ACS APPLIED MATERIALS & INTERFACES

www.acsami.org

Research Article

Gas-Informed Machine Learning Framework for Stage Classification and Early Forecasting of Battery Degradation

Yiqing Lu,[#] Jianli Zou,[#] Liang Zhang, Haoyu Wang,* and Franklin Kim*

Cite This: *ACS Appl. Mater. Interfaces* 2026, 18, 26311−26322

Read Online

ACCESS | Metrics & More | Article Recommendations | Supporting Information

ABSTRACT: Battery health prediction typically relies on electrical signals such as capacity, voltage, and impedance, which are chemically nonspecific and can lag the onset of failure. Here we show that operando gas evolution measured by online electrochemical mass spectrometry (OEMS) provides SOC-resolved chemical signatures that enable earlier degradation-stage identification and long-horizon capacity forecasting. Using intermittently sampled CO, CO_2, and C_2H_4 profiles, we develop two complementary models: (i) a label-free PCA−K-means stage classifier that distinguishes linear and accelerated nonlinear degradation, and (ii) a single-shot forecaster that fuses a short capacity/SOH history with one OEMS-measured cycle to predict future capacity trajectories without autoregressive rollout. On held-out test cells, the stage classifier exhibits a 3.6% misclassification rate (5/136 OEMS-measured cycles) and flags the nonlinear transition 32.0 ± 17.2 cycles before the Bacon−Watts knee (18.0 ± 8.1% of lifetime). On held-out cells closest to the training protocol, the forecaster achieves 1.7−2.9% RMSE at 100-cycle horizons and reduces long-horizon drift relative to capacity-only baselines. These results establish gas evolution as a chemically grounded complement to electrical monitoring and motivate future integration with compact off-gas sensing for cycle-level battery health management.

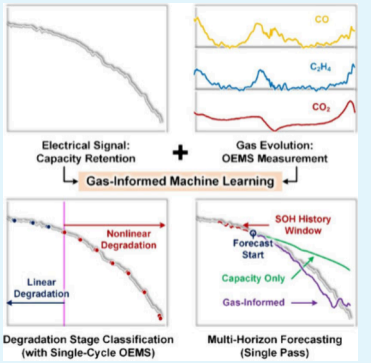

KEYWORDS: *lithium-ion battery, gas evolution, degradation stage classification, degradation forecasting, machine learning*

INTRODUCTION

Reliable operation of lithium-ion batteries depends on timely assessment of their state of health (SOH) and early recognition of degradation pathways.[1−3] Conventional SOH metrics are most often based on capacity retention, yet this measure provides only a limited view of internal chemistry. Distinct failure mechanisms can lead to the same apparent SOH, and as a result the metric alone provides little predictive power for degradation progression. This limitation highlights the need for approaches that can identify internal degradation status while also anticipating the full trajectory of capacity fade.

The capacity retention trajectory of lithium-ion batteries commonly exhibits two phases: an initial, approximately linear regime followed by an accelerated nonlinear regime.[4,5] Detecting the transition, often referred to as the knee point, is critical for timely intervention.[6] Conventional approaches treat knee identification as a curve-fitting problem, including piecewise-linear fits such as the Bacon−Watts model and tangent-based methods.[7,8] These techniques require capacity data that span the knee (often the full life). As a result, the knee is identified only after it has occurred, delaying detection.[7,9,10] In addition, the results are often highly sensitive to the selected fitting range and manual initialization.

To address these limitations, recent studies have turned to machine learning (ML) for earlier knee detection and forecasting.[11−14] Supervised classifiers can identify knee onset earlier than retrospective mathematical fits, but their performance still depends on the quality of the labels used for training.[12] More recently, Transformer-based approaches such as ICFormer have leveraged incremental-capacity (IC) curves to detect knees and degradation modes without explicit curve fitting, although they require high-quality IC data collected over extended cycling windows, which may limit practical deployment.[14] Even when the knee can be identified, forecasting from the pre-knee regime remains challenging. Conventional autoregressive (AR) models accumulate error over long horizons, and the nearly constant slope of the pre-knee region encourages linear extrapolation that can miss the impending inflection.[15−17] Transfer-learning strategies can reduce horizon drift, but accurate localization and anticipation of the knee remain difficult.[18] Advanced sequence models such as the Temporal Fusion Transformer (TFT) have also been explored for single-shot, multihorizon forecasting, yet capacity-history-only forecasters may still have limited early warning capability in the nearly linear regime when precursor information is absent from the input.[19]

Received: February 5, 2026
Revised: April 9, 2026
Accepted: April 12, 2026
Published: April 27, 2026

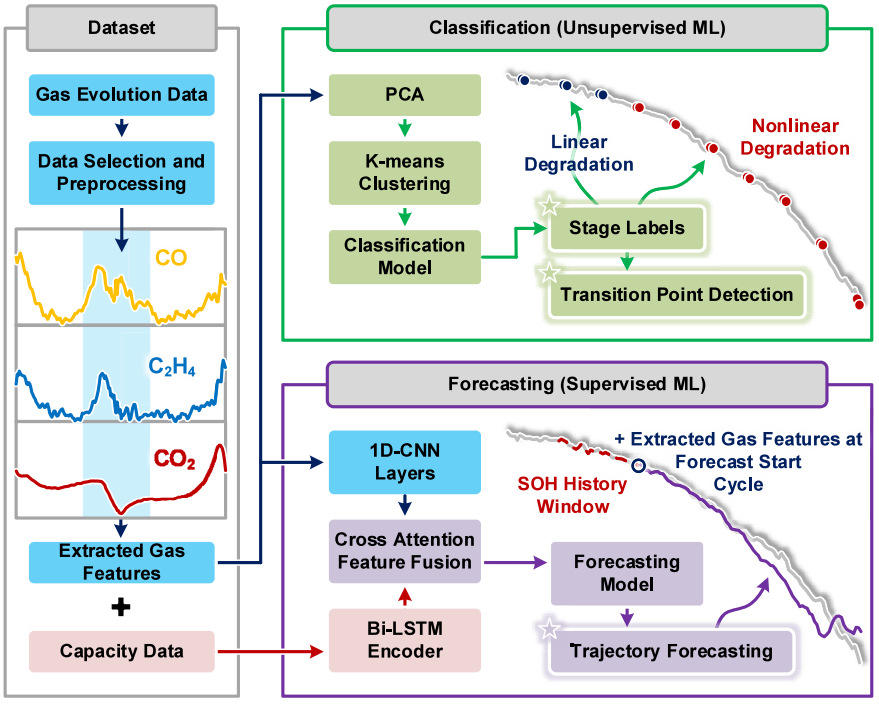

Figure 1. Schematic of the gas-informed machine learning framework for battery degradation analysis. Gas evolution data from individual charge–discharge cycles, collected via online electrochemical mass spectrometry (OEMS), are preprocessed and remapped to the state of charge (SOC) domain. Through correlation analysis, weakly correlated gases (CO, CO_2, and C_2H_4) within specific SOC window that provide distinct chemical insights are selected. The extracted gas serves as the input feature set for both classification and forecasting tasks. For classification, principal component analysis (PCA) reduces the dimensionality of the gas features, followed by K-means clustering to distinguish between linear and nonlinear degradation stages. For forecasting, a supervised model fuses a single-cycle gas profile, processed through a 1D-CNN with channel-wise attention, with a short history of normalized capacity (SOH) encoded by a Bi-LSTM. A cross-attention mechanism integrates these features to predict multihorizon capacity trajectories in a single shot, without autoregressive rollout.

A fundamental limitation of capacity-based approaches is that electrical signals provide only an indirect view of internal chemistry.[20] In contrast, gas evolution offers a direct, mechanistically grounded indicator of degradation.[21,22] The composition and rate of evolved gases are closely linked to key degradation pathways, including solid electrolyte interphase (SEI) instability, solvent oxidation, cathode dissolution, and additive breakdown,[23−25] making gas signals both mechanistically specific and temporally responsive. Unlike capacity, which often changes slowly and is confounded by operating conditions (e.g., temperature or load), gas evolution captures real-time signatures of the underlying reactions.[26,27] In particular, gases such as CO, CO_2, and C_2H_4 are strongly associated with electrolyte degradation and SEI failure, which are two defining features of the transition into nonlinear degradation. These gas-phase signals therefore provide an early window into the onset of nonlinear degradation, enabling predictive diagnostics before capacity changes are detectable. Despite this promise, to our knowledge, no existing method combines chemically grounded signals with long-range forecasting.

To demonstrate the diagnostic potential of gas evolution, we develop two complementary gas-informed ML frameworks: (i) an unsupervised stage-classification model and (ii) a single-shot, multihorizon forecaster for long-range capacity trajectories. These frameworks are illustrated in Figure 1. The models operate on gas-evolution data collected using online electrochemical mass spectrometry (OEMS) acquired intermittently during cycling.[28,29] Through correlation analysis of the OEMS data, we select weakly correlated gases (CO, CO_2, and C_2H_4) within a specific state of charge (SOC) window that is rich in electrochemical information. Gas curves within this window are extracted as features for both classification and forecasting tasks.

First, we present the unsupervised classification pipeline for detecting the onset of nonlinear degradation in NMC/graphite cells. Principal component analysis (PCA) reduces dimensionality, followed by K-means clustering to assign each cycle to either the stable linear or accelerated nonlinear regime. Despite the absence of labeled training data, the method yields a misclassification rate of 3.6% (5/136) on an independent test set. These results show that gas-phase signatures encode stage transitions and enable cycle-level health diagnostics without retrospective capacity trajectories.

Building on the stage classifier, we develop a gas-conditioned capacity forecaster that fuses one SOC-resolved gas profile from a single OEMS-measured cycle with a short SOH fragment to produce single-shot, multihorizon capacity trajectories without autoregressive roll-out. Gas evolution reflects chemically specific parasitic reactions that can intensify while the capacity trajectory still appears linear; conditioning on this chemical context enables the model to anticipate the impending rollover even when the input capacity history shows negligible curvature.

Together, these results establish gas evolution as a chemically grounded, data-efficient complement to electrical monitoring methods such as capacity, incremental capacity, and impedance. The framework enables stage identification from a single OEMS-measured cycle and knee-aware forecasting from minimal historical data, reducing reliance on

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322

dense electrical feature tracking and long post-knee capacity histories. As embedded off-gas sensing technologies mature, this minimal-input design is conceptually compatible with future integration into battery management systems.

EXPERIMENTAL SECTION

Materials

21700-format dry cylindrical lithium-ion cells without the outer stainless-steel case were purchased from Jiangxi Better Way New Energy Technology Co., Ltd. (Jiangxi, China). Each cell consisted of a lithium−nickel−manganese−cobalt oxide (NMC811) cathode and a graphite anode (NMC/graphite), with a nominal capacity of approximately 4.5 Ah.

Cell Preparation and in Situ OEMS Configuration

The dry cylindrical cell assembly was inserted into a MicroElab MS-21700 gastight in situ OEMS test cell configured for a 21700-format cylindrical assembly (schematic shown in Figure S1). The electrolyte was 1.0 M LiPF6 in ethylene carbonate/diethyl carbonate (EC/DEC, 50/50 vol %). After electrolyte addition, the assembled cell was rested for 2 days and then formation-cycled at 0.05 C for two cycles. Subsequent cycling at the specified charge/discharge rates was carried out using a LANHE G340A battery cycler.

Electrochemical Cycling Protocols

Cells were cycled under controlled charge/discharge protocols designed to span meaningful operational variations (e.g., charge/discharge rate, electrolyte volume, and additive usage). A standard cycle consisted of constant-current (CC) charge, constant-voltage (CV) charge, CC discharge, and a voltage hold step, following the protocol used throughout this work. Most cells were cycled at 0.5C/0.5C, with selected cells cycled under modified conditions to evaluate robustness under domain shift. A per-cell summary of conditions and the train/test partition is provided in Table S1.

Operando Gas Measurement by OEMS

Gas evolution during cycling was monitored using a mass spectrometer (Hiden HPR-20 R&D). Fragment ion channels corresponding to the key gas species were recorded (e.g., m/z = 2, 15, 27, 28, 30, 32, 44). Prior to the measurements, OEMS signals were calibrated using certified standard gas mixtures and a pure argon background under the same sampling configuration as the battery experiments. After applying fragmentation corrections, calibration curves were obtained by fitting gas concentration versus ion-current signal for each species. Calibration details are provided in SI Section 1.

For the OEMS measurements, the maximum flow rate in the capillary inlet of the mass spectrometer was set at 0.8 mL/min. Argon gas was used as the carrier gas at a flow rate of 1.5 mL/min with a bypass to maintain stable system pressure. Prior to each measurement, the system was purged with argon for 2 h to stabilize the background. In this study, the fragmented molecular ion peaks at m/z = 2, 15, 27, 28, 30, 32, and 44 were monitored, corresponding to H_2, CH_4, C_2H_4, CO, C_2H_6, O_2, and CO_2 gases, respectively. The C_2H_4 signal at m/z = 27 was corrected by subtracting the contribution from C_2H_6 at m/z = 30. Similarly, the m/z = 28 signal, attributed to CO, was corrected by subtracting contributions from C_2H_4, C_2H_6, and CO_2.

Gas-Signal Preprocessing and SOC Remapping

Raw gas traces were preprocessed to improve comparability across long-term experiments. The following two steps were conducted. 1) Noise reduction: A smoothing filter (e.g., moving average) was applied to attenuate high-frequency fluctuations. 2) Baseline correction: Baseline drift was corrected using an asymmetric least-squares (ALS)/Whittaker-smoothing approach. To compare cycles with different durations as the cell aged, gas signals were remapped from the time domain to the state-of-charge (SOC) domain using coulomb counting based on recorded current and capacity. Each gas profile was resampled on a uniform SOC grid (algorithmic details and parameters for preprocessing and SOC remapping both in SI Section 2).

Feature Construction: Gas Selection and SOC Window Selection

Multiple gases were monitored; the ML feature set was constructed using species that provide complementary (nonredundant) information. Gas selection was performed using inter-gas correlation analysis (Pearson correlation) across cycles, and a minimally correlated subset was selected for downstream learning. The final feature set used CO, CO_2, and C_2H_4. Selection criteria and correlation results are provided in SI Section 3 and Figure S2.

In addition, rather than using the full SOC range, a SOC region of interest was selected to maximize cluster separability of degradation stages. SOC-window boundaries were optimized using an unsupervised criterion (silhouette coefficient of K-means clustering, Figure S3) evaluated over candidate SOC intervals. The selected window was then used consistently for both classification and forecasting. After SOC-window selection, each OEMS-instrumented cycle yields a fixed-length feature vector consisting of the concatenated gas profiles of the selected species sampled on the SOC grid. For the classification pipeline, these SOC-windowed gas profiles are Z-score normalized before PCA and K-means; for forecasting, the gas profiles are not normalized so as to preserve absolute magnitudes and inter-gas ratios. The overall preprocessing and feature-construction workflow is summarized in Figure S4.

Unsupervised Degradation-Stage Classification

An unsupervised pipeline was used to classify OEMS-instrumented cycles into degradation stages (SI Section 5, Figure S5):

Principal component analysis (PCA) was applied (separately per gas species) to reduce dimensionality while retaining degradation-relevant variance. The number of retained PCs per gas was selected by jointly considering explained-variance retention and clustering separability (silhouette coefficient). The quantitative PC-selection procedure is provided in SI Section 6 (including Figures S6 and S7). K-means clustering was then performed in the PCA feature space. The optimal number of clusters was determined using cluster validity indices (silhouette), and the final classifier uses K = 2 to represent linear vs accelerated nonlinear degradation. The robustness of K selection (including evaluation of K = 3) is provided in SI Section 7 and Figures S8 and S9.

PCA transforms and K-means centroids were fit using the training cells only. For a new cycle (test cell), features were preprocessed identically, projected using the training PCA transforms, and assigned to the nearest K-means centroid. To emulate practical monitoring, test samples were evaluated in chronological order. A nonlinear "warning" is emitted if either of two adjacent OEMS-instrumented cycles is classified as nonlinear; a transition is declared when two consecutive tests contain nonlinear warnings. This reduces sensitivity to single-cycle disturbances and measurement variance.

Because the classifier is unsupervised and does not rely on manual labels, evaluation emphasizes deployment-relevant temporal coherence and transition-detection behavior. Where required for comparison, knee points were estimated from capacity trajectories using Bacon−Watts piecewise-linear fitting. Early detection advantage from the unsupervised degradation-stage classification pipeline was quantified by defining the lead cycle as

$$\Delta N = N_{knee} - N_{pred}$$

where N_{knee} is the mathematical knee cycle obtained from the Bacon−Watts fit and N_{pred} is the gas-informed predicted transition cycle, defined by our sequential warning rule (two consecutive tests containing nonlinear warnings). The per-cell values of N_{pred}, N_{knee}, and ΔN are summarized in Table S2. To characterize deployment-oriented reliability under intermittent sampling without external stage labels, we additionally quantify label-free monitoring stability (within-test disagreement, transient pretransition warnings, and post-transition persistence) and summarize per-cell and pooled results in Table S3.

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322

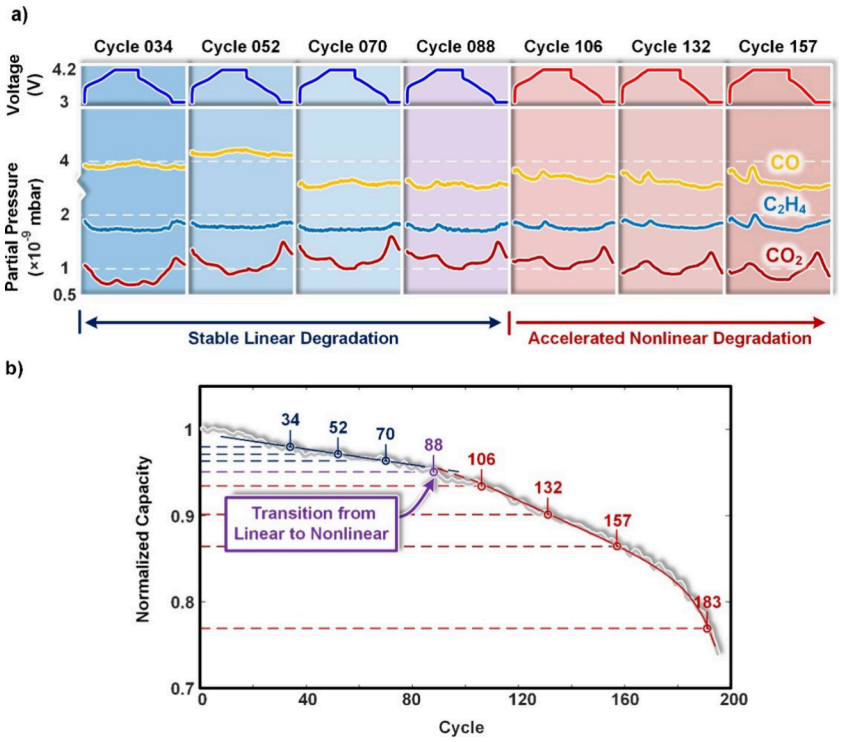

Figure 2. Representative data set used in the ML framework. (a) Voltage profiles and corresponding gas evolution signals of CO, CO_2, and C_2H_4, acquired periodically using online electrochemical mass spectrometry (OEMS) during long-term cycling of an NMC/graphite cylindrical cell. (b) Normalized discharge capacity versus cycle number for the same cell. The degradation trajectory separates into a stable linear stage and an accelerated nonlinear stage; the onset of the nonlinear stage coincides with a pronounced increase in the evolution of CO, CO_2, and C_2H_4. These three gases were used as input features in subsequent ML analyses.

Gas-Informed Capacity-Trajectory Forecasting

A supervised forecaster was developed to predict long-horizon capacity trajectories using (i) a short historical capacity fragment and (ii) a single-cycle gas snapshot from the same end-of-fragment cycle (Figure S10):

- Inputs: a fixed-length SOH/capacity history window and the SOC-windowed gas profiles of $CO/CO_2/C_2H_4$ from the final cycle of that history window.
- Gas encoder: A lightweight 1D-CNN extracts local patterns from the multigas SOC profiles, followed by channel-wise attention to reweight gas channels by relevance.
- SOH encoder: A Bi-LSTM encodes the short capacity history to capture temporal trends.
- Fusion: Cross-attention integrates the gas representation with the SOH representation to produce a chemically informed latent state.
- Decoder/prediction: A single-shot multihorizon decoder predicts the future capacity trajectory without autoregressive rollout.

Architecture details and training hyperparameters are reported in Table S5. To prevent leakage training and test sets were split by cell (cells in the test set were never used for training/validation). Also, within each cell, historical inputs always precede prediction targets in time.

Model development used 5-fold cross-validation on the training cells (stratified to preserve inter-cell representation). Forecast uncertainty was quantified using an ensemble approach: the 5 independently trained models from cross-validation were used as an ensemble, and prediction mean and variance were computed across ensemble members. A 95% uncertainty band was reported as mean $\pm 2\sigma$. Forecast accuracy was evaluated using root mean square error (RMSE) and maximum absolute error (maxAE) over multiple start points and forecast horizons (L_{pred},50 or 100 cycles, or Full), with summary statistics across test scenarios reported in SI Table S6. For representative advanced-baseline comparison, a capacity-only Temporal Fusion Transformer (TFT) forecaster was also trained and evaluated using the same cell-level train/test split, 32-cycle SOH-history input, prediction start points, and forecast horizons; the resulting case-by-case comparison for the two representative held-out cells shown in Figure 8 is reported in SI Table S7

RESULTS

Gas Evolution Characteristics, Preprocessing, and Data Set Definition

In this study, 4.5 Ah NMC/graphite 21700-format cylindrical cells were used as the model system. Gas evolution during extended cycling was measured using OEMS. In our previous work,[27] we demonstrated the use of OEMS to probe gas evolution across aging stages and observed multiple hydrocarbon gases such as ethylene prior to the onset of nonlinear degradation. Here, we expand the data set by incorporating additional cells cycled under diverse conditions (10 cells and 273 cycles in total), providing a broader, more heterogeneous basis for training and validating the ML models.

Gaseous species such as CO, CO_2, C_2H_4, CH_4, and C_2H_6 were monitored approximately every 20 cycles. Representative gas like CO, CO_2 and C_2H_4 are plotted in Figure 2a to show the gas-evolution signals during extended cycling of a single cell. To mitigate baseline drift and high-frequency noise inherent in long-term OEMS measurements, the raw gas evolution data were processed with a two-step pipeline. First, a moving average filter was applied to attenuate transient noise arising from instrumental fluctuations. Next, baseline drift was corrected using the asymmetric least-squares (ALS) method with Whittaker smoothing, which penalizes abrupt changes and

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322

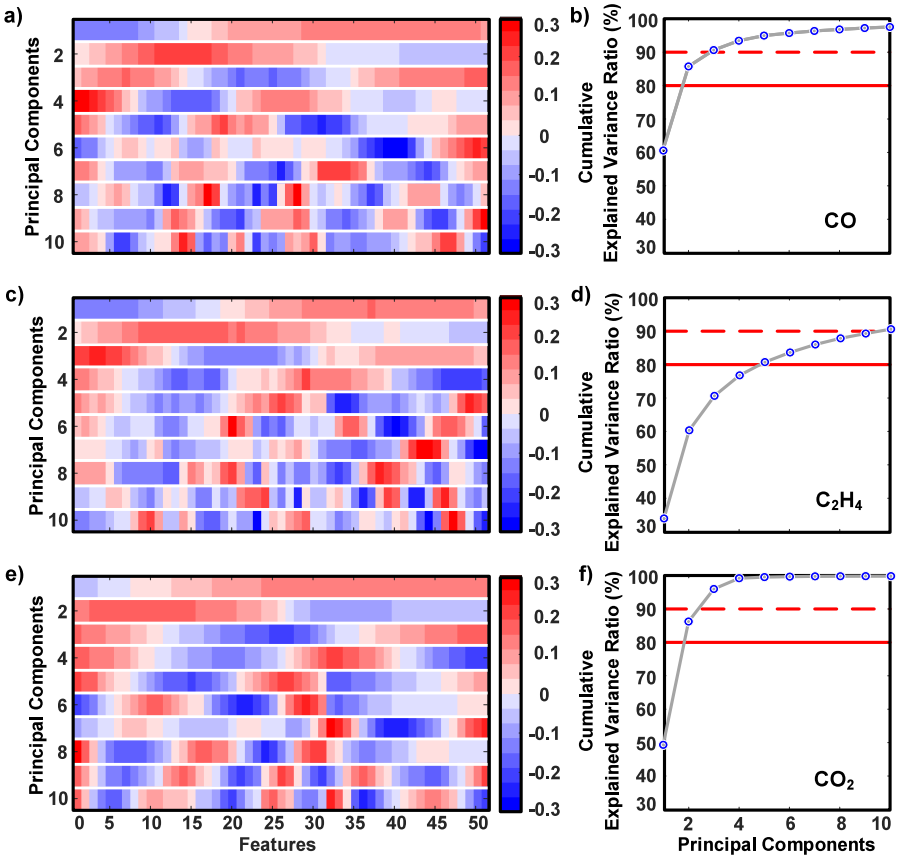

Figure 3. PCA loadings and cumulative explained variance ratio of gas profiles. (a, c, e) Heatmaps of the PCA loading matrices for CO, C_2H_4, and CO_2, respectively. Rows correspond to the top eight principal components, ranked by explained variance ratio, and columns represent the 50 SOC-mapped features derived from the gas profiles. Color intensity reflects the contribution of each feature to a given PC, with darker red and blue indicating stronger positive or negative loadings. (b, d, f) Cumulative Explained variance ratio plots (elbow plots) for the PCs of each gas. The red solid line marks the 80% retention threshold and the dashed line marks the 90% retention threshold that bound the optimal PC search space.

produces a smoothly varying baseline. The ALS algorithm applies asymmetric weighting to ensure that the fitted baseline remains below the signal, thereby preserving peak features.[30]

To ensure comparability across cycles with varying durations and capacities, gas evolution profiles were remapped from the time domain to the SOC domain. Each OEMS gas signal was uniformly sampled at 1% SOC intervals, resulting in 100 points during charging and 100 points during discharging for each gas. This transformation is critical because the durations of the charge and discharge steps shorten as degradation advances, making time-aligned comparisons unreliable. To prepare the data set for machine learning, Z-score standardization was applied to each gas profile, transforming features to zero mean and unit variance. This normalization enhances algorithm stability and mitigates feature scaling bias during clustering.

Although preprocessing steps such as ALS and standardization are sometimes discouraged in ML pipelines due to concerns about bias or overfitting, their use is justified here given the relatively small training set. In this context, baseline correction and noise filtering improve feature quality, reduce variance attributable to noise, and enhance interpretability. Collectively, these steps increase model robustness across different battery cells.

Pearson analysis identifies C_2H_4, CH_4, and C_2H_6 as the most statistically coupled species, prompting our selection of CO, CO_2, and C_2H_4 as the feature set. This choice maximizes information content while minimizing redundancy, making it ideal for unsupervised degradation stage classification (see Supporting Information for detailed selection criteria). Figure 2b shows the normalized discharge capacity versus cycle number. The trajectory separates into two stages: a stable linear stage (blue), with gradual capacity loss, and a nonlinear stage (red), with rapid decline. Notably, the selected gases exhibit marked increases during the nonlinear stage, underscoring their relevance as degradation indicators.[27]

The gas evolution data are high-dimensional, with each cycle generating approximately 200 points per gas. To enhance computational efficiency while retaining degradation-relevant variance, we employ a dimensionality reduction by extracting mechanistically critical segments from each gas profile within a specific SOC window. The optimal SOC window is 70−100% during charge and 100−80% during discharge (see Supporting Information for detailed selection criteria). Within this window, the gas data set is constructed for each sample cycle, comprising 50 points × 3 gases feature vector per cycle. This approach ensures that the data set is both informative and computationally manageable for subsequent machine learning tasks.

The final processed data set for training and evaluation comprised 273 cycles from 10 cells. The training set included Cells 11−1, 14, 15 and 16, each cycled under 0.5C/0.5C charge−discharge conditions, totaling 137 cycles (approximately 50% of the data set). The remaining cells, cycled under varied conditions, comprised the test set (Table S1).

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322

Gas-Based Unsupervised Classification of Degradation Stages

Figure S5 (see Supporting Information) summarizes the unsupervised workflow for per-cycle stage classification using OEMS gas-evolution signals. The workflow begins by acquiring OEMS gas-evolution signals across multiple charge−discharge cycles. The raw time-domain traces are preprocessed to ensure consistency across cycles, including noise reduction and baseline correction, followed by remapping to the SOC domain. The gas features as input for classification are extracted form specific SOC window, yielding 50 data point per gas. Given the high dimensionality of these features for unsupervised classification, PCA is implemented for dimensionality reduction. After PCA, K-means clustering is applied to classify each cycle into two categories as linear or nonlinear.

PCA is applied separately to each gas to reduce dimensionality while preserving variance associated with degradation behavior. Figure 3a, c, and e show heatmaps of the PCA loadings for each gas. These heatmaps represent the weights assigned to each SOC-mapped data point in constructing the first 10 principal components (PCs). Darker colors indicate higher absolute loadings, reflecting greater influence of specific SOC regions on each PC.[31] Figure 3b, d, and f illustrate the corresponding cumulative explained variance ratios, with solid and dashed red lines marking the 80% and 90% variance retention benchmarks, respectively.

Cumulative explained variance quantifies the fraction of total data set variability captured by successive principal components, serving as a critical metric for balancing dimensionality reduction against information retention. The selection of retained PC for each gas is based on the silhouette coefficient of K-means clustering (see Supporting Information for detailed selection criteria). Higher-dimensional feature spaces degrade K-means silhouette coefficients, whereas retaining more PCs monotonically increases cumulative explained variance, necessitating a principled trade-off. Different number of PC combinations were rigorously evaluated via 5-fold cross-validation on the training set. The optimal configuration was identified as that minimizing the composite loss relative to maximum achievable variance retention and silhouette coefficient, respectively. This quantitative optimization converged on a balanced solution requiring 3 PCs for CO, 5 PCs for CO_2, and 10 PCs for C_2H_4.

Figure 4 illustrates 3 key PCs for each gas. In each subplot, the heatmaps (panels a, c, e) show the PCA loading values, where darker red and blue regions indicate SOC intervals with strong positive or negative correlations to the retained PCs. These high-weight regions frequently align with peaks, valleys, or inflection points in the normalized gas profiles (panels b, d, f), confirming that the dimensionality reduction process preserves degradation-relevant signal characteristics.

Notably, some SOC regions are assigned high importance in the PC loadings despite lacking visually prominent features in the gas evolution curves. This suggests that the PCA transformation captures latent structural variations, such as subtle shifts in curve shape or cross-cycle trends that may not be detectable through direct visual inspection. These regions exemplify how PCA can extract meaningful structure from gas evolution profiles, revealing patterns that support more robust classification.

The final 18-dimensional feature vectors, comprising selected principal components from each gas, were clustered using the K-means algorithm with the number of clusters fixed at $K = 2$, corresponding to the two degradation stages. Figure 5 visualizes the training samples (137 cycles) projected onto a three-dimensional subspace defined by the first principal component (PC_1) from each gas. The resulting clusters represent stable linear degradation (blue) and accelerated nonlinear degradation (red), with cluster centroids indicated by star-shaped markers. The clear spatial separation between the two groups indicates that the extracted gas features capture distinct degradation patterns, enabling effective unsupervised classification.

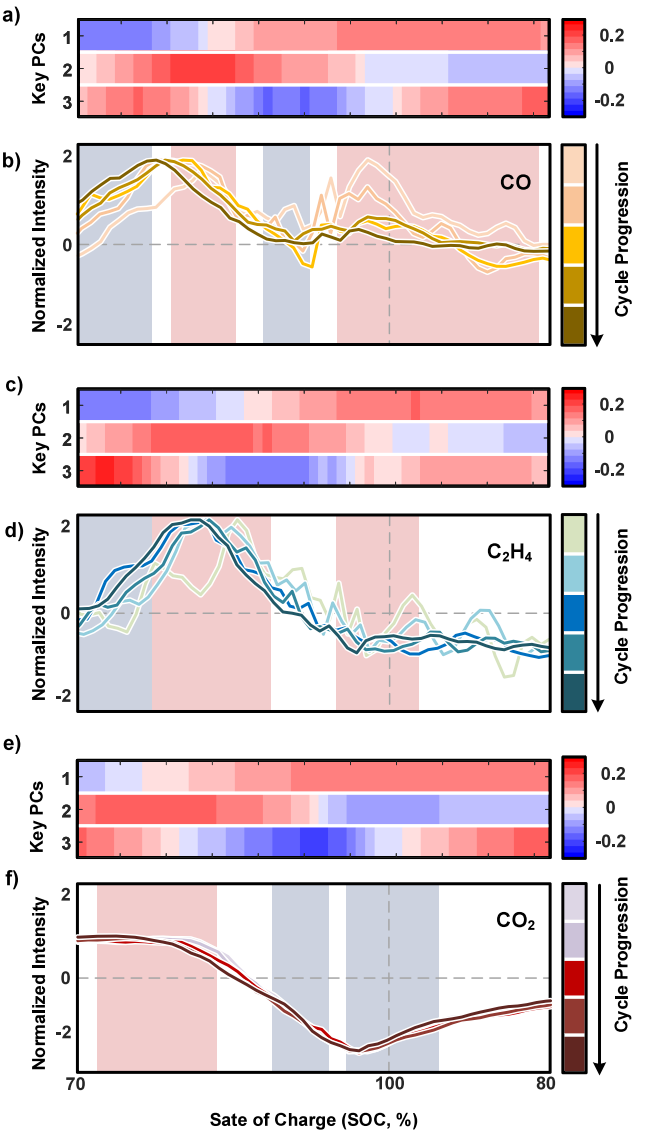

Figure 4. Alignment of PCA loadings with normalized gas profiles across different cycles from a representative cell. (a, c, e) Heatmaps of selected PCs for CO, C_2H_4, and CO_2, respectively, as defined in Figure 3. (b, d, f) Corresponding normalized gas evolution profiles across cycles, color-coded by cycle progression. The heatmaps show PCA loading values resolved along the SOC axis, where darker red and blue regions indicate strong positive or negative correlations. High-weight regions often align with peaks, valleys, or inflection points in the gas profiles, confirming that PCA preserves degradation-relevant signal structure. Additionally, some regions are emphasized in the PC loadings despite lacking visually prominent features, suggesting that PCA captures latent variations not readily apparent through direct inspection.

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322

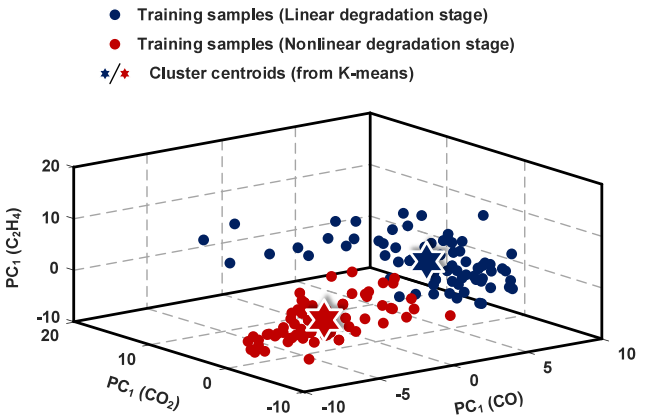

Figure 5. K-means clustering in a reduced PCA feature space. Unsupervised clustering of 137 training samples was performed in an 18-dimensional feature space composed of selected principal components derived from the gas evolution data of CO, C_2H_4, and CO_2. For visualization, the samples are projected into a three-dimensional subspace defined by the first principal component (PC_1) of each gas. The K-means algorithm ($K = 2$) identifies two clusters corresponding to stable linear degradation (blue) and accelerated nonlinear degradation (red), with cluster centroids indicated by star-shaped markers. The clear spatial separation between the clusters demonstrates that the reduced PCA feature space effectively captures degradation-related differences in the gas evolution behavior.

To validate the stage-classification model, gas data from six additional cells were preprocessed, projected into the same PCA space fitted on the training set, and classified based on Euclidean distance to the trained K-means centroids. Each test cycle was assigned to its nearest centroid, enabling stage identification from single-cycle data without a complete cycling history. Unlike curve-fitting methods that rely on full capacity trajectories, this ML framework uses single-cycle gas-evolution signals to classify cycles into linear or nonlinear stages and detect transitions between them. This per-cycle resolution offers a practical route to early stage diagnostics in real-world battery systems, where full degradation histories are often unavailable.

To assess the validity of the unsupervised classification, Figure 6 compares the predicted degradation stages (colored markers) with the capacity retention curves of four representative training cells. Red dashed lines indicate the mathematically determined knee points, based on the Bacon–Watts piecewise linear regression model. In all cases, the unsupervised clustering identifies the onset of accelerated degradation earlier than the fitted knee point. This anticipatory detection suggests that gas-based unsupervised learning can serve as an effective early warning tool, particularly in scenarios where complete lifetime cycling data are unavailable or truncated. The clustering outcomes in Figure 6 further demonstrate that the two degradation stages are well separated, with minimal overlap or ambiguity.

Validation of Degradation Stage Classification Model

In the validation phase, the PCA transformation matrix and K-means cluster centroids obtained from the training set were applied to the held-out test data without refitting. The test set was designed to evaluate generalizability under operating-condition shifts relative to training. Two test cells contained 2 wt % vinylene carbonate (VC) as an electrolyte additive, and one test cell was cycled at a higher charge/discharge rate of 1C/0.5C. In contrast, the training set comprised four cells cycled at 0.5C/0.5C without VC additive. Additional protocol details across cells are summarized in Table S1, and the potential influences of VC and C-rate on gas evolution are discussed in the Supporting Information. Together, these deliberate test-side variations provide a stringent assessment of robustness under differing operational conditions.

A total of 136 OEMS-measured test cycles, spanning both stable linear and accelerated nonlinear degradation regimes,

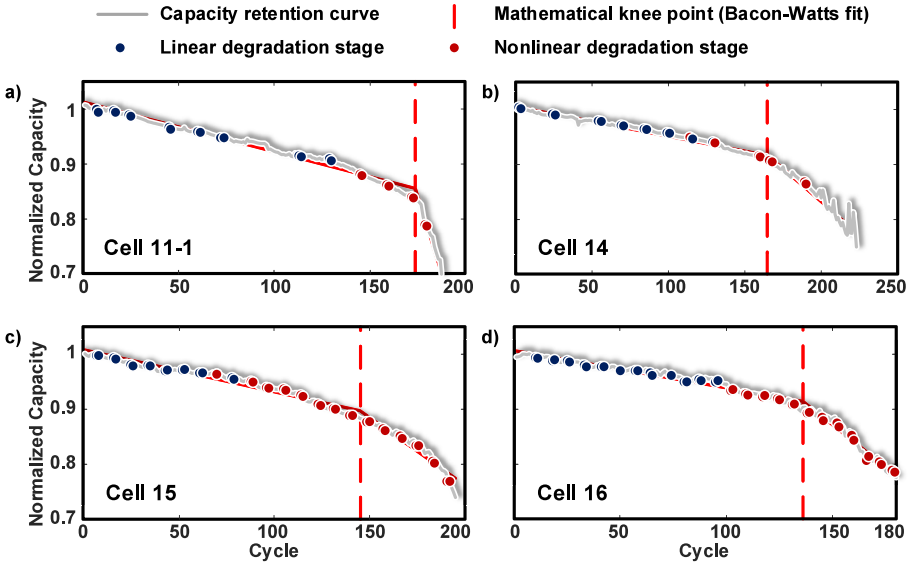

Figure 6. Predicted degradation stages along capacity trajectories for the four training-set cells: (a) Cell 11-1, (b) Cell 14, (c) Cell 15, and (d) Cell 16. Capacity-retention curves are shown as gray traces, and each OEMS-measured cycle is color-coded by its assigned degradation stage based on gas-profile-derived K-means clustering (blue, stable linear stage; red, accelerated nonlinear stage). Red dashed vertical lines indicate the mathematically estimated knee points obtained from Bacon–Watts piecewise-linear fits to the capacity data. In all cases, the unsupervised gas-based classification identifies the onset of accelerated degradation earlier than the fitted knee point, highlighting its potential as an early indicator of rollover without requiring full-cycle capacity data.

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311–26322

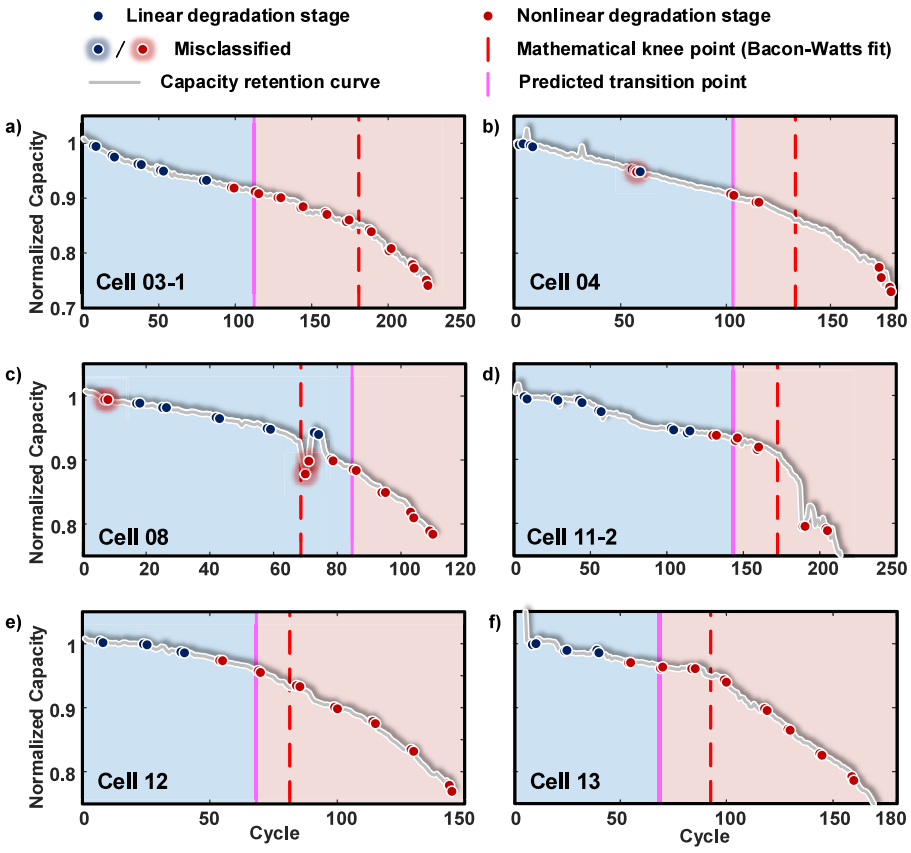

Figure 7. Predicted degradation stages for OEMS-measured cycles across the six held-out test cells: (a) Cell 03-1, (b) Cell 04, (c) Cell 08, (d) Cell 11-2, (e) Cell 12, and (f) Cell 13. Capacity-retention curves are shown as gray traces, and each OEMS-measured cycle is color-coded by gas-profile classification as stable linear degradation (blue dots), accelerated nonlinear degradation (red dots), or misclassified (dots with blurred outlines). The vertical purple line marks the model-predicted transition to rollover. The transition is declared when two consecutive OEMS tests produce nonlinear warnings. Each OEMS test consists of two adjacent OEMS-measured cycles, and a nonlinear warning is emitted if either of the two cycles is classified as nonlinear. Despite variations in charge/discharge rates and electrolyte composition, the model consistently anticipates the onset of nonlinear degradation. In Cell 08 (panel c), a temporary temperature drop near cycle 70 caused a transient capacity dip unrelated to degradation. While conventional methods may misidentify this disturbance as a knee point, the model detects the true transition at cycle 85, with only a single-test misclassification, demonstrating robustness to external disturbances. Label-free monitoring stability metrics derived from these predicted warning sequences are summarized in Table S3.

were classified into the two predefined clusters, enabling stage identification at the OEMS-measured cycle level. To better reflect real-world deployment scenarios, the testing data were processed sequentially in chronological order. Cycles without OEMS measurements are not assigned stage labels in this study. In deployment, the stage estimate is updated only when a new OEMS measurement is available, and the sequential warning rule below provides robustness under intermittent sampling. To reduce sensitivity to single-cycle variance, each OEMS test consists of two adjacent OEMS-measured cycles. If either cycle is classified as nonlinear, the test emits a nonlinear warning. The transition point is identified when two consecutive tests contain nonlinear degradation warnings. Shortening the test interval could enable earlier identification of the transition point but may also increase susceptibility to local variances, such as environmental temperature or operational condition fluctuations.

Figure 7 summarizes the degradation-stage predictions for six held-out test cells, with the model-predicted onset of nonlinear degradation indicated by vertical purple lines. Only OEMS-measured cycles are classified and plotted. Points are color-coded as stable linear (blue), accelerated nonlinear (red), and misclassifications (blurred outlines). Across the 136 OEMS-measured test samples, the misclassification rate is 3.6% (5/136). In all cases, the framework predicts the transition to nonlinear degradation before observable rollover in capacity, demonstrating predictive capability across varied cycling conditions. Notably, the sequential transition rule improves robustness against transient disturbances unrelated to degradation. For example, Cell 08 exhibits a temporary capacity drop around cycle 70 due to a short-term temperature fluctuation. While conventional curve-fitting methods may misidentify this disturbance as a knee, the proposed rule requires two consecutive tests with nonlinear warnings and therefore does not trigger a false transition. As a result, the model identifies the true onset of sustained nonlinear behavior at cycle 85, despite a single-test (two-cycle) misclassification (Figure 7c). This highlights the model's ability to distinguish genuine mechanistic degradation from incidental anomalies.

Beyond the point-wise inconsistency count, we quantified label-free monitoring stability under the same sequential protocol. Table S3 shows low within-test disagreement (2/68 tests), rare transient pretransition warnings (2/27 pretransition tests, occurring only as isolated single-test episodes), and no post-transition reversions (0/35; 100% persistence after transition declaration), supporting stable transition monitoring

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322

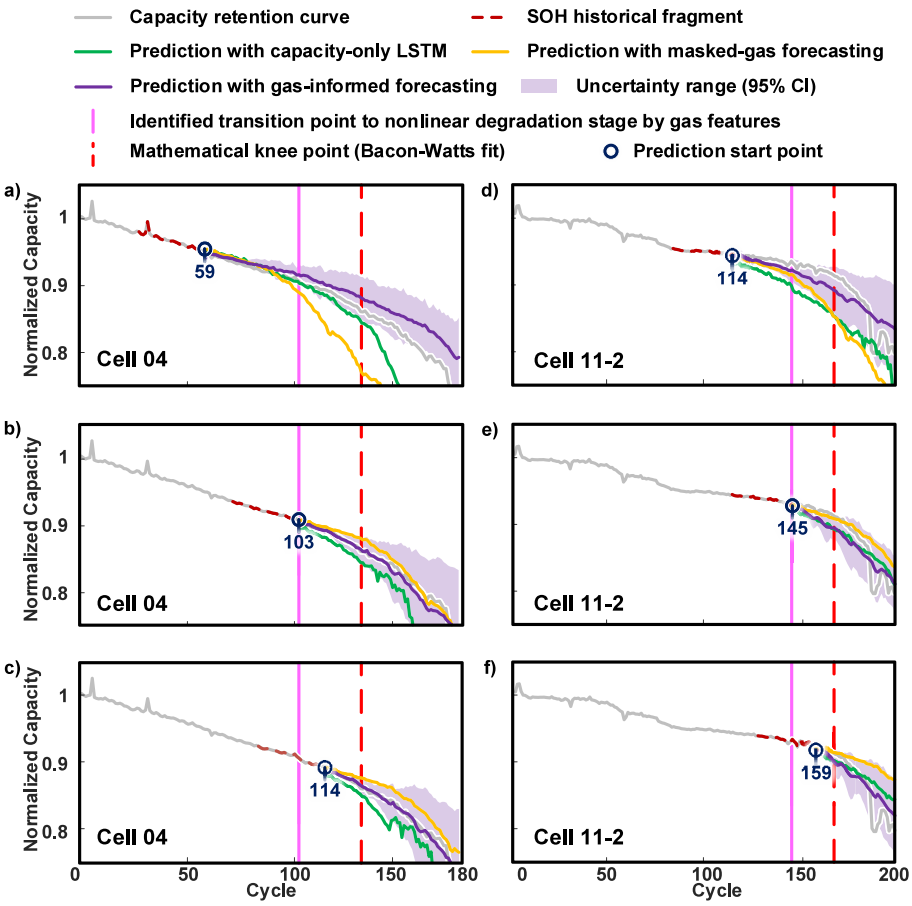

Figure 8. Gas-informed single-shot capacity-trajectory forecasting on held-out Cells 04 (a–c) and 11-2 (d–f), compared with a capacity-only LSTM roll-out baseline and a masked-gas ablation. Forecasts are initialized at three points along the degradation trajectory: (a, d) early in the linear regime, (b, e) near the gas-identified transition to nonlinear degradation (pink vertical line), and (c, f) near the Bacon–Watts knee (red dashed vertical line). The gray curve shows the measured normalized capacity. The dark red dashed segment indicates the 32-cycle SOH history provided as input; the circle marks the prediction start cycle (the end point of the input window and the OEMS-measured cycle supplying the gas snapshot). Predicted trajectories are shown for the gas-informed model (purple), LSTM-only baseline (green), and masked-gas model (yellow); shaded bands denote ensemble uncertainty (mean ± 2σ).

under intermittent OEMS sampling. To further assess robustness beyond this predefined hold-out split, we also performed leave-one-cell-out (LOCO) cross-validation across all 10 cells. The resulting cell-level monitoring metrics are summarized in Table S4 and show similarly stable behavior under the same sequential warning rule, with an overall LOCO misclassification rate of 2.9% (8/273 OEMS-measured cycles). Together, the fixed-split and LOCO results support the robustness of the stage-classification framework at the cell level.

Gas-Informed Capacity Trajectory Forecasting

The unsupervised PCA-based classification above shows that characteristic cycle-resolved temporal patterns within each gas profile are sufficient to distinguish linear from nonlinear degradation regimes. This indicates that gas-evolution features can serve as chemically specific fingerprints of impending rollover. However, PCA operates on normalized fluctuations and therefore discards absolute concentration levels and interspecies stoichiometric relationships, which can provide information critical for quantitative capacity-loss forecasting. To capture these additional degrees of freedom, we develop a supervised, gas-informed forecaster that predicts multihorizon capacity trajectories by conditioning a non-AR (single-shot) model on single-cycle gas features together with historical capacity/SOH trends. In contrast to AR designs, which predict step by step and accumulate error over time, our model outputs all future horizons in one pass, reducing error propagation.

Figure S10 (see Supporting Information) illustrates the end-to-end workflow of the gas-informed forecasting framework. Each input sample comprises a 32-cycle window of SOH fragment ending at an OEMS-instrumented cycle. The corresponding CO, CO_2, and C_2H_4 traces are baseline corrected and mapped to the SOC domain. The processed gas data are passed to a feature-extraction network without normalization to preserve absolute magnitudes and interspecies ratios. The network contains two 1D convolutional layers that operate jointly across the three gas channels, followed by a channel-wise attention module. The convolutions learn temporal motifs within and across gases, and the attention module reweights channels so that the most degradation-sensitive signals contribute more strongly. The resulting representation is flattened into a single gas feature vector.

In parallel, the SOH fragment is encoded with a Bi-LSTM encoder to capture temporal patterns in capacity fade. A multihead cross-attention block uses the projection of gas feature vector as queries and Bi-LSTM hidden states as keys and values, allowing each time step in the SOH history to

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311–26322

consult the chemical context at the end of the fragment. The attention output is fused back into the SOH representation through a residual connection, yielding a feature tensor that encodes both electrical history and chemically informed cues. A time-aware decoder then maps this fused tensor to a multihorizon forecast, producing per-step SOH changes for a user-specified future horizon in a single shot rather than by rolling an AR loop. The decoder outputs two learnable horizon-indexed coefficients which, together with the starting capacity, generate the sequence of per-step capacity decrements without autoregressive recursion.

Figure 8 illustrates the forecasting performance of two representative test cells (Cell 04 and Cell 11-2). These cells were selected as the testing set for forecasting task because they exhibit the minimal configuration shift from the training set (Cells 11-1, 14, 15, 16), isolating "generalization to new cells" from extrapolation to unseen operating regimes. For each cell, predictions with gas-informed model (purple curves) are initiated from multiple points spanning the degradation trajectory (with 95% confidence intervals shown as shaded regions).

The prediction performance of the gas-informed model is compared to baseline LSTM-only (green) and masked-gas (yellow) models. The baseline of a 4-layer LSTM is trained on SOH fragments, with long-horizon predictions generated by iterative roll-out beyond the 32-step output window. The masked-gas model denotes the full architecture with gas inputs zero-masked, isolating the contribution of historical SOH data alone.

The gas-informed model (purple curves) delivers the most reliable long-horizon forecasts when the prediction start point lies in the late-linear regime or near the gas-identified transition. Forecasts initialized farther from the transition are intrinsically more uncertain because the capacity history is still nearly linear, the gas precursors are weaker, and the remaining forecast horizon is longer. Even so, on held-out cells closest to the training protocol (Cells 04 and 11−2), it achieves consistently low error at 100-cycle horizons (e.g., RMSE ≈ 1.7−2.9% at 100-cycle horizons) and shows reduced long-horizon drift relative to capacity-only baselines. In contrast, the capacity-only LSTM baseline performs strongly on short horizons when the start point is far from rollover, but its iterative roll-out tends to accumulate error over longer horizons and can remain biased toward extrapolating the quasi-linear pretransition trend. The masked-gas ablation further degrades long-horizon performance in many cases, indicating that the SOC-resolved gas snapshot provides nonredundant chemical context beyond capacity history alone. Accordingly, the advantage of gas conditioning is more evident at longer horizons and for start points closer to the transition, whereas very early short-horizon predictions can remain competitive with capacity-only baselines.

Under protocol shifts (e.g., VC additive and electrolyte-volume variation), forecasting accuracy becomes more variable, and some early stage short-horizon cases can be competitive with capacity-only baselines, indicating that robust cross-condition forecasting will require broader multicondition training data. Comprehensive performance metrics across test cells, prediction start points, and horizons are provided in Table S6.

For the two representative held-out cells shown in Figure 8 (Cells 04 and 11-2), we additionally evaluated a capacity-only TFT forecaster using the same data split, 32-cycle SOH-history input, prediction start points, and forecast horizons; the resulting case-by-case comparison is summarized in Table S7. This representative advanced-baseline comparison shows that advanced capacity-history-only architectures can be competitive for some start-point/horizon combinations. To further probe robustness beyond these representative cases, we also performed a subset-based leave-one-cell-out sensitivity analysis of the gas-informed forecasting model, summarized in Table S8. This analysis shows that forecasting performance is most stable when the held-out cell remains close to the training domain, whereas errors increase under stronger configuration shifts, particularly electrolyte-volume variation. Accordingly, the benefit of gas conditioning should be interpreted as start-point-, horizon-, and cell-dependent rather than universal.

These results support the interpretation that a single-cycle gas profile captures chemical signals that precede observable capacity curvature, enabling a nonautoregressive, multihorizon forecast without requiring postrollover capacity data in the input window. The additional sensitivity analysis in Table S8 further indicates that the observed forecasting variability is not limited to a single predefined split, but reflects a broader dependence on configuration shift and training-domain mismatch. At the same time, the observed variability under protocol shifts highlights an important scope limitation of the present data set and motivates future expansion toward multicondition training for more uniform cross-protocol forecasting. Nevertheless, the framework demonstrates a practical path to long-horizon forecasting from minimal inputs, requiring only a short SOH history and a single gas measurement.

DISCUSSION

The performance of the proposed framework is rooted in using SOC-resolved gas evolution as a chemically specific, cycle-level readout of parasitic reaction activity that is only indirectly reflected in electrical signals. The selected species (CO, CO_2, and C_2H_4) are well-known products of electrolyte/interphase reactions and provide complementary information about evolving degradation pathways.[32] For example, C_2H_4 is commonly associated with reductive ethylene carbonate (EC) decomposition and concurrent formation of lithium ethylene dicarbonate (LEDC), a major organic SEI component. CO_2 and C_2H_4 can also be released during subsequent LEDC decomposition, which can be accelerated by trace H_2O/HF and accompanied by inorganic byproducts such as LiF and Li_2CO_3.[33−36] More broadly, CO/CO_2 can arise from multiple carbonate decomposition routes and interfacial processes; thus, their SOC-resolved profiles are best interpreted as chemically grounded indicators of changing reaction regimes rather than signatures of a single unique mechanism.

Within this context, the gas-based stage classifier is effective because it tracks how the timing, shape, and relative amplitude of $CO/CO_2/C_2H_4$ evolve within a single charge−discharge cycle. During the early linear regime, SEI restructuring and electrolyte decomposition proceed comparatively slowly, and gas evolution remains modest. As degradation accelerates, electrolyte reduction/oxidation and interphase reconstruction intensify, increasing solvent consumption and gas generation while electrical indicators may still appear close to linear. The resulting SOC-resolved gas profile changes provide a distinct chemical fingerprint of this regime shift. By operating on SOC-aligned profiles and applying a label-free PCA + K-means pipeline, the stage classifier emphasizes robust, cycle-resolved

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322

morphological evolution rather than relying on densely sampled capacity curvature or retrospective knee fitting.

Gas information is also valuable for forecasting because it provides precursor chemical context that capacity histories alone often lack in the linear regime. Capacity-only baselines tend to extrapolate the near-linear trend and may miss the impending acceleration until curvature becomes pronounced. By conditioning the forecaster on a single OEMS-measured cycle together with a short SOH history, the model can incorporate SOC-resolved reaction signatures that begin to evolve before rollover becomes visible in the capacity trajectory. In addition, the single-shot, multihorizon design avoids autoregressive roll-out and therefore reduces long-horizon drift caused by compounding step-ahead errors. This advantage is most evident when forecasts are initialized in the late-linear or near-transition regime; when predictions are started substantially earlier, the gas precursors are weaker and the remaining prediction horizon is longer, so the forecast uncertainty is correspondingly larger.

At the same time, the present results should be interpreted within a limited scope. The robustness demonstrated here is confined to within-platform shifts in electrolyte volume, VC additive, charge/discharge rate, and CV-hold protocol within a single NMC811/graphite 21700 cell family. Stage classification remains comparatively robust across these shifts, whereas forecasting performance is strongest on cells closest to the training configuration and becomes more variable under larger protocol or configuration changes. The current trained model, selected gas subset, and optimized SOC window should therefore be regarded as platform-specific rather than universally transferable. From an application standpoint, OEMS is used here as a research-grade sensing modality to establish the diagnostic value of gas information, rather than as a claim of immediate onboard deployment. The framework is nevertheless data-efficient because it does not require manually curated stage labels for monitoring and reduces reliance on long historical capacity trajectories by leveraging a chemically informative gas snapshot. Translation to embedded sensing will require validation under reduced selectivity and resolution, sensor noise/drift, and cross-sensitivity, together with broader multicondition data sets to establish generality across protocols, chemistries, and cell formats.

CONCLUSION

In this study, we demonstrate that SOC-resolved gas evolution provides a chemically specific signal that can be leveraged for battery health monitoring and prognostics beyond what is accessible from electrical trends alone. We present a gas-informed ML framework with two components: (i) a label-free stage classifier that uses a single OEMS-measured cycle to identify the onset of accelerated nonlinear degradation under intermittent sampling, and (ii) a gas-conditioned, single-shot forecaster that fuses one OEMS-measured cycle with a short SOH history to predict multihorizon capacity trajectories without autoregressive rollout. Together, the results show that incorporating chemically grounded gas information enables earlier transition warnings and more stable long-horizon forecasting from minimal historical input.

This work assumes access to OEMS-quality, SOC-resolved gas measurements acquired intermittently during cycling. Translating the approach to embedded sensing will require robustness validation under sensor-limited resolution, noise/drift, and cross-sensitivity, as well as expanded multicondition data sets to establish generality across protocols and chemistries. As compact off-gas/pressure sensing matures, gas-conditioned models offer a route toward chemistry-aware battery health monitoring and earlier, more actionable degradation prognostics.

ASSOCIATED CONTENT

Data Availability Statement

The code and processed data underlying this study are openly available in GitHub at https://github.com/YiqingLU-ShanghaiTech/Gas-Informed-Machine-Learning-Framework-for-Battery-Degradation-Analysis.

Supporting Information

The Supporting Information is available free of charge at https://pubs.acs.org/doi/10.1021/acsami.6c02666.

OEMS signal calibration and preprocessing; rationale for gas species and SOC-window selection; workflows for label-free degradation-stage classification and gas-informed capacity-trajectory forecasting (PDF)

AUTHOR INFORMATION

Corresponding Authors

Franklin Kim − *School of Physical Science and Technology, ShanghaiTech University, Shanghai 201210, China*; orcid.org/0000-0002-6548-6588; Email: fkim@shanghaitech.edu.cn

Haoyu Wang − *School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China*; Email: wanghy@shanghaitech.edu.cn

Authors

Yiqing Lu − *School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China*

Jianli Zou − *School of Physical Science and Technology, ShanghaiTech University, Shanghai 201210, China*

Liang Zhang − *School of Information Science and Technology, ShanghaiTech University, Shanghai 201210, China*

Complete contact information is available at: https://pubs.acs.org/10.1021/acsami.6c02666

Author Contributions

#Y.L. and J.Z. contributed equally to this paper.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work was supported by the research fund from ShanghaiTech University (Grant No. 2017F0201-000-11) and the National Natural Science Foundation of China (Grant Nos. 52077140 to H.W. and 22275123 to F.K.).

REFERENCES

(1) Harper, G.; Sommerville, R.; Kendrick, E.; Driscoll, L.; Slater, P.; Stolkin, R.; Walton, A.; Christensen, P.; Heidrich, O.; Lambert, S. Recycling lithium-ion batteries from electric vehicles. *Nature* **2019**, *575* (7781), 75−86.

(2) Grey, C. P.; Hall, D. S. Prospects for lithium-ion batteries and beyond—a 2030 vision. *Nat. Commun.* **2020**, *11* (1), 1−4.

(3) Yang, Y.; Okonkwo, E. G.; Huang, G.; Xu, S.; Sun, W.; He, Y. On the sustainability of lithium ion battery industry-A review and perspective. *Energy Storage Mater.* **2021**, *36*, 186−212.

(4) Schuster, S. F.; Bach, T.; Fleder, E.; Müller, J.; Brand, M.; Sextl, G.; Jossen, A. Nonlinear aging characteristics of lithium-ion cells under different operational conditions. *J. Energy Storage* **2015**, *1*, 44−53.
(5) Fang, R.; Dong, P.; Ge, H.; Fu, J.; Li, Z.; Zhang, J. Capacity plunge of lithium-ion batteries induced by electrolyte drying-out: Experimental and Modeling Study. *J. Energy Storage* **2021**, *42*, 103013.
(6) Attia, P. M.; Bills, A.; Planella, F. B.; Dechent, P.; Dos Reis, G.; Dubarry, M.; Gasper, P.; Gilchrist, R.; Greenbank, S.; Howey, D. "Knees" in lithium-ion battery aging trajectories. *J. Electrochem. Soc.* **2022**, *169* (6), 060517.
(7) Zhang, H.; Altaf, F.; Wik, T. Battery capacity knee-onset identification and early prediction using degradation curvature. *J. Power Sources* **2024**, *608*, 234619.
(8) Zhang, J.; Zhang, Y.; Fu, J.; Zhao, D.; Liu, P.; Zhang, Z. Capacity fading knee-point recognition method and life prediction for lithium-ion batteries using segmented capacity degradation model. *Reliab. Eng. Syst. Saf.* **2024**, *251*, 110395.
(9) Haris, M.; Hasan, M. N.; Qin, S. Degradation curve prediction of lithium-ion batteries based on knee point detection algorithm and convolutional neural network. *IEEE Trans. Instrum. Meas.* **2022**, *71*, 1−10.
(10) Zhang, C.; Wang, Y.; Gao, Y.; Wang, F.; Mu, B.; Zhang, W. Accelerated fading recognition for lithium-ion batteries with Nickel-Cobalt-Manganese cathode using quantile regression method. *Appl. Energy* **2019**, *256*, 113841.
(11) Baptista, M. L.; Henriques, E. M.; Goebel, K. More effective prognostics with elbow point detection and deep learning. *Mech. Syst. Signal Process.* **2021**, *146*, 106987.
(12) Fermín-Cueto, P.; McTurk, E.; Allerhand, M.; Medina-Lopez, E.; Anjos, M. F.; Sylvester, J.; Dos Reis, G. Identification and machine learning prediction of knee-point and knee-onset in capacity degradation curves of lithium-ion cells. *Energy AI* **2020**, *1*, 100006.
(13) Sohn, S.; Byun, H.-E.; Lee, J. H. Two-stage deep learning for online prediction of knee-point in Li-ion battery capacity degradation. *Appl. Energy* **2022**, *328*, 120204.
(14) Costa, N.; Anseán, D.; Dubarry, M.; Sánchez, L. ICFormer: A Deep Learning model for informed lithium-ion battery diagnosis and early knee detection. *J. Power Sources* **2024**, *592*, 233910.
(15) Zhang, Y.; Zhao, M. Cloud-based in-situ battery life prediction and classification using machine learning. *Energy Storage Mater.* **2023**, *57*, 346−359.
(16) Tan, R.; Lu, X.; Cheng, M.; Li, J.; Huang, J.; Zhang, T.-Y. Forecasting battery degradation trajectory under domain shift with domain generalization. *Energy Storage Mater.* **2024**, *72*, 103725.
(17) Che, Y.; Lam, V. N.; Rhyu, J.; Schaeffer, J.; Kim, M.; Bazant, M. Z.; Chueh, W. C.; Braatz, R. D. Diagnostic-free onboard battery health assessment. *Joule* **2025**, *9*, 102010.
(18) Zhou, Z.; Liu, Y.; You, M.; Xiong, R.; Zhou, X. Two-stage aging trajectory prediction of LFP lithium-ion battery based on transfer learning with the cycle life prediction. *Green Energy Intell. Transp.* **2022**, *1* (1), 100008.
(19) Gomez, W.; Wang, F.-K.; Chou, J.-H. Li-ion battery capacity prediction using improved temporal fusion transformer model. *Energy* **2024**, *296*, 131114.
(20) Zhang, Y.; Tang, Q.; Zhang, Y.; Wang, J.; Stimming, U.; Lee, A. A. Identifying degradation patterns of lithium ion batteries from impedance spectroscopy using machine learning. *Nat. Commun.* **2020**, *11* (1), 1706.
(21) Jung, R.; Metzger, M.; Maglia, F.; Stinner, C.; Gasteiger, H. A. Oxygen release and its effect on the cycling stability of LiNixMnyCozO2 (NMC) cathode materials for Li-ion batteries. *J. Electrochem. Soc.* **2017**, *164* (7), A1361.
(22) Fang, C.; Tran, T.-N.; Zhao, Y.; Liu, G. Electrolyte decomposition and solid electrolyte interphase revealed by mass spectrometry. *Electrochim. Acta* **2021**, *399*, 139362.
(23) Christensen, J.; Newman, J. Cyclable lithium and capacity loss in Li-ion cells. *J. Electrochem. Soc.* **2005**, *152* (4), A818.
(24) Zhang, Q.; White, R. E. Capacity fade analysis of a lithium ion cell. *J. Power Sources* **2008**, *179* (2), 793−798.
(25) Pinson, M. B.; Bazant, M. Z. Theory of SEI formation in rechargeable batteries: capacity fade, accelerated aging and lifetime prediction. *J. Electrochem. Soc.* **2013**, *160* (2), A243.
(26) Metzger, M.; Gasteiger, H. A. Diagnosing battery degradation via gas analysis. *Energy Environ. Mater.* **2022**, *5* (3), 688−692.
(27) Jiang, Z.; Zhang, L.; Yu, Y.; Zou, J.; Kim, F. The Resurging of Hydrocarbon Gas as Early Sign of Battery Rollover Degradation. *ACS Appl. Mater. Interfaces.* **2025**, *17* (3), 4934−4940.
(28) Mattinen, U.; Klett, M.; Lindbergh, G.; Lindström, R. W. Gas evolution in commercial Li-ion battery cells measured by on-line mass spectrometry-Effects of C-rate and cell voltage. *J. Power Sources* **2020**, *477*, 228968.
(29) Shitaw, K. N.; Yang, S.-C.; Jiang, S.-K.; Huang, C.-J.; Sahalie, N. A.; Nikodimos, Y.; Weldeyohannes, H. H.; Wang, C.-H.; Wu, S.-H.; Su, W.-N.; Hwang, B. J. Decoupling Interfacial Reactions at Anode and Cathode by Combining Online Electrochemical Mass Spectroscopy with Anode-Free Li-Metal Battery. *Adv. Funct. Mater.* **2021**, *31* (6), 2006951.
(30) Paulus, B. C.; Banh, J. K.; Rector, K. D.; Stein, B. W.; Lilley, L. M. Whispering gallery mode resonators in continuous flow: spectral assignments and sensing with monodisperse microspheres. *Anal. Methods* **2022**, *14* (17), 1690−1697.
(31) Wang, J.; Zhang, C.; Zhang, L.; Su, X.; Zhang, W.; Li, X.; Du, J. A novel aging characteristics-based feature engineering for battery state of health estimation. *Energy* **2023**, *273*, 127169.
(32) Metzger, M.; Strehle, B.; Solchenbach, S.; Gasteiger, H. A. Origin of H2 evolution in LIBs: H2O reduction vs. electrolyte oxidation. *J. Electrochem. Soc.* **2016**, *163* (5), A798.
(33) Hobold, G. M.; Khurram, A.; Gallant, B. M. Operando gas monitoring of solid electrolyte interphase reactions on lithium. *Chem. Mater.* **2020**, *32* (6), 2341−2352.
(34) Heiskanen, S. K.; Kim, J.; Lucht, B. L. Generation and evolution of the solid electrolyte interphase of lithium-ion batteries. *Joule* **2019**, *3* (10), 2322−2333.
(35) Liu, T.; Lin, L.; Bi, X.; Tian, L.; Yang, K.; Liu, J.; Li, M.; Chen, Z.; Lu, J.; Amine, K. In situ quantification of interphasial chemistry in Li-ion battery. *Nat. Nanotechnol.* **2019**, *14* (1), 50−56.
(36) Song, Y.; Wang, L.; Sheng, L.; Ren, D.; Liang, H.; Li, Y.; Wang, A.; Zhang, H.; Xu, H.; He, X. The significance of mitigating crosstalk in lithium-ion batteries: a review. *Energy Environ. Sci.* **2023**, *16* (5), 1943−1963.

https://doi.org/10.1021/acsami.6c02666
ACS Appl. Mater. Interfaces 2026, 18, 26311−26322