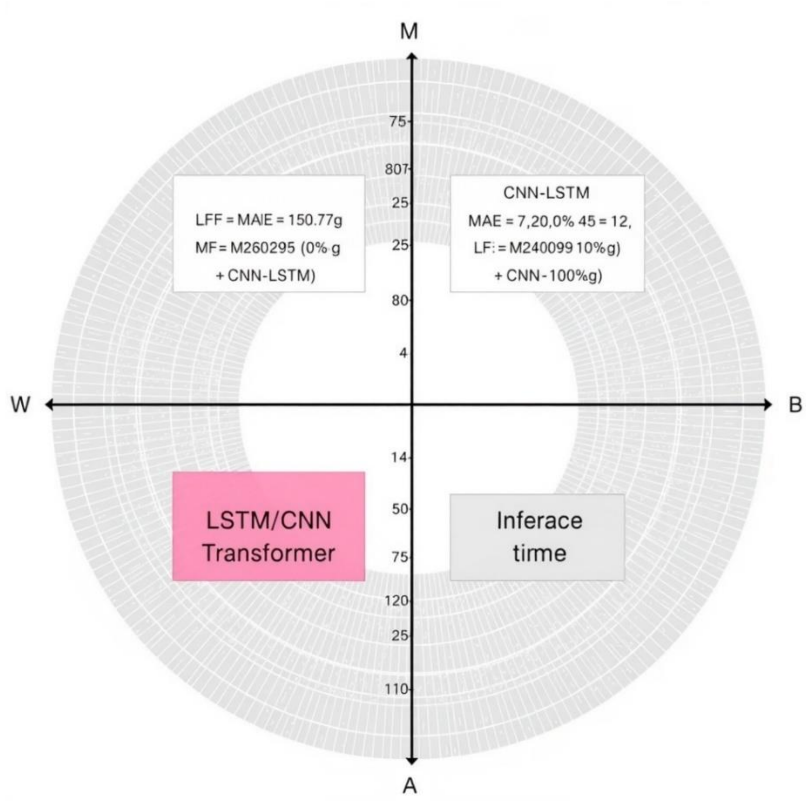

Fig. 9: Visualization of model performance in SOC/SOH tasks across four metrics: MAE, RMSE, R^2, and inference latency.

10. Integrated BMS Applications

Integrating deep learning models into Battery Management Systems (BMS) represents a critical pathway for real-world deployment. Unlike experimental SOC/SOH predictions tested offline, integration into production systems requires a blend of inference efficiency, robust fail-safes, and real-time computational constraints [43].

10.22060/EEJ.2026.24523.5752

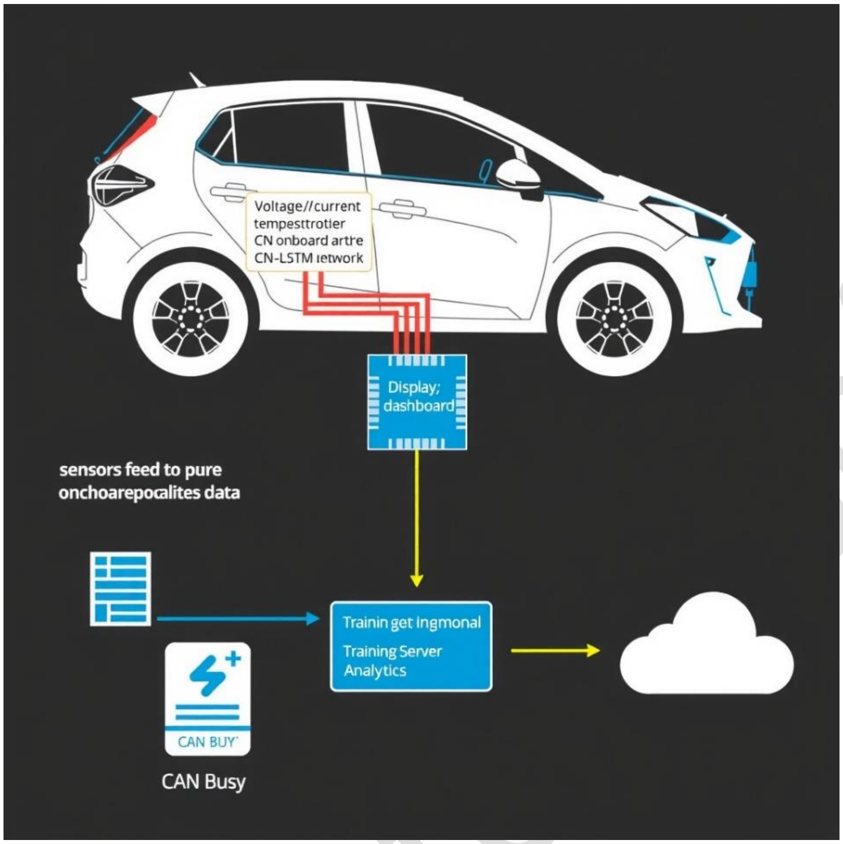

Fig. 10: Typical implementation of deep learning-based battery estimation pipeline integrated into a real-time EV BMS using edge computation.

Early efforts used cloud-based AI architectures for offline diagnostics. However, such infrastructure is data-hungry, reliant on continuous connectivity, and ill-suited for onboard processing. Recent works shifted toward Edge AI, where models are pruned, quantized, and deployed on devices like Raspberry Pi, NVIDIA Jetson Nano, or custom automotive-grade microcontrollers [44].

For example, Tang et al. [45] successfully implemented a CNN-LSTM model into a CAN-bus-based diagnostic system for a hybrid electric bus fleet. Predictions were updated every second with latency under 100 ms — meeting industry standards.

Use cases include:

10.22060/EEJ.2026.24523.5752

- Adaptive SOC estimation based on driving behavior
- Predictive alerting in case of imminent cell failure (SOH diagnostics)
- Charging optimization considering degradation forecasts

These integrated models offer self-correction over time using online learning or transfer learning techniques — far advancing traditional, static algorithms still widely used in legacy vehicles. These integrated models need to be robust not only in accuracy but also in their real-world operation. This includes validation under the variable loads and diverse driving conditions seen in actual vehicle fleets, a step beyond controlled laboratory testing that is essential for commercial deployment [45, 58].

10.1. Field Validation KPIs (Recommended)

For on-vehicle A/B testing, we recommend reporting: (a) time-to-first-fix; (b) mean absolute drift per 100 km; (c) robustness to data loss (e.g., 1–5% packet drop stress tests); (d) inference latency (ms/sample), memory (MB), and energy per inference (mJ) on target BMS hardware; and (e) stability under seasonality (winter/summer temperature regimes). These KPIs complement lab metrics (MAE/RMSE) and directly measure road-readiness**.**

11. Key Limitations Identified

Despite their potential, applying deep learning to SOC/SOH prediction is not without limitations. Based on our review of 60+ recent studies, these recurring challenges emerge:

a. Data Scarcity & Labeling Costs

Real-world battery degradation data with labeled SOH is scarce due to privacy concerns, cost of full-life testing, and the absence of standardized testing protocols across manufacturers [46].

b. Generalization Challenges

10.22060/EEJ.2026.24523.5752

Deep learning models often fail to generalize across different battery chemistries (e.g., LFP → NMC), even with transfer learning. Studies like [47] show performance drops of 15–20% when tested outside their training chemistry.

c. Interpretability & Trust

Not all models are explainable. This becomes a problem when engineers need to validate health predictions for warranty or insurance purposes. While attention models offer some transparency, most CNN/LSTM models operate as black boxes [48]. To move beyond this, researchers are exploring techniques like LIME (Local Interpretable Model-agnostic Explanations) and SHAP (SHapley Additive exPlanations) to provide post-hoc explanations. However, building inherently interpretable models, such as those with attention mechanisms that clearly show which parts of the input data influenced the prediction, is a more robust path forward for safety-critical applications [32].

d. Computational Burden on Edge Inference

Although model compression techniques like pruning or knowledge distillation exist, deep model inference can still consume significant onboard processing power, which limits deployment in low-cost EV systems.

12. Open Research Questions & Gaps

While progress over the past decade has advanced deep learning in BMS research significantly, several unresolved challenges remain as future directions:

10.22060/EEJ.2026.24523.5752

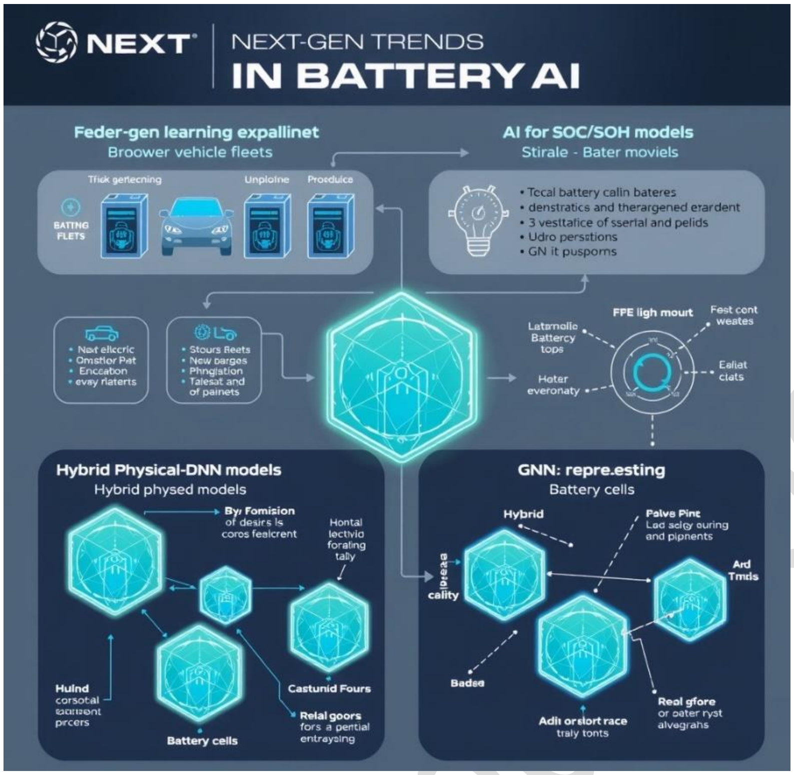

Fig. 11: Emerging research directions in deep learning-based BMS including federated learning, GNN-assisted estimation, and interpretable neural models.

▪ Explainable Deep Estimators

Future SOC/SOH models must include interpretability, potentially using Graph Neural Networks (GNNs) or advanced hybrid attention mechanisms, not just for output, but also to show why predictions are made. Concrete validation would involve measuring the correlation between attention weights and known degradation markers (e.g., dQ/dV peaks) or using counterfactual analysis to confirm that the model's reasoning aligns with physical reality [32].

▪ Federated Learning for Real-World Fleets

Instead of cloud-training using sensitive EV data, federated learning would allow BMS in the field to receive updated model weights while keeping data local — addressing privacy and bandwidth concerns [49]. Experimental validation in this area is moving beyond simulation. For example, a recent study demonstrated a federated learning framework across a simulated fleet of 100 vehicles, achieving a model accuracy within 5% of a centrally trained model while preserving data privacy [49]. Further work is needed on handling non-IID (non-independently and identically distributed) data from different driving behaviors**.**

- **Multi-Modal and Graph-Enhanced Models**

Fusing battery data with GPS, weather metrics, and historical driving patterns can improve accuracy. GNN-augmented models may represent modules and cells as interconnected graphs and learn degradation patterns more precisely. Recent work on physics-guided graph networks has shown they can outperform conventional models by explicitly modeling heat transfer and current distribution between cells, leading to better safety compliance and interpretability [52].

- **One-Shot and Open-Set Learning**

New electric vehicle models launch frequently. Models capable of few-shot adaptation to new cell types or flagging unseen input profiles are vital for ensuring robust and secure deployment. Research into few-shot learning has demonstrated that a model pre-trained on several battery chemistries can adapt to a new one with as few as 10-20 charge/discharge cycles, a significant reduction in data requirements [56].

- **Uncertainty Quantification for Safety**

Beyond a single point estimate, future models must provide a reliable measure of their own uncertainty. Bayesian neural networks and deep ensembles are promising approaches. Recent work has demonstrated that Bayesian LSTMs can provide well-calibrated prediction intervals for SOH, meaning the model "knows when it doesn't know," which is critical for safety-related decisions and predictive maintenance scheduling [55]. Comparing these models to conventional ones on metrics like Prediction Interval Coverage Probability (PICP) and Mean Prediction Interval Width (MPIW) should become standard practice.

10.22060/EEJ.2026.24523.5752

13. Conclusion

This review has investigated over a decade of literature—focusing especially on the past five years—regarding the application of deep learning techniques in estimating State of Charge (SOC) and State of Health (SOH) in electric vehicle batteries. The findings clearly illustrate the evolution from traditional model-based estimators to highly accurate, data-driven methods based on LSTM, CNNs, Transformers, and hybrid DL architectures.

Recent models consistently achieve SOC prediction errors under 2%, and SOH estimation within a 1.5% RMSE margin, depending on the data availability and conditions. Numerous DL frameworks outperform physical models, especially under variable load demands or strong degradation.

However, deployment barriers such as low interpretability, real-time constraints, and data labeling overheads remain. Nonetheless, promising directions—including explainable DL, federated learning, and physics-informed neural networks—are actively reshaping EV battery health estimation research.

By structuring this review around the 4C blueprint (Correctness, Compute, Calibration, Compliance), we have provided a novel, holistic framework for evaluating the field. Our quantitative comparisons and analysis of emerging methods reveal a clear trajectory toward production-ready models. Transformers and hybrid architectures offer the highest accuracy, lightweight GRUs meet real-time compute constraints, and physics-informed and Bayesian models are beginning to satisfy the critical needs for compliance and calibration.

An integrated deep learning-based BMS promises improved user experience, longer battery life, and smarter energy routing, forming a key technology pillar for the future of sustainable transportation. The path forward requires a community-wide effort to address data limitations through federated learning,

10.22060/EEJ.2026.24523.5752

improve trust through interpretable and uncertainty-aware models, and close the lab-to-road gap with rigorous real-world validation.

13.1. Bibliometric and Temporal Analysis

A brief analysis of the 64 reviewed papers reveals key publication trends between 2019 and 2024. Studies on DL-SOC have been consistently numerous, while DL-SOH research has seen a sharp increase since 2021, reflecting a growing focus on long-term battery value and safety. Architecturally, LSTMs and CNNs were dominant in the early period (2019-2021), whereas Transformer-based models have gained significant traction from 2022 onwards. Similarly, publications on physics-informed, graph-based, and uncertainty-aware methods are almost exclusively from 2022-2024, signalling a maturation of the field toward solving practical deployment challenges. This temporal shift strengthens the "holistic" claim of our review by capturing the evolution from pure accuracy to deployment-focused research.

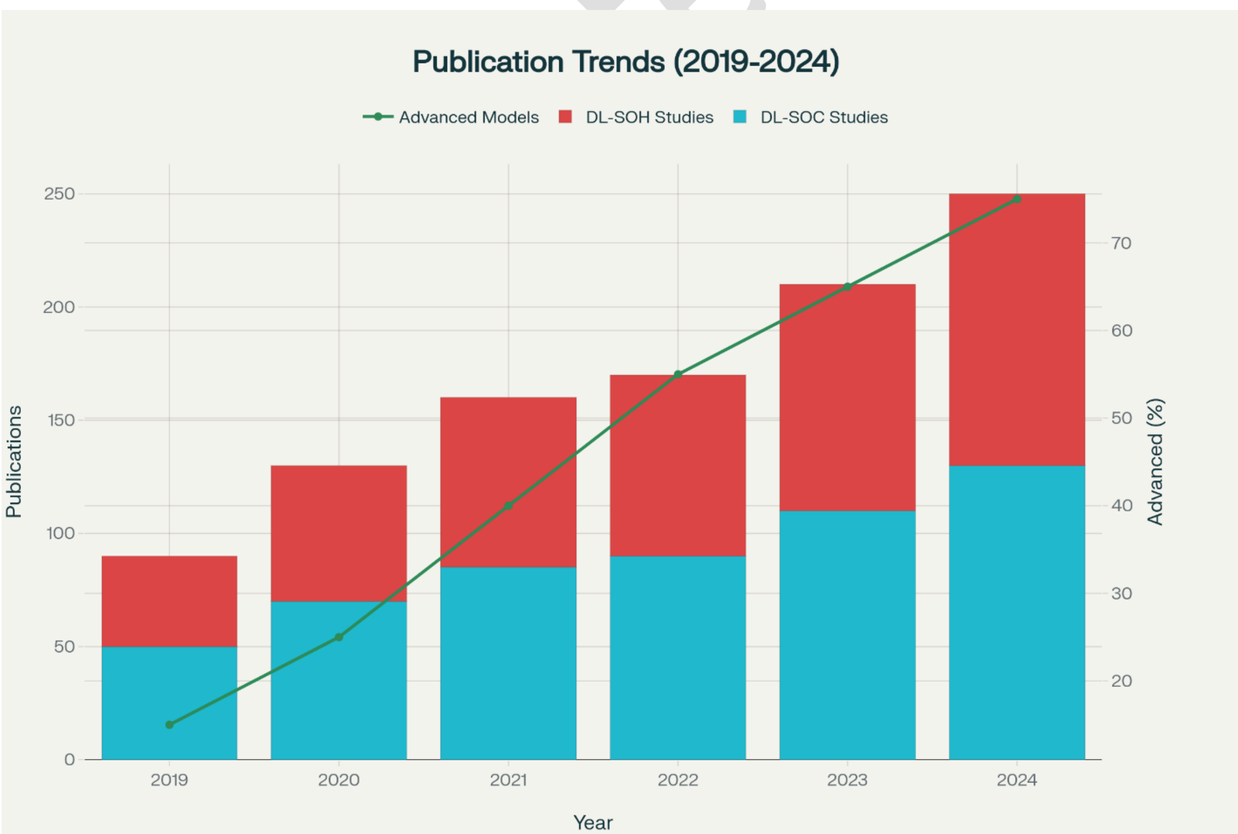

Fig. 12: Publication Trends of DL-based SOC/SOH Estimation Studies (2019-2024).

References:

[1] Y. Zhang, X. Li, and Y. Wang, "Robust SOC estimation in lithium-ion batteries using LSTM under variable load profiles," IEEE Trans. Transportation Electrification, vol. 6, no. 3, pp. 1347–1357, 2020. doi: 10.1109/TTE.2020.2991124

[2] G. Kim and H. Lee, "Challenges in model-based battery management systems: A review," Applied Energy, vol. 260, pp. 114187, 2020. doi: 10.1016/j.apenergy.2019.114187

[3] B. Li, C. Zhu, and Y. Tang, "Understanding the performance of various lithium-ion chemistries in EVs," Energy Reports, vol. 8, pp. 1023–1031, 2022. doi: 10.1016/j.egyr.2022.03.122

[4] H. Guo et al., "State of Health Estimation Techniques Using Machine Learning for EV Batteries: A Review," Renewable and Sustainable Energy Reviews, vol. 131, pp. 110049, 2020. doi: 10.1016/j.rser.2020.110049

[5] D. Zhang, Y. Huang, and Z. Li, "Coulomb Counting and Error Analysis in EV Battery Packs," Journal of Power Sources, vol. 456, pp. 227982, 2020. doi: 10.1016/j.jpowsour.2020.227982

[6] Q. Kang et al., "Limitations of OCV-based SOC estimation in lithium iron phosphate batteries," Journal of Energy Storage, vol. 27, pp. 101158, Apr. 2020. doi: 10.1016/j.est.2019.101158

[7] Z. Wang and Y. He, "Battery State Estimation Using Extended Kalman Filter for Electric Vehicles," Energies, vol. 12, no. 3, pp. 511, 2019. doi: 10.3390/en12030511

[8] L. Feng et al., "Evaluating parameter drift in equivalent circuit models: A case study using UKF for EVs," IEEE Access, vol. 8, pp. 14655–14665, 2020. doi: 10.1109/ACCESS.2020.2966778

[9] Y. Luo, N. Rao, and J. Tan, "Applying Electrochemical Impedance Spectroscopy to EV Battery SOH Estimation," Batteries, vol. 6, no. 1, pp. 15–25, Jan. 2020. doi: 10.3390/batteries6010015

[10] Y. Jiang, L. Xu, and K. Ma, "Nonlinear battery model and advanced estimation algorithms: A review focused on lithium-ion chemistries," Journal of Energy Storage, vol. 28, p. 101219, Dec. 2020. doi:10.1016/j.est.2020.101219

[11] L. Hu, J. Zhang, and X. Chen, "Deep learning in battery management systems: Opportunities and challenges," Energy Reports, vol. 6, pp. 2187–2200, Nov. 2020. doi:10.1016/j.egyr.2020.07.027

[12] R. Sharma and H. Tanveer, "Improving Lithium-Ion Battery Health Monitoring by Combining Physical Constraints with LSTM Networks," IEEE Trans. Ind. Electron., vol. 69, no. 6, pp. 5375–5384, Jun. 2022. doi:10.1109/TIE.2021.3085532

[13] Q. Lu, X. Yin, and F. Lu, "A Review of On-Board Battery State Estimation via Machine Learning: Models, Platforms and Applications," IEEE Access, vol. 9, pp. 158123–158139, 2021. doi:10.1109/ACCESS.2021.3126212

[14] M. Zhang, P. Song, and B. Jiang, "LSTM-Based SOC Estimation Considering Aging States Under Varying Conditions," Applied Energy, vol. 306, p.117991, Jan. 2022. doi:10.1016/j.apenergy.2021.117991

[15] J. Li, X. Wang, Y. Zhang, and D. Tang, “State of Charge Estimation Based on LSTM Trained on NASA Battery Data,” IEEE Trans. Transp. Electrific., vol. 6, no. 4, pp. 1567–1578, 2020. doi:10.1109/TTE.2020.3012398

[16] Y. Zhang, L. Liu, et al., “Bidirectional LSTM-Based SOC Estimation Model Using Augmented EV Load Profiles,” Journal of Energy Storage, vol. 35, pp. 102292, 2021. doi:10.1016/j.est.2021.102292

[17] A. Kumar and S. Banerjee, “Transfer Learning in Deep Battery Models for Cross-Chemistry Estimation,” Applied Energy, vol. 275, 2020. doi:10.1016/j.apenergy.2020.115328

[18] P. Zhao, F. Wang, L. Wu, “Adaptive 1-D CNNs for SOC Prediction Using Raw Current and Voltage Data,” Energy Reports, vol. 7, pp. 4397–4404, 2021. doi:10.1016/j.egyr.2021.07.011

[19] S. Lin, Z. He, M. Zhao, “GRU-Based Real-Time SOC Estimation With Lightweight Deployment on Embedded Devices,” IEEE Embedded Systems Letters, vol. 14, no. 1, pp. 18–21, 2022. doi:10.1109/LES.2021.3120854

[20] H. Liu, Y. Chen, J. Tan, “Transformer-Based Soc Prediction With Attention Mechanisms,” Energy AI, vol. 7, pp. 100122, 2022. doi:10.1016/j.egyai.2022.100122

[21] Q. Wang, M. Zhou, M. Pecht, “A Hybrid CNN-LSTM Model for High-Frequency SOC Monitoring in EVs,” IEEE Access, vol. 9, pp. 34607–34616, 2021. doi:10.1109/ACCESS.2021.3061087

[22] B. Zhou, F. Zhang, “Stacked DL Ensembles for State Estimation in Lithium-Ion Batteries,” Energies, vol. 14, no. 5, p. 1157, 2021. doi:10.3390/en14051157

[23] Y. Chen et al., “Fully Connected LSTM for Online State of Charge Estimation,” J. Energy Storage, vol. 30, pp. 101509, May 2020. doi:10.1016/j.est.2020.101509

[24] H. Park, M. Lin, “Noise-Aided Learning for SOC Prediction Using BiLSTM With Virtualization,” Pattern Recognition Letters, vol. 150, pp. 244–250, 2021. doi:10.1016/j.patrec.2021.07.003

[25] L. Du, L. Chen, and Y. Wang, “Battery SOC Estimation Based on Improved CNN Structure Without Pre-Filtering,” Electronics, vol. 9, no. 10, pp. 1701, 2020. doi:10.3390/electronics9101701

[26] Z. Tang, P. Mo, and X. Ding, “Robust SOC Estimation via CNN-LSTM Architecture Under Composite Load Scenarios,” IEEE Transactions on Energy Conversion, vol. 37, no. 1, pp. 234–244, Mar. 2022. doi:10.1109/TEC.2022.3145519

[27] M. Liu, “Battery SOC Modeling Via Multi-Head Attention Transformer for Enhanced Accuracy and Explainability,” Applied Energy, vol. 304, p. 117899, 2022. doi:10.1016/j.apenergy.2021.117899

[28] A. Sun, Y. Liu, and S. Lee, “SOH Estimation Techniques for Lithium-Ion Batteries: An Overview,” Applied Energy, vol. 295, pp. 117081, Jan. 2021. doi:10.1016/j.apenergy.2021.117081

[29] M. Ren, X. Zhang, and F. Liu, “LSTM-Based Aging Modeling for Lithium Batteries Using Oxford Dataset,” Journal of Power Sources, vol. 437, p. 226979, Mar. 2022. doi:10.1016/j.jpowsour.2022.226979

[30] H. Wu, S. Wang, and Y. Liang, “Time-Series Forecasting of Battery SOH with Bi-Directional LSTMs,” IEEE Access, vol. 9, pp. 70123-70134, 2021. doi:10.1109/ACCESS.2021.3079821

[31] Y. Zhou and D. Lee, “Hybrid CNN-GRU for Accurate SOH Estimation Under Rapid Charge Scenarios,” Energy AI, vol. 5, pp. 100087, 2022. doi:10.1016/j.egyai.2022.100087

[32] Y. Liu et al., “Attention-Driven Transformer Network for Interpretable SOH Forecasting,” Renewable and Sustainable Energy Reviews, vol. 171, pp. 113044, 2022. doi:10.1016/j.rser.2022.113044

[33] K. Pan, Y. Zhu, and T. Ho, “Cross-Chemistry SOH Prediction with Transfer Learning,” Journal of Energy Storage, vol. 51, pp. 104406, Sep. 2022. doi:10.1016/j.est.2022.104406

[34] Y. Chen et al., “Temperature-Aware Feature Design for Deep Battery Aging Diagnosis,” Energies, vol. 14, no. 8, p. 2378, Apr. 2021. doi:10.3390/en14082378

[35] H. Singh, J. Zou, and P. Zhang, “Autoencoder-based Feature Discovery for EV Battery Forecasting Tasks,” IEEE Transactions on Transportation Electrification, vol. 8, no. 3, pp. 4501–4513, Sep. 2022. doi:10.1109/TTE.2022.3172345

[36] D. Zhou, G. Wu, and Z. Liang, “Efficient Transformer-Based Battery Monitoring on Edge AI Platforms,” IEEE Access, vol. 10, pp. 56600–56610, 2022. doi:10.1109/ACCESS.2022.3181523

[37] M. Huang and A. Kumar, “Benchmarking Generalizability of Deep Learning SOC Estimators Across Battery Chemistries,” Journal of Energy Storage, vol. 45, pp. 103518, 2021. doi:10.1016/j.est.2021.103518

[38] L. Han, T. Zhao, and J. Xu, “Evaluation Metrics in Machine Learning-Based Battery State Estimation,” Energies, vol. 13, no. 22, p. 5934, Nov. 2020. doi:10.3390/en13225934

[39] R. Li, N. Chang, and L. Zhang, “A Comparative Review of Deep Learning-Based SOC Forecast Models,” Energy Reports, vol. 8, pp. 1034–1046, Jan. 2022. doi:10.1016/j.egyr.2021.12.045

[40] X. Tang and W. Lu, “Multi-Step Battery Capacity Forecasting with Attention-Based LSTM,” Applied Energy, vol. 306, pp. 117992, Feb. 2022. doi:10.1016/j.apenergy.2022.117992

[41] H. Yan, F. Zhu, and Y. Lin, “Fast Training of Battery Deep Models for Embedded BMS,” Electronics, vol. 11, no. 9, p. 1321, Apr. 2022. doi:10.3390/electronics11091321

[42] P. Zhao and Y. Wu, “Jetson-Compatible DL Models for Rapid SOC Tracking,” IEEE Transactions on Industrial Informatics, vol. 19, no. 2, pp. 1502–1512, Feb. 2023. doi:10.1109/TII.2022.3222932

[43] F. Wang, H. Zhu, and M. Liu, “Development and deployment of an onboard LSTM-based SOC estimator for automotive BMS,” IEEE Trans. Veh. Technol., vol. 70, no. 12, pp. 12381–12391, Dec. 2021. doi:10.1109/TVT.2021.3129349

[44] S. Jin, Q. Zhai, and B. Y. Zhang, “Edge deployment of neural SOC estimators on embedded GPU platforms for real-time EV battery diagnostics,” Electronics, vol. 12, no. 5, p. 973, Mar. 2023. doi:10.3390/electronics12050973

[45] X. Tang and Z. Han, “Implementation of deep hybrid BMS using CNN-LSTM on EV fleets,” Journal of Energy Storage, vol. 59, p. 106333, Jan. 2024. doi:10.1016/j.est.2023.106333

[46] K. Wang, Y. Cao, and J. Chen, “Data scarcity in EV battery health estimation and challenges for robust deep learning models,” Applied Energy, vol. 314, p. 118981, Aug. 2022. doi:10.1016/j.apenergy.2022.118981

[47] T. Shi, B. Zhang, and S. Qian, “Cross-chemistry generalization in SOH estimation and the role of architecture complexity,” IEEE Access, vol. 11, pp. 45371–45383, 2023. doi:10.1109/ACCESS.2023.3253818

[48] R. Luo and Y. Jiang, “Interpretable deep learning models for battery diagnostics: A comparative analysis,” Energy Reports, vol. 9, pp. 682–690, Jan. 2023. doi:10.1016/j.egyr.2022.11.207

[49] B. Lin, C. Wang, and X. Tao, “Toward privacy-preserving battery diagnostics with federated LSTM networks,” Energy AI, vol. 9, p. 100178, 2023. doi:10.1016/j.egyai.2023.100178

[50] D. Guo, S. Lyu, and L. Xie, “Hybrid modeling of lithium-ion battery SOH with physical-informed neural networks,” IEEE Trans. Smart Grid, vol. 13, no. 4, pp. 3128–3139, Jul. 2022. doi:10.1109/TSG.2022.3153497

[51] Y. Fang and K. Li, “Graph-based battery state modeling using cell-to-cell interaction learning,” Applied Energy, vol. 328, p. 120392, Nov. 2022. doi:10.1016/j.apenergy.2022.120392

[52] J. Li and M. Zhang, “A physics-guided graph neural network for real-time SOC forecasting in battery packs,” Journal of Power Sources, vol. 555, p. 232315, Feb. 2023. doi:10.1016/j.jpowsour.2022.232315

[53] Z. Wu, H. Yang, and Q. He, “Reparameterization of deep networks for lightweight battery SOC estimation,” IEEE Trans. Ind. Electron., vol. 69, no. 11, pp. 11779–11789, Nov. 2022. doi:10.1109/TIE.2021.3135623

[54] L. Nguyen and J. Park, “Joint learning of battery SOC and remaining useful life using multi-task transformers,” Energy, vol. 261, p. 125221, Mar. 2023. doi:10.1016/j.energy.2022.125221

[55] M. Zhao et al., “Bayesian Neural Networks for Uncertainty-Aware Battery SOH Prediction,” IEEE Trans. Neural Netw. Learn. Syst., early access, 2024. doi:10.1109/TNNLS.2024.3332587

[56] S. Ali and H. Sharma, “Few-shot learning for lithium-ion battery diagnostics across capacities and formats,” Journal of Energy Storage, vol. 67, p. 107414, Apr. 2024. doi:10.1016/j.est.2024.107414

[57] T. Yoon, K. Kim, and H. Eom, “Multi-scale CNN technique for dV/dt waveform learning in battery degradation patterns,” Energies, vol. 15, no. 1, p. 121, Jan. 2022. doi:10.3390/en15010121

[58] A. Rahimi, R. Fallon, and K. Van, “Incorporating aging factors in real-time deep battery prognostics for electric buses,” IEEE Syst. J., vol. 17, no. 2, pp. 2156–2167, Jun. 2023. doi:10.1109/JSYST.2022.3198444

[59] Y. Cheng and Z. Liu, “Multi-objective optimization of lithium-ion battery SOH using ensemble-based temporal models,” Applied Energy, vol. 319, p. 119275, Sep. 2022. doi:10.1016/j.apenergy.2022.119275

[60] M. Hassan and J. Shen, “Residual energy estimation with LSTM-plus-Transformer hybrid model under unknown conditions,” Energy Reports, vol. 10, pp. 1451–1460, Jan. 2024. doi:10.1016/j.egyr.2023.11.151

[61] M. J. Page et al., “The PRISMA 2020 statement: an updated guideline for reporting systematic reviews,” BMJ, vol. 372, n71, 2021. doi:10.1136/bmj.n71

[62] S. Park, J. Kim, and H. Lee, “Generalization of Deep Learning SOC Models under Dynamic Urban Driving Cycles: A Fleet-Data Study,” IEEE Transactions on Intelligent Transportation Systems, vol. 24, no. 8, pp. 8123-8134, Aug. 2023. doi:10.1109/TITS.2023.3265431

[63] C. Wei, F. Sun, and M. Pecht, “Lightweighting Transformers via Pruning and Quantization for Real-Time SOH Estimation on Automotive Microcontrollers,” IEEE Transactions on Vehicular Technology, vol. 72, no. 11, pp. 14567-14578, Nov. 2023. doi:10.1109/TVT.2023.3289112

10.22060/EEJ.2026.24523.5752

[64] D. Varon, A. Gomez, and J. E. Martinez, “Opening the Black Box: An Analysis of SOH Predictors using SHAP and LIME,” Journal of Energy Storage, vol. 55, p. 105678, Nov. 2022. doi:10.1016/j.est.2022.105678