



Equivalent-circuit-model-guided feature selection for interpretable battery state-of-health estimation

Ying-Zhu Ren^a, Xiang Liu^b, Hai-Yuan Chen^{b,*}, Xiao-Bin Niu^{b,*}

^a School of Materials Science and Engineering, Southwest Jiaotong University, Chengdu, 610097, China

^b School of Materials and Energy, University of Electronic Science and Technology of China, Chengdu, 611731, China

ARTICLE INFO

Publishing editor: Yan Wang

Keywords:

Electrochemical impedance spectroscopy
Equivalent circuit model
Feature selection
Leave-one-battery-out validation
Lithium-ion battery
State of health

ABSTRACT

Reliable state-of-health (SOH) estimation depends not only on the nonlinear mapper but also on whether its inputs are compact, informative, and physically interpretable. We develop an equivalent circuit model (ECM)-guided feature-selection strategy that uses fractional-ECM parameters offline as physical anchors for charging-response and incremental capacity analysis (ICA) features. The predictive evaluation comprises 4259 cycle samples from ten commercial nickel-cobalt-aluminum (NCA) cells tested at 25 °C, 35 °C, and 45 °C under a 0.5 C protocol. A 13-dimensional (13D) full input is compared with five equal-cardinality 8-dimensional (8D) schemes using gated recurrent unit (GRU) and temporal convolutional network (TCN) nonlinear mappers. Ten-fold leave-one-battery-out (LOBO) validation provides the primary evidence. Among the 8D schemes, the Pearson correlation coefficient (PCC-8D) gives the lowest GRU macro root-mean-square error (RMSE; 0.005272 ± 0.002064). Physically anchored ECM-8D gives the lowest TCN macro RMSE (0.005882 ± 0.003804). ECM-8D is therefore competitive rather than universally most accurate. The two retained charging-time ratios are deterministically related, and 7D removal sensitivities reveal battery- and architecture-dependent effects. ECM fitting is required only during offline feature design; inference uses charging, ICA, and temperature variables. These results support a bounded route to interpretable SOH input construction within the tested NCA protocol, not cross-chemistry or wider-regime generalization.

1 Introduction

Lithium-ion batteries support electric mobility, stationary energy storage, and portable electronics because of their high energy density and favorable power capability [1]. Their usable capacity nevertheless declines through interacting aging processes, which also increase polarization and reduce power capability.

* Corresponding author.

E-mail addresses: pm23yr@leeds.ac.uk (Y.-Z. Ren), 202421030325@std.uestc.edu.cn (X. Liu), hychen@uestc.edu.cn (H.-Y. Chen), xbniu@uestc.edu.cn (X.-B. Niu).

<https://doi.org/10.1016/j.jnlest.2026.100374>

Received 28 April 2026; Received in revised form 23 July 2026; Accepted 22 August 2026; Available online 15 September 2026.

1674-862X/© 2026 The Authors. Publishing Services provided by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

State of health (SOH) quantifies this capacity loss and informs safety assessment, lifetime management, and maintenance decisions [2]. Reliable SOH estimation is therefore important, but operating temperature and cell-specific aging trajectories complicate the relationship between measured signals and remaining capacity.

SOH estimation methods commonly include direct testing, model-driven estimation, and data-driven learning. Direct capacity measurement is reliable but requires controlled charge-discharge testing that is impractical for frequent online use. Electrochemical and equivalent circuit models (ECMs) provide physically interpretable descriptions, although their identification and measurement requirements can limit routine deployment [3]. Data-driven methods learn nonlinear signal-to-SOH mappings, but their reliability depends strongly on the input representation [4]. Charging-response descriptors and incremental capacity analysis (ICA) features provide accessible degradation indicators [5]. However, concatenating correlated duration, voltage, and ICA variables can preserve deterministic redundancy and obscure which electrochemical responses support a selected feature [6].

Electrochemical impedance spectroscopy (EIS) resolves frequency-dependent responses associated with ohmic, interfacial, charge-transfer, and diffusion processes [7,8]. EIS descriptors have also been combined with interpretable deep-learning workflows for battery-health estimation [9]. Fractional-order ECMs can represent non-ideal interfacial dispersion more flexibly than matched integer-order circuits [10]. Their fitted parameters can therefore act as physical references for operational-feature screening. This role differs from supplying EIS or ECM parameters directly to the deployed estimator. It also requires bounded validation because a few representative impedance spectra cannot establish universal circuit superiority. Signed associations, rank-based checks, and equal-cardinality prediction comparisons are further needed to distinguish physical direction, selection stability, and predictive performance.

Recent SOH studies also explore more elaborate signal processing, deep architectures, and deployment settings. Battery-oriented examples include a temporal convolutional network (TCN) operating on partial load profiles [11], an extreme-learning-machine-GRU ensemble with a robust loss [12], and a Transformer architecture for high-frequency partial-state-of-charge data [13]. One edge-cloud workflow combines multi-population evolution whale optimization algorithm and variational mode decomposition (MEWOA-VMD) signal decomposition with Transformer-based SOH estimation [14], while a real-world battery-pack study combines TCN and Transformer components with transfer learning for SOH estimation and life warning [15]. These studies illustrate architecture, signal-conditioning, and field-transfer directions, but they do not validate ECM-guided selection or length-one TCN behavior in the present setting. Other physics-guided approaches embed prior knowledge in the learner itself [16]. An unresolved question is whether offline electrochemical characterization can instead guide a compact operational input without requiring EIS during inference.

This study addresses that question for 4259 samples from ten commercial nickel-cobalt-aluminum (NCA) cells cycled at 25 °C, 35 °C, and 45 °C under a 0.5 C protocol. We first construct a 13-dimensional (13D) full input from 12 charging-response and ICA candidates plus temperature. Fractional-ECM parameters then provide offline physical anchors for an 8-dimensional (8D) input containing temperature and seven selected degradation features. A matched integer-versus-fractional comparison on one preselected representative spectrum per temperature assesses the anchor model within that limited scope. Pearson correlation coefficient (PCC-8D), Spearman correlation coefficient (SCC-8D), least absolute shrinkage and selection operator (LASSO-8D), and random forest importance (RF-8D) selections provide equal-cardinality comparators. Signed Pearson and Spearman diagnostics separately examine association direction and ranking sensitivity. Primary evaluation uses ten-fold leave-one-battery-out (LOBO) validation with the held-out cell excluded from fitting, scaling, target-based selection, and stopping-epoch choice. Gated recurrent unit (GRU) and TCN models are evaluated as nonlinear mappers for cycle-level vectors with sequence length one. They are not used as evidence of long-range temporal learning. The study thus tests whether ECM-8D remains physically anchored and competitive across architectures while explicitly examining ratio redundancy and applicability boundaries.

2 Methods

2.1 Dataset, SOH definition, and applicability domain

The predictive dataset was derived from a public aging study of commercial 18650 NCA cells [17]. Ten cells were cycled at 0.5 C and at controlled ambient temperatures of 25 °C, 35 °C, or 45 °C. Each charge

comprised a 1.5-A constant-current (CC) stage to 4.2 V followed by a constant-voltage (CV) stage until the current decreased to 0.05 C; discharge was conducted at constant current to 2.65 V. EIS was performed after full charge and a 30-min rest, using a 20-mV excitation over 50 mHz–10 kHz at 25-cycle intervals. The resulting prediction table contained 4259 cycle-level observations. Each observation retained its battery identity, cycle index, temperature, charging-derived features, and capacity-derived SOH.

SOH was defined as the capacity-retention ratio:

$$\text{SOH}_k = \frac{Q_k}{Q_r} \quad (1)$$

where Q_k denotes the measured discharge capacity at cycle k , and Q_r denotes the rated (initial) capacity of the fresh cell. SOH is therefore dimensionless and is reported on a 0–1 scale in the modeling workflow.

The empirical domain of the predictive evaluation is restricted to these commercial NCA cells, the three tested temperatures, and a 0.5 C cycling rate. Results shown in Fig. 1 do not establish transfer to another chemistry, format, rate, temperature range, or operating protocol.




Fig. 1. Experimental protocol and dataset scope: (a) CC–CV charge profile for cycle 1 of cell 25C11 and (b) SOH trajectories for ten NCA cells at 25 °C, 35 °C, and 45 °C.

2.2 Charging-curve and incremental-capacity features

Twelve candidate degradation features were extracted from the CC–CV charging record [18]. The time-domain candidates were the CC duration (t_{CC}) and CV duration (t_{CV}). Two voltage descriptors were the initial charging voltage (v_{ini}) and the average voltage over the initial charging segment ($v_{\text{ini_ave}}$). ICA [5] provided the magnitudes and voltage positions of three identifiable peaks, represented as p_1 , v_1 , p_2 , v_2 , p_3 , and v_3 . ICA was calculated according to (2), in which the charge increment is evaluated with respect to the corresponding voltage increment. Representative charging-curve and incremental-capacity definitions are shown in Fig. S1 and Fig. S2 in Appendix A, respectively.

$$\text{IC}(V) = \frac{dQ}{dV} \quad (2)$$

where Q denotes capacity and V denotes terminal voltage.

Two additional timing ratios were retained from the original feature protocol: τ_1 is t_{cc} divided by t_{cv} , and τ_2 is t_{cc} divided by total CC–CV time. They are deterministically related by $\tau_2 = \tau_1/(1 + \tau_1)$ and, where $\tau_2 < 1$, by $\tau_1 = \tau_2/(1 - \tau_2)$. This identity held to a maximum absolute residual of 7.27×10^{-10} in the assembled table; thus, the two columns encode one independent degree of freedom rather than two independent mechanisms. Both were nevertheless retained in the frozen main protocol for continuity with the original analysis, and their removal was examined separately.

Ambient temperature was appended to the 12 degradation candidates to form the 13D full input. Temperature was treated as a known operating condition and was fixed in every compact input set; it did not compete with degradation features during ranking.

2.3 Fractional ECM formulation and matched circuit comparison

The fractional-order equivalent-circuit model (ECM) was used as an offline physical reference for feature screening, not as an online input to the capacity predictor. The circuit comprises an inductive element, an ohmic resistance, a film/interfacial branch, a charge-transfer branch, and a diffusion element [19,20]. Four fitted descriptors, R_0 , R_1 , R_2 , and W_1 , were used to represent ohmic, interfacial-film, charge-transfer, and diffusion-related behavior, respectively. The resulting topology and its nominal impedance-region assignment are shown in Fig. 2, where the region labels are functional guides rather than unique mechanistic assignments.




Fig. 2. Fractional-order ECM schematic and impedance-region interpretation. The labels are functional guides rather than unique mechanistic assignments.

The impedance expression of the fractional circuit remains as

$$Z_{\text{FO-ECM}}(\omega) = L + R_{\text{ohm}} + (R_{\text{SEI}} \parallel Z_{\text{CPE,SEI}}) + (R_{\text{ct}} \parallel Z_{\text{CPE,dl}}) + Z_{\omega} \quad (3)$$

where $\parallel$ denotes a parallel connection; L is the parasitic inductance; R_{ohm} is the ohmic resistance; R_{SEI} and R_{ct} are the resistances associated with the the solid-electrolyte interface and charge transfer, respectively; $Z_{\text{CPE,SEI}}$ and $Z_{\text{CPE,dl}}$ are the constant-phase-element impedances of the constant-phase elements (CPEs) representing the non-ideal interfacial and double-layer capacitances; Z_{ω} is the Warburg diffusion impedance. Each CPE has the impedance $Z_{\text{CPE}} = 1/[Q(j\omega)^{\alpha}]$, where Q is the CPE constant and α ($0 < \alpha \leq 1$) is the dispersion exponent. The ideal-capacitor case is recovered at $\alpha = 1$, whereas $\alpha < 1$ describes non-ideal interfacial dispersion.

To test whether fractional elements improved representative-spectrum fitting under a matched protocol, the integer circuit L_0 – R_0 – $(R_1//C_1)$ – $(R_2//C_2)$ – W_1 was compared with the corresponding fractional circuit L_0 – R_0 – $(R_1//\text{CPE}_1)$ – $(R_2//\text{CPE}_2)$ – W_1 . In these circuit notations, $//$ denotes a parallel connection and is equivalent to the $\parallel$ operator in (3). The numbered elements correspond to the symbols of (3) as $L_0 \equiv L$,

$R_0 \equiv R_{ohm}$, $R_1 \equiv R_{SEI}$, $R_2 \equiv R_{ct}$, $CPE_1 \equiv Z_{CPE,SEI}$, $CPE_2 \equiv Z_{CPE,dl}$, and $W_1 \equiv Z_\omega$. Each ideal capacitor contributes one free parameter (C), whereas each CPE contributes two (Q , α); replacing the two ideal capacitors by CPEs therefore increases the parameter count from seven to nine. The two models therefore had seven and nine free parameters, respectively. Both were fitted to exactly the same retained frequency points, with a common rule retaining points for which the measured imaginary impedance was negative, the legacy initial values, non-negative bounds, the same solver tolerances, and a maximum of 100000 function evaluations. The retained point counts were 86, 81, and 79 for the representative 25 °C, 35 °C, and 45 °C spectra, respectively.

Spectrum choice was fixed before this comparison: The seventh impedance record already selected in each of the three NCA workbooks was used. This design tests fit adequacy on one preselected spectrum per temperature and is not an exhaustive circuit-selection study over all spectra or aging states.

For each fit, residuals from the real and imaginary impedance components were stacked without frequency weighting. If a spectrum contained N complex points, the residual count was $n = 2N$, and the component root-mean-square error (RMSE) was defined as the square root of $RSS/(2N)$, where RSS stands for residual sum of squares. Complex R^2 was computed from the same stacked residual sum of squares, with the real and imaginary observations centered by their respective means. Model parsimony was assessed using $AICc = n \ln(RSS/n) + 2k + 2k(k+1)/(n-k-1)$, where k is the number of fitted parameters. Reporting both component RMSE and $AICc$ separates improvement in fit from the penalty for the two additional fractional exponents.

2.4 ECM-guided screening, signed diagnostics, and equal-cardinality comparators

ECM-guided screening used 107 matched charge-EIS checkpoints from six batteries: four NCA batteries (25C04, 25C11, 35C01 and 45C13) and two nickel-cobalt-manganese (NCM) batteries (25C16 and 35C02). The individual checkpoint counts were 15, 8, 24, 25, 8, and 27, respectively. Five batteries were cycled at 0.5 C, whereas 25C04 was cycled at 0.25 C. This deliberately mixed pool was used only to derive an offline physical ranking. It was not used as evidence of cross-chemistry prediction, and the chemistry-specific 35C02 records were not conflated across datasets. Three screening batteries (25C11, 35C01, and 45C13) also belonged to the ten-cell prediction dataset; consequently, ECM-8D was treated as a fixed offline prior rather than an outer-fold-reestimated, prediction-target-only selector.

The screening procedure had two stages [21,22]. First, the PCC matrix among the 12 candidate degradation features was inspected as a descriptive redundancy diagnostic. This step did not remove variables automatically. Second, within each screening battery, the PCC between every candidate feature and each of the four ECM descriptors was calculated at matched checkpoints. Temperature was excluded from this ranking and appended only after feature selection. Representative within-cell Pearson matrices are shown in Fig. S3 in Appendix A.

For candidate feature i in screening battery b , (4) defines the within-battery physical-association score as the equally weighted mean of the absolute PCC values over the four ECM descriptors.

$$S_i^{(b)} = \frac{1}{|\varepsilon|} \sum_{e_k \in \varepsilon} |\rho(f_i, e_k)| \quad (4)$$

where i is the index of the candidate degraded feature, $i = 1, 2, \dots, 12$, corresponding to the 12 degraded features; b is the index of the screening battery; f_i is sequence of values of the i -th candidate degraded feature at the matching checkpoint of battery b ; ε is the set of 4 descriptors of the fractional-order equivalent circuit model (fractional-order ECM); e_k is the k -th ECM descriptor in ε , $k = 1, 2, 3, 4$; $\rho(f_i, e_k)$ is PCC between the candidate feature f_i and the ECM descriptor e_k ; $S_i^{(b)}$ is the equal-weighted average of the absolute PCC between feature i and the 4 ECM descriptors in battery b .

Within each battery, the 12 scores were ordered from the largest to the smallest and converted to Borda points using (5).

$$P_i^{(b)} = 13 - \text{rank}_i^{(b)} \quad (5)$$

where $\text{rank}_i^{(b)}$ denotes in battery b , after sorting features i by $S_i^{(b)}$ from the largest to the smallest, the rank obtained by feature i ($\text{rank}_1 = \text{highest score}$); $P_i^{(b)}$ denotes the Borda points of feature i in battery b . This

transformation prevents batteries with wider raw correlation ranges from dominating the consensus, but it intentionally discards differences in absolute effect size between ranks. Exact ties were resolved by the fixed candidate order in Table 1.

Table 1

Representative-spectrum ECM fit statistics. Each row reports one legacy-preselected spectrum.

Temperature (°C)	Integer RMSE (Ω)	Fractional RMSE (Ω)	Integer R^2	Fractional R^2	Integer AICc	Fractional AICc
25	2.1414×10^{-4}	2.5226×10^{-5}	0.99715	0.99996	-2891.729	-3623.041
35	1.3720×10^{-4}	1.8265×10^{-5}	0.99791	0.99996	-2866.952	-3515.828
45	1.5663×10^{-4}	3.2808×10^{-5}	0.99879	0.99995	-2753.928	-3243.436

The global consensus score was then calculated by equally summing the battery-level Borda points in (6).

$$P_i = \sum_{b=1}^B P_i^{(b)} \quad (6)$$

where B denotes the total number of batteries in the screening pool; P_i denotes the global consensus score of feature i , that is, the equal-weighted sum of the Borda points $P_i^{(b)}$ of the 6 screening batteries: rank₁ $\rightarrow$ 12 points, rank₁₂ $\rightarrow$ 1 point.

The sum in (6) spans all six screening batteries, so each battery contributes the same maximum number of points regardless of its checkpoint count. The seven highest-ranked degradation features were t_{cc} , τ_1 , τ_2 , p_1 , v_1 , v_{ini} , and p_2 . Appending temperature produced ECM-8D. This subset retains the deterministic ratio pair by design; the resulting redundancy is disclosed rather than interpreted as two independent physical contributions.

Signed association was evaluated separately from the absolute-value selection score. For the representative 0.5 C NCA screening batteries at 25 °C, 35 °C, and 45 °C (25C11, 35C01, and 45C13; 8, 24, and 25 checkpoints), signed PCC and SCC matrices were reported. PCC describes signed linear association, whereas SCC is a rank-based check of monotonic robustness; SCC is not a general test of nonlinear dependence. As a supplementary sign-coherence diagnostic, the absolute value of the mean signed coefficient over the four ECM descriptors was calculated separately for PCC and SCC for each feature and battery. A low value can reveal cancellation that is hidden by a high mean absolute coefficient. Neither diagnostic replaces the prespecified selection rule, and no significance, causal direction nor universal sign claim were inferred from these small, sparsely sampled batteries.

Equal-cardinality alternatives were constructed for the battery-disjoint analysis. Each compact set contained temperature plus seven degradation features. ECM-8D was fixed by the offline procedure above. PCC-8D and SCC-8D ranked the 12 candidates by the absolute Pearson or Spearman association with SOH, respectively, using only the current outer-training batteries. LASSO-8D fitted standardized LASSO models over 80 logarithmically spaced regularization strengths from 10^{-6} to 10^{-2} ; the strength was chosen by three-fold group cross-validation by battery within the outer-training set, and features were ranked by absolute coefficient. RF-8D ranked impurity-based importance from a 500-tree random forest fitted only to the outer-training batteries. Candidate order in Table 1 broke exact selector ties. The same compact-network configuration was used for all five 8D sets, so the comparison changed the feature-selection rule rather than input dimensionality or model capacity. Two 7D sets obtained by removing τ_1 or τ_2 from ECM-8D were analyzed separately as redundancy sensitivities and were not treated as equal-cardinality competitors.

2.5 Cycle-level GRU and TCN nonlinear mappers

Two established neural architectures were used as nonlinear regression mappers: A GRU network [23] and a TCN [24]. GRU-based estimators have been applied to charging-derived battery-health features [12,18], while TCNs have been used with partial battery load profiles [11]. We therefore retained the two architectures already implemented in the original study to test whether feature-selection conclusions were robust across two model families, not to claim a new architecture or an exhaustive architecture search.

For both models, each row was one cycle-level feature vector and was reshaped with seq_len = 1. No preceding cycles, rolling window or future cycle was supplied. Accordingly, the GRU did not learn long-range

degradation dynamics: A linear projection and layer normalization were followed by activation and dropout, a unidirectional multilayer GRU operated on a single step, and the sole output entered a regression head. The TCN applied the corresponding projection followed by residual temporal blocks with weight-normalized causal dilated convolutions, cropping, activation, and dropout, then global average pooling and the same regression-head pattern. At sequence length one, these convolutional blocks also act as nonlinear mappings of the current-cycle vector rather than as estimators of temporal history. The final sigmoid bounded the predicted SOH to the 0–1 interval.

Networks were optimized with AdamW against the stabilized loss $\sqrt{\text{MSE} + 10^{-6}}$, using the frozen warm-up/cosine learning-rate schedule and early stopping defined in the experiment protocol. Four fixed configurations were maintained: one each for GRU and TCN with the 13D full input and one compact configuration per architecture for the 8D and 7D inputs. The original search spaces and frozen settings are reported in Tables S2 and S3 in Appendix A, respectively. Search variables were learning rate (LR), batch size, weight decay, warm-up, gradient clipping, hidden width/depth, and dropout; the TCN search also included kernel size. The final sigmoid bounded the predicted SOH to the 0–1 interval. Both architectures are summarized in Fig. 3.




Fig. 3. GRU and TCN model schematics: (a) GRU and (b) TCN. The implemented predictors receive a single cycle-level vector (seq_len = 1).

2.6 Predictive evaluation protocols

2.6.1 Original 75/10/15 seen-cell evaluation

The original evaluation was retained as a historical, within-battery interpolation benchmark. Within each temperature group, cycle-level rows were shuffled and assigned approximately 75%, 10%, and 15% to training, validation, and test partitions, respectively. All ten batteries occurred in every partition. The feature scaler was fitted on the training partition only, after which validation and test rows were transformed with the training statistics. Models were trained on the training rows, the validation rows controlled early stopping, and performance was calculated on the test rows. Historical split counts are reported in Table S1 in

Appendix A.

This protocol tests interpolation among observations from already observed batteries; it is not an unseen-battery test. Moreover, the pre-revision hyperparameter search used Optuna [25] with shuffled, ordinary five-fold sample-level cross-validation over all 4259 rows before the final 75/10/15 split. Fifty trials were run for each architecture and input configuration, with at most 50 epochs per trial (Table S2 in Appendix A). Thus, samples later assigned to the nominal test partition influenced selection of the fixed architecture and optimization settings (Table S3 in Appendix A). Seen-cell performance, leave-one-feature-out (LOFO) results, and SHapley additive exPlanations (SHAP) explanations are therefore reported as descriptive fixed-protocol analyses rather than as an independent estimate of deployment generalization.

2.6.2 Strict battery-disjoint LOBO evaluation

Unseen-battery performance was assessed by a ten-fold LOBO protocol using seed 2026 and deterministic training settings. In each outer fold, one complete battery was held out and the remaining nine batteries formed the outer-training pool. The held-out battery supplied neither training rows, validation rows, scaler statistics, target-based feature ranks, nor stopping decisions. Across the ten folds, every one of the 4259 observations received exactly one out-of-fold prediction for each evaluated model and input set.

All data-dependent operations were repeated within the current outer-training pool. PCC-8D, SCC-8D, LASSO-8D, and RF-8D were reselected without the held-out battery; ECM-8D remained the prespecified offline set. During inner validation, the predictor scaler was fitted only on the inner-training partition and was applied unchanged to the corresponding inner-validation partition. For the final outer-fold refit, a new scaler was fitted on all nine outer-training batteries and applied unchanged to the held-out battery. LASSO's internal standardization and group cross-validation likewise used only outer-training batteries.

Stopping duration was selected without inspecting the outer test battery. For each architecture, the nine outer-training batteries were divided into three shuffled group folds by battery. Training could run for at most 200 epochs, with patience 20 and a minimum improvement threshold of 10^{-5} . The best epoch from each inner fold was recorded, and the rounded median of the three values, with a minimum of one epoch, became the final duration. A separate duration was obtained for the 13D full input and for the compact ECM-8D family. Within an outer fold, the compact duration was then shared across ECM-8D, PCC-8D, SCC-8D, LASSO-8D, RF-8D, and the two ratio-removal sets to avoid giving one compact selector a larger stopping budget.

After duration selection, each model was freshly initialized and refitted on all nine outer-training batteries for the prescribed number of epochs. The held-out battery was used only for final prediction and metric calculation for each prespecified model-feature combination. The architecture and optimizer hyperparameters were not re-searched in the outer folds; they were the four configurations inherited from the all-data pre-revision search described in Subsection 2.6.1. The LOBO analysis therefore provides battery-disjoint training, scaling, target-based selection, and epoch control, but it is a fixed-configuration generalization sensitivity, not a fully nested hyperparameter evaluation. Its inference remains confined to held-out NCA batteries within the observed 25–45 °C, 0.5 C domain.

2.7 Auxiliary LOFO and SHAP analyses and performance metrics

LOFO analysis was performed only under the original seen-cell 75/10/15 protocol. Starting from ECM-8D, each of its eight inputs—including temperature—was removed in turn. GRU and TCN models were retrained with the corresponding frozen compact configuration and the same training and stopping workflow as the full ECM-8D model. The archived result table stores model-specific signed differences, $\Delta\text{RMSE} = \text{RMSE}_{\text{ablated}} - \text{RMSE}_{\text{baseline}}$; positive values indicate deterioration after removal, whereas negative values indicate a lower RMSE in that realized seen-cell analysis. The original composite plot clipped non-increases to zero, but revised Fig. S6 in Appendix A reports the archived signed values without that clipping. Because batteries were shared across partitions and no repeated-run inferential analysis was performed, the signed LOFO differences were treated only as descriptive fitted-model sensitivities, not as causal feature effects or evidence for unseen-battery deployment.

SHAP [26] was also applied only to the trained seen-cell ECM-8D models and did not participate in feature screening or model fitting. Feature importance was summarized by mean absolute SHAP magnitude, with

signed distributions retained for directional inspection. These values describe the fitted model around its analysis background; they are not causal effects, physical parameter estimates or evidence of extrapolation. Temperature-stratified SHAP summaries are provided in Fig. S4 in Appendix A.

Predictive error was quantified in SOH units. RMSE was the square root of the mean squared prediction residual, MAE was the mean absolute residual, and R^2 was one minus the residual sum of squares divided by the total sum of squares about the observed mean. In LOBO, all three metrics were first calculated separately for each held-out battery. The primary summary was the unweighted macro mean across the ten batteries, with the sample standard deviation across batteries used as a descriptive variability measure. Metrics calculated after pooling all out-of-fold observations were reported as secondary summaries because they weight batteries in proportion to their sample counts. No inferential test was attached to the across-battery means.

3 Results and discussion

3.1 Representative-spectrum comparison supports the fractional ECM as an offline physical anchor

We first compared the fractional-order ECM used to define the physical references with a topology-matched integer-order ECM. The comparison was restricted to the legacy-preselected seventh impedance record in each of the 25 °C, 35 °C, and 45 °C workbooks. These spectra had already been fixed in the original analysis and were not selected in response to the fit results. After identical $\text{Im}(Z) < 0$ filtering, 86, 81, and 79 complex impedance points were retained at 25 °C, 35 °C, and 45 °C, respectively. Both circuits were fitted with the same unweighted residual vector formed by stacking the real and imaginary impedance components. The integer circuit contained seven fitted parameters, whereas the fractional circuit contained nine because the two ideal capacitors were replaced by CPEs.

For each representative spectrum, the fractional circuit followed the measured Nyquist response more closely (Fig. 4 and Table 1). Here, $Z = Z' + jZ''$, where Z' and Z'' denote the real and imaginary parts of the impedance, respectively. At 25 °C, the RMSE decreased from $2.1414 \times 10^{-4} \Omega$ for the integer circuit to $2.5226 \times 10^{-5} \Omega$ for the fractional circuit, an 88.2% reduction, while the complex coefficient of determination (R^2) increased from 0.99715 to 0.99996. The corresponding RMSE values were 1.3720×10^{-4} and $1.8265 \times 10^{-5} \Omega$ at 35 °C, an 86.7% reduction, and $1.5663 \times 10^{-4} \Omega$ and $3.2808 \times 10^{-5} \Omega$ at 45 °C, a 79.1% reduction. The R^2 values increased from 0.99791 to 0.99996 at 35 °C and from 0.99879 to 0.99995 at 45 °C. The improvement was retained after accounting for the two additional CPE exponents: The corrected Akaike information criterion (AICc) was lower for the fractional circuit by 731.31, 648.88, and 489.51 at 25 °C, 35 °C, and 45 °C, respectively.




Fig. 4. Representative integer- and fractional-order ECM fits: (a) 25 °C, (b) 35 °C, and (c) 45 °C.

The consistent preference after the AICc penalty is compatible with the fractional circuit accommodating non-ideal dispersion in these three spectra, rather than the lower residual being explained solely by parameter count. This result provides a bounded adequacy check for using the fitted fractional-ECM parameters as offline physical anchors during feature screening. It does not establish that the fractional circuit is universally preferable, identifies a unique electrochemical mechanism, or validates its fit across

every spectrum in the aging dataset.

3.2 Signed association analysis reveals stable overlaps, directional heterogeneity and deterministic redundancy

The frozen offline screening pool comprised 107 impedance-matched checkpoints from six physical cells: Four NCA cells and two NCM cells, including one cell cycled at 0.25 C and five at 0.5 C. This mixed-chemistry pool was used only to screen charging and ICA descriptors against the four fractional-ECM references, R_0 , R_1 , R_2 and W_1 . The original PCC procedure ranked features within each cell by the mean absolute association across these four references and then combined the cell-level ranks using a Borda-style score. Its seven selected degradation features were t_{CC} , τ_1 , τ_2 , first ICA-peak magnitude (p_1), first ICA-peak voltage (v_1), v_{ini} , and second ICA-peak magnitude (p_2). Adding temperature produced the frozen ECM-8D input.

Replacing Pearson association with the SCC provided a monotonic robustness analysis. The SCC ranking selected p_1 , t_{CC} , v_1 , p_3 , v_{ini} , τ_1 , and p_2 . Thus, six of the seven selected descriptors were shared, with τ_2 unique to the PCC list and p_3 unique to the SCC list. The top-seven Jaccard similarity was 0.75, but the Spearman correlation between the complete PCC and SCC rankings was only 0.615. The overlap therefore shows partial robustness of the compact subset, whereas the moderate full-ranking association argues against describing the selection as method invariant.

Signed matrices further separated association strength from direction (Fig. 5). Across the six-cell screening pool, CC time, both ratio features, and the ICA-peak magnitudes were generally negatively associated with R_0 , R_1 , and R_2 , whereas v_{ini} and v_p were generally positively associated. In particular, the first p_1 was negatively associated with the charge-transfer-related R_2 by both PCC and SCC in all six screening cells; for 25C11, the coefficients were -0.842 and -0.452 , respectively. This direction is physically compatible with a lower and broader ICA peak accompanying increased polarization and charge-transfer impedance during ageing [5,19,20]. It is not a unique mechanistic assignment because ICA-peak evolution can combine loss-of-lithium-inventory, loss-of-active-material, and kinetic effects, while fitted R_2 remains model dependent. The sign is therefore interpreted only as a descriptive physical-consistency check. Associations with W_1 frequently changed direction among cells, preventing a universal directional interpretation. A separate sign-coherence score, defined as the absolute signed mean across R_0 , R_1 , R_2 , and W_1 , was compared with the prespecified mean-absolute score. The absolute-score and coherence-score rankings retained six of the same top seven features for both PCC and SCC, with Jaccard similarity of 0.75; their full-rank correlations




Fig. 5. Signed feature-ECM associations: (a) Pearson, 25 °C (25C11), (b) Spearman, 25 °C (25C11), (c) Pearson, 35 °C (35C01), (d) Spearman, 35 °C (35C01), (e) Pearson, 45 °C (45C13), and (f) Spearman, 45 °C (45C13).

were 0.748 and 0.832, respectively. t_{cv} illustrates why the signed diagnostic is needed: Its global coherence-to-magnitude ratio was only 0.109 for PCC and 0.0277 for SCC because positive and negative associations with different ECM references were nearly cancelled. Mean absolute association remains useful for screening strength, but it cannot by itself demonstrate a common physical direction.

The ratio audit exposed a separate structural limitation. Across all 4259 formal modeling rows, $\tau_2 = \tau_1/(1 + \tau_2)$, with a maximum numerical residual of 7.27×10^{-10} , Pearson correlation of 0.99687, and Spearman correlation of 1.00000. The two variables therefore contribute only one independent degree of freedom. Both were retained in the frozen ECM-8D definition to preserve continuity with the original analysis, but the main input cannot be described as redundancy free. Their performance consequences are examined separately through 7D removal sensitivities in Subsection 3.4. Finally, neither the signed matrices nor the mixed NCA/NCM screening pool constitutes prediction validation, a significance test, or evidence of cross-chemistry generalization.

3.3 Seen-cell interpolation results provide a secondary comparison with the original analysis

We retained the original sample-random results as a secondary, historical analysis. The 4259 samples were stratified by temperature and divided in a 75:10:15 ratio, but samples from all ten cells occurred in the training, validation, and test subsets. Moreover, the original hyperparameter search had access to the full sample collection. These estimates therefore quantify interpolation among observations from already seen cells and are not an unbiased test of performance on an unseen battery.

Within this seen-cell setting, compact inputs improved RMSE relative to the 13D input, but the ordering depended on architecture. For GRU, the 13D, ECM-8D and PCC-8D inputs yielded RMSE values of 0.001969, 0.001734, and 0.001622, respectively. For the 13D and ECM-8D inputs, respectively, the corresponding MAEs were 0.001429 and 0.001246, and the R^2 values were 0.9991 and 0.9993. For TCN, the 13D, ECM-8D, and PCC-8D RMSE values were 0.003162, 0.002274, and 0.002348, respectively. For the 13D and ECM-8D inputs, the MAEs were 0.002319 and 0.001730, and the R^2 values were 0.9976 and 0.9987, respectively. Thus, PCC-8D gave the lowest GRU RMSE, whereas ECM-8D gave the lowest TCN RMSE. These results motivated a battery-disjoint re-evaluation, but they do not support claims about transfer to unseen cells. Seen-cell prediction agreement is shown in Fig. S5 in Appendix A.

3.4 LOBO validation shows architecture-dependent performance on unseen NCA cells

The primary evaluation used LOBO validation across ten NCA cells measured at 25 °C, 35 °C, and 45 °C under the present 0.5 C protocol. Each cell served as the outer test set in one fold. In every fold, the held-out cell was excluded from input scaling, target-related feature selection, stopping-epoch determination, and weight fitting. The completed evaluation contained 160 outer-test results, covering ten cells, two model architectures, and eight input configurations; all recorded data-separation and result-integrity checks passed. The principal comparison comprised five equal-cardinality 8D schemes, each combining temperature with seven degradation descriptors selected by ECM guidance, PCC, SCC, LASSO, or RF importance. The 13D input was retained as a separate baseline rather than included in the equal-dimensional ranking.

Performance rankings differed between architectures. For the GRU, PCC-8D had the lowest macro RMSE among the five 8D schemes at 0.00527 ± 0.00206 , followed by SCC-8D at 0.00528 ± 0.00209 , ECM-8D at 0.00566 ± 0.00320 , RF-8D at 0.00625 ± 0.00305 , and LASSO-8D at 0.00665 ± 0.00249 . For the TCN, ECM-8D ranked first at 0.00588 ± 0.00380 , followed by RF-8D at 0.00600 ± 0.00309 , PCC-8D at 0.00604 ± 0.00320 , SCC-8D at 0.00630 ± 0.00316 and LASSO-8D at 0.00644 ± 0.00322 . The separate 13D baseline gave macro RMSE values of 0.00595 ± 0.00206 for the GRU, and 0.00991 ± 0.00594 for the TCN. These are unweighted macro means $\pm$ sample standard deviations across the ten held-out cells, so each cell contributes equally despite unequal numbers of checkpoints. The standard deviations describe battery-to-battery variation, not variation among random seeds or confidence intervals. Complete cell-wise LOBO metrics are provided in Table S4.

The controlled ranking resolves two distinct points. First, the ECM-guided subset remained competitive under battery-disjoint evaluation, ranking third for the GRU and first for the TCN. Second, no selector was uniformly best: The lowest GRU error was obtained with PCC-8D, whereas the lowest TCN error was obtained with ECM-8D. The close PCC-8D and SCC-8D results for the GRU are also consistent with their six-

feature overlap, but the different TCN ordering shows that overlap in selected variables does not guarantee identical predictive behavior.

Because the two ratio variables are deterministic transforms, we evaluated their removal only as a 7D sensitivity analysis, not as part of the equal-cardinality ranking. Removing τ_1 or τ_2 from ECM-8D reduced the GRU macro RMSE to 0.00445 ± 0.00171 and 0.00446 ± 0.00162 , respectively. Relative to ECM-8D, these removals lowered the battery-level RMSE in 7 of 10 and 8 of 10 folds. For the TCN, the corresponding macro RMSE values were 0.00563 ± 0.00341 and 0.00579 ± 0.00323 , and each removal lowered RMSE in 6 of 10 folds. The lower macro averages support removing deterministic duplication in a future input definition, but the heterogeneous fold-level changes do not show that either removal benefits every cell. No inferential comparison was performed.

3.5 Interpretability diagnostics, deployment implications, and limitations

Auxiliary feature-attribution and LOFO diagnostics clarify how the trained models used the compact input, without providing a physical validation of the features. In the archived seen-cell LOFO table, removing t_{cc} produced the largest positive RMSE change for both architectures, whereas several other removals had negative model-specific changes (Fig. S6 in Appendix A). The negative values show that the original clipped composite concealed direction, but they describe only single realized GRU and TCN analyses with batteries shared across partitions. They do not establish a reproducible benefit, a causal feature effect, or unseen-battery behavior. The separate LOBO ratio-removal sensitivity in Subsection 3.4 provides the battery-disjoint evidence. SHAP distributions provide a post-hoc description of the fitted ECM-8D predictors. Both diagnostics are descriptive properties of the fitted models and do not validate an electrochemical mechanism or show that large attribution corresponds to physical relevance.

The workflow separates offline physical guidance from online prediction. EIS measurements and fractional-ECM fitting are used only to construct and freeze the candidate feature definition. Once that definition is fixed, inference uses temperature and features derived from the charging and ICA curves; online EIS acquisition and online ECM fitting are not required. Reducing the nominal input from 13D to 8D removes 38.5% of the variables and may simplify data handling. We did not benchmark acquisition time, computation time, or energy consumption, so no operational speed or resource-saving claim is warranted. Furthermore, because the two retained ratios encode one independent quantity, the nominal dimension reduction overstates the number of independent inputs in the frozen ECM-8D scheme.

Several limitations bound the generalization claim. The LOBO prediction evidence contains only ten NCA cells, three temperatures from 25 °C to 45 °C, and one 0.5 C laboratory protocol. The two NCM cells and the 0.25 C NCA cell occurred only in the six-battery, 107-checkpoint offline screening pool and do not provide cross-chemistry or variable-rate prediction evidence. The study has not tested a wider temperature range, field duty cycles, an external dataset, or real-world deployment. The representative ECM comparison covers only three spectra and is not an exhaustive circuit-selection analysis.

The modeling design imposes additional constraints. Both GRU and TCN implementations used a sequence length of one, so the present comparison evaluates two nonlinear mapping architectures rather than long-history temporal learning. Testing only these two architectures does not establish superiority over multilayer perceptrons, tree-based regressors, or other model classes. Hyperparameter configurations were inherited and frozen from the original search rather than reoptimized within every outer training partition; the LOBO results are battery disjoint, but they are not a fully outer-nested hyperparameter-optimization estimate. Finally, maximal information coefficient (MIC) selection was not evaluated. It remains a candidate nonlinear dependence analysis for a larger, harmonized dataset with matched impedance and charging observations, alongside external validation and a prospective nonredundant input definition.

4 Conclusions

This study develops an ECM-guided route to compact and interpretable SOH input construction. Evidence covers 4259 samples from ten commercial NCA cells tested at 25 °C, 35 °C, and 45 °C under a 0.5 C protocol. For the preselected seventh impedance spectrum at each temperature, the fractional-order ECM achieved lower RMSE and AICc than the matched integer-order circuit. These representative-spectrum results support offline physical anchors, not universal circuit superiority. In the primary ten-fold LOBO evaluation, ECM-8D remained competitive across two length-one nonlinear mapping architectures. PCC-8D produced the lowest

GRU macro-averaged RMSE (0.005272 ± 0.002064), whereas ECM-8D produced the lowest TCN value (0.005882 ± 0.003804). This supports physical anchoring with architecture-dependent performance, not uniform accuracy dominance. The two charging-time ratios are deterministically related and do not provide independent degrees of freedom. Removing either ratio improved the macro average relative to ECM-8D in both architectures, although the effect varied among held-out cells and models.

EIS acquisition and ECM fitting are confined to offline feature design, while inference uses charging, ICA, and temperature variables. No timing benchmark was performed, so the study does not establish faster or lower-energy inference. Generality remains limited to this NCA laboratory protocol. Because GRU and TCN receive sequence-length-one vectors, the experiments do not demonstrate long-range temporal learning. Previously reported architecture configurations were frozen for LOBO, with only stopping epochs selected inside the outer-training batteries. The evaluation is therefore battery-disjoint but not a fully nested hyperparameter-optimization study. Future validation should use harmonized battery-disjoint datasets spanning additional chemistries, temperatures, rates, field duty cycles, and genuinely longitudinal inputs before broader deployment claims are made.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jnlest.2026.100374>.

CRedit Authorship Contribution Statement

Ying-Zhu Ren: Writing—original draft, Software, Formal analysis, Data curation. **Xiang Liu:** Software, Formal analysis. **Hai-Yuan Chen:** Writing—review & editing, Conceptualization, Supervision. **Xiao-Bing Niu:** Resources and Supervision.

Declaration of competing interest

The authors declare no conflicts of interest.

References

- [1] C.-Y. Liu, H. Li, K.-X. Li, et al., Deep learning for state of health estimation of lithium-ion batteries in electric vehicles: a systematic review, *Energies* 18 (6) (2025) 1463.
- [2] F. von Bülow, T. Meisen, A review on methods for state of health forecasting of lithium-ion batteries applicable in real-world operational conditions, *J. Energy Storage* 57 (2023) 105978.
- [3] L.-P. Chen, X.-Y. Bao, A.M. Lopes, et al., State of health estimation of lithium-ion batteries based on equivalent circuit model and data-driven method, *J. Energy Storage* 73 (2023) 109195.
- [4] F.-J. Wang, Z. Zhai, Z.-B. Zhao, et al., Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis, *Nat. Commun.* 15 (1) (2024) 4332.
- [5] Y.-J. Zhang, Y.-J. Liu, J. Wang, et al., State-of-health estimation for lithium-ion batteries by combining model-based incremental capacity analysis with support vector regression, *Energy* 239 (2022) 121986.
- [6] A. Lanubile, P. Bosoni, G. Pozzato, et al., Domain knowledge-guided machine learning framework for state of health estimation in lithium-ion batteries, *Commun. Eng.* 3 (1) (2024) 168.
- [7] B. Jiang, J.-G. Zhu, X.-Y. Wang, et al., A comparative study of different features extracted from electrochemical impedance spectroscopy in state of health estimation for lithium-ion batteries, *Appl. Energy* 322 (2022) 119502.
- [8] Q.-M. Zhang, C.-G. Huang, H. Li, et al., Electrochemical impedance spectroscopy based state-of-health estimation for lithium-ion battery considering temperature and state-of-charge effect, *IEEE T. Transp. Electr.* 8 (4) (2022) 4633–4645.
- [9] B.-H. Shao, J. Zhong, J. Tian, et al., State-of-health estimation of lithium-ion batteries based on electrochemical impedance spectroscopy features and fusion interpretable deep learning framework, *Energies* 18 (6) (2025) 1385.
- [10] L.-P. Chen, S.-Q. Xie, A.M. Lopes, et al., A new SOH estimation method for lithium-ion batteries based on model-data-fusion, *Energy* 286 (2024) 129597.
- [11] S. Bockrath, V. Lorentz, M. Pruckner, State of health estimation of lithium-ion batteries with a temporal convolutional neural network using partial load profiles, *Appl. Energy* 329 (2023) 120307.
- [12] J.-S. Xue, W.-T. Ma, X.-Y. Feng, et al., Stacking integrated learning model via ELM and GRU with mixture correntropy loss for robust state of health estimation of lithium-ion batteries, *Energy* 284 (2023) 129279.
- [13] X.-P. Li, M.-H. Zhao, S.-S. Zhong, et al., BMSFormer: an efficient deep learning model for online state-of-health estimation of lithium-ion batteries under high-frequency early SOC data with strong correlated single health indicator, *Energy* 313 (2024) 134030.
- [14] Y. Chen, X.-H. Huang, Y.-G. He, et al., Edge–cloud collaborative estimation lithium-ion battery SOH based on MEWOA-

- VMD and Transformer, *J. Energy Storage* 99 (2024) 113388.
- [15] Y. Chen, D.-Y. Li, X.-H. Huang, et al., Exploring life warning solution of lithium-ion batteries in real-world scenarios: TCN-transformer fusion model for battery pack SOH estimation, *Energy* 335 (2025) 138053.
- [16] C.-S. Lin, L.-X. Wu, X.-G. Tuo, et al., A lightweight two-stage physics-informed neural network for SOH estimation of lithium-ion batteries with different chemistries, *J. Energy Chem.* 105 (2025) 261–279.
- [17] J.-G. Zhu, Y.-X. Wang, Y. Huang, et al., Data-driven capacity estimation of commercial lithium-ion batteries from voltage relaxation, *Nat. Commun.* 13 (1) (2022) 2261.
- [18] Y. Xie, K. Luo, L.-H. Zheng, et al., A state of health estimation method for lithium-ion batteries based on initial charging segment and gated recurrent unit neural network, *J. Power Sources* 638 (2025) 236607.
- [19] G.-J. Qian, Y.-J. Zheng, X.-Y. Li, et al., State of health estimation for lithium-ion batteries using impedance-based simplified timescale information, *Appl. Energy* 382 (2025) 125272.
- [20] Y.-L. Zhu, B. Jiang, J.-G. Zhu, et al., Adaptive state of health estimation for lithium-ion batteries using impedance-based timescale information and ensemble learning, *Energy* 284 (2023) 129283.
- [21] L. Yao, J.-S. Wen, Y.-Q. Xiao, et al., State of health estimation approach for Li-ion batteries based on mechanism feature empowerment, *J. Energy Storage* 84 (2024) 110965.
- [22] X.-L. Xia, Y. Chen, J.-W. Shen, et al., State of health estimation for lithium-ion batteries based on impedance feature selection and improved support vector regression, *Energy* 326 (2025) 136135.
- [23] K. Cho, B. van Merriënboer, C. Gulcehre, et al., Learning phrase representations using RNN encoder–decoder for statistical machine translation, in: Proc. of the Conf. on Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 2014, pp. 1724–1734.
- [24] S.-J. Bai, J.Z. Kolter, V. Koltun, An empirical evaluation of generic convolutional and recurrent networks for sequence modeling [Online]. Available, <https://arxiv.org/abs/1803.01271>, March 2018.
- [25] T. Akiba, S. Sano, T. Yanase, et al., Optuna: a next-generation hyperparameter optimization framework, in: Proc. of the 25th ACM SIGKDD Intl. Conf. on Knowledge Discovery & Data Mining, Anchorage, USA, 2019, pp. 2623–2631.
- [26] S.M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, in: Proc. of the 31st Intl. Conf. on Neural Information Processing Systems, Long Beach, USA, 2017, pp. 4768–4777.



Ying-Zhu Ren was born in Chengdu, China in 2005. He is currently pursuing the B.S. degree with the School of Materials Science and Engineering, Southwest Jiaotong University, Chengdu, China. His current research interest mainly focuses on battery health management.



Xiang Liu was born in Sichuan, China in 2002. He received the B.S. degree from Guizhou University, Guiyang, China in 2022. He is currently pursuing the M.S. degree with the School of Materials and Energy, University of Electronic Science and Technology of China (UESTC), Chengdu, China. His current research interests include battery health management, state-of-charge estimation, and deep neural networks.



Hai-Yuan Chen received the Ph.D. degree from UESTC in 2018, where he is currently an Associate Professor. His research interests include AI for materials discovery and first-principles computation.



Xiao-Bin Niu received the Ph.D. degree from University of California, Los Angeles, USA in 2008. He is currently a Professor with UESTC. His research interests include data-driven materials design, materials informatics & database construction.