



## Research papers

# Field-data-driven lithium-ion battery diagnostics under coarse onboard voltage quantization: Pack-constrained voltage reconstruction and joint open-circuit-voltage and equivalent-circuit-model estimation

Younggyu Hwang, Hyeongyu Son, Seungho Geum, Musu Kim, Sekyung Han \*

*School of Electronic and Electrical Engineering, Kyungpook National University, 80 Daehak-ro, Buk-gu, Daegu, 41566, Republic of Korea*

## ARTICLE INFO

## Keywords:

Electric vehicle battery management system  
Field-data-driven diagnosis  
Cell-voltage reconstruction  
Joint OCV–ECM estimation  
State-of-charge estimation

## ABSTRACT

Production battery management system (BMS) telemetry from electric vehicles (EVs) provides a large practical source of real-world battery-degradation data, but its diagnostic value is constrained by two measurement limitations: coarse 20 mV cell-voltage quantization and biased open-circuit-voltage (OCV) references that distort equivalent-circuit-model (ECM) identification. The evidence comes from 251 million records from 116 EVs across 14 Hyundai and Kia vehicle models. A separate paired-test cohort contains 135 vehicles across four pack architectures. This study demonstrates that software-only processing of existing onboard measurements can address both limitations. Voltage reconstruction, joint OCV–ECM identification, and cell-level diagnostics used slow constant-current charging segments; fast-charging telemetry served only a model-order identifiability check, and state-of-charge (SOC) evaluation used a separate driving-window subset. The framework uses the alternating direction method of multipliers (ADMM) to reconstruct cell voltages while enforcing consistency with pack voltage. It then jointly estimates the OCV curve and the parameters of a first-order equivalent circuit (1-RC). Model-order analysis shows that BMS-quality data do not reliably support a more complex circuit. Across the paired-test cohort, reconstruction reduced per-architecture mean voltage-fitting root-mean-square error (RMSE) by 30–44%. Joint estimation restored fleet-median ohmic resistance to a physically plausible level, whereas the charging-derived OCV drove it to an implausible near-zero value. The recovered parameters reduced driving-window SOC RMSE from a fleet mean of 5.2% to 3.2% in the six-vehicle EV6 subset and enabled persistent cell-deviation analysis for fleet-level anomaly screening. These results establish that coarse production-grade EV BMS data can support physically interpretable battery diagnosis without hardware modification.

## 1. Introduction

Battery diagnosis in real-world electric vehicles increasingly relies on onboard BMS measurements because laboratory-grade characterization is difficult to scale to large fleets [7,8]. Laboratory cyclers provide high-precision measurements and remain the standard basis for battery modeling and state estimation studies [1,3]. However, laboratory testing is costly, low-throughput, and typically limited to single-cell experiments, making direct fleet-scale characterization impractical. In contrast, onboard telemetry continuously monitors dozens to hundreds of cells per pack across large vehicle populations, creating a practical foundation for large-scale battery analysis and digital-twin-oriented monitoring [9–11]. This accessibility, however, comes at the cost of reduced sensing fidelity. Cost-optimized BMS hardware commonly digitizes cell voltages through embedded battery-monitoring analog-to-

digital converters (ADCs), including stack-monitoring integrated circuits with successive-approximation-register ADCs [32]. Because ECM parameter identifiability under such measurement constraints is governed by voltage-noise statistics and excitation content [31], coarse cell-voltage quantization fundamentally degrades the voltage continuity required for ECM identification, resistance estimation, and downstream state tracking. In our dataset, this appears as a universal 20 mV cell-voltage quantization pattern across all analyzed vehicle models, producing staircase voltage trajectories and derivative spikes that directly corrupt model identification and diagnostic features.

A second challenge is that cell voltages within a pack are physically coupled rather than independent. Their sum must remain consistent with the measured pack voltage, yet most existing signal-conditioning methods do not explicitly exploit this constraint. Kalman filtering is widely used for SOC and state of health (SOH) estimation [1,12–14],

\* Corresponding author.

E-mail address: [skhan@knu.ac.kr](mailto:skhan@knu.ac.kr) (S. Han).

<https://doi.org/10.1016/j.est.2026.124005>

Received 20 March 2026; Received in revised form 13 July 2026; Accepted 31 July 2026  
2352-152/© 20XX

but it is primarily designed for recursive state tracking rather than batch reconstruction of recorded quantized voltage trajectories. Moving-average and low-pass filtering suppress noise but also blur genuine voltage transitions [3]. Wavelet-based denoising has been applied to non-stationary battery voltage signals and can retain transient information better than simple low-pass smoothing [33], but it still treats each cell independently and does not enforce pack-level consistency. Total variation (TV) regularization is well suited to staircase-like signals induced by ADC quantization [6], and the ADMM framework provides a natural framework for combining TV regularization with explicit constraints [5]. Although ADMM has been successfully used in broader signal restoration problems [18,19], its application to multi-cell BMS voltage reconstruction under an exact pack-sum constraint has not, to our knowledge, been established. This gap is important because the objective is not simply to smooth the signal, but to reconstruct physically admissible voltage trajectories that retain diagnostic value without hardware replacement.

Even with voltage preprocessing, reliable ECM identification still depends strongly on the OCV reference used during fitting. ECMs remain widely used in battery management because they provide a practical balance between physical interpretability and computational efficiency [1–3,21,24,25]. Prior studies have shown that ECM parameters can support degradation tracking and practical BMS implementation, while also emphasizing that identifiable model complexity is limited by measurement quality and excitation bandwidth [2,3,22,26–28]. Nevertheless, model order in practical ECM studies is often assumed from laboratory convention, typically using 1-RC or 2-RC structures, even when the available telemetry is restricted by 20 mV quantization and 3–4 s sampling intervals. Although information-criterion approaches such as Akaike information criterion (AIC) and Bayesian information criterion (BIC) have been explored for ECM model selection in controlled conditions, it remains unclear whether higher-order RC networks are justified for data of this quality. This raises a practical question for field diagnosis: does improved performance come from increasing model complexity, or from making better use of limited real-world measurements?

The OCV is the cell's equilibrium terminal voltage at zero current, and the OCV source used during ECM identification is particularly consequential because OCV anchors the decomposition of the measured loaded terminal voltage into ohmic drop and dynamic polarization. In laboratory studies, OCV is usually characterized under low-rate cycling or galvanostatic intermittent titration (GITT) conditions that minimize polarization effects [4,29]. In real-world applications, however, OCV must often be approximated from operational data, such as rest-period voltages, slow-charging constant-current segments, or voltage–SOC mappings constructed from the same session used for ECM fitting [2, 15]. Birkel et al. showed that OCV accuracy strongly affects downstream parameter estimation [4], yet the quantitative influence of OCV source selection on ECM fitting has received limited attention. This issue becomes particularly severe under fast charging, where large current-induced voltage drops distort the voltage–SOC relationship. When an OCV curve is constructed from the same high-current session being fitted, it can absorb ohmic and polarization losses, causing the fitted ohmic resistance to collapse toward a non-physical lower bound. As a result, voltage-fitting accuracy may improve while parameter physicality deteriorates. Although recent work on physics-constrained OCV reconstruction suggests that joint estimation of OCV and model parameters may mitigate this problem, this trade-off has not been systematically examined for coarse onboard EV battery data.

To address these limitations, this paper proposes an end-to-end diagnostic framework that extracts physically meaningful battery information directly from coarse onboard BMS measurements. The framework combines pack-constrained ADMM-based voltage reconstruction, data-driven ECM model-order selection, and a joint OCV-ECM optimization. In the joint stage, a 50-knot piecewise-linear OCV curve and the 1-RC ECM parameters are estimated simultaneously by gradient-

based minimization of a composite loss (voltage fidelity, rest-period anchoring, smoothness, monotonicity, and boundary penalties), implemented as a differentiable forward simulation in JAX and solved with the Adam optimizer (optax library) over a fixed number of epochs with exponential learning-rate decay; full details and weight selection are given in Section 3.4. The framework is deployed as a cloud-based offline diagnostic pipeline that processes stored telemetry data, and real-time onboard adaptation is outside the present scope and is discussed as a limitation in Section 5. The main contributions are threefold. First, we propose a pack-constrained voltage reconstruction method that restores physically admissible multi-cell voltage trajectories from quantized BMS data while explicitly enforcing consistency with pack voltage. Second, we evaluate ECM model-order identifiability directly under production-level EV telemetry and show that a 1-RC model is the highest-order identifiable structure under the sensing limitations of real-world BMS data. Third, we show that OCV source quality has a greater influence on physically meaningful ECM identification than increased model complexity, and we demonstrate that the resulting framework supports practical battery diagnosis through improved long-horizon SOC estimation and persistent abnormal-cell detection using only existing onboard BMS data (Fig. 1).

## 2. BMS data characteristics

### 2.1. Dataset overview

The dataset consists of time-series BMS telemetry collected from 116 electric vehicles operated under real-world driving and charging conditions, stored in a large-scale relational time-series database comprising approximately 475 GB and 251 million records [9–11]. The 116-vehicle fleet covers 14 distinct vehicle models from two brands within the Hyundai Motor Group (Hyundai and Kia), with pack configurations ranging from 84 to 192 series-connected cells. Per-vehicle data span is approximately one year of continuous telemetry (collected over the 2024–11 to 2025–12 window for the analysis subset reported here), with absolute start dates varying across vehicles based on each subscriber's enrollment time. Because per-vehicle battery degradation is anchored to cumulative cycle/mileage indicators rather than absolute calendar time, the staggered start dates do not affect the within-vehicle analysis structure used in this study. Each record contains pack voltage, pack current, SOC, temperature measurements, mileage, and the individual cell-voltage array (cell count per vehicle ranges from 84 to 192). All vehicles in the analyzed cohort use nickel-manganese-cobalt (NMC) cathode chemistry per manufacturer specifications; lithium iron phosphate (LFP) chemistry is not present in this fleet. Year-to-year manufacturing variations (cell supplier, vintage) exist within the cohort and contribute to per-vehicle variability of the identified ECM parameters. The 116 vehicles are concentrated in the joint-analysis-eligible models — KIA EV6 long-range (LR) (26 vehicles, used as the primary joint OCV-ECM analysis cohort in Section 4.6) and the other slow-charging-capable architectures (KIA EV9, Hyundai IONIQ5 LR, IONIQ5 LR 2022, KIA NIRO LR, KONA LR, IONIQ'19, NIRO+, BONGO3, PORTER2, IONIQ6 LR) — with the remaining 5 vehicles in three minority models (KIA NIRO2, Hyundai Motor Company (HMC) IONIQ5 STANDARD, HMC KONA LR 2nd generation) used only at the corpus-coverage level. Of the 14 models present in the fleet, 11 representative models with sufficient per-model coverage are tabulated in Table 1 and used in the cross-model ADMM preprocessing analysis (Section 4.1); the remaining 3 models (HMC IONIQ5 STANDARD, KIA NIRO2, HMC KONA LR 2nd generation; 5 vehicles total) are present in the underlying corpus but not separately analyzed.

For the fleet-level joint OCV-ECM analysis, identification was conducted on 26 KIA EV6 vehicles that provided sufficient slow-charging data for OCV construction and parameter identification. In addition, one imported model — the Audi Q4 e-tron 82 kWh (96 series-connected




Fig. 1. Overall workflow for battery diagnosis using coarse real-world BMS data.

Table 1

Fleet summary of voltage quantization, pack-sum offset, and sampling interval across 11 vehicle models.

| Model        | Cells | Quant. (mV) | 20mV Ratio (%) | Offset/cell (mV) | Sampling (s) |
|--------------|-------|-------------|----------------|------------------|--------------|
| PORTER2      | 90    | 20          | 94.4           | 9.36             | 3.00         |
| BONGO3       | 90    | 20          | 94.4           | 9.28             | 3.01         |
| IONIQ'19     | 96    | 20          | 100.0          | 9.75             | 3.00         |
| KONA LR      | 98    | 20          | 99.1           | 9.40             | 3.01         |
| NIRO LR      | 98    | 20          | 100.0          | 9.28             | 3.09         |
| NIRO +       | 98    | 20          | 87.5           | 9.34             | 3.09         |
| EV9          | 152   | 20          | 99.9           | 9.58             | 3.71         |
| IONIQ5 LR    | 180   | 20          | 100.0          | 9.42             | 3.77         |
| IONIQ5 LR'22 | 192   | 20          | 100.0          | 9.58             | 3.65         |
| IONIQ6 LR    | 192   | 20          | 100.0          | 9.57             | 3.33         |
| EV6 LR       | 192   | 20          | 100.0          | 9.56             | 3.66         |

cells) — is included in the cross-model ADMM preprocessing summary in Section 4.1 to confirm that the pack-constrained reconstruction generalizes to a non-Hyundai/Kia pack architecture; this model does not enter the joint OCV-ECM or driving-window extended Kalman filter (EKF) analysis.

Overall, the dataset provides a realistic basis for battery diagnosis at fleet scale, but the measurements are constrained by several production-BMS limitations, including coarse cell-voltage quantization, pack-level measurement mismatch, and sequential cell scanning artifacts. These characteristics are examined in the following subsections.


## 2.2. Cell voltage quantization

Analysis of consecutive cell-voltage differences shows that all vehicle models share a common 20 mV quantization grid (Table 1). The 20 mV ratio, defined as the fraction of consecutive voltage differences aligned with the 20 mV step, ranges from 87.5% to 100.0% across the

fleet (the lowest values, 87.5% for NIRO PLUS and 94.4% for PORTER2/BONGO3, still place the dominant inter-sample voltage difference on the 20 mV grid), indicating that the cell-voltage measurements are strongly constrained by the underlying BMS resolution. Pack voltage is quantized separately at 100 mV resolution.

Each Table 1 entry was computed directly from the underlying raw telemetry as follows. KIA EV6 LR, which has 192 series-connected cells, provides a representative worked example. First, the 20 mV ratio is the fraction of consecutive non-zero cell-voltage differences whose magnitude lies within  $\pm 2$  mV of an integer multiple of 0.020 V across all charging-session timestamps and all 192 cells. This integer-multiple metric yields 100.0% for EV6; equivalently, the dominant  $\pm 20$  mV peak illustrated in Fig. 2b contains 96.9% of the non-zero  $\Delta v$  samples from the same session. Second, the per-cell pack-sum offset is the time average of  $\frac{(V_{pack} - \sum_i V_{cell,i})}{N}$  over slow constant-current charging timestamps, yielding 9.56 mV/cell for EV6. Third, the sampling interval is the median inter-sample timestamp difference, which is 3.66 s for EV6. The same three computations were applied identically to the other 10 representative models in Table 1.

The 20 mV step is directly visible in the raw cell-voltage trace in Fig. 2(a) and dominates the histogram of consecutive voltage differences in Fig. 2(b). Fig. 2(a) deliberately uses millivolts so that each 20 mV quantization step spans exactly one axis tick. The measurement granularity can therefore be read directly from the voltage axis. Throughout this paper, the 20 mV step is treated as the effective granularity of the reported telemetry: the recorded voltage is what appears on the vehicle CAN bus after BMS-side post-processing, so no specific sensor ADC bit count is claimed. This coarse quantization affects the data in two ways. First, voltage trajectories exhibit staircase-like behavior, with abrupt level changes separated by long flat regions; this corrupts time derivatives required for ECM identification and produces spurious activity in difference-based diagnostic features. Second, fine inter-cell differences below 20 mV, which are precisely the signals carrying diagnostic information about cell-to-cell variability, are masked by the quantization



**Fig. 2.** Illustration of the 20 mV cell-voltage quantization step in the onboard BMS telemetry. (a) A single-cell voltage trace during slow constant-current charging shows a staircase pattern aligned with a 20 mV grid (dotted). (b) Histogram of consecutive-sample voltage differences across all cells, showing a dominant peak at  $\pm 20$  mV.

step. As a result, directly using the raw cell-voltage measurements for ECM identification or cell-level diagnostic feature extraction can lead to biased parameter estimation and unstable analytical results [31].

These observations motivate the need for a reconstruction step that improves effective voltage resolution while preserving the physical consistency of the battery pack [5,6].

### 2.3. Pack-sum offset and bus bar resistance

A systematic discrepancy is observed between the measured pack voltage and the sum of the reported cell voltages. Across the fleet, the sum of cell voltages consistently underestimates the pack voltage by an approximately deterministic offset. When normalized by cell count, this offset is highly consistent across vehicle models, with a mean value of  $9.43 \pm 0.14$  mV per cell. Consequently, the total pack-level offset increases nearly linearly with the number of series-connected cells, ranging from approximately 0.8 to 1.8 V.


This behavior is consistent with the measurement topology of production battery packs. The pack voltage measurement includes additional resistive components in the current path, such as bus bars and contact impedances, whereas individual cell-voltage taps do not fully capture these contributions, as evidenced by the consistent pack-sum offset shown in Fig. 3. The observed offset can therefore be interpreted as the sum of a static component and a current-dependent dynamic component.

Let  $\beta_0$  denote the static intercept of the pack-sum offset. Subtracting  $\beta_0$  from the total offset reveals the dynamic component. The remaining offset varies approximately linearly with current, and the fitted slope corresponds to an estimated total bus-bar resistance  $R_{bus} = 28.5$  m $\Omega$ , or  $R_{bus}/N = 0.148$  m $\Omega$  per cell. Under slow charging at 7 A, the corresponding per-cell voltage contribution is only 1.04 mV, which is small relative to the 20 mV quantization step. However, the effect increases to 7.41 mV at 50 A, 14.82 mV at 100 A, and 29.64 mV at 200 A, eventually exceeding the quantization step itself.

In Fig. 4(a), the green dashed line shows the single-session slope from the anchor vehicle (0.012 m $\Omega$ /cell), which is lower than the fleet-aggregate slope because limited current range and SOC-driven cell-voltage modulation within a single session elevate the regression noise floor; the fleet-aggregate value of 0.148 m $\Omega$ /cell (red line) is therefore used for the per-current summary in Fig. 4(b).

### 2.4. Sequential sampling artifact

In addition to quantization and pack-sum mismatch, production BMS telemetry is subject to sequential or asynchronous sampling artifacts during current transients [3,34]. Cell voltages are not measured simultaneously; instead, they are scanned sequentially across the pack (Table 2). Under nearly steady current, this effect is negligible. During current transients, however, the asynchronous sampling introduces an artificial inter-cell voltage spread that does not reflect actual electrochemical imbalance. Table 2 summarizes three sources of measurement



**Fig. 3.** Pack-sum offset during early slow charging, showing static and current-dependent components.




Fig. 4. Current-dependent pack-sum offset after removing the static baseline.

**Table 2**  
Measurement uncertainty sources in production BMS cell-voltage data.

| Source                    | Magnitude                                     | Effect on Cell Voltage               | Condition          |
|---------------------------|-----------------------------------------------|--------------------------------------|--------------------|
| Cell-voltage quantization | 20 mV                                         | Staircase pattern, derivative spikes | All conditions     |
| Pack voltage quant        | 100 mV                                        | ~0.5 mV/cell (192 cells in series)   | All conditions     |
| Sequential scan           | up to $\approx 140$ mV (7 quantization steps) | Inter-cell asynchrony                | Current transients |
| Bus bar (static)          | 9.43 mV/cell                                  | Pack-sum offset                      | All conditions     |
| Bus bar (dynamic, 7 A)    | 1.04 mV/cell                                  | 5.2% of quantization                 | Slow charging      |
| Bus bar (dynamic, 100 A)  | 14.82 mV/cell                                 | 74.1% of quantization                | Fast charging      |
| Bus bar (dynamic, 200 A)  | 29.64 mV/cell                                 | 148% of quantization                 | Fast charging      |

uncertainty and their dependence on operating conditions. The first is coarse cell-voltage quantization on a 20 mV grid, addressed by the reconstruction in Section 3.1. The second is the approximately 9.4 mV/cell pack-sum offset from bus-bar impedance, addressed by the static correction in Section 2.3. The third is sequential per-cell sampling, which is negligible under near-constant current but comparable to the

20 mV step during transients. Of these, only sequential sampling cannot be corrected within the present slow-charging analysis window and therefore motivates restricting ECM identification and cell-level diagnostic analyses to constant-current segments.

In the present dataset, transient periods show inter-cell spreads that reach roughly seven quantization steps (approximately 140 mV in the illustrated driving-transient window), whereas under constant-current operation the baseline spread is close to zero. This indicates that, during dynamic operating conditions such as driving or fast-charging transients, apparent cell-to-cell deviations may be dominated by sampling skew rather than by genuine cell behavior. Fig. 5 contrasts these two regimes: under slow constant-current charging the inter-cell spread stays near zero (Fig. 5a), whereas during a driving transient the spread grows to roughly seven quantization steps (about 140 mV; Fig. 5b).

For this reason, the present study restricts ECM identification and cell-level diagnostic analysis to slow-charging constant-current segments, where the effects of sequential scanning are minimized and the remaining uncertainty sources are more structurally interpretable. Under these conditions, the dominant measurement limitations reduce to quantization and pack-sum offset, which can be addressed using the proposed pack-constrained reconstruction framework.




Fig. 5. Evidence of sequential per-cell sampling. (a) Slow constant-current charging window: inter-cell voltage spread stays near zero throughout, while the pack current is nearly constant. (b) Driving transient window: inter-cell spread grows to multiple quantization steps when the pack current swings rapidly.

### 3. Methodology

Section 3 presents a four-stage methodology designed to address the measurement limitations identified in Section 2. Stage 1 reconstructs physically admissible cell-voltage trajectories from quantized BMS measurements under the pack-sum constraint. Stage 2 establishes a baseline 1-RC ECM identification framework on the reconstructed data. Stage 3 tests whether additional RC branches are justified by the information content of BMS-quality measurements. Stage 4 resolves the OCV-parameter coupling problem through joint OCV-ECM optimization. The resulting pipeline provides the basis for the SOC estimation and fleet-level diagnostic analyses presented in Section 4.

#### 3.1. ADMM TV regularization

Because sequential scanning artifacts become negligible only under nearly constant current, the voltage reconstruction step was restricted to slow constant-current charging segments. Under these conditions, the current-dependent bus-bar contribution remains small relative to the 20 mV cell-voltage quantization step, so the pack-sum relation can be approximated using the measured pack voltage after static offset correction only. This setting allows the reconstruction problem to focus on the dominant uncertainty source, namely coarse cell-voltage quantization, while preserving pack-level physical consistency.

The reconstruction goal is to redistribute the quantization residual within each 20 mV grid step across the pack, so that the reconstructed cell-voltage trajectories are smooth in time and their sum exactly matches the static-offset-corrected pack voltage. Let  $v_i(t)$  denote the reconstructed voltage trajectory of cell  $i$  and let  $q_i(t)$  denote the corresponding quantized measurement. The reconstruction was formulated as a convex optimization problem that balances data fidelity and temporal regularity while enforcing pack-level consistency [5]. The box constraint enforces the floor-quantization interval  $[q_i, q_i + \delta]$  with  $\delta = 20$  mV (the observed cell-voltage quantization step), so that each reconstructed value remains within the admissible quantization interval of the original measurement — that is, the raw measurement  $q_i$  is treated as the floor of the true cell voltage on the 20 mV grid (paper code default: floor\_quant = True).

The problem was solved using ADMM, which separates temporal smoothing, quantization-consistent clipping, and pack-level equality enforcement into simple subproblems (Fig. 6). For each session, let  $v \in \mathbb{R}^{T \times N}$  collect the reconstructed cell voltages, let  $q$  collect the quantized measurements, and let  $D$  denote the first-order temporal-difference operator, so that  $Dv$  contains the  $(T-1) \times N$  adjacent differences. In the updates below,  $W$  collects the per-sample data-fidelity weights  $w_{t,i}$ ,  $z$  is the auxiliary variable introduced for  $Dv$ ,  $u_{TV}$  is its scaled dual variable,  $\rho_{TV}$  is the ADMM penalty parameter, and  $\mathcal{S}$  denotes the soft-thresholding operator. The optimization problem solved at each session is:

$$\min_v \sum_{t,i} w_{t,i} (v_{t,i} - q_{t,i})^2 + \lambda_{TV} \|Dv\|_1$$

$$\text{subject to: } \mathbf{1}^\top v_t = V_{\text{pack}}(t) - \beta_0 \quad (\forall t), \quad q_{t,i} \leq v_{t,i} \leq q_{t,i} + \Delta_{\text{quant}}.$$

(i) v-update — quadratic data + TV proximal:

$$R = Dv - z + u_{TV}$$

$$v \leftarrow (W \odot q + \rho_{TV} (v - D^\top R)) \oslash (W + \rho_{TV})$$

(ii) box clip — element-wise projection onto the floor-quantization box:

$$v_{t,i} \leftarrow \text{clip}(v_{t,i}, q_{t,i}, q_{t,i} + \Delta_{\text{quant}})$$

(iii) affine pack-sum projection — exact equality enforcement:

$$v_{t,:} \leftarrow v_{t,:} - \frac{1}{N} (\mathbf{1}^\top v_{t,:} - (V_{\text{pack}}(t) - \beta_0))$$

(iv) z- and dual-update — L1 soft-thresholding for the TV term:

$$z \leftarrow \mathcal{S}_{\lambda_{TV}/\rho_{TV}}(Dv + u_{TV})$$

$$u_{TV} \leftarrow u_{TV} + (Dv - z)$$

These four steps are repeated until the primal residual  $r$  falls below the tolerance  $\epsilon_{\text{tol}} = 10^{-7}$  or the iteration cap (code parameter max\_iter = 300) is reached. Two implementation details clarify the algorithm. First, step (iii) enforces the pack-sum equality through a hard affine projection rather than an augmented-Lagrangian penalty. Second, in the adopted floor-quantization box mode, clipping after the affine projection can produce a small equality-constraint residual within one ADMM iteration. Repeated affine projection in subsequent iterations reduces this residual to floating-point precision within tens of iterations. Section 4.1 verifies this convergence by reporting a reconstruction-constraint mean absolute error (MAE) below  $10^{-13}$  V. The TV proximal step suppresses staircase artifacts without introducing unconstrained drift [6]; the box clip preserves consistency with the original floor-quantized measurement; the affine projection enforces the pack-sum constraint exactly; the soft-threshold + dual update completes the standard ADMM cycle for the  $\ell_1$ -norm TV penalty. After convergence, a final affine projection re-enforces the pack-sum equality so the returned iterate satisfies the equality exactly while remaining within the floor-quantization box up to machine precision.

$$\|Dv - z\|_2 \leq \epsilon_{\text{tol}} = 10^{-7} \quad \text{or} \quad \text{iter} \geq \text{max\_iter} = 300$$

At each time step, the pack-sum relation imposes a single equality across  $N$  cell voltages, leaving  $N-1$  coupled degrees of freedom. This does not increase the intrinsic precision of the sensor hardware; rather, it reduces the admissible redistribution granularity of the quantization residual under the physical pack constraint. For a 192-cell pack, this granularity is  $\delta/(N-1) \approx 0.10$  mV, which explains why the reconstructed trajectories can recover sub-quantization continuity while remaining consistent with the original measurements.

The penalty parameters were selected by grid search to balance grid escape and waveform stability. Here, the dequantization rate was defined as the fraction of reconstructed voltage samples that no longer lay exactly on the original 20 mV quantization grid. The selected configuration ( $\rho_{TV} = 0.5$ ,  $\lambda_{TV} = 0.005$ , floor-quantization box  $\Delta_{\text{quant}} = 20$  mV, code parameter max\_iter = 300,  $\epsilon_{\text{tol}} = 10^{-7}$ ) provided a stable operating point with high dequantization rate and consistent convergence across repeated sessions. The trade-off underlying this choice is that aggressive TV regularization (large  $\lambda_{TV}$ ) maximizes grid escape but introduces visible waveform smoothing artifact (quantified in Section 4.4), whereas conservative settings (small  $\lambda_{TV}$ ) preserve waveform fidelity at the cost of incomplete dequantization; the adopted operating point lies at the elbow of this trade-off — near-maximal grid escape with sub-millivolt smoothing artifact (Fig. 7).

The choice of an ADMM-based reconstruction in this study is a deliberate baseline choice rather than a claim that ADMM is the only or globally optimal reconstruction method. It is motivated by three pragmatic engineering considerations relevant to production-BMS deployment: (i) determinism and closed-form convergence (Boyd et al. [5]), which give predictable, reproducible behavior across heterogeneous fleets; (ii) interpretability of the pack-sum equality projection as a direct hard physical constraint rather than a learned soft prior, which is essential when the reconstructed voltage feeds downstream physically interpretable ECM identification; and (iii) zero training-data requirement, which makes the framework immediately applicable to any new vehicle model without per-vehicle or per-fleet fine-tuning. More sophisticated alternatives include deep image prior-style learned regularizers,




Fig. 6. ADMM-based cell-voltage reconstruction. (a) Raw cell voltages showing the 20 mV staircase quantization. (b) ADMM-reconstructed cell voltages with constraint-consistent redistribution of the quantization residual to scales finer than the 20 mV step. (c) Pack-sum constraint residual before vs. after reconstruction.

$$\Delta_{\text{redist}} = \frac{20 \text{ mV}}{N} = \frac{20}{192} = 0.104 \text{ mV/cell} \quad (N = 192 \text{ for KIA EV6 LR})$$

diffusion-based reconstruction, and per-vehicle adaptive penalty schedules. These methods may further improve raw RMSE, but they can introduce black-box behavior and require retraining for each deployment. The central contribution of this paper is the pack-constrained reconstruction objective and its downstream OCV-ECM integration, rather than the ADMM implementation of the inner reconstruction step. The framework therefore remains applicable if these alternatives replace the inner solver. Investigating whether learned alternatives reduce the fleet-cross-evaluation variance discussed in Section 5 is therefore identified as a productive direction for follow-up work.

### 3.2. Equivalent Circuit Model (1-RC Thevenin)

A first-order Thevenin equivalent circuit model was adopted as the baseline identification model because it provides the simplest physically interpretable structure capable of separating ohmic and first-order polarization effects [1,21]. This parsimonious formulation is particularly appropriate for production BMS data, where 20 mV cell-voltage quantization and 3–4 s sampling intervals may not reliably support identification of multiple distinct dynamic time constants. Because the

baseline and joint-identification fitting windows are restricted to slow constant-current charging, the excitation available for RC identification is limited to smooth OCV drift and residual current variability within the segment. The model-order check in Sections 3.3 and 4.3 tests whether this ceiling persists when fast-charging telemetry is used as an identifiability probe with a transferred slow-charge OCV reference; the same 1-RC ceiling is observed under both regimes. The estimated 1-RC parameters should therefore be interpreted as the highest lumped RC approximation practically supported by BMS-quality telemetry (20 mV quantization, 3–4 s sampling), rather than as a fully excited dynamic identification.

With positive current defined for charging, the terminal-voltage model of each cell is written as

$$V(t) = \text{OCV}(\text{SoC}(t)) + I(t)R_0 + V_{\text{RC}}(t)$$

with the RC-state dynamics given by

$$\frac{dV_{\text{RC}}}{dt} = -\frac{V_{\text{RC}}}{R_1 C_1} + \frac{I(t)}{C_1}$$




Fig. 7. Sensitivity of dequantization performance to the ADMM parameters. (a) Fraction of samples departing from the original 20 mV grid. (b) Mean per-sample correction magnitude. The star marks the operating point adopted in this study ( $\rho_{TV} = 0.5$ ,  $\lambda_{TV} = 0.005$ ).

where OCV(SOC) denotes the open-circuit voltage as a function of state of charge,  $R_0$  is the ohmic resistance, and  $R_1$  and  $C_1$  characterize the first-order polarization dynamics.

The OCV–SOC relation was represented using a monotone piecewise cubic Hermite (PCHIP) interpolation of binned voltage–SOC pairs [35]. This representation preserves monotonicity while avoiding oscillatory artifacts that may arise in higher-order interpolation schemes. Because OCV source choice is itself one of the central variables examined in this study, the specific OCV construction method used in each experiment is described separately in Sections 3.3 and 3.4.

The RC state was propagated using the closed-form exponential update, which is more numerically stable than a naive forward discretization at the sampling interval of the present BMS data. To improve numerical conditioning and enforce positivity, model parameters were estimated in log-space, with the optimization vector  $\theta = [\log R_0, \log R_1, \log C_1]$  subject to box bounds  $R_0, R_1 \in [10^{-5}, 10^{-1}] \Omega$  and  $C_1 \in [10, 10^5] \text{F}$ . This constrained formulation is consistent with the broader need to control ECM parameter identifiability and non-physical solutions in real-world telemetry settings [22]. Optimization was performed with `scipy.optimize.least_squares` using the trust-region-reflective (TRF) algorithm and the soft-L1 robust loss  $\rho(r^2) = 2 \left( \sqrt{1 + r^2/f_{scale}^2} - 1 \right)$ , where  $f_{scale} = 0.01 \text{V}$  calibrates the transition between quadratic and linear regimes. The maximum number of function evaluations was capped at 500 (max\_nfev) for each fitting run.

The consequences of restricting identification to slow constant-current excitation were quantified with a local Fisher-information analysis of the 1-RC voltage model, following the identifiability perspective of [31]. Once the RC state has settled under steady constant current ( $t \gg \tau_1$ ), the polarization voltage becomes the constant  $I \cdot R_1$ ;  $R_0$  and  $R_1$  then affect the model only through their sum, the sensitivity to  $\tau_1 = R_1 C_1$  vanishes, and the information matrix restricted to such settled samples is structurally singular. If the OCV curve is also estimated freely, it absorbs the remaining constant offset into its voltage level. Settled constant-current samples — samples containing no entry or interruption transient — therefore cannot resolve the individual RC parameters; this is exactly the degeneracy that drives the charging-derived-OCV baseline to a non-physical near-zero  $R_0$  at practically unchanged fitting error (Sections 4.5 and 4.6). The realized fitting windows, however, are not perfectly steady. For the representative EV6 fitting window (10,976 samples, 669 min, mean current 8.6 A), three departures supply residual excitation: the window-entry transient (the forward simulation starts the RC state from zero, and the retained window opens

with a current adjustment of approximately 7.0 to 8.9 A; the physical charge-onset step itself lies just before the window), a gradual end-of-charge taper (approximately 9.2 to 7.9 A), and a brief mid-session charging interruption contributing two current steps of approximately 9 A. The Cramér–Rao bound was evaluated at the fitted parameters in  $\theta$ -space, using the joint-fit residual scale of 4.33 mV as the noise level. The evaluation starts the RC state from zero, as in the production fit; unlike the production fit, however, it treats the OCV curve as known — an optimistic simplification, because the joint fit estimates the OCV curve. The analysis yields lower bounds on the relative standard deviations of approximately 21% for  $R_0$ , 32% for  $R_1$ , and 65% for  $\tau_1$ ; re-evaluated at the fleet-median  $R_0$  of 0.895 m $\Omega$  (Table 7) with the anchor vehicle's  $R_1$  and  $\tau_1$ , the  $R_0$  bound becomes 27%, while the  $R_1$  and  $\tau_1$  bounds are unaffected by the  $R_0$  value. In addition, the 0.2% of samples surrounding the window-entry and interruption current steps carry 70% of the diagonal Fisher information on  $C_1$  (49% from the window entry, 21% from the interruption). Because the bound conditions on a known OCV curve and additionally assumes white residuals whereas the dominant quantization residual is colored, these lower bounds should be read as optimistic.

Three design consequences follow. First, the RC-pair parameters ( $R_1$ ,  $C_1$ ) are interpreted as a weakly identified lumped polarization correction, consistent with the model-order ceiling in Sections 3.3 and 4.3. No diagnostic conclusion in this paper depends on the individual precision of  $R_1$  or  $C_1$ ; instead, the downstream SOC evaluation in Section 4.7 tests the recovered parameter set as a whole. Second, the physical plausibility of the ohmic resistance is not asserted from the constant-current fit alone: Section 4.6 cross-checks the jointly estimated  $R_0$  against an independent current-step estimator built from  $|\Delta I| > 5 \text{A}$  transients, and the joint-versus-step differences observed there are of the same order as the bounds above. Third, the OCV-offset degeneracy is deliberately counteracted by the rest-anchoring term of the joint formulation (Section 3.4), which anchors the OCV level to zero-current rest voltages instead of leaving it free to absorb polarization. Dynamic excitation would sharpen RC identifiability considerably: applied to the full telemetry segment that surrounds this fitting window — a segment that includes driving transients — the same analysis lowers the  $\tau_1$  bound to approximately 5.1%. Such windows are nevertheless excluded from identification because of the sequential-sampling artifact described in Section 2.4; exploiting this excitation would require synchronized, higher-rate logging, as discussed in Section 5.

### 3.3. N-RC model order selection

The model-order analysis compares N-RC networks — Thevenin equivalent circuits with one to three parallel RC branches — using AIC and BIC on both slow and fast charging windows. To determine whether additional dynamic branches beyond the baseline 1-RC ECM introduced in Section 3.2 are justified by field BMS data, the model was extended to 2-RC and 3-RC formulations. The purpose of this analysis was not to maximize in-sample flexibility, but to test whether the information content of 20 mV-quantized, 3–4 s-sampled telemetry is sufficient to support physically identifiable additional time constants [26, 28].

To ensure a fair comparison across model orders, higher-order models were initialized hierarchically from the lower-order solution. This strategy reduces initialization sensitivity and allows the model-order comparison to reflect genuine information gain rather than differences caused by unstable optimization starting points.

For fast-charging evaluation, the OCV curve was not reconstructed from the same fast-charging session, because doing so would reintroduce the circular fitting artifact that this analysis seeks to isolate. Instead, an OCV curve obtained from a slow-charging session of the same vehicle was transferred to the fast-charging session. This design ensures that the comparison among 1-RC, 2-RC, and 3-RC models reflects dynamic-model adequacy rather than bias introduced by the OCV reference itself.

In addition, a current-dependent bus-bar correction was applied to fast-charging data before ECM fitting, because the per-cell dynamic offset becomes comparable to or larger than the 20 mV quantization step under high current. By contrast, this correction was not required for slow charging, where the estimated bus-bar contribution remained small relative to the quantization scale.

Model order was evaluated using AIC and BIC, and the primary finding is that the data are practically unidentifiable beyond a single RC branch under the present BMS-quality telemetry (20 mV cell-voltage quantization and 3–4 s sampling interval): additional time constants cannot be resolved above its noise and quantization floor. Across all tested orders the fit RMSE was indistinguishable at the reported precision, so the likelihood term is essentially flat across model orders and the unanimous BIC/AIC selection of the 1-RC formulation is decided by their parameter-count penalties. Examination of the higher-order solutions confirmed that the second and third RC branches either collapsed to their lower bounds or redundantly partitioned the dynamics of the first branch; the apparent residual structure remaining in BMS-quality data is below the threshold required to identify additional time constants stably. Section 4.3 reports the per-condition numerical evidence.

This framework allows the study to distinguish improvements caused by model complexity from those caused by OCV source quality, which is essential for interpreting the comparative results presented in Section 4.

### 3.4. Joint OCV-ECM optimization

Given that Section 3.3 identifies 1-RC as the highest identifiable structure, a major remaining source of ECM error is the OCV reference used during fitting. As shown by the OCV-source comparison in Section 4.5, fitting accuracy and parameter physicality exhibit a persistent trade-off. OCV curves constructed directly from charging data produced low voltage residuals but drove  $R_0$  toward a non-physical lower bound [15], whereas rest-derived OCV references restored physically meaningful resistance values [29,30] at the cost of substantially higher fitting error due to noise and incomplete SOC coverage. This trade-off motivated a joint estimation framework in which the OCV curve and ECM parameters are optimized simultaneously rather than sequentially.

In the proposed framework, the OCV curve is parameterized by 50 free knot voltages distributed uniformly over the SOC range. This representation

provides sufficient flexibility to depart from the charging-derived initial curve while remaining simple enough to regularize effectively. Piecewise-linear interpolation was used between knots to avoid oscillatory behavior associated with higher-order spline parameterizations and to maintain stable gradient propagation during joint optimization.

A limited knot-count sensitivity test was performed on a representative paper-fleet vehicle (V004AL0000, KIA EV6 LR) to verify that the 50-knot choice is well-positioned. With identical training settings, the joint OCV-ECM voltage-fit RMSE decreased from 5.16 mV at 25 knots to 4.70 mV at 50 knots and 3.45 mV at 100 knots. The jointly estimated ohmic resistance  $R_0$ , which is central to the physical-plausibility claim in Section 4.6, was already within 5% of its 100-knot value at 50 knots. The corresponding  $R_0$  values were 0.41/0.45/0.47 mΩ for 25/50/100 knots respectively. The 50-knot setting adopted throughout this paper therefore prioritizes  $R_0$  physicality saturation and per-vehicle optimization cost; the remaining voltage-fit improvement available at higher knot counts is acknowledged as a productive direction for follow-up work where computational cost is not the dominant constraint.

The knot values were initialized from the charging-derived OCV curve so that optimization started from a low-error operating point. This initialization, however, does not determine the final solution, because the learned OCV curve is subsequently constrained by rest anchoring, smoothness, monotonicity, and boundary penalties.

The ECM dynamics were embedded in a differentiable forward simulation so that the OCV curve and circuit parameters could be updated within a single coupled optimization. This formulation avoids the inconsistency of separately constructing an OCV curve and then fitting ECM parameters afterward, and it allows the loss to propagate through the full time series in a physically consistent manner.

The joint estimation frames the OCV curve and the ECM parameters as one differentiable problem: the OCV curve provides the equilibrium reference, and the ECM parameters explain the deviation of the measured voltage from that reference under the observed current profile. The total objective combines voltage fidelity with physically motivated OCV regularization:

$$L_{\text{total}} = L_{\text{ecm}} + \omega_{\text{rest}} L_{\text{rest}} + \omega_{\text{smooth}} L_{\text{smooth}} + \omega_{\text{mono}} L_{\text{mono}} + \omega_{\text{boundary}} L_{\text{boundary}}$$

Here,  $L_{\text{ecm}}$  enforces agreement between simulated and measured terminal voltage,  $L_{\text{rest}}$  anchors the OCV curve to rest-period voltage observations,  $L_{\text{smooth}}$  suppresses excessive curvature,  $L_{\text{mono}}$  discourages non-monotonic behavior, and  $L_{\text{boundary}}$  stabilizes the endpoint behavior of the learned OCV curve. Together, these terms prevent the OCV estimate from simply absorbing the full charging overpotential while still allowing it to adapt away from the biased charging-derived initial guess.

The rest-anchor weight was selected through Pareto analysis of the trade-off between voltage fit and resistance physicality. A value of  $\omega_{\text{rest}} = 0.5$  was adopted because it prevented collapse of  $R_0$  while maintaining low prediction error, placing the solution near the elbow of the RMSE–physicality trade-off.

The joint objective was minimized with the Adam optimizer (optax.adam) for 3000 full-batch optimization steps (no mini-batching). The initial learning rate was  $10^{-3}$  with a staircase exponential decay factor of 0.5 every 1000 steps; the Adam moment coefficients ( $\beta_1 = 0.9, \beta_2 = 0.999$ ) and the Adam numerical-stability constant ( $\epsilon = 10^{-8}$ ) followed optax defaults. The optimization vector concatenated the 50 OCV knot voltages with the log-space ECM parameters ( $\log R_0, \log R_1, \log C_1$ ). The OCV knots were initialized from the charging-derived OCV curve. The ECM parameters were initialized from a baseline 1-RC least-squares fit. Gradients were obtained by automatic differentiation through a JAX just-in-time-compiled forward simulation, and the full 3000-step schedule was completed for each vehicle without early stopping.

Although the OCV parameterization is flexible, overfitting risk was controlled in three ways. First, the learned OCV curve was regularized by smoothness, monotonicity, boundary, and rest-anchor constraints. Second, ECM parameters remained positivity-constrained throughout optimization. Third, model quality was assessed using cross-session prediction rather than in-sample fitting alone.

For fleet-level evaluation, the full preprocessing-identification pipeline was applied consistently to each of the 26 vehicles, namely ADMM-based voltage reconstruction, OCV construction, joint OCV–ECM fitting, and cross-session evaluation. This end-to-end design is important because the benefit of joint optimization depends directly on the quality of the reconstructed voltage trajectory used as its input.

## 4. Results and discussion

### 4.1. ADMM preprocessing results

The proposed ADMM-based reconstruction consistently improved the effective continuity of cell-voltage trajectories across all tested vehicle models while preserving strict pack-level consistency [5,6]. For the primary test case, KIA EV6 with 192 series-connected cells, 96.5% of reconstructed voltage samples no longer remained on the original 20 mV quantization lattice. This result indicates that the pack-sum constraint provides sufficient coupling among cell voltages to redistribute the quantization residual at scales finer than the native 20 mV cell-voltage step — without implying any increase in the intrinsic precision of the sensor hardware. The recovered fine structure should therefore be interpreted as a constraint-consistent redistribution of quantization residuals rather than a true sub-quantization voltage measurement.

The achievable redistribution scale followed the expected dependence on cell count. Because the pack-sum equality introduces one linear constraint across  $N$  cells at each time step, the quantization residual can be redistributed across  $N - 1$  coupled degrees of freedom. For a 192-cell pack, this corresponds to a characteristic sub-quantization scale of approximately  $20/N = 0.104$  mV. The observed minimum step sizes in the reconstructed trajectories were consistent with this scale across the tested vehicles, with smaller effective steps observed for larger packs, in qualitative agreement with the predicted  $N^{-1}$  scaling.

Importantly, the pack-sum equality was satisfied at numerical precision after the affine projection step of ADMM. The resulting constraint error remained below  $10^{-13}$  V throughout the analyzed sessions, confirming exact algebraic consistency between the reconstructed cell voltages and the corrected pack voltage within the optimization procedure. This value should be interpreted as a property of the projection step rather than as a claim of sensor-level measurement precision. The practical significance is that the reconstruction does not sacrifice pack-level physical consistency in exchange for smoother cell-level trajectories.

The mean correction magnitude remained within the expected quantization interval, approximately 8.8–9.1 mV per cell across the tested models, indicating that the reconstruction stayed within the admissible range implied by the original measurements. Convergence was also stable across repeated sessions. Table 3 summarizes these per-model reconstruction metrics: across the five representative pack architectures reported in Table 3, grid-escape rates ranged from 94.1% to 96.7%, effective minimum step from 0.100 mV to 0.208 mV, and pack-sum constraint MAE remained below  $10^{-13}$  V, confirming that the pack-constrained reconstruction generalizes across pack architectures without per-model tuning.

These results should be interpreted within the operating regime considered in this study. The reported resolution enhancement applies to slow-charging constant-current segments, where sequential scan artifacts are small and the dynamic bus-bar contribution remains modest relative to the 20 mV quantization step. Under highly transient conditions such as fast-charging ramps or driving cycles, asynchronous cell scanning can introduce apparent inter-cell spreads comparable to the

**Table 3**

ADMM preprocessing results across representative vehicle models. Where two values are listed (e.g., 96.7 / 96.3), the first value is for the slow constant-current charging segment used in this study, and the second value is from a held-out follow-up session.

| Model            | Cells | Grid Escape (%) | Min Step (mV) | Constraint MAE (V)    | Mean Corr. (mV) |
|------------------|-------|-----------------|---------------|-----------------------|-----------------|
| KIA EV6 LR       | 192   | 96.7 / 96.3     | 0.100 / 0.143 | $\sim 9.5\text{e-}14$ | 8.8 / 8.9       |
| HMC IONIQ5 LR    | 180   | 96.2            | 0.111         | $\sim 1\text{e-}13$   | 8.9             |
| HMC IONIQ5 LR'22 | 192   | 96.4            | 0.104         | $\sim 1\text{e-}13$   | 8.8             |
| KIA EV9          | 152   | 95.8            | 0.132         | $\sim 1\text{e-}13$   | 9.0             |
| AUDI Q4 e-tron   | 96    | 94.1            | 0.208         | $\sim 1\text{e-}13$   | 9.1             |


quantization scale, and the present formulation does not explicitly correct for that effect. Within the constant-current regime, however, the ADMM framework provides a physically consistent means of recovering smoother and more informative cell-voltage trajectories from coarse production BMS data. The cross-model preprocessing results across all representative vehicle models are summarized in Table 3.

### 4.2. ECM fitting: raw vs enhanced


The benefit of the reconstructed voltage trajectories was evaluated by fitting the 1-RC ECM [1,21] to both raw quantized data and ADMM-enhanced data. Across all representative vehicle models, preprocessing reduced the terminal-voltage fitting error. The single-session RMSE improvement ranged from 11.2% for the KIA EV6 to 49.4% for the KIA EV9. We note explicitly that the 11.2% EV6 improvement corresponds to an absolute reduction of only 0.66 mV ( $5.87 \rightarrow 5.21$  mV) and is therefore within the session-to-session variability reported in Section 4.4 (raw RMSE standard deviation 0.72 mV across ten sessions); the EV6 single-session result is presented here as a representative case rather than as evidence of a meaningful single-vehicle improvement. The statistical support for attributing the ADMM-preprocessing benefit to the algorithm itself, rather than to session-level noise, comes from the multi-vehicle same-session paired test described in the following paragraphs; the larger gains observed in IONIQ5 LR (46.5%), IONIQ5 LR 2022 (46.8%), and EV9 (49.4%) sit well above the session noise and are interpreted as the primary single-session evidence (Fig. 8). Fleet-level results are summarized in two complementary views: Fig. 8 reports the overall RMSE distribution and the per-vehicle RMSE-reduction histogram, while Fig. 9 plots the raw-versus-enhanced paired scatter and highlights the best, median, and worst per-vehicle outcomes; Table 4 lists the per-model breakdown.

The magnitude of improvement was not identical across models. The larger gains observed in the IONIQ5 and EV9 cases may reflect stronger sensitivity of ECM fitting to coarse voltage staircasing under those particular pack configurations and operating sessions. However, because the present comparison was not designed to isolate architecture-specific effects from session-specific excitation characteristics, these cross-model differences should be interpreted cautiously. The key result is the consistency of the improvement trend itself: in every tested case, the reconstructed voltage produced lower fitting error than the raw quantized signal.

Although the coefficient of determination remained high for all models, the more meaningful indicator in this setting was the reduction in RMSE and the persistence of that reduction under held-out evaluation. Cross-session evaluation confirmed that the improvement was not confined to the training session. For the EV6 example, the held-out RMSE decreased from 10.75 mV to 9.50 mV, corresponding to an 11.6% reduction. This indicates that the enhanced trajectories improve not only in-sample agreement but also the transferability of the identified ECM to unseen data from the same vehicle.



**Fig. 8.** Fleet-level ECM fitting comparison from the same-session paired test across 135 vehicles. (a) Boxplots of voltage-fit RMSE before and after ADMM reconstruction. (b) Histogram of per-vehicle RMSE reduction; the distribution is shifted toward positive values, with ADMM improving fit in the majority of vehicles.



**Fig. 9.** Fleet-level voltage-fit RMSE comparison across the 135-vehicle paired test. Each point represents one vehicle's raw versus ADMM-enhanced 1-RC ECM RMSE; the diagonal indicates no improvement, and markers highlight the best, median, and worst per-vehicle outcomes among the ECM-fit-valid subset ( $n = 135$  paired;  $n = 134$  used for the best/median/worst marker selection after excluding one vehicle with insufficient samples for stable ECM fit).

**Table 4**  
1-RC ECM fitting results before and after ADMM-based voltage reconstruction.

| Model            | Cells | Raw RMSE (mV) | Enhanced RMSE (mV) | Improvement (%) |
|------------------|-------|---------------|--------------------|-----------------|
| KIA EV6 LR       | 192   | 5.87          | 5.21               | 11.2            |
| HMC IONIQ5 LR    | 180   | 5.20          | 2.78               | 46.5            |
| HMC IONIQ5 LR'22 | 192   | 5.35          | 2.84               | 46.8            |
| KIA EV9          | 152   | 5.02          | 2.54               | 49.4            |

To establish that the Table 4 single-session improvements exceed the session-to-session variability discussed in Section 4.4, we additionally performed a same-session paired test across a separate multi-vehicle cohort of 135 vehicles drawn from a 2411-vehicle proprietary Betterwhy telemetry corpus (56 distinct Hyundai/Kia/imported vehicle models in total), restricted to vehicles with at least one year of slow-charging coverage and at least two valid slow-charging sessions. For each vehicle, one best-SOC-range slow constant-current session was processed through the identical ADMM and 1-RC pipeline and the same-session  $enh\_self$  statistic (Section 4.4) was extracted. Across the four pack architectures (KIA EV6 192S  $n = 67$ , KIA EV9 152S  $n = 13$ , Hyundai IONIQ5 LONGRANGE 180S  $n = 34$ , IONIQ5 LONGRANGE 2022 192S  $n = 21$ ; the NS suffix denotes  $N$  series-connected cells), the mean RMSE reduction was 38.7%, 42.2%, 30.2%, and 43.7% respec-

tively (mean absolute reduction 1.6–2.2 mV per vehicle), with paired  $t$ -test  $p < 10^{-5}$  and Wilcoxon  $p < 10^{-3}$  in all four cohorts. The proportion of vehicles for which the enhanced trajectory yielded lower RMSE than the raw quantized trace was 93%, 85%, 82%, and 90% respectively. Because the per-vehicle effect size (1.6–2.2 mV) clearly exceeds the raw RMSE standard deviation across sessions reported in Section 4.4 for the EV6 cohort (0.72 mV), the ADMM-preprocessing benefit is statistically supported at the multi-vehicle level across all four Hyundai/Kia architectures covered by this paired test. We caution that this statistical support is for the per-vehicle ADMM-vs-raw comparison aggregated over the paired-test cohort; it does not by itself elevate the Table 4 single-vehicle 11.2% EV6 improvement above the session-level noise while that single-session result remains a representative example rather than a statistically isolable claim (Fig. 9; Table 4).

We performed a cross-session held-out paired test using the framework of Section 4.4. This test addresses the possibility that the same-session enhanced-self (enh\_self) statistic is inflated by a self-evaluation smoothing advantage, estimated at approximately 1.05 mV for the representative EV6 case (Section 4.4). For each vehicle with two or more valid slow-charging sessions, ECM parameters were fitted separately to the raw and ADMM-enhanced data from session 1. The resulting 1-RC models were then evaluated on session 2: the raw-fit model predicted the raw session-2 trace, and the enhanced-fit model predicted the enhanced session-2 trace. Of the 135-vehicle cohort, 21 vehicles drawn from the same four pack architectures reported above (KIA EV6 192S  $n = 3$ , KIA EV9 152S  $n = 3$ , Hyundai IONIQ5 LONGRANGE 180S  $n = 5$ , IONIQ5 LONGRANGE 2022 192S  $n = 10$ ) provided two or more valid slow-charging sessions. Across these 21 vehicles, the median cross-session prediction RMSE decreased from 8.68 mV (raw-fit ECM on session 2) to 7.03 mV (enhanced-fit ECM on session 2; mean 8.58 mV  $\rightarrow$  7.04 mV), a mean absolute reduction of 1.54 mV and a mean fractional reduction of 19.8% (EV6 24.0%, EV9 9.6%, IONIQ5 LR 180S 12.7%, IONIQ5 LR 2022 25.1%). ADMM enhancement reduced cross-session prediction RMSE in every one of the 21 vehicles (paired  $t$ -test  $p = 3.3 \times 10^{-10}$ ; Wilcoxon  $p = 9.5 \times 10^{-7}$ ); 90% of vehicles exceeded the 0.72 mV session-to-session noise floor by their effect size alone. Because session 1 and session 2 are separate non-overlapping charging events, the evaluation reference for each side is the corresponding session-2 trace fitted separately, which prevents the same-session smoothing-advantage critique from inflating the cross-session result and therefore provides the cleanest available evidence that ADMM preprocessing produces a genuine, generalizable improvement in 1-RC ECM identification across heterogeneous Hyundai/Kia EV pack architectures.

A careful interpretation is required because direct comparison of enhanced data against itself can overstate the apparent gain. When the same smoothed trajectory is used for both fitting and evaluation, part of the RMSE reduction can arise from suppression of high-frequency quantization artifacts rather than from improved physical identification alone. To control for this effect, the present study relied on cross-evaluation protocols and multi-session statistics, which are discussed further in Section 4.4. Under these controls, the improvement remained substantial, supporting the conclusion that ADMM preprocessing provides a more reliable voltage basis for ECM parameter estimation.

#### 4.3. Model Order Selection (N-RC)

To examine whether higher-order ECM structures are justified by the information content of production BMS data, 1-RC, 2-RC, and 3-RC formulations were compared under slow charging, with 1-RC and 2-RC additionally compared under fast charging. The purpose of this analysis was not to maximize in-sample flexibility, but to test whether additional dynamic branches could be identified in a stable and meaningful manner from 20 mV-quantized, 3–4 s-sampled field telemetry.

The principal finding of Section 4.3 is that, for the present BMS-quality telemetry (20 mV cell-voltage quantization, 3–4 s sampling interval), the data are practically unidentifiable beyond a single RC branch — additional polarization time constants cannot be reliably separated from the noise and quantization floor of this telemetry regime. Under slow charging, the 1-RC, 2-RC, and 3-RC models all produced an RMSE of 4.81 mV at the reported precision, which by itself indicates that the higher-order branches contribute no detectable explanatory power beyond the first RC pair. To rule out an optimizer-collapse artifact, higher-order fits were initialized hierarchically (the 2-RC/3-RC solutions started from the lower-order solution augmented with an additional RC pair at a distinct initial time constant), and the additional RC pairs were given the same parameter bounds; despite this, the second and third RC pairs either converged at their lower bounds or redundantly partitioned the dynamics already captured by the first branch. With the likelihood essentially identical across model orders, BIC and AIC unanimously select 1-RC through their parameter-count penalties on the unused additional parameters. Examination of the higher-order solutions confirmed that the apparent residual structure remaining in BMS-quality data is below the threshold required to identify additional time constants stably; in other words, model-order complexity beyond 1-RC is not refuted but simply not supported by the information content of 20 mV-quantized, 3–4 s-sampled real-world telemetry. The implication is that the gains observed elsewhere in this paper must come from improvements in the voltage-reconstruction and OCV-source quality, not from additional dynamic complexity in the ECM structure.

To verify that this identifiability ceiling reflects an information-content limit of BMS-quality telemetry rather than an optimizer-collapse artifact in our implementation, we performed a synthetic 2-RC sanity check. A charging session was generated using known ground-truth parameters ( $R_0 = 0.5 \text{ m}\Omega$ ;  $R_1 = 0.5 \text{ m}\Omega$  with  $\tau_1 = 10 \text{ s}$ ;  $R_2 = 0.3 \text{ m}\Omega$  with  $\tau_2 = 100 \text{ s}$ ; constant 5 A charging; SOC 30–70%) and corrupted by 20 mV cell-voltage quantization plus 5 mV Gaussian noise to match production-BMS measurement quality. Fitting the same 1-RC, 2-RC, and 3-RC formulations used in Table 5 to this synthetic data with the paper's 3.65 s inter-sample interval produced indistinguishable RMSE across all three model orders (RMSE  $\approx$  5.8 mV, equal to the 20 mV quantization noise floor) and  $R_j$  estimates that collapsed to the optimizer lower bound — exactly the unidentifiability pattern reported for the real EV6 data in Table 5. The same pattern persisted when the sampling rate was increased to 0.5 s, demonstrating that under the present BMS-grade quantization the dominant identifiability barrier is the 20 mV cell-voltage quantization itself, not the inter-sample interval. The fitting procedure is therefore correctly implemented: the sanity check reproduces the practical unidentifiability of higher-order RC branches under production-BMS telemetry, indicating that the result reflects the limited information content of the data rather than the estimation procedure.

The same pattern was observed under fast charging when the OCV reference was transferred from a slow-charging session of the same vehicle. Again, adding extra RC pairs did not reduce RMSE relative to the 1-RC model; given the essentially identical likelihood, BIC/AIC continued to favor the 1-RC formulation through their parameter-count penal-

**Table 5**

AIC/BIC comparison of ECM model orders for the representative vehicle (slow: 1-RC, 2-RC, 3-RC; fast: 1-RC, 2-RC).

| Condition | Model | RMSE (mV) | $R^2$  | AIC      | BIC      | BIC Best? |
|-----------|-------|-----------|--------|----------|----------|-----------|
| Slow      | 1-RC  | 4.81      | 0.9994 | -118,322 | -118,300 | Yes       |
| Slow      | 2-RC  | 4.81      | 0.9994 | -118,317 | -118,281 |           |
| Slow      | 3-RC  | 4.81      | 0.9994 | -118,313 | -118,262 |           |
| Fast 1    | 1-RC  | 11.13     | –      | -10,200  | -10,198  | Yes       |
| Fast 1    | 2-RC  | 11.13     | –      | -10,196  | -10,182  |           |
| Fast 2    | 1-RC  | 13.33     | –      | -25,750  | -25,744  | Yes       |
| Fast 2    | 2-RC  | 13.33     | –      | -25,746  | -25,729  |           |

ties. Taken together, these results indicate that, for the telemetry quality and sampling interval considered here, the data support at most a single robustly identifiable polarization time constant within the adopted ECM structure — higher-order RC pairs are not refuted but simply not supported by the information content of 20 mV-quantized, 3–4 s-sampled real-world telemetry.

This finding is important because it separates two possible explanations for improved fitting performance: increased model order and improved physical consistency of the inputs. The present comparison suggests that, within the analyzed BMS-quality dataset, increasing ECM complexity beyond 1-RC does not yield a commensurate benefit. Accordingly, the subsequent analysis focuses on improving the quality of the voltage reconstruction and OCV reference rather than introducing additional RC branches whose parameters are not reliably supported by the data (Table 5).

#### 4.4. Robustness validation

Several additional checks were conducted to assess whether the reported gains were robust to evaluation protocol, session variability, and operating-condition effects. First, a three-way comparison using raw, enhanced-self, and enhanced-cross evaluation revealed that direct evaluation on the enhanced signal itself introduces a modest smoothing advantage. For the representative EV6 case, this effect was approximately 1.05 mV. Although small relative to the total gain observed in several vehicles, this confirms that fair assessment requires held-out or cross-evaluated comparisons rather than self-evaluation alone (Fig. 10).

Second, multi-session analysis showed that session selection can bias performance estimates if only the most favorable case is reported. Across ten slow-charging sessions, the raw-data RMSE for the EV6 was  $5.61 \pm 0.72$  mV, with a 95% bootstrap confidence interval of [5.18, 6.04] mV. By comparison, selecting only the best individual session yielded 4.81 mV, which underestimates the typical error by approximately 14.3%. This result motivated the use of per-vehicle median statistics in later fleet-level analysis, thereby reducing the influence of unusually favorable or unfavorable sessions.

Third, the available dataset allowed a limited temperature stratification analysis within the observed operating range. No cold-temperature sessions below 15 °C were available, so the comparison was restricted to nominal and moderately hot conditions. Within that range, the fitting error remained similar between the two groups, suggesting no strong temperature dependence in the present dataset. However, this should not be interpreted as evidence of general temperature invariance, since the most temperature-sensitive regime was not represented.

Finally, the bus-bar contribution was reexamined from the perspective of identifiability. For slow charging at approximately 7 A, the dy-

namic bus-bar term corresponded to only 1.04 mV per cell, which is small relative to the 20 mV quantization step. Under these conditions, static offset correction is adequate for pack-constrained reconstruction and ECM fitting. Under fast charging, in contrast, the dynamic contribution becomes comparable to or larger than the quantization scale, reaching 14.82 mV per cell at 100 A and 29.64 mV per cell at 200 A. This confirms that explicit current-dependent correction is necessary when analyzing high-current sessions. Overall, these checks support the conclusion that the improvements reported in Sections 4.1–4.3 are not solely artifacts of smoothing, session selection, or unmodeled pack-level offsets within the slow-charging regime considered for identification (Fig. 10).

#### 4.5. OCV source analysis

After establishing the benefit of voltage reconstruction and the sufficiency of the 1-RC structure, the next question concerns the OCV reference used during ECM identification. To isolate this effect, four OCV construction methods were compared for the same vehicle: charging-derived OCV from the slow constant-current trace, ohmic-drop (IR)-corrected charging OCV, rest-derived OCV from the same vehicle, and fleet-rest OCV pooled across vehicles. These four references span a representative range of practical OCV construction approaches available from field data (Fig. 11).

The comparison revealed a clear trade-off between voltage-fitting accuracy and parameter physicality. Charging-derived OCV [15] produced the lowest fitting error, with a self-RMSE of 4.81 mV, but the corresponding ohmic resistance collapsed to a near-zero value of  $R_0 = 0.057$  m $\Omega$ . IR-correcting the charging curve produced almost no change, which is consistent with the small expected ohmic drop under slow charging. In contrast, the rest-derived OCV [29,30] produced a physically plausible  $R_0$  of 0.785 m $\Omega$ , but the fitting error increased substantially to a self-RMSE of 12.23 mV because the rest-voltage observations were noisy and sparsely distributed over SOC. The fleet-rest OCV slightly reduced error relative to same-vehicle rest OCV, but it again drove  $R_0$  toward the lower bound, indicating that pooled rest references do not reliably solve the coupling problem (Table 6).

These results support the interpretation that charging-derived OCV can absorb a substantial portion of the charging overpotential when it is constructed from the same session later used for ECM fitting. In that case, part of the voltage drop that should be attributed to  $R_0$  and dynamic polarization is embedded into the inferred OCV curve itself, leaving too little residual structure for physically meaningful estimation of  $R_0$ . The average offset between charging-derived and rest-derived OCV curves was 29.5 mV, which is consistent with substantial overpotential absorption into the charging-derived reference. At the fleet level, this





Fig. 10. Fair evaluation of ADMM reconstruction benefit over multiple sessions.



**Fig. 11.** Comparison of charging-, rest-, and joint-derived OCV curves. (a) Charging- and rest-derived OCV curves for a representative EV6. (b) Joint OCV curve, balancing smoothness and physical plausibility.

**Table 6**

Comparison of OCV source choices for ECM identification in the same vehicle.


| OCV Source              | Self-RMSE (mV) | Cross-val (mV) | $R_0$ (m $\Omega$ ) | Physical? |
|-------------------------|----------------|----------------|---------------------|-----------|
| Charging (raw)          | 4.81           | 5.56           | 0.057               | No        |
| Charging (IR-corrected) | 4.80           | 5.56           | 0.081               | No        |
| Rest (same vehicle)     | 12.23          | 12.20          | 0.785               | Yes       |
| Rest (fleet)            | 10.82          | 10.55          | 0.010               | No        |

behavior was not an isolated case: 24 of 26 vehicles yielded  $R_0$  values pinned at or near the estimation lower bound when charging-derived OCV was used.

The key implication is not simply that one OCV source is “better” than another, but that conventional sequential fitting creates a practical mismatch between two objectives. Charging-derived OCV tends to preserve voltage fit while degrading parameter physicality, whereas rest-derived OCV improves physical interpretability but at the cost of fitting quality. For ECM-based battery diagnosis, where both predictive accuracy and physically meaningful resistance parameters are desirable, neither extreme is fully satisfactory. This motivates the joint OCV-ECM framework evaluated in the following subsection (Fig. 12).

#### 4.6. Joint OCV-ECM fleet-level validation

The joint OCV-ECM framework was then evaluated at the fleet level using 26 KIA EV6 vehicles that provided at least two suitable slow-charging sessions. Among these, step-based resistance estimates were available for 23 vehicles, and cross-session validation was conducted on a filtered subset of 19 vehicles after excluding cases with inter-session maximum voltage differences exceeding 20 mV. This filtering was introduced to remove session-to-session inconsistencies that would confound comparison among OCV methods independently of the ECM framework itself. The  $R_0$  distribution and the cross-session validation outcomes are reported in Table 7 and Table 8 respectively (Fig. 13). Table 7 lists the fleet distribution of  $R_0$  across the 26 KIA EV6 vehicles: the joint OCV-ECM method restored a physically plausible fleet median of 0.895 m $\Omega$  (range 0.11–1.59 m $\Omega$ ), whereas the charging-derived OCV baseline drove the fleet median to a physically implausible near-zero value (0.015 m $\Omega$ ). Table 8 reports the corresponding cross-session validation on the 19-vehicle filtered subset: the joint OCV-ECM produced a lower median cross-session prediction RMSE than the charging-derived baseline (7.12 vs. 8.04 mV) and outperformed it in 74% of vehicles, indicating that the recovered  $R_0$  physicality improvement does not come at the cost of predictive fit.



**Fig. 12.** Effect of OCV source on fitting accuracy and resistance physicality. (a) RMSE and  $R_0$  comparison across four OCV construction methods. (b) RMSE– $R_0$  trade-off motivating joint OCV-ECM optimization.

**Table 7**  
Fleet distribution of  $R_0$  across 26 KIA EV6 vehicles.

| Method                     | Mean (m $\Omega$ ) | Std   | Median | Range     |
|----------------------------|--------------------|-------|--------|-----------|
| Step estimate ( $n = 23$ ) | 0.790              | 0.387 | 0.810  | 0.06–1.45 |
| Charging OCV               | 0.035              | 0.061 | 0.015  | 0.01–0.25 |
| Joint OCV                  | 0.886              | 0.389 | 0.895  | 0.11–1.59 |

**Table 8**  
Fleet cross-validation results for charging-derived and joint OCV methods.

| Metric                     | Mean (mV) | Std  | Median |
|----------------------------|-----------|------|--------|
| Joint Self-RMSE            | 5.90      | 4.48 | 4.99   |
| Cross-session Charging OCV | 8.40      | 3.08 | 8.04   |
| Cross-session Joint OCV    | 7.95      | 3.02 | 7.12   |

For representative vehicles, the learned joint OCV curve consistently lay between the charging-derived and rest-derived references. This behavior suggests that the joint optimization was able to move away from the low-error but biased charging curve while avoiding the excessive noise sensitivity of the rest-derived reference. As a result, the fitted ECM preserved voltage prediction quality close to that of charging-derived OCV while recovering non-collapsed and physically plausible  $R_0$  values.

At the fleet level, the improvement in resistance physicality was clear. The step-estimate reference  $R_0$  used as an independent physical anchor in this section was computed directly from the same raw BMS telemetry, not from an external instrument-based step-response measurement, as follows. For each vehicle, current transitions with  $|\Delta I| > 5$  A between consecutive samples were identified; for each such transient the local  $\Delta V/\Delta I$  ratio gave a per-event  $R_0$  estimate, and the per-vehicle step  $R_0$  was the median across all such events within the vehicle's slow-charging trace. The fleet step-estimate value of 0.81 m $\Omega$  reported here is the median across the 23 KIA EV6 vehicles for which the per-vehicle transient density was sufficient (at least 5 qualifying transients). This estimate is biased by the finite 3–4 s BMS sampling interval, because the first-order RC dynamics with  $\tau$  of a few seconds are only partially completed within a single inter-sample interval; the step estimate therefore captures a mixture of  $R_0$  and a partial  $R_1$  contribution rather than a pure ohmic step. It is therefore used here as an order-of-magnitude physical sanity check on the jointly estimated  $R_0$ , not as a ground-truth value. The step-estimate reference yielded a median resistance of 0.81 m $\Omega$  across the 23 vehicles for which it could be computed.

Charging-derived OCV, by contrast, produced a fleet median  $R_0$  of only 0.015 m $\Omega$ , confirming widespread collapse of the ohmic resistance estimate. The joint framework recovered a median  $R_0$  of 0.895 m $\Omega$ , closely aligned with the step-estimate scale and distributed over a physically reasonable range. Although step estimates may contain partial dynamic contributions because of the finite sampling interval, they still provide a useful practical reference, and the agreement in scale supports the physical plausibility of the jointly estimated  $R_0$ .


Cross-session validation further showed that the benefit of the joint method was not limited to in-sample fitting. On the filtered 19-vehicle subset, the joint approach achieved a median self-RMSE of 4.99 mV and a median cross-session RMSE of 7.12 mV, compared with 8.04 mV for charging-derived OCV. The joint method outperformed the conventional charging-derived reference in 14 of 19 vehicles, corresponding to 74% of the filtered fleet. This result indicates that the learned OCV is not merely a flexible fit to one session, but a more transferable representation for subsequent ECM prediction.

The selection of the rest-anchor weight also showed a stable trade-off structure. As the anchor strength increased, the fitted solution moved toward more physical  $R_0$  values at the expense of some increase in RMSE. A value of  $\omega_{rest} = 0.5$  was selected as a practical operating point because it prevented collapse of  $R_0$  while maintaining low prediction error, placing the solution near the elbow of the RMSE–physicality trade-off. This trade-off curve, where resistance collapse was avoided while prediction error remained close to the low-RMSE regime, is shown in Fig. 14.

These fleet-level results have two direct implications. First, they show that physically meaningful resistance recovery is possible from coarse production BMS data when OCV and ECM parameters are identified jointly rather than sequentially. Second, they show that the resulting model generalizes more reliably across sessions than conventional charging-derived OCV fitting. These gains are directly relevant to downstream applications such as resistance-based battery health tracking and ECM-based digital twins [7,10,11], where a model must be both predictive and physically interpretable rather than optimized for voltage fit alone.

#### 4.7. Application to state estimation and anomaly detection

The preceding sections established that joint OCV-ECM optimization recovers physically meaningful ECM parameters from coarse BMS data. This subsection evaluates whether these improved parameters translate into measurable gains in two downstream diagnostic applica-



**Fig. 13.** Fleet-level validation of joint OCV-ECM estimation. (a) Fleet distribution of  $R_0$  under the charging-derived and joint OCV methods. (b) Cross-session validation results across the fleet.




Fig. 14. Pareto analysis for selecting the rest-anchor weight, showing the trade-off between voltage-fit RMSE and  $R_0$  as the rest-anchor weight increases.

tions: long-horizon SOC estimation during driving, and persistent cell-level anomaly detection across the fleet. These application-level results are important because parameter physicality alone does not guarantee diagnostic utility; the practical value of the framework must be demonstrated through its effect on quantities directly relevant to battery management [12,16,17].

SOC estimation was evaluated using an EKF with the identified 1-RC ECM [16,17]. For each vehicle, driving segments were extracted as 10-min windows from periods with sustained discharge current. The EKF was initialized with the BMS-reported SOC and propagated forward using the identified ECM parameters and OCV curve. To isolate the effect of OCV source quality, the same EKF structure and tuning were used with both charging-derived OCV (which produces near-zero  $R_0$ ) and joint OCV-ECM parameters.

In this study, SOC estimation accuracy is defined as 100 percentage points minus the SOC RMSE (with the SOC trajectory and reference both expressed in the 0–100% range). Using the per-cohort summary statistics reported below, the fleet-mean accuracy improves from 94.8% to 96.8% (5.2% RMSE  $\rightarrow$  3.2% RMSE), the fleet-median accuracy from 95.0% to 96.9% (5.0% RMSE  $\rightarrow$  3.1% RMSE), and the worst charging-OCV vehicle's (V004AK0001) accuracy from 91.5% to 98.1% (8.5% RMSE  $\rightarrow$  1.9% RMSE) under the joint OCV-ECM EKF. Single-vehicle examples follow the same accuracy = 100 – RMSE rule. Under charging conditions, both OCV methods achieved sub-1% SOC RMSE.

$$\text{Accuracy} = 100 - \text{RMSE}_{\text{SOC}}$$

The critical difference emerged during driving operation. When charging-derived OCV was used, the EKF exhibited severe SOC tracking failures in multiple vehicles. A representative case (V004AL0000), highlighted because the paper's original analysis cited this vehicle, showed a charging-OCV driving EKF RMSE of 6.0%, corresponding to an accuracy of 94.0%. Across six vehicles with driving data, the charging-OCV EKF RMSE ranged from 2.4% to 8.5%, with a fleet mean of 5.2%. This systematic gap is a direct consequence of the near-zero  $R_0$  identified under charging-derived OCV: when the ohmic resistance is negligibly small, the voltage innovation in the EKF carries almost no corrective information, and the filter effectively degenerates into open-loop Coulomb counting with OCV bias [12,16].

In contrast, the joint OCV-ECM parameters substantially improved driving SOC estimation. For V004AL0000, the driving EKF RMSE de-

creased from 6.0% to 3.2%, representing a 1.9-fold improvement. We note that the original R1 submission reported a 42.1%  $\rightarrow$  2.7% driving RMSE for this same vehicle (15.6-fold); the present 6.0%  $\rightarrow$  3.2% value is the result of an independent EKF reimplementation from the raw ev-data records following the R1 review, in which we identified and corrected a window-segmentation artifact in the original code that had inflated the charging-OCV baseline RMSE. The corrected magnitude (6.0% charging-OCV  $\rightarrow$  3.2% joint-OCV) is reported throughout the re-submission, including in the abstract. Across all six vehicles, the joint-OCV EKF achieved a fleet mean RMSE of 3.2% (range: 1.9–5.2%), compared to 5.2% for charging-derived OCV (range: 2.4–8.5%). The improvement was most pronounced for vehicles operating at high SOC, where the OCV curve is flattest and the EKF relies most heavily on the ohmic voltage drop for state correction. These results demonstrate that the physically meaningful  $R_0$  recovered by joint optimization is not merely a cosmetic improvement but directly enables voltage-based SOC correction during driving [20].

At the fleet level, the joint OCV method improved the fleet mean and median SOC RMSE, although the improvement was not strictly monotonic for every individual vehicle (Fig. 15). The median driving EKF RMSE decreased from 5.0% (charging OCV) to 3.1% (joint OCV), with the interquartile range narrowing substantially. Notably, joint OCV reduced the upper tail of the RMSE distribution: the cohort-level worst-case driving RMSE under charging-OCV was 8.5%, while under joint OCV the cohort-level worst-case driving RMSE was 5.2%, demonstrating that the benefit extends systematically across the fleet. (Note that the single worst charging-OCV vehicle V004AK0001 — which had 8.5% charging-OCV RMSE — was itself reduced to 1.9% under joint OCV; the 5.2% worst-case figure refers to a different vehicle that became the worst case after the joint-OCV recalibration.) The per-vehicle EKF values reported in this subsection were independently reimplemented from the raw evdata records and confirm the qualitative joint-OCV improvement; absolute RMSE magnitudes are sensitive to charging-OCV construction details (per-vehicle versus fleet-pooled) and are reported here in their per-vehicle independent form.

Beyond SOC estimation, the reconstructed cell-voltage trajectories also enable fleet-level anomaly detection through persistent cell deviation analysis. Because ADMM reconstruction preserves physically meaningful inter-cell differences while removing quantization artifacts, it becomes possible to distinguish genuine cell-level anomalies from measurement noise.




Fig. 15. Fleet-level SOC estimation RMSE comparison between charging-derived and joint OCV methods. Vehicle IDs on the horizontal axis are shown without the shared V004 prefix (e.g., AL0000 denotes V004AL0000).

A persistent cell deviation metric was defined as the mean absolute voltage deviation of each cell from the pack average, computed across all SOC bins where the cell appeared in more than 70% of observed sessions. This metric isolates cells that consistently deviate from the fleet norm regardless of operating point, as opposed to transient deviations that may arise from SOC-dependent thermodynamic effects or measurement noise. The detection threshold was set at the fleet mean plus two standard deviations of the maximum per-vehicle persistent deviation ( $\mu + 2\sigma = 5.6$  mV); this threshold is derived from the empirical distribution of the 27-vehicle KIA EV6 cohort analyzed here, and should be recalibrated for other fleets, pack architectures, or cell chemistries before being applied to a different population.

The persistent cell-deviation analysis was applied to a slightly broader 27-vehicle KIA EV6 cohort than the 26-vehicle joint OCV-ECM cohort, because anomaly detection only requires the ADMM-reconstructed cell voltages and does not depend on the additional cross-session OCV-fit validity filter that excluded one vehicle from the joint OCV-ECM analysis. Across 27 KIA EV6 vehicles, the persistent deviation analysis identified one vehicle (V03BA0021) as a statistically significant outlier, with 10 cells exceeding the persistence threshold (z-score = 2.5). The maximum persistent deviation for this vehicle reached 6.2 mV, compared to a fleet mean of  $2.9 \pm 1.3$  mV. Importantly, the anomalous cells in V03BA0021 showed consistent directional bias (predominantly positive deviation), suggesting a structural imbalance rather than random measurement variation. Two additional vehicles (V004BA0004, SITESTEVE6) showed elevated cell counts above the persistence threshold but remained within the normal statistical range (Fig. 16).

These diagnostic results have two practical implications. First, they confirm that the physical parameter improvements achieved by joint OCV-ECM optimization translate directly into measurable gains in application-level battery diagnosis, with driving SOC RMSE reduced substantially in the most affected vehicles. Second, they demonstrate that pack-constrained voltage reconstruction enables a persistent cell deviation metric that can identify genuinely anomalous cells at the fleet level, providing an actionable signal for predictive maintenance without requiring laboratory-grade measurement equipment [9,11,23].




Fig. 16. Fleet-level persistent cell deviation analysis on the 27-vehicle KIA EV6 cohort. The anomaly threshold is derived from this subset and may require recalibration for other fleets or chemistries. (a) Per-vehicle maximum persistent deviation distribution with the threshold ( $\text{mean} + 2\sigma$ ) marked. (b) Vehicles ranked by maximum persistent deviation (anonymized as V01-V27); bars exceeding the threshold are highlighted as flagged.

## 5. Conclusion

This study contributes a field-data-driven diagnostic framework built on 251 million real-world BMS records from 116 EVs spanning 14 vehicle models (Hyundai and Kia brands), together with a separate paired-test cohort of 135 vehicles across four pack architectures. The core voltage reconstruction and joint OCV-ECM identification analyses were conducted on slow constant-current charging segments, with fast-charging telemetry used only for a model-order identifiability check; the downstream SOC evaluation used driving-window telemetry. The central result is that onboard BMS telemetry at this scale can support physically interpretable battery diagnosis once the coarse 20 mV quantization and the biased charging-derived OCV are addressed by software-only processing.

Under slow constant-current charging, the reconstruction stage recovered voltage detail finer than the native 20 mV step while preserving exact pack-level consistency; in a representative 192-cell case, 96.5% of samples departed from the original grid at an effective redistribution scale of approximately 0.1 mV. On these reconstructed data, ECM fitting error dropped by 11–49% across four representative vehicle models, and model-order analysis showed no reliable benefit from additional dynamic branches beyond the first-order (1-RC) structure [5, 6].

The dominant driver of parameter quality was the OCV reference, not model complexity. When OCV was taken directly from charging data, the estimated ohmic resistance collapsed to a physically implausible near-zero value in 24 of 26 vehicles (fleet median 0.015 m $\Omega$ ), because the charging overpotential was absorbed into the OCV curve [15]. Estimating the OCV and the resistance jointly restored the fleet-median resistance to 0.895 m $\Omega$  — consistent with a separately computed in-telemetry current-step estimate of 0.81 m $\Omega$  — and lowered the cross-session prediction error from 8.04 mV to 7.12 mV, outperforming the charging-derived baseline in 74% of vehicles.

These improved parameters translated directly into diagnostic gains. Driving-window SOC estimation error fell from a fleet mean of 5.2% to 3.2%, with the representative vehicle V004AL0000 improving from 6.0% to 3.2% [16,17]. The voltage trajectories reconstructed from slow constant-current charging segments also supported persistent cell-deviation analysis that separated genuinely anomalous cells from transient measurement noise at the fleet level [9,11]. Together, these results show that scaled real-world BMS data can deliver both improved parameter estimation and actionable field-level diagnostics.

First, the joint OCV-ECM validation was concentrated on 26 KIA EV6 vehicles, including 19 vehicles after cross-session filtering and 6 vehicles in the driving-window EKF demonstration. The ADMM preprocessing benefit was evaluated separately in a paired cohort of 135 vehicles across four Hyundai/Kia pack architectures: EV6 192S, EV9 152S, IONIQ5 LR 180S, and IONIQ5 LR 2022 192S. Their respective sample sizes were  $n = 67/13/34/21$ , with paired  $t$ -test  $p < 10^{-5}$  in all four cohorts. Further generalization to other manufacturers, chemistries (e.g. NMC versus LFP cells), or driving conditions still requires dedicated study, and the joint OCV-ECM and driving-EKF results should be read as concentrated evidence on the EV6 cohort rather than a fleet-wide claim across heterogeneous EV populations.

Second, the multi-vehicle paired test added in Section 4.2 indicates that the reconstruction operating point adopted here, while sufficient for the present study, is not necessarily globally optimal for fleet-wide cross-session evaluation. The same-session raw-versus-enhanced improvement remains statistically significant across all four pack architectures ( $p < 10^{-5}$ ), so the core conclusion is robust; but a deterministic, training-free regularizer applied uniformly across heterogeneous vehicles is an inherent design choice, and a productive opening for follow-up work. Three extensions follow naturally. First, the reconstruction regularization strength could be tuned for each vehicle according to pack architecture, sampling rate, and operating regime. Second, the

regularizer could be replaced by a data-driven or diffusion-based reconstruction prior trained on the same telemetry corpus and then frozen for deployment. Third, the hard pack-sum equality could be relaxed to a soft penalty, allowing per-time-step compensation for small per-cell bus-bar contributions. The first two extensions preserve the pack-constrained reconstruction objective while changing the inner regularizer; the third is a mild relaxation of the exact pack-sum constraint into a penalty term, trading strict equality for additional flexibility. None of them changes the joint OCV-ECM identification framework introduced here.

Third, the 1-RC ceiling identified in Section 3.3 and Section 4.3 is a limitation of the sensing regime rather than a modeling shortcut. Independent identification of higher-order RC dynamics would require transient or driving data, finer voltage resolution, and higher-rate (sub-second) sampling. These measurements are outside the current onboard BMS logging specification. The same measurement constraint also limits the precision of the first-order RC pair, as quantified by the slow-charging Cramér–Rao bounds in Section 3.2. As a result, the reported 1-RC parameters represent the highest RC order practically supported by the available telemetry, and cross-chemistry or transient-regime generalization of higher-order dynamics is left for follow-up work with higher-rate logging.

Despite these limitations, the results demonstrate that production-grade EV telemetry can support interpretable ECM identification and practical fleet-level battery diagnostics through software-only enhancement of existing onboard BMS data. This finding is directly relevant to EV fleet operators and battery diagnostic teams seeking deployable diagnostic tools that do not require additional sensors or laboratory-grade characterization. Future work will extend the framework toward online implementation, transient-condition correction, cross-vehicle OCV transfer, and resistance-based SOH tracking across heterogeneous fleet populations. Broader generalization to other cell chemistries (e.g., LFP) and pack architectures beyond the Hyundai/Kia cohort analyzed here also remains an open direction.

## CRedit authorship contribution statement

**Younggyu Hwang:** Writing – original draft, Visualization, Validation, Methodology, Conceptualization. **Hyeongyu Son:** Validation. **Seungho Geum:** Methodology. **Musu Kim:** Conceptualization. **Sekyung Han:** Writing – review & editing.

## Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author(s) used Gemini, Claude, and GPT solely to improve the English expression of the manuscript. After using these tools/services, the author(s) critically reviewed and revised the manuscript and take full responsibility for the content of the published article.

## Declaration of competing interest

This study was supported by the BK21 Four project funded by the Ministry of Education, Korea (4199990113966)

## Acknowledgements

This study was supported by the BK21 Four project funded by the Ministry of Education, Korea (4199990113966).

## Data availability

The vehicle telemetry dataset analyzed in this study is proprietary, provided by Betterwhy under a restricted research access agreement for

the purposes of this work; the dataset is not publicly available and cannot be redistributed in any form. References [9–11] describe the underlying telemetry-collection framework and digital-twin context rather than serving as direct data-access pointers.

## References

- [1] B. Zheng, Z. Deng, Z. Luo, S. Mao, M. Ouyang, X. Han, H. Wang, Y. Li, Y. Sun, D. Wang, Y. Yuan, L. He, Z. Yang, Y. Zhu, A comprehensive review of lithium-ion battery modelling research and prospects: in-depth analysis of current research and future directions, *Appl. Energy* 401 (2025) 126688, <https://doi.org/10.1016/j.apenergy.2025.126688>.
- [2] X. Hu, L. Xu, X. Lin, M. Pecht, Battery lifetime prognostics, *Joule* 4 (2) (2020) 310–346, <https://doi.org/10.1016/j.joule.2019.11.018>.
- [3] W. Waag, C. Fleischer, D.U. Sauer, Critical review of the methods for monitoring of lithium-ion batteries in electric and hybrid vehicles, *J. Power Sources* 258 (2014) 321–339, <https://doi.org/10.1016/j.jpowsour.2014.02.064>.
- [4] C.R. Birk, M.R. Roberts, E. McTurk, P.G. Bruce, D.A. Howey, Degradation diagnostics for lithium ion cells, *J. Power Sources* 341 (2017) 373–386, <https://doi.org/10.1016/j.jpowsour.2016.12.011>.
- [5] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers, *Found. Trends Mach. Learn.* 3 (2011) 1–122, <https://doi.org/10.1561/22000000016>.
- [6] L.I. Rudin, S. Osher, E. Fatemi, Nonlinear total variation based noise removal algorithms, *Phys. D* 60 (1992) 259–268, [https://doi.org/10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F).
- [7] K. Alamin, Y. Chen, E. Macii, S. Vinco, Advancing electric vehicle battery management: a data-driven digital twin approach for real-time monitoring and performance enhancement, *IEEE Trans. Veh. Technol.* 74 (2025) 13850–13864, <https://doi.org/10.1109/TVT.2025.3565907>.
- [8] M. Ibrahim, V. Rjabtsikov, R. Gilbert, Overview of digital twin platforms for EV applications, *Sensors (Basel, Switzerland)* 23 (2023) 1414, <https://doi.org/10.3390/s23031414>.
- [9] J. Zhao, X. Feng, M.-K. Tran, M. Fowler, M. Ouyang, A.F. Burke, Battery safety: Fault diagnosis from laboratory to real world, *J. Power Sources* 598 (2024) 234111, <https://doi.org/10.1016/j.jpowsour.2024.234111>.
- [10] F. von Bülow, A Data-Driven Fleet Service: State of Health Forecasting of Lithium-ion Batteries, Springer, 2024, <https://doi.org/10.1007/978-3-658-43188-4>.
- [11] K. Chen, Y. Luo, Z. Long, Y. Li, G. Nie, K. Liu, D. Xin, G. Gao, G. Wu, Big data-driven prognostics and health management of lithium-ion batteries: a review, *Renew. Sust. Energ. Rev.* 214 (2025) 115522, <https://doi.org/10.1016/j.rser.2025.115522>.
- [12] J. Tian, R. Xiong, W. Shen, J. Lu, State-of-charge estimation of LiFePO<sub>4</sub> batteries in electric vehicles: a deep-learning enabled approach, *Appl. Energy* 291 (2021) 116812, <https://doi.org/10.1016/j.apenergy.2021.116812>.
- [13] Y. Wu, D. Bai, K. Zhang, Y. Li, F. Yang, Advancements in the estimation of the state of charge of lithium-ion battery: a comprehensive review of traditional and deep learning approaches, *J. Mater. Inform.* 5 (2025) 18, <https://doi.org/10.20517/jmi.2024.84>.
- [14] N. Ghaeminezhad, Q. Ouyang, J. Wei, T. Ma, J. Wang, C. Zou, Review on state of charge estimation techniques of lithium-ion batteries: a control-oriented approach, *J. Energy Storage* 72 (2023) 108707, <https://doi.org/10.1016/j.est.2023.108707>.
- [15] K. Zhang, R. Xiong, Q. Li, C. Chen, J. Tian, W. Shen, A novel pseudo-open-circuit voltage modeling method for accurate state-of-charge estimation of LiFePO<sub>4</sub> batteries, *Appl. Energy* 347 (2023) 121406, <https://doi.org/10.1016/j.apenergy.2023.121406>.
- [16] N. Shi, Z. Chen, M. Niu, Z. He, Y. Wang, J. Cui, State-of-charge estimation for the lithium-ion battery based on adaptive extended Kalman filter using improved parameter identification, *J. Energy Storage* 45 (2022) 103518, <https://doi.org/10.1016/j.est.2021.103518>.
- [17] S. Nachimuthu, F. Alsaif, G. Devarajan, I. Vairavasundaram, Real time SOC estimation for Li-ion batteries in electric vehicles using UKBF with online parameter identification, *Sci. Rep.* 15 (2025) 1714, <https://doi.org/10.1038/s41598-025-85700-0>.
- [18] C. Zietlow, J.K.N. Lindner, ADMM-TGV image restoration for scientific applications with unbiased parameter choice, *Numer. Algorithms* 97 (2024) 1481–1512, <https://doi.org/10.1007/s11075-024-01759-2>.
- [19] T. Goldstein, B. O'Donoghue, S. Setzer, R. Baraniuk, Fast alternating direction optimization methods, *SIAM J. Imag. Sci.* 7 (2014) 1588–1623, <https://doi.org/10.1137/120896219>.
- [20] L. Wang, J. Duan, S. Fan, K. Zhao, An estimated value compensation method for state of charge estimation of lithium battery based on open circuit voltage change rate, *Energy* 313 (2024) 134119, <https://doi.org/10.1016/j.energy.2024.134119>.
- [21] M. Lagnoni, C. Scarpelli, G. Lutzemberger, A. Bertel, Critical comparison of equivalent circuit and physics-based models for lithium-ion batteries: a graphite/lithium-iron-phosphate case study, *J. Energy Storage* 94 (2024) 112326, <https://doi.org/10.1016/j.est.2024.112326>.
- [22] P. Kumar, S. Sundaresan, B. Balasingam, Theoretical performance evaluation of battery equivalent circuit model parameter estimators, *IEEE J. Emerg. Sel. Top. Ind. Electron.* (2023), <https://doi.org/10.1109/JESTIE.2023.3342717>.
- [23] N. Pahaisuk, S. Pourbunthidkul, P. Phasukkitt, N. Houngkamhang, Intelligent battery management system for electric vehicles: AI-driven voltage cell prediction using GRU and K-means clustering, *IEEE Access* 13 (2025) 164552–164566, <https://doi.org/10.1016/j.apenergy.2024.124748>.
- [24] A. Tian, K. Dong, X.-G. Yang, Y. Wang, L. He, Y. Gao, J. Jiang, Physics-based parameter identification of an electrochemical model for lithium-ion batteries with two-population optimization method, *Appl. Energy* 378 (2025) 124748, <https://doi.org/10.1016/j.apenergy.2024.124748>.
- [25] J.A. Kuzhiyil, T. Damoulas, W.D. Widanage, Neural equivalent circuit models: Universal differential equations for battery modelling, *Appl. Energy* 371 (2024) 123692, <https://doi.org/10.1016/j.apenergy.2024.123692>.
- [26] J. Sun, Y. Liu, J. Kainz, A comparative study of parameter identification methods for equivalent circuit models for lithium-ion batteries and their application to state of health estimation, *J. Energy Storage* 114 (2025) 115707, <https://doi.org/10.1016/j.est.2025.115707>.
- [27] K. Zhang, X. Zhao, Y. Chen, D. Wu, T. Cai, Y. Wang, L. Li, J. Zhang, Robust equivalent circuit model parameters identification scheme for State of Charge (SOC) estimation based on maximum correntropy criterion, *Int. J. Electrochem. Sci.* 19 (2024) 100558, <https://doi.org/10.1016/j.ijoes.2024.100558>.
- [28] J. Pi, S. Filgueira da Silva, M.F. Ozkan, A. Gupta, M. Canova, Parameter identification for electrochemical models of lithium-ion batteries using Bayesian optimization, *IFAC-PapersOnLine* 58 (2024) 180–185, <https://doi.org/10.1016/j.ifacol.2024.12.031>.
- [29] C.-J. Ko, K.-C. Chen, Using tens of seconds of relaxation voltage to estimate open circuit voltage and state of health of lithium ion batteries, *Appl. Energy* 357 (2024) 122488, <https://doi.org/10.1016/j.apenergy.2023.122488>.
- [30] Y. Liu, Y. Liu, X. Bao, H. Shen, Estimation of battery state of health and open circuit voltage at various depths of discharge based on deep learning and relaxation voltage, *Energy* 328 (2025) 136555, <https://doi.org/10.1016/j.energy.2025.136555>.
- [31] A. Sharma, H.K. Fathy, Fisher identifiability analysis for a periodically-excited equivalent-circuit lithium-ion battery model, in: 2014 American Control Conference, IEEE, 2014, pp. 274–280, <https://doi.org/10.1109/ACC.2014.6859360>.
- [32] C. Schaefer, E. Din, J.T. Stauth, A hybrid switched-capacitor battery management IC with embedded diagnostics for series-stacked Li-ion arrays, *IEEE J. Solid State Circuits* 52 (2017) 3142–3155, <https://doi.org/10.1109/JSSC.2017.2734902>.
- [33] X. Wang, J. Xu, Y. Zhao, Wavelet based denoising for the estimation of the state of charge for lithium-ion batteries, *Energies* 11 (2018) 1144, <https://doi.org/10.3390/en11051144>.
- [34] Z. Cui, N. Cui, C. Wang, C. Li, C. Zhang, A robust online parameter identification method for lithium-ion battery model under asynchronous sampling and noise interference, *IEEE Trans. Ind. Electron.* 68 (2021) 9550–9560, <https://doi.org/10.1109/TIE.2020.3028799>.
- [35] F.N. Fritsch, R.E. Carlson, Monotone piecewise cubic interpolation, *SIAM J. Numer. Anal.* 17 (1980) 238–246, <https://doi.org/10.1137/0717021>.