Journal of Science and Technology – Smart Systems and Devices
Volume x, Issue x, Month Year, ppp–ppp

Sensitivity of Lithium-Ion Battery State-of-Health Estimation from Partial Charging Data to Measurement Density under Cell-Wise Validation

Thanh Le Son[1], Tuyen Nguyen Duc[1,*], Ali Moradi Amani[2], Syed M. Nawazish Ali[2]

[1]*School of Electrical and Electronic Engineering, Hanoi University of Science and Technology, Ha Noi, Vietnam*
[2]*School of Engineering, RMIT University, Melbourne, Australia*
[*]*Corresponding author email: tuyen.nguyenduc@hust.edu.vn*

Abstract

Partial-charge state-of-health (SOH) estimation relies on restricted segments of charging records, yet its performance may depend on measurement density within the observed window. This study evaluates the sensitivity of partial-charge SOH estimation to reduced measurement density using cell-wise validation on unseen battery cells. Using the Huazhong University of Science and Technology (HUST) lithium iron phosphate/graphite dataset (77 cells, 146,122 cycles), we analyzed a fixed voltage window ($W_0 = [2.8955, 3.5962]$ V) across native observations and elapsed-time retention thresholds of 10, 20, 40, and 80 s. Four charging descriptors were recomputed at each density level. Model family and hyperparameters were selected via nested grouped cross-validation on native-density data across 33 prespecified configurations and held constant, with separate estimators refitted per density using shared cell partitions. Thinning reduced the median number of retained observations from 250 at native density to 17 at the 80 s threshold. Cell-macro mean absolute error (MAE) increased nonmonotonically from 0.012725 at native density to 0.012772, 0.012895, 0.013089, and 0.013000 across the 10–80 s thresholds. Paired, native-referenced MAE differences evaluated via 10,000 physical-cell bootstrap replicates produced simultaneous 95% uncertainty bands that all spanned zero, and no equivalence margin was prespecified. Consequently, while measurement density should be treated as an explicit factor in partial-charge SOH modeling, these results establish neither practical equivalence to native density nor definitive performance degradation across the reduced-density conditions under deterministic downsampling.

Keywords: Cell-wise validation, lithium-ion battery, measurement density, partial charging, state-of-health estimation.

1. Introduction

Partial-charge measurements have been used as an observable basis for lithium-ion battery state-of-health (SOH) estimation, with health-related information extracted from charging voltage and current trajectories rather than requiring a complete charge or discharge cycle. Prior studies have developed charging-curve descriptors, restricted charging windows, and learning frameworks for this purpose [1–3]. In such settings, the observation design is determined not only by the selected partial-charge window and descriptor definitions, but also by how densely measurements are retained within that window. This motivates treating measurement density as an explicit experimental factor when evaluating a partial-charge SOH representation.

Existing work has examined several related forms of sampling and information reduction, but these interventions are not equivalent. Shu *et al.* varied temporal sampling time within a fixed charging-voltage range, while Ling *et al.* evaluated SOH performance over multiple temporal downsampling levels [4, 5]. Khaleghi *et al.* instead reduced point density through voltage-domain sampling increments in a partial-charge representation and also evaluated untrained cells or regimes [6]. Lai *et al.* used separated cells and datasets with different native sampling periods, although sampling density was not isolated as a controlled within-cohort factor [7]. Other studies have varied the available partial-charge window or the number of partial-charge events used for estimation, illustrating that reduced information can arise along dimensions other than within-event temporal density [2, 3]. Recent review evidence likewise emphasizes the need to report observation-window, measurement-quality, and sampling conditions explicitly when assessing battery-health models under partial or sparse observability [8].

Prior studies have examined partial-charge SOH estimation and several related dimensions of sampling or data availability, but the studies reviewed here did not use the same combination of design elements adopted in this work [3, 4, 6–8]. The contribution lies in the study design rather than in introducing a new regression family or treating sampling sensitivity as an unexplored problem. The design combines prespecified within-window temporal measurement-density reduction, common physical-cell validation partitions across density conditions, model selection using native-density data only and then held fixed across densities, paired cell-level performance contrasts, and simultaneous bootstrap uncertainty.

p-ISSN 3093-3285
e-ISSN 3093-3315
https://doi.org/10.51316/jst.xxx.saad.xxxx.xx.x.x
Received: Sep 5, 2026; Revised: Oct 1, 2026;
Accepted: Oct 9, 2026; Online: Nov 23, 2026

In press

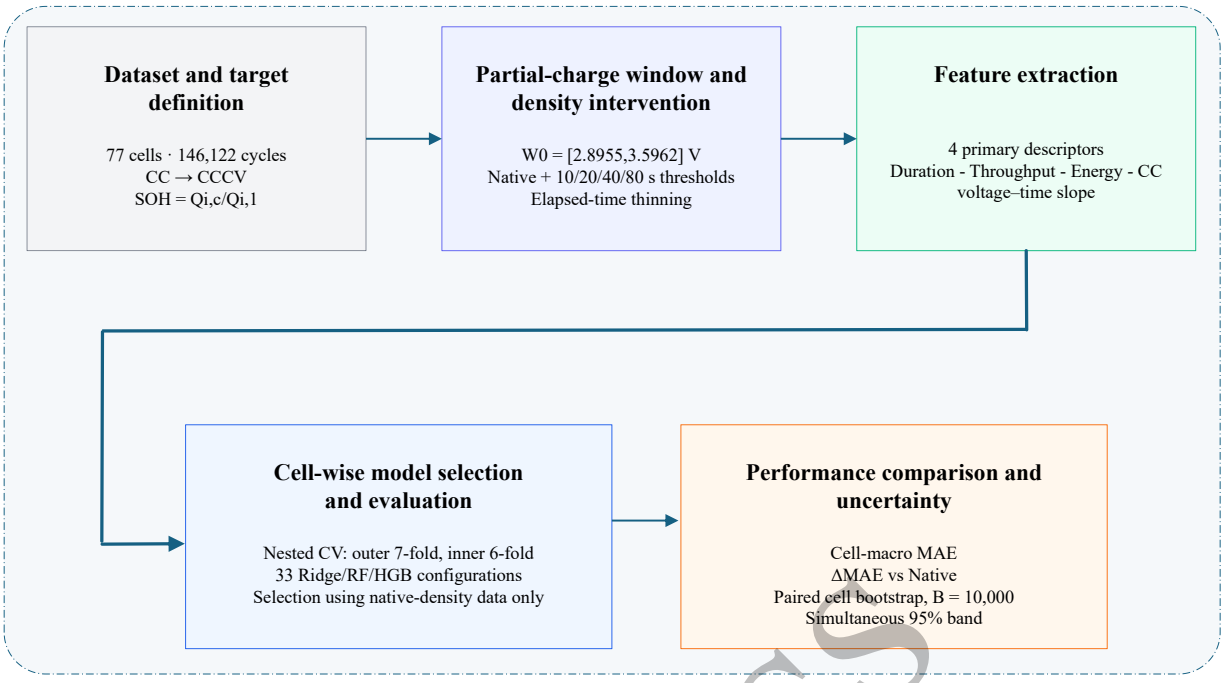


Fig. 1. Overall study design and cell-wise evaluation workflow. SOH, state of health; CC, constant current; CCCV, constant current constant voltage; CV, cross-validation; RF, random forest; HGB, histogram gradient boosting; MAE, mean absolute error; B, number of bootstrap replicates. Native denotes the original unthinned observations

Accordingly, this study examines the sensitivity of partial-charge SOH estimation to reduced measurement density under cell-wise validation on previously unseen cells. The objective is to quantify how the retained measurement fraction and estimation performance vary across the evaluated density conditions. The analysis applies a common partial-charge representation across native and progressively reduced-density conditions, recomputes the same primary descriptors after each density intervention, and compares the resulting predictions using matched physical-cell partitions. The design separates the density intervention from condition-specific model selection and evaluates paired cell-level contrasts with simultaneous uncertainty. These comparisons characterize measurement-density sensitivity but are not intended to identify an optimal sampling interval, establish equivalence to native density, or support a hardware-level sampling recommendation.

2. Materials and Methods

The overall study design and cell-wise evaluation workflow is summarized in Fig. 1.

2.1. Data Source and Operational SOH Target

The data originate from the lithium iron phosphate/graphite degradation dataset reported by Ma *et al.* [9] and released through Mendeley Data (Version 2) [10]. The analytical cohort and exclusion counts used in this study are summarized in Table 1.

Within each cycle, the charging region was identified from the exact contiguous sequence of raw status labels corresponding to constant-current (CC) charging followed by constant-current constant-voltage (CCCV) charging (CC→CCCV).

For cell i at cycle c, the discharge capacity quantity $dq_{i,c}$ for that cycle, stored in mAh, was converted to ampere hours as $Q_{i,c} = dq_{i,c}/1000$. The operational state of health (SOH) target was defined relative to the discharge capacity of the first observed cycle of the same cell, denoted $Q_{i,1}$:

$$\mathrm{SOH}_{i,c} = \frac{Q_{i,c}}{Q_{i,1}}, \qquad Q_{i,c} = \frac{dq_{i,c}}{1000} \tag{1}$$

Thus, $\mathrm{SOH}_{i,1} = 1$ by construction. Values above unity and local increases were retained; no clipping or monotonic correction was applied.

2.2. Partial Charge Window and Measurement Density Design

Partial charging was represented by one global window defined by observed voltage thresholds, $W_0 = [2.8955, 3.5962]$ V. The numerical bounds were derived before primary model-performance evaluation from a structural coverage analysis of the contiguous CC→CCCV charging records. The resulting interval was the common ordered-support domain shared by all 146,122 cycles from all 77 cells, and a subsequent assessment confirmed that it remained strictly partial relative to the full charging block in every cycle by both observed-sample count and duration. The window

In press

was therefore fixed before primary model-performance evaluation and applied unchanged across all density conditions; no SOH target or model-performance information was used to select or refine it. Within the contiguous charging region, the segment began at the first observed sample with $V \geq 2.8955$ V and ended at the first subsequent observed sample with $V \geq 3.5962$ V, inclusive. Native row and time order were preserved, with no interpolation or voltage sorting.

The native data and four prespecified conditions with reduced measurement density were evaluated; their numerical thresholds are listed in Table 1. For each reduced condition, the first observed W_0 sample was retained, each subsequent observed sample was retained once the elapsed time from the most recently retained sample reached the specified threshold, and the final observed W_0 sample was preserved if not already selected. The thresholds therefore define retention criteria based on elapsed time rather than exact uniform sampling periods. No synthetic uniform grid, selection of every kth sample by index, or forced retention of the internal CC→CCCV boundary was used. Calculations involving time used the actual retained timestamps and adjacent Δt values.

2.3. Predictors and Candidate Regression Models

Four predictors were recomputed independently for every cycle and density condition. Partial charge duration was the elapsed time between the first and last retained W_0 samples. Charge throughput was obtained by composite trapezoidal integration of retained current over the actual retained timestamps and expressed in Ah; charge energy was obtained analogously from the product of retained voltage and current and expressed in Wh. The voltage versus time slope in the constant current region was estimated by ordinary least squares using only retained samples carrying the status label for constant current charge. These descriptor families are consistent with prior feature engineering from charging curves for battery SOH estimation [1]; the exact definitions and rule based on status labels were specific to this study.

The candidate models were ridge regression, random forest (RF) regression, and histogram gradient boosting (HGB) regression. Ridge used standardization fitted using the training data only within the applicable partition and density condition, whereas RF and HGB used the predictors without external scaling. The size of the candidate set is reported in Table 1. Ridge used an SVD solver with $\alpha \in \{10^{-2}, 10^{-1}, \ldots, 10^{6}\}$. RF used 300 trees with maximum depth $\{8, 16\}$, minimum leaf size $\{1, 10, 50\}$, and maximum feature fraction $\{1.0, 0.5\}$. HGB used maximum leaf nodes $\{7, 15, 31\}$ and minimum leaf size $\{20, 100\}$ with learning rate and iteration schedules $(0.10, 100)$ and $(0.05, 200)$; the prescribed iteration counts were used without early stopping. The candidate set was fixed before model evaluation. Random states were fixed for the stochastic models to ensure reproducibility, whereas Ridge was deterministic.

2.4. Nested Validation and Model Selection by Physical Cell

Generalization was evaluated strictly across physical cells using nested grouped cross-validation. The outer and inner fold counts and their corresponding train and evaluation cell counts are given in Table 1. One predefined group of physical cells was held out in each outer fold; the remaining groups formed the outer training set. Inner validation rotated only among those outer training groups. The same physical cell partition was reused for every measurement density condition, so no cell appeared in both training and evaluation sets and density comparisons used matched evaluation cells.

Model family and hyperparameter selection was performed separately for each outer fold using only native density data. Every candidate configuration was scored by the mean of cell MAEs from inner validation, with equal weight assigned to each cell, and the lowest score was selected. Ties were resolved deterministically using an ordering defined before evaluation. After selection, the chosen family and hyperparameters were held fixed across all density conditions; a fresh estimator was refitted separately for each density using the complete outer training set and evaluated once on the corresponding evaluation cells. Ridge scaling was therefore fitted using outer training data from the same density condition only. Model family and hyperparameters were not retuned separately for individual density conditions.

2.5. Performance Measures and Uncertainty Analysis

The primary performance measure was mean absolute error (MAE) macroaveraged across physical cells, so each cell received equal weight irrespective of its number of evaluated cycles. For density condition d,

$$\mathrm{MAE}_d = \frac{1}{N_{\mathrm{cell}}} \sum_{i=1}^{N_{\mathrm{cell}}} \left[\frac{1}{n_i} \sum_{c} \left| y_{i,c} - \hat{y}_{i,c,d} \right| \right] \tag{2}$$

Here, $N_{\mathrm{cell}} = 77$, n_i is the number of evaluated cycles for cell i, $y_{i,c}$ is the operational SOH target, and $\hat{y}_{i,c,d}$ is the prediction for cycle c of evaluation cell i at density d. Secondary measures were root mean square error macroaveraged across physical cells and cell-balanced R^2; signed bias macroaveraged across physical cells was retained as a diagnostic. Cell-balanced R^2 was computed from all cycle-level predictions with each physical cell assigned equal total weight: each cycle from cell i received weight $1/(N_{\mathrm{cell}} n_i)$, and the corresponding weighted mean of observed SOH values was used in the total sum of squares. This metric was not computed as the mean of per-cell R^2 values.

In press

Table 1. Core quantitative study and validation parameters

Design block	Quantity	Value
Data	Physical cells	77
	Cycle records	146,122
	Excluded whole cells / cycles	0 / 0
Representation	W_0 bounds	2.8955–3.5962 V
	W_0 width	0.7007 V
	Measurement density conditions	5
	Reduced density thresholds	10, 20, 40, 80 s
	Primary predictors	4
Validation	Outer grouped CV	$K = 7$; 66 train / 11 evaluation cells
	Inner grouped CV	$K = 6$; 55 train / 11 validation cells
Models	Candidate configurations	33 total (Ridge 9; RF 12; HGB 12)
Statistics	Primary ΔMAE contrasts	4
	Bootstrap replicates	10,000
	Bootstrap resampling units	77 physical cells

The primary density effects were the four paired MAE contrasts relative to the native condition,

$$\Delta\mathrm{MAE}_d = \mathrm{MAE}_d - \mathrm{MAE}_{\mathrm{native}} \tag{3}$$

Positive values indicate a larger MAE than the native condition; negative values were retained without post hoc correction.

Uncertainty was quantified from predictions for the evaluation cells by paired bootstrap resampling of physical cells without model refitting or retuning. Accordingly, the resulting intervals are conditional on the out-of-sample cross-validated predictions and quantify variability under resampling of the 77 evaluated physical cells; they do not capture the additional uncertainty that would arise from repeating model fitting and model selection on a new cohort. The number of replicates and resampling units are summarized in Table 1. The same sampled cell indices were applied simultaneously to all density conditions to preserve pairing, with a fixed random state used for reproducibility. Marginal 95% percentile intervals used linear quantiles. For the four primary ΔMAE contrasts, simultaneous 95% uncertainty used one common half width equal to the 95th percentile of the bootstrap distribution of the maximum absolute centered deviation across the four contrasts. Interpretation therefore emphasized paired effect estimates and simultaneous bootstrap bands; no primary null hypothesis significance tests or p-value thresholds were used, and the outer folds were not treated as independent experimental replicates.

3. Results

The elapsed-time thinning procedure progressively reduced the number of observations retained within the partial charging window. Across cycles, the median retained count decreased from 250 observations at native density to 125, 63, 32, and 17 observations at the 10, 20, 40, and 80 s retention thresholds, corresponding to median retained fractions of 1.000, 0.500, 0.253, 0.130, and 0.068, respectively. Table 2 summarizes the realized density reduction and predictive performance.

Cell-macro MAE was 0.012725 at native density and 0.012772, 0.012895, 0.013089, and 0.013000 at the 10, 20, 40, and 80 s thresholds, respectively. All four reduced-density conditions therefore had higher MAE point estimates than the native condition, but the trajectory was nonmonotonic: MAE increased through the 40 s condition and then decreased slightly at 80 s. Cell-macro RMSE likewise had its largest point estimate at 40 s, while cell-balanced R^2 had its lowest point estimate at that condition.

Relative to the native condition, the paired ΔMAE estimates were 0.000047, 0.000170, 0.000364, and 0.000275 for the 10, 20, 40, and 80 s thresholds. The corresponding simultaneous 95% bootstrap bands were $[-0.000386, 0.000480]$, $[-0.000263, 0.000604]$, $[-0.000069, 0.000797]$, and $[-0.000158, 0.000708]$, respectively; all four bands included zero. Thus, although all four ΔMAE point estimates were positive, the simultaneous uncertainty did not establish a consistent increase in MAE across the four reduced-density conditions. Because no equivalence margin was prespecified, inclusion of zero was not interpreted as evidence of equivalence. Fig. 2(a) visualizes these four native-referenced contrasts and their simultaneous uncertainty. Fig. 2(b) provides signal-level context by overlaying the cycle-1 voltage trajectories within W_0 for all 77 cells; the highlighted trajectory corresponds to the representative cell used in Fig. 3.

Within the prespecified set of 33 Ridge, Random Forest, and Histogram Gradient Boosting configurations, native-only inner cross-validation selected a Random Forest configuration in all seven outer folds. This selection result is restricted to the candidate set evaluated in this study.

In press

Table 2. Measurement density realization and predictive performance. Retained counts are medians across cycles, and retained fractions are medians of cycle-specific fractions relative to the native retained count of the same cycle. MAE and RMSE are cell-macro summaries, and R^2 is cell-balanced. The 10–80 s conditions denote elapsed-time retention thresholds rather than uniform sampling periods. The final column reports the paired change in cell-macro MAE relative to native density with the simultaneous 95% bootstrap band

Condition	**Median retained *n* (fraction)**	**MAE**	**RMSE**	R^2	ΔMAE **[simultaneous 95% band]**
Native	250 (1.000)	0.012725	0.016182	0.9412	—
10 s	125 (0.500)	0.012772	0.016250	0.9406	+0.000047 [−0.000386, 0.000480]
20 s	63 (0.253)	0.012895	0.016485	0.9387	+0.000170 [−0.000263, 0.000604]
40 s	32 (0.130)	0.013089	0.016647	0.9370	+0.000364 [−0.000069, 0.000797]
80 s	17 (0.068)	0.013000	0.016541	0.9372	+0.000275 [−0.000158, 0.000708]

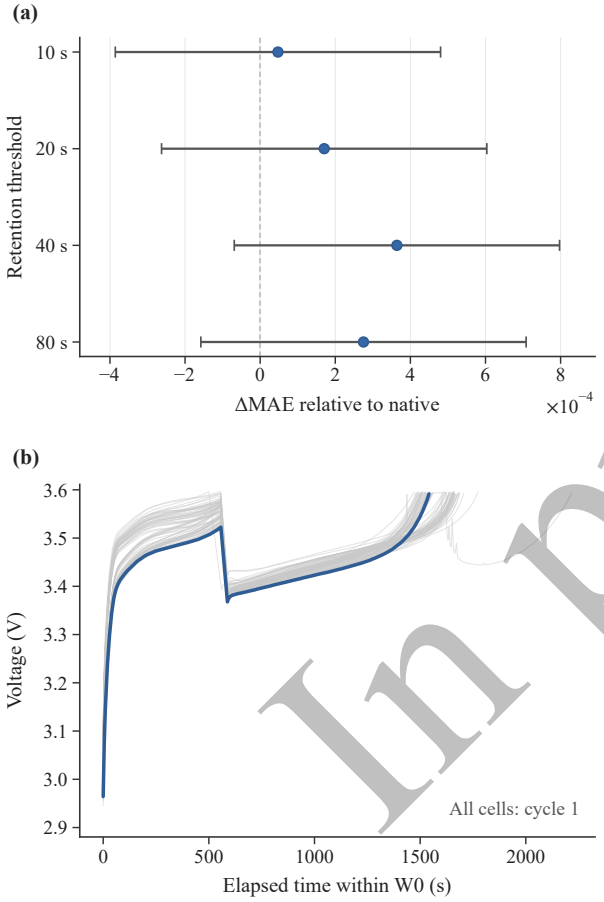


Fig. 2. Measurement-density contrasts and charging-curve context. (a) native-referenced ΔMAE estimates with simultaneous 95% bootstrap bands; the dashed line marks zero. (b) Cycle-1 voltage trajectories within W_0 for all 77 cells; the highlighted trace is the representative cell used in Fig. 3

For an illustrative single-cell view, the representative cell was selected deterministically as the cell whose mean cell-level MAE across all five density conditions was closest to the cohort median. Fig. 3 shows all available cycles for that same cell under all five measurement-density conditions. Table 3 quantifies the visible late-cycle pattern using cycle 1200 as a descriptive division. MAE was higher over cycles 1201–1690 than over cycles 1–1200 in all five conditions, and the mean late-cycle error was positive throughout. Relative to Native, the mean absolute prediction difference over cycles 1201–1690 increased from 0.00256 at 10 s to 0.00619 at 80 s, whereas the signed mean difference varied in direction across conditions. Because the cycle-1200 division was not prespecified and the comparison concerns one representative cell, it is treated as descriptive; population-level inference remains based on the paired 77-cell analysis reported in Table 2 and Fig. 2(a).

4. Discussion

Under cell-wise validation, all four prespecified reduced-density conditions yielded higher cell-macro MAE point estimates than the native condition, but the trajectory was nonmonotonic, and the simultaneous 95% bootstrap bands for all four native-referenced ΔMAE contrasts included zero (Table 2 and Fig. 2). Although every ΔMAE point estimate was positive, the simultaneous bands including zero do not establish a consistent increase in MAE across the four reduced-density conditions. Conversely, they do not demonstrate the absence of a measurement-density effect; equivalence to native density was not tested because no equivalence margin was prespecified.

The MAE point estimates increased from the native condition through 40 s and then decreased slightly at 80 s, showing a nonmonotonic response rather than a progressive increase in error as measurement density was reduced. Previous partial-charge SOH work has reported nonmonotonic sensitivity when sampling time was varied, while other battery-SOH studies have evaluated performance over multiple temporal downsampling levels [4, 5]. Because those studies differ in datasets, feature representations, models, sampling operators, and evaluation procedures, the comparison is qualitative rather than quantitative. Although the 40 s condition had the highest MAE point estimate, the observed pattern does not identify 40 s as a knee, threshold, or optimal retention interval.

In this design, physical-cell grouping ensures that the cells used for outer-fold evaluation remain separate from the corresponding training and model-selection data. Within each outer fold, model family and configuration were selected using native-density data

In press

Table 3. Cycle-range diagnostics for representative cell 7-1

Condition	MAE 1–1200	MAE 1201–1690	Bias 1201–1690	mean$\lvert\hat{y}_d - \hat{y}_N\rvert$ 1201–1690	mean$(\hat{y}_d - \hat{y}_N)$ 1201–1690
Native	0.00613	0.02309	+0.02230	0.00000	0.00000
10 s	0.00635	0.02375	+0.02313	0.00256	+0.00082
20 s	0.00646	0.02448	+0.02339	0.00348	+0.00109
40 s	0.00637	0.02462	+0.02219	0.00488	−0.00012
80 s	0.00664	0.02318	+0.01981	0.00619	−0.00249

Note: Bias is mean$(\hat{y} - y)$. The final two columns compare each reduced-density prediction with Native (N) over cycles 1201–1690. The cycle-1200 division is descriptive and was not prespecified.

only; the selected specification was then held fixed across density conditions, with a fresh estimator refitted at each density. Thus, differences between density conditions were evaluated without condition-specific model or configuration selection. Random Forest was selected in all seven outer folds within the prespecified set of 33 Ridge, Random Forest, and Histogram Gradient Boosting configurations. This selection pattern is specific to the evaluated candidate set and should not be interpreted as evidence that Random Forest is generally superior.

Across the tested density range, the median number of retained observations within W_0 decreased from 250 at native density to 17 at the 80 s threshold, while the same four primary descriptor definitions and evaluation protocol were retained across all five conditions. The four-descriptor representation could therefore be evaluated at every tested density, but these results do not establish predictor invariance, preservation of predictive information, or a minimum sufficient measurement density. Fig. 2(b) places the charging trajectory of the representative cell in the context of the full cohort, whereas Fig. 3 and Table 3 provide a trajectory-level complement to the cohort-level contrasts. The higher late-cycle MAE and positive bias under Native as well as all reduced-density conditions indicate a prediction difficulty common to this portion of the representative trajectory. Temporal thinning additionally altered the predicted trajectory relative to Native: the mean absolute difference increased with the retention threshold, but the signed difference changed direction. At 80 s, the largest mean absolute departure from Native (0.00619) was accompanied by a negative signed shift (−0.00249), which partly offset the positive Native bias and yielded a late-cycle MAE (0.02318) close to the Native value (0.02309). Thus, the effect of a density-related trajectory perturbation on absolute error depended on both its magnitude and direction, providing a trajectory-level explanation for the nonmonotonic MAE pattern in this cell. Because the cell was selected by a deterministic cohort-median rule and the cycle-range comparison was not prespecified, this evidence remains illustrative rather than a basis for population-level inference.

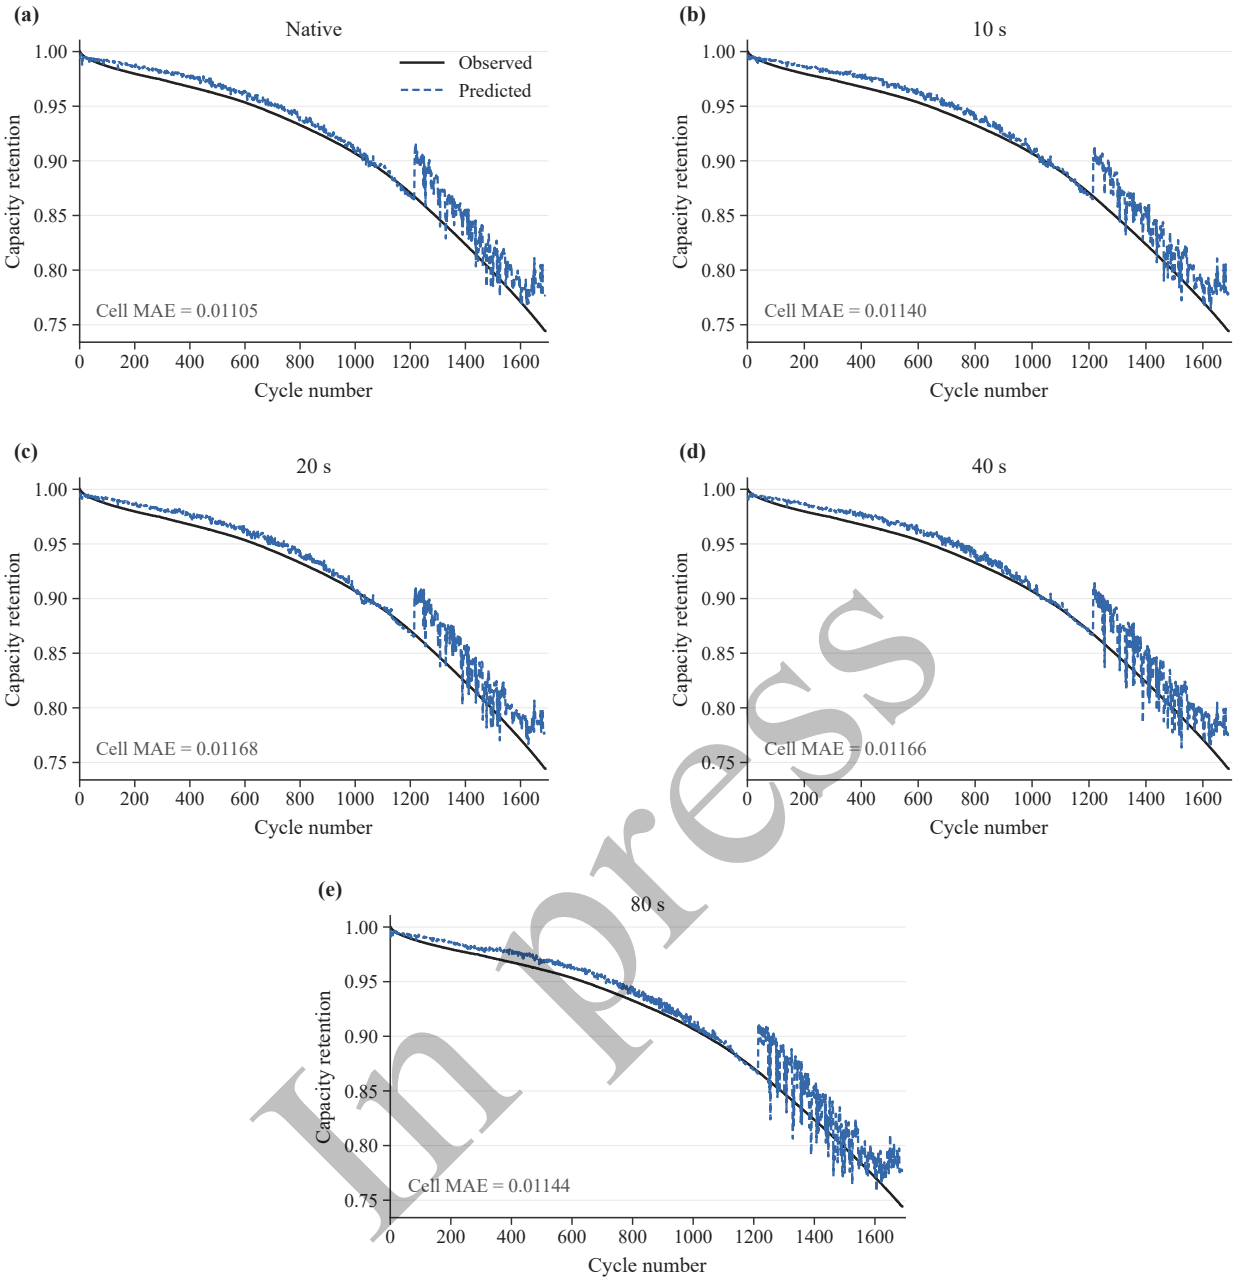
The findings are limited to the tested cell cohort, one prespecified partial-charge window, one prespecified four-descriptor representation, and a finite candidate set spanning three regression families. Measurement density was varied by deterministic elapsed-time thinning of previously acquired observations rather than by collecting data under different hardware sampling regimes. Accordingly, the 10–80 s conditions should be interpreted as retention thresholds applied offline to previously acquired data, not as experimentally evaluated hardware sampling intervals. The results therefore do not directly establish a hardware sampling requirement, communication requirement, or acquisition strategy for a battery management system. Storage, communication bandwidth, processor time, and power savings were not evaluated in this study. All four simultaneous bands included zero, and no equivalence margin was defined, so the results support neither practical equivalence to the native condition nor a firm degradation conclusion across all reduced-density conditions.

The sensitivity pattern observed here should be tested across additional cell populations, partial-charge windows, descriptor sets, model families, and hardware-level acquisition conditions. Prior work has examined other forms of reduced information, including changes in the available partial-charge window and in the number of partial-charge events used for SOH estimation [2, 3]. These constraints differ from within-event measurement density and warrant separate evaluation. Studies intended to determine practical interchangeability between reduced-density and native acquisition should prespecify application-relevant equivalence or noninferiority margins and use an uncertainty framework designed for that purpose. The present findings support treating measurement density as an explicit experimental factor rather than assuming either a monotonic response or a negligible effect across settings.

5. Conclusion

Under cell-wise validation, all four prespecified reduced-density conditions yielded higher cell-macro MAE point estimates than the native condition, but the trajectory was nonmonotonic, and the simultaneous 95% bootstrap bands for all four native-referenced ΔMAE contrasts included zero. These intervals do not support a conclusion of no measurement-density effect, and equivalence to the native condition was not tested.

In press



In press

Fig. 3. Representative SOH prediction trajectories. Observed and predicted capacity retention for the same deterministically selected cell under (a) Native, (b) 10 s, (c) 20 s, (d) 40 s, and (e) 80 s conditions. All available cycles are shown; the single-cell display is illustrative rather than a basis for the primary population-level inference

These findings support treating measurement density as an explicit evaluation factor in partial-charge SOH estimation rather than assuming a monotonic response or a negligible effect. These conclusions apply to the tested cohort and to deterministic elapsed-time thinning of existing observations, rather than to hardware acquisition under altered sampling regimes. Studies intended to assess practical interchangeability should prespecify application-relevant equivalence or noninferiority margins and use an uncertainty framework suited to that objective.

Acknowledgment

This research is funded by the Vietnam National Foundation for Science and Technology Development (NAFOSTED) under grant number 31/2025/TN.

References

[1] D. Roman, S. Saxena, V. Robu, M. Pecht, and D. Flynn, Machine learning pipeline for battery state-of-health estimation, Nature Machine Intelligence, vol. 3, no. 5, pp. 447–456, May 2021. https://doi.org/10.1038/s42256-021-00312-3

[2] A. Falai, T. A. Giuliacci, D. A. Misul, and P. G. Anselma, Reducing the computational cost for artificial intelligence-based battery state-of-health estimation in charging events, Batteries, vol. 8, iss. 11, Nov. 2022, Art. no. 209 .
https://doi.org/10.3390/batteries8110209

[3] S. B. Vilsen and D.-I. Stroe, Sensitivity analysis for reliable state-of-health estimation based on battery partial charging, Cell Reports Physical Science, vol. 6, iss. 6, Jun. 2025, Art. no. 102646 .
https://doi.org/10.1016/j.xcrp.2025.102646

[4] X. Shu, G. Li, J. Shen, Z. Lei, Z. Chen, and Y. Liu, A uniform estimation framework for state of health of lithium-ion batteries considering feature extraction and parameters optimization, Energy, vol. 204, Aug. 2020, Art. no. 117957 .
https://doi.org/10.1016/j.energy.2020.117957

[5] L. Ling, S. Yang, and S. Tong, Lithium-ion battery state of health estimation based on statistical features derived from voltage and temperature probability density functions under realistic randomized cycling conditions, Journal of Energy Storage, vol. 137, Nov. 2025, Art. no. 118496 .
https://doi.org/10.1016/j.est.2025.118496

[6] S. Khaleghi, D. Karimi, S. H. Beheshti, M. S. Hosen, H. Behi, M. Berecibar, and J. Van Mierlo, Online health diagnosis of lithium-ion batteries based on nonlinear autoregressive neural network, Applied Energy, vol. 282, Jan. 2021, Art. no. 116159 .
https://doi.org/10.1016/j.apenergy.2020.116159

[7] R. Lai, X. Li, and J. Wang, Flexible deep learning-based state of health estimation of lithium-ion batteries with features extracted from partial charging curves, Batteries, vol. 10, iss. 5, May 2024, Art. no. 164 .
https://doi.org/10.3390/batteries10050164

[8] S. Rizanov, A. Stoynova, and G. Mihov, Toward trustworthy and transferable SOH/RUL estimation for lithium-ion batteries: A critical review and multi-fidelity validation framework from laboratory cells to real-world packs, Batteries, vol. 12, iss. 7, Jul. 2026, Art. no. 255 .
https://doi.org/10.3390/batteries12070255

[9] G. Ma, S. Xu, B. Jiang, C. Cheng, X. Yang, Y. Shen, T. Yang, Y. Huang, H. Ding, and Y. Yuan, Real-time personalized health status prediction of lithium-ion batteries using deep transfer learning, Energy & Environmental Science, vol. 15, no. 10, pp. 4083–4094, Oct. 2022.
https://doi.org/10.1039/D2EE01676A

[10] Y. Yuan, G. Ma, and S. Xu, The Dataset for: Real-time personalized health status prediction of lithium-ion batteries using deep transfer learning, Mendeley Data, Version 2, May 2022.
https://doi.org/10.17632/nsc7hnsg4s.2

In press