energies

MDPI

Article

Single-Cycle Identification of State-of-Charge Windows in Large-Format LFP Cells

Ibrahim Eski [1], Hikmet Yigit [1], Ali Toprak [2,3,*] and Ali Rifat Boynuegri [1,2]

[1] BATLAB Research Center, Clean Energy Technologies Institute, Yildiz Technical University, Istanbul 34220, Türkiye; ibrahim.eski@yildiz.edu.tr (I.E.); hikmet.yigit@yildiz.edu.tr (H.Y.); alirifat@yildiz.edu.tr (A.R.B.)
[2] Department of Electrical Engineering, Yildiz Technical University, Istanbul 34220, Türkiye
[3] Department of Electrical and Energy, Plato Vocational School, Istanbul Topkapi University, Istanbul 34087, Türkiye
* Correspondence: alitoprak@topkapi.edu.tr

Abstract

The state-of-charge (SoC) window used during cycling influences battery behaviour and is relevant to the assessment of cells considered for second-life use. This study investigates whether the SoC window in which a cell was cycled can be identified from the electro-thermal response of a single reference cycle that is identical for every cell, without cycling data recorded within that window. Twenty commercial 102 Ah prismatic lithium iron phosphate (LFP) cells were cycled in five fixed SoC windows with periodic reference cycles. Sixty-five features describing capacity, efficiencies, rest voltages, the voltage curve and surface temperatures were derived from each reference cycle. Under leave-one-cell-out cross-validation, a regularised logistic-regression classifier achieved a per-cycle balanced accuracy of 0.75, and aggregating each cell's reference cycles correctly classified 19 of 20 cells. With feature-group, model and hyper-parameter selection nested within each fold, the per-cycle balanced accuracy was 0.61 and 16 cells were correctly classified. Misclassifications were concentrated between closely related windows, with a quadratically weighted κ of 0.88 and a depth-of-discharge mean absolute error of 2.0 percentage points. Rest voltages after the full discharge provided the strongest discrimination. The results establish single-cycle window identification for the tested cells and conditions, providing a basis for second-life assessment.

Keywords: lithium-ion battery; LFP; battery diagnostics; machine learning; second-life batteries; state-of-charge window identification

Check for updates

Academic Editor: JongHoon Kim

Received: 1 October 2026
Revised: 24 October 2026
Accepted: 28 October 2026
Published: 3 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Electric-vehicle adoption is increasing the number of lithium-ion batteries that will eventually leave their first application. An approximately 80% state-of-health (SoH) threshold is often used to describe retirement from electric-vehicle duty [1], but cells may also be withdrawn from their first service while retaining higher capacity. Both groups merit assessment before reuse or recycling. Managing them matters because discarded cells retain useful energy and contain materials that require responsible handling [2,3].

Recycling is essential for recovering the raw materials needed to manufacture new cells, but on its own it falls short of resolving the waste-battery problem: mature recycling routes are not yet available for every cell chemistry, collection networks are sparse, processing costs are high, and facility capacity is limited [4]. Moreover, a cell retired from its first

use typically still holds a large fraction of its energy and power capability, so sending it straight to recycling, where only the valuable metals are recovered, discards most of that remaining value. Reusing retired cells as second-life batteries (SLBs) extends their useful life, uses material resources more effectively, lowers the effective cost of stationary storage, and reduces the cradle-to-grave environmental footprint [5]. Encouraging SLB use, and designing cells with it in mind from the outset, is therefore an important strategy for a more sustainable and economically viable battery ecosystem [2].

SLBs are deployed across a broad range of stationary and lower-rate applications, including grid-scale energy storage, power-backup and uninterruptible-supply systems, the integration and smoothing of intermittent renewable generation, support for electric-vehicle charging, and various smaller off-grid uses [5].

Because cells considered for SLB use differ in their condition and prior duty, individual assessment can help match them to an appropriate application [1,6].

How a cell ages depends on its depth of discharge (DoD), charge/discharge rate (C-rate), temperature and other operating conditions [1]. The DoD is set by the state-of-charge (SoC) window, the lower and upper SoC limits between which a cell is cycled: 20–80% corresponds to 60% DoD, whereas 0–100% corresponds to 100% DoD. Where operational records are incomplete, information about a cell's past SoC window could complement present SoH during cell assessment. Most data-driven diagnostics instead treat the operating conditions as known inputs and estimate capacity, SoH or remaining useful life from the measured response. Here, the SoC window in which a cell was previously cycled is the classification target, inferred from a reference cycle that is the same for every cell.

Most related studies map electrochemical impedance spectroscopy (EIS) to a present or future health target. Zhang et al. [7] used full EIS and Gaussian-process regression (GPR) to predict capacity and remaining useful life (RUL) using measurements from 12 cells. Li et al. [8] instead estimated capacity with an enhanced GPR model fed by circuit parameters derived from partial-frequency EIS; their public dataset contained four cells. For SLB SoH, Faraji-Niri et al. [9] trained GPR on selected EIS features, SoC and temperature from 30 cells and verified it on five additional cells. Camboim et al. [10] coupled EIS with principal component analysis (PCA) and a multilayer perceptron across 431 cells. Wu et al. [11] combined selected EIS, domain-adversarial transfer and GPR using EIS measurements from ten cells. They then performed an auxiliary voltage-and-temperature check on eight Oxford cells. Both datasets were drawn from existing sources. Sun et al. [12] identified degradation modes with EIS, simulation data and an attention-enhanced deep neural network, verified on 24 cells. Li et al. [13] predicted RUL from the imaginary component of the EIS spectrum with Conv1D-SAM using nine cells from published data, while Si et al. [14] used EIS and XGBoost to estimate capacity and the knee point across 16 cells. Guo et al. [15] addressed a different target by reconstructing full EIS from a partial charging-voltage curve with GPR, using seven cyclically aged and nine calendar-aged cells.

Cycling-data methods retain the same forward direction. Shen et al. [16] mapped partial-charge voltage, current and charge capacity to cell capacity with a deep convolutional neural network using ensemble and transfer learning (DCNN-ETL). Their existing datasets comprised eight source cells and 20 target cells. Lin et al. [17] mapped observed capacity paths to future degradation trajectories and RUL with a multiple-change-point linear model, an autoregressive model and a Cox proportional-hazards model, evaluating the approach on eight experimental and 19 published cells. Nagulapati et al. [18] used GPR and support-vector machines (SVMs) to estimate capacity and SoH from discharge voltage alone or from voltage, current and temperature; the NASA dataset contained four cells. Severson et al. [19] used early-cycle discharge-voltage features and an elastic-net model to predict cycle life across 124 cells.

The closest retrospective studies reconstruct missing data or assess retired cells with unknown prior use. Tang et al. [20] used incremental-capacity (IC) features together with a neural network and a migration model. Three capacity measurements calibrated the recovery of labels for six historical trajectories. Xu et al. [21] reconstructed labels for 96 retired cells from voltage and IC sequences, using three labels per cell. Their ensemble model combines a recovery convolutional neural network with an exponential model through Kalman filtering. Both methods therefore require archived trajectories and sparse labels rather than identifying a cell's past SoC window from a single reference cycle. Wang et al. [22] followed a different route: using 108 features from four constant-current, constant-voltage and relaxation curves measured in one full-charge test, they grouped 86 retired cells into future-consistency groups by *k*-means with derivative-based refinement, while prior history remained context rather than the target.

Table 1 summarises the selected studies; experimental verification denotes author-generated cell-test data, while studies using only third-party or public datasets receive a cross. The "2nd-life" column marks studies relevant to second-life assessment or use, including work on cells that retain capacity above the conventional 80% SoH retirement threshold. "EIS-free" denotes inference without a measured EIS spectrum as an input, even when spectra serve as training targets. "Past prediction" denotes estimation of past battery data or operating conditions. "History-free inference" denotes inference without the target cell's archived multi-cycle data or capacity calibration. Prior work has estimated health, reconstructed missing labels or spectra, and grouped cells by consistency; this study identifies one of five controlled cycling windows from the efficiency, voltage-curve and thermal response of a single reference cycle common to all cells, without EIS.

Table 1. Taxonomy of sixteen selected data-driven battery-diagnostics studies and this work.

Work	Input Modality	Model	Prediction Target	Past Prediction	2nd-Life	EIS-Free	History-Free Inference	Experimental Verification
[7]	EIS	GPR	Capacity/RUL	×	×	×	✓	✓
[8]	EIS	GPR	Capacity	×	×	×	✓	×
[9]	EIS	GPR	SoH	×	✓	×	✓	✓
[10]	EIS	MLP NN	SoH	×	✓	×	✓	✓
[11]	EIS	DaNN + GPR	SoH	×	×	×	✓	×
[12]	EIS	DNN	Degradation modes	×	×	×	✓	✓
[13]	EIS	Conv1D-SAM	RUL	×	×	×	✓	×
[14]	EIS	XGBoost	Capacity + knee point	×	×	×	✓	✓
[15]	Charging curve	GPR	Impedance spectrum	×	×	✓	✓	✓
[16]	Cycling data	DCNN	Capacity	×	×	✓	✓	×
[17]	Cycling data	Change-point + AR + Cox	Degradation path/RUL	×	×	✓	×	✓
[18]	Cycling data	GPR/SVM	Capacity/SoH	×	×	✓	✓	×
[19]	Cycling data	Elastic net	Cycle life	×	×	✓	×	✓
[20]	Charging IC curve	NN + migration	Historical capacity labels	✓	×	✓	×	✓
[21]	V & IC curves	RE-CNN + KF	Historical capacity labels	✓	✓	✓	×	✓
[22]	Single charge curve	*k*-means + DIRM	Consistency grouping	×	✓	✓	✓	✓
This work	Efficiency, V curve & thermal	LogReg	Past SoC window	✓	✓	✓	✓	✓

A checkmark (✓) or a cross (×) shows whether a study has the property named in each column.

Motivated by this gap, this paper identifies a cell's SoC cycling window using the efficiency, voltage-curve and temperature signature of one reference charge–discharge cycle that is identical for all cells. Figure 1 outlines the approach. The core contributions are as follows:

- A five-class operating-window identification task is formulated using the electro-thermal signature of one reference cycle applied identically to all cells after window cycling, distinct from prediction of battery health or retrospective recovery of missing capacity labels from archived signals.
- A multi-modal single-cycle signature is designed from 65 features of the reference cycle: step-level capacity, efficiency and voltage summaries, voltage-curve features including incremental-capacity and differential-voltage peaks, and surface-temperature features.
- A pool of classifiers is compared under cell-wise leave-one-cell-out (LOCO) cross-validation and assessed with a group-respecting permutation test. Regularised logistic regression, the best-performing model under LOCO, is retained.
- The selected classifier is examined through permutation-importance analysis to identify which measured signals discriminate among the controlled windows and to inform future SLB-oriented assessment studies; the rest voltages after the full discharge provide the strongest discrimination.

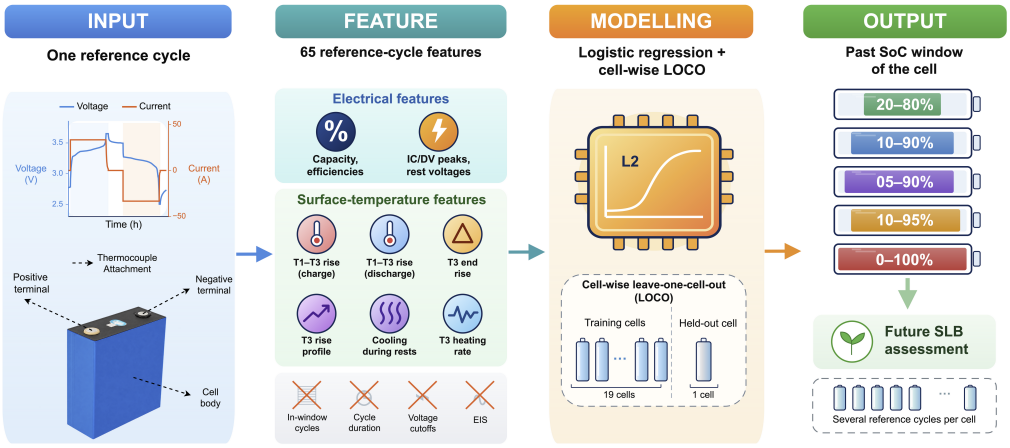

Figure 1. The proposed inverse-problem workflow for identifying the SoC window in which a cell was previously cycled from a single reference charge–discharge cycle.

The remainder of this paper is organised as follows. Section 2 describes the experimental setup, feature construction, models and cell-wise validation. Section 3 presents classification results, the measured feature relationships and their implications for SLB assessment. Section 4 concludes and outlines the next experimental and modelling steps.

2. Materials and Methods

2.1. Experimental Setup and Instrumentation

All ageing experiments were carried out at the BATLAB Research Center, whose laboratory infrastructure is described by Akkuş et al. [23]. The integrated bench in Figure 2 combines a regenerative charge–discharge system, temperature logging, an impedance analyser and a programmable climate chamber.

The cells were cycled by a Chroma 17010 system with 17212M-6-100S modules (Chroma ATE Inc., Taoyuan, Taiwan). A Hioki 8423 DataBox (HIOKI E.E. Corporation, Ueda, Japan) recorded three surface temperatures per cell, and a Binder MKF1020 chamber (BINDER GmbH, Tuttlingen, Germany) maintained a 25 °C and 40% relative humidity set point. EIS

characterisation used a Hioki BT4560-50 analyser and SW1002 multiplexer (HIOKI E.E. Corporation, Ueda, Japan).



Figure 2. The integrated battery test system at the BATLAB Research Center.

The charge–discharge system has twenty-four channels, each rated up to 100 A over 0–5 V. Its specified output and measurement accuracy is up to ±0.01% of full scale, with precision up to ±0.001% of full scale. A 24-bit analogue-to-digital converter and internal 1 ms sampling support integration of charge and energy; the per-step summaries used below are generated from these measurements.

For every cell the bench logs per-step summaries, namely voltage, current, capacity, energy and three surface temperatures, together with periodic impedance spectra and a raw stream sampled at 100 ms; in this work the per-step summaries and the 100 ms stream of each reference cycle are used downstream.

Figure 3 shows the four-terminal (Kelvin) wiring for DC cycling and AC impedance measurement. The two paths were kept separate, and dedicated sense leads reduced cable and contact-resistance contributions. The EIS path was zero-adjusted with a short-circuit fixture matched to the cell terminal spacing.

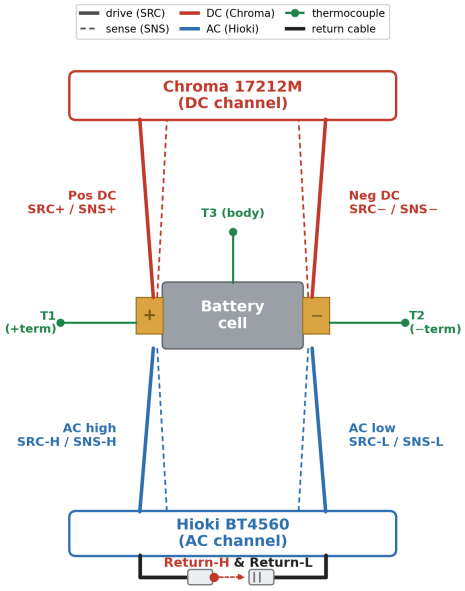

Figure 3. Representative four-terminal (Kelvin) wiring of a cell for combined DC cycling and AC impedance measurement.

Three thermocouples are bonded to each cell: one at the positive-terminal busbar, one at the negative-terminal busbar and one on the cell body, as shown in Figure 3. They provide the surface-temperature channels T_1, T_2 and T_3 used to construct the thermal features in Section 2.2.

2.2. Data and Data Preprocessing

The dataset comprises 20 commercial prismatic cells of 102 Ah nominal capacity with lithium iron phosphate (LFP) cathodes and graphite anodes. Each cell was cycled in one assigned SoC window for about two months, with periodic reference cycles. Before this campaign, each cell had been cycled under the same window assignment. The SoC window assigned to each cell is the prediction target. The first reference cycle of each cell precedes the window cycling of this campaign and is used only to characterise the cells; the classification uses the 69 subsequent reference cycles, whose distribution over the five classes is summarised in Table 2.

Table 2. Class distribution used in this study.

SoC Window	DoD (%)	Cell Count	Reference Cycles
0–100	100	6	21
10–95	85	3	11
05–90	85	3	11
10–90	80	3	12
20–80	60	5	14
Total		20	69

In the 0–100% SoC group, three cells were cycled at 0.98C and three at 0.49C; in the 20–80% group, three were cycled at 0.98C and two at 0.49C. The remaining groups were cycled at 0.98C. These rates correspond to currents of 100 and 50 A, set by the 100 A limit of each test channel described in Section 2.1. Each regular charge and discharge transferred the window's DoD multiplied by the capacity measured in the cell's most recent reference cycle, and a 15 min rest separated the charge and discharge steps. Before each reference cycle, each cell was discharged to the cutoff voltage and rested for 60 min. The reference cycle followed: a constant-current–constant-voltage (CC–CV) charge to 3.65 V, a 60 min rest and an impedance measurement, a CC–CV discharge to the cutoff voltage and a 30 min rest. These steps and the preceding discharge ran at 0.33C (33.3 A) with a 1.67 A current cutoff, and the discharge cutoff voltage was 2.0 V in the first three weeks and 2.5 V afterwards. A charge to the upper window limit then resumed regular cycling.

Because every cell undergoes the same reference cycle, its voltage limits, current profile and termination criteria are identical for all cells, and its discharge capacity measures the cell's state rather than the protocol. The discharge capacity is therefore used as a feature; the step durations, which follow from the capacity at a fixed current, are not.

The 65 features summarised by group in Table 3 are extracted from each reference cycle. Seven are computed from the per-step records: the discharge capacity, the Coulombic and energy efficiencies, the mean charge and discharge voltages, computed as energy divided by charge, and the voltages at the end of the rests after charge and discharge. The Coulombic and energy efficiencies are defined in Equations (1) and (2), respectively.

$$\eta_{\mathrm{C}} = \frac{Q_{\mathrm{dis}}}{Q_{\mathrm{chg}}} \tag{1}$$

$$\eta_{\mathrm{E}} = \frac{E_{\mathrm{dis}}}{E_{\mathrm{chg}}} \tag{2}$$

Here, Q_{chg} and Q_{dis} denote the charge and discharge capacities (Ah) of a cycle, and the Coulombic efficiency η_C is the fraction of inserted charge that is recovered. The quantities E_{chg} and E_{dis} denote the corresponding energies (Wh), and the energy efficiency η_E additionally reflects the charge–discharge voltage gap and associated losses. Their different sensitivities may help distinguish windows. In the reference cycle, both ratios describe the state of the cell after its window history.

Table 3. Reference-cycle features by group.

Group	Features	Count
Step summary	Discharge capacity; Coulombic and energy efficiency (Equations (1) and (2)); mean charge and discharge voltage; voltage at the end of the rest after charge and after discharge	7
Voltage curve	Voltage quantiles: voltage at 10, 25, 50, 75 and 90% of the step capacity (charge, discharge)	10
	IC peaks: voltage and height of two charge peaks and one discharge peak	6
	DV peaks: position and height of two peaks in each step	8
	Rest voltage and CV share: voltage 1 and 5 min into each rest; constant-voltage share of charge and discharge	6
Surface temperature	Temperature rise (T_1–T_3): mean and maximum rise of T_1, T_2 and T_3 (charge, discharge)	12
	T_3 end rise and heating rate: end-of-step rise and heating rate over the first 10 min of T_3 (charge, discharge)	4
	T_3 rise profile: T_3 rise at 25, 50 and 75% of the step capacity (charge, discharge)	6
	Cooling during rests: change of T_1, T_2 and T_3 during the rests after charge and discharge	6
Total		65

Thirty features describe the voltage recorded at 100 ms during the charge, the discharge and the rests that follow them: the voltage at 10, 25, 50, 75 and 90% of the step capacity in the constant-current part; the voltage and height of the main incremental-capacity (dQ/dV) peaks, two in the charge and one in the discharge; the position and height of two differential-voltage (DV, dV/dQ) peaks in each step; the voltage 1 and 5 min into each rest; and the share of the step capacity delivered in the constant-voltage (CV) phase. The incremental-capacity curves of the charge and the discharge are computed on a 1 mV voltage grid after smoothing, and their peaks are then located. The remaining 28 features describe the surface temperatures T_1, T_2 and T_3 from the terminal and body thermocouples of Section 2.1, each relative to its value at the start of the step: the mean and maximum rise during the charge and the discharge, the end-of-step rise and the heating rate over the first 10 min for T_3, the T_3 rise at 25, 50 and 75% of the step capacity, and the temperature change during the rests after charge and discharge. Because every thermal feature is referenced to the cell's own starting temperature, the chamber temperature is not required. These variables are measured under the common chamber set point, rest steps and thermocouple positions of the present campaign. C-rate is not supplied to the classifier; its effect will be examined in the controlled follow-up study described in Section 4.

Reference cycles were identified in the per-step records from their sequence of steps and rests. No feature values are missing, so no imputation is needed. Feature standardisation is fitted within each training fold and for each discharge cutoff.

As shown in Figure 4a–e, EIS was recorded at 100% SoC before the campaign; these spectra characterise cell condition and are not classifier inputs, nor is the impedance measured within the reference cycles. Each curve in Figure 4a–e is the spectrum of one cell. Figure 4f shows the state of health, defined as the reference-cycle discharge capacity divided by the 102 Ah nominal capacity, at the first and last reference cycles of each cell, marked by open and filled markers, respectively, and joined by a vertical line.

The high-frequency intercepts of these spectra, which approximate the ohmic resistance, lie between 0.39 and 0.52 mΩ in all window groups. The reference-cycle capacities in Figure 4f place the partial-window cells at 98–100% of the 102 Ah nominal capacity at the first reference cycle and at 94–97% at the last. The 0–100% cells started at 95–96%; those cycled at 0.49C ended at about 95% and those cycled at 0.98C at 79–88%, consistent with their higher C-rate. This health range is relevant to cells replaced before or shortly after reaching the conventional 80% end-of-life threshold and considered for a subsequent application.

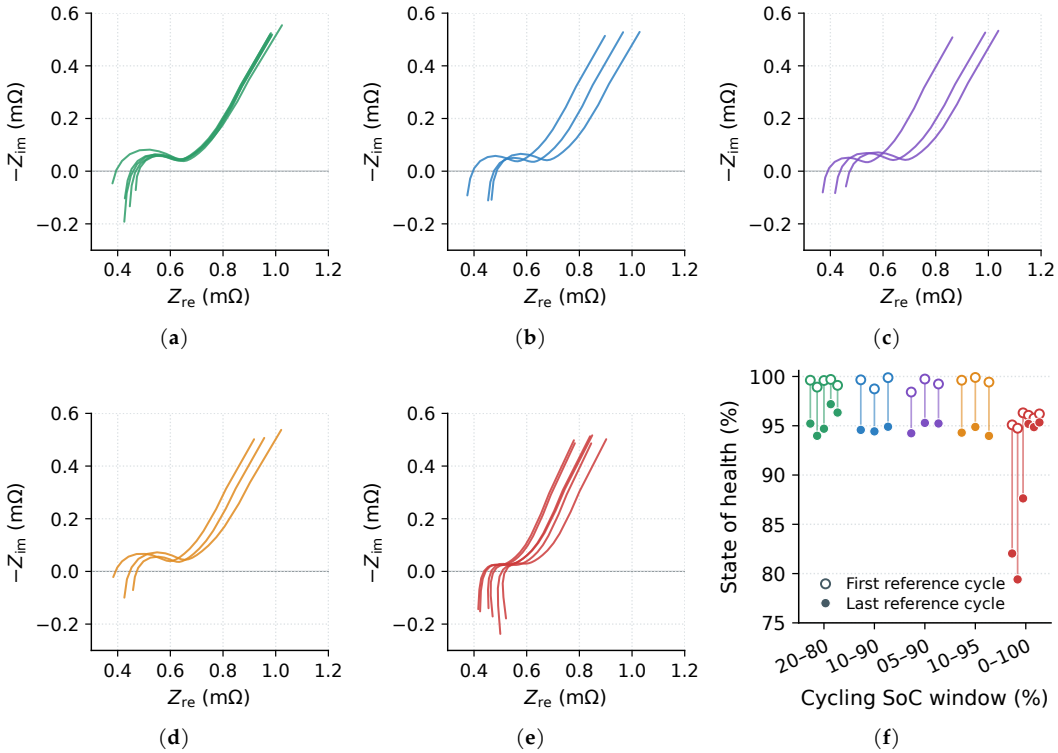

Figure 4. State of the twenty tested cells. Nyquist impedance spectra at 100% SoC for the cycling windows (**a**) 20–80%, (**b**) 10–90%, (**c**) 05–90%, (**d**) 10–95% and (**e**) 0–100%; (**f**) state of health at the first and last reference cycles.

2.3. Models and Validation

The SoC window in which the cell was previously cycled is treated as a five-class target. The model pool comprises a histogram-based gradient-boosted ensemble (HistGradientBoosting, HGB), ℓ_2-regularised logistic regression, random forest, extra trees, k-nearest neighbours, a linear support-vector machine and a majority-class dummy classifier. All non-dummy models except k-nearest neighbours use class-balanced weighting. Tree depth and minimum leaf size are constrained to control ensemble complexity, as listed in Table 4. The pipeline is implemented in Python 3.9.6 with scikit-learn 1.6.1 [24]. Figure 5 summarises the model selection and evaluation described below.

The primary model is regularised logistic regression. For a standardised feature vector $\mathbf{x} \in \mathbb{R}^{65}$ formed as in Section 2.2, its multinomial formulation assigns a probability to each SoC window $c \in \{1, \ldots, 5\}$ according to Equation (3).

$$P(y = c \mid \mathbf{x}) = \frac{\exp(\mathbf{w}_c^\top \mathbf{x} + b_c)}{\sum_{k=1}^{5} \exp(\mathbf{w}_k^\top \mathbf{x} + b_k)} \tag{3}$$

Here, $\mathbf{w}_c$ and b_c are the weight vector and intercept of class c. The parameters are estimated by minimising the class-balanced cross-entropy under an ℓ_2 penalty, as defined in Equation (4).

$$\min_{\{\mathbf{w}_c, b_c\}} -\sum_{i=1}^{N} \nu_{y_i} \log P(y_i \mid \mathbf{x}_i) + \frac{1}{2C} \sum_{c=1}^{5} \|\mathbf{w}_c\|_2^2 \tag{4}$$

Here, N is the number of training reference cycles, $\mathbf{x}_i$ and y_i are the features and window of cycle i, ν_{y_i} are the inverse-frequency class weights that compensate for the imbalance of Table 2, and C is the inverse regularisation strength. The penalty shrinks the coefficients towards zero to reduce overfitting, which is important with only twenty independent cells. The logistic-regression formulation gives a linear decision function for the 65 standardised features [25]. Because many features are strongly correlated, their individual coefficients are not used to rank the features; feature relevance is assessed with the permutation importance described in Section 2.4.

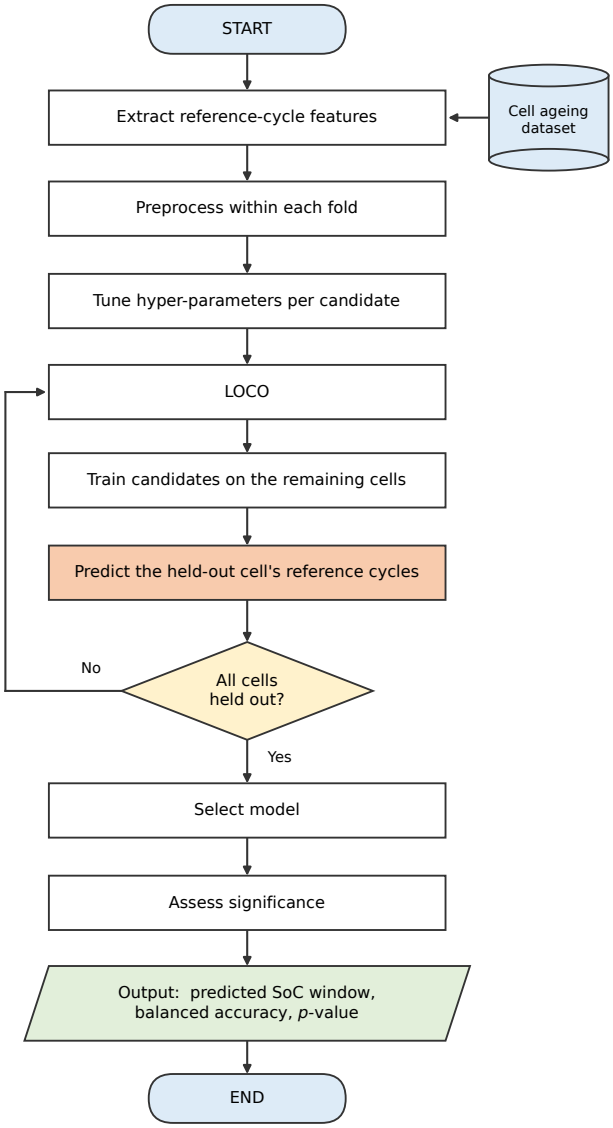

Figure 5. Cell-wise model-selection and evaluation flowchart of the reference-cycle analysis.

Table 4. Model-pool comparison (per-cycle balanced accuracy, cell-wise LOCO).

Model	Balanced Acc.	Best Hyper-Parameters
ℓ_2-Logistic Regression	0.75	$C = 10$
Linear SVM	0.66	$C = 1$
Extra-Trees	0.56	leaf 3
Random Forest	0.55	depth 4, leaf 3
HistGradientBoosting	0.47	learning rate 0.1, depth 2, 300 iterations, leaf 3
k-Nearest Neighbours	0.44	$k = 5$, distance
Dummy (majority)	0.20	n/a

Single-cycle inference is history-free: the classifier uses only the 65 features of one reference cycle, without preceding cycles, archived trajectories or capacity-calibration labels from the target cell. Soft voting aggregates separate reference-cycle predictions and supplies no historical features to the classifier.

Per-cycle balanced accuracy is the primary performance metric, reflecting inference from one reference cycle. All data splits are grouped by cell because cycles from the same cell are dependent; uncertainty is also evaluated at the cell level.

In the fixed-model comparison, each candidate's hyper-parameters are first tuned on all twenty cells by grid search with four-fold cell-wise GroupKFold and balanced accuracy as the scoring metric. Neighbour counts and minimum leaf sizes in the grids are scaled to the number of reference cycles. The tuned candidates are then compared under LOCO: each cell is held out in turn, and the model is trained on the remaining nineteen cells and applied to every cycle of the held-out cell, as illustrated in Figure 6. Table 4 reports the selected hyper-parameters.

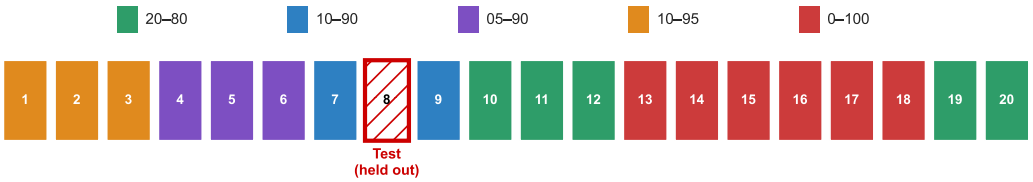

Figure 6. Leave-one-cell-out (LOCO) validation.

In a complementary fully nested analysis, the feature group, which can be the step summary, the voltage curve, the surface temperature or any combination of them, the model family and the hyper-parameters are selected within each LOCO fold using only its nineteen training cells, by four-fold GroupKFold over reduced versions of the same grids. This procedure is reported separately from the selected logistic-regression model in Section 3.1.

2.4. Performance Metrics

Because the classes are imbalanced, as shown in Table 2, performance is reported as balanced accuracy (BA). This metric is the mean of the per-class recalls and is defined in Equation (5).

$$\mathrm{BA} = \frac{1}{K} \sum_{c=1}^{K} \frac{\mathrm{TP}_c}{\mathrm{TP}_c + \mathrm{FN}_c} \tag{5}$$

Here, K denotes the number of classes and takes the value five in this study; TP_c and FN_c are the true-positive and false-negative counts for class c, respectively. The chance level for five balanced classes is $1/5 = 0.20$. Balanced accuracy is reported per cycle and after aggregation by cell. The per-cycle score measures how well a single cycle identifies the window. For the cell-level soft vote, predicted class probabilities are averaged over all cycles

of a cell, and the class with the highest mean probability is selected. Cell accuracy is the number of correctly classified cells out of twenty. Because the cells contribute unequal numbers of reference cycles, from two to five, a cell-weighted balanced accuracy is also reported. Recall is first computed over each cell's own cycles, then averaged equally across cells within each window and finally across the five windows. Thus, a cell with many cycles cannot dominate the score. Further ordinal-aware metrics summarise the structure of the errors, since the windows are naturally ordered by depth of discharge. The windows are ordered by increasing DoD, and the two windows that share an 85% DoD are ordered by their lower SoC limit, which places 05–90 before 10–95, so that the ordinal scale used throughout is 20–80, 10–90, 05–90, 10–95 and 0–100.

The quadratically weighted Cohen's κ [26] rewards agreement while penalising a misclassification in proportion to the squared distance between the true and predicted windows on this ordinal scale. The coefficient and its quadratic agreement weights are defined in Equations (6) and (7), respectively.

$$\kappa_w = \frac{P_o - P_e}{1 - P_e} \tag{6}$$

$$w_{ij} = 1 - \frac{(i-j)^2}{(K-1)^2} \tag{7}$$

Here, $P_o = \sum_{i,j} w_{ij}\, p_{ij}$ and $P_e = \sum_{i,j} w_{ij}\, p_{i+}\, p_{+j}$ are the weighted observed and chance-expected agreements. The quantity p_{ij} is the proportion of samples with true class i predicted as class j, while p_{i+} and p_{+j} are the row and column marginals. For the quadratic weighting, w_{ij} is the quadratic agreement weight and $K = 5$ is the number of windows, so that confusing adjacent windows costs little and confusing distant ones costs much. Two windows with equal DoD differ in their SoC limits rather than their depth. A second, window-distance-weighted κ is therefore computed using the form in Equation (6), with weights based on the distance between the windows' SoC limits. This distance is $d_{ij} = (|\Delta_{\text{low}}| + |\Delta_{\text{high}}|)/2$, expressed in percentage points. The weights are $w_{ij} = 1 - (d_{ij}/d_{\max})^2$, where $d_{\max} = 20$ is the largest distance in the experimental design. On this scale, confusion between windows with nearly coincident limits carries a small penalty regardless of their nominal DoD. The depth-of-discharge mean absolute error (DoD-MAE) maps each window to its DoD value and averages the absolute difference between the true and predicted DoD, as defined in Equation (8).

$$\text{DoD-MAE} = \frac{1}{N}\sum_{i=1}^{N}\left|\text{DoD}(y_i) - \text{DoD}(\hat{y}_i)\right| \tag{8}$$

Here, N is the number of evaluated reference cycles and $\hat{y}_i$ is the predicted window of cycle i; the error is expressed directly in percentage points of depth of discharge. Because the DoD-MAE assigns no error to a confusion between the two equal-depth windows, it is complemented by a boundary mean absolute error (boundary MAE) that averages d_{ij} between the true and predicted windows over all cycles in the same way, so that a mistake between 05–90 and 10–95 is counted as a five-point error. A cluster bootstrap resamples the twenty cells 2000 times and recomputes the selected model's per-cycle, cell-aggregated and cell-weighted balanced accuracies, and the fully nested per-cycle score, from held-out predictions. The 2.5th and 97.5th percentiles define the 95% confidence intervals (CIs). For permutation testing, cell labels are shuffled while keeping each cell's cycles together, 200 times for the fixed logistic-regression pipeline and 100 times for the fully nested procedure, in which every permutation repeats the complete nested selection. LOCO evaluation is repeated for each permuted labelling; nested feature-group, tuning and model-family selection remain within the nineteen training cells of each outer fold. Each p-value is the proportion of permuted scores at least as high as the observed score,

with the +1 correction. Because the 10–90 window lies inside both the 05–90 and the 10–95 windows, the evaluation is repeated without the three 10–90 cells as a sensitivity analysis reported in Section 3.1. Individual features are shuffled 15 times to measure the decrease in balanced accuracy. This permutation importance is computed for the selected logistic regression fitted to all twenty cells; it describes this fitted model rather than cross-validated performance. The features of each subgroup in Table 3 are also shuffled jointly. As a robustness check, the importance is also determined for extra trees, the best tree ensemble in Table 4, fitted to all twenty cells in the same way. Because this ensemble averages many trees, its in-sample predictions hardly change when a feature or subgroup is shuffled; its impurity-based importance, the mean decrease in Gini impurity attributed to each feature, is therefore used and summed over the features of each subgroup. For comparison, the subgroup importances of each model are expressed as shares of their sum.

3. Results and Discussion

3.1. Model Selection and Performance

Table 4 compares the models under cell-wise LOCO and lists the hyper-parameters selected by the grid search.

As Figure 7 shows, logistic regression attains a per-cycle balanced accuracy of 0.75, followed by the linear SVM at 0.66, while the tree ensembles and *k*-nearest neighbours score 0.44–0.56, above the 0.20 chance reference. The two linear models therefore lead the comparison. The grid search selects an inverse regularisation strength $C = 10$ for logistic regression, and this model is carried forward. In Figure 7, the selected logistic regression is shown in orange, the other candidates in blue and the majority-class dummy in grey.

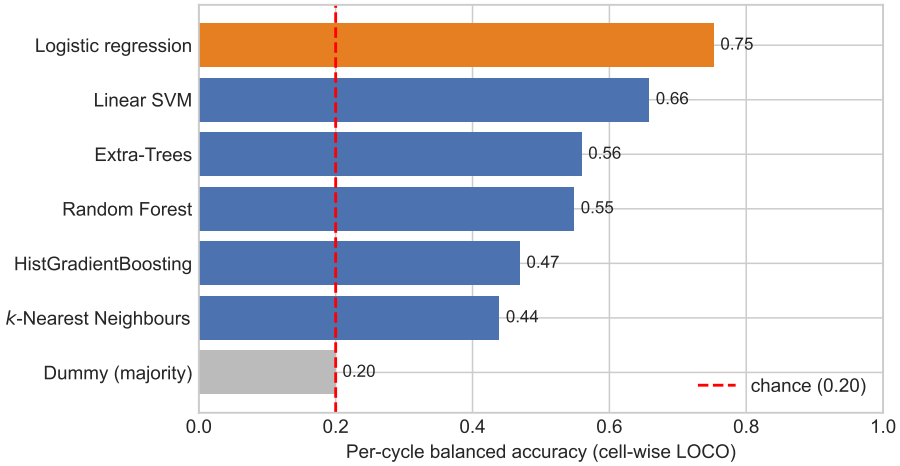

Figure 7. Model comparison under cell-wise LOCO.

Table 5 summarises the evaluation. The selected logistic regression achieves a per-cycle balanced accuracy of 0.75 and a cell-aggregated balanced accuracy of 0.93, correctly assigning 19 of 20 cells. The complementary fully nested procedure, which also selects the feature group within each fold, achieves a per-cycle balanced accuracy of 0.61, a cell-aggregated balanced accuracy of 0.73 and 16 correct assignments; the difference reflects that the fixed comparison selects the model and its hyper-parameters on all twenty cells, whereas the nested procedure makes these selections, and the feature-group selection, within each fold. The fixed model and the fully nested procedure return $p = 0.005$ and $p = 0.010$ in their group-respecting permutation tests; Table 5 provides the corresponding chance means, cell-bootstrap intervals and cell-weighted balanced accuracy.

Table 5. Cell-wise LOCO evaluation of the selected logistic regression and fully nested selection, with and without the 10–90 window.

Metric	All 20 Cells	17 Cells, 10–90 Excluded
Selected logistic regression		
Per-cycle balanced accuracy [95% CI]	0.75 [0.67, 0.83]	0.89 [0.69, 1.00]
Cell-aggregated balanced accuracy [95% CI]	0.93 [0.80, 1.00]	0.92 [0.75, 1.00]
Correctly assigned cells	19/20	16/17
Cell-weighted balanced accuracy [95% CI]	0.76 [0.68, 0.83]	0.87 [0.69, 1.00]
Permutation p-value (chance mean)	0.005 (0.161)	0.005 (0.201)
Quadratically weighted κ (per cycle)	0.88	0.95
Window-distance-weighted κ (per cycle)	0.92	0.93
DoD-MAE (per cycle, percentage points)	2.0	1.4
Boundary MAE (per cycle, percentage points)	1.1	0.9
Fully nested selection within each fold		
Per-cycle balanced accuracy [95% CI]	0.61 [0.49, 0.73]	0.77 [0.63, 0.93]
Cell-aggregated balanced accuracy	0.73	0.75
Correctly assigned cells	16/20	14/17
Permutation p-value (chance mean)	0.010 (0.153)	–

The cell-aggregated score exceeds the per-cycle score because, once the class probabilities are averaged over a cell's reference cycles, its correctly classified cycles outweigh the misclassified ones. The cell-aggregated confusion matrix in Figure 8 shows that all 20–80, 10–90, 10–95 and 0–100 cells are assigned correctly; the only misassigned cell is a 05–90 cell assigned to 10–90. The window-distance-weighted κ of 0.92 reflects the small penalty assigned to such confusion between nearly coincident SoC limits and is read alongside the classification results and other error measures in Table 5.

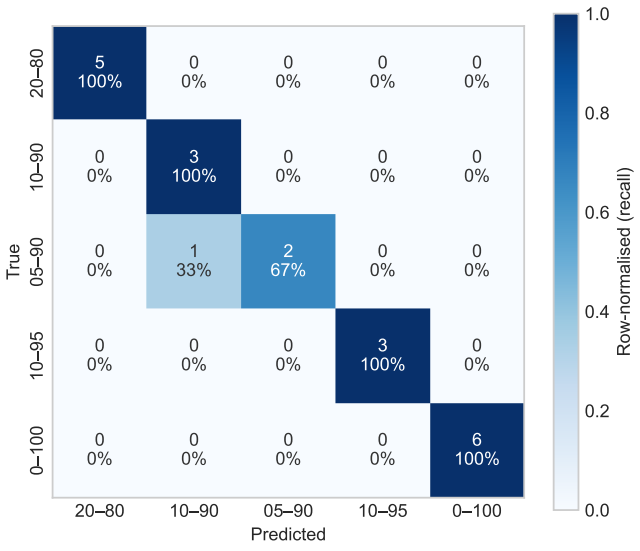

Figure 8. Cell-aggregated confusion matrix for the logistic regression under LOCO.

Figure 9 shows the recall of each cell over its own reference cycles, with one point per held-out cell. Ten of the twenty cells, all 20–80 cells and five of the six 0–100 cells, are correctly classified in every reference cycle. Per reference cycle, recall is 1.00 for 20–80, 0.90 for 0–100, 0.73 for 10–95, 0.58 for 10–90 and 0.55 for 05–90, so most misclassified reference cycles fall among the 10–90, 05–90 and 10–95 windows. Ten of the 15 misclassified reference cycles are confusions among these three windows.

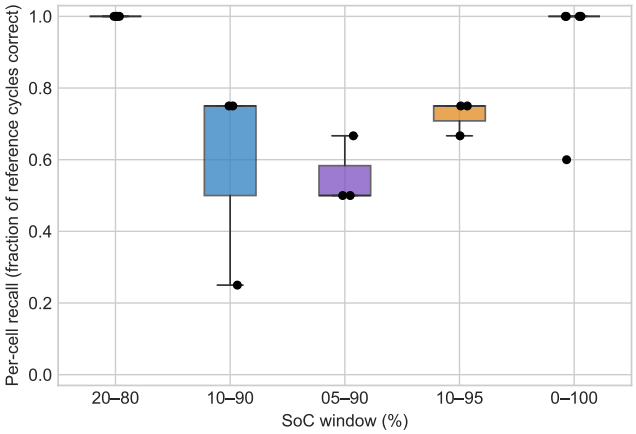

Figure 9. Per-cell recall of the logistic regression under LOCO by SoC window.

Across the four rounds of reference cycles, recorded after about two, three, five and six to eight weeks of the campaign, per-cycle balanced accuracy is 0.67, 0.67, 0.93 and 0.76. The rounds group the reference cycles by recording time and differ in size and composition: only 11 cells have a reference cycle in all four rounds, and the rounds contain 12, 20, 15 and 22 reference cycles from 12, 20, 15 and 19 cells. These rounds therefore represent time in the campaign rather than a matched SoH.

Figure 10a shows why these three windows are the hardest to separate. The 10–90 window, whose limits are marked by the dashed lines, lies inside both of the other two: it shares its lower limit with 10–95 and its upper limit with 05–90, so it differs from each of them by 5 SoC points at one end only, whereas 05–90 and 10–95 differ at both ends. Because the LFP voltage curve is flat over most of this range, the limits of the three windows also lie at closely spaced voltages: in regular cycling, the charge ended at median voltages of 3.51 V for both 10–90 and 05–90 and 3.56 V for 10–95, and the discharge at 3.05, 2.97 and 3.04 V, respectively. Accordingly, in the standardised feature space of the cell medians, the 10–90 class centre lies at distances of 4.4 and 5.6 from the 05–90 and 10–95 centres, which are 6.2 apart, and these distances are comparable to the spread within a class, 4.5–4.8. To quantify how much of the error this nested window causes, the evaluation is repeated without the three 10–90 cells. The per-cycle balanced accuracy then rises to 0.89: the recall of 05–90 rises to 0.73 and that of 10–95 to 0.91, whereas the recall of 20–80 and 0–100 is unchanged, as the hatched bars in Figure 10b show. As listed in Table 5, 16 of 17 cells are then correctly assigned. Most of the error of the five-window classification is therefore associated with the nested window rather than spread across all windows; the five-window result remains the primary result.

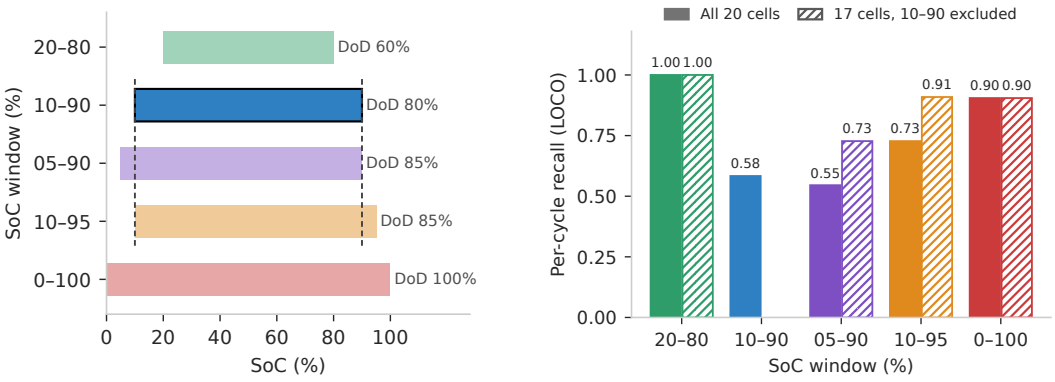

Figure 10. The nested 10–90 window: (**a**) limits and DoD of the five SoC windows; (**b**) per-cycle recall of the logistic regression under LOCO with and without the 10–90 window.

3.2. Feature Importance and Physical Interpretation

As shown in Figure 11, the rest voltages after the full discharge rank first for logistic regression: shuffling the subgroup of rest voltages and constant-voltage shares lowers the in-sample balanced accuracy by about 0.26, which is 41% of the summed subgroup importance. The step summary, which contains the voltage at the end of the same rest, follows at 0.11, the surface-temperature rises reach 0.07, and the incremental-capacity peaks, cooling during rests and differential-voltage peaks 0.05 each. The extra-trees ensemble confirms the leading role of the first two subgroups: the step summary and the subgroup of rest voltages and constant-voltage shares receive the largest shares of its impurity-based importance, 0.18 and 0.17, followed by the incremental-capacity and differential-voltage peaks at 0.14 and 0.12. Its importance is spread more evenly than that of logistic regression because impurity-based importance is divided among correlated features, but the two rankings of the nine subgroups agree with a Spearman correlation of 0.87. In Figure 11, the subgroups are those of Table 3, and the error bars show the standard deviation of the logistic-regression importance over the 15 shuffles.

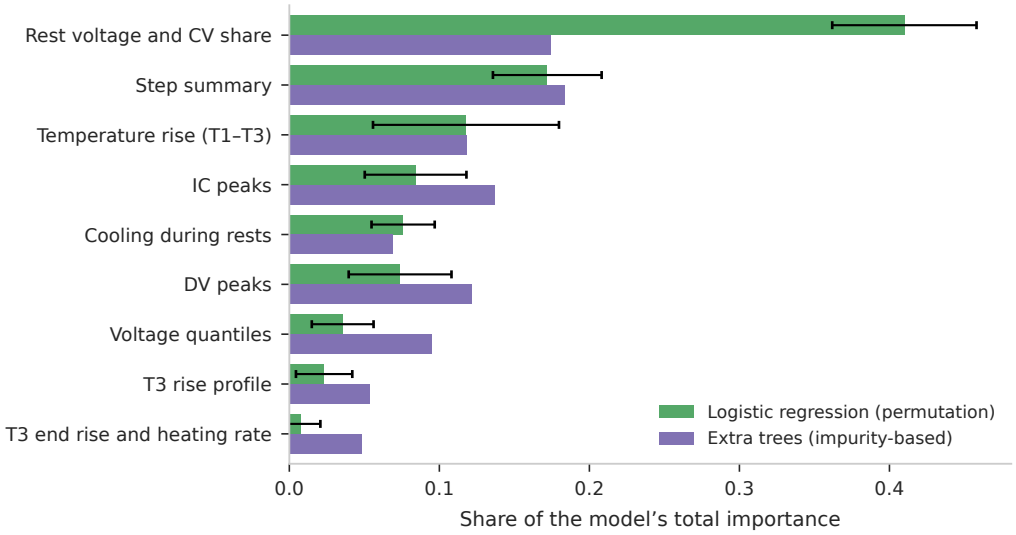

Figure 11. Feature importance by feature subgroup for logistic regression and extra trees, as a share of each model's total.

Individually, the voltage 1 min into the rest after discharge is the most important feature for logistic regression, with an importance of 0.26. It is followed by the voltage 5 min into this rest at 0.21 and the voltage at the end of the rest at 0.07, whereas the next feature, the mean rise of T_1 during the discharge, reaches only 0.02. For extra trees, the same three rest voltages are among the four features with the largest impurity-based importance, together with the height of the second incremental-capacity peak of the charge. The discharge capacity, which reflects the state of health, has little importance: 0.001 for logistic regression and 0.016 for extra trees, about the average share per feature. Figure 12 shows how these and the other features are correlated with one another and with the target. For readability, the 65 features are divided into four panels of closely related features, two electrical and two thermal. The last row of each panel represents the target by the DoD of the SoC window, which orders the windows but does not distinguish the two 85% windows. All correlations are computed after standardisation for each discharge cutoff.

Every pair of features with $|r| \geq 0.7$ lies within one panel, and the mean $|r|$ is 0.55 within the panels but only 0.17 between them. Over all 65 features, the mean $|r|$ is 0.89 among the voltage quantiles, 0.27–0.46 within each of the other subgroups of Table 3 and only 0.18 between the electrical and the thermal features. Among the most important features, the

three rest voltages after discharge in Figure 12a are strongly correlated with one another, with r between 0.74 and 0.94, and the end-of-rest voltage also correlates with the Coulombic efficiency at $r = 0.75$. The same features are the most strongly correlated with the target: r is between −0.78 and −0.87 for the three rest voltages after discharge and −0.78 for the Coulombic efficiency, whereas no thermal feature exceeds $|r| = 0.28$.

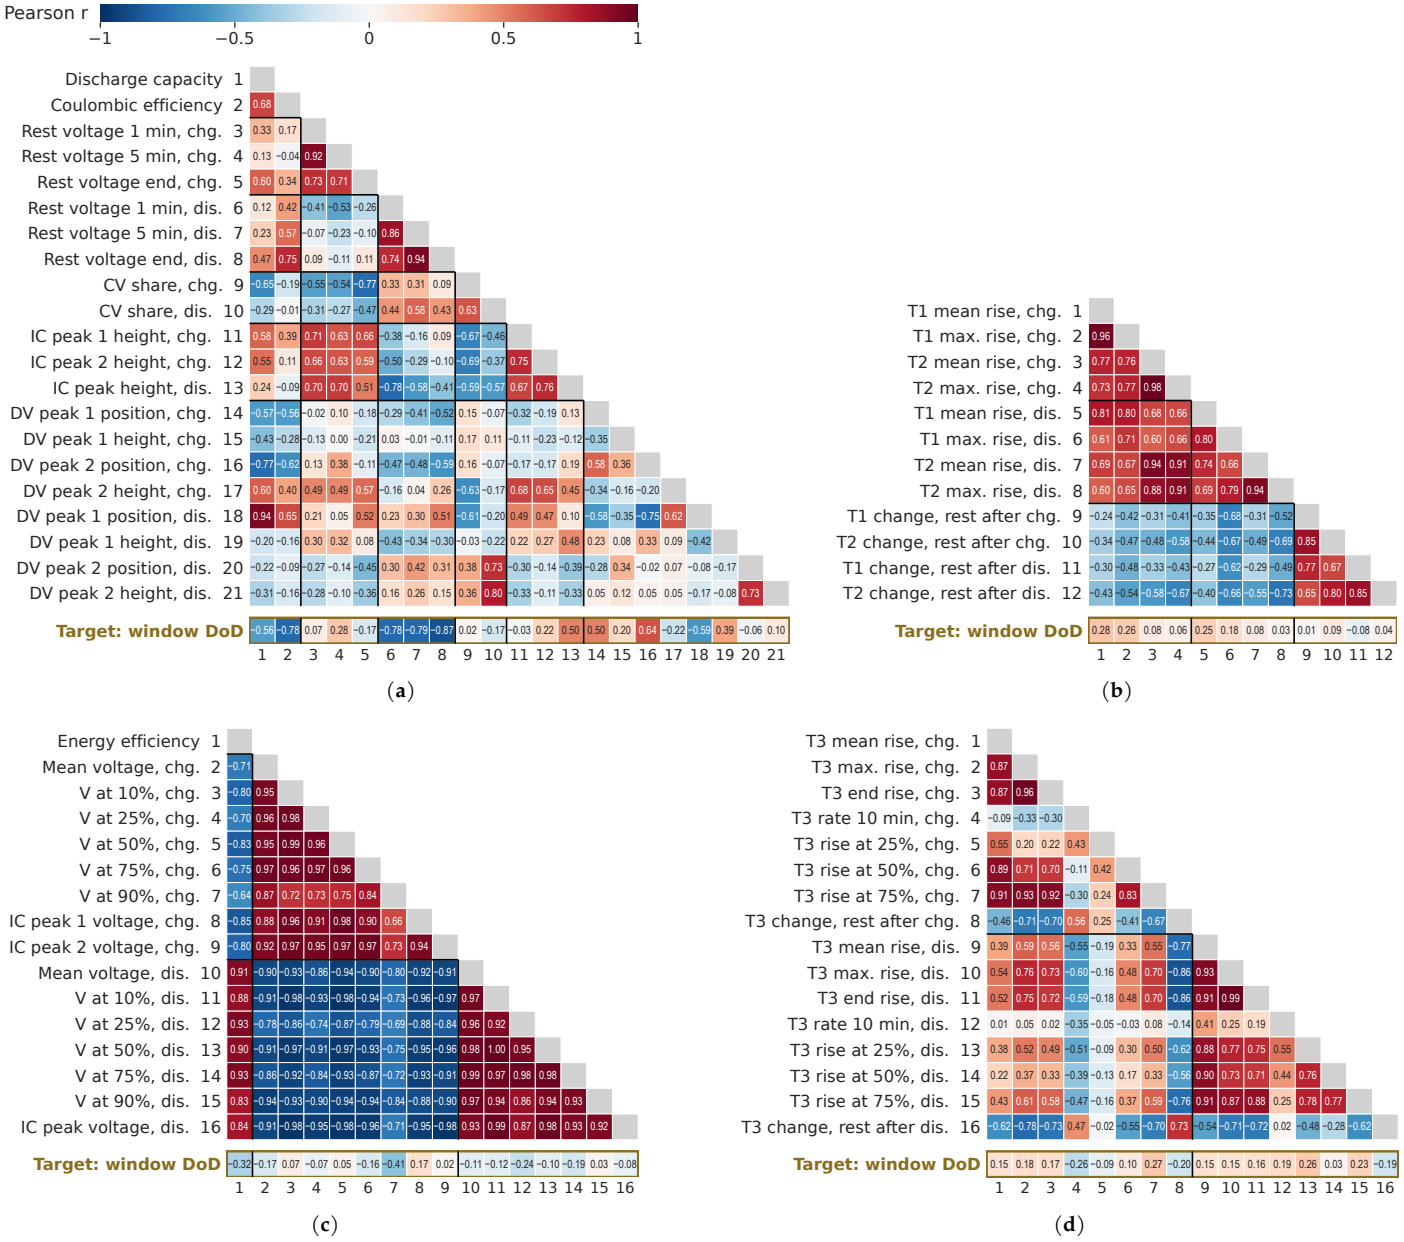

Figure 12. Pearson correlation among the 65 reference-cycle features and with the target: (**a**) capacity, Coulombic efficiency, rest voltages, CV shares, IC peak heights and DV peaks; (**b**) terminal temperatures T_1 and T_2; (**c**) energy efficiency, mean voltages, voltage quantiles and IC peak voltages; (**d**) body temperature T_3.

Figure 13 shows the distributions of these features for the reference cycles with the 2.5 V cutoff. The rest voltages after the full discharge follow the windows. At the end of the 30 min rest, the median voltage is 2.733 V for 20–80, 2.726 V for 10–90, 2.723 V for 05–90, 2.721 V for 10–95 and 2.702 V for 0–100; the voltages after 1 and 5 min keep the same order with smaller differences. The order follows the depth of discharge: the 60% DoD cells relax to the highest voltage and the full-range cells to the lowest, with the two 85% windows close together between them. This order also holds among the partial windows, whose

SoH overlaps at 94–97% at the last reference cycle, as Figure 4f shows. The reference-cycle Coulombic efficiency follows a similar order, from 1.003 for 20–80 to 0.998 for 0–100. The mean discharge rise of T_1, with medians of 0.72–1.13 °C, does not follow the windows monotonically and contributes in combination with the electrical features.

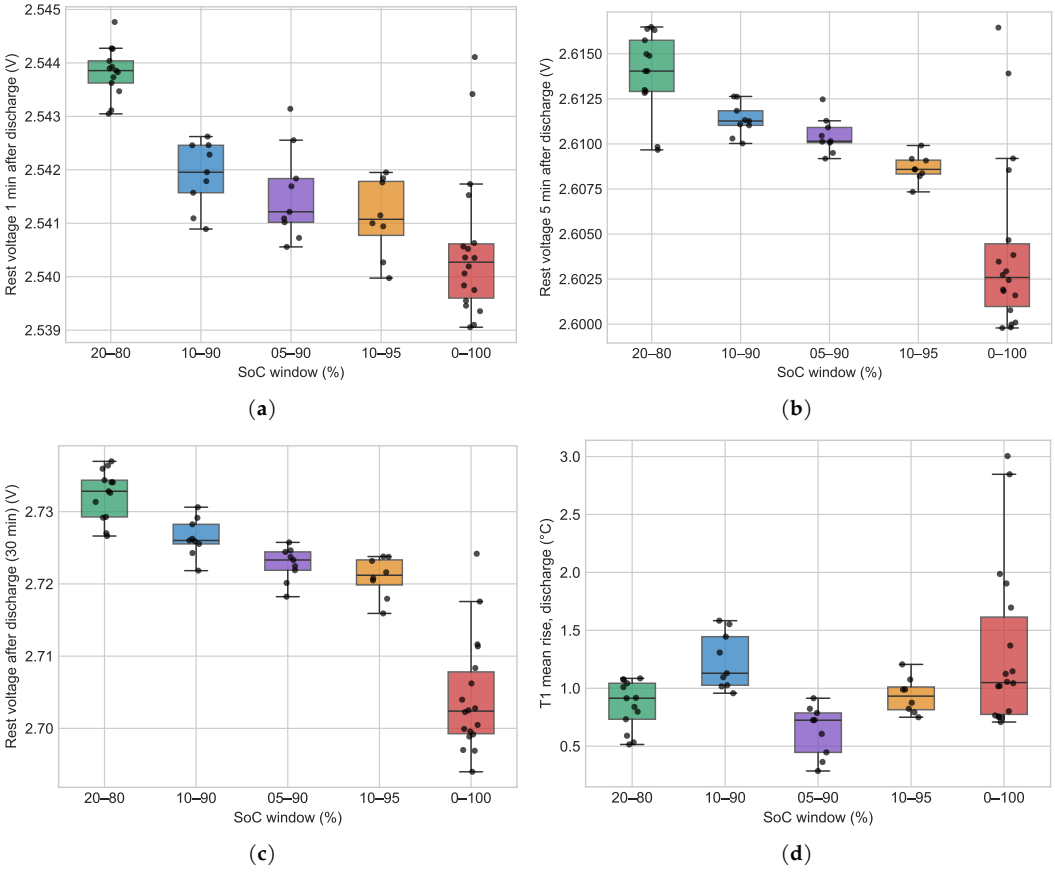

Figure 13. Distribution of the four most important features of logistic regression by SoC window: (**a**) voltage 1 min into the rest after discharge; (**b**) voltage 5 min into this rest; (**c**) voltage at the end of this rest; (**d**) mean rise of T_1 during the discharge.

One possible interpretation of these associations involves the distribution of lithium in the anode. Lithium exchange between the active anode and its overhang, the part of the anode without an opposing cathode, depends on the preceding SoC history and changes the extractable capacity, which rises when lithium returns from the overhang [27]. If window cycling leaves different amounts of lithium in the overhang, different amounts may return to the active anode at the end of the full discharge of the reference cycle. This would change the discharge capacity relative to the charge capacity, and thus the Coulombic efficiency, as well as the voltage to which the cell relaxes after the discharge, consistent with the common window order of these quantities and with their correlation of $r = 0.75$. Other history-dependent processes in the electrodes may also contribute.

Taken together, these electrical and thermal associations provide a single-reference-cycle descriptor of the previously applied SoC window that could complement present-state characterisation in SLB screening and sorting. The inferred quantity is the SoC window in which the cell was previously cycled, distinct from capacity, SoH or future-consistency estimates. As the prediction targets, measurement inputs and cell populations differ across the studies in Table 1, the comparison focuses on task definition and measurement requirements rather than ranking the reported accuracy values of EIS-based capacity or SoH regressors against the present classification score. Transfer to cells from first-life service and

downstream SLB grading will be examined through the future experimental and modelling work outlined in Section 4.

The application domain evaluated here comprises the five defined SoC windows, the tested 102 Ah prismatic LFP format and the controlled laboratory conditions, including the common 25 °C chamber set point and the cycling protocols described in Section 2.2. The inferred window is the one in which the cell was cycled before the reference cycle. Practical implementation requires the same reference cycle, a 0.33C CC–CV charge to 3.65 V and discharge to 2.5 V with 60 and 30 min rests, with voltage recorded at 100 ms, together with surface-temperature measurements at comparable locations. Application to other chemistries or cell formats, variable ambient conditions or more heterogeneous cell states requires representative calibration and validation for those conditions. The future experimental programme described in Section 4 will examine these factors and establish how the measurement procedure and calibration can be adapted for second-life assessment.

4. Conclusions

This study identified one of five controlled SoC windows from the electro-thermal response of a single reference cycle that was identical for every cell. At inference, only this common full 0–100% reference cycle is needed: the classifier uses no cycling data recorded within the cell's SoC window and no other records of the cell's history. Using 65 step-summary, voltage-curve and surface-temperature features, the selected classifier achieved a per-cycle balanced accuracy of 0.75 under cell-wise validation, and 0.61 when feature-group, model and hyper-parameter selection was nested within each fold. Aggregating predictions over each cell's reference cycles yielded a balanced accuracy of 0.93 and correctly assigned 19 of 20 cells. The strongest discrimination came from the rest voltages after the full discharge, which, like the reference-cycle Coulombic efficiency, follow the depth of discharge of the windows. These findings establish window identification for the tested 102 Ah LFP format and controlled cycling conditions.

The cells retained 94–97% SoH after the campaign, except the 0–100% cells cycled at 0.98C, which reached 79–88%, a range relevant to cells removed from a first application before or shortly after reaching the conventional 80% threshold and considered for another use. The present results provide a basis for developing assessment methods for such cells; extending the inference to cells from first-life service tested under the same reference cycle is the principal next step.

Future experimental work will extend the assessment to cells nearer conventional end-of-life SoH, cells withdrawn from first service, other cathode chemistries and cell formats, and more realistic non-repeating duty profiles. Controlled changes in C-rate, ambient temperature, rest duration and temperature-sensor placement or contact will determine how these factors affect the reference-cycle signature. In particular, a dedicated experimental design will examine the C-rate dependence of the voltage and thermal response of the reference cycle. The effect of cell placement within the chamber will be evaluated in a separate study. Within-cell tests that apply different SoC windows to the same cells will also examine window-dependent responses separately from cell-to-cell variation.

Model development will follow these experimental tests. Candidate directions include improved feature selection and calibration, models that represent uncertainty across cells, and transfer across chemistries and operating conditions, all assessed with cell-disjoint validation. The periodic EIS data recorded by the bench can also be evaluated as an additional modality against the present EIS-free baseline. The broader-population experiments will further support testing inferred windows as inputs to downstream SoH, remaining-life and ageing-knee models, and to SLB grading and pricing. Further applications can be explored

in battery digital twins with incomplete duty logs, warranty assessment, fleet management, battery-passport schemes and adaptive battery-management systems.

Author Contributions: Conceptualization, I.E. and A.R.B.; methodology, I.E. and H.Y.; software, I.E. and H.Y.; validation, I.E., H.Y., A.T. and A.R.B.; formal analysis, I.E. and A.R.B.; investigation, I.E.; data curation, I.E. and A.T.; writing—original draft preparation, I.E., H.Y. and A.T.; writing—review and editing, A.R.B.; visualization, I.E. and H.Y.; supervision, A.R.B.; project administration, A.R.B.; funding acquisition, A.R.B. All authors have read and agreed to the published version of the manuscript.

Funding: This study was supported by Istanbul Topkapi University within the scope of its publication incentive program and funded by the Presidency of Türkiye, Presidency of Strategy and Budget, under grant number 2024K12-223110.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The authors do not have permission to share the underlying data.

Acknowledgments: The authors appreciate the support of the Presidency of Strategy and Budget. During the preparation of this manuscript, the authors used Claude (Opus 5.5, Anthropic) for the purposes of improving language clarity, ensuring consistency of meaning and checking grammar. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Li, J.; He, S.; Yang, Q.; Wei, Z.; Li, Y.; He, H. A comprehensive review of second life batteries toward sustainable mechanisms: Potential, challenges, and future prospects. *IEEE Trans. Transp. Electrif.* **2023**, *9*, 4824–4845. [CrossRef]
2. Hossain, E.; Murtaugh, D.; Mody, J.; Faruque, H.M.R.; Sunny, M.S.H.; Mohammad, N. A comprehensive review on second-life batteries: Current state, manufacturing considerations, applications, impacts, barriers & potential solutions, business strategies, and policies. *IEEE Access* **2019**, *7*, 73215–73252. [CrossRef]
3. Nováková, K.; Pražanová, A.; Stroe, D.-I.; Knap, V. Second-life of lithium-ion batteries from electric vehicles: Concept, aging, testing, and applications. *Energies* **2023**, *16*, 2345. [CrossRef]
4. Lai, X.; Huang, Y.; Gu, H.; Deng, C.; Han, X.; Feng, X.; Zheng, Y. Turning waste into wealth: A systematic review on echelon utilization and material recycling of retired lithium-ion batteries. *Energy Storage Mater.* **2021**, *40*, 96–123. [CrossRef]
5. Dong, Q.; Liang, S.; Li, J.; Kim, H.C.; Shen, W.; Wallington, T.J. Cost, energy, and carbon footprint benefits of second-life electric vehicle battery use. *iScience* **2023**, *26*, 107195. [CrossRef] [PubMed]
6. Roslan, M.F.; Satpathy, P.R.; Prasankumar, T.; Ramachandaramurthy, V.K.; Mansor, M.; Walker, S.L. Second-life battery energy storage system for energy sustainability: Recent advancements, key takeaways and future perspectives. *J. Energy Storage* **2025**, *123*, 116808. [CrossRef]
7. Zhang, Y.; Tang, Q.; Zhang, Y.; Wang, J.; Stimming, U.; Lee, A.A. Identifying degradation patterns of lithium ion batteries from impedance spectroscopy using machine learning. *Nat. Commun.* **2020**, *11*, 1706. [CrossRef] [PubMed]
8. Li, Y.; Ye, M.; Wang, Q.; Lian, G.; Kemény, M.; Xia, B.; Zhang, B. Capacity estimation of lithium-ion battery through interpretation of electrochemical impedance spectroscopy combined with machine learning. *Measurement* **2025**, *243*, 116374. [CrossRef]
9. Faraji-Niri, M.; Rashid, M.; Sansom, J.; Sheikh, M.; Widanage, D.; Marco, J. Accelerated state of health estimation of second life lithium-ion batteries via electrochemical impedance spectroscopy tests and machine learning techniques. *J. Energy Storage* **2023**, *58*, 106295. [CrossRef]
10. Camboim, M.M.; Moreira, A.C.; Rosolem, M.d.F.N.C.; Beck, R.F.; Arioli, V.T.; Omae, C.; Ding, H. State of health estimation of second-life batteries through electrochemical impedance spectroscopy and dimensionality reduction. *J. Energy Storage* **2024**, *78*, 110063. [CrossRef]
11. Wu, J.; Meng, J.; Lin, M.; Wang, W.; Wu, J.; Stroe, D.-I. Lithium-ion battery state of health estimation using a hybrid model with electrochemical impedance spectroscopy. *Reliab. Eng. Syst. Saf.* **2024**, *252*, 110450. [CrossRef]
12. Sun, Y.; Xiong, R.; Wang, P.; Li, H.; Sun, F. A deep learning approach for enhanced degradation diagnostics of NMC lithium-ion batteries via impedance spectra. *J. Energy Chem.* **2025**, *107*, 894–907. [CrossRef]

13. Li, J.; Zhao, S.; Miah, M.S.; Niu, M. Remaining useful life prediction of lithium-ion batteries via an EIS based deep learning approach. *Energy Rep*. **2023**, *10*, 3629–3638. [CrossRef]
14. Si, Q.; Matsuda, S.; Ando, Y.; Momma, T.; Tateyama, Y. Capacity estimation and knee point prediction using electrochemical impedance spectroscopy for lithium metal battery degradation via machine learning. *Adv. Sci*. **2025**, *12*, 2502336. [CrossRef] [PubMed]
15. Guo, J.; Che, Y.; Pedersen, K.; Stroe, D.-I. Battery impedance spectrum prediction from partial charging voltage curve by machine learning. *J. Energy Chem*. **2023**, *79*, 211–221. [CrossRef]
16. Shen, S.; Sadoughi, M.; Li, M.; Wang, Z.; Hu, C. Deep convolutional neural networks with ensemble learning and transfer learning for capacity estimation of lithium-ion batteries. *Appl. Energy* **2020**, *260*, 114296. [CrossRef]
17. Lin, C.-P.; Cabrera, J.; Yang, F.; Ling, M.-H.; Tsui, K.-L.; Bae, S.-J. Battery state of health modeling and remaining useful life prediction through time series model. *Appl. Energy* **2020**, *275*, 115338. [CrossRef]
18. Nagulapati, V.M.; Lee, H.; Jung, D.; Paramanantham, S.S.; Brigljevic, B.; Choi, Y.; Lim, H. A novel combined multi-battery dataset based approach for enhanced prediction accuracy of data driven prognostic models in capacity estimation of lithium ion batteries. *Energy AI* **2021**, *5*, 100089. [CrossRef]
19. Severson, K.A.; Attia, P.M.; Jin, N.; Perkins, N.; Jiang, B.; Yang, Z.; Chen, M.H.; Aykol, M.; Herring, P.K.; Fraggedakis, D.; et al. Data-driven prediction of battery cycle life before capacity degradation. *Nat. Energy* **2019**, *4*, 383–391. [CrossRef]
20. Tang, X.; Liu, K.; Li, K.; Widanage, W.D.; Kendrick, E.; Gao, F. Recovering large-scale battery aging dataset with machine learning. *Patterns* **2021**, *2*, 100302. [CrossRef] [PubMed]
21. Xu, Y.; Zeng, T.; Peng, W.; Tian, J.; Yi, X.; Xu, Q.; Zhu, S.-P. Regenerate large-scale retired second-life battery datasets via recovered capacity labels-based deep learning. *Energy Storage Mater.* **2026**, *86*, 104947. [CrossRef]
22. Wang, S.; Gao, F.; Tian, H. Deep sorting of reused batteries for enabling long-term consistency grouping with unknown prior conditions. *Cell Rep. Phys. Sci*. **2025**, *6*, 102657. [CrossRef]
23. Akkuş, M.; Eski, İ.; Boynueğri, A.R. Establishing a pioneering laboratory for second-life battery applications: Enhancing energy storage and reducing environmental impact. *Clean Energy Technol. J*. **2025**, *3*, 62–72. [CrossRef]
24. Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res*. **2011**, *12*, 2825–2830.
25. Hosmer, D.W.; Lemeshow, S.; Sturdivant, R.X. *Applied Logistic Regression*, 3rd ed.; Wiley: Hoboken, NJ, USA, 2013; pp. 269–278. [CrossRef]
26. Yilmaz, A.E.; Demirhan, H. Weighted kappa measures for ordinal multi-class classification performance. *Appl. Soft Comput*. **2023**, *134*, 110020. [CrossRef]
27. Lewerenz, M.; Münnix, J.; Schmalstieg, J.; Käbitz, S.; Knips, M.; Sauer, D.U. Systematic aging of commercial $LiFePO_4$|Graphite cylindrical cells including a theory explaining rise of capacity during aging. *J. Power Sources* **2017**, *345*, 254–263. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.