

Interpretable Battery Aging without Extra Tests via Neural-Assisted Physics-based Modelling

Yuan Qiu¹, Wei Li¹, Wei Zhang^{1,*}, Yi Zhou¹, Fang Liu², Jianbiao Wang³, Zhi Wei Seh³

¹*Singapore Institute of Technology (SIT), Singapore, 828608*

²*Singapore University of Social Sciences (SUSS), Singapore, 599494*

³*Institute of Materials Research and Engineering (IMRE), Agency for Science, Technology and Research (A*STAR),*

2 Fusionopolis Way, Innovis #08-03, Singapore 138634, Republic of Singapore

{yuan.qiu, wei.li, wei.zhang, zhouyi}@singaporetech.edu.sg, liufang@suss.edu.sg, {wang_jianbiao, sehzw}@a-star.edu.sg

Abstract—State of health (SoH) is widely used for battery management, but it is a single scalar and offers limited interpretability. Two batteries with similar SoH can exhibit very different degradation behaviors and the lack of interpretability hinders optimal battery operation. In this paper, we propose IBAM for interpretable battery aging modelling with a neural-assisted physics-based framework. IBAM outputs a 2-D aging fingerprint without extra diagnostic tests and uses only routine logs from the battery management system. The fingerprint offers great interpretability by capturing a battery’s curve-wide polarization voltage loss and the tail loss near the end-of-discharge. IBAM first creates a physics-based battery model based on a fractional-order equivalent-circuit model, and then extracts per-cycle fingerprints from the model using a two-stage least-squares method. IBAM further anchors fingerprints on the SoH axis with physics-guided regression, where the per-cycle SoH is estimated via a bidirectional gated recurrent unit with customized multi-channel voltage features. Across batteries with short-, medium-, and long-lifespans, IBAM consistently yields the best physics-model fidelity at different aging stages, and provides clear interpretations of degradation mechanisms and fingerprint patterns about batteries of different lifespans. The resulting fingerprints support interpretable battery health assessment and can inform battery control choices.

Index Terms—Battery health, state of health, machine learning, industrial artificial intelligence, interpretability, explainability

I. INTRODUCTION

Electric vehicles (EVs) are becoming a mainstream transport option. The International Energy Agency (IEA) reports that EV sales exceeded 17 million globally in 2024, reaching a sales share of over 20% of new cars worldwide [1]. The battery is one of the most important and expensive components in an EV. Its performance and degradation therefore become practical concerns, as battery aging directly affects range, power capability, charging limits, etc. Battery operation is managed by the battery management system (BMS), where health monitoring is foundational. A widely used health indicator is the state of health (SoH), which is commonly defined as the ratio of the current usable capacity to the rated beginning-of-life capacity. However, SoH is not directly observable during

operation and remains challenging to estimate and forecast accurately and robustly under diverse battery usage conditions.

Most battery health studies focus primarily on the accuracy of SoH modelling and prediction. A common trend is to adopt data-driven machine learning (ML) algorithms, with recent efforts increasingly favor large and complex ML models to maximize accuracy. Basic ML models are often considered in early studies, e.g., regression in [2] and ensemble learning in [3]. Hybrid ML architectures combine complementary backbones, e.g., Transformer and bidirectional gated recurrent unit (BiGRU) in [4] and Informer and gated recurrent unit in [5]. Recent advancements mainly follow two notable directions. One direction integrates domain knowledge and physics to improve generalization and reduce reliance on purely ML-based modelling, e.g., guiding liquid neural networks with thermal-related entropy in [6] and expanding temporal convolutional networks with the features from the equivalent circuit model (ECM) in [7]. The other direction explores generative AI (GenAI) for battery health modelling, e.g., [8] uses diffusion modelling for estimating remaining useful life and [9] regulates GenAI with battery knowledge for robust SoH prediction. While these approaches can be highly accurate, they typically output a single health scalar and provide limited visibility into degradation mechanisms and the evolution of these mechanisms. Such a lack of explainability and interpretability make it difficult to translate a prediction of battery health into actionable battery management decisions, and motivates interpretable battery analytics.

A small but growing line of work argues that accuracy alone is insufficient and battery health models should also provide explanations that battery operators and users can trust and act on. Most existing studies follow a post-hoc explainable AI paradigm on ML-based models. A unified pipeline that combines post-hoc explainability with available tools such as Shapley additive explanations (SHAP) and local interpretable model-agnostic explanations is proposed in [10] for feature-level interpretation and confidence-aware diagnostics. Several recent studies also extend SHAP-based interpretation from explaining a single health prediction to summarizing degradation behavior over time. For example, SHAP is aggregated across sliding-window sequences in [11] to enable identification of dominant features, and used to interpret multiple ML models

This work was supported in part by A*STAR under its MTC Programmatic (Award M23L9b0052), MTC Individual Research Grants (IRG) (Award M23M6c0113), and MTC Programmatic (Award M24N6b0043).

* Corresponding author. Email: wei.zhang@singaporetech.edu.sg.

in [12]. Another group of studies improves explainability by introducing additional sensing modalities that are directly coupled with batteries' internal electrochemical and mechanical changes. For example, a study extracts ultrasound features as inputs for SoH models and then uses interpretable models and explainability AI tools to visualize the relationship between the sensing features and battery degradation [13]. Overall, current explainable battery studies mainly improve transparency by attaching explanation tools to ML models or by adding extra sensing modalities. The proposed explanations are also not explicitly physics-guided so the explanations may not reflect real degradation behavior.

In this paper, we aim to enable interpretable battery health analytics from routine BMS logs without introducing additional battery diagnostic tests. We propose IBAM (pronounced “eye-bam”), an interpretable battery aging monitoring framework. IBAM models battery operation dynamics with a fractional-order ECM (FOECM), which is physics-based and encodes battery knowledge in a compact set of interpretable parameters. The physics model includes two resistance-like parameters R_{dyn} and R_W . The former captures the broad under-load voltage penalty, where the discharge voltage curve of a battery shifts downward and drops faster as the battery ages. The latter captures the late-stage voltage drop that emerges near the end-of-discharge and can trigger earlier low-voltage cutoff. These two parameters summarize complementary degradation effects in a compact and physically meaningful way, so IBAM defines (R_{dyn}, R_W) as an aging fingerprint for interpreting a battery's health beyond the commonly used SoH. IBAM first fits the model to each discharge cycle using a two-stage least squares method (LSM) to extract per-cycle raw fingerprints associated with two voltage-loss mechanisms. IBAM then estimates cycle-wise SoH using a BiGRU model with multi-channel features extracted from the same discharge voltage trace. Finally, IBAM performs a physics-guided mapping that projects the cycle-wise fingerprints onto the SoH domain via isotonic regression, so that the fingerprints are anchored to the industry-standard health axis. This yields physically consistent SoH-indexed fingerprints that can be queried online to support BMS decision-making. For example, a high R_{dyn} suggests dominant polarization loss and motivates conservative high-rate discharging, while a high R_W indicates pronounced end-of-discharge voltage loss and suggests stricter voltage gating policies. In summary, we have the following main contributions in this paper.

- We formulate interpretable battery health monitoring from routine BMS logs, and interpret SoH using a 2-D aging fingerprint that captures polarization and tail losses.
- We propose IBAM, a physics-based framework that reliably identifies this fingerprint and constructs physically consistent SoH-indexed fingerprints for online querying.
- We validate IBAM on short-, medium-, and long-life batteries. IBAM achieves consistently improved physics-model fidelity and reveals lifespan-aligned fingerprint patterns that yield actionable degradation insights.




Fig. 1. An illustration of IBAM architecture.

The rest of the paper is organized as follows. We present the technical details of IBAM in Section II. In Section III, we conduct an experimental study and present the experiment results. Finally, we conclude this paper in Section IV.

II. IBAM METHODOLOGY

An overview of IBAM's system architecture is in Fig. 1, and technical details are provided in the following subsections.

A. Problem Formulation and Interpretation Targets

This paper adopts a *no-extra-tests* setting. We use only the signals and sensor measurements that are routinely logged by a typical BMS during battery operation. We do not assume additional cycling procedures, diagnostic tests, or sensor installation for practical deployments. We consider a battery observed over n_c charge/discharge cycles, and the BMS log for cycle c is given by,

$$\mathcal{D}_c = \left\{ (t_i, \mathbf{x}_c(t_i)) \right\}_{i=1}^{T_c}, \quad (1)$$

where t_i denotes the i -th timestamp within cycle c and T_c is the number of timestamps in the cycle. The log of each timestamp is recorded as a 2-tuple, e.g., $(t_i, \mathbf{x}_c(t_i))$, where $\mathbf{x}_c(t)$ stacks the BMS-available channels such as terminal voltage, current, and temperature when available. We assume that the logged channels are time-aligned by the BMS, e.g., signals can be resampled or interpolated onto a common time grid. In this study, we primarily rely on voltage trajectory $V_c(t)$, which carries the dominant information about cycle-to-cycle degradation. Other BMS channels in $\mathbf{x}_c(t)$ such as current and temperature are used when available for standard state calculation and calibration. For example, current $I_c(t)$ is needed for state of charge (SoC) computation. It is approximately constant in many practical settings, and it contributes limited additional variability compared to the voltage trace.

In this paper, we have two complementary goals for battery health monitoring. Our first goal is to quantify the battery's health for each cycle, and the second goal is to interpret that health estimate by providing physically meaningful interpretations of the degradation effects. First, we follow common practice and use SoH as the primary cycle-level indicator of




Fig. 2. Illustration of the polarization loss and tail loss for a representative battery cell across different discharge cycles.




Fig. 3. Illustration of the physics-based models for ECM and FOECM.

degradation. SoH measures the remaining usable capacity relative to a fresh cell and is typically defined as $Q_t/Q_0 \times 100\%$ as a percentage for the SoH at time t , where Q_t denotes the usable capacity at t and Q_0 is the rated capacity at the beginning of life. While SoH provides a concise scalar summary of degradation, it does not reveal degradation mechanisms underlying a specific health level. As such, our second goal, also the main goal of this paper, is to go beyond SoH and interpret it using factors with consistent physical meaning.

In practice, two batteries can share a similar SoH yet behave quite differently in operation. Motivated by common observations in battery health studies, degradation shown in routine discharge voltage curves typically appears in two distinct ways. One is the curve-wide under-load loss, where battery aging often presents as a global downward shift across discharge curves from different cycles. This indicates increasing under-load voltage loss and decreasing power capability, often described as polarization loss, as shown in Fig. 2. As this loss grows, the battery becomes more prone to voltage drop and heat generation under high current, since resistive heating scales with current squared. The other is the tail loss. Aging can be more pronounced near the end-of-discharge per cycle, where the voltage drops below the minimum required voltage and approaches the low-voltage cutoff earlier. Part of the remaining energy becomes infeasible to use under practical voltage constraints. As such, a single SoH value can appear optimistic even though end-of-discharge usability deteriorates.

Accordingly, we represent battery health as a compact 2-D aging fingerprint $\mathbf{z} = (z^1, z^2)$, where z^1 and z^2 quantify the polarization loss and tail loss, respectively. We use two views of this fingerprint. In the cycle domain, we form a discrete sequence $\{(z_c^1, z_c^2)\}_{c=1}^{n_c}$ as the fingerprints across different cycles. In the SoH domain, we re-parameterize the fingerprints

as health-conditioned curves $\mathbf{z}(s) = (z^1(s), z^2(s))$ with SoH s as the independent variable. Fingerprints in this domain enable noise-robust interpretations across the aging trajectory.

B. Physics-based Model for Fingerprint Extraction

To extract an interpretable aging fingerprint from routine BMS logs, we adopt a physics-based battery model. Such a model provides a structured decomposition of terminal voltage into components with clear physical meaning, which directly supports our goal of quantifying the two voltage-related degradation mechanisms. Moreover, because the model structure is grounded in established battery domain knowledge, its key parameters can be identified from standard operating data without requiring extensive historical data or additional diagnostic experiments. Specifically, we introduce ECM and its fractional-order extension in this study. ECM first serves as a physics-grounded baseline for describing the discharge voltage under load [7]. It represents the terminal voltage as an OCV component minus an ohmic drop and a set of polarization or relaxation drops, as shown in Fig. 3(a), as,

$$V(t) = V_0(\text{SoC}(t)) - I(t)R_0 - \sum_{i=1}^m v_i(t), \quad (2)$$

where $V_0(\cdot)$ is the open-circuit voltage (OCV), which changes as the battery is discharged and depends on SoC. SoC can be computed by standard bookkeeping, e.g., $\text{SoC}(t) = \text{SoC}(t_0) - \frac{1}{Q} \int_{t_0}^t I(\tau) d\tau$ where t_0 is the start time and Q is the battery's usable capacity for SoC normalization. In addition, R_0 is the ohmic resistance, often modeled as a constant [14]. Each $v_i(t)$ corresponds to the voltage drop across the i -th resistor-capacitor pair. For each pair, the dynamics follow the standard relation as,

$$\frac{dv_i(t)}{dt} = -\frac{1}{R_i C_i} v_i(t) + \frac{1}{C_i} I(t), \quad (3)$$

where $I(t)$ is the applied current at time t , and R_i and C_i control how quickly the pair voltage responds and relaxes under current changes. The multiple resistor-capacitor pairs capture relaxation dynamics under the applied current $I(t)$, and their aggregate contribution, as shown in the last term of Eq. (2), forms a load-dependent voltage loss that shapes the terminal voltage over a broad portion of the discharge trajectory. While ECM is partially interpretable and widely used, representing multi-timescale relaxation such as tail loss may require many resistor-capacitor pairs and can still be insufficient for capturing aging behaviors. This motivates FOECM, which provides a compact yet flexible description of both the curve-wide polarization loss and the end-of-discharge tail loss. Compared with Eq. (2), FOECM preserves the same high-level voltage decomposition, i.e., OCV minus under-load losses, but refines the polarization-related voltage drop and introduces an explicit term for the tail loss. We therefore rewrite the terminal-voltage model as shown in Fig. 3(b) as,

$$V(t) = V_0(\text{SoC}(t)) - I(t)R_0 - v_{\text{dyn}}(t) - v_W(t), \quad (4)$$

where $v_{\text{dyn}}(t)$ denotes the polarization-related loss term and $v_W(t)$ denotes the tail-loss term. In FOECM, the term $v_{\text{dyn}}(t)$

is captured by a polarization pair composed of a resistor R_{dyn} in parallel with a constant phase element (CPE). $v_{\text{dyn}}(t)$ is defined as the voltage across the parallel resistor and CPE paths and we denote the currents through the two paths by $I_{\text{R}}(t)$ and $I_{\text{CPE}}(t)$, for the resistor and CPE, respectively. By Kirchhoff's laws, we have,

$$I(t) = I_{\text{R}}(t) + I_{\text{CPE}}(t). \quad (5)$$

Since the resistor and CPE are in parallel, they share the same voltage $v_{\text{dyn}}(t)$. Based on Ohm's law, we have,

$$I(t) = v_{\text{dyn}}(t)/R_{\text{dyn}} + I_{\text{CPE}}(t), \quad (6)$$

which makes the role of R_{dyn} explicit. For a given applied current $I(t)$, a large R_{dyn} requires a large voltage to sustain the resistor path current, and increases the polarization-based voltage drop for the terminal voltage. We therefore use R_{dyn} as the FOECM parameter that quantifies the polarization loss of the aging fingerprint. Note that CPE captures a phase-lag response and therefore does not react instantaneously to voltage changes. At this stage, it suffices to note that $I_{\text{CPE}}(t)$ is also driven by the same voltage $v_{\text{dyn}}(t)$ and our formulation does not depend on the CPE's exact form.

In FOECM, we model the tail loss term with a Warburg-type diffusion element to capture the fact that, near the end-of-discharge, charge carriers such as lithium ions cannot transport fast enough to sustain the load, which causes an extra voltage drop $v_{\text{W}}(t)$. Because the element lies on the main current path, the current through the tail element equals the terminal current, i.e., $I_{\text{W}}(t) = I(t)$. Here, we parameterize the tail element with a resistance-like magnitude R_{W} and a diffusion timescale τ_{W} . Although IBAM mainly uses time-domain battery trajectories, it is convenient to specify the Warburg element using its standard frequency-domain impedance form as,

$$Z_{\text{W}}(j\omega) = R_{\text{W}} / \sqrt{j\omega\tau_{\text{W}}}, \quad (7)$$

where $Z_{\text{W}}(j\omega)$ is the Warburg impedance, $j = \sqrt{-1}$ is the imaginary unit, and ω is the angular frequency. Here $Z_{\text{W}}(j\omega)$ specifies the frequency-domain impedance of the tail element, and describes the relationship between a current component at frequency ω and a corresponding voltage drop across the element. Equivalently, in the time domain, this means the tail voltage $v_{\text{W}}(t)$ depends on the current history, rather than only the instantaneous current. The equation highlights the distinct roles of R_{W} and τ_{W} , where the former directly scales the magnitude of the voltage drop $v_{\text{W}}(t)$ under a given current profile and the latter sets the diffusion timescale, e.g., slow diffusion with large τ_{W} . Consequently, R_{W} scales with the tail loss and we therefore use it to quantify the loss. In summary, FOECM instantiates the two voltage-based battery degradation effects with two role-specific parameters R_{dyn} and R_{W} . The parameters govern the curve-wide polarization loss and end-of-discharge tail loss, respectively, and we therefore instantiate (z^1, z^2) in the 2-D fingerprint with the FOECM's physically-interpretable $(R_{\text{dyn}}, R_{\text{W}})$.

C. Per-cycle Fingerprint Identification with Two-Stage LSM

We identify the cycle-wise FOECM fingerprint parameters from routine BMS logs on a per-cycle basis. In this study, the aging fingerprint is instantiated by two FOECM parameters R_{dyn} and R_{W} , and the rest of the FOECM parameters such as $V_0(\cdot)$, $\text{SoC}(t)$, and R_0 are obtained from standard calibration or treated as fixed inputs. Given a discharge cycle c with measured voltage and current $\{V_c(t_i), I_c(t_i)\}_{i=1}^{T_c}$, our goal is to find parameter values that make the FOECM-predicted voltage best match the measured voltage. We achieve this by LSM, which plays the role of physics-based parameter fitting. Because the polarization loss and tail loss affect different regions of the voltage trajectory, we adopt a two-stage strategy to separate the contributions of the two loss terms. In the first stage, we fit the polarization term over the entire discharge curve where the tail term is temporarily disabled. In the next stage, we fit the tail loss using the remaining voltage residual after subtracting the stage-1 predicted polarization voltage, so as to focus on the end-of-discharge behavior.

In the first stage, we estimated R_{dyn} by fitting the polarization voltage drop assuming $v_{\text{W}}(t) = 0$. For simplicity, we treat FOECM as a cycle-level forward model $\mathcal{F}(\cdot)$, which takes the measured current trajectory of a discharge cycle and produces a predicted terminal voltage trajectory. Specifically, we have,

$$\hat{\mathbf{V}}_c^1(R_{\text{dyn}}) \triangleq \mathcal{F}(\mathbf{I}_c; R_{\text{dyn}}, \Theta^1), \quad (8)$$

where the model specification includes the following components. We use the measured current time series $\mathbf{I}_c = \{I_c(t_i)\}_{i=1}^{T_c}$ of cycle c . R_{dyn} is the only cycle-wise parameter estimated at this stage. Θ^1 collects all fixed or known settings such as $V_0(\cdot)$ and R_0 shared across cycles. Any internal dynamic states of FOECM such as internal polarization/CPE states are initialized at the start of each cycle and then calculated forward under the input current. The model output is the voltage curve $\hat{\mathbf{V}}_c^1(R_{\text{dyn}})$ predicted by FOECM at the first stage when the polarization loss is parameterized by R_{dyn} . Essentially, we aim to apply LSM to estimate the cycle-wise fingerprint parameter R_{dyn} by comparing the model output voltage $\hat{\mathbf{V}}_c^1(R_{\text{dyn}})$ against the measured voltage $\mathbf{V}_c = \{V_c(t_i)\}_{i=1}^{T_c}$. The optimal R_{dyn}^c can be derived by,

$$R_{\text{dyn}}^c = \arg \min_{R_{\text{dyn}}} \left\| \mathbf{V}_c - \hat{\mathbf{V}}_c^1(R_{\text{dyn}}) \right\|_2^2, \quad (9)$$

where $\|\cdot\|_2^2$ is the sum of squared errors and the optimization is equivalent to minimizing the mean squared error.

In the second stage, we estimate R_{W} by fitting the tail loss term. We utilize the same BMS logs for current $\mathbf{I}_c$ and voltage $\mathbf{V}_c$ as the first stage. As tail loss is most pronounced near the end-of-discharge, we define a low-voltage gate to select tail-region timestamps for cycle c as,

$$\mathcal{T}_c \triangleq \{i \mid 1 \leq i \leq T_c \wedge V_c(t_i) \leq V_g\}, \quad (10)$$

where V_g is a fixed low-voltage threshold. We use $\mathbf{I}_c[\mathcal{T}_c]$ and $\mathbf{V}_c[\mathcal{T}_c]$ to denote the sub-vectors restricted to these gated timestamps for current and voltage, respectively. With the

polarization parameter fixed to the optimal R_{dyn}^c estimated in the first stage, we activate the tail term in FOECM by removing the constraint of $v_W(t) = 0$ and define an updated voltage predictor as,

$$\hat{V}_c^2(R_W) \triangleq \mathcal{F}(\mathbf{I}_c; R_{\text{dyn}}^c, R_W, \Theta^2), \quad (11)$$

where R_W is the only cycle-wise parameter estimated at this stage. Θ^2 contains the same fixed or known terms as Θ^1 with minor additions to support the tail-region fitting. We then apply LSM again by comparing the measured and predicted terminal voltage curves on the gated timestamps as,

$$R_W^c = \arg \min_{R_W} \left\| \mathbf{V}_c[\mathcal{T}_c] - \hat{V}_c^2(R_W)[\mathcal{T}_c] \right\|_2^2, \quad (12)$$

which yields the cycle-wise estimated R_W^c for cycle c . Altogether, we have the cycle c fingerprint $(R_{\text{dyn}}^c, R_W^c)$ to reflect the battery's degradation mechanisms for cycle c . As the battery continues operating, we apply the same per-cycle fingerprint identification online and maintain the fingerprint history up to the current cycle, i.e., $\{(R_{\text{dyn}}^i, R_W^i)\}_{i=1}^c$.

D. Physics-guided SoH-domain Fingerprint Mapping

While our two-stage LSM provides cycle-wise aging fingerprints, these estimations can fluctuate from cycle to cycle due to measurement noise and operating variability. Moreover, the cycle index itself is not a industry-standard battery health scale. To obtain a more interpretable and transferable representation, we map the fingerprints from the cycle domain to the SoH domain as $(R_{\text{dyn}}(s), R_W(s))$ for different SoH $s \leq 100\%$. Because SoH is not directly available in BMS logs, we predict it from the discharge voltage using a BiGRU-based estimator. With the predicted SoH, we then learn a physics-guided mapping that converts the raw cycle-wise fingerprints into physics-guided and interpretable SoH-domain curves.

1) *BiGRU-based SoH Modelling*: To enable SoH-domain fingerprinting, we first estimate a battery's SoH from routine discharge voltage data. We adopt a lightweight sequence regressor based on a two-layer BiGRU model which maps each cycle's voltage curve to a scalar SoH. Specifically, we have voltage $\{V_c(t_i)\}_{i=1}^{T_c}$ for cycle c and we construct a 4-channel time-series input. Since different cycles may have different sampling rates and lengths, we resample the raw trajectories onto a fixed-length time grid and construct the 4-channel input on this standardized grid via linear interpolation. The channels include the terminal voltage $V_c(t)$, the relative terminal voltage to OCV, the voltage change rate, and the timestamps. Each BiGRU layer uses 96 hidden units per direction with dropout and the output at the final timestamp is used as a cycle-level representation. A lightweight regression head with two linear layers and a ReLU activation in between maps the representation to a latent value which is processed by a Sigmoid function to generate a valid SoH value. During battery operation, once cycle c is observed, the BiGRU model outputs the predicted SoH $\hat{s}_c$, which serves as the health coordinate for the subsequent SoH-domain fingerprint mapping. Note that BiGRU is used as a practical SoH model

for IBAM as it offers a good balance between lightweight deployment and competitive modelling capability. Our team's research portfolio includes more sophisticated ML-based SoH modelling architectures, e.g., [6], [7], [9], which, however, are not adopted to keep the overall IBAM pipeline efficient.

2) *Physics-guided Mapping*: At the current cycle c , we have accumulated cycle-wise SoH estimates $\{\hat{s}_i\}_{i=1}^c$ from the BiGRU model and cycle-wise raw fingerprint estimates $\{R_{\text{dyn}}^i\}_{i=1}^c$ and $\{R_W^i\}_{i=1}^c$ from the two-stage LSM. Because the cycle index is not a comparable health scale across batteries, we convert these cycle-wise points into SoH-domain mappings. We form paired samples $(\hat{s}_i, R_{\text{dyn}}^i)$ and $(\hat{s}_i, R_W^i)$ for cycle $1 \leq i \leq c$. For intuitive understanding, these pairs form c points in a two-dimensional coordinate system, e.g., SoH as the horizontal axis and R_{dyn}/R_W as the vertical axis.

Our first task is to construct an SoH-domain mapping $R_{\text{dyn}}(s)$ that follows a physically consistent monotonic trajectory and stays close to the obtained points. Starting from the paired samples $\{(\hat{s}_i, R_{\text{dyn}}^i)\}_{i=1}^c$, we do not treat all points equally because the per-cycle identification of R_{dyn}^i can be affected by measurement noise and operating variability. We therefore associate each cycle i with a scalar error computed from the per-cycle FOECM voltage fit, and assign a large weight to a point when this error is small. Let $\varepsilon_i \geq 0$ denote the per-cycle squared error for cycle i , and we set the weight as $w_i = 1/(\varepsilon_i + \varepsilon_0)$, where $\varepsilon_0 > 0$ is a small constant for numerical stability. These weights allow the fitting procedure to automatically trust low-error points more. In this study, we perform isotonic regression and let $\text{iso-reg}(\cdot)$ be the regressor model. The regressor can be regulated by a constraint, which enforces the monotonicity of the curve $R_{\text{dyn}}(s)$ as a function of SoH s . To facilitate the regression, we sort the input points in the ascending order of SoH. Then, the regression process can be represented as,

$$\tilde{R}_{\text{dyn}}(s) = \text{iso-reg}(\{(\hat{s}_i, R_{\text{dyn}}^i), w_i\}_{i=1}^c; \theta_{\text{dyn}}), \quad (13)$$

where θ_{dyn} is the Boolean parameter of the monotonicity constraint. The regressor optimizes based on the weighted squared errors, and the model produces a non-decreasing fitted curve $\tilde{R}_{\text{dyn}}(s)$ in the SoH-domain.

The SoH-domain mapping for $R_W(s)$ follows the same pipeline as $R_{\text{dyn}}(s)$, except the weights calculation. Because R_W captures tail loss that is observed primarily near end-of-discharge, we assign a high (small) weight w'_i to cycle i in which the low-voltage tail occupies a large (small) fraction of the discharge duration. The tail and fraction can be determined by the same low-voltage gate during per-cycle LMS-based R_W identification. This allows the monotonic fit to trust R_W^i more when the tail effect is better supported by the data of cycle i . Here, the regression process for R_W is given by,

$$\tilde{R}_W(s) = \text{iso-reg}(\{(\hat{s}_i, R_W^i), w'_i\}_{i=1}^c; \theta_W), \quad (14)$$

where θ_W is the Boolean parameter of monotonicity.

After fitting, we produce a compact representation that can be queried online via a lookup table. We uniformly sample the SoH axis over the valid operating range and configure k




Fig. 4. IBAM learned degradation fingerprints for representative batteries in the short-life (Bat039 and Bat025), medium-life (Bat116), and long-life (Bat100 and Bat124) categories; each battery’s lifespan (in cycles) is indicated in the legend. Overall, short-life batteries exhibit stronger tail loss.

reference SoH points, e.g., $k = 25$ evenly spaced values over SoH from ideal condition 100% to end-of-life (EoL) 80%. During operation, given an inferred SoH s , e.g., by BiGRU, we locate the two neighboring reference points s_i and s_{i+1} and compute $R_p(s)$ by standard linear interpolation between $R_p(s_i)$ and $R_p(s_{i+1})$ in the lookup table for $p \in \{\text{dyn}, W\}$.

III. EXPERIMENTAL STUDY

A. Experiment Setup

In our experiments, we utilize the MIT-Stanford battery dataset [2] which is commonly used in battery studies. The dataset contains commercial cylindrical lithium-ion batteries with nominal capacity 1.1 Ah and voltage 3.3 V. The batteries were tested in a 30°C chamber via a cycler and discharged at a uniform 4C rate. For each battery and cycle, we focus on the battery operation scenarios and the discharge segment only. We standardize the start time of discharge as 0 and resample the entire discharge duration with 150 evenly distributed timestamps for consistency. We create three categories for representative batteries with short-, medium-, and long-lifespans. ML models are trained and evaluated using the processed discharge sequences, and optimized with the AdamW optimizer. IBAM is lightweight and experiments are implemented in Python 3.13 with PyTorch 2.7.1 and run on a CPU-only machine with Intel Core i5-220H and 24 GB RAM.

B. Interpretable Fingerprint Insights

In this part, we apply IBAM to extract degradation fingerprints of the tested batteries for health interpretation. For visualization clarity, we present representative batteries from each lifespan category. We select two short-life batteries, Bat039 and Bat025, which reach EoL at cycles 648 and 941, respectively. The medium-life category is represented by Bat116 with 1,155 cycles. The long-life category is represented by Bat100 and Bat124, which reach EoL at cycles 1,636 and 1,798, respectively. We visualize the learned fingerprints by plotting R_{dyn} versus SoH, R_W versus SoH, and R_W versus R_{dyn} in Figs. 4(a), 4(b), and 4(c), respectively. These visualizations allow us to compare the trends of degradation mechanisms and we have the following observations.

First, we observe that both fingerprint components R_{dyn} and R_W increase as the battery ages. As shown in Fig.

4(a), R_{dyn} rises as SoH decreases, which indicates a growing polarization loss. Similarly, Fig. 4(b) shows that R_W increases with aging, indicating an increasingly severe tail loss during the end-of-discharge. Importantly, both mechanisms can be observed throughout the tested health range. Even at high SoH, the batteries already exhibit measurable polarization loss and a non-zero tail loss. The fingerprint trends suggest that battery aging is not a single-mode process and different losses intensity concurrently. Among the two losses, R_W exhibits more significant lifespan discriminability than R_{dyn} , which suggest that R_W can be a key differentiator of early EoL.

Our next observation is a consistent fingerprint trade-off that aligns with lifespan. Long-life batteries are polarization dominant and relatively tail-healthy, whereas short-life batteries suffer more from tail loss. At a comparable health level, e.g., 90% SoH, the two short-life batteries Bat039 and Bat025 exhibit low R_{dyn} of 0.071 and 0.070 Ω, respectively, and 0.038 and 0.028 Ω, respectively, for R_W . In comparison, the long-life batteries Bat100 and Bat124 show higher R_{dyn} , e.g., 0.081 and 0.079 Ω, yet substantially lower R_W , e.g., 0.017 and 0.012 Ω at the same SoH. The trajectories of R_{dyn} and R_W of the medium-life battery Bat116 generally lie between the trajectories of the short- and long-life batteries. These results suggest that, at similar health levels, long-life batteries primarily incur a broad and curve-wide polarization penalty for discharge voltage while maintaining a strong end-of-discharge performance. Short-life batteries develop a pronounced tail bottleneck that can limit practical energy delivery even when SoH remains comparable. Importantly, a tail bottleneck has a threshold-like operational impact. As tail loss grows, the voltage near the end-of-discharge drops faster under the same load, so the battery reaches the cutoff voltage earlier and discharge terminates sooner. This means that usable battery energy is increasingly constrained by early cutoff. In contrast, long-life batteries’ low R_W suggests that they sustain the end-of-discharge voltage for longer, so the batteries’ degradation shows more as distributed loss rather than abrupt tail losses.

Our third observation is that the proposed 2-D fingerprint offers an informative and intuitive visualization of degradation mechanisms, as shown in Fig. 4(c). Compared with tracking $R_{\text{dyn}}(s)$ and $R_W(s)$ separately along the SoH axis, the joint

TABLE I

THE PERFORMANCE OF DIFFERENT PHYSICS-BASED BATTERY MODELS, INCLUDING ECM, FOECM_{base} WITHOUT THE TAIL COMPONENT R_W , AND IBAM'S FOECM IN TERMS OF PER-CYCLE VOLTAGE-FITTING RMSE. FOUR REPRESENTATIVE BATTERY OPERATION STAGES ARE CONSIDERED, INCLUDING 100%, 67%, 33%, AND 0% REMAINING LIFESPAN BASED ON 100% SOH AND EoL. A SMALL VALUE INDICATES A GOOD PERFORMANCE.

Optimized Physics Battery Model	100% Lifespan			67% Lifespan			33% Lifespan			0% Lifespan		
	Short	Medium	Long	Short	Medium	Long	Short	Medium	Long	Short	Medium	Long
ECM	53.067	48.350	46.723	39.073	34.607	29.777	42.073	37.133	32.140	39.813	38.577	32.663
FOECM _{base} (w/t R_W)	52.740	48.053	46.053	37.997	33.890	28.350	41.167	36.340	30.763	38.627	37.850	31.293
FOECM (IBAM, ours)	45.457	43.547	43.303	34.790	31.340	27.400	36.577	33.387	28.310	34.050	34.200	28.727

view in the (R_{dyn}, R_W) plane provides a compact health portrait that distinguishes batteries with different lifespans more clearly. In particular, long-life batteries concentrate in a region with relatively high R_{dyn} but low R_W , whereas short-life batteries extend toward much larger R_W . Moreover, the SoH-colored trajectories reveal distinct aging trajectories. Short-life batteries follow a relatively vertical path with rapid growth in R_W , while long-life batteries evolve more uniformly. Overall, the 2-D fingerprints enable an efficient comparison of degradation mechanisms beyond a single SoH.

A practical implication of the 2-D fingerprints is that they suggest different battery operation priorities even at the same level of battery health. For batteries dominated by tail loss, a main risk is early cutoff near end-of-discharge. Thus, it suggests conservative battery operations when the battery capacity is low, e.g., avoid deep discharge and apply a higher practical cutoff voltage under load. In contrast, different operation aspects shall be considered for polarization-dominant batteries, e.g., limiting peak current through non-fast charging/discharging or current derating to minimize the associated battery stress. Overall, IBAM provides a mechanism-aware basis with interpretability for optimizing battery operations.

C. Physics-Model Fitting Performance

In this part, we evaluate the performance of different physics-based battery models. We identify and optimize their key parameters with IBAM and the models' performance is reflected by comparing the models' per-cycle output voltage and the measured voltage from routine BMS logs. Specifically, we consider a ECM with one resistor-capacitor pair, FOECM_{base} without the Warburg diffusion element and R_W , and a full-version FOECM with both R_{dyn} and R_W . We report the models' voltage-fitting error in terms of root mean square error (RMSE), evaluated at four lifespan stages, i.e., 100%, 67%, 33%, and 0% of the lifespan, the first and the last of which denote the cycle with 100% SoH and the EoL cycle, respectively. The results are presented in Table I.

First, IBAM's optimized FOECM achieves the best fidelity. As shown in the table, FOECM yields the lowest fitting error in all reported settings covering four lifespan stages compared to ECM. For the short-lifespan batteries at their EoL cycles, ECM has an RMSE of 39.8 and the RMSE of FOECM is 34.1, which means a 14.5% error reduction. When we consider the best health with 100% SoH for the same battery category, FOECM can reduce the error by 14.3%. In addition,

FOECM_{base}, i.e., without capturing tail loss R_W , can reduce the errors of ECM, but still shows significant performance gaps compared to FOECM. For the two settings above, the errors of FOECM_{base} are 38.6 and 52.7, respectively, with 3.0% and 0.7% error reductions over ECM, which are much lower than the percentage reductions achieved by FOECM. The results show that simply switching from ECM to a basic FOECM offers limited gains, and the full IBAM design contributes the dominant improvement in modelling fidelity.

We also observe that IBAM's performance advantages are more evident when batteries are close to their EoL. This is visible in the table when IBAM FOECM is compared against ECM and FOECM_{base} in different lifespan stages. For short-life batteries at 67% remaining lifespan, IBAM improves over ECM and FOECM_{base} with 11.0% and 8.4% error reductions, respectively. At the batteries' EoL cycles, the reductions enlarge to 14.5% and 11.8%, respectively. Similar patterns can be observed for medium- and long-life batteries, e.g., the percentages increases from 8.0% and 3.4% at 67% remaining lifespan to 12.1% and 8.2% at the EoL cycles. The results show that the benefit of IBAM FOECM is not a constant, and the solution becomes more valuable when batteries are more aged, where voltage behavior can be harder to fit. Meanwhile, FOECM_{base} performs relatively similar to ECM, which implies that the late-life performance gain of IBAM does not come from just switching ECM to FOECM. As the voltage behavior distorts near EoL, IBAM's explicit modelling of the tail terms becomes increasingly important.

Besides, voltage fitting is more challenging for short-life batteries than long-life ones and the errors of short-life batteries are higher than the rest in most of the reported settings. For example, ECM reduces the errors from 53.1 for short-life batteries to 46.7 for long-life batteries when the batteries are new. For aged batteries, ECM can still reduce the errors from 39.8 to 32.7. IBAM's FOECM achieves 43.3 RMSE for long-life batteries, lower than 45.5 for short-life batteries at 100% lifespan. When one-third of the lifespan is left, IBAM's error is 28.7 for long-life batteries, smaller than 34.1 for short-life batteries. The results imply that short-life batteries exhibit voltage dynamics that is more difficult to explain using the same model and parameter identification process. This is consistent with our previous observations about short-life batteries' increased tail effect and potentially strong variability. On the contrary, long-life batteries relatively follow regular degradation patterns that are easier to be modeled.

TABLE II
IMPACT OF 4-CHANNEL FEATURE CONSTRUCTION ON THE BiGRU SoH ESTIMATOR USED IN IBAM. RESULTS ARE REPORTED IN MAE AND RMSE, WHERE SMALL VALUES INDICATE GOOD PERFORMANCE.

BiGRU Input Battery Lifespan	MAE		RMSE	
	1-Channel	4-Channel	1-Channel	4-Channel
Short	0.0141	0.0039	0.0174	0.0050
Medium	0.0184	0.0053	0.0216	0.0060
Long	0.0111	0.0056	0.0132	0.0066

D. Impact of Multi-Channel Inputs on BiGRU SoH Modelling

In IBAM, SoH is not only a prediction target but also the health coordinate used to construct and query the SoH-domain fingerprints. Therefore, errors in SoH estimation propagate directly into the fingerprints. This can blur the SoH-fingerprint relations across cycles and reduce the quality of interpretation. As such, we conduct an ablation study on the BiGRU-based SoH estimator, specifically on the impact of our designed 4-channel feature construction to justify the design for IBAM. Here, we keep the BiGRU backbone and training protocol fixed and vary the input representation. The baseline model takes the raw discharge voltage sequence of each cycle as a 1-channel feature as input. In IBAM, we construct a 4-channel input with the raw voltage, relative voltage, voltage change rate, and timestamps. We evaluate SoH modelling performance using mean absolute error (MAE) and RMSE and report the results in Table II for batteries in different lifespan categories.

First, we can see that the proposed 4-channel feature construction consistently improves BiGRU SoH modelling accuracy over using a single raw-voltage channel. For example, MAE drops from 0.014 with 1-channel input to 0.004 with 4-channel input for short-life batteries, where RMSE decreases from 0.017 to 0.005. For medium-life batteries, 4-channel input helps reduce the errors by 71.2% and 72.2% for MAE and RMSE, respectively. The improvements are expected because the multi-channel features expose complementary battery health patterns that a single voltage trace is not able to provide. For example, the absolute voltage reflects the overall discharge profile and the voltage change rate captures local dynamics of battery operation. Together, multiple channels provide a more informative representation and allow the lightweight BiGRU to learn health-related patterns effectively. The accurate SoH estimates directly benefit IBAM, by serving as accurate coordinates of the health fingerprints.

We also observe a mild trend that SoH modelling becomes more challenging for batteries with longer lifespans. IBAM's SoH modelling MAE increases from 0.004 for short-life batteries to 0.006 for long-life batteries, and RMSE increases from 0.005 to 0.007. A possible reason is that long-life batteries tend to exhibit slow degradation, which brings challenges to distinguish nearby SoH levels between per-cycle voltage features. Besides, long degradation trajectories may accumulate battery operation variability over time, which increases diversity of the long-life category and thus modelling errors even with the same ML backbone. Nevertheless, the errors remain low

in IBAM and are consistently small, suggesting that IBAM remains robust even for more challenging SoH modelling.

IV. CONCLUSION

In this paper, we propose IBAM for interpretable battery aging modelling from routine BMS logs. IBAM constructs a physics-based battery model and extracts 2-D aging fingerprints from the model to quantify polarization and tail losses. To obtain the fingerprints accurately, IBAM performs per-cycle parameter fitting with a two-stage LSM. It then estimates cycle-wise SoH via BiGRU with multi-channel voltage features, and produces physically consistent SoH-indexed fingerprints. Experiments on batteries with different lifespans show that IBAM consistently improves modelling fidelity and offers interpretable fingerprint patterns that are not visible from SoH alone. Overall, IBAM bridges accurate battery health estimation and actionable interpretability by turning routine BMS logs into degradation mechanism indicators, and can improve BMS operations and enable safe battery usage.

REFERENCES

- [1] IEA, "Iea global ev outlook 2025," <https://www.iea.org/reports/global-ev-outlook-2025>, 2025, accessed: 2026-01-29.
- [2] K. A. Severson, P. M. Attia, N. Jin *et al.*, "Data-driven prediction of battery cycle life before capacity degradation," *Nature Energy*, vol. 4, no. 5, pp. 383–391, 2019.
- [3] K. Kokkalis, C. Chronis, E. Politi *et al.*, "Predicting the lifespan of lithium-ion batteries using machine learning, parameter tuning and model ensembles," in *2024 IEEE 10th International Conference on Big Data Computing Service and Machine Learning Applications (Big-DataService)*. IEEE, 2024, pp. 122–129.
- [4] Z. Liu, Y. Liu, Y. Zhang *et al.*, "Data-driven lithium-ion battery soh prediction: A novel shmm-transformer-bigru hybrid neural network method," *Measurement*, p. 118579, 2025.
- [5] S. Sameer, W. Zhang, X. Lou *et al.*, "Ginet: Integrating sequential and context-aware learning for battery capacity prediction," *arXiv preprint arXiv:2501.04997*, 2025.
- [6] W. Li, W. Zhang, and Q. Yan, "Entrolnn: Entropy-guided liquid neural networks for operando refinement of battery capacity fade trajectories," *arXiv preprint arXiv:2601.06195*, 2026.
- [7] S. Sameer, W. Zhang, K. D. Dharshini *et al.*, "Pace: Physics-aware attentive temporal convolutional network for battery health estimation," *arXiv preprint arXiv:2512.11332*, 2025.
- [8] H. Eivazi, A. Hebenbrock, R. Ginster *et al.*, "Diffbatt: A diffusion model for battery degradation prediction and synthesis," *arXiv preprint arXiv:2410.23893*, 2024.
- [9] Z. Kwan, W. Zhang, A. B. Ng *et al.*, "Knowdiff: Knowledge-regulated dual-stream diffusion for robust battery prognostics," *Available at SSRN 5871665*.
- [10] E. I. H.-E. Lim and N. H. L. Wong, "A unified framework for interpretable and uncertainty-aware battery state of health estimation using deep neural networks," in *2025 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2025, pp. 1392–1397.
- [11] O. Gahramanov, H. Hokmabad, E. Ginzburg-Ganz *et al.*, "Detecting battery degradation factors using explainable ai," in *2025 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe)*. IEEE, 2025, pp. 1–5.
- [12] T. Tetik, "Feature engineering and explainable artificial intelligence for state of health estimation of lithium-ion batteries," *Journal of Energy Storage*, vol. 144, p. 119873, 2026.
- [13] K. Liu, Y. Liu, S. Zhao *et al.*, "An ultrasonic wave-based method for efficient state-of-health estimation of li-ion batteries," *IEEE Transactions on Industrial Electronics*, 2024.
- [14] Y. Xie, W. Guo, T. Zhou *et al.*, "An impedance-based electro-thermal model integrated with in-situ lithium-plating criterion for ac heating at low temperatures," *Applied Energy*, vol. 391, p. 125930, 2025.